

System MTBF and High Availability in a Fully Redundant NENA-Compliant PSAP Controller for Disruption-Free 911 Dispatch Centers

Geoffrey Hungerford, Ph.D. and William O'Neil ^{a)}
Experient Corporation, 295 Clayton Street, Suite 200, Denver, Colorado 80206

(prepared 4 January, 2004)

Abstract: The mean time between failure (MTBF) and availability of Experient's fault-tolerant PSAP Controller is estimated using a Monte Carlo simulation. PSAP Controller features and functions are discussed in terms of their impact, as both liabilities and assets, to system MTBF. Traditional single-point-of-failure subsystem components that have adversely affected reliability, such as single CPU, operating system and classic interface communications, are avoided through redundancy, remote monitoring and disruption-free service and replacement. The simulation estimates system availability at better than five nines (99.999%) and system MTBF to be about two million hours.

I. INTRODUCTION

In mission critical applications, end users frequently ask equipment vendors for some estimation of how long their equipment will perform, on average, in a fully functional condition. For individual printed circuit cards or integrated circuits the mean time before failure (MTBF), also known as hardware MTBF or component MTBF, has been a primary focus. Ideally, MTBF is calculated based on failure rates from field data over many thousands or tens of thousands of identical components operating collectively for many millions of hours under conditions representative of the user environment.

When this data is not available, or there is not enough of it, hardware MTBF can be estimated. At the component or card level this is done according to two widely recognized failure rate models: The first is MIL-HDBK-217, the U.S. Military Handbook for "Reliability Prediction of Electronic Equipment" and is published by the Dept. of Defense. The second model is the Bellcore Reliability Prediction Procedure, TR-NWT-000322. Bellcore's original charter as research arm for the regional holding companies (RHC's) diminished with deregulation and it was bought out by Science Applications International Corp. in 1997 to become Telcordia

Technologies under SAIC. The Bellcore standard is available today as Telcordia SR 322¹.

Recently, sophisticated systems have become increasingly complex with hot spares, component redundancy and more comprehensive maintenance agreements. As a result, the concept of hardware MTBF has become less useful. In fault-tolerant systems of redundant components, individual subsystem component failures are invisible to end users because they do not result in a disruption to the system function. If covered under a blanket maintenance agreement, such as the Experient controller, failures have no financial impact on the end user as well. It is for these reasons that storage servers and advanced embedded applications commonly report a *system MTBF*². Although a system of redundant components may be too complex to apply practically the Bellcore parts count method, the system itself can still be viewed as a single black-box component with an effective system MTBF³ (expressed in hours) and system availability⁹ (expressed as a fraction or percentage). These measurements are far more useful since they reflect

^{a)} Author to whom correspondence should be addressed; electronic mail: bill.oneil@experient.com



all disruptions detectable by the customer and exclude all those that are not.

In this paper, we explore estimation of both the system MTBF and availability using a mathematical simulation based on the reported MTBF of subsystem components. It is referred to as a Monte Carlo simulation because it is based on random number generation for subsystem component failures and yields useful estimations of MTBF over very large numbers of events.

II. PSAP CONTROLLER DEFINITION

911 dispatch and emergency call centers, known as public safety answering points (PSAP), all contain a common component referred to by the National Emergency Number Association (NENA) as a PSAP Controller⁴. It is also commonly referred to as an ALI controller, and its function can be described in the system architecture diagram in figure 1.

The primary communication interfaces that define a PSAP Controller are summarized as follows:

- 1) **ANI link**: Provides the PSAP Controller with Automatic Number Identification (ANI) of an inbound emergency call over a CAMA trunk or caller ID (CLID) enabled POTS line. In some cases, this is a single RS-232 serial connection. The Experient PSAP Controller typically interfaces with a Tone Commander (formerly Proctor Assoc.) ANI-Link.
- 2) **ALI Database**: Redundant paths to the multiple street address guide (MSAG), or ALI database, allow access to caller's current billing information and dynamic update of cell phone location. Availability of this data to dispatchers constitutes the E911 service level.
- 3) **CAD port**: NENA standard, RS-232 serial interface to a computer-aided dispatch (CAD) and mapping subsystem used by the dispatcher.
- 4) **LAN interface**: Allows for networking of 911 call center computers and services, including local

databases, alarm condition monitoring and remote maintenance. It also provides the foundation for mobile or wireless dispatcher positions as well as remote networked PSAPs and call overflow management.

