

Building The Always-On, Always-Available Enterprise

by Stephanie Balaouras and Rachel A. Dines, June 11, 2014

KEY TAKEAWAYS

So-Called “Disasters” Are Not Rare Occurrences

For those who think, “It won’t happen to me” -- think again. According to the latest Forrester/DRJ survey, one in three companies has declared a disaster during the past five years. And overall, only 31% of firms say they have never declared a disaster.

Follow Three Steps To Build The Always-On, Always-Available Enterprise

Your path for improving the resiliency of your enterprise should include: 1) understanding the costs of downtime for critical services; 2) measuring availability appropriately; and 3) matching the right technologies to support business objectives.

Downtime Is Inevitable -- Be Prepared

Since some downtime is inevitable, it’s important for you to shift your attitude from reacting to downtime toward proactive planning, good processes, and preventive efforts. You may not be able to achieve 100% uptime, but you can at least strive to make services available when your customers most need them.

Building The Always-On, Always-Available Enterprise

Three Steps To Get You On The Path To Running 24x7x365 Services

by [Stephanie Balaouras](#) and Rachel A. Dines
with [Doug Washburn](#) and Andrew Hewitt

WHY READ THIS REPORT

Your extended enterprise consists of customers, partners, suppliers, and employees, and every stakeholder in this ecosystem expects your services to be available 24x7x365 in a seamless fashion. But running an always-on, always-available enterprise can become very complex and very expensive in the blink of an eye. This report will explore three steps that infrastructure and operations (I&O) professionals can take to support the always-on, always-available extended enterprise, including calculating the cost of downtime, measuring availability effectively, and matching business needs to the appropriate technologies.

Table Of Contents

2 **There's Never A Good Time For Downtime**

Today's I&O Professionals Are Supporting The Extended Enterprise

2 **Three Steps To Build Your Always-On, Always-Available Enterprise**

Step 1: Understand The Costs Of Downtime For Critical Services

Step 2: Ensure That Every Component In Customers' End-To-End Service Is Available

Step 3: Match Business Objectives To The Right Mix Of Technologies

RECOMMENDATIONS

8 **Downtime Is Inevitable — Be Prepared**

Notes & Resources

In developing this report, Forrester drew from a wealth of analyst experience, insight, and research through advisory and inquiry discussions with end users, vendors, and industry experts.

Related Research Documents

[The State Of Business Technology Resiliency, Q2 2014](#)

May 12, 2014

[Execute The Five Key Processes Of The Business Technology Resiliency Life Cycle](#)
February 24, 2014

[Elevate Your Business Technology Resiliency Program With An In-Depth Strategic Plan](#)

April 8, 2013



THERE'S NEVER A GOOD TIME FOR DOWNTIME

Imagine the following: You work in the I&O department at a large retailer with a significant eCommerce presence, and at noon today, a critical component of your infrastructure failed. While you scramble to find a solution, your company's website that brings in tens of thousands of dollars a day is greeting all of your potential customers with an error message and the social networks are starting to buzz. But to make matters worse, today is not a normal day — it's Black Friday, one of the largest shopping days of the year in the US.¹

This nightmarish scenario is an extreme example of downtime occurring at the worst possible moment for a business. But the truth of the matter is that there is never a good time for downtime, even planned downtime. As more and more employees become mobile and working-from-home policies relax, the concept of the 9 to 5 workday has eroded.² And as companies become more global, with employees, customers, partners, and suppliers spread across every time zone, it becomes increasingly difficult to accommodate even planned outages for essential maintenance, patches, and upgrades.

For those who think, "It won't happen to me" — think again. According to the latest Forrester/DRJ survey, one in three companies has declared a disaster during the past five years. And overall, only 31% of firms say they have never declared a disaster.³

Today's I&O Professionals Are Supporting The Extended Enterprise

Forrester defines the extended enterprise as an enterprise for which a business function is rarely, if ever, a self-contained workflow within the infrastructure confines of the company.⁴ Most enterprises operate in an extended ecosystem of customers, partners, and employees who rarely operate within the four walls of your organization. Today's I&O pros must meet the availability demands of all these stakeholders, and these constituencies are more similar than you think: They expect to perform their jobs anytime and anywhere. Partners are no different, especially if they rely on you for a critical function or to help them generate revenue. If you're down, they're down. All of your stakeholders share the same expectations regarding their ability to purchase, receive support, or access your data and systems.

