


RELIABILITY, PERIOD.

**CALL CENTERS SHOULD UNDERSTAND THE
CRITICALITY OF THEIR INFORMATION RESOURCES.**

By John Leonard, Stratus Technologies



The time is 8 a.m. on a weekday. The place is a quiet suburban town; the situation, an anonymous phone call to the high school warns that an armed gunman is headed for the building.

The response: the district's schools are locked down. The police chief sounds an alert summoning police, fire and medical personnel from across the region. Parents, the media and the public are kept informed, but out of harm's way at a community hall nearby. In a smoothly coordinated response, every resource needed to manage the developing emergency is in place within sixty minutes. The crisis is resolved without incident before the noon bell rings.

But what if the departments and agencies involved didn't have instant access to the information they needed to be effective? What if their CAD (computer-aided dispatch) systems had been down? Quite a different story would have unfolded and the outcome of the situation could have been very, very different.

The mere possibility of the second scenario is what keeps the systems administrator—and everyone else—in the PSAP or communications center awake at night. Nobody wants to explain to the chief why the CAD or RMS (records management system) server crashed during an emergency. In fact, many departments and agencies use some combination of computer server redundancy and data replication to protect against downtime. These days, when our days are color-coded to the latest Homeland Security Advisory, public safety professionals have become too keenly aware of how critical their information resources really are—it's all about *reliability, period*.

Police, fire and other first responders depend on reliable information and communications, which are driven by computer servers in the backroom of the PSAP or communications center. These servers deliver emergency dispatch information consolidated from phone calls, radio, satellite and other sources and make it available on the dispatcher's desktop, mobile data communications devices in vehicles on the street and personal digital assistants (PDAs) in the hands of officers. Most people take for granted that such computer servers operate dependably. Few may be aware that sustaining reliability in servers and applications often requires constant vigilance on the part of the systems administrator, who may be too accustomed to being called in to solve problems that have a way of escalating during off-duty hours.

Computer Server Availability Categories

Could most municipalities' departments be better prepared to safeguard the reliability of CAD, RMS, and other essential information technology (IT) systems? Understanding the different approaches to computer server availability and evaluating new options is a good place to find out. The technology research firm IDC (www.idcresearch.com) defines four categories of server availability, from lowest to highest:

Availability level 1 (AL1)—A conventional computer server replicates data using disk mirroring or RAID (redundant arrays of independent disks); a log-based or journal file system is used to identify and recover incomplete requests or activities in progress. In the event of a server failure, work stops and users lose access to their applications and information.

Availability level 2 (AL2)—Users' work is transferred to backup components; multiple servers have access to disks where data resides. This means users are interrupted during a server crash, but can quickly log back on. Users may notice slower performance while the faulty server is out of operation, and may have to rerun some transactions using a journal file.

Availability level 3 (AL3)—The system handles automatic failover that transfers the user's session and workload to backup components; multiple system connections to disks keep data available. The advantage is that the user stays online during a system crash. The disadvantages are that the current transaction

may have to be restarted and users may notice slower performance.

Availability level 4 (AL4)—This approach is based on 100 percent component and functional redundancy. If one of the system components fails, the server continues to process normally. Comprehensive redundancy ensures that users don't experience any interruption, no transactions or data are lost and there is no degradation in performance. Fault-tolerant servers fit this category best.

How to Achieve Server Reliability

The convention that IT folks use is nines to quantify uptime reliability delivered by different approaches to achieving server availability. The gold standard long has been five nines, or 99.999 percent availability, which works out to about five minutes of downtime per year. One nine more or less may seem trivial on first glance, but say the CAD center experiences 95 percent availability per year; that translates to twenty-eight days of downtime. Each additional nine indicates an exponential improvement in server availability (see **Chart 1**). With differences like these, it stands to reason that most communications centers would opt to exclusively purchase servers that deliver a long string of nines.

Reliability of key servers and applications in the PSAP and communications center is imperative and becoming even more so.

Chart 1 The Nines

Availability measure	Approximate down time each year
99 percent	87.6 hours—more than two average workweeks
99.9 percent	8.8 hours
99.99 percent	.9 hours
99.999 percent	0.09 hours—about five minutes
99.9999 percent	59 seconds

Chart courtesy of Stratus Technologies.

Fault-Tolerant Servers

Simpler operation is another prime advantage of a fault-tolerant solution, and it's an attribute that pays back for as long as the community owns the servers. A fault-tolerant server design

provides availability protection by simply running off-the-shelf applications; no software changes or recoding are necessary.

There also is the matter of day-to-day administration. Today's fault-tolerant servers can be managed like ordinary Windows servers. When servers are easier to administer and there are fewer of them, they demand less attention and time to manage. As server configuration and system management become more complex, the potential for human error also increases. Such mistakes rob a server of its uptime availability, and shorten the string of nines delivered by the solution (99.5 to 99.95 percent in the case of a high-availability cluster).

Self-monitoring, online problem resolution and collaborative support are among the advanced maintenance capabilities that fault-tolerant servers bring to the Windows-based CAD environment. These servers can be designed to monitor their own

By understanding the various approaches to computer server availability, the uptime expected and the options in current server technology, municipalities have more tools ready to protect and serve.

