WHY THE THREAT OF DOWNTIME SHOULD BE KEEPING YOU UP AT NIGHT

Stratus® Technologies

For an **Always-On** World

www.stratus.com

Your Plan B Just Isn't Good Enough. Learn Why—and What to Do About It.

Server downtime is an issue that many organizations struggle with. But when it comes to building security and automation systems, you simply can't afford to risk any downtime. Naturally, a security system can only provide protection if it's up and running 24/7/365. If video monitoring systems, access control, or other building security systems go down, it can mean costly, dangerous, and potentially life-threatening consequences.

Perhaps you think if your server goes down, you can easily manage security concerns by redeploying security staff. But how will you know there is an emergency without a working alert system or video cameras? How will you open automated doors in a fire? How will you explain if people lose their lives because you didn't know a crisis was occurring?

This is why relying on a plan B in an emergency just isn't good enough. It's crucial to be sure your servers are up and running all the time with a reliable availability solution. In this paper you'll learn all about availability, the exact risks of server downtime, the less obvious costs you can incur, and the different types of availability solutions you can use to protect your organization.

What Exactly Is "Availability"?

The term "availability" refers to your server reliability. Basically, it's defined as the percentage of time that your applications are operational and accessible to users. When talking about servers, there are three different types of availability:

1. Backups and restores: This is achieved by having basic backup, data-replication, and failover procedures in place, which are included with conventional servers. This method delivers approximately 99 percent availability, which equates to an average of 87.5 hours of downtime per year, or more than 90 minutes of uninvited downtime per standard work week.
2. High availability (HA): With an HA solution, applications are accessible a very high percentage of the time. HA solutions include cluster systems (which focus on quick recovery after a failure) and software solutions (which aim to prevent downtime from happening). High availability translates into 99.95 percent – 99.99 percent uptime, or 4.5 hours – 52 minutes of downtime per year.
3. Continuous availability (CA): Also called an always-on solution, a CA solution's goal is to eliminate downtime altogether through software or a specialized server system. CA solutions are the ultimate in availability (99.999 percent) and result in just one to five-and-a-half minutes of downtime in a year.

This goes to show that a solution that delivers 99 percent availability may sound good at first. But when you need security systems to be operating continuously, 99 percent is nowhere near good enough.

When you need security systems to be operating continuously, 99 percent is nowhere near good enough.

Four Reasons You Simply Can't Risk Downtime

It may seem that it's not crucial to have building security systems servers up and running all the time. What's the likelihood of an security incident occurring when your server is down? Probably not too high. But the question is—what if? It's your job to protect against that what-if scenario. And that's impossible if your server is down.

The fact is, your security systems are fulfilling an important need. If there is a security incident and your server is down, you're not going to be able to respond in the way you must. The risks involved are probably much greater than you realize:

1. **You will not be able to immediately react when there is an incident.** Most building security solutions have distributed systems. Information that is collected at local devices, such as card readers or access control units, is stored locally and sent to the server at specific intervals. A server outage doesn't affect the data collected.

   What a server outage does affect is your ability to redirect the resources of the distributed system. For example, a manager may need to change security access, reposition cameras, or adjust fire control when an emergency situation arises. If the servers are down, the manager is unable to make these changes at the moment he needs to do so. If there is a fire and he needs to override the access control, he can't do it. The disaster scenarios are unfortunately endless.

2. **If systems go down, your liability goes up.** If you don't have control over all access points in your building, what could happen? If you are a hospital, newborn tracking devices will stop working and the children's safety will be compromised. If you are a power plant, unauthorized individuals can get into your facility. Whatever type of organization you are, having a nonfunctional security system is a major and potentially life-threatening liability.

3. **The potential financial costs are considerable.** In many cases security systems are in place not only to keep people safe, but to comply with government regulations. Utility companies have to meet demanding requirements for safety and security. Airports need to meet transportation regulations for security. Healthcare and government facilities have strict regulations, too. For these and other organizations, downtime can mean high penalties or even losing licensure.

   Also consider that if you're protecting research facilities, they need to be guaranteed highly stable environments. Even the slightest variability in conditions such as temperature, humidity, air flow, and air quality can skew scientific findings and render ongoing research work useless. No scientist wants years' worth of work—and grant dollars—to be lost when a environmental-monitoring system goes offline.

4. **Disasters happen, and you have to be ready.** Whether it is a natural disaster or a man-made emergency, we've all seen the news headlines about catastrophes that claim lives and damage or destroy buildings. The possibility of these events means that disaster recovery is a crucial element of building security.

   Your security systems should have back-up solutions in place in geographically-separated sites so that if a catastrophic event occurs, your assets are protected. Think of the prevention of downtime as one element of a comprehensive disaster recovery plan. In times where you can't prevent a catastrophe from taking out your building (and therefore your servers) you can mitigate its impact on your IT systems with off-site backup in place.

