



## **BUSINESS CONTINUITY** **PLANNING GUIDE**

---

Written by Steven Shepard, President,  
Shepard Communications Group, LLC  
Commissioned by XO Communications

# XO Communications

# BUSINESS CONTINUITY PLANNING GUIDE

---

## Table of Contents

---

<b>Summary of Guide</b>	<b>4</b>
<b>Why Business Continuity Is Needed</b>	<b>4</b>
Sobering Facts	5
<b>Elements of Success:</b>	<b>5</b>
<b>Determining Resiliency Requirements</b>	
Identification of Mission-Critical Systems	6
Prioritizing Services into Risk Tolerance Tiers	6
Conducting a Systematic Risk Assessment	7
Elements of the Plan	7
Identifying Points of Failure	7
Assessing the Pros and Cons of a Business Continuity Partner	8
Selecting a Vendor Partner	8
#1: Breadth and Richness of Services	9
#2: Solution Functionality	9
#3: Financial Viability	9
#4: Cost	9
Project Assessment	10
<b>Summary of Recommendations</b>	<b>10</b>
XO Communications as a Business Continuity Partner	11
<b>Conclusion</b>	<b>11</b>

## Summary of Guide

Business continuity, also referred to as disaster recovery, is one of the most-talked about – and worrisome – topics in industry today. Whether you are a five-person small business determined to protect the integrity of the day's receipts, or a multinational corporation concerned about Sarbanes-Oxley compliance, shareholder comfort and revenue assurance, the issues are similar.

The factors that lead to the disruption of business activities – disasters, as they are often called – fall into two general categories: natural disasters, which include uncontrolled flooding, fire, earthquakes, hurricanes, typhoons, mudslides, disease, lightning strikes, and biological infestations; and man-made disasters, which include disruptions brought on as the result of human error, strikes, intentional sabotage, burglary, theft, IT-related faults stemming from viruses, Trojan horses, worms, denial-of-service attacks, and other breaches of computer and network security; and increasingly, disruption brought about by terrorist acts.

For the most part, natural disasters are unavoidable; dealing with them tends to be a matter of preparedness and recovery as well as an avoidance strategy. Man-made disasters, on the other hand, can for the most part be avoided (or their impact minimized) through prudent IT management practices and adherence to a well-designed Business Continuity Plan. In aviation, the phrase that is drilled into every pilot's head as their operating mantra is "Always plan your flight, and always fly your plan." The same rules apply in business continuity: Build a plan, and stick to it. Your discipline will be well-rewarded.

This paper is a practical guide for businesses concerned with ensuring their own continuity in the event that an unavoidable disruption challenges their ability to operate in a business-as-usual fashion. We begin with a rather sobering discussion about the need for business continuity planning, followed by an in-depth section about the steps involved in creating a plan: identifying mission-critical systems; conducting risk assessments; creating "strata" of risk tolerance for business systems; locating single points of failure; selecting partners for business continuity resources; and finally, identifying the steps required to ensure effective business process continuity and data recovery.

## Why Business Continuity Is Needed

It was a business-as-usual day in the switching center. The building served well over 100,000 subscribers, six hospitals, 11 firehouses, three post offices, a police station, nine schools and three universities. Like most central offices, this one was completely invisible to the public. It was just a big, windowless structure.

Just after midnight, a relatively inconsequential piece of power equipment in the cable vault shorted and spit a few sparks into the air. One of the sparks fell on a piece of insulation and began to smolder. The insulation melted and began to burn, changing from a smoldering spot on the surface of a cable to a full-blown fire.

The fire burned its way up the cables, spreading from floor to floor. Soon the building was engulfed, and the world was about to see the single worst disaster that any American telephone company had ever experienced.

Emergency services vehicles converged on the now evacuated building. They flooded the lower floors with hundreds of thousands of gallons of water, severely damaging whatever equipment remained functional.

