

Realtime  
publishers

# The Essentials Series

## Making High Availability Pay For Itself

*sponsored by*

**MARATHON**  
Run to Infinity

*by Ron Barrett*

---

Article 1: High Availability’s Spectrum of Cost and Capabilities.....	1
Trends that Are Transforming HA .....	1
Measures to Determine HA Solution Needs .....	3
Calculating ROI .....	5
The Availability Scale .....	5
Coming Up Next.....	7
Article 2: Matching High-Availability Technology with Business Needs .....	8
Availability vs. Affordability .....	8
Matching Technology to Availability Needs .....	9
HA Technologies.....	10
Hot-Swappable Components .....	10
Snapshots.....	11
Replication and Failover .....	11
HA (Failover) Clusters.....	11
Synchronized Systems .....	11
Virtualization .....	12
Coming Up Next.....	12
Article 3: Relating High-Availability Metrics to Business Value.....	13
How to Explain Your Metrics and Show Business Impact .....	13
Translating the Technology of HA to Useful Business Reports.....	14
Summary .....	16

---

## Copyright Statement

© 2009 Realtime Publishers. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtime Publishers (the “Materials”) and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtime Publishers or its web site sponsors. In no event shall Realtime Publishers or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtime Publishers and the Realtime Publishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtime Publishers, please contact us via e-mail at [info@realtimepublishers.com](mailto:info@realtimepublishers.com).

# Article 1: High Availability's Spectrum of Cost and Capabilities

---

One of the most difficult achievements in technology is gaining the resolve to spend on the possibilities of what if. This is understandable especially in lean economic times such as we are now experiencing. However, the idea of losing data, services, or the ability to access an application that is usually readily available highlights the need to keep these IT services up and running at all costs (sort of). But what is high availability (HA) and how do you determine whether you have a need for it in your organization?

HA is sometimes mistakenly defined as a means of *recovering* from an unexpected loss or outage within your organization. Although recovery plays a part in HA technologies, HA is more about making sure applications, services, and data are always available to end users at all times regardless of what is taking place behind the scenes. Advances in technology have improved the way we work and made us an 'always on' society, which naturally means that the technology running this society must itself be always on.

It was not long ago that when we spoke about HA, it was in the context of wants. Businesses wanted to be able to have mission-critical applications and data available with almost no downtime. And there are some industries that have been working in this realm of HA for a long time: The need for consumers to access banking, surf the Internet, shop, or execute a stock trade any time of the day or night means that for these companies HA has become mission critical in and of itself. Lately, HA has moved beyond these specific industries and has become more widely adopted in what might seem to be non-traditional businesses. What trends have moved HA out of the nearly unattainable and into the mainstream?

## Trends that Are Transforming HA

The need to be available all the time has transformed over the years; it is no longer specific industries that need to think in terms of HA. We have gone from the 9×5 local or regional businesses to a 24×7 global marketplace. Sectors that could not touch a large marketplace are suddenly able to market, produce, and sell anywhere in the world. And that paradigm shift has changed how we now think, work, and purchase goods and services. The perfect storm of communications, mobility, and affordability has transformed HA from a want to a need.

As these trends have changed, so has the consumers' view of what is an acceptable level of service. It has also changed the business' view of meeting the expectations of their clients. This naturally has driven demand for technology to meet the needs of both. These trends have not only changed the view of HA but also the composition of it. In the past, we thought of HA as an all-or-nothing technology. The complexity and expense of implementing an HA solution meant that this was only a solution for companies with deep pockets. And even then, it meant they could apply HA only to 'mission-critical' applications.

The technology to meet the needs of business has become both less complex and much more affordable. This has given HA a broader appeal and now has all sorts of organizations looking at how they can be both more competitive and more responsive to the needs of their clients and even their employees.

As business has become a 24×7 operation, the workforce behind that business has also changed. The need to access email, documents, data, and so on means that back-end systems need to be just as available as the front-end-facing systems for an organization. For some businesses, even the normally-accepted downtime for backup and maintenance is no longer an option, and this is where HA solutions shine.

