WHITE PAPER

Empowered by Innovation NEC

Building More Security into Building Security: How Fault Tolerant Servers Reduce Risk and Lower Costs for Access Control

In this white paper, we will examine the key concepts of and concerns associated with access control systems. Most importantly, we will explain the concepts inherent in high availability access control systems, and show how fault tolerant (FT) servers provide an innovative method of keeping access control systems running reliably, without interruption, and at a lower cost than other options.

Table of Contents

Executive Summary

In another time, a building could be secured with a few locks and a dependable security guard. But today's threats are sophisticated, and the security to fight them must be even more so. High tech access control systems—featuring card readers, intelligent video surveillance, electronic locking devices, and the computers that control them—are becoming standard for building security.

In the wake of September 11, businesses and landlords have been compelled to invest in this new generation of access control system, and learn the technology that comes with it. Facilities managers and security professionals don't have to be computer scientists. But they do have to have a basic understanding and appreciation for the technology that drives 24/7 building security.

Building Security: The Emerging Imperative

Security professionals dread this scenario, but it happens every day: an access control system malfunctions for a few seconds. Impatient employees, anxious to get to their offices and meetings, prop open an access-controlled door.

The actual system downtime is only a few seconds. But the server takes minutes to reboot, and the door stays open even longer before a security guard arrives on the scene to correct the problem.

Dozens, perhaps a hundred people, enter the propped-open door without verified credentials. Is just one of them an unwanted intruder?

Noticing that access control is not reliable, employees lose confidence in the system. Worse, they take security less seriously. They hold doors for strangers without cards, or override a lock to go outside for a cigarette. The result: a thief, spy, terrorist, or other undesirable person only needs to watch and wait for the next opportunity to enter.

Before September 11, 2001 strict access control was typically found at military, government or R&D facilities. Most other businesses got away with the ubiquitous visitor sign-in sheet and the accommodating and helpful security guard.

Since that day, it has become expected and even routine to see access control systems at transportation hubs, power stations, hospitals and laboratories, hotels, campuses, manufacturing plants, and corporate offices. Businesses that have adopted access control systems include financial institutions, legal and professional services firms, high-tech companies, news and entertainment companies, political organizations, and any other groups that could be a target of theft, terrorism or espionage at any level.

As organizations become more aware of their need for secure facilities, IT departments and building managers are becoming newly familiar with access control technologies, and the need for high availability computers—with virtually zero downtime and flawless operation—at the center of this mission critical function.

The Nerve Center of Access Control

Computer hardware is truly the nerve center, or backbone, of the new-generation access control system. Each component works as part of a unified sub-network to verify credential data, record events, and trigger notifications. Because the safety and security of people and property are on the line, the continuous and consistent operation of this network is crucial.

Components of Access Control: 24/7 Operaion

The components of an access control system depend on constant communication with the central computing system in order to function properly. Uninterrupted and reliable operation is of particular importance in the following subsystems:

- **Alarm Monitoring System**
  As the visual window into the building's access control status, this system is the core access control application. Security staff depends on alarm monitoring for real-time notifications of all events that impact building security.

- **Live Video Monitoring**
  An effective live monitoring systems uses breakthrough alarms to push live video to monitoring stations as events occur. These systems must be always online to guarantee a real-time response.

- **Video Verification**
  With the ability to compare live video of a person to an image stored in a computer database, security staff can make informed and timely decisions before granting access.

- **Intrusion Detection and Fire Alarms**
  For an alarm, any downtime is a severe security breach. Alarms must operate without interruption.

The Central Server:

The Backbone of Effective Access Control

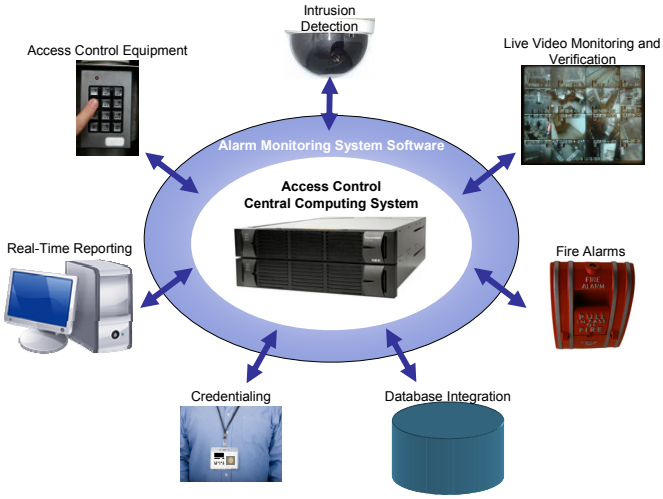

The effectiveness of an access control system is dependent on the uninterrupted service of the central computing system. Without a reliable backbone, security is compromised.

