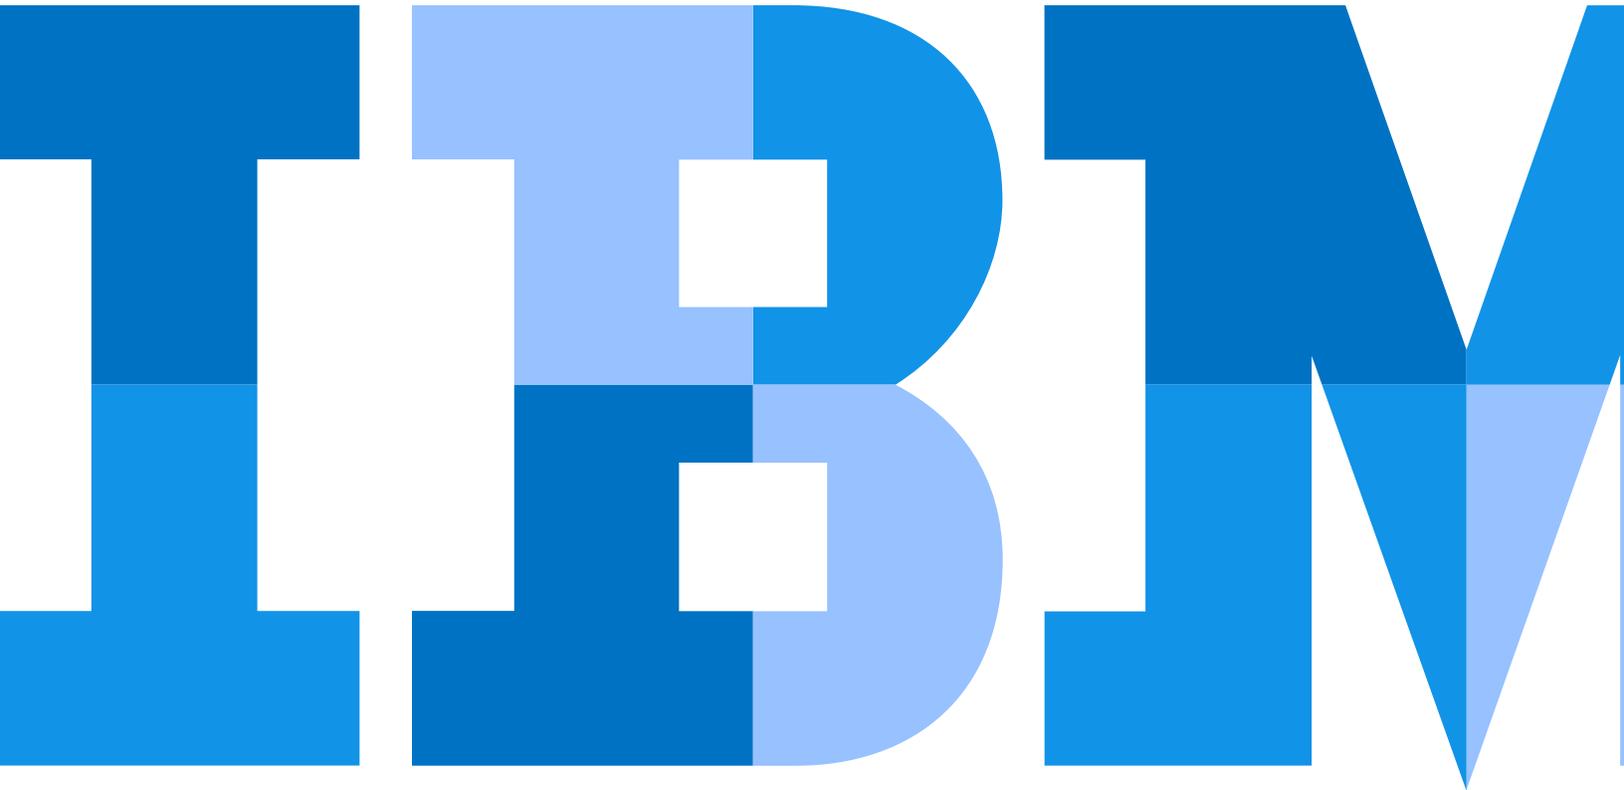


Your journey to always-on in four steps

An always-on platform for an always-on world



Contents

- 2 Executive summary
- 2 Industry trends
- 3 Enterprise needs and challenges
- 5 Step 1. Assess and evaluate
- 7 Step 2. Plan and design
- 10 Step 3. Implement and test
- 10 Step 4. Manage and sustain
- 11 How IBM can help
- 11 For more information