THREE STEPS TO BUILD YOUR ALWAYS-ON, ALWAYS-AVAILABLE ENTERPRISE

There is no "easy" button when it comes to running always-on, always-available services; a blend of a mature and stable process, people, and, of course, technologies is required. For companies that have matured their approach to high availability and disaster recovery to the point where they are one and the same — a concept that Forrester refers to as continuous availability — it has taken years of refining policies, adapting responses to downtime, and securing the appropriate levels of investment.

While you can't transform your organization overnight into an always-on, always-available enterprise, these three initial steps will get you on the right path: 1) Understand the costs of downtime for critical services; 2) measure availability appropriately; and 3) match the right technologies to support business objectives.

Step 1: Understand The Costs Of Downtime For Critical Services

The vast majority of companies use some sort of technology to communicate and interact with the extended enterprise. Whether it's a website, call center, checkout line, or mobile phone app, technology has become a critical piece of the way we work with our external customers, partners, and suppliers. Because of this, downtime today means lost or deferred revenue, strained partner relationships, damaged reputation, and lower customer retention rates, not to mention lost employee productivity, lowered morale and confidence in IT's ability, and noncompliance penalties over service-level agreements (SLAs).⁵ Beyond that, downtime can also give away competitive advantage.

Securing investment in the capabilities required to run an always-on, always-available enterprise can be difficult, especially if you don't know your hourly cost of downtime. Because it's such a complex task, Forrester finds that the majority of companies have not calculated it for their critical services. Although trying to calculate the impact of an outage on reputation and customer retention can be a daunting task, just calculating revenue losses or productivity losses can be a worthwhile exercise (see Figure 1).

Remember that not all outages are created equal: Timing and duration have a significant impact on the costs of downtime. In our earlier example, the outage was perfectly timed to affect the largest number of potential customers and thus have the largest business impact. What if this outage occurred at 3 a.m. Eastern time instead of noon Eastern time? Or what if it happened on a different day? Or, what if, instead of the website being down for 4 hours straight on a single day, it was down for 30 minutes on eight different days? Shorter duration outages tend to be less disruptive than longer ones. All of this must be taken into account when calculating the impact of an outage.

Don't try to tackle the entire infrastructure all at once; break down your calculations on a service-by-service basis, starting with the most critical business services. Understanding the costs of downtime will guide the appropriate level of investment in downtime prevention for these services.

Figure 1 What Does A Cost Of Downtime Calculation Look Like?

Messaging services (email, messenger) outage		Commentary
Hours of service impact	1	In hours, enter the duration of the service impact.
Employee productivity losses		
Number of employees affected	1,000	Enter the number of employees who experienced a productivity hit as a result of the impact.
Average burdened salary per employee per hour	\$50	Enter the burdened hourly rate for the average employee. The burdened rate includes salary plus benefits.
Approximate employee productivity impact	20%	As a percentage, enter the impact to employee productivity. For example, the average employee is 20% less productive without messaging services.
Total productivity losses	\$10,000	
Revenue losses		
Annual revenue generation of the service	\$10 million	Estimate the revenue generated or supported on an annual basis. This example assumes that 20% of deals come in through email.
Hours of service operation	8,760	Enter the total number of hours the service is expected to be available on an annual basis.
Hourly revenue generation of the service	\$1,142	
Total revenue losses	\$1,142	
Customers lost		
Number of customers affected	100	Enter the average number of customers who correspond via email per hour.
Expected customer churn	30%	Enter the percentage of customers whose business you have permanently lost. Enterprises that have a transactional relationship with customers and/or operate in an industry with a number of competitive options should enter a higher percentage.
Total number of customers lost	30	
Total cost of downtime for messaging services outage	\$11,142	Plus 30 permanently lost customers


Figure 1 What Does A Cost Of Downtime Calculation Look Like? (Cont.)

