

READY YOUR INFRASTRUCTURE FOR THE ALWAYS-ON WORLD



Is Your Infrastructure Ready for the Always-on World?

Solutions To Prevent Downtime Rather Than Recover From Problems That Have Already Caused Damage



We live in an always-on world. Everything is connected — not just our computers, phones, and other portable devices, but nearly every aspect of daily life. From the gas pumps that fill up our cars (plus the cars themselves) to the card reader for

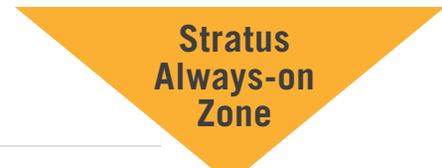
entering the office, the checkout at the grocery store, and the television streaming a movie at night. Services across banking, transportation, healthcare, government, public safety, utilities, manufacturing, telecommunications, and building systems are all expected to be accessible day and night, 365 days a year. That's why organizations in these industries must consider the availability of their IT infrastructures.

Much is at stake in terms of business profitability, protection of people and property, and reputation. How much downtime can an organization and its employees, constituents, and customers tolerate per year? Three days? Eight hours? Five minutes? Often, the answer comes down to the cost of downtime. According to research firm Aberdeen Group, one hour of downtime costs the average company \$163,674.14. It does not take long for this to mushroom into millions of lost dollars.

Despite these numbers, most organizations underestimate how well their systems are protected. Emphasis is typically placed on data backups and disaster recovery. While both these strategies are critical aspects of protection, they focus on damage control rather than preventing downtime from occurring in the first place. Moreover, they rarely account for everyday “disasters” such as server crashes and memory meltdowns. Such events affect many areas within an organization — not just “mission-critical” or “business-critical” systems, but also research and development, messaging, human resources, and websites. All these parts of the organization must be kept running without disruption to maintain productivity, employee satisfaction, competitiveness, and a quality customer experience.

Organizations are addressing the availability issue in a variety of ways, including the use of conventional servers, virtualization, and clustering. These approaches offer varying degrees of effectiveness in addressing downtime (Figure 1).

According to research firm Aberdeen Group, one hour of downtime costs the average company \$163,674.14. It does not take long for this to mushroom into millions of lost dollars.



Bottom Line on Nines: Stratus defines always on

	 Conventional, Unmanaged	 Typical Cloud Service Level	 Conventional Clusters, VMs	 Stratus Advantage
Always on Level	99%	99.9%	99.95%	99.999%
Downtime/Year	87.6 hours	8.76 hours	4.38 hours	5¼ minutes
Annual Cost	\$ 14,688,242	\$ 1,468,824	\$ 734,412	\$ 14,688

Assumption: one hour of downtime costs the average company \$167,674*. Calculation: 365 days x 24 (hours in a day) x 60 (minutes in an hour) x 60 (seconds in a minute) = 31,536,000 seconds in a year. At 99.999% uptime, there is 315.36 seconds of downtime a year or 5.25 minutes.

*Source: Aberdeen Group, Analyst Insight, June 2013

Figure 1: Always-on Effectiveness Matters

Conventional Servers

Some organizations simply rely on the availability levels of conventional x86 servers (Figure 2), which typically experience 87.6 hours of downtime per year.

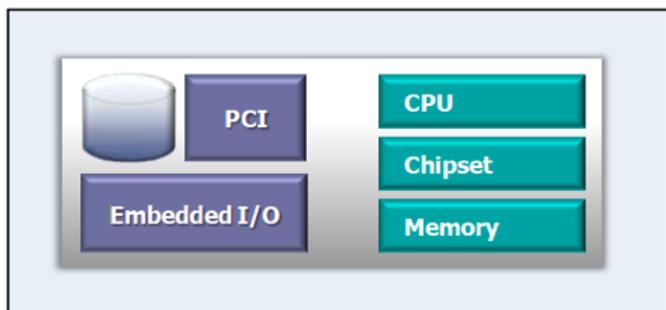


Figure 2: Conventional Standalone Server

Data is usually housed on storage systems configured with RAID (redundant arrays of independent disks). However, any server failure will stop all processing and users will lose access to their applications and information. As a contingency for more serious problems, organizations may deploy a cold standby server loaded with the same operating system and applications

as the primary server. But the standby is kept turned off unless needed, which is a costly waste of resources. Moreover, failing over to the standby server can take hours and requires skilled IT staff to be onsite.

Virtualization

Increasingly, organizations are embracing virtualization and cloud solutions, expecting the agility and data mobility in these environments to provide higher levels of availability. Virtualization adds a layer of abstraction above the hardware, allowing multiple virtual machines (VMs) to share the resources of a single physical server. If a virtual machine crashes, it can be independently rebooted without affecting the other VMs. VMs can also be moved easily from one physical device to another, which is convenient for hardware upgrades and maintenance. However, if the physical server fails, all the VMs running on it also go down. Virtualization is at the core of cloud computing, which suffers on average 8.76 hours of downtime per year. With outages of high-profile cloud providers recently in the news, clearly availability in the cloud remains a major issue.