Executive summary

Today's information-infused society hinges on cloud, analytics, mobile and social. On-the-go consumers are almost always within reach of an internet-connected device. Users expect quick, consistent access to their business-critical data and applications around the clock. That's why outages, whether they're planned or unplanned, can negatively affect revenue, consumer trust and your reputation.

“Less than 10 percent of companies report that their IT infrastructure is fully prepared to meet the demands of mobile technology, social media, big data and cloud computing.”¹

An always-on platform that doesn't compromise the end-user experience or allow for service disruption is a strong competitive differentiator for your business. What do we mean by an always-on platform? Always-on is about achieving continuous operations so the business service being consumed is not interrupted, making downtime transparent to the end user. Always-on should be able to:

- Withstand component failures affecting networking components, hardware, operating systems and so on
- Withstand catastrophes such as hurricanes, tornadoes, floods and other natural disasters
- Introduce changes — like scheduled updates and security patches—without disruption

But planning, designing, implementing and managing an always-on platform that can meet these goals can be challenging. Your organization needs to have the right approach that involves the full spectrum of people, processes, applications and technology. For always-on to be successful, it needs to be an integral part of your organization and requires enterprise-level commitment.

The purpose of this white paper is to describe the current market trends and enterprise requirements that are influencing the need for an always-on platform. We then provide four time-tested steps that detail how your organization can shift from your legacy high-availability (HA) and disaster recovery (DR) solution to an always-on design.

Industry trends

Expansion of digital forms of engagement drives demand for always-on

In the past, enterprises relied solely on systems of record to manage and store corporate data. But with the introduction of big data, cloud, mobile and social, it's clear that new digital systems of engagement are influencing the way businesses function and generating data at an exponential rate. Employees, consumers and partners move freely among various devices, such as smartphones, tablets, notebooks and desktops. With constant access to business applications, digital channels are core to the entire lifecycle of sales, solutions and support.

Here we define several trends that are shaping business processes and solutions today. They include:

- **Cloud.** Cloud computing supports converged infrastructures and shared services over the internet. Cloud can be deployed on-premises, off-premises or hybrid.
- **Social media channels.** Social media channels such as Facebook, Twitter and Instagram support real-time interaction among users worldwide.
- **Mobile apps and rate of change.** Consumers are accustomed to getting regular updates pushed to their phones, and they expect an exceptional experience anytime, anywhere —making the dependency on mobile applications more prevalent.
- **Software-defined environments (SDEs).** SDEs can optimize storage, compute and network. They allow IT organizations to deliver services to the line of business (LoB) in one of the most efficient ways possible.

- **Continuous delivery/development and operations (Dev/ops).** This process allows agile developers to do continuous iterative releases with no downtime, rather than large “waterfall” releases that require planned outage windows. This process can bring development and operations teams closer together and support fast-paced change.
- **Continuous operations.** Organizations must maintain operational business services all the time without being impacted by planned or unplanned downtime.

To keep pace with these trends, organizations must evolve IT services to support consumer and business demands with more agility. It's time to start managing systems of engagement differently—and more efficiently—than legacy systems of record.

Enterprise needs and challenges

Based on new digital trends and systems of engagement as characterized in the previous section, the implications of an outage's effect on your business are growing. Digital channels are taking on a larger market share and are driving revenue generation. Losing services to an outage can have many consequences, such as revenue and productivity loss—and today, downtime isn't tolerated. Customers will quickly abandon you and use competitive services to fulfill their needs.

Putting the consumer's needs first

User perception of availability is a key tenet of your business. For example, if users want to access a mobile banking application to make deposits, they expect it to be available and functional no matter what time they want to log in. And in today's increasingly global economy, there can no longer be “planned maintenance windows” that interrupt services. If users attempt to log into their services and discover that the system is down, they may become dissatisfied with the vendor. They might log on to their social network and complain about the situation to their friends—and news of downtime spreads quickly over social media.

