



SUNGARD[®]
Availability Services

Mastering Recovery

Exploring the options for getting
your business back in business.

Table of Contents

Chapter 1	
The Changing Face of Disaster	5
Chapter 2	
Assessing Where You Are – And Where You Need to Be	9
Chapter 3	
Overcoming the Obstacles	15
Chapter 4	
Someone’s Looking Over Your Shoulder	19
Chapter 5	
Review Your Options: Tape-based Strategies	23
Chapter 6	
Review Your Options: Electronic Backup Strategies	26
Chapter 7	
Review Your Options: Advanced Recovery Strategies	29
Chapter 8	
You Can’t Be Too Careful	33
Chapter 9	
Starting at the Very Beginning	36
Chapter 10	
Make the Grade with Testing	41
Chapter 11	
Choosing the Right Provider	43
Chapter 12	
Who We Are	46
Appendix	
Get Started Today	48
Notes:	49

Exploring the Options for Getting Your Business Back in Business

A Foreword from SunGard

Organizations everywhere are discovering that it's not easy managing IT requirements these days. Technology is moving forward at an unprecedented rate and our dependence on it continues to grow. Management is under the gun to evaluate and achieve best-in-class solutions for everything from system maintenance to ongoing operations. Not only that, but the technical infrastructure is the very foundation for most routine aspects of business, so you simply can't function without your computing and telecommunications resources. However, budgets remain flat, forcing everyone to do more with less.

Despite our increasing reliance on computing and telecommunications in nearly every aspect of business, most corporations remain ill prepared to recover from the loss of their IT infrastructure. By remaining unprepared, they're playing the odds hoping disaster won't strike them.

But those odds are mounting against them, every day.

Just consider the volatile storm season of 2005, one of the busiest on record. For the first time ever, the National Hurricane Center ran through its preapproved list of names for hurricanes and tropical storms before the season's end. And businesses up and down the Eastern seaboard and along the Gulf Coast went into disaster mode as a result of the devastating aftermath of these storms.

Let's look at the results of an *Economic Outlook Survey*, taken in the wake of Hurricane Katrina by Business Roundtable, an association of chief executive officers of leading corporations. **Strongly to moderately negative effects** were felt by **65 percent** of the businesses. Of that group, 32 percent believed that they would be feeling these effects for fewer than three months, while **63 percent** responded that they would be **feeling the effects for three to 12 months**. 27 percent of the respondents pointed to "damage to infrastructure and facilities in the affected region" and 19 percent cited "supply disruptions" as issues

that would bear on their businesses. Of course, natural disasters like hurricanes aren't the only disruptive forces that can impact your organization. There are scores of others, ranging from blackouts to acts of terror to wild fires, which could separate you from your computer resources and data.

Are you ready for such disasters? Think about your response. Will you:

- Reach for the antacids or the telephone?
- Hit the panic button or the ground running?
- Send out your resume or the call to action?
- Lose sleep over what happens next or rest assured that your business will stay up and running?

Unfortunately, many organizations will opt for the former alternative in each case—and they're going to have a bumpy ride. That's because they need access to data, to systems, to applications to get through the day. To keep the production environment humming along and do what the business is supposed to be doing. To make sure that employees stay on the job and are productive. To stay in business.

Even in today's age of computer system-dependence, far too many enterprises aren't adequately prepared to deal with disaster. Maybe they're unaware of the options open to them. It could be they're informed but don't have the budget, resources or time to initiate a preparedness program. Or perhaps they're thinking that these things only happen to others.

Whatever the reason; the truth is, situations such as those above (and others like them), can happen to anyone. As a matter of fact, the chances of an interruption occurring within the work environment are increasing every day.

That's why we have created this primer on "Mastering Recovery." It covers the various and multiplying forms of disaster, the risks you and your organization face, how to prepare for a disaster, the costs of DR planning—and of no planning. This primer is decidedly not an all-inclusive treatise on the topic. Frankly, we'd need to write a much longer document and take up far too much of your reading time to cover everything you need to know. Rather, our guide is an elementary introduction designed to get you thinking—and to take action now.

Chapter 1

Mastering Recovery: The Changing Face of Disaster

Before we begin to describe the changing face of disaster, let's try to define "disaster."

A few years back, in a book called *Action Plan for Disaster*, SunGard published a serviceable definition. We said a disaster is: "*Any unplanned, extended loss of critical business applications due to lack of computer processing capabilities for more than a 48-hour period.*"

Is there anything wrong with that definition? These days, the answer is a resounding "Yes!" How many businesses today would suffer severe losses well before 48 hours passed without computer processing capabilities? That traditional recovery window is now more of a peephole. Many companies who formerly operated on the premise that recovery within 48 hours was acceptable have since narrowed that

to within 24 hours, or to 12 hours, and, in some cases, to minutes or seconds.

That's because, today, the price of losing computing or telecommunications capabilities is incredibly high:

- A *Computerworld* article reported that it is “estimated that 50% of all firms go out of business within three years of a major disaster.”
- The *Washington Post* wrote that the Institute for Business and Home Safety, an insurance industry trade group, estimated that 25 percent of businesses that close during a disaster will not re-open.
- *Computer Technology Review* informed readers that the “estimated costs of systems failures would be nearly fatal to some companies and can range to nearly \$3 million/hour of downtime in certain industries.” The publication also pointed out that a figure of 99 percent availability seems to be pretty good, unless you realize that this number translates into **5,000 minutes of unavailability** per year.

It's probably tempting to think that damage on that scale could be caused only by some cataclysmic event. Unfortunately, disasters come in some pretty ordinary shapes and sizes. (On the bright side, however, options for disaster recovery are as diverse as the disasters themselves. More on this in Chapters 5 through 7.)

You've probably been around the industry long enough to have seen some of the types of disasters that can affect computer operations. For years, such lists were confined to disasters that might strike a traditional glass house, the home of a mainframe, which of course was the heart of each large computing environment.

Mainframes still are the core of large computing environments, but critical computing and telecommunications operations, along with the list of potential disasters, have expanded well beyond the glass house throughout the enterprise: to the remote office, to the desktop, even to the BlackBerry®.

If you haven't already done so, take a few moments to consider the new forms disaster could take at your own company, some that might not have worried you just five or three or even one year ago.