Your security systems should have back-up solutions in place in geographically-separated sites so that if a catastrophic event occurs, your assets are protected.

Four Availability Options for Protecting Your Building Security Systems

Now that you understand exactly why you need reliable availability for your building security systems, let's delve into four types of solutions. All four deliver high availability or continuous availability. Which you choose depends upon your desired availability, IT resources, server preferences, and more.

High-availability clusters

This is a solution that aims to recover from downtime as quickly as possible, as opposed to preventing it from occurring. HA clusters are a custom-built system composed of two or more servers that are running with the same configuration. They are connected with cluster software to keep the application data updated on both/all servers.

Servers in a high-availability cluster communicate with each other by continually checking for a heartbeat that confirms other servers in the cluster are up and running. If a server fails, another server in the cluster (designated as the failover server) will automatically take over, ideally with minimal disruption to users.

While high-availability clusters improve availability, their effectiveness is highly dependent on the skills of specialized IT personnel. Clusters can be complex and time-consuming to deploy, and they require programming, testing, and continuous administrative oversight. And because there are two or more servers, you need to license software on each of them. As a result, the total cost of ownership (TCO) is often high.

It is also important to note that downtime is not eliminated with high-availability clusters. In the event of a server failure, all users who are currently connected to that server lose their connections. Therefore, in-flight data is lost.

High-availability software

This type of software is designed to prevent downtime, data loss, and business interruption from occurring. HA software is equipped with predictive features that automatically identify, report, and handle faults before they become problems and cause downtime.

Two important features of high-availability software are that it works with standard x86 servers and doesn't require the skills of highly advanced IT staff to install or maintain it. HA software is designed to configure and manage its own operation, making the setup of application environments easier and more economical than clusters.

There is a key difference between high-availability clusters and high-availability software: The software continuously monitors for issues to prevent downtime from occurring, whereas cluster solutions are designed to recover after a failure has already occurred. The most effective HA solutions provide more than 99.99% availability, which translates to less than one hour of unscheduled downtime per year.

Continuous-availability software

This type of availability solution not only prevents downtime from occurring, but also provides the cost-saving benefits of an always on solution based on standard x86 servers. Each application lives on two virtual machines. If one machine fails, the applications continue to run on the other machine with no interruptions or data loss. If a component fails, it's replaced by the healthy component from the second system.

Continuously availability software can also offer disaster recovery and split-site capabilities. It prevents data loss, is simple to configure and manage, requires no special IT skills, and delivers upwards of 99.999 percent availability—all on standard servers.

Continuous-availability servers

A continuous availability server system is truly turnkey in that its hardware, software, and services are all integrated for easy management. This type of solution relies on specialized servers that are purposely built to prevent failures from ever occurring.

These servers are managed just like standard servers, so they don't require specialized IT personnel to manage them. The sophisticated technology simply runs in the background and isn't apparent to system administrators.

These solutions include redundancy of components and error-detection software. Automatic fault detection and correction is engineered into the design so that most errors are resolved without you even knowing they existed. Plus, continuous availability servers can run in a virtualized environment. You can expect greater than 99.999 percent availability if you choose this option.

In the End, Always-On Is the Only Option

There are many applications which organizations rely upon that can tolerate hours of downtime. But building security and automation systems simply can't. When people's safety, proprietary technology, or entire businesses are at stake, reliable availability with minimal or no downtime is truly needed. A system that prevents downtime from happening in the first place is the ideal solution, and with the range of options available today, should be a natural choice for those who manage building security systems.

When people's safety, proprietary technology, or entire businesses are at stake, reliable availability with minimal or no downtime is truly needed.

About Stratus

Stratus Technologies is the leading provider of infrastructure-based solutions that keep your applications running continuously in today's always-on world.

Stratus always-on solutions can be rapidly deployed without changes to your applications. Our platform solutions provide end-to-end operational support with integrated hardware, software and services. Our software solutions are designed to provide always-on capabilities to applications running in your chosen environment – physical, virtualized or cloud. Our approach and our people enable us to identify problems that others miss and prevent application downtime before it occurs. Multiple layers of proactive diagnostic, monitoring and self-correcting services are backed by a global team of engineers who provide immediate support no matter where in the world your system is located.

If **always-on** is an application requirement, Stratus Technologies has a solution that fits.

Stratus Technologies

www.stratus.com

Specifications and descriptions are summary in nature and subject to change without notice.

Stratus and the Stratus Technologies logo are trademarks or registered trademarks of Stratus Technologies Bermuda Ltd. All other marks are the property of their respective owners. ©2014 Stratus Technologies Bermuda Ltd. All rights reserved. 140407