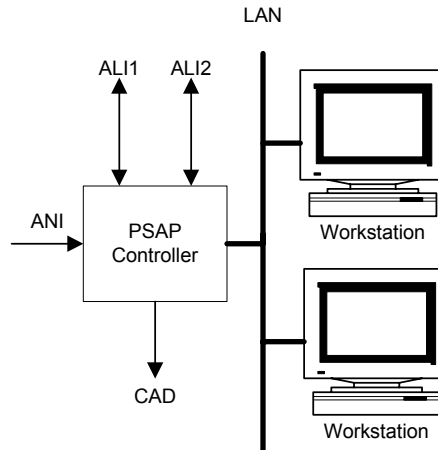


Figure 1
PSAP Controller Definition

III. DEMAND FOR INCREASED RELIABILITY

The importance of reliability and continuous, disruption free service of a PSAP controller is addressed specifically by NENA⁵. It is more attainable than ever before by selection of hardware and software components developed over the past few years.

Many existing (legacy) PSAP Controllers use a single PSAP Controller CPU and contain single points of failure that significantly reduce availability and system MTBF. Experient's dual PSAP Controller architecture, failover cluster network and device server interfaces eliminate single points of failure and minimize sources of PSAP disruption, increasing system MTBF and availability.

In addition to recent advances in reliability of individual off-the-shelf subsystem components, such as hard disk, power supply, network interface card, device servers and operating system, the cost of these

components has dropped as well, so that component redundancy is now more cost effective than ever and provides for sharply increased system MTBF and availability of the PSAP Controller.

IV. FEATURES AFFECTING RELIABILITY

The Experient PSAP Controller achieves extended periods of disruption-free service by eliminating single points of failure. The key features are shown schematically in figure 2 and their impact on reliability is explained as follows:

Dual Linux-based Servers

Dual Dell PowerEdge blade servers allow individual component replacement or software updates on any single server while the other manages all PSAP functions. While both controllers are identical and fully functional, one acts as the primary service provider, giving the PSAP the appearance of having a single server in the back end. However, both controllers monitor the PSAP environment, as well as each other, at all times. As a result, the secondary controller can quickly and seamlessly become the primary in the event either detects performance difficulties in the original primary controller.

Continuous Remote Monitoring

Alarm conditions are monitored both locally and remotely. Experient's maintenance agreement provides for automated notification of alarms, errors and system status sent directly to e-mail capable cell phones carried by maintenance technicians. As a result, exception conditions in the PSAP Controller environment can be addressed within minutes, greatly reducing the probability of PSAP disruption. In addition, each of the two Linux controllers can be securely accessed by maintenance technicians via the internet for collection of more extensive event logs and system data when troubleshooting.

Remote Software Update

At present, the software update process most familiar to PSAP supervisors involves bringing down key servers for some period of time. Experient's dual/redundant controllers eliminate this disturbance by allowing its servers to be upgraded and rebooted one at a time from a remote service center. This can be safely done, even at busy call times, since one controller alone fully services the entire PSAP.

Watchdog Timer/Systematic Reboot

A hardware watchdog in both controllers automatically protects the servers from any software anomaly that would disturb the normal flow of execution by performing a reset and reboot cycle, restoring the operating system and application software. In addition, the PSAP Controller software itself can initiate an orderly power down and reboot sequence in the same manner as a desktop PC user might reboot his system when unusual behavior is suspected. The application software does this when it detects any condition that might be corrected or diagnosed in a reboot cycle.

Temperature Monitor and Fan Sensors

Experient Controllers are equipped with temperature monitors and fan sensors that allow PSAP supervisors and remote maintenance technicians to be notified when inadequate cooling conditions exist. This is done automatically via e-mail enabled cell phones so that technicians can begin the corrective action process within minutes of detecting any cooling problem.

Serial Device Servers

In order to allow access to single serial communication links by multiple controllers, serial device servers are used to convert serial data to and from