Website outage		Commentary
Hours of service impact	2	In hours, enter the duration of the service impact.
Productivity losses		
Number of employees affected	1,000	Enter the number of employees who experienced a productivity hit as a result of the impact.
Average burdened salary per employee per hour	\$50	Enter the burdened hourly rate for the average employee. The burdened rate includes salary plus benefits.
Approximate employee productivity impact	2%	As a percentage, enter the impact to employee productivity. For example, the average employee is 2% less productive without messaging services.
Total productivity losses	\$2,000	
Revenue losses		
Annual revenue generation of the service	\$200 million	Estimate the revenue generated or supported by the service on an annual basis. This example assumes that, on average, 30% of revenue is through the website.
Hours of service operation	8,760	Enter the total number of hours the service is expected to be available on an annual basis.
Hourly revenue generation of the service	\$22,831	
Total revenue losses	\$45,662	
Customers affected		
Numbers of customers affected	1,000	Enter the average number of hourly website visitors.
Expected customer churn	98%	Enter the percentage of customers whose business you have permanently lost. Enterprises that have a transactional relationship with customers and/or operate in an industry with a number of competitive options should enter a higher percentage.
Total number of customers affected	1,960	
Total cost of downtime for website outage	\$47,662	Plus 1,960 affected customers

Step 2: Ensure That Every Component In Customers' End-To-End Service Is Available

Think back to our original scenario about the eCommerce website that failed on Black Friday. What caused this outage? Was there a server failure? Did the network fail? Or was there a database error? Or was the whole data center a smoking hole? In the end, the results are the same: Customers are not receiving the service that they expected. Few companies track a single service's uptime end-to-end, meaning from every infrastructure and software component that works together to deliver a single service (see Figure 2). This, however, is the single most important thing that an I&O pro can track because it is the metric that gets closest to the actual customer experience. Every component in the chain is critical. This outside-in perspective is necessary in the age of the customer where businesses compete and differentiate themselves on the experience of technology-enabled business processes and transactions more than ever.⁶

Figure 2 Ensure That Every Component In The Customer's End-To-End Service Is Available



61151

Source: Forrester Research, Inc.

Step 3: Match Business Objectives To The Right Mix Of Technologies

Once you've calculated your cost of downtime and shifted your focus to end-to-end availability, the next step is to select the right technologies to support your critical services. There are many technologies that are available today, but the difficult part is finding the right technologies that support your availability objectives but also match what the business is willing to pay to protect critical services (see Figure 3).⁷

It's important to remember that none of these technologies can stand alone, and most enterprises invest in a mix of many different approaches to match the service needs for availability, recovery times, and data retention with the appropriate technology. While there are many more technologies to consider, Forrester finds that our enterprise clients mention the following technology approaches to deliver always-on, always-available services:⁸

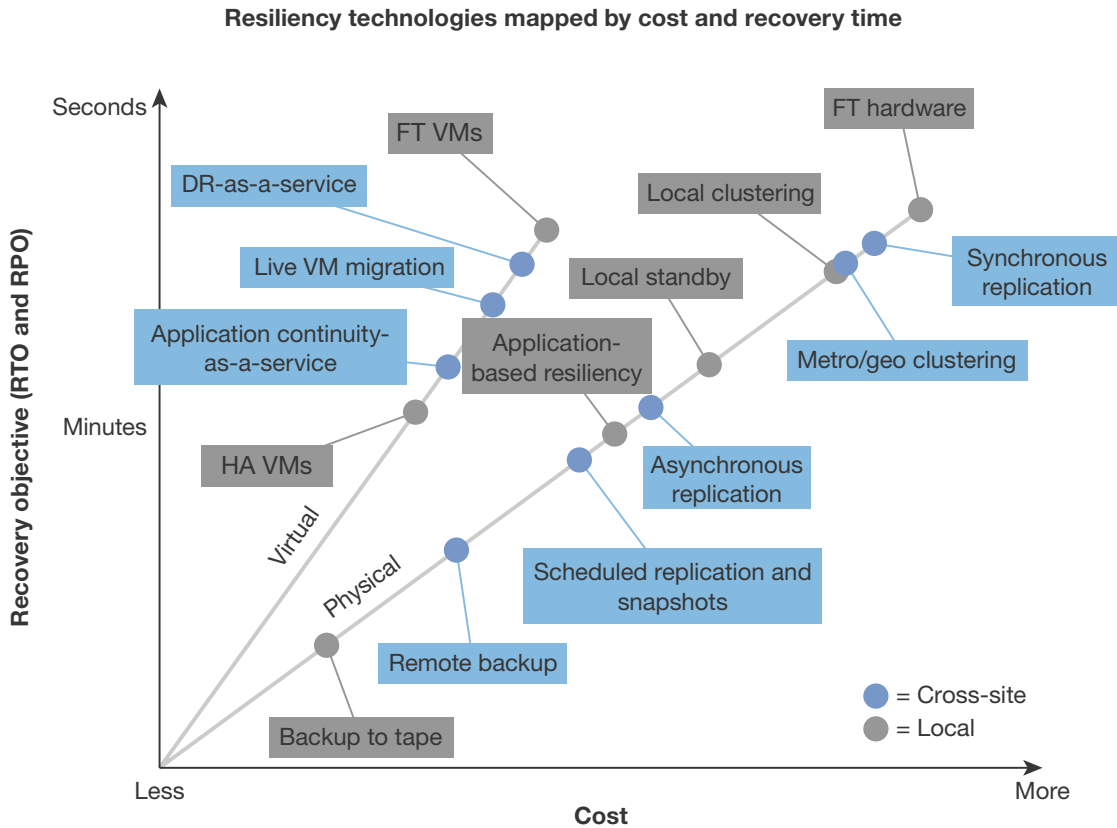
- **Active-active architectures.** The ultimate panacea for the always-on, always-available infrastructure for many companies is an active-active architecture. Active-active configurations come in many flavors, but ultimately, the architecture includes two sets of infrastructure, often