Disasters A-Z

Acts of God–air conditioning failure–arson–Blackouts–blizzards–boiler explosion–bomb threats–bridge collapse–brownouts–brush fires–Chemical accidents–civil disobedience–communications failure–computer crime–corrosive materials–Disgruntled employees–denial of service–Earthquakes–embezzlement–explosions–extortion–Falling objects–fires–floods–Hardware crash–high winds–heating or cooling failure–hostage situations–human error–hurricanes–Ice storms–interruption of public infrastructure services–Kidnapping–Labor disputes–lightning strikes–Malicious destruction–military operations–mismanagement–mud slides–Personnel non-availability–plane crashes–phishing–public demonstrations–Quirky software–Radiology accidents–railroad accidents–Sabotage–sewage backups–snow storms–software failure–sprinkler failure–Telephone problems–theft of data or computer time–thunderstorms–tornados–transportation problems–Unexpected, the–Vandalism–viruses–Water damage–worms–Xenon gas leaks–Yellow fever outbreak–Zombie, attack of the (yes, that really is a hacker attack!)

Finished? Now, consider a few of the changing aspects of disaster you might not have thought of:

- **There's a Technology Assimilation Gap:** This is a concept developed by Anderson Consulting, which identifies the problems created by technologies that evolve faster than we can assimilate the changes.
- **Multiple Platforms Abound:** As recently as a few years ago, it would have been sufficient in the event of a disaster to fly your tapes to a hot site, commandeer a 600J and restart your operations. You probably can't do that anymore. Today, operations typically have several platforms to recover—everything from mainframes to distributed processors like IBM® iSeries, Sun, or HP, to blade servers, to PC's, to telecommunications systems. It's harder than ever to keep all the balls in the air.
- **Security Breaches are Escalating:** Now that most companies are increasing their exposure to security risks through such windows as distributed computing and Internet use, the number and cost of security breaches continues to spiral upward. The 10th Annual Computer Crime and Security Survey reported that the cost of

security breaches “averaged \$204,000” per respondent, with unauthorized access and denial of services as two of the most significant concerns. And the coverage generated by several prominent companies that “misplaced” consumer data has worked to keep security a top-of-mind concern.

- **Compliance is Ever-Present and Growing:** Whether it's Sarbanes-Oxley, Healthcare Insurance Portability and Accountability Act (HIPAA), Gramm- Leach-Bliley, or marching orders from somewhere else, you're required to safeguard the availability, integrity and privacy of information and the systems on which it resides. And you better have the planning in place as well as the documentation to prove it.

As you can see, it's a complicated scenario and one that's only likely to worsen. That's why you need to determine where you stand with your disaster recovery planning.

Chapter 2

Mastering Recovery: Assessing Where You Are – And Where You Need to Be

People who manage their company's computer operations fall into one of four general categories in defining their level of disaster preparedness. There are those who:

- Understand and are ready
- Understand and are not ready
- Understand but don't want disaster recovery planning
- Don't understand

Into which category do you fall?

The question is an important one, because knowing where you are today is the first step, however small, toward ensuring that you are prepared to recover from a disaster. So, let's find out where you are today.

The introduction to this book asks how you will respond to a disaster scenario. If you had a rational, effective answer, you're probably either in the first group or the fourth—either you **“Understand and are ready”** for disaster or you **“Don't understand”** and don't realize you need to worry. To find out which, ask yourself:

- Do I have a disaster recovery plan?
- Is my recovery plan detailed, down to describing who does what, when, and how, so that we can recover from a disaster without reference to any other document?
- Is my recovery plan comprehensive, covering all technology-dependant units within the company for all types of configurations: WANs, LANs, client servers, voice and data networks, remote sites, and last, but not least, distributed platforms and mainframes?
- Are business, people and processes included?
- Is my plan up to date?

If you answered, “yes” to all four questions, you are probably in the **“Understand and are ready”** for disaster category. Four “no” answers exclude you from the first category, but not from the others, so let's probe a little further. Do any of the following statements apply to you?

- I know that if my company is not sufficiently prepared to recover from a disaster, I can be fired and/or sued.
- I know that 100% of companies who report computer downtime also report lost revenues.
- I know that estimates for lost revenue during unscheduled computer downtime range from hundreds of thousands to millions of dollars, depending on my business.
- Thinking about the three preceding points sometimes keeps me awake at night.

If you answered, “yes” to all four, you probably fall into the **“Understand and are not ready”** group. If you confidently replace the fourth bullet point with...

“What, me worry?”

...you probably fall into the **“Understand but don’t want disaster recovery planning”** category.

Finally, if you read the first three points and after each one said to yourself, “I didn’t know that,” you probably belong to the fourth group: **“Don’t understand.”**

For the members of that group—those who really aren’t aware of the implications arising from the threat of disaster, there’s an excuse. It’s an oversight. But the others don’t have this excuse. For them, there isn’t a cogent argument for operating without a comprehensive up-to-date disaster recovery plan. So why do so many companies do just that? Some are simply hesitant to take that first step.

There are reasons we’ve heard from customers and prospects that have been leery about moving forward, and we will look at them in Chapter 3—Overcoming the Obstacles. But for now, let’s ask another, very pointed question:

Can you live without Disaster Recovery Planning?

You can begin to answer that question by asking yourself a different one: What would the loss of your computing and/or telecommunications capabilities cost your organization?

Certainly, it varies according to your size, industry and data dependence. One study by Infonetics Research polled 80 large organizations and found that **“overall downtime costs averaged an astounding 3.6% of annual revenue.”** Another survey sponsored by AT&T and the International Association of Emergency Managers, and reported on by CNN/Money found that fully two-thirds of companies that suffered a disaster lost business, with 16 percent losing between \$100,000 and \$500,000 per day.

There’s a quick and easy way to roughly determine your costs. Ask yourself what are the organization’s key business functions (sales,

service, supply chain, etc.). Now list the IT and business resources that support them. Next, assume that you were no longer able to maintain the continuity of one of these functions because an unplanned event disrupted one of these resources.

Then, estimate the financial impact this loss would have on your business based on how long this resource was unavailable—a minute, an hour, a day or worse. Multiply this number for every business process you can't maintain, and you start to see the magnitude of the problem.

Of course, there are other costs to consider:

- lost or compromised data that can't be replaced;
- missed sales opportunities and untapped revenues, gone forever;
- downtime that ripples throughout the enterprise to affect other users and departments;
- customers, stockholders, business partners and other end-users who are dissatisfied—or worse;
- kinks in the supply chain that affect everyone from raw material suppliers to distributors; and
- fines and penalties from regulatory agencies.

We recommend this self-evaluation because most people tend to regard the cost of a disaster as equal to only the cost of replacing the damaged or malfunctioning equipment. Thus, if a PC crashes, the cost of the “disaster” is merely the cost of replacing the defective hardware—probably no more than a few hundred dollars. But, as you can see from the above examples, there are far too many collateral costs.

So, this disaster preparedness self-evaluation may be a fairly sobering exercise, perhaps even more so if you consider all the key risk areas listed. Obviously, there is quite a bit at stake.