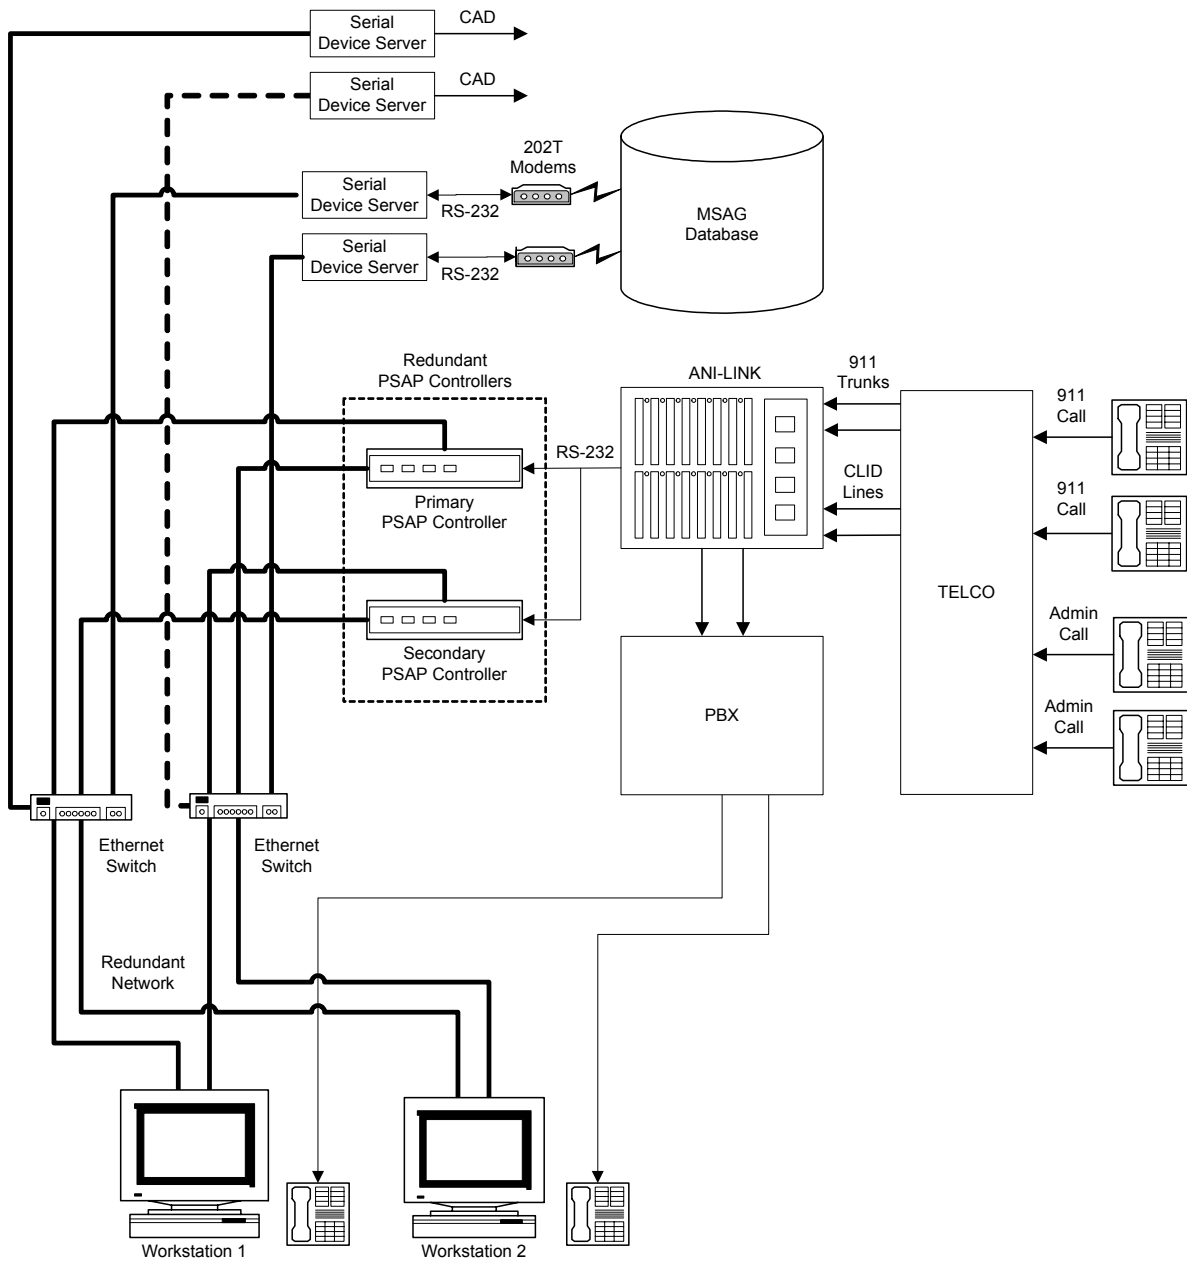


Figure 2. Fully Redundant PSAP Controller with dual LAN and CAD ports. This configuration may support any number of workstations or nodes on the LAN.



LAN packets. These are accessible to all CPU's on the network. This increases overall PSAP reliability by allowing multiple PSAP Controllers access to serial links such as the ANI-Link, ALI database or CAD ports. However, if the PSAP Controller has only a single CAD port, the corresponding serial device server becomes part of a single point of failure and can effectively determine the MTBF of the PSAP Controller as a whole. For this reason, Experient uses high MTBF device servers.

