MTR 05B0000021

MITRE TECHNICAL REPORT

End-to-End Houston Emergency System Performance and Process Assessment

MARCH 2005

Robert Blount, Jr., MITRE
Shin-Jyh Frank Chang, MITRE
Fernando Galdos, MITRE
Rene Gonzales, DC Services
Brian Jones, L. Robert Kimball & Associates, Inc.
Shaun P. Morrissey, MITRE
Shu Glen Nakamoto, MITRE
Diane Z. Shaffer, MITRE
Bryan Smith, L. Robert Kimball & Associates, Inc.
Rick Taormina, L. Robert Kimball & Associates, Inc.

Sponsor:	City of Houston	**Contract No.:**	56479
Dept. No.:	G037	**Project No.:**	14058001-1A

NOTICE

This is the copyright work of The MITRE Corporation, and was produced for the City of Houston Government under Contract Number 56479, and is subject to the terms and conditions stated therein. No other use other than that granted to the City of Houston Government, or to those acting on behalf of the City of Houston Government, under the Contract is authorized without the express written permission of The MITRE Corporation. For further information, please contact The MITRE Corporation, Contracts Office, 7515 Colshire Drive, McLean, VA 22102-7508, (703) 883-6000.

© 2005 The MITRE Corporation.

MITRE

Corporate Headquarters
Bedford, Massachusetts

MITRE Project Approval: _____
Robert Blount, Jr.

Executive Summary

The City of Houston began to consolidate its public safety emergency systems in September 2000 when the City Council approved a lease/purchase agreement formalizing the creation of the Houston Emergency Center (HEC) facility. This consolidation was part of Mayor Brown's management improvement initiatives in 1998[1]. The facility was designed to house the personnel and some, but not all, of the systems that supported the Houston Police Department (HPD), Houston Fire Department (HFD), and Emergency Medical Service (EMS) call takers and dispatchers operations. The Computer Aided Dispatch (CAD) system was one of the systems located in the HEC facility that supported common and shared call taking and dispatching operations. The new CAD was acquired through an upgrade to the existing HPD CAD system. This new system replaced both the HPD and HFD CAD systems and provided interfaces to external systems including the Greater Harris County 911 emergency network, the mobile data terminals (MDTs), and HPD Record Management System (RMS).

The CAD has experienced several major outages prior to and since system acceptance. These outages have led to concerns with the performance of this new system. The City of Houston executed a contract with The MITRE Corporation to conduct an end-to-end performance and process assessment of the new system. The scope of The MITRE Corporation effort was to analyze the performance and processes of the public safety data systems located at the HEC. The other data and radio systems were not included in the contract. The MITRE Corporation assessed the following:

- Existing contracts and other documents that defined system performance and whether these performance requirements were met.
- Technical design of the system and overall end-to-end performance.
- Existing processes that support the system performance.
- Technical solutions and engineering processes that were needed to improve performance.

The team assessed the report which initially described the new call taker and dispatcher operations written by Arthur Andersen in 2002. Arthur Andersen was engaged by the City of Houston to provide a Technology/Management Plan for the new consolidated Houston Emergency Center. The principal purpose of this engagement was to prepare an organization structure, combining the related organizations in a unified command concept, and to prepare a budget[1]. The plan showed the need for a new system to support the recommended consolidated operations. The decision was made not to replace all of the voice, data, network and computer systems at once. Instead, based on budget and other constraints, the decision was made to upgrade the central components,

[1] Houston Emergency Center Technology Management Plan, 26 March 2002.

the CAD and RMS. The operations and management of the upgraded system was assigned to the city organization called the HEC. With the exception of the internal computer network within the HEC, the other public safety data and voice systems remained under the responsibility of the departments that operated and maintained them.

In general, MITRE's team findings focus on two constant themes. First, the public safety system needs additional resources and staffing to provide end-to-end management, sustainment and maintenance support. The team noted the high degree of customized code that was needed to support the identified operations and to provide the capability to interface to the external systems. The team also recognized many large cities and counties procure customized dispatch systems. However, customization typically requires long-term funding and resources, which have not been sufficiently provided for the City of Houston's public safety system. Second, the overall public safety system is not operated and maintained as a single homogeneous end-to-end system. HPD, HFD, HEC, and Information Technology Department (ITD) maintain separate systems that comprise the overall public safety communications system and departments work together to resolve critical issues. A homogeneous system would contribute significantly to performance enhancements.