“Although overall service quality has improved for mission-critical applications, such as systems of record, service availability still isn't ‘good enough’ for most enterprises because the cost of downtime continues to escalate.”²

Addressing the cost challenge

The need for an always-on solution is clear. But there are several cost-related inhibitors at stake. Organizations may have skewed understandings of the cost involved or have difficulty quantifying the value that can be attained through always-on. Without proper planning and a lifecycle-based approach, building an always-on platform ad hoc can lead to high maintenance needs and unattainable service-level agreements (SLAs). Always-on is much more than just hardware and software. The facility costs and resource requirements often exceed the initial implementation costs. With the public cloud, facility costs are no longer an inhibiting factor.

To help address the cost challenge, always-on needs to be viewed in terms of total cost of ownership (TCO) and organizations must consider the risks of not having always-on. Impacts related to downtime may include loss of revenue or e-commerce sales, regulatory fines for downtime or deadlines missed, customer compensation, reputation damage, lost productivity for employees and so on.

Questions to consider

As your organization considers business drivers today, it's important to look closely at your own requirements. Some key questions to consider include:

1. Do you want zero outage changes to avoid planned maintenance windows?
2. Do you want at least fast failover instead of disaster recovery for your critical services?
3. If there's a catastrophic disaster, do your customers expect your services to be always-on?
4. Do you need the ability to continue service to a customer, uninterrupted?
5. Are component failures causing extended downtime and disruption to business services?
6. Are your digital channels growing, increasing your digital presence and driving revenue?
7. Are LOBs driving increased changes within the IT platform causing operational disruption?

If any of these questions apply to your organization, it's time to change your approach. Based on IBM's experience in this field, we provide an always-on lifecycle methodology that brings together people, processes, information and technology in a structured manner, designed to help maintain consistent synchronization of IT and the business. Our methodology consists of four phases, starting with assess and evaluate and then moving through the subsequent plan and design, implement and test, and manage and sustain phases. You'll need to continually revisit each phase regularly to check that the always-on platform remains synchronized with evolving business requirements and rapidly changing technology.

“Based on an IDC study of Fortune 1000 organizations across multiple industries, the average cost of a critical application failure per hour is USD 500,000 to USD 1 million.”³

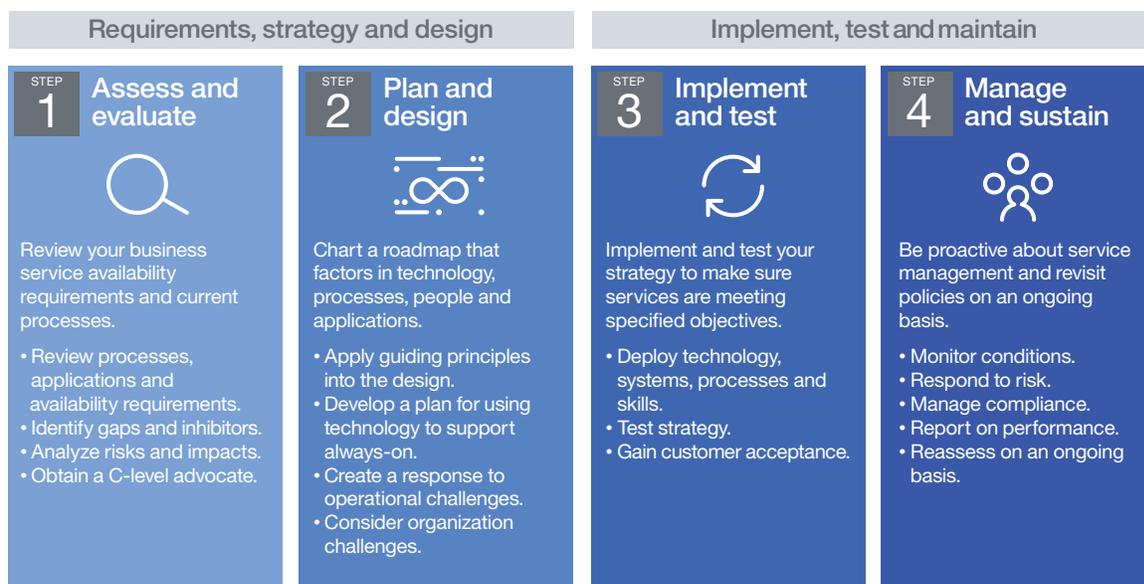


Figure 1. The journey to always-on in four steps.