Two days later the fire was finally out and engineers were able to enter the building to assess the damage. On the first floor, the 240-foot main distribution frame was reduced to a puddle of iron. Water had ruined the power equipment in the basement. Four switching offices on the lower floors were completely destroyed. Cable distribution ducts were deformed and useless. Carrier equipment on the second floor was destroyed. And 170,000 telephones were out of service.

Within an hour of discovering the fire the company mobilized its forces to restore service to the affected area, and resources converged on the building to coordinate the effort that would last just under a month and cost more than \$90 million.

Within 24 hours, service had been restored to all medical, police and fire facilities. The day after the fire, a new main distribution frame had been located and was shipped to the building. Luckily, the third floor had been vacant and was therefore available as a staging area to assemble and install the 240-foot-long iron frame. Under normal circumstances,

from the time a frame is ordered, shipped and installed in an office, six months elapse. This frame was ordered, shipped, installed, wired and tested - in four days.

It is almost impossible to understand the magnitude of the restoration effort, but this may help. 6,000 tons of debris were removed from the building and 3,000 tons of equipment were installed including 1.2 billion feet of underground wire, 8.6 million feet of frame wire, 525,000 linear feet of exchange cable, and 380 million conductor feet of switchboard cable. 5 million underground splices hooked it all together. And 30 trucking companies, 11 airlines, and 4,000 people were pressed into service. And just after midnight on March 21st, 22 days following the fire, service from the building was restored.

How was this possible? Two factors contributed to the success of the restoral effort. First and foremost, the company had resources it could call on when needed, and the ability to deploy them.

Second, they understood the nature of their business: What was critical and what wasn't. The fact that they restored emergency services first and consumer telephones second said a great deal about the degree to which they had though the scenario through. They knew which services posed the highest risk and dealt with them first.

Business continuity imperatives haven't changed much since this disaster took place; technology has evolved, businesses have become more sophisticated, the global nature of the world's economy now dictates that businesses operate around the clock, but one singular mandate remains the same: be prepared.

## SOBERING FACTS

Every day, 30 billion text messages, 40 billion e-mail messages, 19 billion mobile minutes, 30 billion PSTN minutes travel the world's telecommunications networks. As if that weren't enough, eight exabytes (that's 18 zeroes) of global IP traffic are generated every month, a number that's expected to increase to a zettabyte (21 zeroes) by 2015. The problem with those numbers is not their magnitude: It's the nature of their intent. A large percentage of those messages are destined for enterprise networks, and a small percentage of those aren't friendly – they are intended to disrupt business operations.

Consider these simple facts. The economic impact of a single intrusion, such as a denial-of-service attack, is roughly \$1 million, an expense that could be prevented enterprise-wide through the deployment of a simple firewall. It gets worse: in 2004, five years ago, the average large business spent \$55 million per year to protect itself from viruses, Trojan horses, and worms, \$26 million against denial of service attacks, and a comparatively paltry \$12 million to guard against physical theft . And why are they willing to spend so much on cyber-protection? Because 93% of companies that lose access to the information in their data centers for 10 days or more will file for bankruptcy protection within a year of the loss. Can they afford not to spend the money?

Forrester Research has extensively studied Business Continuity practices and there is much to be learned from them . Most companies recognize the need to improve their preparedness for disruptive events, and when asked why they are doing so the list narrows to six primary drivers: the increasing cost (and impact) of system downtime; the need to be online, available and therefore competitive in an increasingly globalized (and therefore around-the-clock) market; the need to satisfy financial responsibilities; the need to take steps to reduce increasing levels of enterprise risk; the desire to increase the availability (and survivability) of a critical business application; and finally, the requirement to ensure compliance with regulatory or legal mandates such as Sarbanes-Oxley, HIPAA and CALEA.

So why is business continuity planning required? Besides the obvious answer – the need to protect critical business resources in order to protect revenue-generating customer services – it is because “chance favors the prepared mind.” The better prepared an organization is for the possibility of a disruptive event, the more rapidly the event will be resolved.

