

the *Availability Digest*

www.availabilitydigest.com
[@availabilitydig](https://twitter.com/availabilitydig)

911 Systems Experiencing Unacceptable Availability

March 2014

Around the world, we depend upon emergency numbers for critical police, fire, and medical support. Virtually every country has one or more emergency numbers. In the United States, Canada, Mexico, and many Latin American countries, the emergency number is 911. In Europe, Russia, Ukraine, and other countries, the emergency number is 112. In Afghanistan, you dial 119 for police or fire emergencies or 102 for medical aid. Most GSM mobile phones will automatically dial the local emergency number.



Emergency calls are answered by telephone operators or dispatchers. If a call is answered by a telephone operator, the operator will determine the appropriate dispatcher to which to route the call. It is the dispatcher's responsibility for getting the appropriate emergency services to the site.



Stratus Technologies

When a dispatcher receives a call, he or she must be provided with all of the emergency information and the location of the nearest available emergency vehicles so that appropriate dispatch instructions can be issued. Most emergency call systems (certainly the larger ones) are dependent upon computer-aided dispatch (CAD) systems. These can be large server farms (or virtualized farms) providing multiple applications, such as computer-based telephony, computerized radio control, automatic number verification, location mapping, automatic emergency vehicle location, and real-time access to dispatch logs and police, fire, rescue, motor-vehicle, and court records.

If any one of these applications should fail, the dispatching of emergency services can be severely hampered, with potential loss of life or property. It is imperative that all emergency call centers be configured for the maximum availability of their computer services. Unfortunately, there is much evidence that this criterion is not being met, and often not being met by a large margin. We look at some of that evidence in this article.

911 Services

In the United States, Canada, and Mexico, the emergency number is 911. The first 911 system went into operation in the United States in 1968. Since then, the term "911" has become synonymous with public safety. The formal name for an emergency call center like the U.S. 911 centers is Public Safety Answering Point (PSAP). There are currently over 6,000 PSAPs in the United States, handling an estimated 240 million 911 calls per year.

A classic case of 911 failures occurred when New York City upgraded its 911 system in mid-2013.¹ It upgraded its 40-year old green-screen mainframe CAD system with a new, web-based system that was

¹ [New York City's New 911 System Goes Down Four Times](#), *Availability Digest*, June 2013.

being used by 2,500 other emergency call centers around the world. The system was exhaustively tested for six months before being put into service at the end of May, 2013.

Even with all of the testing, the system failed on the first day, and then failed again ... and again ... and again. It failed four times in three days, often for up to a half hour. During this time, operators taking emergency calls had to write the emergency details onto slips of paper and give them to runners who ran with the slips to the dispatchers. This is not the kind of 911 operations we expect.

Next-Generation 911 Services

The current 911 services are based on telephone communications. The receiving operator or dispatcher obtains the details of the emergency from the caller and arranges for emergency services to be dispatched. If the call is from a land line, the address of the caller is generally known. If the call is from a mobile phone, typically only the approximate location is known; and the actual address must be obtained from the caller.

However, with the advent of smart phones, much more information is now available. Callers can communicate via text. They can take still photos and videos showing more detail of the emergency. It is estimated that currently over one-third of 911 calls originate from wireless devices, and the expectation is that this percentage will continue to grow.

The next generation of CAD systems, dubbed NG911, will allow CAD systems to receive, capture, and route digital communications such as photos and text messages received from virtually any communications device to dispatchers and emergency responders. For instance, a caller could take a picture of a tanker carrying hazardous material that is involved in an accident. The photo could be relayed to dispatchers so that they know in advance the type of hazardous material they will be dealing with and can dispatch the appropriate emergency equipment to the site.

NG911 systems are necessarily going to be more complex and fault-prone than today's systems. Configuring for availability will be even more important.

Stratus Technologies' Public Safety PSAP Survey

Stratus Technologies has just published its second annual survey of PSAP's, the 2013 Public Safety PSAP Survey Results.² It received responses from over 900 North American PSAP professionals on topics ranging from population covered, call volume, staffing, and NG911 migration to statistics on system availability. Its findings on PSAP availability were appalling.

Disaster Preparedness

Even though there is a nationwide focus on PSAP availability, the survey found that 20% of all PSAP's had no disaster backup site for their CAD systems. This means that a disaster could take down the entire 911 service. This is exactly the time when a 911 service is most critical. During a severe earthquake, flood, or storm (such as Superstorm Sandy), people in distress would not be able to call for emergency help.

Even more disconcerting was that almost 20% of PSAP's did not even have a disaster recovery plan. If one of these areas lost its PSAP, there was no plan for alternative emergency communications.

Dave LeClair, Senior Director of Strategy at Stratus Technologies, stated in an interview that:

"The fact that 20 percent of PSAPs do not have a physical backup PSAP location in case of natural disaster or catastrophic outage made sense possibly due to the cost of maintaining backup locations.

¹ http://www.availabilitydigest.com/public_articles/0806/nyc_911.pdf

² <http://www.stratus.com/~media/Stratus/Files/Library/WhitePapers/WP-PSAP-Survey-Results.pdf>

However, the fact that one out of every five PSAPs does not have a contingency plan in place was alarming. We're planning a follow-up question for the next survey. We'd like to know why they don't have a plan in place. Budget? Lack of expertise?"