The MITRE analysis began with the identification of performance requirements. The public safety system is comprised of the various systems managed by HEC, HPD, HFD and the ITD shown in Figure 1. There is not a single source document that specifies end-to-end performance requirements for all of these systems. With the exception of the CAD/RMS system, no formal requirement document exists for the systems. The majority of them are legacy systems that have been sustained by the City for a period of years. The HEC is responsible for managing the agreement with Northrop Grumman for the CAD, RMS, Message Switching System (MSS) and Storage Area Network (SAN). The MITRE team conducted an in-depth review of the requirements in the contract between the City of Houston and Northrop Grumman. The analysis of the contract showed the following:

- The majority of requirements contain: configuration specifications for the equipment, functional specifications to support call takers and dispatch operations, and specifications for network interfaces to the various voice, radio, and data legacy systems.
- The performance requirements primarily apply to the initial system design and to the acceptance test criteria. Thus, user response, system reliability, system monitoring, and engineering process requirements do not exist to sufficiently validate the current CAD/RMS performance against baseline requirements.
- The performance requirements primarily apply to the CAD application performance.

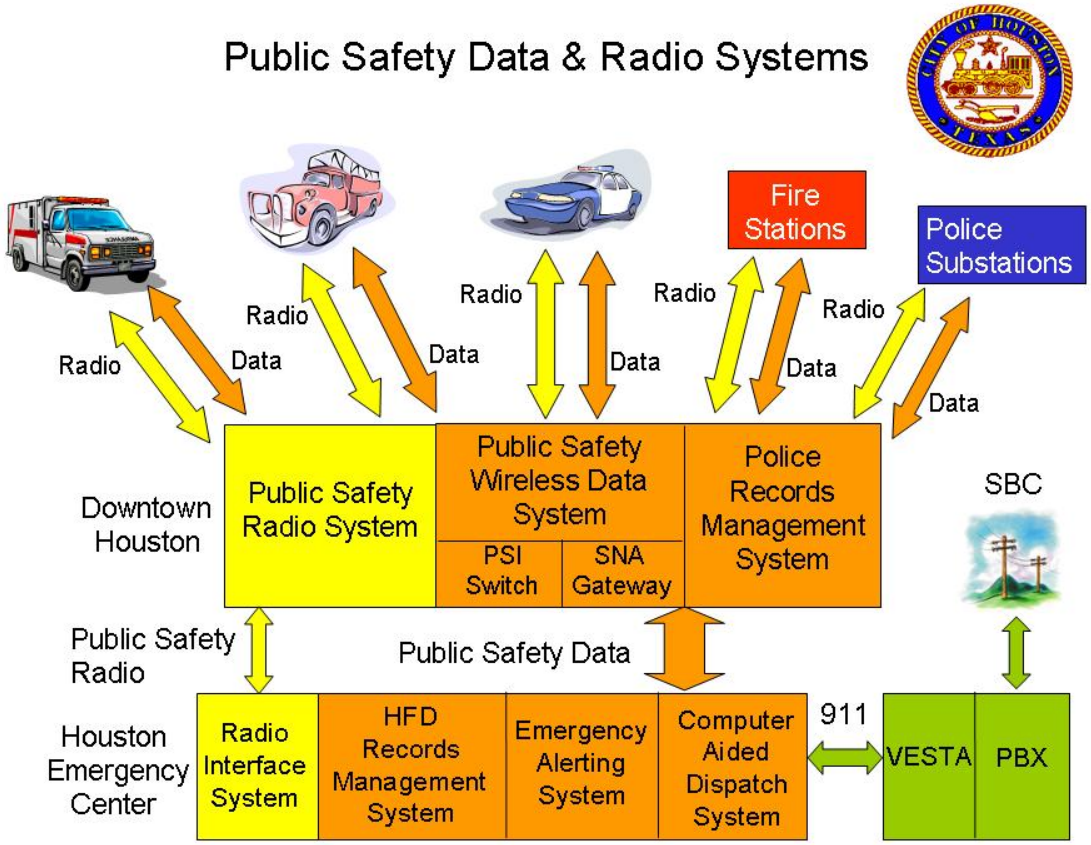

Figure 1. Public Safety Data and Radio Systems[2]