## Elements of Success: Determining Resiliency Requirements

One element of business continuity planning is disaster recovery preparedness. And while disaster recovery is an important element of any prudent IT strategy, disaster avoidance – the process of taking steps to avoid the impact of a disastrous event when it occurs - is at least as important. Because these disruptions are largely unanticipated, a collection of up-front efforts can reduce the downtime – and cost

– associated with disastrous events. These efforts, which we will describe in the sections that follow, include:

- Identification of mission-critical systems
- Prioritizing services into risk tolerance tiers
- Conducting a systematic risk assessment
- Identifying points of failure
- Assessing the pros and cons of dependency on a single business continuity provider
- Determining the steps involved in selecting a provider

### IDENTIFICATION OF MISSION-CRITICAL SYSTEMS

Mission-critical systems are those resources – hardware, software, operational processes, operating procedures – that are absolutely necessary for the business to operate. Fundamentally speaking, if these resources become unavailable or severely impaired for any reason and for any length of time, business operations are jeopardized. It is critical, therefore, to identify these systems and take steps to ensure that they are appropriately protected.

The first step in the process, therefore, is to identify these systems. In reality, though, this is not the first step – it is the third step. The first step is to map out the day-to-day operations of the business to ensure a clear understanding of exactly how the business operates. This means taking inventory of business processes, and then relating those business processes to the critical systems that underlie them. Sounds silly? Take this simple test: If you were to lose your laptop tomorrow, what would the actual impact on you be? Let's assume that your data is backed up (a major assumption). Do you have a list of the applications that reside on the laptop that you will have to reinstall to restore the machine to full functionality? You probably have the CDs and DVDs of the applications you've purchased in a box somewhere, but what about the installers for the applications you've purchased online? And once they're installed, how long will it take you to download all of the updates required to bring the applications up-to-date? And what about all of the serial numbers and product keys required to securely install the software? If you have to change operating systems because of an update that came out since you bought the laptop, how long will it take you to be comfortable with it to ensure functionality?

We're talking about a single laptop here – not the far more complex processes and resources of an entire business – even a small one. It is fundamentally important, then, to build an “aerial view” of the ebbs and flows of the business, so that all business processes can be identified, evaluated, and ranked according to their perceived levels of criticality to business operations.

Once the processes have been identified, then the underlying systems can be identified as well. This process must include a comprehensive hardware inventory across all data centers, office environments, call centers, operations centers, and any other locations where a significant amount of human-to-machine (and therefore system) interaction takes place. At the same time, an application inventory should be conducted to identify all of the ways that people rely on the hardware to do their jobs. Keep in mind that the hardware is nothing more than a very large heat engine and real estate consumer without applications running on it; all data flows and all application-to-application flow-through should be identified at the same time as part of a comprehensive systematic analysis of critical IT resources.

### PRIORITIZING SERVICES INTO RISK TOLERANCE TIERS

Once the mission-critical systems have been identified, including all hardware, software, processes and procedures, the next step in the business continuity planning process is to rank them according to their relative importance to the business. This is an extraordinarily difficult part of the overall process for one simple reason: it tends to become emotionally charged as different organizations, if given the chance, argue for the relative importance of their critical systems. For this reason alone, it is important that the ranking of these processes be conducted by a team made up of representatives from two organizations: IT and corporate finance. IT must be involved because they understand the inner workings of the systems and can speak to the complex interworking that often exists between seemingly unrelated applications and hardware. Finance must be involved because in the final analysis, the ultimate “tie-breaker” is the degree to which one mission-critical process affects the ability of the corporation to generate revenue, relative to other systems. After all, how do you compare the relative importance of an order tracking application vs. a strategic data collection application used by marketing? Remember,

we're not talking about which one lives and which one dies: We're talking about a hierarchy of recovery to ensure that the applications most necessary for revenue protection and cost control are brought back to life first. As long as this is understood by all concerned, the emotional challenges tend to fade in favor of more rational arguments.