Step 1. Assess and evaluate

The first step is to assess and evaluate your current environment. You need to get a better understanding of the business services that are required to be continually available, and then assess your current capabilities to meet business objectives. Here are some examples of how to get started.

- **Review your current processes.** Your assessment should include an evaluation of your organization's current processes and requirements for availability. You should also:
 - Assess and classify applications based on availability and continuous operations requirements.

- Assess the maturity of your operational practices. CMM Level 3 is required.
- Evaluate the organizational maturity level in availability, for example, accountability.
- Consider costs based on business impact.
- **Conduct deep dives into each application domain.** Carefully review each application domain and determine what an application can do. For example:
 - Active/standby. This method is the traditional recovery method in IT that has a “warm” standby environment ready in case of failure. Modern organizations use this pattern for fast failover, A/B application releases and planned platform changes.

- Partitioned active. Each site application cluster runs independently, as do the databases. A percentage of some user activity is directed to one application cluster, and others are directed to another application cluster.
- Active/query. Each site application cluster is live. Reads are performed from a local database and writes are performed on a primary database, system of record, only.
- Active/active. This method is considered the best-case scenario where both data centers are providing the same services, where data can read and write at either location, and both locations are synchronized.
- **Consider ongoing services and upgrades.** You should define and document availability requirements at the plan and design step for new applications and services.
- **Filter the noise.** Not all applications require always-on availability, so you should determine these applications during the assessment. For example, some internal applications might not be the target for these advanced requirements, at least not initially. At the same time, it's generally believed that most applications will move to this model over time, based on organizational priorities. It's not feasible from a cost perspective to enable every business service application to become always-on. For example, a business service application must be modernized to function in an always-on platform, similar to cloud native modernization, which may be a disruptive leap requiring funding, testing and an integrated operations model.

- **Identify gaps and inhibitors.** Identify application, operational and organizational gaps that may inhibit moving to the always-on platform.
- **Analyze risks.** Analyze current and potential risks.

During this step, it's important to build confidence within the technical community and obtain a C-level advocate who can educate other decision makers about creating an always-on point of view (POV) for the company. Being able to drive conversations with C-level executives, such as the chief marketing officer (CMO) and chief financial officer (CFO), can strengthen the relationship between business stakeholders and IT. Remember that business is driving the requirements and needs, while IT is delivering the desired end result. Part of these conversations might include showing business value, for example, considering possible outcomes for revenue generation and improving the digital experience that drives incremental revenue. The conversations might also include discussing the cost and reputational repercussions of not providing always-on. We propose a positive approach, where organizations put their customers' experiences first. By choosing this approach, continuous availability becomes a competitive differentiator.

You can also initiate a discussion about the way users consume services, resulting from the shift to mobile and the expectations that are associated with highly available applications and data. With these considerations in mind, your organization may be ready to start planning a roadmap.

Step 2. Plan and design

Designing for always-on requires a shift in architectural thinking. You need to think in terms of business service maintaining availability all the time—during changes, planned outages and unplanned outages. In the plan and design step, it's time to take the information gathered from the assess and evaluate Step 1 and use it to create the groundwork for your actionable plan. We recommend that you begin with general always-on concepts, supported by Brewer's CAP Theorem.

Brewer's CAP Theorem on distributed systems limits the technology solution to providing only two of the three guarantees:

1. **Consistency:** All distributed nodes have a single up-to-date copy of all data at all times.
2. **Availability:** Every request receives a success or failure response.
3. **Partition tolerance:** The system continues to run despite arbitrary message loss or a failure of part of the system. For example, the network stops delivering messages between data centers or cloud regions.