In the previous chapter, we reviewed the changing face of disaster—the many and diverse ways you can lose your computer and telecommunications operations. And now you have a fairly clear picture of everything you stand to lose if disaster strikes—the true cost of disaster.

If disaster can come in practically any shape and happen to practically anyone, and if the cost of going without disaster recovery planning can be so high, why would you—why would anyone—take that chance? Well, the cost of going without has a natural complement, but disaster recovery planning isn't free, which begs the next question...

Can you live with Disaster Recovery Planning?

No two information systems configurations are alike. If configurations were standardized beyond the physical aspects of hardware and commercial software packages, then recovery planning might be a very different thing from what it is. However, because no two are alike, any ballpark figure on the cost of recovery planning must be accompanied by a caveat: prices may vary by method and vendor.

According to META Group research, "more than 80% of Global 2000 companies will significantly increase their business continuity/disaster recovery budget by 2006—from the current average of about 2-3% of total IT budgets to 5%+ of their budget." This includes everything: auditing, planning, testing, etc. And there's no question that this is real money, especially in these times of cutbacks and stagnant spending.

The expenditure is big enough on its own, but has been magnified by the trend in corporations to downsize, to focus on core competencies, and to trace everything to the bottom line. In fact, companies that use Shareholder Value Added (the practice that tests every dollar spent for maximum return) have been given higher earnings multiples by investors.

In this environment, with its intense focus on the bottom line, the chief financial officer and the chief information officer can be your biggest allies. That's because they, better than anyone else, understand the full ramifications of their decision: should we spend money on disaster recovery or leave our business exposed?

Making a successful pitch will require getting the CFO and the CIO, or their counterparts in your organization, on board. Don't take them for granted, though, because they may resist; the CFO because the costs don't contribute directly to profits and the CIO because this would consume dollars that might otherwise be spent on improving day-to-day operations or other strategic IT projects.

Note that neither the CFO nor the CIO actually likes these costs. After all, disaster recovery is a non-operating expense—a non-operating expense that could save your business from catastrophe, to be sure, but a non-operating expense nonetheless. Still, both understand that the value received from disaster recovery planning is equal to, and greater than in many cases, the potential losses avoided. (This, as we saw earlier in this chapter, can be enormous.) So, to borrow from an old saying, disaster recovery is the ounce of prevention that is worth the pound of cure.

What's really the hold-up? Let's take a look at some valid reasons we've heard from customers and prospects on why they're reluctant to move ahead with a comprehensive recovery plan.

Chapter 3

Mastering Recovery: Overcoming the Obstacles

Change is scary, and making a big change with impact felt across your enterprise can be downright intimidating. For many, it's the fear factor that prevents them from developing and implementing a disaster recovery strategy. We've routinely encountered five common fears that customers tell us are typical.

1) It will cost a bundle.

Disaster recovery doesn't come out of petty cash, that's for sure, so you may be thinking that it's more economical to develop an in-house capability. But that's not so. Upgrading systems and infrastructure—and doing it twice for redundancy—requires substantial capital investment, considerable time and skilled resources. It also necessitates ongoing expenses for managing, monitoring and upgrading the network, systems and applications, as well as training staff in unfamiliar or emerging technologies.

However, taking advantage of a third-party DR provider's investment—in both technology and people resources—can reduce these expenses while ensuring that you have all the needed resources at your disposal.

2) I'll have to make a commitment.

Yes, that's true and there's no sidestepping this one. Disaster planning and recovery must be a management priority, starting at the top. Boards of directors have to be made to understand the reality of threats to IT resources as well as the risks if these resources are interrupted. They also need to be encouraged to allocate sufficient funding for a disaster recovery program.

As mentioned in the previous chapter, C-level management also needs to "green light" the project, leading the cause internally, legitimizing it and spreading the word throughout the organization. And remember, this isn't always an easy sell: senior managers may not be willing at first because of allocating money, time and personnel to activities that don't contribute directly to your profitability or strategic direction.

So, yes, it does require a commitment, but, for many, many reasons, it's one you won't regret. Consider the alternative when a disaster threatens your IT systems and you have no planning: you fail to maintain critical operations—and perhaps fail to stay in business, even for as few as five minutes. A little commitment goes a long way in preventing this from happening.

3) I could make a bad decision.

We can alleviate most concerns in this area by stating that almost any planning you do to prepare for a disaster is a good thing. However, having said that, there are some choices that may not be the best for your organization.

- You could focus on the IT end of your operation and neglect the business end of things. Bad idea. The data center is vitally important, but work processes and people drive the business. Cover the technical and business aspects in your planning and you can't go wrong.
- You could direct planning efforts toward recovering from disaster

instead of maintaining ongoing operations. That's a myopic view. You want to create a strategy that helps avoid losing systems, devices, networks and applications in the first place. But, should a part of your IT infrastructure fail, you also want to make sure that you're able to maintain business as usual and keep the rest of the world unaware of that failure.

- You could forget to factor logistics into your planning. Alternate work spaces, telephones, PCs, emergency notification procedures, communications strategies, they'll all figure in how successful you are at keeping things up and running.

How do you avoid making the wrong choices? Expert advice is readily available from trustworthy and reliable sources, including disaster recovery vendors, business continuity consultants, planning software, trade journals and industry associations. Helpful information exists so you don't have to let #3 hold you back.

4) It's going to create more work for me.

Yes. Or no. Actually, disaster recovery can be as much or as little hands-on work as you like. Of course, if you take on everything in-house, you're going to have to allocate significant IT staff and budget resources to this effort. However, there is no shortage of third-party vendors providing managed solutions these days, and you can look to one of them to offload selected areas of responsibility. This option takes advantage of the fact that they've already built the infrastructure, installed the systems, devices and applications, and trained the specialists who keep everything running.

5) I don't know where to start.

There is a lot to consider (and we'll cover the details more fully in Chapter 9). Do you analyze technical or business needs first? Maybe you should look at the network and security. What do you do about critical applications and how do you decide which are really the most critical? Which ones require high availability? What kinds of interruptions are likely to occur?

Luckily for you, disaster recovery is a well thought-out science. If you're unsure of where to begin, simply look where others have gone before you and benefit from their experience.

There are advanced planning tools and sound methodologies to use when developing, implementing and validating business continuity plans. They come with templates, checklists and step-by-step instructions that guide you along the way. And they've been used by all sizes and types of organizations so they've been recovery-tested and proven to work.

You can also tap into consulting expertise to work side-by-side with you. Outside professionals can objectively analyze what you have in your technical and business areas as well as what you need. They can then help develop thorough, multi-level plans that take into account your unique requirements and prepare you for interruptions ranging from simple outages to worst-case disasters. We hope we've adequately addressed the basic questions or concerns you may have when it comes to committing to a disaster recovery strategy.