geographically separated, that can load balance a single workload so that if one site is lost, the other continues to deliver the service.⁹ While they can be complex to manage and configure and won't be possible for all types of applications, active-active architectures are nonetheless extremely effective at preventing downtime and improving resiliency.

- **Rapid virtual machine reboot.** Several new technologies on the market today are dramatically blurring the lines between what is traditionally known as disaster recovery and high availability (HA). Rapid virtual machine rebooting, which can occur within the data centers or across data centers in geographic regions, is one of these technologies. Configured correctly, the downtime experienced when using automated rapid virtual machine rebooting solutions is literally the amount of time it takes to boot a virtual machine (VM) — usually minutes.
- **Application and service monitoring.** How do you know when an outage or a service impairment is occurring? Most companies wait for an employee or customer to complain, but what if you could detect failures and problems and fix them before a customer is even affected? Application and service monitoring solutions constantly monitor the performance of critical applications and can send alerts when something isn't performing according to service levels.¹⁰ Many I&O departments have invested in infrastructure management software, but few have taken the customer-centric approach, which is to manage the service or application itself. Usually the longest portion of mean time to repair (MTTR) is actually the time to detect a problem, or mean time to know (MTTK), something that can be shortened with application or service monitoring.¹¹
- **System management and automation tools.** This is a broad bucket that includes tools for faster patching, live kernel patching, and quick system rollback that help I&O professionals complete maintenance within a shorter service window. While the intent of these tools is not for resiliency, they are nevertheless important because, as any I&O pro would tell you, most downtime is actually planned downtime, and these days, stakeholders don't often make a distinction between planned and unplanned downtime.
- **Cloud and SaaS services.** These days, it seems like cloud is the answer to everything, which leaves many people walking away with the false notion that cloud-based services come with inherent resiliency. This is not a valid assumption. In reality, however, you can often pay an additional fee to the service provider to make the solution more resilient. An example is Amazon Elastic Compute Cloud (Amazon EC2)'s multiple availability zones. In this model, customers pay an additional fee for the ability to operate out of multiple zones in case of disaster in any single zone. During Amazon EC2's infamous outage during spring 2011, companies that rely almost entirely on Amazon services, like Netflix, were able to maintain service availability.¹²

If you aren't ready to outsource your critical business and technology services to the cloud, there are a variety of cloud-based disaster recovery services and application continuity-as-a-service solutions on the market today that can provide high levels of resiliency for lower costs than traditional DR services.¹³

Figure 3 Technologies That Can Support The Always-On, Always-Available Extended Enterprise



61151

Source: Forrester Research, Inc.