The team next evaluated the architecture and design of the current system. The evaluation showed that the systems' design reflects the requirements within the agreement between the City of Houston and Northrop Grumman as well as the current legacy systems architectures, the upgraded CAD/RMS, and SBC network. The team first assessed whether the failure of any equipment could lead to an outage that would prevent a large number of users from being able to access and use the system, referred to as a single point of failure. Several single points of failure were identified and analyzed. Where appropriate, technology or process changes were recommended to reduce the risk of outages due to failure of the equipment or system.

The team analyzed the architecture and design to identify equipment that may be at the end of its life. End of life means replacement parts cannot be obtained or the vendor has stated that the equipment will no longer be supported. The team identified equipment and software at its end of

[2] Drawing provided by HEC.

life that needed to be replaced. The team analyzed the security posture of the system by reviewing the security assessment conducted by Strategic Network Consulting (SNC) and assessing security vulnerabilities based on configuration information. [redacted] A back-up capability does exist to support call taking and dispatching voice and paper operations at another location. A concept of operations describes the planned operations and it identifies the need for data access to data resources.

After gaining an understanding of the requirements and the architecture, the MITRE team analyzed the performance of the system. The analysis included:

- Evaluation of the times that the system failed (i.e., outages).
- Determination of the system availability and reliability.
- Analysis of user and data performance.

The MITRE team was provided detailed summaries of the outages to the CAD that occurred prior to and after acceptance of the system. Outage was defined as the public safety dispatch system becoming unavailable to the majority of call takers and dispatchers. It did not include failures in the radio system. Ten of these outages occurred prior to acceptance and seven occurred after. The team noted that the frequency of the outages has decreased but that the time period of recovery of the outages increased. The analysis showed that the outages occurred for various reasons including equipment failure and human error. Further, the analysis showed that some incidents did not start from CAD directly, but they still caused CAD to be unavailable for operational use. The team recommended the following to prevent future outages:

- Elimination of major single points of failure.
- Expand system monitoring to identify and correct potential problems
- Increase staff skill base and training.

MITRE assessed the system availability and reliability of the HEC portions of the public safety system using the system outage data. The analysis calculated the system availability for the end-to-end system and the major subsystems, CAD, RMS, MSS, and SANs. The calculations showed that the system availability is approximately 99.7 while the CAD availability is higher at 99.8. Of the major systems, the SAN appears to have the highest rate of failure. The team provided suggestions for enhancing the overall and subsystems availability through better isolation of the CAD from other subsystems and component failures as well as improving the stability of the SANs. The team next calculated the daily availability figures for the major subsystems. This calculation showed that on seven days since system acceptance the availability has been less than 100%. Further, on three days the system was less than 80% available. The team recommends that a risk mitigation strategy be

developed to minimize system failures that can cause daily outages below the expected 99.9 threshold. The team recommends that the strategy include both equipment upgrades, such as disaster recovery, and process improvements. Finally, the team recommends that the availability of all major systems meet or exceed 99.99. Some systems such as the HPD RMS and the telephone switch exceed this statistics availability number.

The MITRE team attempted to assess the user and data performance in a variety of ways. The team measured end-to-end and subsystems performance using performance monitoring tools. These tools were limited to only a temporary snapshot of data performance of call takers and dispatchers transactions. MITRE also interviewed users and reviewed performance data provided by Northrop Grumman. The interviewees stated general satisfaction with the system performance but noted that at times performance was perceived as slow. The performance data could not be used to back up these concerns because the data was based on a monthly time period as opposed to an hourly or shorter time period. MITRE recommends that system monitoring, focused on the application layer, be performed on a more frequent basis and that feedback from users on performance be captured through formal processes.