## CONDUCTING A SYSTEMATIC RISK ASSESSMENT

The “tiering” process is usually performed according to a set of standardized measures by which priorities can be assigned. These measures, the Recover Point Objective (RPO) and Recovery Time Objective (RTO), are used to determine the relative value of systems and applications based on their importance to the business and based on a set of agreed-upon performance objectives. The Recovery Point Objective is the specific point in time to which you must recover data in order to resume effective business operations. It does not imply 100% recovery of all data. The RPO is generally considered to be a measure of acceptable data loss following a major disruptive event. RPO is serious business: If your RPO is determined to be four hours but a full recovery will take seven, the last three hours are deemed to fall into the category of acceptable loss and business operations will be resumed from the four-hour recovery point, with attempts to manually recover the lost activity carried out later, as time permits. In the IT world, RPO is the technical equivalent of medical triage – the steps necessary to stabilize the patient.

Closely related to RPO is Recovery Time Objective (RTO). RTO is the amount of time within which a critical business process must be restored following a systemic disruption to avoid unacceptable consequences in terms of the ability to satisfy customer needs and meet critical business objectives. Bear in mind that the RTO is associated with the recovery of critical business processes, not with the recovery of the underlying systems: that's the domain of the RPO. Businesses must address both if they are to develop a holistic plan for assured operational continuity. When a company's Web site is the victim of a denial-of-service attack, and the company relies on their Web site for revenue generation, the recover priority (RTO) is much higher than it would be for a company whose Web site is used as nothing more than an “online business card.”

## Elements of the Plan

The plan, once completed, typically has five elements:

1. An analysis of the business, as described earlier;
2. A description of the various solution elements that have been put into place to deal with the potential for a significant disruption of business operations;
3. A plan for the actual implementation of the identified procedures, should they be needed;
4. A process for routinely testing the procedures to ensure their functionality and currency; and finally,
5. A maintenance program to ensure that the plan is constantly updated according to changes in the business.

One of the main reasons that business continuity plans fail is that they are not routinely updated to model the changing profile of the business, leaving some critical systems, processes or infrastructure inadequately protected.

There's an old saying in business: “That which gets measured gets managed.” RPOs and RTOs are two of the important success metrics of a business continuity plan. Once they are understood, put into place and socialized throughout the business, they must be mapped to the systems and critical infrastructure that they are designed to protect. The next step in the process is to identify failure points.

## IDENTIFYING POINTS OF FAILURE

Every complex system has its weak links, and part of the business continuity planning process is to identify them so that they can either be fixed or appropriately protected. Creating a map of critical process dependencies (described earlier) and then establishing metrics (RPOs and RTOs) represents phase one of this process; phase two is the examination procedure designed to identify areas where attention is required. Areas that tend to be at risk, and that should be carefully examined, include:

- Power feeds to critical buildings: Is the power coming into critical buildings (call centers, data centers, operations centers) fed redundantly and from different points on the grid? Is there uninterruptible power (UPS) in place? Is it tested routinely, and under load?
- High-bandwidth connections to critical buildings: Are the circuits redundant? Is that redundancy accom-

plished over facilities that are physically separated from one another? Is there a high-speed wireless connection that could be used in the event of an event that disrupts the physical media?