And that's very important, especially in light of an issue that has become a heavyweight these days: compliance. It has become a very compelling reason for organizations to implement DR planning.

To buttress that argument, consider that CIO Update, in citing an AMR Research study, predicted that the cost of compliance through the year 2010 will reach the \$80 billion mark, and that organizations would "spend close to \$15.5 billion on compliance-related activities in 2005."

With corporate accounting scandals, as well as new laws and industry directives affecting how you plan to protect the safety, integrity and privacy of information, compliance is becoming a force that is driving many organizations to begin to take a critical look at their disaster preparedness.

Chapter 4

Mastering Recovery: Someone's Looking Over Your Shoulder

GLBA. HIPAA. SOA. NYSE 446. COOP and COG.

If you're familiar with any or all of these terms then you probably know that you need to look at regulatory compliance as part of your disaster recovery planning. In fact, depending on your industry, size, type of business, even your value chain, you're required to meet certain provisions contained in one or more of these— and maybe other— directives.

Compliance is certainly not a new issue, but has become more stringent and more complicated in recent years, thanks to developments such as increasing technology dependence, multi-industry deregulation, business consolidation and the impact of globalization.

We can also point to the increasing number and types of today's threats as another factor in the compliance equation. While we still encounter predictable, quantifiable disasters like fires, floods and power outages, we are now experiencing threats that are more random and often malicious in nature, such as cyber-crime, acts of terror, and threats to human capital and the public infrastructure. Significant events in recent history, such as 9/11, the August 2003 blackout and the 2005 personal information breaches that have made headlines, simply underscore this situation.

With these issues converging, it's no small wonder that regulations have grown in number and complexity, and that penalties for non-compliance have increased. Consider the following:

- five Wall Street firms were levied fines totaling more than \$8 million for e-mail retention infractions;
- a major securities lending firm incurred a \$10 million penalty for document production failures;
- the boards of directors of two very prominent companies involved in accounting scandals faced personal liability as a result of their actions.

In addition to the financial consequences, these organizations also met with widespread negative publicity as well as a loss of confidence among their customers, shareholders, employees and the general public.

It is becoming readily apparent that taking a well-defined compliance posture is simply the cost of doing business today. But the quest is not an easy one. Just when you think you've managed to effectively deal with one compliance regulation, it seems as though another materializes.

So what do you do to work toward compliance and incorporate it within your disaster recovery and business continuity plan?

Industry experts point to three broad areas that you must address: reliability, accountability and privacy. While we will address them in this chapter, please note that we do not offer legal advice and do not include the specifics of what your company needs to do to achieve

compliance. We are providing general recommendations, food for thought if you will, to help you begin to tackle this formidable challenge in your planning. The rest will be up to you.

Reliability

As a major issue to focus on, reliability translates into having the planning in place to make sure that your business can continue to operate despite an interruption. Not only does the planning need to cover technology resources, it also must include people and processes.

Consultants have identified key elements for planning that are based on current compliance directives from sources such as the Interagency White Paper on Clearing & Setting and the SEC & NASD/NYSE policies for security firms, among others.

The list is comprehensive and may include components that surprise you. Of course, it includes elements like data backup, recovering mission-critical systems, and financial and operational assessments. But it also mentions alternate communications between customers and the organization, as well as employees and the organization; alternate physical location of employees; the impact on critical business constituents, banks and other third parties; and regulatory reporting.

Accountability

Your second challenge to contend with in DR planning for compliance is accountability, which boils down to making sure that company records and documents are properly maintained and accessible, as needed.

Accountability has been mandated in the financial sector for a number of years but has recently become an issue for all types of organizations—including those in the healthcare, education and defense industries, as well as for public companies. This is because of the widespread transition from legacy, manual and paper-based processes to digital and electronic systems, which has changed the guidelines and best practices for records management.

To begin to work toward a compliant system of records management and retention, you need to take action now. Experts will advise you that the first phase involves reviewing your DR plan to make sure it

supports effective, efficient information management throughout its lifecycle.

Beyond the more obvious areas of data continuity and recovery, you also need to examine new areas entering the fray, notably e-mails and instant messages (IMs). Last, you must define rules, strategies and evaluations for records retention and make sure your infrastructure—both the technical and non-technical components—can support them.

Privacy

A third area of compliance to address is privacy, since a number of today's regulations have provisions to ensure the confidentiality of an individual's personal information. This holds true for many industries. In the case of healthcare organizations, for instance, there are both civil and criminal penalties to be imposed if a covered healthcare entity—or one of its business associates that has access to sensitive data—fails to safeguard a patient's confidential information.

You'll find that consulting organizations have best-practices guidelines to help you in this mission. Among the first steps they'll recommend is to make sure that your current disaster recovery plans account for all technology within the organization, right down to the end-user. Second, professionals will advise you to regularly investigate and review the privacy and security guidelines that could affect your organization, and to update policies and procedures to keep up with changes. Finally, their recommendations will include incorporating change management into your strategies and plans so that they stay current with changes in your organization.

In light of all these compelling reasons to move forward with DR planning, what do you, as an IT manager, actually do to help maintain the continuity of your technical and business resources across the enterprise? You take a good, hard look at the options available and make some informed decisions on what will work best within your organization. The next few chapters will help shed some light on this task.

Chapter 5

Mastering Recovery: Review Your Options: Tape-based Strategies

Not too many years ago, disaster recovery options were like Hobson's choice:

Thomas Hobson was the keeper of a livery stable who required customers either to take a horse from the nearest stable door or none at all. In other words, Hobson's choice was no choice at all. Likewise, in disaster recovery you either had a reciprocal agreement, a hotsite contract, or no backup at all.

Today, the story is different, with a number of backup alternatives available:

- Internal reciprocal agreements between departments or data centers
- External reciprocal backup agreements
- Facilities purchase at time of disaster
- Redundant standby facilities
- Commercial hotsite, warmsite or coldsite
- “Quick Shipment” of critical equipment
- Mobile site
- Replacement equipment

Once you get into the planning stage of your disaster preparation, you’ll be able to make an informed decision about which is right for your operation. During the plan development stage, you will evaluate your assets, your critical applications, and your overall recovery needs, and this will help you narrow the list of recovery options. If the cost of duplicate systems and locations is out of your reach—as it is for many users—then you will probably turn to a collection of commercial disaster recovery options.

Top commercial disaster recovery vendors will offer a variety of choices beyond basic recovery facilities. A few have mobile recovery units, which are fully equipped hotsites on wheels that can drive to your site of choice and get you up and running. Some also might offer “business recovery centers,” ready-made workspaces with telecommunications equipment, LANs, PCs, terminals and other business basics that let displaced end users set up shop and connect to a distant hotsite.

