

73%
of organizations
that experienced
a systems failure
were down for
30 minutes
or more.

- Vision Solutions
2017 State of
Resilience Report

How to Ensure IT and Business Continuity

START HERE: Basic DR

For today's small and mid-sized businesses, the risk and effects from any unplanned outage (downtime) grow with each additional critical application, network enhancement or system upgrade. We don't have to look back very far to see the consequences of sudden or unexpected disasters affecting the IT infrastructure of major cities and businesses of every description.

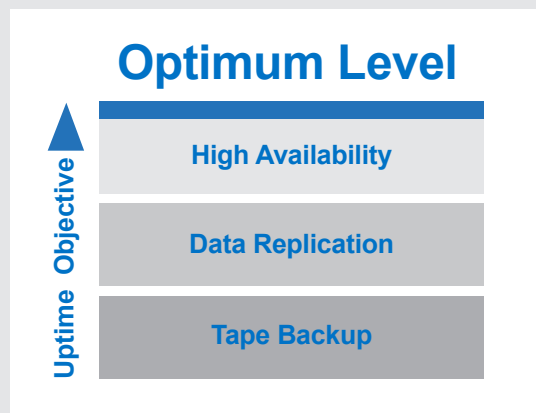
Consequently, IT managers have been—or soon will be—tasked to find ways to mitigate, eliminate or minimize as cost-effectively as possible the risks and effects of unplanned outages on the business. And, even more important, executives will want assurances that information assets—data and applications—can remain available no matter what happens.

This white paper will help you ensure business continuity and survival by leading you through three essential steps—from understanding the concepts of disaster recovery and high availability to calculating the business impact of downtime.

We will discuss in general terms the concepts of business continuity and disaster planning. We will focus primarily on specific IT strategies you can easily and affordably implement.

Bottom line:

Disaster recovery plans and data replication alone are not enough. You will want to look for the most effective way to ensure the optimum level of uptime for your organization to match your specific objectives with the best availability choice.



STEP 1: Getting Started

Before you begin reviewing the available technologies that support disaster recovery, you first must consider the business. You need to identify which business processes are most important to keeping your organization operational.

Once you have identified the most critical business processes, work with the business units to determine their availability requirements for each process. Document the requirements in an internal SLA that specifies the availability goals for each process and articulates the costs of not meeting the goals.

For example:

At **company A**, the order entry and shipping departments require that their information infrastructure processes must be functional 24 hours every day of the year except corporate holidays. If this requirement is not met, the company loses 80 percent of its productivity, which translates to \$10,000 (US) per hour plus penalties of \$100,000 per hour for every hour the processes are unavailable.

At **company B**, the payroll department requires that their information infrastructure processes must be functional from 8 a.m. to 6 p.m. Monday through Friday. Not meeting this requirement costs the company 50 percent of their productivity, which translates to \$25,000 (US) per hour of downtime.

Another organization, **company C**, has the need to comply with strict information availability requirements due to government regulations and has made it imperative that its applications remain available even during routine backup processes.

Documenting the cost of not meeting availability requirements helps you determine the value of a software investment used to improve availability. This information also helps you prioritize the processes to analyze. After documenting the service levels required, you can start analyzing the availability needs of each business process technology by technology.

Understanding Downtime and Availability

Professionals know recovery plans are critical, but confidence in them is low:

84%

are less than 100% confident that their DR plan is complete and tested.

- Vision Solutions
2017 State of
Resilience Report

Most organizations have availability objectives somewhere along a continuum between multiple hours of downtime with significant data loss to real-time 24/7 uptime with zero data loss.

Your objectives depend on your business needs, your data and application requirements and your organizational structure. The goal, however, should be to prevent the inevitable system downtime from affecting business uptime.


There are two types of downtime: unplanned and planned.

Unplanned Downtime

Surprisingly, unplanned downtime represents less than 5-10 percent of all downtime.

These events include security violations, corruption of data, power outages, human error, failed upgrades, natural disasters and the like. Some forms of unplanned downtime, such as hardware failure, pose a lessening threat to availability as most servers today offer exceptional reliability.