The MITRE team assessed the processes that currently exist or are needed to support the system performance. The team noted that city-wide engineering processes are fragmented and not consistent among all departments. The HEC has adopted informal processes for problem resolution and system enhancement which support the public safety data system. However, these processes were not documented. The team identified the need for the City of Houston to implement city-wide risk management and configuration management processes for the public safety system. These processes will ensure a balance in addressing the technical performance, budget requirements, and enhancements to the public safety system.

The MITRE team recommendations were then structured to support the overall assessment. The first group of recommendations were drafted to identify solutions to reduce the number and risk of major outages that have been experienced in the past. These actions include:

- Establish staff positions responsible for end-to-end system management and integration of the public safety data and radio system.
- Eliminate single points of failure and establish effective automatic fail over.
- Increase system maintenance scope and time periods to provide a tiered 7x24 support team (technicians and public safety system help desk).
- Enhance HEC system performance monitoring and analysis.
- Enhance security [redacted]
- Document current processes and incorporate formal configuration management and risk management processes.

The next recommendations were intended to identify how the general system performance can be improved. They include:

- Measure and monitor the system's end-to-end availability.
- Develop end-to-end performance monitoring and analysis.
- Replace obsolete equipment and software. Establish a tighter control and tracking of equipment and software expected life through a formal configuration management process. At a minimum, the equipment identified as end-of-life in this report should be replaced.
- Enhance testing capabilities and processes.
- Identify and measure user and system performance statistics.

The last recommendations are those that are needed to support the system throughout its operational and sustainment life.

- Determine appropriate Contractor and City of Houston system operations and sustainment model.
- Develop end-to-end public safety strategic plan, architecture, and roadmaps.
- Incorporate disaster recovery system and processes.
- Develop a strategy to decrease application customization.

Table of Contents

1 Introduction 1-1

1.1 Background 1-1

1.2 Houston Departments Roles 1-1

1.3 Purpose 1-2

1.4 Approach 1-2

2 Performance Requirements Analysis 2-1

2.1 Strategic Vision 2-1

2.2 Performance Requirements 2-2

2.3 City of Houston Scope of Services 2-2

3 Architecture 3-1

3.1 Old System 3-1

3.2 Current System 3-1

3.2.1 [redacted] 3-6

3.2.2 RMS-HPD System From/To CAD 3-6

3.2.3 Storage Area Network Architecture 3-7

3.2.4 Integrated Database – Integrating CAD and RMS-Fire 3-7

3.3 Systems Operations and Support 3-8

4 Performance Analysis 4-1

4.1 Analysis of Outages and Errors 4-1

4.2 System Availability Calculations 4-8

4.2.1 Availability of the Overall System 4-9

4.2.2 Availability of CAD/RMS 4-10

4.2.3 Availability Since Acceptance 4-10

4.2.4 Result Summary of System Availability 4-11

4.2.5 Confidence Level and Confidence Limit for Availability Estimates 4-12

4.2.6 Monthly and Daily Availability 4-13

4.2.7 System Availability Enhancement 4-14

4.3 Workload and Performance Assessment 4-15
4.4 Scope of Services and Performance 4-20
4.5 Reliability 4-22
4.5.1 End of Life 4-22
4.5.2 Network Configuration Analysis 4-24
4.6 System Performance Monitoring 4-26
4.7 Security 4-30
5 Process Analysis 5-1
5.1 Outage Cause Summary 5-1
5.2 Critical Processes 5-3
5.2.1 Problem Resolution Process 5-5
5.2.2 System Enhancement Process 5-6
5.2.3 Change Control Process 5-7
5.2.4 Recommended Engineering Processes 5-9
5.3 Training 5-10
5.3.1 IT/Support Staff Training 5-10
5.3.2 User Training and Training Processes 5-11
5.3.2.1 Issues 5-12
5.3.2.2 Processes 5-13
5.4 Testing 5-14
6 Recommendations 6-1
Appendix A Referenced Documents A-1
Appendix B Operations of Call Takers and Dispatchers B-1
Appendix C System Availability Concepts and Calculations C-1
Appendix D HEC Call Volume Statistics D-1
Distribution List DI-1