As you can see, there are many choices for **how you recover** from a disaster. And increasingly, there are also significant options for **how you prepare to recover** from one: specifically, **how you back up** your data.

In the relatively short history of disaster recovery, users at one time had the single option of backing up their data to tapes and then physically shipping those tapes to an off-site storage company for safekeeping. If a disaster occurred, users then had to ship the tapes

to a recovery facility, reload their system(s), and restore applications and data from the backup point before they could get up and running. The process could take anywhere from 24-48 hours—or longer—before a company could restore its production environment and get back to business.

Of course, back then, it was the only game in town, so it had to suffice. (And it certainly was a cut above no back-up at all.) However, end-of-day tape backups came with many drawbacks. Consider that:

- tapes get damaged and data becomes corrupt or unusable;
- there is limited accessibility to stored tapes;
- control over data is given up, leading to security issues;
- tapes get lost in transit;
- human errors enter the picture;
- there are a day's worth of lost transactions between the points of backup and failure;
- on-site staff is required for time-consuming nightly backups;
- the organization loses time in getting tapes shipped to recovery site;
- recovery times are delayed;
- there are hardware and software costs involved; and
- there are ever-increasing amounts of data to back up.

The list goes on. Still, some companies, particularly small to mid-sized businesses, find that it works just fine for them. And there have been improvements to the process.

For instance, some vendors allow users to store standby copies of the operating system or distributed systems at their data center. It then becomes a matter of transferring operations to the recovery site and returning to normal operations in a quicker fashion. Still, this option retains many of the problems inherent in the tape medium itself.

Other options have entered the picture to eliminate many of tape's drawbacks, enabling a quicker, more efficient recovery. Electronic vaulting is one such alternative. Let's take a look at this technology and how it offers benefits over tape-based backup methods.

Chapter 6

Mastering Recovery: Review Your Options: Electronic Backup Strategies

Electronic vaulting is another data backup approach that's become more widely used, largely due to its advantages over the traditional tape methods. Many organizations are finding it to be a reliable way to preserve data, regardless of where it is located.

Essentially, vaulting provides automated backup for capturing data at scheduled frequencies and times, even if a customer's servers are located outside the data center. And it does so without the need for tapes or IT staff involvement, so that servers in remote offices, without onsite technical personnel, get the same level of protection as those in the primary data center.

How does it work? In an elementary example, the vaulting software captures changed file or database information, then compresses and transmits this data to a remote vaulting facility. While data is typically transferred over the Internet, it also can travel over a dedicated communications circuit to utilize higher bandwidth connections, which can handle greater data volumes. Some vendors also offer optional encryption of data before it is transmitted, to ensure privacy and security.

The vaulted data is stored offsite, preferably in highly secure Tier 1 data centers. Some providers will also save the backups to tape and store them at a second secure facility, as an extra precautionary measure.

If and when an event threatens the production environment, the organization has accurate information, updated to the point of its last backup, readily available to be retrieved and restored at the data facility.

What are the advantages of electronic vaulting over tape backup? There are plenty. Vaulting:

- simplifies the logistics of backup;
- reduces the chance for human error;
- eliminates ongoing costs of hardware, software and media required for data backup;
- reduces the bandwidth and time needed to back up and retrieve data;
- enables quicker recovery of corrupted or lost information;
- avoids time-consuming retrieval and shipment of backup data during a disaster;
- automates and improves access to data across the enterprise;
- reduces on-site labor costs;
- shortens the recovery window; and
- reduces downtime and its costs.

You may be wondering why more organizations don't turn to this technology, given the many benefits it offers. Note that there are some drawbacks to the technology.

In the past, the technology was relatively expensive, putting it out of the reach of many users. However, pricing has dropped in recent years, thanks to lower telecommunications costs.

And while this technology offers backups that are closer to the point of failure and, therefore, provides faster access and recovery, there are changes or transactions that occur between the regular backups. As a result, data is not updated in real-time or near real-time.

With this approach, when you experience a disaster, you have to re-create those transactions posted since the most recent transmission. As a result, you might lose several hours of transactions—which is still far better than a full day's worth.

There are organizations, however, that can't afford that lag: online retailers, healthcare providers, government agencies and financial services organizations are some that come to mind. To meet their business and regulatory requirements, they need backed up data as well as real-time, or near real-time, data availability in case of an interruption. And for them, there's another solution—replication services. The next chapter explores this technology, as well as its pros and cons.

But before we move onto replication, note that there are other recovery offerings among your choices, chiefly networking. Some vendors can provide advanced capabilities in this area, an important consideration. Chances are, one aspect or another of your business depends on data or telecommunications for day-to-day operations. Internet usage in particular, whether it's an e-commerce site, an e-mail system or a Web-based application, requires reliable connectivity, as do any voice communications such as phone systems or automated call centers.

For virtually all organizations, it's far too expensive and cumbersome to have a private backup network at the ready. This is one of the situations where looking to a recovery vendor is almost a necessity. These providers have invested considerable time and money in designing, building and maintaining advanced networking infrastructure that is redundant, diverse and recoverable. And most can also offer the specialized network configurations users need, whether it's a dedicated, point-to-point, frame relay/ATM, Internet VPN or private MPLS.

Chapter 7

Mastering Recovery: Review Your Options: Advanced Recovery Strategies

SunGard maximizes application and data availability

For some organizations, operating in a recovery mode simply doesn't work, particularly if the cost and consequences of downtime are unacceptable. Brokerage firms with online trading, hospitals with critical-care delivery systems, telecommunications companies with vital network connections—none can afford to be long without their systems, applications and data.

Take a look at the mission-critical applications that run your business. Will losing one or more of them, even for a few seconds, result in real losses? Maybe it's call-center or tech-support lines that go dead,

patient records that can't be accessed in an emergency, financial reporting that won't make a deadline... you begin to see why you have to keep them available and accessible all the time. Today, however, there are advanced recovery services that can help you ensure availability across the enterprise, for your systems, applications and data. Through these services, you can compress recovery timeframes; keep information available, updated and protected; maintain applications and work processes; and make sure that end-users stay connected, despite disruptions within the production environment.

These advanced offerings go beyond traditional methods, leveraging leading technologies for solutions that help you protect data and applications, while maximizing their availability and performance.

Storage Replication

One such technology is storage replication, the process of copying, or mirroring, data from a source to a target. It's finding favor as a quick, reliable solution for organizations that need to keep their information up-to-the-minute.

Replication services help provide critical data, regardless of the host environment, instantly or in near real time. In a replication scenario, your information is stored in a local storage array, the source. Specialized replication software and other components work together to continuously duplicate this information, independent of the host server, at a target-storage replication site, transmitting via a high-bandwidth connection.