CAUSES OF UNPLANNED DOWNTIME



Source: Gartner Group, Inc.

For example, IBM Power System servers running IBM i provide more than 99.9 percent documented reliability and average 61 months between failures— more than five years of server uptime.

Unplanned downtime can strike at any time from any number of causes. Although natural disasters may appear to be the most devastating cause of IT outages, application problems are the most frequent threat to IT uptime. According to Gartner, people and process problems cause an estimated 80 percent of unexpected application downtime. Human error, such as not performing a required task, performing a task incorrectly (such as mis-configuring software), overburdening a disk drive or deleting a critical file, play havoc with applications.

Planned Downtime

While unplanned events tend to attract the most attention, planned downtime actually poses a bigger challenge to business uptime. Routine daily/weekly maintenance to databases, applications or systems usually lead to interrupted services. Studies show that system upgrades, performance tuning and batch jobs create more than 70-90 percent of downtime.

Although companies must be concerned about natural disasters, the inherent daily threat posed by application problems and human error should be a major focus of your efforts. This is especially true when the exposure of software applications to unplanned downtime is aggravated by a host of other business and IT issues such as:

- The need to retain, protect and audit email, financial and other records under regulatory compliance mandates.
- The acceleration of security risks from both inside and outside the business including viruses, worms, hacker attacks and industrial espionage.

The average total cost per minute of an unplanned outage increased from **\$5,617 in 2010** to **\$7,908 in 2013** to **\$8,851 in 2016**.

- Emerson Network Computing

- Distributed applications that are accessed, maintained and updated by different classes of users and business partners.
- Multiple platform IT environments in which applications operate interdependently to accomplish business critical tasks.
- Fewer IT personnel and labor hours available to maintain.

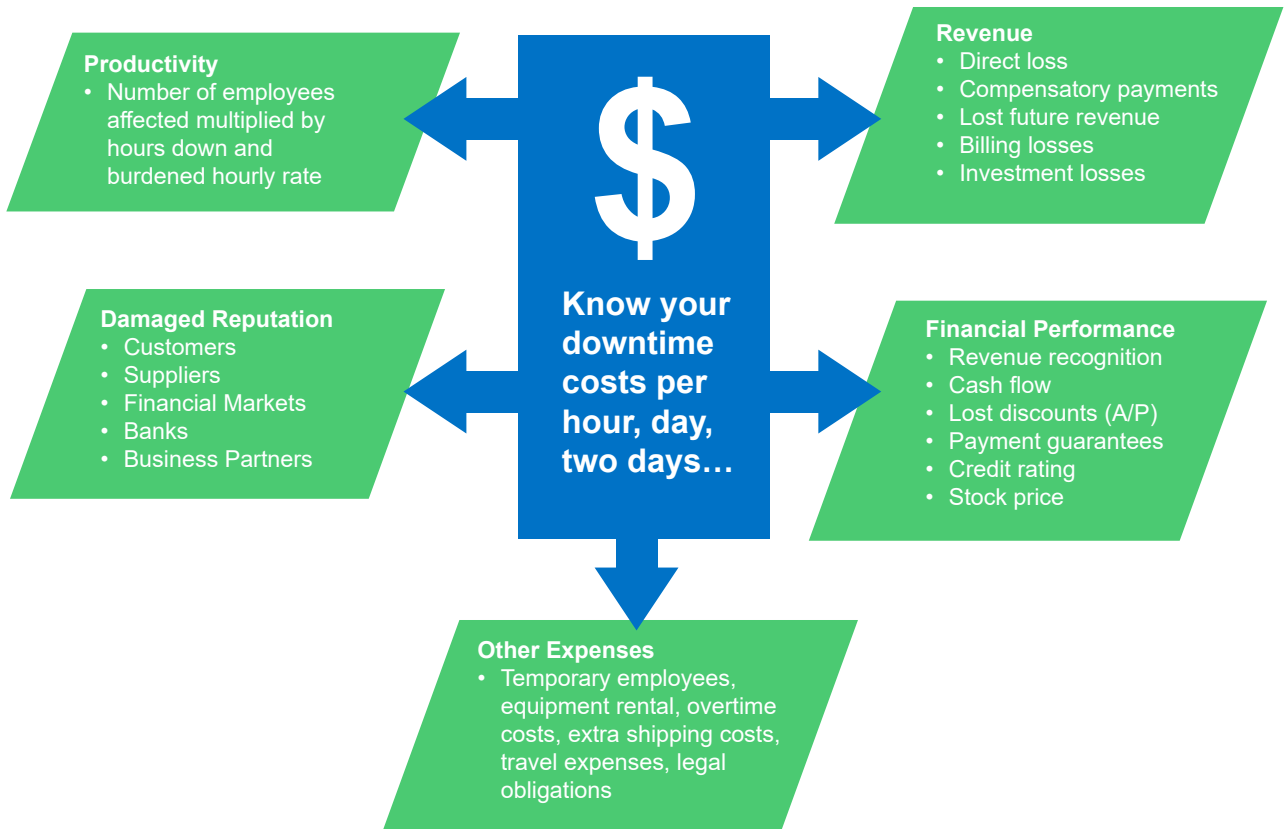
STEP 2: Assess the Financial Impact— Calculate the Cost of Downtime