List of Figures

Figure 2-1. Strategic Vision 2-1

Figure 3-1. Old Public Safety System Architecture 3-1

Figure 3-2. [redacted] 3-2

Figure 4-1. Progressive Point Estimates of MTBF 4-15

Figure 4-2. Call Volume Statistics and Outages 4-16

Figure 4-3. Performance Analysis Layers 4-18

Figure 4-4. System Utilization - Snapshot 4-19

Figure 4-5. Public Safety System Life-Cycle 4-23

Figure 4-6. [redacted] 4-25

Figure 5-1. Problem Resolution Process 5-6

Figure 5-2. System Enhancement Process 5-7

Figure 5-3. Change Control Process 5-8

Figure 6-1. End-to-End Portfolio 6-2

List of Tables

Table 3-1. [redacted] 3-3

Table 3-2. [redacted] 3-4

Table 3-3. [redacted] 3-5

Table 4-1. HEC System Outages 4-2

Table 4-2. Counts of Incidents by Problem Type 4-5

Table 4-3. Classification of Incident Types 4-6

Table 4-4. Summary of SIRT Problems 4-9

Table 4-5. Results of System Availability with Different Assessment Periods 4-11

Table 4-6. Confidence Limits of Availability With 95% Confidence Level 4-12

Table 4-7. Monthly Availability (Percentage of Uptime) After Acceptance 4-13

Table 4-8. Daily Availability (Percentage of Uptime) After Acceptance 4-14

Table 4-9. Measurement Results 4-21

Table 5-1. Process and Practices Relevant to the Outage Problems 5-2

Table 5-2. HEC Processes 5-4

1 Introduction

1.1 Background

The City of Houston public safety mission was served by separate and distinct public safety systems. PRC (hereinafter referred to as Northrop Grumman) installed the system used for police dispatching in 1987. The Fire Department dispatch and records management systems were developed in-house in 1985. The systems operated by these departments reached the end of their operations and sustainment life, and the City of Houston began to combine the public safety call taking and dispatching operations into a single facility. In September 2000, the City Council approved a lease/purchase agreement formalizing the creation of the Houston Emergency Center (HEC) at 5320 North Shepard. This facility housed the personnel and some, but not all, of the public safety system that supported the Houston Police Department (HPD), Houston Fire Department (HFD), and Emergency Medical Service (EMS) operations.

Northrop Grumman was contracted to upgrade the existing Police dispatch system and to expand the system to support integrated Fire/EMS dispatch and records management. The City of Houston also created a new organization, called the HEC, [redacted]

During the acceptance test period and since its acceptance, the public safety data system encountered technical problems that resulted in several system outages. Outages means the data system is unavailable to a majority of call takers and dispatchers. These outages led to concerns by the City of Houston and the general public on the reliability and performance of the upgraded system.

The MITRE Corporation was requested to conduct an end-to-end performance and process analysis of the public safety data system located at the HEC to address the following questions:

- Does the system perform in accordance with the agreement established between the City of Houston and Northrop Grumman?
- How can the system performance be improved?
- What processes can be implemented to improve performance?
- How does the system design and operations compare to other cities and counties with consolidated operations?

1.2 Houston Departments Roles

Several City of Houston departments, other organizations, and contractors operate and maintain the public safety system. The City of Houston departments include the HPD, HFD and ITD as well as the HEC organization. The other organizations include Greater Harris County 9-1-1 Emergency Network. [redacted] The description of the systems that the City of Houston, other organizations, and contractors operate, manage, or maintain are described in Section 3-2.

1.3 Purpose

The purpose of this analysis is to conduct an end-to-end performance and process assessment of the public safety data system located at the HEC.

1.4 Approach

The scope of The MITRE Corporation effort specifically encompassed the operational performance, processes, and architecture of the public safety data system; with primary focus on the design, acquisition, use and maintenance of the data systems located at the HEC. The public safety system includes data and radio systems that are external to the HEC and part of other departments within the City of Houston and Greater Harris County.