- Database hardware (disks, servers): Are the servers, processors and disk arrays that host and make available critical applications configured in such a way that they have “hot spares” available should they be needed? In the event of a major failure of the physical site, is there access to a redundant site to ensure business continuity? Is the actual access redundant, i.e., is there (for example) point-to-point wireless available to ensure connectivity in the event of a disruption of the physical plant?
- Backup processes: Is there a clearly-documented backup strategy that ensures that all critical applications (and the data they create) are archived on a regularly-occurring schedule? Are multiple copies created, and if so, are they stored in physically different locations? Is there a process in place to gain access to backup copies in the event of a business disruption? For many corporations, backup practices rise to levels that appear to be one step removed from organizational paranoia. The need to backup customer data and the applications that depend on that data cannot be overstated, however. Some of the methodologies that companies consider include:
  - Tape backups that are transported to offsite facilities on a daily basis
  - Disk backups that are performed locally but copied to a remote backup facility
  - Real-time data replication to an offsite archive
  - Support for remote workers: In the event that a significant number of remote workers would be affected by a catastrophic system failure, are there alternatives available that they can use to gain access to backup critical systems? Is there adequate security in place to ensure that remote access to sensitive information assets can be granted without threatening those assets?
- Application availability: Are critical applications available from multiple servers? In other words, in the event

that a primary application server should die, is there a backup system in place and online to ensure that business can continue as required?

- System security: Are there procedures in place for auditing security policies and practices? These procedures should include examination of physical security that protects buildings, outside plant, and personnel, as well as logical security practices that protect software and data from worms, viruses, Trojan horses, denial-of-service attacks, malware, spam, social engineering, and other intentionally disruptive attacks.

This is by no means intended to be a complete list, but it does represent a valid sample of the key issues that should be included in the development of a business continuity protection strategy. It can also serve as the basis for identifying weak points in the strategy: If questions arise as the result of the process described above, there are weaknesses in the overall plan that should be examined and rectified.

#### ASSESSING THE PROS AND CONS OF A BUSINESS CONTINUITY PARTNER

Once the metrics have been decided upon, are accepted, and have been logically linked to the critical elements of the corporate information infrastructure, a priority-based strategy for resource recovery can be created, based on the steps described earlier. In large corporations it has been a standard practice to perform most recovery processes internally, using the knowledge resources of the IT or Information Systems organization. Increasingly, however, companies are relying on outsourced business continuity providers to handle this aspect of protecting the business. The primary reasons for this are two-fold: expertise, and strategic diversity. Business continuity organizations have considerable expertise in their field and are therefore valuable partners. And because they offer service diversity in terms of diverse network routing, distributed database backup and hardware redundancy, their role as a value-added provider is a valid one.

#### SELECTING A VENDOR PARTNER

The vendor selection criteria must be carefully crafted to not only ensure that the vendor’s product is up to the task of addressing the needs of the business, but that the vendor

is in a position to provide ongoing, long-term support. This is an important and often misunderstood part of the vendor selection process, and companies have been known to select a vendor based solely on the viability of their technology when in fact the technology is perhaps the least important of all criteria. Some companies, for example, have a policy that dictates that if their company's purchases represent more than 20% of a vendor's annual revenues, they won't buy from that vendor. It doesn't seem to make logical sense until you think about the decision strategically: If your company represents 20% of that company's revenues and you have a bad year and cut spending, you could put the vendor in financial jeopardy, which would make it impossible for them to provide the level of support you require. It is important therefore to consider vendor selection from multiple, seemingly unrelated points-of-view.

Vendors are typically evaluated in four key areas: (1) breadth and richness of services; (2) overall solution functionality; (3) financial viability; and (4) overall cost.

### #1: Breadth and Richness of Services

Services are represented not only by the capabilities of the vendor's solution but also by the added value that the vendor overlays on top of and after the sale. For example, a multi-stage, hosted backup and archival solution is a complex sale and an even more complex implementation. Is your expectation that the vendor will work in lockstep with you before, during and after the implementation? Will they rely entirely on their own resources for implementation or will they also rely on contracted capability? What are the qualifications of the people who will be your primary implementation resources? What kind of post sales support will you receive? Will there be resources on site for a period of time following project completion or will you be dependent on remote resources that must be dispatched if they are needed?