How much does downtime cost your business? The answer may not be as obvious as you think. Unexpected IT outages can unleash a procession of direct and indirect consequences both short term and far reaching.

The widely varying dollar amount that can be assigned to each hour of downtime depends on the nature of your business, the size of your company and the criticality of your IT systems to primary revenue generating processes. For instance, a global financial services firm may lose millions of dollars for every hour of downtime, whereas a small manufacturer that uses IT primarily as an administrative tool would lose only a margin of productivity. Also, government agencies with downtime issues would find it very difficult to assign a value to not delivering promised services to citizens.

Common effects and costs of downtime include:

Tangible/Direct Costs	Intangible/Indirect Costs
• Lost transaction revenue	• Lost business opportunities
• Lost wages	• Loss of employees and/or employee morale
• Lost inventory	• Decrease in stock value
• Remedial labor costs	• Loss of customer/partner goodwill
• Marketing costs	• Brand damage
• Bank fees	• Driving business to competitor
• Legal penalties	• Bad publicity/press



Source: Gartner Group, Inc.

Most businesses cannot function without computer support, and most businesses that suffer catastrophic data loss or an extended IT outage go out of business. On average, enterprises lose between \$337,020 and \$474,480 (US) for every hour of IT system downtime according to estimates from studies and surveys performed by IT industry analyst firms. In addition, financial services, telecommunications, manufacturing and energy organizations are also high on the list of industries with a high rate of revenue loss during IT downtime. Here is a brief sampling of typical U.S. dollar downtime costs per hour by industry:

- Brokerage Service - \$6.48 million
- Energy - \$2.8 million
- Telecom - \$2.0 million
- Manufacturing - \$1.6 million
- Retail - \$1.1 million
- Health Care - \$636,000
- Media - \$90,000

Sources: Network Computing, the Meta Group and Contingency Planning Research. All figures U.S. dollars.

Consequences of Downtime

No matter what the cause, downtime impacts more than day-to-day interactions. It can impact the integrity of your databases and the applications that use them. For example, a disaster recovery strategy that relies on once-a-day, nightly tape backups risks a whole day's worth of data should an unplanned event occur and crash IT systems a few hours or minutes before the backup process kicks off.

Some businesses could survive that kind of data loss. Others, dependent on electronic data interchange, required to archive information for legal reasons, deploying a global workforce expected to collaborate around the clock, or using eCommerce to make sales and deliver customer service 24/7, will suffer the effects for a long time into the future.

5 SIGNS Downtime May Be a Major Threat

Many organizations may believe that they remain unaffected by downtime issues. After all, few users complain; customers seem happy. Some important signals, however, may indicate that your current situation or availability solution may need re-evaluation.