The analysis provided by MITRE followed the critical thread of performance to and from the HEC systems to the extent the external system status appeared to warrant further investigation and to the degree that information could be obtained. The detailed analysis focused on assessing performance and enhancements for the HEC systems. In addition, process and general system engineering performance assessments were conducted that are applicable to all portions of the public safety system. Thus, the analysis produced recommendations that extend to areas outside of the strict technical boundaries of the HEC and its component systems.

The MITRE team conducted the assessment through the review of documents and data listed in Appendix A; interviews of City of Houston, Greater Harris County, and Northrop Grumman staff; and by gathering performance data from the system. This process permitted the team to gather information that was indirectly and directly related to the performance assessment. The focus of the assessment was to identify alternatives and solutions to improve the performance of the system.

The team also contacted and gathered information from other cities and counties regarding the procurement, operations, and maintenance of their public safety systems. This information is contained within the analysis and served to validate the findings and recommendations. The City of Chicago provided data and information that was considered the most pertinent to the City of Houston's current environment. MITRE recommends that the City of Houston continue communications with the City of Chicago to share lessons learned.

2 Performance Requirements Analysis

2.1 Strategic Vision

The Arthur Andersen report titled "Houston Emergency Center Technology Management Plan" establishes a documented framework for the strategic vision of the HEC consolidated operations. The report focused on establishing a new organization structure, the budget necessary to initially support the organization, and consolidation of Neutral 911, HPD, and HFD/EMS call taking and dispatching operations. Figure 2-1 shows the key functions that the MITRE team believes are critical to the achievement of the strategic vision.

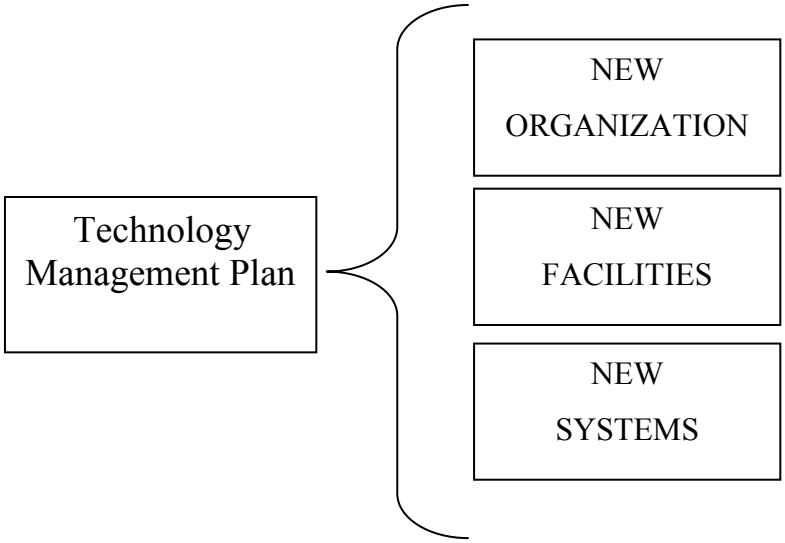

Figure 2-1. Strategic Vision

The report identified the following as the major benefits of the consolidation initiative:

- Break the current space barriers that exist between the different emergency offerings.
- Allow for open communications between the different emergency services.
- Increase overall visibility and understanding of all emergency service processes.
- Upgrade and standardize the technology supporting of the 911 system and reduce maintenance and support costs.
- Replace antiquated and crowded facilities with a state-of-the-art emergency communications center.
- Realize increased efficiencies by eliminating the three tiered system of 911 call taking/dispatching by consolidating the 911 neutral call takers with those in police and fire and eliminating numerous positions.

- Provide more effective emergency management through closer proximity and communication with the call center receiving citizen emergency service requests.
- Improve public safety through better and more timely response to emergency situations.

The Andersen report defined a clear business strategy and approach for the City of Houston's public safety system. While there appears to be high-level agreement and understanding between the various stakeholders (Mayor's Office, HPD, HFD, and HEC) involved in Houston's public safety service, there is not a clear end-to-end alignment of the public safety system across all of these organizations in regards to the implementation of the Andersen report business strategy.