Multiple CAD Ports

The Experient PSAP Controller has the ability to support multiple CAD ports. Although such a configuration is not specified in NENA standards⁵, it overcomes a single point of failure in the conventional CAD serial interface specification.

Multiple ANI Interfaces

As with the CAD interface, the serial ANI-Link interface traditionally provides a sole communication path that is itself a single point of failure. However, the Experient PSAP Controllers can support multiple ANI-Link interfaces that overcome this limitation. When only one ANI-Link serial interface is available, downstream redundancy can be achieved by simply splitting the serial RXD line (receive data) since communication is unidirectional, going only from the ANI-Link to the PSAP Controller.

Multiple MSAG Interfaces

In its simplest configuration the PSAP Controller communicates with an ALI database through 202T modems over a single pair of leased lines, often operating at 1200 baud. For larger call centers with multiple dispatcher positions, and where baud rates cannot be increased, 1200 baud is not a high enough data transfer rate for MSAG data during busy call periods. To overcome this, multiple pairs of leased lines are supported by the Experient PSAP Controller in order to increase bandwidth or to support multiple MSAG databases using different record formats. These serial links are accessed using serial device servers so either PSAP Controller can access them on the network.

High Availability/Failover Cluster Network

PSAP Controllers support two LAN ports for each CPU, providing redundant TCP/IP communication paths at both the software and hardware level. The switches are Dell PowerConnect unmanaged 24-port ethernet switches with an MTBF of 5 years¹⁰.

Redundant Synchronized SQL Databases

All PSAP event and call data are stored in two identical SQL databases arranged in a mirrored RAID-1 architecture across dual network nodes. The databases are synchronized continuously so records are always accessible even when one PSAP Controller is down for maintenance or upgrade. As a result, no data will be lost in the event a hard disk needs to be replaced. After installation of a new disk or PSAP Controller, the new system will synchronize automatically to contain an original copy of the SQL database, along with the complete PSAP history. The SQL database is managed in such a way so that random loss of power does not jeopardize data integrity. This behavior is a significant advancement over existing database management and is discussed by Solin⁶.

Disruption-Free Hardware Replacement

Hardware components in the Experient PSAP Controller are replicated in the redundant architecture so that any one component can be powered off and replaced without any loss of service in the PSAP Controller.

20-min. UPS Backup

An uninterruptible power supply provides for continuous PSAP Controller service during periods of power outages that last for less than twenty minutes. NENA⁷ specifies the minimum allowed UPS backup period as 15 minutes.



NENA-Compliant XML Format

Experient’s PSAP Controllers support XML based data exchange as specified in the NENA Recommended Formats and Protocols⁴. This communication protocol is detected automatically, so there is no disruption in service when a particular MSAG, CAD or ANI-Link upgrades to Phase-II XML based communications.

Web Browser Interface

The Experient PSAP Controller supports a standard HTML web interface so that it can be configured, maintained and its database queried through any generic web browser. PSAP supervisors are also provided extensive reporting services through this interface over their LAN; no special software or hardware is required.

V. SIMPLIFIED MTBF ESTIMATION

Since critical functions in the PSAP Controller are replicated and managed as independent redundant resources, exposure to PSAP disruption becomes limited to periods when both primary and backup systems are simultaneously not functional. Since Experient PSAP Controllers are independent and receive identical inputs, the probability of PSAP disruption is further minimized by reducing the time taken to identify and restore any single failed component, called the Mean Time To Repair (MTTR).

System MTBF is strongly related to system architecture. Dell Computer Corp. performed a study similar to ours using a redundant network of PowerEdge 6450 servers (in this study we used the PowerEdge 650). In his work, Shetty⁸ models the availability and system MTBF of the *high-availability cluster* configuration. This is also known as a failover cluster and is identical to the Experient PSAP Controller network configuration, except the PSAP controller has additional subsystem components to manage PSAP specific I/O. Shetty’s work demonstrates the mathematical impact of both parallel and serial redundant components on system

MTBF and availability. A reliability block diagram for the Experient PSAP Controller is shown in fig. 3.