Failover Clusters

Augmenting conventional servers or VMs with clustering improves availability by a few hours annually, but at the price of increased complexity and cost. With clustering, two or more physical servers are connected in a single network (Figure 3).

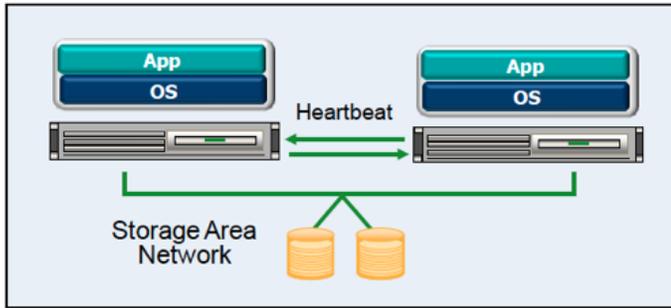


Figure 3: Failover Cluster

Servers in a cluster communicate with each other by continually checking for a heartbeat, which confirms that other servers in the cluster are up and running. If one server in the cluster fails, another will automatically take over with access to network resources. Clustering, however, requires highly sophisticated software, along with specialized expertise to write and test failover scripts. The result is not just high initial cost, but also significant ongoing expenses to update these items — requiring additional on-site expertise. Moreover, clusters often rely on shared storage, which becomes a single point of failure and potential source of downtime.

The Stratus Solution: Preventing Downtime Simply and Cost-Effectively

Stratus takes a completely different approach to availability by preventing downtime from occurring rather than recovering from an outage. Stratus always-on solutions deploy quickly and easily without the need for specialized expertise or changes to applications. Redundancy and fault tolerance are built in and ready to transform your current IT infrastructure into an always-on infrastructure out of the box. As a result, downtime is extremely rare — typically less than five minutes per year.

Stratus takes a completely different approach to availability by preventing downtime from occurring rather than recovering from an outage.

Software Solutions

Stratus® everRun® Enterprise downtime prevention software supports Windows® and Linux environments deployed on standard Intel®-based servers, including those from IBM, Dell and HP. This always-on software solution provides comprehensive downtime prevention, enabling true business continuity for everyday faults, as well as widespread outages or natural disasters (Figure 4).

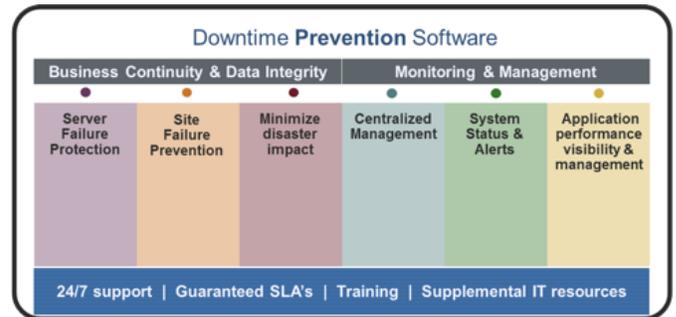


Figure 4: Stratus Downtime Prevention Software

To prevent server failure, everRun Enterprise incorporates the Stratus Availability Engine, which mirrors an application on two physical servers (Figure 5). With the Stratus Availability Engine, if one server fails, the application continues to run on the other server with no interruptions or lost data. It preserves all in-flight transactions, including data in memory and cache, so no restarts are necessary.

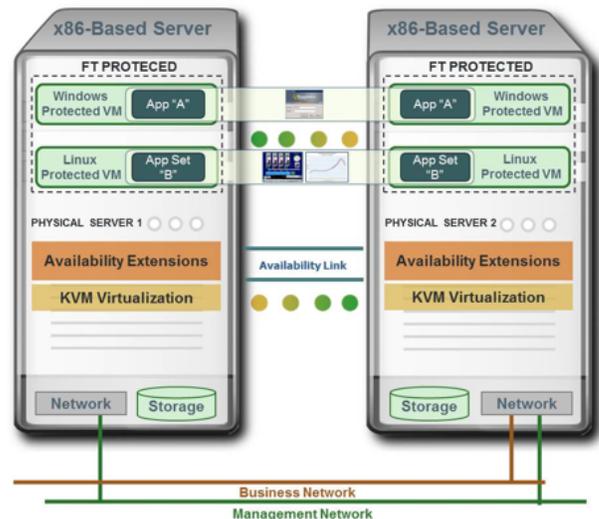


Figure 5 – Stratus Availability Engine Architecture

To protect against localized site loss due to human error, fire, or power outages, Stratus also offers add-on solutions to enable synchronous replication of the protected infrastructure across sites in a metropolitan area. For greater geographic separation to protect against major disasters, Stratus also provides integrated asynchronous replication between sites over a wide area network connection.