This target-storage subsystem can quickly be connected to an available server within the target-storage facility in case of an interruption at your site, to maintain continued production service levels. Production at your facility can resume almost instantaneously—avoiding the time delays, manual processes, data integrity and security issues that come with other methods.

A storage replication solution is more expensive than simple tape backup offerings, with costs depending on the levels of availability required and the complexity of the technical infrastructure. But the payoff is undeniable, as it can:

- minimize data loss for mission-critical applications inherent in unplanned outages;

- reduce RTOs to a few hours, even a few minutes in some cases;
- support multiple mirroring options to meet varying data-protection requirements and distances between source and target storage systems;
- avoid unplanned demands on IT staff; and
- simplify recovery logistics.

Server Replication

Server replication is another service that helps ensure the continuity of your operation. With this type of offering, you can maintain the availability of mission-critical data and applications, through a fail-over database server that's equipped to pick up the processing load should an interruption occur. A server replication solution comprises two essential components, a source (also known as a publisher), which is a server that makes data available to other servers for replication; and a target (or subscriber), the server that receives updated transactions and data from the source. These servers can be connected via LANs, WANs or the Internet.

Broadly, vendors offer three types of replication services today. While the specific functionality of each of these three options may vary, their essential capabilities are as follows:

Snapshot replication is the simplest type of offering. As the name implies, the source takes a snapshot of all data at regular intervals and sends it to the target. Snapshot replication is best used as a method for replicating data that changes infrequently as well as for small volumes that need to be replicated.

Transaction replication, better suited for replicated data that changes frequently, uses an intelligent agent to continuously monitor the source for updates and then transmits them to the target.

Merge replication, the most complex of the three, allows both target and source to make changes to the database, regardless of whether or not they are connected. Once they are re-connected, an intelligent replication agent checks both sides for differences and synchronizes the data as needed, reconciling any conflicts using a predetermined set of rules.

Some server replication offerings can initiate fail-over capabilities that are automatically or manually activated to provide seamless continuity, as well as the ability to failback from the source to master server, again with no loss of availability.

Organizations that opt for storage or server replication services will find that they:

- minimize data loss;
- minimize application recovery time—especially important for databases, e-mail and file servers where a 24-hour window is unacceptable;
- reduce RTO to fewer than 30 minutes when using failover;
- speed up recovery and testing processes—utilizing dedicated or shared servers;
- help reduce the chance of regulatory non-compliance and its financial penalties;
- retain the cost benefits of hotsite processors, peripherals and networks for applications not requiring replication for their recovery time-frame;
- minimize internal IT support responsibilities; and
- enhance load-balancing capabilities.

Of course, we've reduced these technologies to their essence; there is much more than the elemental descriptions we've outlined. This chapter could also go on for several more pages discussing other DR options for your organization. It's best is to consult with a third-party provider and learn more about those that are best-suited for your recovery needs.

For this endeavor, there's no time like the present. It seems like each week brings news of new threats to your systems and data, threats that could cause everything to come to a grinding halt.

Chapter 8

Master Recovery: You Can't Be Too Careful

40 Mastering Recovery 41 Just because you think everyone's out to get you doesn't mean they're not. These days, you can pretty much count on the fact that someone or something is testing the limits of your security measures—and trying to break through. As your dependence on systems and data grows, so do the threats to both.

As a matter of fact, the security landscape is changing quite a bit, becoming more complicated and more sinister as Internet use and misuse pervade everyone's lives. We still have the worms, viruses and Trojan horses of the past out there. And unless your security software and patches are updated, they still remain a threat.

But the old-school, run-of-the-mill hackers, who caused random havoc for fun and bragging rights, are being replaced. Now, there are cyber-criminals, whose motive is taking someone's personal information and, ultimately, their money. Their attacks are deeper, more sophisticated and more focused, and they can come from outside the organization as well as from within.

In the quest for securing systems and data, we've entered into an arms race of sorts, with the bad guys trying to develop bigger and better offensive weapons, and the good guys just as quickly trying to come up with ways to defend against them.

The bottom line of all of this is that companies are losing money, real money, through both fraud and through having to spend heavily to defend themselves.

What needs to be done? Security consultants, particularly those designated Certified Information Systems Security Professionals (CISSPs), will tell you that assessing both your IT and business resources, and determining where they are exposed to risk from both the physical and cyber varieties, is the place to start.

As part of this, it's important to analyze your current environment—technical and non-technical—as well as your security planning efforts. The findings will serve as a baseline to benchmark your existing procedures against industry standards and those used by similar-sized organizations. Essentially, you need to see where you stack up against leading practices.

Your goal is to locate your weaknesses, plug the holes and eliminate vulnerabilities so that you bolster security, including the privacy, safety, reliability and integrity of data and systems, throughout the enterprise.

And you need to establish a defense-in-depth strategy, with security experts recommending a firewall as the place to start. IT experts know that a firewall is hardware or software, or a combination of the two, that prevents access to or from the network. Most often, it's used to stop unauthorized Internet users from getting into private networks that are Web-connected. It's a great first line of defense, but more is needed.

For instance, security professionals can help you set up an intrusion detection service that proactively monitors and identifies hacking and other suspicious activities on-line and in real-time, extending the protection of the firewall. This type of offering provides alerts when someone tries to breach your system and blocks threats as they become known. As a security strategy, it can really enhance the integrity of your information with high levels of authentication, access control, and confidentiality.

Other offerings can block threats, such as spam or viruses, at the perimeter level. They help keep these unwanted e-mails away from your mailboxes and systems, before they have a chance to cause performance problems.

Of course, you may have some budget concerns about these, and other types of defensive measures that may be recommended. From experience, we can say that when it comes to safeguarding your technical assets, it's unequivocally true that an ounce of prevention is well worth the pound of cure.

Chapter 9

Mastering Recovery: Starting at the Very Beginning

Years of DR-industry experience show unequivocally that:

Companies Fail to Keep Operations Running Only When They Fail to Plan

Yet planning how you will recover from a disaster may appear to be a huge sprawling undertaking. You'll have to cover at least your computing operations, but more likely your entire business. Planning for distributed processing is one thing, but what about planning for work group recovery, for WANs, LANs, clientservers, and PCs? How about alternate sites and office equipment? You have an information technology network that's scattered across your organization. You have small platforms that need to be backed up. You have users who need to access alternate computers.

Consultants will tell you that in the last decade, nearly all their business was planning for computer systems. Nowadays, however, you can't separate business recovery planning from computer recovery planning. Indeed computing has spread the length and breadth of the organization, but the business operations folks simply don't have the mindset or the knowledge to put together a disaster recovery program.