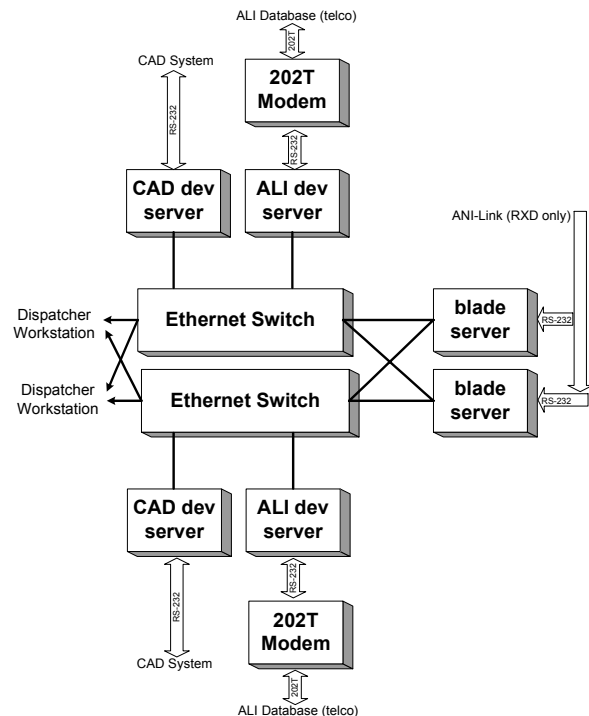


Figure 3. Reliability block diagram for Experient PSAP Controller.


The probability a component will perform, without failure, over a specific period of time t , is given by²

$$R(t) = e^{(-t / \text{MTBF})} \quad (1)$$

Equation (1) reveals an interesting aspect of MTBF. If a component has an MTBF of 4 years (35,000 hours), what is the probability of the component being operational after 4 years? We might be tempted to say it is a 50% probability if we are misled by the word “mean” in MTBF. In fact, it is less than 37% because this probability is related to MTBF in an exponential distribution.

The simple non-redundant component configuration shown in figure 4 can be modeled as a system according to

$$R_{\text{sys}}(t) = R_{\text{server}}(t) R_{\text{switch}}(t) \quad (2)$$






Figure 4. Reliability block diagram for non-redundant system of two network components.

because it is a serial configuration of components such that any individual component failure results in a system failure. As a result, the reliability of the system is less than that of any individual component.

The simple redundant component configuration shown in figure 5 has been modeled as a system according to⁸

$$R_{\text{sys}}(t) = R_{\text{sw1}}(t) + R_{\text{sw2}}(t) - R_{\text{sw1}}(t) R_{\text{sw2}}(t) \quad (3)$$

because it is a parallel configuration of redundant components such that a failure of either component leaves system functionality unaffected. As a result, the reliability of this system is much higher than that of any individual component.




Figure 5. Reliability block diagram for redundant system of two network components.

Equation (3) is determined by first expressing the probability of a system failure as a function of time

$$F_{\text{sys}}(t) = F_{\text{sw1}}(t) F_{\text{sw2}}(t) \quad (3a)$$

where the failure probability for any component i , is given by

$$F_i(t) = 1.0 - R_i(t) \quad (3b)$$

By substituting equation (3b) into equation (3a), and solving for $R_{\text{sys}}(t)$ we get equation (3).

However, equation (3) does not take into consideration MTTR. Although in our case the component repair process does not itself require a system disruption, a system failure can occur if a second component fails during the MTTR period. Meaningful estimation of the MTBF for a system of redundant components must depend on how long a failed component remains out of service before it is replaced. The reason that the component restoration time affects system MTBF is that the system will temporarily have a single point of failure for this period. Minimizing MTTR decreases the window of opportunity where a second component can fail and increases system MTBF.

Equation (3) is useful to demonstrate the effect of component redundancy on MTBF, but is not adequate to model a maintained system of redundant components because system MTBF is a function of MTTR. However, the MTBF can be estimated using a Monte Carlo simulation while viewing the PSAP Controller as a *discrete-event system* that easily accommodates the effect of MTTR.

For most computer systems, Steady State Availability is the most commonly quoted availability metric. It is, on average, the fraction (or percentage) of time that the system is fully functional throughout its lifetime. When sufficient field service data is available, it is calculated in terms of downtime per year

$$A_{\text{sys}} = \frac{(\text{total system hours}) - (\text{downtime hours})}{(\text{total system hours})} \quad (4)$$

The closer availability is to 1.00 (100%) the better. Fault-tolerant, highly available systems are

sometimes referred to as “four-nines” or “five-nines” which indicates the number of consecutive leading nines in the availability measurement: A system whose availability is 0.99995 (or 99.995%) is a four-nines system and is down, on average, only 26 minutes per year. Five-nines are considered very good. In fact, Motorola¹², in August 2000, claimed to have the first Linux-based telecom application to make five nines availability.