### #2: Solution Functionality

Related to these basic questions are questions around the offered product, service or solution. How well does the vendor understand your business, where you are, and where you want to go? How broad are the considerations it addresses – does the proposal go beyond pure technology to include economic, human capital, end user, application and competitive concerns? How effectively does the vendor “push

back” in project discussions to raise concerns about potential issues? Do they ask as many questions as they attempt to answer? Have they created and offered a technique for assessing the overall effectiveness of the implementation and distinct, well-defined intervals? Similarly, to what degree has the vendor identified management concerns?

### #3: Financial Viability

Mentioned earlier, this particular criterion is of utmost importance, particularly in business continuity implementations that are deemed “mission-critical.” If the vendor does not strike you as being viable, don't consider them in the final running. Keep in mind that the implementation of a complex system does not end the day it is turned up: in essence, it never ends and requires various levels of vendor support throughout its life cycle. Critical questions to ask, then, should revolve around the vendor's business model, survivability, cash position, vision, and commitment to product longevity. For example, to what degree is the vendor committed to building products around open standards? What does the vendor's future vision look like for themselves and their products? Can they tell a compelling and logical story about where they see themselves in three-to-five years? What customer testimonials will the vendor make available to you? Who are their customers, and what sector do they hail from? Are they exclusively in one sector or another, or do they span a broad range and therefore have the ability to draw from a broad experience base? Do they work with channels, and if so, what's the nature of the relationship? How long have they been in business? What is their cash position?

### #4: Cost

Overall cost is a measure of flexibility and operational effectiveness, and the manner in which the vendor presents their solution is a measure of how well they understand your “pain points” and business drivers. For example, does the vendor go to significant lengths to understand your current situation and take whatever steps are available to protect any pre-existing investments that you would like to protect? Does the vendor offer a variety of solutions or packages that provide “silver, gold and platinum” options if you request them? Does the vendor offer anything in the way of modeling tools that will help you make a buy decision?

## Project Assessment

When filling out the paperwork for a work-related accident, regardless of how minor, one of the last questions asked on the form is “How could this accident have been prevented?” If a business continuity solution is to be put into place and it is to deliver on its promises of cost reduction, efficiency and enhanced workplace effectiveness, project managers must create an assessment methodology that monitors project progress, compares it to expectations, and performs qualitative analysis of the results relative to those expectations. There is no such thing as a perfect, problem-free project, but there is such a thing as a well-managed project that experiences minimal problems – typically because of a well-designed assessment and intervention process.

Elements of such a process should include, at a minimum, cost, identified benchmarks, a clear list of anticipated (desirable) outcomes, solid, extensive project documentation, well-crafted flowcharts that identify the interactivity among all logical and physical elements of the project, and a well-designed training program to ensure that personnel responsible for operating, maintaining and using the new system can do so with maximum levels of capability. There should also be an audit process in place to track those things that went well, those things that didn't and discussions about alternative paths that could/should have been taken.

There is nothing wrong with reliance on a single vendor for business continuity services, as long as the selected vendor satisfies the concerns listed in the prior section. If they do, then they are behaving more as a business partner than they are a vendor, and they take their relationship with you seriously – your success is their success – but just as important, your failure is theirs, as well. A professional services provider in this space understands the business criticality of the services they deliver, and will do everything they can to make the customer comfortable with the degree to which they take that relationship seriously.

## Summary of Recommendations

Dentists often have a poster in their examining rooms that says, “You don't have to floss all of your teeth – only those that you want to keep.” That same message applies to the world of business continuity. You don't have to protect all

of your systems – only the ones that provide services that matter. And since most businesses only run systems that matter, those that they do run typically provide services to employees or customers that are deemed mission-critical. If they fail to operate, the business fails to operate, often with disastrous results.

We began this paper talking about the variety of events, some predictable, some not, some accidental, some not, that can cause a serious disruption of your organization's ability to get the job done. Natural disasters, such as flood, fire, earthquakes, hurricanes, mudslides, disease, and lightning strikes cannot, for the most part, be prevented; they can only be recovered from, the degree to which an organization prepares for their occurrence is a measure of the degree to which the impact of these events can be minimized.