Outage Experience

The survey found that 70% of PSAP's had experienced at least one outage in the last twelve months. Almost half had two or more outages. Sixteen percent had five or more outages. Seven percent had more than ten outages. Outages seem to be a fact of life for PSAP's.

73% of the smallest PSAP's (serving a population under 20,000) experienced outages, while 65% of the largest PSAPs (serving a population over 200,000) suffered outages. This indicates that the frequency of outages seems to be relatively independent of the size of a PSAP.

Only 11% of outages were recovered rapidly (under one minute). 58% of all outages lasted for more than fifteen minutes, and 28% lasted for more than an hour. The average number of missed calls as a function of call volume and outage duration is shown in the following table:

Number of 911 Calls Per Year						
	10,000	25,000	50,000	100,000	250,000	500,000
1 minute	0	0	0-1	0-1	0-1	1
15 minutes	0-1	1	2	3	7	14
30 minutes	1	2	3	6	14	29
1 hour	2	3	6	12	29	57

Hourly outages were worst for PSAPs handling large call volumes. 34% of PSAPs with call volumes over 100,000 calls per year had outages exceeding one hour. For a PSAP that handles 500,000 calls per year, a one-hour outage means that an average of 57 calls may be missed. How many lives is this? How many destroyed buildings is this?

The Movement to Managed Services

Many PSAPs are moving to integrated environments that provide management of large server farms. One such move is to virtualization.

32% of all PSAPs are now operating their CAD systems in a virtualized environment, and 19% are planning to virtualize. 49% of PSAPs are not planning to move to virtualization. However, the survey shows that 53% of PSAPs that are running virtualized environments experienced one or more outages last year exceeding 15 minutes. Clearly, virtualization is not a solution to high availability.

15% of PSAPs are currently running in a cloud, and 13% plan to do so. But based on current cloud availability of three 9s (8 hours of downtime per year), the cloud is not likely to prove to be a solution to the PSAP availability problem either.

High Availability Options

Looking at different approaches to availability inevitably leads us to a discussion of "9s." If an application is up 99.9% of the time, we say that it has an availability of three 9s.

The number of 9s is directly related to the amount of time that a system is down. If a system is up 99.9% of the time, it is down 0.1% of the time. Since there are 8,760 hours in a year (excluding leap years), three 9s means that the system will be down 8.76 hours per year. The relation of 9s to downtime is as follows:

Nines	Uptime	Average downtime per year
2	(99%)	87h 36m
3	(99.9%)	8h 46m
4	(99.99%)	53m
5	(99.999%)	5m
6	(99.9999%)	0.5m

A proper value for PSAP availability should be at least five 9s (down five minutes per year). If we look at various approaches to PSAP configurations, field experience has shown the following availabilities:

Conventional unmanaged commodity servers	99% (down 87.5 hours/year)
Hardened servers (some redundancy)	99.9% (down 8.8 hours/year)
Clusters	99.95% (down 4.5 hours/year)
Virtualization	99.95% (down 4.5 hours/year)
Clouds	99.9% (down 8.8 hours/year)

None of these solutions achieve the desired availability of five 9s or better. This gets us into the realm of fault-tolerant systems. There are two types of fault-tolerant systems – software fault tolerance and hardware fault tolerance. These are systems that provide two processing systems, with one system ready to take over or with both systems active in the processing. The only way for a fault-tolerant system to fail is for both systems to fail simultaneously (more likely, the second system fails after the first system has failed but before it has been recovered). Thus, their availabilities are very high. Typical availabilities for fault-tolerant systems are:

Software fault tolerance	99.995% (down 26 minutes/year)
Hardware fault tolerance	99.999% (down 5 minutes/year)

With software fault tolerance, two systems are provided. Each is fully configured to run the application, but only one is actively in production. The other system is an active standby, ready to take over on a moment's notice. Its database is kept in synchronism with the production system via real-time data replication. Should the production system fail, all activity is rerouted to the standby system, which can resume processing rapidly. Software fault-tolerant systems have an availability of 99.995% (almost five nines), exhibiting a yearly downtime of about 26 minutes.

Hardware fault-tolerant systems employ two systems that are actively processing all activity. If one should fail, the other system continues processing. Field experience has shown that hardware fault-tolerant systems have availabilities exceeding five 9s.

An example of a commercially available software fault-tolerant system is everRun Enterprise from Stratus Technologies, with an availability of 99.995% (down 26 minutes/year). Stratus' ftServers are hardware fault-tolerant systems with continuing field measurements showing availabilities in the order of 99.9997% - almost six nines (1.5 minutes of downtime per year).

Another option for larger systems is HP NonStop systems. These systems employ a mix of hardware and software fault tolerance and achieve five 9s of availability (5 minutes of downtime per year).

What is the value of life? Is it worth the investment in fault-tolerant systems? This is a critical decision that each town, city, county, or state must make.

Summary

The reliability of today's emergency call systems is far from what it should be. Statistics show that the availability of these systems is, at best, dismal. Classic approaches to high availability such as clusters miss the mark by a wide margin, as do virtualized systems and cloud deployments. All of these approaches result in hours of downtime per year, during which life and property are at great danger.

Fault-tolerant systems are commercially available that reduce downtimes to just a few minutes per year. Are the lives saved worth this investment?