Availability is also related to MTBF and MTTR (mean time to repair)⁹:

$$A_{\text{sys}} = \frac{\text{MTBF}_{\text{sys}}}{(\text{MTBF}_{\text{sys}} + \text{MTTR}_{\text{sys}})} \quad (5)$$

If we have enough field service data, we can determine the MTBF of any component *i*, (including the entire system MTBF) in terms of its failure rate, FR, which is given by

$$\text{FR}_i = (\text{total failures})_i / (\text{total hours})_i \quad (5a)$$

where

$$\text{MTBF}_i = 1 / \text{FR}_i \quad (6)$$

when large amounts of data are available. MTTR can also be calculated directly from field service data:

$$\text{MTTR} = \text{total downtime hours} / \text{total failures} \quad (7)$$

Field service data includes three things: Which components failed, how long each failure took to repair, and the total system run time over which components had an opportunity to fail. With this data, we can calculate MTTR from equation (7), system MTBF from equations (5a) and (6) and availability from either equation (5) or (4).

When we don't have field service data, system MTBF can still be estimated using equation (6) if the field service data is simulated. In this study, the PSAP Controller was modeled using an asynchronous simulation, advancing from one time event to the next and describing the changes in state at these moments. The event time interval chosen for this simulation was 3 minutes, which is

approximately the time it takes for a Linux-based controller to reboot. It is also the shortest period for failure detection and restoration of any redundant PSAP Controller component, so that there is no significant mathematical accuracy gained by additional granularity. The duration of each simulation was over one billion hours, enough time for the MTBF to converge.

PSAP disruption is defined as any combination of failed components that would render the PSAP Controller dysfunctional in any way, for any period of time. This is not limited just to cases where two identical components fail. It also considers parallel redundancy as any two events leaving both blade servers dysfunctional, and serial effects such as simultaneous failure of the 202T modem on one leased line and failure of the device server on the other (see fig. 3).

Table 1 reveals the logic behind the simulation, which is coded in C and is freely available from Experient Corp. as open source code. When there are no single points of failure, any individual component can fail and be replaced under the maintenance agreement without a system failure. If the simulation shows that two or more components are out of service simultaneously, then table 1 is consulted to determine if these failed components constitute a system failure.

Another advantage of a Monte Carlo simulation (as opposed to complex analytical system models) is that the system configuration can be very easily changed, or made quite complex, and the model can be re-run by simply adjusting or expanding table 1. One disadvantage of the Monte Carlo model is that it takes substantial number crunching to run each simulation: For some scenarios, it took more than 10 hours.

Table 2 shows the subsystem components in our model with nominal values of MTBF and MTTR. The hard disk, microprocessor, memory, network interface cards and cooling have been bundled into the Dell PowerEdge 650 server as a single item with an estimated MTBF¹⁰ of 3.3 years (29,000 hours). Of

	software1	software2	blade srv1	blade srv2	switch1	switch2	ALI dev srv1	ALI dev srv2	CAD dev srv1	CAD dev srv2	ALI modem1	ALI modem2
software1		F		F								
software2	F		F									
blade srv1		F		F								
blade srv2	F		F									
switch1						F	F		F		F	F
switch2						F	F		F		F	F
ALI dev srv1						F	F					F
ALI dev srv2						F	F				F	
CAD dev srv1						F				F		
CAD dev srv2						F			F			
ALI modem1						F	F					F
ALI modem2						F	F				F	

Table 1. System failure conditions for two or more simultaneously failed components. An "F" at the intersection of the two components means the system is not functional while these two components remain out of service.

particular interest is the software component, recognized here as a potential disruption affecting system MTBF. Although software is not normally measured in terms of MTBF, it is useful when simulating a system of redundant components. Consider how a suspicious behaving desktop PC can be restored when rebooted. All software-based systems have some probability of software failure. Our model recognizes this and allows us to study the impact of this phenomenon by varying software component MTBF while leaving all else constant.

Software components have short MTTR because their failure mode can be corrected with a reboot cycle. During day-to-day PSAP Controller operation, application software, operating system (OS) and SQL database performance is closely monitored such that abnormalities are identified in a matter of seconds.

The server experiencing anomalies will be rebooted either by the application software itself or via the watchdog timer, which automatically cycles power on an affected CPU.

COMPONENT	MTBF	MTTR
Software	12 months	3 min
Ethernet Switch	5 years	48 hours
202T Modem	5 years	48 hours
Dell PE-650 blade server	29,000 hrs	48 hours
Serial device server	134,000 hrs	48 hours

Table 2. Inputs for system MTBF estimation.

MTTR for hardware components is a conservative estimation. It is based on the worst-case MTTR in the Experient maintenance agreement providing for failed component replacement in 48 hours or less. The time to repair (TTR) is actually a distribution whose maximum is 48 hours, so in practice it should take less time, on average, to replace a failed hardware component than this maximum. By assuming MTTR is 48 hours our model errs on the side of conservatism, predicting lower system MTBF and availability than it will probably achieve. Most components in the system can be identified as faulty and replaced in a few hours more than it takes to ship the replacement part by some form of overnight express. The system is designed so that all hardware components can be hot-swapped in just a few minutes without any disruption to PSAP Controller function, even during busy call periods.