2.2 Performance Requirements

The identification of documented performance requirements was the first task undertaken by the MITRE team to conduct the assessment. Specifically, MITRE requested request for proposals (RFPs) and other requirements documents for each of the major systems that comprise the public safety system. The only requirement document identified in response to this request was the contract between Northrop Grumman and the City of Houston for the CAD and RMS.

There is not a single source document that specifies end-to-end performance requirements for all of the public safety data and radio systems. With the exception of the CAD, RMS, MSS, and SAN, no formal requirement document exists for the other systems. The majority of them are legacy systems that have been sustained by the City for a period of years. The HEC is responsible for managing the agreement with Northrop Grumman. The MITRE team conducted an in-depth review of the requirements in the contract between the City of Houston and Northrop Grumman.

2.3 City of Houston Scope of Services

Based on staff interviews and review of the contract, the intent of the agreement between the City of Houston and Northrop Grumman was to identify specific system requirements were needed for the upgrade at the existing police CAD system to the new system. Therefore, the scope of service was drafted to specify the following: the preferred hardware configuration, functional requirements to support 911, HPD, and HFD operations, and specifications, to deliver interfaces to the legacy systems. The City of Houston also decided to pursue a sole source agreement with Northrop Grumman as opposed to releasing an RFP for a new system. The MITRE team did not assess the reasons for this decision but noted that it did impact how the requirements in the scope of services were written; i.e., the requirements were written for a known system, Altaris®.

The MITRE team initially reviewed the requirements individually to determine if they were adequate service level requirements and to identify which of them were performance based requirements. The analysis then focused on the performance based requirements to determine how they impact the current operations and performance of the overall system. These results are

documented in the report "Service Level Agreement Review and Assessment for The City of Houston, Texas", prepared by L. Robert Kimball and Associates.

The requirements in the "Scope of Services for the Houston Public Safety Dispatch System" address the delivery, testing, and maintenance of the hardware, software, application of the CAD, RMS, MSS and SAN. Application customization was needed by Northrop Grumman to meet the requirements in the scope of services for two major reasons: first, the separate and distinct operations of the HPD and HFD call takers and dispatching functions; and second, to develop and maintain the external interfaces to all of the other systems that were not being upgraded in sequence with the CAD and RMS.

The team recognized the need for the customization but noted the following long-term concerns. First, customizing of any system leads to increased maintenance and support costs for several reasons: the vendor's ability to leverage its resources when problems or changes become more difficult and the customized solution is further away from the vendor's base offering which impacts the degree of testing and training that must be maintained to support continual customization. Therefore, the City of Houston needs additional resources to support changes made to the baseline and to maintain the customized code. Further, the City needs a strategy that analyzes the need for and provides out year costs for continual system operations and support.

Sections 2 and 4, "CAD Upgrade Services" and "Workstation Requirements," have an impact on the current system performance. These sections specify the equipment requirements. These requirements do not specify who is responsible for upgrading the equipment nor the process for these changes to occur. This is not a major problem because the City of Houston and Northrop Grumman, through practice, will and has used change orders for equipment upgrades. The warranty for the equipment is also specified in this section. The warranty requires support during business days and allows for a four-hour response time. Because of the criticality of some hardware components, the MITRE team recommends that the warranty period be changed to 7 days a week and 24 hours a day (7x24) for major components and systems.

In Section 3, "Functional Requirements," the City requested software modifications and customizations. As noted in comments above, software customization is needed to meet these requirements. This customization will support the HPD and HFD operations. It also supports the following existing systems that the CAD is required to interface with: HPD RMS, HPD and HFD MDTs, [redacted], HFD Alerting and Paging, and [redacted]

Sections 5, 6, 7, 8, and 9 identify database conversion requirements and provide details for external interfaces. These requirements provide a good measure for monitoring current and future conversion and interface requirements.

Section 10, "Application Development Tools," the MITRE team was not clear in reviewing this section on its original intent. Based on feedback from interviews, this section was included to allow the City of Houston to provide city maintenance programming services, as had been done previously on the prior HPD and HFD systems. This section provides options for continual

maintenance of the system by City of Houston staff if needed. During interviews and review of processes, the team was not clear on whether the City of Houston and Northrop Grumman actually intends for the City or the contractor to maintain and monitor the system.