THE NINES

Achieving true 24/7 uptime (100 percent of the time) is virtually impossible, if for no other reason than that at certain times, an upgrade such as a service pack must be applied to a production server. This is where the concept of nines comes into play. The percentage of uptime all

companies should strive for is some variation of 99.x percent, where x is a specified number of nines. Five nines is 99.999 percent uptime for your systems, which, to many, is considered to be the ultimate in availability. This may be difficult to achieve, since it means only about five

total minutes of downtime in a calendar year. Three nines is a more practical number to shoot for. Three nines is just short of nine hours of downtime per year—a very respectable number.
—Microsoft TechNet
(www.microsoft.com/technet/)



GET MORE INVOLVED IN THE MAGAZINE!

E-mail
christina@ctipublishing.com
for more details.

EDITORIAL ADVISORY BOARD

CASE STUDY

operation, isolate hardware and software errors, run self-diagnostics, phone home to notify the system vendor of a problem and allow authorized service engineers to investigate and resolve the problem without stepping foot in the communications center. The server's replicated components keep the system and application running without any interruption. Not only does this ensure reliability, but it also gives the over-worked systems admin the time to uncover the root cause of a recurring issue and correct it once and for all.

Upgrading Infrastructure

With the federal government strengthening homeland security, PSAP and communications center managers may find themselves upgrading infrastructure sooner and more often. Federal funding is anticipated for first-responder planning, training, exercises and equipment. Technology upgrades will be needed if state and local agencies are to participate fully in linking systems and sharing information to fight terrorism. A useful central resource about funding opportunities is The Catalog of Federal Domestic Assistance (CFDA), where interested parties can find a database of federal grant programs for state and local governments (www.cfda.gov).

Reliability of key servers and applications in the PSAP and communications center is imperative and becoming even more so. By understanding the various approaches to computer server availability, the uptime expected and the options in current server technology, municipalities like Kansas City have more tools ready to protect and serve.

John Leonard is public safety segment manager at Stratus Technologies (Maynard, MA). He can be reached via e-mail at John.Leonard@stratus.com.

Kansas City, MO, has a keen understanding of the need for continuous reliability. Kansas City's Information Technology Department recently launched an RFP process to pick a continuously available hardware and software solution for the city's emergency CAD and records management systems. On average, Kansas City's 9-1-1 emergency dispatch center receives 420,000 emergency and general dispatch calls a year. This includes police, fire and EMT services. The city also needed a platform for its records management application, which tracks and stores nearly 128 million transactions or police, fire and other yearly updates to emergency services records.

The city turned to CAD software solutions vendor Tiburon Software (Fremont, CA) and its Central Region office in Austin, TX, for the project. Together with its fault-tolerant server partner Stratus Technologies (Maynard, MA), Tiburon was able to meet and exceed the stringent reliability demands required to support Kansas City's Public Safety Information Systems. The parameters Kansas City needed Tiburon and Stratus to satisfy are typical of what most municipalities should expect when evaluating CAD and other emergency services application platforms: application functionality and integration; customer support, warranty and maintenance; and vendor experience and resources.

"Municipalities and public safety providers expect these systems to perform flawlessly 24/7, especially when the needs placed on the system are extraordinary," says Rick Brisbin, Kansas City's Project Manager for the Public Safety Information Systems Project. "Vendor experience and top quality hardware and software components are at the top of the priority list when choosing a technology partner. For us, Tiburon's extensive experience with similarly sized cities, its reputation of top-quality support and maintenance, and the Stratus reputation for reliability and solid performance made Tiburon the consensus choice for the project."

FIVE NINES ENOUGH?

Many aerospace, medical, automotive and telecommunications customers find a 1-in-100,000 downtime [equivalent to five nines] unacceptable. If, for instance, the system that runs the avionics of a fly-by-wire commercial airliner suddenly becomes unavailable when the airliner is one hundred feet over the runway, the pilot may lose control of the flaps and the throttle with catastrophic results. Because there are about 100,000 commercial airline flights each day, using a 99.999% uptime RTOS averages out to one commercial airliner crash per day, which is clearly impractical for aircraft flight systems.

The average commuter uses the micro-processor-controlled antilock brakes in his or her car about one hundred times per day. If the high-availability system that controls the antilock brakes of a car is unavailable, the car may rear-end the car in front of it, fail to make the curve in the road, or sail through a red light into cross traffic. With more than 100 million commuters in the United States, each hitting the brakes one hundred times per day, using a high-availability system could result in 100,000 product liability law suits per day.

The range of such safety-critical applications goes far beyond the obvious. If a telecommunications switch fails, an entire city can lose telephone service, causing lifesaving 9-1-1 calls to be lost. If traffic signals malfunction, cars will collide in intersections. If power distribution systems fail, the resulting blackout will cause traffic signals, telephone switches and emergency response systems to fail.

—Dan O'Dowd, CEO, Green Hills Software, from www.techonline.com.