- **Access Control Equipment**
  Door locks, motion detectors, request-to-exit control switches, glass break sensors, and other access equipment must be constantly and continually connected to the central server.

- **Credentialing**
  Whether to add new employees, change access privileges, or remove personnel from the system, most credential changes are needed immediately, without the risk of delay or interruption. System downtime can result in terminated employees entering a building, bottlenecks in traffic for visitors, or a number of other security headaches.

- **Reporting**
  Security staff and managers rely on a number of real-time reports, such as entry/exit logs, drill muster reports, and other event status information. Without reliable real-time data, security teams lack the information they need to account for personnel and property when critical decisions have to be made.

Access Control Requires Real-Time Integration

Access control systems work best when they are integrated with other business processes. Consider what happens when an employee is terminated. Many things must happen quickly: computer logins must be deleted, credit cards cancelled, and network and facilities access must be denied, just to name a few. Without seamless communication between the HR system and the access control system, the company's assets, and perhaps even the safety of its personnel, are at risk. When the systems that manage these functions are integrated, all of these operations, and more, can be executed with one operation; the Human Resources (HR) department needs only to submit one transaction, and all other systems record the change.

In addition to HR, other applications that are most often integrated with access control include time and attendance, meal plan management, vending, and visitor management. Because of the time-sensitive nature of building security and access control, these integrations usually require a single master database to serve these multiple applications (as opposed to multiple databases that are replicated and/or synchronized).

Reliable and easy to manage systems are required to execute these instructions immediately and flawlessly, with no delay caused by computer downtime.

High Availability: Why Does It Matter?

The high performance, always-on demands of an access control system necessitate high availability computer hardware for the backbone of the system. High availability describes a system designed and implemented to ensure a certain absolute degree of operational continuity. Availability is generally measured as a percentage of uptime. Since uptimes generally vary from 99% to 99.999%, availability is commonly expressed in terms of "nines." An average availability of "five-nines," or 99.999%, represents the optimal performance of today's high availability computing technologies.

"Five Nines" is State-of-the-Art for High Availability

Nines	Availability	Downtime
1	90%	36.5 days/year
2	99%	3.65 days/year
3	99.9%	8.76 hours/year
4	99.99%	52 minutes/year
5	99.999%	5.25 minutes/year

For access control, the difference between "four nines" and "five nines" is significant. For example, a system with only 99.99% uptime could be down for as much as one minute per week. Under these conditions, only a few outages during peak traffic periods could severely degrade the effectiveness of building security.

What the “nines” do not represent, however, is the length of time of each outage, or a system’s ability to recover from an outage. The same “four-nines” system could be down for about one minute every week, or for one hour during one event in the course of a year. And, of course, these are averages; a system rated at “five nines” could actually perform better or worse than 99.999%.

The Science of High Availability

A more scientific approach to the varying types and levels of availability has been developed by the analyst firm IDC. The following table is adapted from IDC’s “Availability Spectrum,” which it uses to aid hardware buyers in selecting an appropriate level of availability. [1]

Clearly, AL4 is the superior level of availability for access control applications, since it offers an environment with no interruption of service, even through a hardware failure. However, IT departments and facilities managers are unfortunately enticed, due to budget constraints, to purchase AL3 or even AL2 systems for access control.

High Availability: Not Just “Nice-to-Have” for Building Security

High availability is imperative to the basic effectiveness of an access control system. In terms of average downtime, the difference between a few minutes per week and a few minutes per year is critical: it can make the difference between unauthorized entry and a secured environment—perhaps making the difference between life and death.

IDC’s Availability Spectrum

	Impact of Component Failure on Priority User	System Protection Features
Availability level 4 (AL4)	Transparent to user; no interruption of work; no transactions lost; no degradation in performance	100% component and functional redundancy
Availability level 3 (AL3)	Stays online; current transaction may need restarting; may experience performance degradation	Automatic fail over transfers user session and workload to backup components; multiple systems connections to disks
Availability level 2 (AL2)	User interrupted, but can quickly re-log on; may need to rerun some transactions from journal file; may experience performance degradation	User work transferred to backup components; multiple system access paths to disks
Availability level 1 (AL1)	Work stops; uncontrolled shutdown; data integrity ensured	Disk mirroring or RAID, and a log-based journal file system for identification and recovery of incomplete in-flight transactions

[1] IDC, Worldwide and U.S. High-Availability Server 2006–2010 Forecast and Analysis, Doc #204815, Dec 2006

Because the uninterrupted service of the central computing system is crucial to effective access control, IT facilities management departments attempt to “cobble together” systems using off-the-shelf computing hardware to achieve high availability through various forms of clustering technology. Unfortunately, this approach:

- Requires considerably more effort to install and stabilize
- Adds unnecessary complexity to a security a environment where stability is critical
- Requires advanced training of security and IT staff to understand operation
- Introduces third-party service organizations into sensitive security access control areas for server maintenance

Because of the above mentioned limitations, AL4 fault tolerant systems that deliver 99.999% uptime have rapidly become the high availability standard for security access control systems. These systems ensure successful computing transactions with card readers, monitoring devices, alarms, locks, and many other devices of the access control system.