Given Brewer's limits, it's extremely important to accurately assess the business requirements of every business service to determine which architectural pattern to apply to each specific service based on business requirements. In addition to the business requirements, the application architecture and platform must be assessed to help determine how to mitigate whichever of Brewer's three guarantees can't be met.⁴

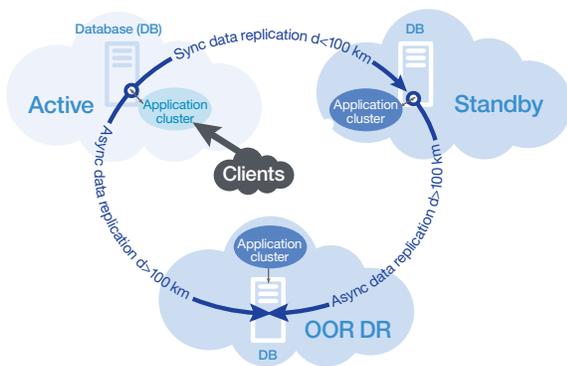
We recommend that the plan and design phase incorporates the following activities:

- **Apply guiding principles into the design architecture of always-on.**
 - Think differently—everything breaks. This principle is what guides the practitioner, architect or engineer to think beyond the traditional HA and DR practices they've learned in the past, and instead embrace innovative concepts that enable resiliency and flexibility. Applications become responsible for resiliency.
 - For example, your organization may need to include systems and infrastructure, such as out-of-region (OOR) data centers that were once only considered strictly for use in case of disaster. But by moving to always-on, OoR data centers may be integrated into live workflow and used for resiliency when performing changes, rather than planned maintenance outage.
 - For more information, review the 21 guiding principles discussed by IBM in the following IBM® Redpaper™ report: [Always On: Assess, Design, Implement, and Manage Continuous Availability](#).

- **Develop a plan for how technology is used to achieve always-on.**
 - Determine product capabilities. For example, review Java applications that support multisite clustering through grid technologies.
 - Consider application design principles as the applications must enable resiliency. Examples include: service-oriented architecture (SOA), loosely coupled systems of engagement from systems of record through SOA or an application programming interface (API), and cloud-enabled and cloud-native applications.
- **Create a response to operational challenges.**
 - Support continuous deployment.
 - Integrate automation, through the stack. Automated testing helps improve stability and quality, remove manual steps and accelerate the delivery cycle.
 - Create a plan for rollback, which is necessary to go back to a workable version following an outage—hopefully an outage that was transparent to users. In addition, developers need the ability to recreate an environment to fix a defect.
 - Facilitate concurrent versioning. This process allows organizations to support phased rollouts of new application versions. Often, the new version of the application is deployed gradually to production to reduce potential issues to the production environment in the event a bug is introduced, also known as “Canary” releases.
 - Establish relationship linkages with the LOB application development and the platform provider.
- **Factor in organizational challenges.**
 - You need to identify change agents, such as senior executive sponsors and IT leaders who embrace change and can lead the evolution.
 - You also need to improve your understanding of skills across the enterprise. For example, even if the application infrastructure allows for always-on and nondisruptive upgrades, application owners aren’t able to use it unless they change the way applications are designed and built. Nonfunctional requirements become as important as functional requirements.
 - Support end-to-end accountability through adopting site reliability engineering practices: engagement—steady state.
 - Facilitate continuous improvement.
 - Focus on the availability of the business services—not the availability of the parts.
- **Plan for the four patterns of near-continuous or always-on availability.** To evolve service to always-on, it’s important to consider all parts of the current architecture design. For example, consider including an OOR data center in your always-on design. IBM suggests four patterns of development to evolve from current resiliency architectures to always-on. See examples and Figure 2.

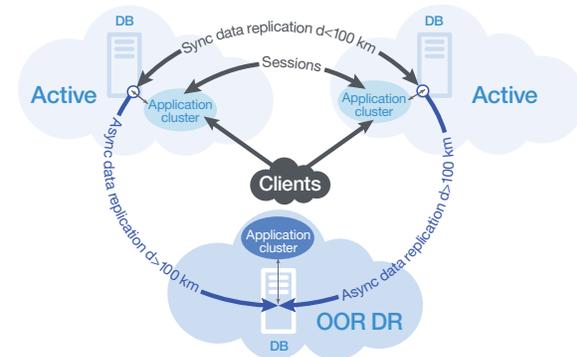
State 1: Active/standby with out-of-region (OOR) disaster recovery (DR)

This pattern is the standard and traditional active/standby where a disaster must be declared to switch to OOR. It's used here to indicate a possible starting point of an organization's journey toward continuous availability.