So here we are at the crossroads of want and need again but with a twist: organizations that **want** to meet their business objectives **need** to have some kind of viable HA solution. Thankfully, meeting that need is no longer about available or not available but rather about levels of availability. What needs to be available? What is acceptable downtime (if any) for a particular application, service, or data set? And what solutions can help you meet your business objectives for HA? To find the answers to these questions, we need a way to measure the cost of the solution versus the cost of downtime—or what we will refer to as the *availability scale*.

## Measures to Determine HA Solution Needs

When considering which services, data, and applications should be part of an HA solution, you need to determine a few key factors. Which services, data, and applications do I need to have available? And what is the cost to me if they are not available? Although they may seem over simplified, these two questions are at the core of all the others that you will ask yourself. So if you are going to find out what needs to be available and what the possible loss is for any downtime to that service, data, or application, you need to have a means of measuring the two factors. You also need a baseline to measure against. The baseline in this case should be a predetermined service level agreement (SLA) based on Recovery Time Objectives (RTOs) and Recovery Point Objectives (RPOs):

- **SLA**—The common understanding about services, priorities, responsibilities, and guarantees of services. These include both a target level of service as well as a minimum acceptable level of service. SLAs often include a promise of percentage of uptime (that is, 99.9999%). This is the standard you look to achieve or exceed in creating an HA solution.
- **RTO**—The duration and service level that needs to be restored after a disruption to avoid negative consequences on the business.
- **RPO**—The point in time in which services and data must be restored; this is also defined as the acceptable amount of lost data.

Once you have established your baseline SLAs, you then use them to compare against the actual performance of your systems. A set of metrics that work well for measuring against the SLA are the ITILv.3 Availability Management key performance indicators (KPIs).

The **ITIL v.3 Availability Management KPIs** provide metrics for the following:


- **Service availability**—Measures the availability of services compared with the established SLAs
- **Number of IT service disruptions**—Measures the total number of service, data, or application interruptions in an organization; these can be further segmented into root causes (that is, hardware failure, customer action, security issue, and so on)
- **Duration of IT service disruptions**—Measures the accumulative duration of service, data, or application disruptions in an organization

### Note

Opinions vary on whether to include both planned and unplanned disruptions into the calculations of availability. Some argue that excluding planned downtime provides higher-than-actual uptime values. When planning for an HA solution, the right answer probably depends on what you are discussing. An application that is being used 24×7 should have planned and unplanned downtime calculated. However, there may be services that do not require 24×7 access; for these, the exclusion of planned downtime would be acceptable when determining an HA solution.

- Availability monitoring—Determines a percentage of all services and infrastructure components that have or will be monitored
- Availability measures—Implements measures with the objective of increased availability

To illustrate the process, let's consider Application A and Service B, which both have an SLA of 100% uptime. Using the current in-place solution (even if that solution is no solution), apply these metrics to gain a baseline of availability such as the one shown in Figure 1.



**Figure 1: ITIL Availability Management KPIs as measured against existing SLAs.**

As you can see from Figure 1, the baseline SLA measurements of 100% uptime (left) are significantly different from the reality of the actual measurements of Application A and Service B. Once you have metrics to determine the actual availability of your systems, you can measure them against the RTO and RPO in your SLAs to get a better idea of where you need HA protection.

However, there is more to consider than the purely technical. HA is a business process as much as it is a set of technologies. So, naturally, the cost impact of these metrics weighs heavily on the decision of what solution to implement; this is where ROI comes into play.

## Calculating ROI

To calculate the ROI, you need to first determine the fiscal impact of disruptions. To do so, you need to consider direct/indirect costs such as:

- Loss of business operations—Revenues from client transactions, orders, production, shipping, and so on
- Lost employee productivity—Employee salaries multiplied by lost hours
- Business reputation—Value of lost business due to the disruption
- Operation costs—Loss of revenue due to penalties, compensations, lost interest, discounts, stock price losses, and so on

Once you understand the fiscal impact, you can take the metrics for current service level, look at your RTO and RPO, then consider the HA solutions that will help you achieve the levels of availability for that IT service. Taking the cost of the solution and subtracting it from the fiscal impact report will provide a ROI for the HA solution.

