

Downtime Prevention Buyer's Guide

6 QUESTIONS

to help you choose the right availability
protection for your applications




For an **Always-On** World

6 questions to help you choose the right availability protection for your applications

Prevention of unplanned downtime is a growing concern in today's always-on world. You know you need a way to keep critical applications up and running, but with so many options on the market, how can you determine which availability solution is right for your organization?

This buyer's guide presents a series of questions you should ask vendors when evaluating solutions to protect your applications against costly downtime. It highlights key considerations and provides valuable insights into the strengths and limitations of various availability approaches. Vendors' responses to these questions will enable you to compare solutions and identify those that best meet your availability requirements, recovery time objectives, IT management capabilities, and return on investment goals, while integrating seamlessly within your existing IT infrastructure.



Question

1



What level of uninterrupted application processing can your solution guarantee?

Downtime Prevention Buyer's Guide

There are a variety of availability solutions on the market today, each of which delivers a different level of application uptime. When evaluating solutions, it is helpful to ask vendors how many “nines” of availability their offerings provide — a figure that represents the average amount of uptime their customers should expect per year. This is an important first step in determining which solution best meets your organization’s specific requirements.

If your availability requirements are relatively low, you may be able to get by using a standard server with duplicate internal components. These servers typically deliver two nines — 99% — or more of availability for the applications running on them, which can result in as much as 87.6 hours of unplanned downtime per year. Continuous data replication delivers three nines — 99.9% availability — which equates to 8 hours and 45 minutes of downtime annually.


For those with more rigorous availability requirements, traditional high-availability clusters, which link two or more physical servers in a single, fault-resilient network, get you to 99.95% availability or 4.38 hours of downtime per year. Virtualized high availability software solutions deliver four nines of availability — 99.99% — which reduces unplanned downtime to 53 minutes per annum.

Fault-tolerant solutions are often described as providing continuous availability because they are designed to prevent downtime from happening in the first place. Fault-tolerant software and hardware solutions provide at least five nines of availability — 99.999+% — for minimal unplanned downtime of between two and a half and five and a quarter minutes per year. While fault-tolerant hardware and software solutions both provide extremely high levels of availability, there is a trade-off: fault-tolerant servers achieve high availability with a minimal amount of system overhead to deliver a superior level of performance while fault-tolerant software can be run on industry-standard servers your organization may already have in place.



For an **Always-On** World

Question



In the event of a server failure, what is the process to restore applications to normal processing operation and how long does it take?

With most availability solutions, there will be some level of system interruption in the event of a server outage. Therefore, when evaluating solutions, it is important to understand what is involved in restoring applications to normal operations and how long the process takes.

If you rely on standalone servers, your recovery time could range from minutes to days given the high level of human interaction required to restore the applications and data from backup — provided you've been backing up your system on a regular basis. With high availability clusters, processing is interrupted during a server outage and recovery can take from minutes to hours depending on how long it takes to check file integrity, roll back databases, and replay transaction logs once availability is restored. If the cluster was sized correctly during the initial planning stages, users should not experience slower application performance while the faulty server is out of operation; they may, however, need to rerun some transactions using a journal file once normal processing resumes.

Fault-tolerant solutions proactively prevent downtime with fully replicated components that eliminate any single point of failure. Some platforms automatically manage their replicated components, executing all processing in lockstep. Because replicated components perform the same instructions at the same time, there is zero interruption in processing — even if a component fails. This means that, unlike a standalone server or high availability cluster, the fault-tolerant solution keeps on functioning while any issue is being resolved.

Question



How does your solution protect against loss of in-flight data?

Downtime Prevention Buyer's Guide

When a system outage occurs all data and transactions not yet written to disk are at risk of being lost or corrupted. In the case of some applications, this risk may be tolerable. But when you consider systems that automate functions like building automation and security, public safety, financial transactions or manufacturing processes, the loss of in-flight data can have serious consequences ranging from a scrapped batch or lost revenue to compliance issues or even loss of life.