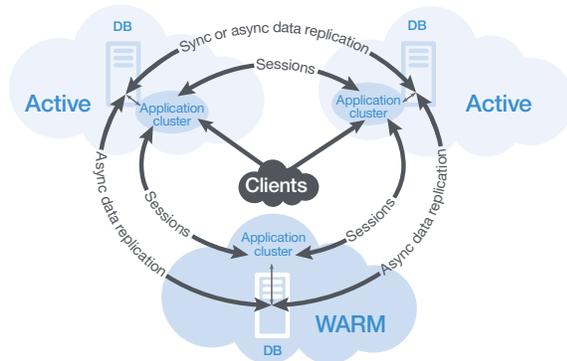
State 2: Two active with out-of-region (OOR) disaster recovery (DR)

This pattern is the common pattern that's typically seen in the mature financial sector where continuous availability is required during business hours and the recovery point objective (RPO) equals 0 or data consistency requirements are atomicity, consistency, isolation and durability (ACID) data consistency.



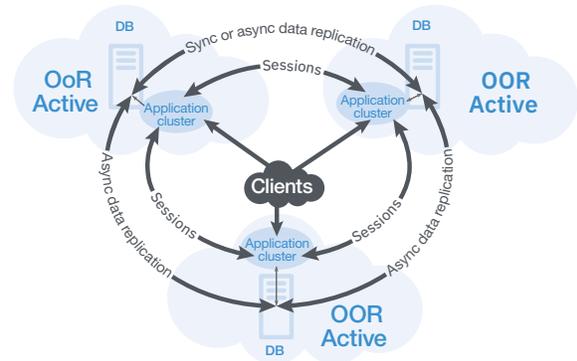
State 3: Two active/query out of region (OOR)

This pattern is a more mature version of the previous two active with OOR DR patterns. Rather than using the OOR "cloud" or data center in a DR scenario, it's integrated into day-to-day operations and can be used for analytics, reporting, batch processing, read-only queries, and in fact might be used as an active component when maintenance is required on a component affecting both the active pairs within the metro.



State 4: Three active with out of region (OOR)

This pattern is the three-active model that has been used to keep ibm.com always-on since June 2001. The key business decision enabling this pattern is that of eventual data consistency. Data can be written to any of the three "clouds," it's captured at the source, and it's applied to its two peers with a replication delay based on the distance between the data centers.



→ One-way data transfer ↔ Two-way logical data transfer | inner circle (sessions) = systems of engagement | outer circle = systems of record

Figure 2. A high-level perspective of the four near- and continuous-availability patterns.

Throughout this process, it's important to consider lifecycle operations and how your approach changes over time. Design availability and continuous delivery functionality into the application and infrastructure architecture. Your roadmap should address the steps required to advance the organizational maturity and iterate as needed.

Step 3. Implement and test

After your plan is in place, it's time to get started. One of the important factors at this step is to think about the end results—realistic testing and being able to “run the business” when faced with a major service disruption. We recommend that you:

- **Begin implementation of your always-on plan.**
 - Implement required technology, systems and processes to achieve an always-on plan. For example, build one cloud or test site at a time, then interconnect them.
 - Invest in and retain skills to support continuous delivery and availability in line with the current and envisioned enterprise architecture.
 - Think about continuous operations, which may expand to include continuous delivery, continuous operations and zero outage changes; and continuous deployment, releasing on the platform for production use.
 - Consider lifecycle operations in transition and readiness testing.
- **Test the always-on strategy and gain customer acceptance.**
 - Know how it works, know how it fails, so operations can master and automate mitigation techniques.
 - Implement agile incident management processes as part of production failure testing.
- Architect and execute tests to make sure services meet specified objectives.
- Ensure that the strategy is supporting always-on services.
- Execute tests of perform walkthrough drills to identify weaknesses for improvement and preparedness.

Step 4. Manage and sustain

Achieving always-on isn't a “once and done” project. Proactive service management is a valuable guiding principle for successful implementations.