How Do Fault Tolerant Servers Deliver High Availability?

A fault tolerant server is built from the ground up to perform at AL4 with “five nines” availability. It is a fully internally-redundant system that goes far beyond the traditional server cluster in terms of reliability and cost-effectiveness. It’s true that fault tolerant servers are more expensive than off-the-shelf servers. But fortunately, advances in computing technology are driving down the prices of fault tolerant servers, even though the demand for high availability computing is increasing. IT and facilities managers who plan access control implementations, no longer have to make difficult choices when balancing security and budget considerations.

A Quick Guide to High Availability Server Technologies

RAID: (Redundant Array of Independent Disks): a data storage scheme using multiple hard drives to share or replicate data among the drives.

Cluster: a group of loosely coupled computers that work together closely so that, in many respects, they can be viewed as though they are a single computer.

Failover: the capability to switch-over automatically to a redundant or standby system upon the failure or abnormal termination of the previously active system. Failover happens without human intervention.

Lockstep: describes a **fault tolerant** machine that uses replicated elements operating in parallel. At any time, all the replications of each element should be in the same state. The same inputs are provided to each replication, and the same outputs are expected.

The overarching benefit of an FT server, aside from its superior AL4 high availability, is that the features that ensure high performance are built into the hardware, so they do not have to be installed, configured, and maintained using software.

Here are some major features of FT servers—particularly NEC's Express5800/ft series of servers—that set this technology apart from other high availability technologies.

Redundant Components

All of the memory, processors, and other components of the FT server are redundant to each other, and physically configured to operate in lockstep. Therefore, one FT server is the equivalent of two conventional servers performing the exact same processes at the same time. In the event that one component fails, its counterpart continues functioning with no interruption in the operation of the system. Since failover is virtually instantaneous, there is no single point of failure, and downtime is nearly eliminated.

■■■ "The emergence of true fault tolerant solutions for standard Windows- and Linux-based server environments from vendors including Stratus, NEC, and others, will provide customers with alternatives to the complexities of clustered-server solutions."

- Matt Eastwood, Vice President, Enterprise Server Research, IDC

Conventional vs. Fault Tolerant Servers

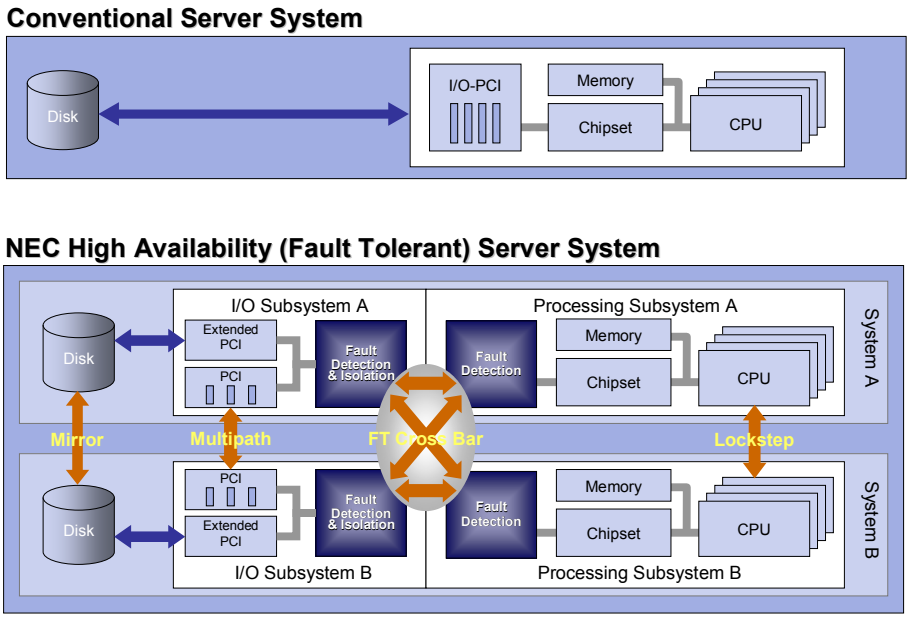

Th
sa
au
co
ma
particularly those in companies who are new to the need for high availability computing, the
effe
tecl

Fault tolerant servers are superior to conventional servers, because there is no single point of failure, no switchover, and a single logical server.

Hot-Swappable Hardware

When a hardware component in an FT server does need replacing, repairs can be made with no downtime, while the overall system is still "hot" and online. Even routine hardware maintenance does not require any planned downtime, and can be performed without compromising the access control system.