Other events, such as failures due to human error, labor strikes, intentional sabotage, burglary, theft, IT-related faults stemming from physical failures, viruses, Trojan horses, worms, and denial-of-service attacks, can largely be prevented through the deployment of carefully thought-out, prudently designed business continuity practices.

Preparation of these practices involves a logical set of steps that include identifying mission-critical systems; conducting a series of granular and highly-detailed risk assessments; creating “strata” of risk tolerance for business systems; locating single points of failure; selecting partners for business continuity resources; and finally, identifying the steps required to ensure effective business process continuity and data recovery. These steps lead collectively to a metrics-driven business continuity strategy that provides guidance for the protection of both logical and physical resources in the corporation. That strategy is then shared throughout the organization to ensure that all employees understand its role, its importance and their responsibilities relative to protecting critical organizational assets.

The final section of the paper dealt with the process of selecting a business continuity service provider suited to the task of implementing the continuity strategy. Assessing and selecting a vendor-partner should be based on a number of factors, and while cost is certainly a consideration, other factors are equally important – if not more so. They include the depth and scope of their offered services, the financial viability of the company, the degree to which they truly

understand the nature of your situation, the appropriateness of their solution for you, and a number of other factors. Don't rush through the process – take your time, assess well, and select based on these factors. You'll thank yourself later.

## XO COMMUNICATIONS AS A BUSINESS CONTINUITY PARTNER

---

XO Communications is one of many companies helping its customers with business continuity solutions. The company stands out in the crowd for one reason: they understand the meaning of the word solution.

Like all service providers, XO has been in the telecom business for a long time, and offers a complete array of telecommunications services designed to handle the access and transport requirements of the most demanding customer applications.

Rather than create technology and then search out applications for it – typical of many players in the industry – XO studies the market, develops an understanding of your business needs in terms of critical information infrastructure – and then crafts solutions around those needs based on the best technology available. We stand behind our solutions because we use them internally to protect our own business resources.

XO's portfolio of business continuity solutions addresses the spectrum of enterprise demands. To ensure high-bandwidth connections to critical buildings, XO's Broadband Wireless solution delivers on the promise of network resiliency. Dedicated Internet Access guarantees network path redundancy and the ability to re-route around network trouble spots on-demand. High-Bandwidth Network Transport ensures the availability of adequate network bandwidth for the most demanding applications.

To ensure peace-of-mind relative to the survivability and availability of database hardware (disks and servers) that might be affected following a disruptive event, XO's Managed Server solution, in concert with our Collocation offering, guarantees access to redundant hardware to support critical operational processes.

Related to hardware availability is application availability. XO offers a number of services that address this critical need. They include Applications Performance Management, Collocation, Hosted Exchange Services, and Web Hosting.

To address the very real requirements of backup solutions, XO's Managed Backup program ensures that enterprise data is safely archived in multiple locations to guarantee survivability in the event of a disruption of business operations.

XO offers Voice Re-Routing and Redundancy Services as well as a collection of Teleworking solutions that ensure functionality for remote workers who could otherwise be isolated during a disruption of normal business operations.

Finally, XO has developed a number of services that address the important demands for system security. These include our MPLS IP-VPN solution, which guarantees the availability of a completely secure, high-bandwidth connection that is extremely flexible and cost-effective; a Managed Security solution, which take care of system and network security as an outsourced service; and Perimeter Email Protection, which dramatically reduces spam and improves network efficiency.

## Conclusion

---

Sad to say, disruptive events do happen – and in most cases, they are a 'when' rather than an 'if' question. Businesses that assume something will occur that will disrupt their ability to do business, and then plan for that eventuality before it happens, will be well ahead of the curve in terms of their ability to survive the event and continue to service customers. Developing a well-planned business continuity plan should be a matter of highest priority for all businesses, regardless of size, structure or function. Remember, chance favors the prepared business: Plan now, and breathe easy later. You won't regret it.