What you need to remember is that all applications, services, and data that are considered for HA need to be put through these measurements to determine a baseline. That baseline along with several other factors will help you begin to create an availability scale.

## The Availability Scale

Rather than thinking of HA in linear terms, you need to forego the all-or-nothing approach and weigh potential solutions in terms of need, technology, and cost. HA is not about trying to fit all your business needs into your solution; it is about your solution fitting your business needs. After you weigh the impact of an application, service, or data disruption, you can apply that weight to your availability scale. The availability scale uses a weighted number (1 to 10) for each IT service, desired level of availability, and cost.

We usually think of availability as having five levels:

- Unprotected—IT services that are not part of any HA solution; these may use some kind of redundancy, but by and large they are unprotected
- Reliable—Applications and services or hot-swappable components that can eventually be recovered; they are not business impacting
- Recoverable—Redundant infrastructure components that have some automatic recoverability built in; some downtime is acceptable
- Highly Available—Mission-critical IT services that drive the business; demand for availability is high, so these need high redundancy at the hardware, software, and communications levels
- Always (continuously) Available—Few if any of these IT services exist in most organizations, as they require absolute zero acceptable downtime; these are triple or even quadruple redundant



The idea behind this scale is to simplify the decision-making process. IT services with a weight of 10, desired availability of 8, and a cost of 8 make sense to implement at the desired level. IT services with a weight of 5, a desired availability of 9, and a cost of 8 might need a reevaluation as to the desired availability level. Does the cost and weight of the service justify that level of availability? The scale helps to take a balanced approach to planning for HA solutions. Table 1 provides a sample of the availability scale. A simple way to make sure the three elements are in sync is to use the following formula:

$$IT\ Svc.\ Weight * Desired\ Availability / Cost = Availability\ Factor$$

	IT Service Weight	Desired Availability	Cost	Availability Factor
Email	10	8	8	10
Web Site	5	9	8	6

**Table 1: A sample availability scale.**

In the first instance, the email application has a service weight of 10 and a desired availability of 8. Multiplying the two factors and dividing them by the cost gives you an availability factor of 10. In contrast, a Web site with a service weight of 5 and a desired availability of 9, once divided by the cost of 8, gives you an availability factor of 6 (rounded up). You can then take this information and match an availability factor with an HA solution within the availability scale.

Availability Factor	Availability Level	HA Solutions
Factor 0 to 1	Unprotected	No Availability Solution
Factor 2 to 3	Reliable	Hot Swappable Components Snapshots
Factor 4 to 6	Recoverable	Replication & Failover
Factor 7 to 9	Highly Available	Virtualization, HA Clustering, Synchronized Systems
Factor 10	Continuously Available	System Fault Tolerance

**Table 2: Comparing availability factors to HA solutions.**

## Coming Up Next...

The second article of this series will look at how to match availability wants to business needs and will use the availability scale to determine the best HA solution to meet the common ground of these two factors.