- 1 Shrinking Backup Windows**—eBusiness and supply-chain processes are putting the squeeze on backup windows. The Gartner Group reports a 66 percent per year decline in the available time for quality backups.
- 2 Expanded Internet Dependence**—As you exploit the internet to improve customer satisfaction and reduce costs, your dependence on internet enabled availability grows exponentially. When email is integrated into business functions to improve customer communications, your dependence becomes even greater. And the risk that downtime poses to the business increases.
- 3 Globalized Computing**—Access to critical data from anywhere in the world improves collaboration and enables faster, more informed decisions. Such dependence requires continuous access to information and applications; therefore the impact of downtime will be enormous.
- 4 Distributed Applications**—New applications are now running across multiple servers simultaneously, enabling them to capitalize on the servers' varying strengths; however, should one server experience downtime, the entire critical application may go down.
- 5 Server Consolidation**—Server, storage, and data center consolidation projects drive down IT and application costs, but with fewer points of failure a consolidated environment poses a greater downtime risk.

Don't Forget the Added Burden of Compliance

Many regulations require companies to support more stringent availability standards. Several acts and regulations, directed at specific industries or a broad cross-section of companies, mandate the protection of business data and system availability. Businesses may incur financial or legal penalties for failing to comply with these data or business availability requirements.

Health Insurance Portability and Accountability Act (HIPAA)—ensures that only properly authorized individuals have access to confidential patient health data and provides long-term guidelines to secure confidential information. HIPAA mandates a five day maximum turnaround on requests for information. A recent amendment also qualifies data storage companies, with access to protected health information, as business associates and liable for compliance with HIPAA standards.

Sarbanes-Oxley Act—stipulates that CEOs and CFOs attest to the truthfulness of financial reports and to the effectiveness of internal financial controls. Sarbanes-Oxley mandates a required timeframe in which to report financial results—each quarter and at year-end. Failure to make these deadlines can result in financial penalties.

New Basel Capital Accord (Basel II) to The Third Basel Accord (Basel III)—requires financial institution capital reserves to include operational and credit risks and includes IT security risk as a principal operational risk. Basel II also requires business resiliency standards for any financial institution doing business in the European Union.

Gramm-Leach-Bliley Financial Services Modernization Act—limits access to non-public information to those with a “need to know” and requires safeguarding of customer financial information. Loss of important data can lead to penalties for the financial institution.

The USA Freedom Act—defines what information can be made available to federal and local authorities for those suspected of terrorism or terrorist-related activities.

Cost of Downtime Calculator: How Much Will Downtime Cost Your Business?

Any application that plays a role in developing, creating, manufacturing, supporting, or distributing a product or service to buyers will significantly impact the organization during an outage event.

- META Group, Inc.

To determine how much an hour of unplanned downtime will cost your business, you need to ask a series of questions regarding the real world impact it will have on your customers, partners, employees and your ability to process transactions, such as:

- How many transactions can you afford to lose without significantly impacting your company?
- Do you depend upon one or more mission critical applications such as ERP or CRM software?
- How much revenue will you lose for every hour your critical applications are unavailable?
- What will the productivity costs be for the loss of available IT systems and applications?
- How will collaborative business processes with partners, suppliers and customers be affected by an unexpected IT outage?
- What is the total cost of lost productivity and lost revenue during unplanned downtime?

STEP 3: Uptime and Business Resiliency— It's All About Recovery

Determine the RPO and RTO Requirements for Your Business

Following any unplanned outage, how quickly must you have the organization up and running as close to normal business operations as possible? Remember, every minute costs you—take a look at your downtime cost per hour.

Your recovery will depend on two objectives: your recovery time and your recovery point. These two measures will determine the optimum availability your organization will need.