We suggest that you:

- **Monitor current conditions to detect and respond to risks**
 - It's important to monitor and trend events to detect abnormalities before they create incidents and problems. Shifting from reactive to proactive service management can save costs and time in the long run.
- **Manage compliance with regulatory requirements.**
- **Report on performance.**
 - Reporting dashboards and metrics help demonstrate readiness and the results of business investments.
- **Reassess.**
 - By revisiting your policies on an ongoing basis, you can consistently define and document availability requirements and costs based on the business impact as part of the business requirement process. And you can update availability requirements as part of an ongoing business impact and assessment process.

How IBM can help

IBM has designed, implemented and managed always-on environments for the past 19 years. We can draw on years of expertise to help you move to an always-on platform. We can identify your most critical business services by understanding the business service requirements for availability, reviewing the current business service application architecture, understanding the current operational and organizational structure, and determining the relationships between business and IT.

Using the information gathered through assessment activity, IBM specialists can analyze and identify the inhibitors of achieving always-on and begin to develop roadmaps to close. Through interactive workshops, we use current guiding principles and patterns to align business services to the most appropriate pattern. This experience helps provide insight into which business services can and can't achieve always-on; identify investment that's needed to align applications, infrastructure, people and processes; and define your individual journey to always-on business services.

With more than 40 years of business continuity and disaster recovery experience, IBM is highly qualified to help you create an effective resilience testing program. IBM's approach to helping clients build resilient businesses is robust and includes coordinating and managing thousands of resiliency test exercises in our recovery and managed services centers around the world. Furthermore, our own corporate portal, ibm.com, has been always-on since June 2001. We've achieved this accomplishment because of our operational model with people and processes, architectural model and ability to mandate application requirements to the development teams that want to board their business applications on our robust platform.

Our consultants use tested methods and tools to help you define, implement and manage a feature-rich resiliency solution.

For more information

To learn more about IBM Business Continuity and Resiliency Services, please contact your IBM representative or IBM Business Partner, or visit ibm.com/services/continuity.

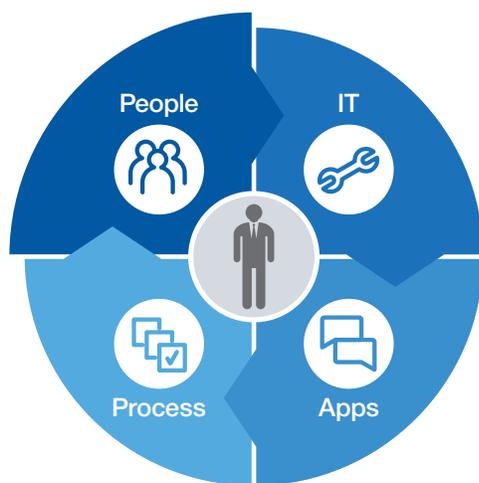


Figure 3. An always-on approach requires the full spectrum of people, processes, applications and technology.



© Copyright IBM Corporation 2018

IBM Global Technology Services
IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the United States of America
June 2018

IBM, the IBM logo, ibm.com, Redbooks, and Redpaper are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at ibm.com/legal/copytrade.shtml.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

It is the user’s responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs. THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

IBM Global Financing offerings are provided through IBM Credit LLC in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates and availability are based on a client’s credit rating, financing terms, offering type, equipment and product type and options, and may vary by country. Non-hardware items must be one-time, non-recurring charges and are financed by means of loans. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice and may not be available in all countries.

¹ “The IT infrastructure conversation,” IBM Institute for Business Value, July 2014, <http://www-03.ibm.com/systems/infrastructure/us/en/it-infrastructure-matters/it-infrastructure-report.html>

² “Hype Cycle for IT Service 1 Continuity Management, 2014,” John P. Morency, Carl Claunch and Pushan Rinnen, Gartner, 10 September 2014, #G00263215

³ “DevOps and the Cost of Downtime: Fortune 1000 Best Practice Metrics Quantified,” Stephen Elliot, IDC, December 2014, IDC #253155

⁴ “Always On: Assess, Design, Implement, and Manage Continuous Availability,” Herbie Pearthree, IBM Redbooks®, <http://www.redbooks.ibm.com/abstracts/redp5109.html?Open>



Please Recycle