RECOMMENDATIONS

DOWNTIME IS INEVITABLE — BE PREPARED

The goal of the always-on, always-available enterprise is 100% service availability. While there are a few companies that have come very close to delivering 100% uptime via active-active architectures, sustaining true 100% uptime for any extended period of time is virtually impossible — there are too many things that can go wrong, from the infrastructure to the applications to natural disasters, human error, or even planned maintenance.

Since some downtime is inevitable, it’s important to shift your attitude from reacting to downtime toward proactive planning, good processes, and preventive efforts. You may not be able to achieve 100% uptime, but you can at least strive to make services available when your customers most need them and have rapid response measures in place to make sure services are brought back online as quickly as possible.

ENDNOTES

- ¹ Black Friday is the day following the Thanksgiving holiday in the US. It has earned this reputation because it's often the largest single shopping day of the year and thus has potential to put retailers "in the black" (i.e., to turn a profit).
- ² According to Forrester survey data, 29% of global information workers work from multiple locations on multiple devices with many applications. See the November 4, 2013, "[Workforce Personas And The Mobile App Gap](#)" report.
- ³ For more information on the state of business technology resiliency, see the May 12, 2014, "[The State Of Business Technology Resiliency, Q2 2014](#)" report.
- ⁴ Successful businesses don't work in isolation. Today's businesses must constantly create new products and services, expand their geographic presence, streamline operations, and deliver top-notch customer services. To do this, your business will increasingly use third-party and cloud services to reduce cost and increase speed-to-market. Your business will unleash the creativity of your employees and customers with mobile, social, and rich media technologies. More and more devices come equipped with microprocessors, which mean cameras, cars, home electronics, and even musical instruments will all become conduits for businesses to deliver services and engage customers. To stay relevant, your enterprise must continuously extend itself to include new peripherals and meet new business scenarios. To capture this trend and its implications for security and risk professionals, Forrester defines the notion of the extended enterprise this way: An extended enterprise is one for which a business function is rarely, if ever, a self-contained workflow within the infrastructure confines of the company. See the November 9, 2011, "[The Extended Enterprise: A Security Journey](#)" report.
- ⁵ Calculating the cost of downtime and data loss is not a simple exercise: Some effects are more easily quantifiable than others. Still, the exercise is essential because it's the cornerstone of the business case. See the June 11, 2012, "[Justifying The Business Case For Business Technology Resiliency](#)" report.
- ⁶ Increasingly powerful customers push all institutions, especially businesses, into the age of the customer, which Forrester defines as a 20-year business cycle in which the most successful enterprises will reinvent themselves to systematically understand and serve increasingly powerful customers. For more information, see the October 10, 2013, "[Technology Management In The Age Of The Customer](#)" report.
- ⁷ Many enterprises find it useful to group services or applications into tiers of criticality and assign standard recovery time objectives (RTOs) and recovery point objectives (RPOs) as well as SLAs for availability. Organizations can then map appropriate technologies to the tiers of criticality using the business requirement. See the June 14, 2011, "[Develop Your Road Map For Business Technology Resiliency Tools](#)" report.
- ⁸ There is a host of new IT service continuity management technologies in the market, in addition to the dozens of mature offerings. To help I&O professionals select the best technologies to meet business needs and to plan their next decade of investments in ITSCM technologies, Forrester investigated the current state of the 18 most important technologies in the space. We examined past research and interviewed vendors and experts in the field. We used the collected data to assess four factors: 1) the current state of the technology; 2) the technology's potential impact on customers' businesses; 3) the time experts think the technology will need to reach the next stage of maturity; and 4) the technology's overall trajectory — from minimal success to significant success. See the June 14, 2011, "[Develop Your Road Map For Business Technology Resiliency Tools](#)" report.