In addition, Stratus provides a centralized administrative console to manage the entire always-on infrastructure. Through a single pane of glass, administrators can build and deploy virtual machines, designate servers for fault-tolerant pairs, select application availability mode, and configure servers across multiple deployments. Multiple fault-tolerant pairs can also be managed through this single console. It is a powerful way to ensure always-on availability while minimizing administrative time and effort. Stratus also provides an added layer of protection through its support center, which monitors systems 24/7/365 and automatically notifies administrators if an issue is detected.

Platform Solutions

Stratus platform solutions provide worry-free operation with an integrated package of hardware, software, and service. These platforms are built on Intel® Xeon® technology and run Windows, Linux®, and VMware® operating environments.

Stratus always-on platforms have two of everything — including CPUs, memory, I/O, disks, and power supplies — all working in synch as a single system (Figure 6). In the event of a component malfunction, the twin component takes over instantly without causing performance degradation, lost data, or downtime. There is no need to reboot or fail over.

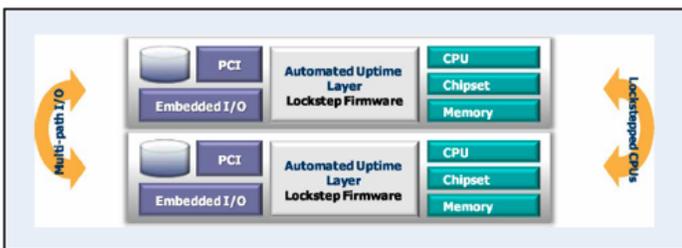


Figure 6 – Stratus Availability Engine Architecture

Since the network, operating system, and applications see a single system image, there is no need to modify applications or license multiple copies of software. This makes Stratus platforms easy to set up and maintain — especially compared to clusters — and often don't even require onsite IT support.

Software built into the Stratus platform detects and prevents many causes of downtime and data loss. It also simplifies monitoring and management of the server and enables remote service and support. A unique capability provided by Stratus is the ability to diagnose potential problems down to the operating system level.

This comprehensive monitoring and analysis is always active, not just after an event has already caused trouble. In fact, it handles most errors transparently, enabling Stratus platforms to operate through transient errors, protect against device driver failures, and prevent hardware and software failures from causing downtime.

If the onboard software can't fully resolve a problem, the Stratus platform automatically “phones home” to alert the Stratus Support Center. Stratus availability experts remotely monitor Stratus platforms 24/7/365 to proactively resolve problems before they impact business operations. These highly trained technicians can resolve most issues within 30 minutes. Less than 1% of system issues require an onsite support call.

Should a platform component fail, a replacement unit is automatically dispatched for next day arrival. Thanks to Stratus's unique hot-swappable design, the new unit can easily be replaced without tools or computer skills. It's as simple as sliding the failed unit out and the new unit in. There is no disruption to active processing on the platform during this replacement and no loss of data.

Stratus infrastructure solutions are designed from the ground up for this always-on world.

Conclusion

Today, with nearly every aspect of daily life connected and dependent on IT systems, continuous availability and downtime prevention are necessities. Stratus infrastructure solutions are designed from the ground up for this always-on world.

Unlike other approaches designed to recover from a fault or outage, such as cold standby servers and failover clusters, Stratus solutions avoid downtime altogether. While most organizations prepare for large natural disasters or disruptive geopolitical events, they overlook everyday server faults and drive crashes. Stratus simply prevents them and avoids problems that lead to lost productivity, lost data, and lost revenue.

Stratus solutions are built on a 30-year history of delivering high availability and fault tolerance for many of the world's most demanding applications. Stratus solutions are:

- **Flexible** – supporting a variety of operating systems and designed to run in physical, virtualized or cloud environment
- **Easy** – deploying in minutes and allowing applications to run without changes or specialized expertise
- **Trusted** – running the most critical applications by Global Fortune 500 companies and small-to-medium sized businesses alike

In manufacturing environments, Stratus solutions are the key to maintaining productivity and reducing waste. For retailers, they ensure that transaction processing systems are always on to maintain sales targets. In building security, Stratus ensures continuous protection of premises and individuals from internal or external threats. First responders are assured of access to life-saving systems. Financial services firms avoid costly downtime when managing thousands of transactions per second. And healthcare organizations enable 24/7/365 access to vital patient records while ensuring compliance.

For any environment in today's always-on world, Stratus has the right solution to prevent downtime for everyday system issues and ensure business continuity in the face of major disasters.

About Stratus Technologies

In today's always-on world, applications run under increasingly demanding circumstances. With these escalating demands comes greater pressure to prevent even the smallest amount of application downtime. Companies are responding to this need for always-on solutions by searching for technologies that either conform to or enhance their current IT infrastructures.

Stratus Technologies' solutions enable rapid deployment of always-on infrastructures, from enterprise servers to clouds, without any changes to your applications. Stratus products (software and servers) combined with Stratus people, enable customers to prevent downtime before it occurs, ensuring uninterrupted 24/7/365 performance of essential business operations.

To learn more, visit stratus.com

For any environment in today's always-on world, Stratus has the right solution to prevent downtime for everyday system issues and ensure business continuity in the face of major disasters.