## Article 2: Matching High-Availability Technology with Business Needs

---

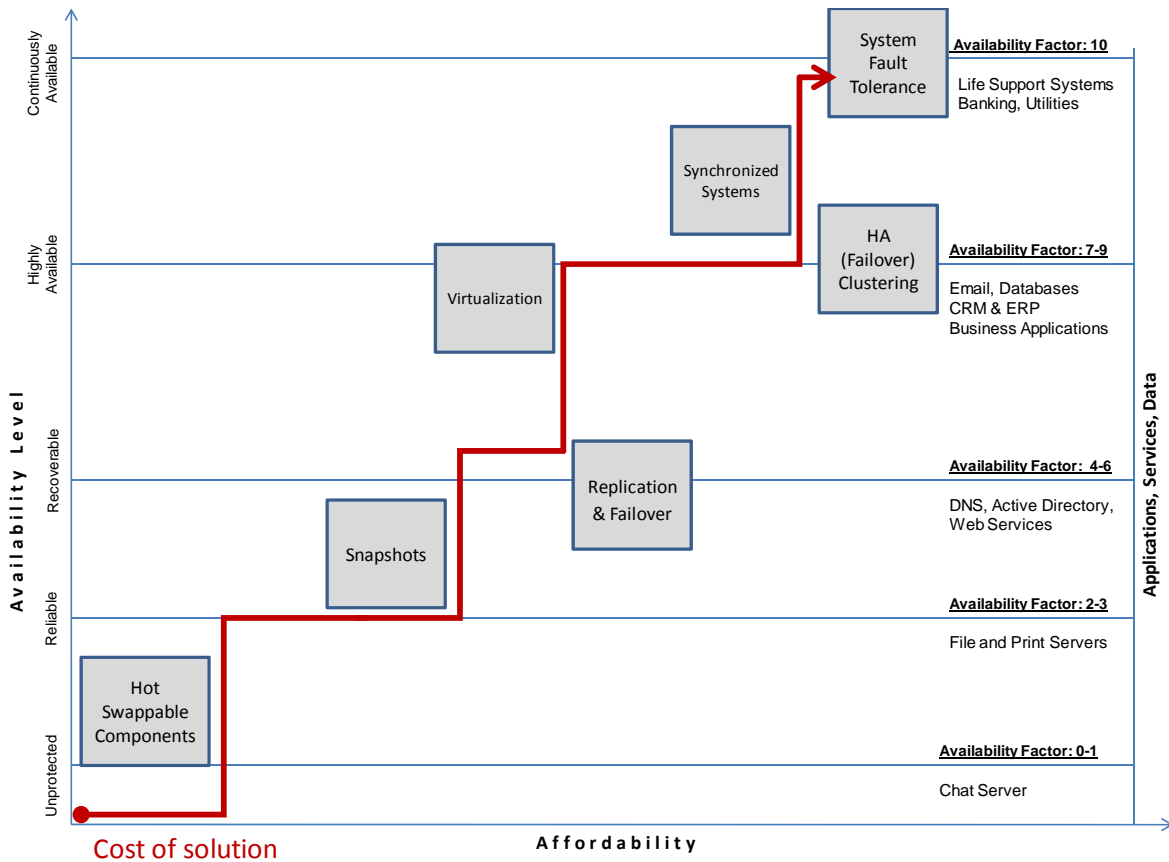
As we discussed in Article 1, there are a few ways to tackle the issue of high availability (HA). Finding a solution—or more appropriately a *set* of solutions—is about matching the available technologies to the needs of the business. The first step was to look at the applications, services, and data that need to be available. Then we used the metrics to determine the current reliability of your environment and the fiscal impact of a disruption.

After using the availability scale to provide a weighted metric for the IT service, desired availability, and cost, you need to next look at availability versus affordability. This is where you can discern whether the HA solution fits both the business needs and budget.

### Availability vs. Affordability

Ideally, we would like to see all our IT services run at the level of continuous availability. After all, that is what end users and clients expect. The reality is that we are not likely to fit our needs into a single HA solution. Rather a set of solutions aligns better with how we actually do business. In any typical organization, there exists a hierarchy of job duties; not all are equally important, although all are necessary to get the job done. The same is true of IT services—all of them are necessary to keep the whole organization running but not all carry the same importance in reaching business objectives.