- ⁹ There are several ways you can approach the development of truly active-active data center configurations. It's not an all-or-nothing approach, and it's more of an evolution than a revolution. See the October 5, 2007, "[Maximizing Data Center Investments For Disaster Recovery And Business Resiliency](#)" report.
- ¹⁰ In today's challenging economic times, business productivity takes precedence over many other considerations. To achieve a quality of service that's in line with this goal, IT infrastructure and operations professionals usually monitor the weakest links in the delivery chain. Over time, the focal point has moved from networks to systems to application code. But the current complexity of business services is such that issues can spring from anywhere in the service delivery chain. Consequently, application performance management (APM) must now monitor, detect, and identify any abnormal behavior across the whole hardware and software infrastructure. See the December 7, 2011, "[Market Overview: Application Performance Management, Q4 2011](#)" report.
- ¹¹ Mean time to know consumes excessive time, energy, and cost largely because of the scenario best described as the "all the king's horses and all the king's men" problem. This line from the popular nursery rhyme Humpty Dumpty effectively captures the nature of the chaotic response that characterizes an all-too-common response to incidents, especially those with high profiles. Once an incident has been identified (mean time to identify), IT staff members from multiple functional groups are assembled into an ad hoc team to resolve it. The intent of such a team is admirable, but its effectiveness is doubtful. With insufficient visibility into the true nature of the incident, the team scatters in different directions, none with enough information to attack the root cause. Eventually, someone either stumbles upon the cause, locates it through good detective work, or the incident clears up on its own. Only one thing is certain with "all the king's horses" responses — they're expensive. See the June 30, 2011, "[Enhance Service Visibility To Reduce MTTR](#)" report.
- ¹² On Thursday, April 21, Amazon experienced a large outage in AWS US-East. This outage was highly publicized because it took down or severely hampered a number of popular websites that depend on AWS for hosting. Source: Adrian Cockcroft, Cory Hicks, and Greg Orzell, "Lessons Netflix Learned from the AWS Outage," The Netflix Tech Blog, April 29, 2011 (<http://techblog.netflix.com/2011/04/lessons-netflix-learned-from-aws-outage.html>).
- ¹³ The ability to provide low-cost, low-effort disaster recovery using the cloud is an appealing prospect for many organizations. DRaaS providers aim to give companies the ability to quickly fail their on-premises production systems to virtual machines hosted in the cloud. This technology is still emerging, so there are few mature offerings so far. While it's relatively easy to send incremental data to the cloud, bringing it all back quickly in the event of a primary site failure is a big challenge, which is why using the service provider as a replication target and a failover site is an attractive option that could see significant growth
- Application continuity-as-a-service represents the opportunity for enterprises to leverage cloud services for their continuity and resiliency needs but to also maintain control over their production environment by keeping it in-house and on-premises. This model will be very appealing for companies that want to use cloud economics but not give up control over their production environment. Some providers offer active-active solutions with near-instant failover and load balancing. Thus far, these solutions, such as Mimecast Email Continuity and Symantec MessageLabs, are mainly focused on email, but they will soon expand to

other applications such as collaboration platforms (e.g., SharePoint) and line-of-business tools (e.g., customer relationship management, human capital management). As organizations become more comfortable with sending their data to the cloud, application continuity-as-a-service will start to gain momentum. See the June 14, 2011, [“Develop Your Road Map For Business Technology Resiliency Tools”](#) report.

About Forrester

A global research and advisory firm, Forrester inspires leaders, informs better decisions, and helps the world's top companies turn the complexity of change into business advantage. Our research-based insight and objective advice enable IT professionals to lead more successfully within IT and extend their impact beyond the traditional IT organization. Tailored to your individual role, our resources allow you to focus on important business issues — margin, speed, growth — first, technology second.

FOR MORE INFORMATION

To find out how Forrester Research can help you be successful every day, please contact the office nearest you, or visit us at www.forrester.com. For a complete list of worldwide locations, visit www.forrester.com/about.

CLIENT SUPPORT

For information on hard-copy or electronic reprints, please contact Client Support at +1 866.367.7378, +1 617.613.5730, or clientsupport@forrester.com. We offer quantity discounts and special pricing for academic and nonprofit institutions.

Forrester Focuses On Infrastructure & Operations Professionals

You are responsible for identifying — and justifying — which technologies and process changes will help you transform and industrialize your company's infrastructure and create a more productive, resilient, and effective IT organization. Forrester's subject-matter expertise and deep understanding of your role will help you create forward-thinking strategies; weigh opportunity against risk; justify decisions; and optimize your individual, team, and corporate performance.

« IAN OLIVER, client persona representing Infrastructure & Operations Professionals