1. Recovery Time Objective (RTO). RTO defines how quickly you need to restore applications and have them fully functional again. The faster your RTO requirement, the closer you move to zero interruption in uptime and the highest availability requirements.
2. Recovery Point Objective (RPO). RPO defines the point at which the business absolutely cannot afford to lose data. It points to a place in each data stream where information must be available to put the application or system back in operation. Again, the closer you come to zero data loss and continuous real-time access, the higher availability you will require.

You may have different RTOs and RPOs for each of your business critical applications. For example, a supply chain application that feeds a production plant may require a recovery time of only a few minutes with very minimal data loss. A payroll system that is updated weekly with only a few records may only require a recovery time of 12 hours and a recovery point of 24 hours or more before the impact will affect the business.

Small to mid-sized businesses can implement information availability in several different ways.

Matching Uptime Requirements to Availability Solutions

How do you best meet the availability requirements of each system in your organization and achieve the optimum RTO and RPO appropriate for your organization? Some organizations, or some particularly critical applications within an organization, may require an exceptionally high level of availability.

Any availability solution you select must ensure that information and applications remain as accessible and available as needed to continue to drive revenue, profitability and productivity at acceptable levels no matter what planned or unplanned events occur.

The availability solution you choose should:

- Protect your data, applications, and systems to a level that meets your business requirements and RTO and RPOs.
- Manage business uptime as automatically as possible to streamline operations and save time.
- Assure the integrity and quality of your environment during interruptions and when it returns to full operations.

Small and mid-sized businesses that face the potentially devastating consequences of unplanned downtime can protect themselves against the loss of time and money with a high availability solution. Depending upon their particular business and IT needs, small to mid-sized businesses can implement information availability in several different ways, including replicating data to a secondary server to maintain continuous application availability or replicating data to a repository at a remote location for disaster recovery in the event of a total facility loss at the production site.

Let's look at some of the options to protect your business from the consequences of downtime.

Tape Backup/Archiving Solutions

Tape-based backup and recovery solutions are the oldest form of disaster protection. Tape solutions offer relatively low cost and high portability. You probably rely on tape for a once-a-day backup of your data now. Because it represents a relatively low-cost way to archive information for the long term, tape will no doubt play some role in the IT infrastructure for some years to come. For example, even in high RTO and RPO businesses, where more advanced availability solutions are also used, tape can still play a role in protecting and backing up non-critical applications. Due to its own limitations, however, tape will be unable by itself to provide RPOs or RTOs of seconds, minutes or even a few hours. Since many organizations have a substantial investment in tape storage solutions, an information availability software solution should act as a complement to your tape strategy, making it much more flexible to use.

Disk-Based Backup and Snapshots

Disk-based backup and snapshot solutions provide readily available access to and protection of your business data with RTOs and RPOs in the range of hours. By performing frequent data backups to a secondary server or partition, it provides businesses with the ability to recover from an unexpected outage without the loss of large amounts of data or days or weeks of labor restoring the production environment. When the backup server is placed in a remote location, it also serves as a disaster recovery solution.

The information protection requirements in the Disaster Recovery regulatory demands will allow only short gaps of missing data or information due to an event. This means that doing daily back-up to tape is no longer sufficient

- META Group, Inc.

Continuous Data Protection

Continuous data protection, or CDP, is a flexible disk-based technology that enables businesses to quickly and easily recover their data to any point-in-time. For example, it's not uncommon for a user to accidentally delete a critical file. Or for a virus to corrupt business data. These actions render the data unusable, even though the server or other hardware resources continue to work as expected. CDP enables you to recover a version of your data to a point-in-time just prior to the accidental deletion or virus corruption. This earlier version of the data can then be restored to the production environment.

High Availability

High availability solutions deliver continuous uptime with zero data loss so your applications and business data are always available when you need them. A backup server with a current replica of your application environment is always available to failover or switchover to replace your production server with an RTO of seconds to minutes and an RPO of zero. High availability dramatically reduces the risks and costs of business interruptions. Innovations in automation and the inclusion of CDP capabilities make it an increasingly agile and easy-to-manage strategy for ensuring business continuity.