There are thousands of considerations, large and small, and the degree of thoroughness with which you address them will determine the ultimate success of your plan.

Feeling daunted? You probably should. It's a big job, but it's certainly not an insurmountable job. Don't let the challenge paralyze you. You can take comfort from the fact that thousands have gone before you and that they have left a pretty good trail.

Step #1

Where do you start?

You can start by taking the lay of the land. Conduct a business impact analysis; develop your "what if" scenario. An in-depth business impact analysis will help you pinpoint the areas that would suffer the greatest financial or operational pain in the event of a disaster. And the analysis will identify the resources already in place that would help mitigate that pain.

To help you conduct a meaningful business impact analysis, don't hesitate to call on consultants. Consultants do impact analyses for a living and they have more experience. As we've hinted before, there's no need to re-invent the wheel. In your analysis, you can start with lost revenues. Other surveys of large disaster recovery users showed that the average revenue impact for a system shutdown is roughly \$6 million an hour. That's not chump change in anyone's budget. For most operations—even very big ones—that's real money.

These surveys went on to estimate that companies lose one percent of the market share for every eight hours they can't ship because of a system failure. They also calculated that it would take three years to recover half a percent loss of the market share. That's roughly one year of consequences for every six hours of downtime.

These types of findings are important to your objective because unless senior management is aware of the real impact of a shutdown, they have no reasonable basis on which to make funding decisions to support recovery strategies.

Step #2

Secure senior management commitment

We've said it before and we'll say it again. The reality is that until the person sitting in the corner office says "go," you're not going anywhere with your disaster recovery planning. In fact, the number one agenda item for disaster recovery coordinators is securing management's commitment to a robust and realistic business continuity plan. One DR coordinator put it this way:

"Developing and maintaining an effective crisis management/emergency response program is like trying to light wet wood—it can only be done if you have a steady and hot flame. Senior management has to be that flame— no substitute will do!"

It's always been that way, and probably always will be. So, how should you approach senior management? It may help you to know the things that keep you awake some nights are some of the same things that keep senior management awake some nights.

According to one disaster recovery survey, these are 10 of the most frequently cited reasons for senior management insomnia:

1. Will disaster recovery capabilities keep up with raw growth?
2. Can we balance the need for cost reductions and the overhead requirements of ongoing business, including disaster recovery?
3. Do we have the people—available and skilled enough—to recover from a disaster? In a regional disaster, people are concerned about their families; they just may not be available to work the recovery.
4. Can we find the people to recover our mainframes? In the mainframe area, knowledgeable people are "scarcer than hen's teeth."
5. What about the project management team? What if they're so stretched on a day-to-day basis that they can't get the adrenaline pumping to do the job in a recovery mode?

6. Disaster plans are in place, but what happens when agendas get heavy? Disaster recovery tends to fall off the plate.
7. Disaster recovery has to be done, but who's going to do it? It's non-revenue producing. Can we outsource some parts, but not others?
8. What if some non-mainframe systems are carrying key information, but they have no plan?
9. How do we meet the needs of customers who have an absolute dependency on delivery of the product and service—mission-critical applications like telecommunications services or equipment—and therefore the need for uninterrupted service?
10. In the event of multiple disasters or a large geographic disaster, will disaster recovery providers have the capacity to handle us? How much capacity is enough?

These concerns don't make your job any easier, but at least you know that you're thinking along some of the same lines as senior management. Once you secure senior management commitment, you're ready for the third step in disaster recovery planning: Defining your MARC.

Step #3

Define your MARC

Your MARC—Minimum Acceptable Recovery Configuration—applies to computer equipment, communications support, furniture, fixtures; the whole shooting match.

Fortunately, there are consultants and even commercial software packages out there to help walk you through all of these considerations. The top comprehensive business recovery planning software offers such indispensable features as:

- Standardized nomenclature with pick lists for data entry;
- MARC;
- Spreadsheet functionality;
- True WYSIWYG;
- CBR (content-based retrieval) librarian;
- Full incident-management planning; and
- Wizards to assist users.

Step #4

Observe the three C's: Complete, Comprehensive and Current

By complete, we mean it must be detailed enough—it must spell out each and every recovery step—so that the plan can be followed just by reading the playbook. It's a mistake to rely on the availability of people familiar with your operation. When a real disaster strikes, they may not be around to help. The plan must stand on its own.

The plan must also be comprehensive, so that it covers all critical business units that require technical support to function. In other words, your plan has to cover everything inside the glass house, as well as out.

Finally, you must keep the plan current. If your company is changing quickly, if it's downsizing, upsizing, merging, or acquiring, if it's buying new platforms, adding client servers or new employees, expanding LANs, WANs, and so on, all of those changes must be reflected in your plan. If not, your plan is for a configuration and a company that no longer exists.

Chapter 10

Mastering Recovery: Make the Grade with Testing

Testing is where the rubber meets the road. Virtually any recovery provider worth its salt will tell you to develop a testing schedule, to be executed annually. You need to know if your plan will work before the disaster strikes. By testing, you can prove to yourself that every facet of your recovery plan is well thought out and buttoned up. Accomplishing that, though, isn't as simple as it may sound.

In today's technical world, with IT staffs pushed to the limits and budgets squeezed dry, many organizations find it difficult to muster the resources and time for traveling to a recovery site for testing. One vendor reported that nearly 40% of the tests scheduled at its hot site have to be re-scheduled. Biggest reason: staff availability.

Fortunately, DR vendors have seen the writing on the wall and have responded by adding more flexible testing options. (You may recall some of them listed as recovery alternatives in Chapter 5). Today, you have choices at your disposal that include:

- **Hotsite:** Traveling to a disaster recovery provider's computer center to conduct a test, using your own staff in cooperation with the provider's support staff.
- **Remote Site:** This is a conveniently located site set up for a remote link to the command center at a hot site. The disaster recovery provider's support staff act as your "hands" at the hot site.
- **Remote Testing:** This form of test can be conducted from your site with a laptop, a modem, and communications software. Remote testing not only eliminates travel expense, but also keeps your staff in-house.
- **Turnkey Testing:** The disaster recovery industry's leading provider offers a testing solution that frees your staff to stay at their normal jobs, yet still have the advantage of a complete and aggressive test. The provider performs an entire disaster recovery test for you—from loading the system and starting applications to performing basic network tests—and then forwards the test results to you.

Remote recovery testing of all kinds will undoubtedly become much more automated and much simpler as technologies improve, new tools are created and as telecommunications costs continue to decline. By the end of the century, perhaps all tests will be conducted remotely. Perhaps you'll never see the inside of a hot site again. Now that's something to look forward to!

By this point, we hope you're convinced that you need a sound, tested disaster recovery plan in place, ready to implement at a moment's notice. But where do you go from here? It's time to evaluate the DR providers and find the one that's best suited for your needs.