Keeping the balance between the two is a matter of picking the best availability to match affordability. This is not just about having the dollars to throw at a solution. It makes no business sense to make an internal Web server that has no direct business impact highly available. The nature of the service would in most cases allow for a recoverable solution with some downtime. Likewise, it would make no business sense to take a core business application and classify it as recoverable when in fact it should be continuously available. Continuous availability is an expense, so in some cases, the business just cannot support that expense. The next logical step is to opt for HA; you should not accept a recoverable system to save a few dollars. If in fact this core application should fail, the ROI will be greater than the cost of the HA solution. Figure 1 provides a matrix of HA solutions, availability levels, and relative affordability. It also classifies examples of IT services that fall into those levels of availability.



**Figure 1: Example matrixes highlighting availability vs. affordability.**

Deciding which solution works best is important in making sure that the HA solution pays for itself. Remember that in many organizations, an HA solution is not sold to decision makers based on the technology itself. It is typically sold based on meeting business objectives at a particular price.

### Matching Technology to Availability Needs

In line with choosing the right availability and affordability is understanding what each HA technology brings to the table. Let’s review the levels of availability and define how they meet (or fail to do so) the business objectives:

- Unprotected—These applications, services, and data do not have an impact on business objectives. A disruption means a complete loss of usability, and the Recovery Time Objective (RTO) is typically days or unrecoverable.
- Reliable—These applications, services, and data do not have an impact on business objectives. A disruption means a complete loss of usability, and the RTO is typically days.

- Recoverable—These applications, services, and data typically are infrastructure components. Recovery is sometimes automatic. A disruption means some impact at the user level while the systems replicate and failover to redundant systems. The RTO is typically 1 day.
- Highly Available—These applications, services, and data are core or mission-critical applications, services, and data. Disruptions are minimal and typical brief. Due to the nature of these IT services, the RTO is typically 1 to 4 hours.
- Always (continuously) Available—These applications, services, and data are highly critical and include utilities, trading systems, banking, telecom systems, and so on. Disruptions are not an option, and due to the nature of these IT services, the RTO is 0.

Understanding the levels of availability and the potential impact on business objectives helps you to better relate how HA pays for itself. It is important to ‘weigh’ the service against the desired availability level, then measure the potential cost of the solution against both. Using the availability scale, you then create a mix and match approach. This will keep critical systems up while reducing overall downtime. Thus, you can meet or exceed the SLA you have put into place. It also means you can control costs.

To make a determination of which HA solution matches your needs, you need to consider the various technologies behind HA. Once you determine the availability need of an application, service, or data set, you must understand what technologies will help you reach those objectives.

## HA Technologies

We will wrap up this article by defining the various technologies that can be used to help you meet your availability goals. These definitions can be useful in presenting an HA solution.

### Hot-Swappable Components

Hot-swappable components provide the ability to change disk drives, controllers, power supplies, and so on without turning off the system. This allows the system to be available without much downtime (certain RAID configurations may require a rebuild, which can degrade service levels although not necessarily service itself). This solution offers no protection from sudden interruptions or disruptions due to software factors. This option is relatively inexpensive to implement because the cost of hardware is relatively low.

## Snapshots

Snapshots provide a point-in-time copy of data and files that can then be backed up to various media. Using snapshots to take a copy of the data at timed intervals and moving snapshots to external storage such as a NAS/SAN will make the data and files recoverable if there is a service disruption. This solution again may provide a reliable or recoverable system but not a highly-available one because this method requires restoring from a snapshot in the case of a disruption or failure. This option can be a fairly inexpensive solution and requires usually only a licensing fee from most SAN/NAS vendors.

## Replication and Failover

Replication and failover solutions provide a complete replica of the primary system and replicate both the hardware and software elements of the systems. This solution usually can detect imminent disruptions and move the workload over to a failover system to ensure availability. These systems are usually failed over automatically with minimal impact on data and users. These solutions are a bit more costly compared with the previously mentioned options because they require hardware to replicate and failover to should a disruption take place.

## HA (Failover) Clusters

HA clusters not only replicate systems and provide automatic failover but also go an extra step in attempting to eliminate any single point of failure (SPOF). These solutions usually will detect a hardware or software fault and start the application on another node immediately. Clusters require at least two nodes (computers) but can consist of dozens of nodes. This setup provides a very high level of availability but comes with a very high cost because redundancy is required not only at the system levels but also within the entire infrastructure. However, in situations where the impact of downtime on the business is significant, the higher costs can be easily justified.