Disaster Recovery as a Service

The growth of cloud adoption for business applications, and advances in disaster recovery software that allow it to exploit the flexibility and scalability of the cloud, have resulted in Disaster Recovery as a Service (DRaaS) offerings for simple, cost effective protection for data and servers. The range of DR services offered can vary greatly from provider to provider, however most are offered under a plan where the DR software is hosted by the provider and licensed to users on a subscription basis. The services provided can vary greatly between providers in terms of coverage, RPO and RTO. Most providers present options that include replication services and the ability to recover data and servers. How frequently the copy of the data or servers is updated, the method for recovery, and especially how quickly recovery can be accomplished, varies by the kind of service being contracted.

DRaaS offerings vary from provider to provider in terms of coverage, RPO and RTO. When considering a DRaaS solution, keep these seven must-have features and capabilities in mind:

- 1 **Multi-platform support** — Make sure that all of your physical, virtual and cloud-hosted production servers can be protected. If you don't already have a presence in the cloud, the provider should help migrate to the cloud.
- 2 **Multi-cloud support** — You should be able to use more than one cloud service or platform concurrently to mitigate the risk of your sole cloud service suffering a major outage. This also frees you to switch from one cloud service vendor to another in the future.
- 3 **Recovery into the cloud** — Rather than just having a recoverable backup image in the cloud, you want to be able to actually switch operations to your cloud backup servers should a disaster occur.

By 2020,
a corporate
“no-cloud” policy
will be as rare as
a ‘no-Internet’
policy is today.

- Gartner, Inc.

- 4 **Flexible licensing** — Your technology provider should offer DR with subscription based, service-oriented licensing and billing options.
- 5 **Real-time replication** — True real-time replication captures changes as they happen, eliminating the risk of losing critical data. All other solutions have windows of varying sizes where data/changes will be lost if an outage occurs.
- 6 **Scalability** — Your DRaaS solution should be able to grow as you grow, whether you are a small business with just a few servers or your datacenter is expanding to thousands of physical and virtual machines.
- 7 **Reporting and administration** — You should be able to participate in monitoring your solution if desired and frequent, detailed reporting should be available regarding your level of protection and readiness to failover.

Take the Next Step: Ensure IT and Business Survival

When the real world costs of unplanned downtime are taken into account, an availability solution is a cost-effective strategy for protecting businesses from serious injury. In particular, small to mid-sized businesses can benefit significantly from availability solutions because they are generally more vulnerable to severe damage from unexpected outages and have fewer resources to stage a recovery.

An availability solution shouldn't be hard work or beyond your budget. There are affordable, easy-to-manage solutions that provide significant benefits to small and mid-sized businesses by minimizing the risks and consequences posed by unexpected IT outages. An availability solution:

- Lowers the risk of significant costs to business such as lost revenue, productivity, legal penalties and brand damage caused by unplanned downtime.
- Protects business relationships with customers, partners and suppliers by ensuring that applications and data will be available to satisfy their needs and unique schedules.
- Enforces service level agreements by maintaining predictable RTOs and RPOs in the event of an IT outage.
- Enhances ROI on existing resources by assuring they will be available to generate revenue and support business processes.
- Ensures compliance with government and trade regulations by securing email and record retention requirements and protecting the availability of business data and reporting processes.

Easy. Affordable. Innovative. Vision Solutions.

Vision Solutions is a leading provider of business resilience solutions—high availability, disaster recovery, migration and data sharing—for IBM Power Systems. For more than 25 years, customers and partners have trusted Vision to protect and modernize their IT environments, whether on-premises or in the cloud.

Visit visionsolutions.com and follow us on social media, including Twitter, Facebook and LinkedIn.

Also find us on:

Facebook: <http://www.facebook.com/VisionSolutionsInc>
Twitter: http://twitter.com/VSI_Power
YouTube: <http://www.youtube.com/VisionSolutionsInc>
Vision Solutions Blog: <http://www.visionsolutions.com/blog>



15300 Barranca Parkway
Irvine, CA 92618
1.949.253.6500
1.800.683.4667

visionsolutions.com