Chapter 11

Mastering Recovery: Choosing the Right Provider

Engaging a disaster recovery provider is a considered purchase, one to be made with care. After all, you're putting your organization's survival into someone else's hands. There are the standard requirements to look at, to be sure, but beyond that there are some that are hard to measure—do you get along, do you share a single philosophy, and the like.

However difficult it may be to measure, the disaster recovery philosophy may be the best place to start. Your principle goal in the event of a disaster is to avoid having systems go down in the first place but, if that happens, it's to recover operations quickly and seamlessly. To do that—and to do it as efficiently as possible—you're going to need a disaster recovery provider who shares a similar goal. When you begin studying disaster recovery vendors, you'll probably find more than one—perhaps even several—that meet your hardware requirements.

Of these, two or three may also meet your requirements for price and location. You still need to narrow down the providers, though, so this is the place to ask:

Which provider shares my principle goal? Which can bring total solutions to my business? Which can offer turnkey testing? Which is focused? Fast? Flexible?

If there is an ideal disaster recovery provider, that provider must meet all of those criteria. Can it offer a testing solution that frees your staff to stay at their normal jobs, yet still have the advantage of a complete and aggressive test? Can that provider perform an entire disaster recovery test for you? Is that provider focused on disaster recovery? Remember, if disaster recovery support service is a sideline, it will be treated as a sideline, and when disaster strikes, your provider will have a principle goal that is different from your own. Disaster recovery must be at the top of your provider's page.

There are other important factors to consider in your search. In fact, there are five failsafe criteria you should bear in mind to help narrow your choices:

1) Highest levels of availability. When you need to maintain continuous operations, no matter what, you need to be especially careful in your choice of vendor. Quite a few don't have the breadth of offerings or the advanced technologies to provide you with real-time or near real-time availability.

2) Flexible user-control. The perfect DR service provider is a cross between an outsourcer and a managed services company. It lets you keep as much control as you want and doesn't ask you to change to fit its systems or infrastructure. Whether you're a large company with a skilled IT staff, ample budget and robust infrastructure, or a small business strapped on all fronts with only the most basic resources, the right vendor has all the services you need but lets you decide which, at what levels, are best.

3) Platform-independence. Your ideal provider has the systems and infrastructure, as well as the technical expertise, to match whatever you're running throughout your enterprise. It's tied to no vendor so you can keep working with what works best for you. Ten years ago, one of the industry's leading providers supported only one manufacturer's

mainframe. Today, that same provider supports more than 20 hardware platforms in all sizes and configurations.

4) Enterprise-wide offerings. It may seem like a no-brainer, but you need a provider that covers all your requirements across the enterprise: infrastructure, people, data, networking and platforms. Traditional DR companies, outsourcers and other vendors have some of the pieces in place, but not all. You can't, however, achieve Information Availability without a complete solution that encompasses everything.

5) Low TCO (Total Cost of Ownership). We've covered the fact that doing it yourself is an expensive proposition, even if it does offer some benefits. You will find that most third-party vendors have the resources in place to help reduce your expenses. You'll simply have to weigh their strengths in other key areas when making your decision.

Finally, you'll want to think about some tried-and-true criteria for making any vendor selection by checking on their experience. How long have they been in business? How many recoveries have they supported? What's their success rate? Do they have experience in my industry? Remember: it's no longer simply recovering from disaster—it's about helping to keep your business in business in the face of any interruption.

Chapter 12

Mastering Recovery: Who We Are

Information Availability

At SunGard Availability Services, our mission is keeping people and information connected despite situations that threaten to interrupt their place of business. Since our early days as the trailblazer in the disaster recovery industry, we have evolved to become a leader in Information Availability.

We understand that your Information Availability strategy must provide for the protection of information, as well as the ability for individuals to access that information. Utilizing people, technology and infrastructure, we help organizations design a strategy and create the solutions to provide the access to information their business requires.

Our services are broader than ever, from assessing availability requirements to delivering Managed IT, Professional and Disaster Recovery Services. It's this mix that provides for a true Information Availability solution designed to meet your organization's needs.

With our global reach and through our worldwide network of hardened, state-of-the-art facilities, SunGard Availability Services:

- helps over 10,000 customers worldwide achieve uninterrupted access to their mission-critical data and systems
- employs more than 2,500 professionals with extensive Information Availability experience
- supports all major platforms, including IBM mainframe, AS/400, iSeries, System/3x, Data General, DEC, Hewlett-Packard, RS/6000, pSeries, Sequent, Sequoia, Stratus, Sun Microsystems, NCR, Tandem, Unisys, Prime, Texas Instruments, and others.
- provides 4 million square feet of secure operations space
- operates data centers and end-user recovery centers in over 60 facilities in North America and Europe, including Atlanta, Boston, Chicago, Cleveland, Dallas, Detroit, Denver, New York, Northern New Jersey, Northern Virginia, Orlando, Philadelphia, Phoenix, Scottsdale, St. Louis, San Diego and Seattle
- has 41 mobile recovery units staged in strategic locations
- manages over 25,000 miles of dedicated network backbone
- has successfully supported more than 100,000 tests and more than 2,300 recoveries
- assists customers with approximately 9,000 tests each year
- has written more than 100,000 business continuity plans
- has business continuity consultants working over 350,000 hours annually

Backed by incredible resources, we feel confident in saying that SunGard is the only Information Availability provider that can provide **comprehensive support for your infrastructure, people and data** to cover every phase of your Information Availability needs.

Appendix

Mastering Recovery: Get Started Today

“Show me the money,” you may say. Fair enough. We have something that can do just that. Just use this link:

<http://www.availability.sungard.com/getstarted>

to take advantage of an online benchmarking tool and estimation form that will help you realize just how much you can lower your Total Cost of Ownership (TCO) with SunGard Availability Services.

Notes:

About SunGard Availability Services

SunGard Availability Services provides disaster recovery services, managed IT services, information availability consulting services and business continuity management software to more than 10,000 customers in North America and Europe. With four million square feet of datacenter and operations space, SunGard assists IT organizations across virtually all industry and government sectors prepare for and recover from emergencies by helping them minimize their computer downtime and optimize their uptime. Through direct sales and channel partners, we help organizations ensure their people and customers have uninterrupted access to the information systems they need in order to do business. To learn more, visit www.availability.sungard.com or call 1-800-468-7483.

For more information about
SunGard Availability Services, visit
www.availability.sungard.com or
call 1-800-468-7483.

SUNGARD® | Keeping People
Availability Services | and Information
Connected.®

680 East Swedesford Road | Wayne, PA 19087
800-468-7483 | www.availability.sungard.com