## Synchronized Systems

Synchronized systems combine two Windows servers and present them to the world as a single server. The servers are monitored and tested by the synchronization software, and if one server fails, the other is available as an exact duplicate to continue providing the applications to support the business. The two servers are completely synchronized at the OS, application, and data levels by ensuring that changes happen to both servers simultaneously. The advantage of this approach is that it is a single solution that can be implemented without a high degree of skill with the individual components that make up systems.

## Virtualization

This newest technology has the ability to create a highly-available solution. Virtual machines can be used to create a failover cluster without the need of a physical machine to place that node onto. However, it is important not to underestimate the hardware needs for using virtualization. You must also account for the SPOF: While an OS failure will have no impact, placing the nodes on the same hardware once again presents a SPOF for the cluster. Costs can be lowered by using virtualization, but the solution needs to be planned so as not to inadvertently create a less-available solution.

## Coming Up Next...

We'll next look at how you can use the metrics you have gathered to gain approval for your HA solution. We will look at how to explain the business impact of those metrics as well as, how to translate the technology into business reports. In addition, we will explore where HA has real business value overall.

## Article 3: Relating High-Availability Metrics to Business Value

---

Once you have a grasp on your availability needs, you need to sell that solution to the decision makers. The reality of this process is that we, as technology experts, completely understand that one technology will meet the business objectives better than another. The trick now is to sell that to the rest of the organization, which will be both less enthusiastic and frankly not as well versed in high availability (HA) as you are. This is where metrics come into play; they serve as a translation from technology solution to business value.

Looking at the metrics from the baselines you developed and adding to that the ROI analysis, you can put together a good business picture that is relatable to those who are not necessarily technical. This is an important part of selling the HA solution because the impact of a service disruption affects the entire organization.

### How to Explain Your Metrics and Show Business Impact

We looked at the metrics to measure the need for HA, now let's look at what these metrics tell you about your services. This is important because many times we can find ourselves presenting our findings while believing that the numbers alone will move people to action. Decision makers want to know not only the numbers but also the meaning behind them. So let's review the metrics for HA and how you can explain these when planning a solution.

As Article 1 discussed, your SLA is your baseline for an IT service. This is the expectation of service set by IT, the business, or both. The metrics show you the reality of how you are doing in comparison with expectations. So how do you use those metrics to show the impact of not reaching your established SLA baseline?

This process requires an information translation. You know what a disruption means to you as an IT department. It may also be apparent to those end users who are directly impacted by the loss. Thus, system-wide HA solutions will be generally easier to sell. But what if the disruption is to a specific department? What if no other business divisions are affected by that disruption? How do you make the case for providing an HA solution for that service? Actually, the same approach works for both types of disruptions. Simply stated, you need to show what those disruptions mean in lost dollars and productivity.




## Translating the Technology of HA to Useful Business Reports

The most important point to focus on when creating reports for establishing a need for an HA solution is that HA is not primarily about the technology. It is primarily about business value. Therefore, you must translate the technology and metrics into a report for management that makes business sense more than technological sense.

The reasoning for this approach is that while you already understand or have come to understand the benefits of HA, many times, those on the business side look to the magic box sitting in front of them and simply expect it to work. Not many business departments outside of technology understand hot-swappable parts, clustering, heartbeat discovery, or many of the other technology-centric terms. That makes it your responsibility to make the report you submit clear in business terms. Using real-time examples of productivity and fiscal losses is an excellent way to make a solid argument for an HA solution. Cost and lost productivity are universally understood in all areas of business. Thus, you need to justify your HA solution based upon loss of money and productivity.