An interesting aspect of simulating field service data is that the MTBF of individual components can be calculated from the results. When individual component MTBF is calculated according to equation (6), the values are close to those in table 2 for large numbers of simulated events. This gives us confidence that the model estimates system MTBF accurately, since individual component MTBF and system MTBF are both calculated in the same way. It also gives us an idea of how long to run the simulation. When all individual component MTBFs converge to values very close to their nominal values given by the individual component manufacturers, then we assume the system MTBF has also converged to something very close to its true value.

When an individual test case is run for more than one billion hours, these conditions are easily satisfied.

VI. RESULTS AND DISCUSSION

If system availability or system MTBF is estimated by any technique that only considers hardware components, then the estimation contains an embedded assumption that system software components have zero probability of failure. We suggest that a complete system model should accommodate this probability. However, for purposes of comparison with results reported by others, we will first consider the case where the probability of software component failure is zero. Our Monte Carlo simulation of the system shown in figure 3, with the component MTBFs shown in table 2, estimates that the Experient PSAP Controller system availability is slightly better than 0.99999 or 99.9999%. The estimated system MTBF is about 2,000,000 hours. All given that the system is regularly maintained such that failed components are promptly identified and replaced, and either

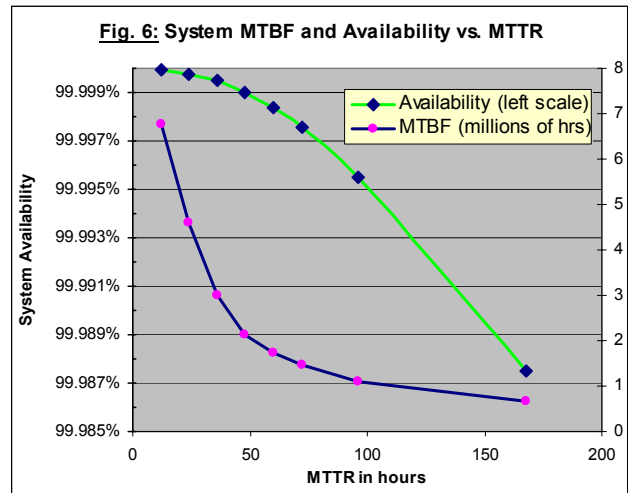
- 1) The CAD port can be made redundant, as in figures 2 and 3.

-OR-

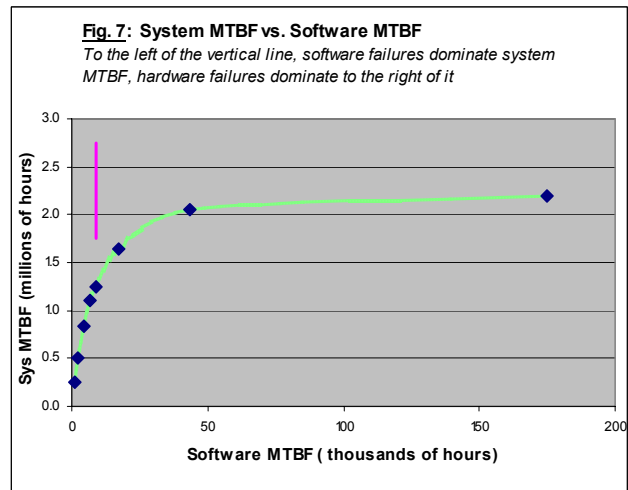
- 2) The Experient PSAP Controller runs in a network where there is no serial CAD port.

If the CAD port cannot be made redundant, then the MTBF of the PSAP Controller will be dominated by the single serial device server for this serial port, which means system MTBF should drop to somewhat less than 134,000 hours¹¹. When the model was configured with a single CAD port, it estimated MTBF at 129,000 hours (about 14 years).

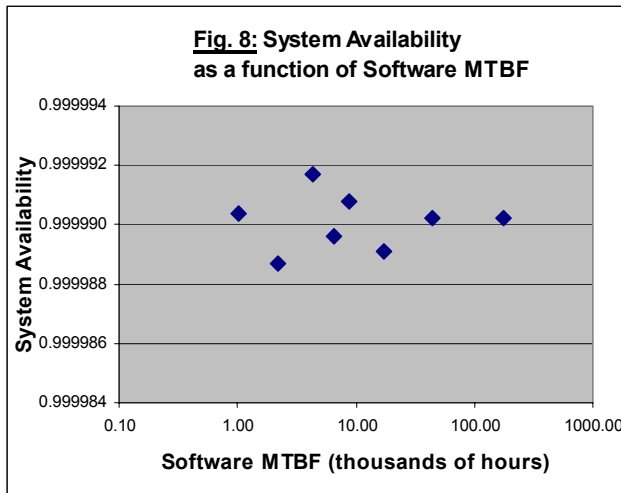
The model reported that the system MTBF is most sensitive to simultaneous failure of both blade servers, which is not surprising since this subsystem component has the lowest individual MTBF and highest MTTR.