Many availability solutions are not designed to ensure transaction and data integrity in the event of a system failure. Depending on how the hardware is configured, standalone servers and high availability clusters can typically preserve the integrity of database transactions, but any in-memory data not yet written to disk will be lost upon failure. Fault-tolerant solutions are built from the ground up to provide higher levels of data integrity. Fully replicated hardware components and mirrored memory ensure that all in-flight transactions are preserved — even when a hardware component fails.




For an **Always-On** World



Can your solution integrate seamlessly into existing computing environments with no application changes required?

Some availability solutions integrate more easily into existing computing environments than others. Certain solutions may require that you make changes to your existing applications — a process that is time-consuming and typically requires specialized IT expertise. For example, high availability clusters may need cluster-specific APIs to ensure proper failover. If ease of deployment and management are top priorities for your organization, you may want to consider a fault-tolerant solution that allows your existing applications to run without the risk and expense associated with modifications, special programming, and complex scripting.

Question



Does your solution require any specialized skills to install, configure, and/or maintain?

Downtime Prevention Buyer's Guide

In addition to a solution's recovery times and ease of integration, it is important to understand exactly what is involved in deploying and managing various availability alternatives. Some are simple to implement and administer while others demand specialized IT expertise and involve significant ongoing administrative effort. For example, deployment of high availability clusters requires careful planning to eliminate single points of failure and to properly size servers. Plus, whenever you make changes to hardware or software within the cluster, best practices suggest that you update and test failover scripts — a task that can be both time consuming and resource intensive. Some planned downtime is typically required to conduct the tests and ensure that the environment is working correctly.

Other solutions provide a more plug-and-play approach to availability. Today's fault-tolerant approaches prevent downtime without the need for failover scripting, repeated test procedures, or any extra effort required to make applications cluster-aware. With fault-tolerant solutions, your applications run seamlessly with no need for software modifications or special configuration changes. Fault-tolerant servers even provide a "single system view" that presents and manages replicated components as one system image, thereby simplifying installation, configuration, and management.

Before investing in a fault-tolerant solution to protect your critical applications against downtime, take serviceability into account, too. Ask about features like 24/7 system monitoring and automatic problem diagnosis, automated identification of failed components and replacement part ordering, customer-replaceable units with automatic system resynchronization features — all of which help ensure continuous operations and eliminate the need for specialized IT expertise.



For an **Always-On** World

Question



Is your solution future-ready and what is the lifetime value of the investment?

Downtime Prevention Buyer's Guide


When you invest in an availability solution, it makes good business sense to consider longevity and total cost of ownership. As more organizations rethink their server refresh schedules, they're looking for platforms that can truly go the distance to maximize return on investment. Therefore, when evaluating solutions, it makes sense to ask vendors about the average life span of their products.

Research has shown that standard servers tend to experience a marked increase in failure rates, downtime, and support costs between years four and five, prompting organizations to refresh on a four-year cycle. Fault-tolerant servers, however, offer significantly longer life spans — many averaging seven years — without notable performance degradation or higher maintenance costs.

Before making your purchase decision, you should also enquire about customer satisfaction ratings and retention rates to verify vendors' claims and make sure they deliver on their promises.



For an **Always-On** World



Conclusion

When it comes to protecting business-critical applications against unplanned downtime, you can't afford to leave anything to chance.

Make sure you ask these key questions as you evaluate vendors' availability solutions:

1. What level of uninterrupted application processing can your solution guarantee?
2. In the event of a server failure, what is the process to restore applications to normal processing operation and how long does it take?
3. How does your solution protect against loss of in-flight data?
4. Can your solution integrate seamlessly into existing computing environments with no application changes required?
5. Does your solution require any specialized skills to install, configure, and/or maintain?
6. Is your solution future-ready and what is the lifetime value of the investment?

Asking vendors the right questions up front will streamline the evaluation process and guide you in selecting a best-fit solution to keep your applications up and running in today's always-on world.

About Stratus Technologies

Stratus Technologies is the leading provider of infrastructure-based solutions that keep applications running continuously in today's always-on world. Stratus enables rapid deployment of always-on infrastructures, from enterprise servers to clouds, without any changes to applications. Stratus' flexible solutions – software, platform and services – prevent downtime before it occurs and ensure uninterrupted performance of essential business operations.





5 Mill and Main
Maynard, MA 01754

www.stratus.com