The previous articles spoke about using the availability scale to determine the balance between services and solution. You next need to look at how to take those determinations and turn them into a useful report. A powerful representation for this is the use of charts to visually explain the impact of failure. Let’s look again at the availability chart from Article 1. This chart plainly shows that neither Application A nor Service B is meeting the SLA. But what is the ultimate business impact of not meeting these service levels. By adding lost productivity and total cost of downtime, you begin to show how these service interruptions impact the entire business (see Figure1).



**Figure 1: Availability management metrics showing the business impact of disruptions.**

To show how an HA solution would improve both productivity and cost of downtime, you need to show (in real numbers) what productivity and cost of downtime represent and then how the cost factor of the HA solution does in fact pay for itself. Table 1 shows the productivity and costs of the Figure 1 chart as compared with the cost of a potential HA solution.

This table shows an example for a 75-user organization. Application A is used by all 75 users. Service B is used only by a single department of 15 employees.

Product Measured	Number of Disruptions per Month	Avg. Duration of Disruptions	Monthly Loss of Productivity (Hrs.)	Annual Cost of Downtime	HA Solution Cost	ROI
Application A	3	3 minutes	11.25	\$20,250	\$12,000	\$8,100
Service B	5	1.5 minutes	1.875	\$5,625	\$2,600	\$775

**Table 1: Calculating the ROI of an HA solution.**

#### Note

It is critical when considering the cost of downtime to include in that cost not only the loss of business productivity but also the cost of the IT personnel hours used to restore those services. When presenting findings, we often forget to include the intangibles (although not really billing for the hours used, these hours are being spent on restoring services rather than other IT tasks).

Table 1 represents what you have already considered in your availability scale. Application A has an IT Service Weight of 10. Although the number and duration of disruptions do not seem to have an impact, upon further analysis, the total monthly interruptions cause an annual loss of \$20,250 because this application affects the entire 75-person organization.

Service B has an IT Service Weight of 5 and the number and duration of disruptions is more frequent. However, this service is used by only 15 employees, resulting in an annual loss of only \$5,625. Thus, it does not support a legitimate cause for the same HA solution as Application A.

This report takes us back to what we discussed earlier: HA is not a one-size-fits-all solution. From what we can see in the table, Service B does not present the same business impact as Application A. Ignoring the loss is not a solution either; rather you need to find a solution where the cost creates a positive ROI.

So how do you calculate the numbers seen in Table 1?

- Use the ITIL metrics to measure Number of Disruptions and Average Duration of Disruptions. Add these together to calculate total minutes lost.
- Multiply the total disruption time by the number of users affected. Doing so will give you the Monthly Loss of Productivity value.
- Multiply the Monthly Loss of Productivity value by the average hourly cost of lost production for your organization (this example used \$150 per hour for Application A and \$250 per hour for Service B), then multiply that number by 12 for your Annual Cost of Downtime.
- Use the cost estimate for the solution to determine your HA solution cost.
- Add the expected RTO of the HA solution (this example used 1 hour annually for Application A and 9 hours for Service B) to the HA solution cost to reach your total ROI.

It is important to run the charts and tables for each service being considered for an HA solution. When going through the process, it makes sense to include all your applications, services, and data sets. At first, some may seem unworthy of an HA solution. If you look at the ROI you gain after implementing HA, you might determine that it is more costly not to implement HA. Again, to make this determination, use the real numbers for lost productivity and cost of the disruption.

## Summary

Ultimately, an HA solution will benefit the entire organization. Let the metrics and the numbers do the talking—they offer a simple and straightforward look at where you stand. They also offer a roadmap for reaching your availability goals. Thus, you must be sure to include any potentially impacted business units to make selling the solution easier. If all areas of business are involved, HA will be a company-wide solution that has been worked on cooperatively.

HA is a technology that is flexible enough to meet the needs of any size business from the smallest organization to the largest enterprise. Rather than a one-size-fits-all solution, HA is a set of solutions that you can customize to create a good fit for each application and service. This perspective will enable you to sell the solution to the business units that you want to support that solution, benefitting the entire organization.