Before hot-swappable components were available in FT servers, clusters had the advantage in this regard, since it is relatively easy to take one machine in a cluster offline for maintenance. Now with hot-swappable FT components, companies do not have to compromise a level of high availability to maintain these servers.

ActiveUpgrade™ for Software Maintenance

Active Upgrade, a feature of NEC's FT servers, allows for software updates with minimal interruption of service. One module can be administered offline while the second module continues to handle the operational load. When updates are completed, the two modules synchronize data and return to full redundant operation. A rollback feature is integrated into Active Upgrade in order to return the server to its previous state in the event there is a problem with the software update.

"Although clustering has been the predominant vehicle for improving server availability characteristics, the complexity that is often associated with these types of solutions in certain market segments has hampered more widespread acceptance…. IDC believes that the complexity that is often associated with some of these implementations must be further hidden from the end customer to preserve ease of use for IT shops that do not specialize in clustering technology and skill-sets (e.g. scripting to create cluster-aware applications)."

- Matt Eastwood, Vice President, Enterprise Server Research, IDC

What Makes an FT Server the Best Value for Access Control?

Application is lower, by every metric, than that of other high availability technologies. The combination of a reduced initial hardware expense, coupled with lower operating costs, makes FT servers a superior value over any other high availability technology. And in the high-stakes implementation of an access control system, there is simply no reason to compromise security.

- **One NEC Express5800 FT server is less expensive than two conventional servers.** It is no longer viable to try saving money by purchasing two less expensive conventional servers and clustering them together. Like so many other computer hardware products, FT servers are becoming more affordable.

- **FT servers require less configuration and maintenance.** IT departments already struggle with high maintenance expenses—in fact, analysts who survey IT groups find that many spend 50%, even 70% of their budgets on maintaining the systems they have, as opposed to building and improving their operations.

FT servers are less expensive to operate than clusters or RAID arrays, because the redundancies and configurations between the components are built into the hardware, instead of configured via software. Fewer specialized IT skills are required to implement FT servers, so it is no longer necessary to hire or contract workers with special skills in order to implement a high availability system.

- **One FT server requires only one software license.** Unlike a cluster, which requires multiple licenses of the operating system, databases, server applications and other software (one for each server in the cluster), an FT server requires only one license of each. Software costs can be dramatically reduced, along with the effort required to install and maintain multiple instances of each software product.

- **Less downtime means less maintenance.** Because downtime is dramatically reduced by virtue of the high-availability hardware, maintenance time is also reduced. IT staff can concentrate on other critical tasks.

Total Cost of Ownership of High Availability Hardware Alternatives

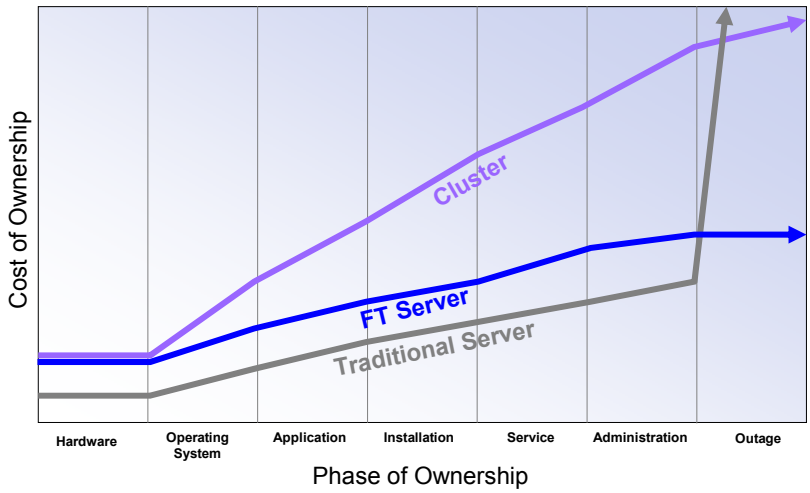

Fault tolerant servers provide lower total cost of ownership over the long haul—especially in the event of an outage.

Conclusion

A fault tolerant server is clearly the must-have hardware for an effective access control system. Organizations can ill-afford to gamble with the security of their property and the safety of their people—and now they don't have to.

The decreasing up-front cost of FT servers, combined with lower total cost of ownership throughout the life of the computer, makes the purchase of an FT server easy to justify for any organization that takes seriously the responsibility of facilities access control.

NEC Fault Tolerant Servers

NEC offers a full line of Express5800 fault tolerant servers that offer up to 99.999% uptime. We are a leading supplier of server technology to the access control industry. If you would like more information, call: 866-632-3226 or visit www.necam.com/ft.

NEC Corporation of America
Departmental Servers Division

2880 Scott Boulevard
Santa Clara, CA 95050

www.necam.com/DynamicIT
DynamicIT@necam.com
(866) 632-3226

NEC Corporation of America