Since the primary failure mode requires hardware replacement, the simulation was run again at various maintenance response scenarios to investigate system MTBF sensitivity to service response. The results are shown in figure 6 with system MTBF plotted as a function of MTTR.



The impact of software reliability was investigated by simulating the system with various values of software MTBF while holding all else constant. Overall PSAP



Controller MTBF was found to be sensitive to software MTBF and is plotted in figure 7. If the software MTBF is less than about one year (indicated as a vertical line in fig. 7), the simulation showed that a software reboot cycle was involved in most system outages. At higher levels of software reliability, most system failures were a result of hardware component failures exclusively. However, the software outages have a short duration of 3 minutes so that system availability is not sensitive to software MTBF, as shown in figure 8.

VI. CONCLUSIONS

Experient Dual redundant Linux-based embedded controllers and associated redundant hardware components allow for dramatically extended periods of disruption-free operation. Automated, remote alarm and status messaging allow Experient maintenance technicians to minimize the mean time to repair (MTTR), which increases the system MTBF of the PSAP Controller. System architecture allows hardware and software upgrades with no PSAP disruption. Redundant, synchronized SQL databases protect call event data as a RAID-1 storage server.

System availability of 99.999% can be achieved through component redundancy and service response of 48 hours or less.

The CAD port, if it exists, will be the system component that will dominate the PSAP Controller MTBF if it is not made redundant. Redundancy could be accomplished through a second CAD serial port, but could more effectively be implemented by having the CAD subsystem interface with the PSAP Controller over the LAN, eliminating the serial CAD port altogether. The Experient PSAP Controller will support either of these CAD enhancements.

System MTBF and availability both show sensitivity to component repair times, increasing as service response time becomes shorter.

System MTBF is sensitive to software reliability, but availability is not sensitive to software reliability because the software failure mode modeled can be corrected with a short reboot cycle.

ACKNOWLEDGEMENTS

- ¹ <http://telecom-info.telcordia.com/site-cgi/ido/newcust.pl?page=idosearch&docnum=S R-332&>
- ² Zzyzx Peripherals, Inc., *Mean Time Between Failure (MTBF) and Availability*, (June 13, 2001), http://www.zzyzx.com/products/whitepapers/pdf/MTBF_and_availability_primer.pdf
- ³ InformIT, *Network Storage Evaluations Using Reliability Calculations* (August 16, 2002), http://www.informit.com/isapi/product_id~{A9DA8ADB-2179-467E-9ADE-7731B216D87C%7D/content/index.asp
- ⁴ National Emergency Number Association (NENA), *Recommended Formats and Protocols for ALI Data Exchange, ALI Response & GIS Mapping*, NENA-02-010 (January, 2002).
- ⁵ National Emergency Number Association (NENA), *Recommended Generic Standards for E9-1-1 PSAP Equipment*, Issue 2 (August 23, 2000), Section 10 QUALITY AND RELIABILITY.



⁶ Daniel Solin, et al, *The MySQL Transaction and Replication Handbook*, Wrox Press Inc; 1st edition (March 25, 2003)

⁷ op. cit., *NENA Generic Standards*, section 3.17.3, POWER, page 37

⁸ Santosh Shetty, *Determining the Availability and Reliability of Storage Configurations*, Dell Computer Corp., (August, 2002)

http://www1.us.dell.com/content/topics/global.aspx/power/en/ps3q02_shetty?c=us&cs=RC956904&I=en&s=hied

⁹ Enrique Vargas, *High Availability Fundamentals*, Sun Microsystems; (November, 2000)

<http://www.sun.com/blueprints>

¹⁰ Dell Computer Corp., private communication; The MTBF for the PowerEdge 650 is approx. 29,000 hours. The MTBF of the 2124 unmanaged switch is 5 years.

¹¹ Neteon private communication; The MTBF for Moxa's DE-211 is 134,000 hours.

¹² D. H. Brown Assoc., Inc.; *Telecom's First Linux Offerings with 5 Nines Availability*, Motorola Computer Group; (August 2000)

<http://mcg.motorola.com/us/article/article0008.pdf>