Section 14, "Maintenance and Technical Support," system availability of 99.9% is specified. This requirement is not usually acceptable for a high availability system solution. Usually 99.99% is the acceptable industry standard for "High Availability" and 99.999% for "Fault Tolerant" mission critical systems. This section also limits availability requirements to the CAD and RMS applications. While not a part of this contract, the availability requirements for hardware and other systems not included in the scope of services should be included in future change orders or other contracts.

In Section 15, "Installation, Testing and Acceptance," performance requirements regarding the system availability, transaction performance, system failover, and testing are contained. These requirements provided a level of detail for transaction performance that is not contained in other sections of the scope of services, i.e., they may only be applicable to the acceptance testing. Thus, the MITRE team could not determine if these requirements apply to the post-acceptance period and, therefore, recommends measures be specified for the current system to establish the expected baseline for the system performance.

3 Architecture

3.1 Old System

Figure 3-1 shows the architecture for the systems supporting call takers and dispatchers prior to the new public safety system. As shown in Figure 3-1, each department operated and maintained their own system.

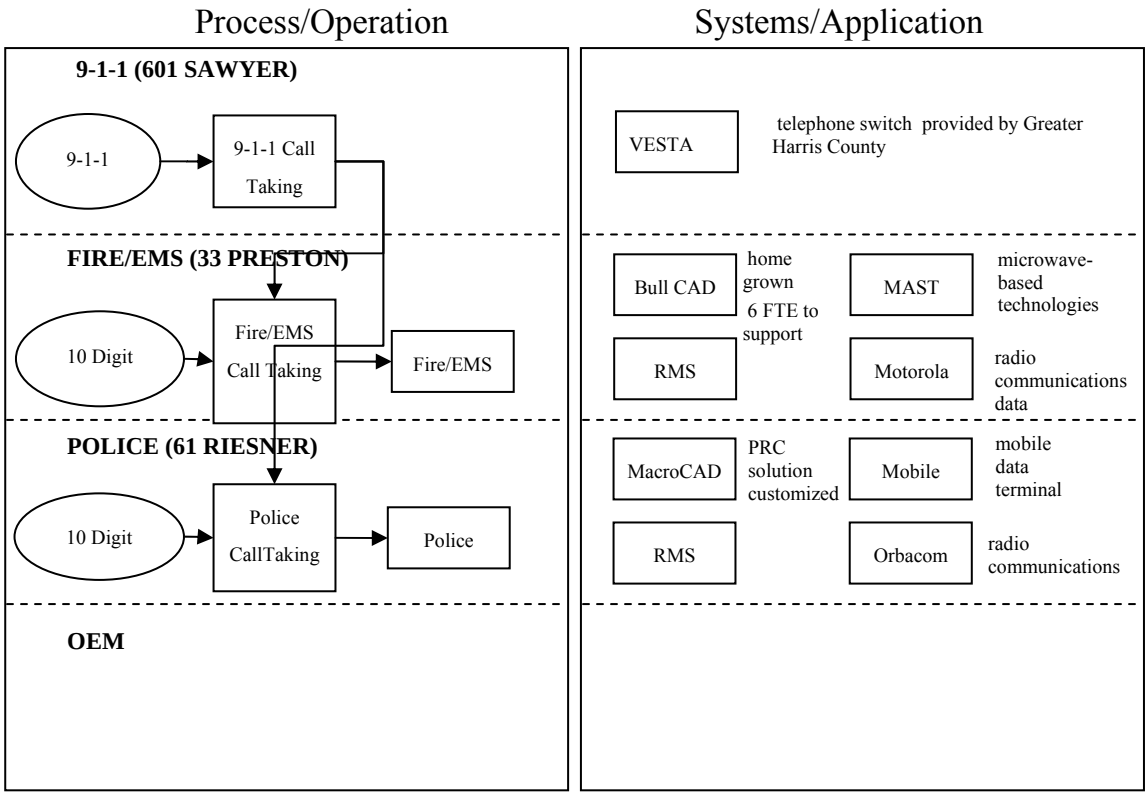

Figure 3-1. Old Public Safety System Architecture

3.2 Current System

The current City of Houston public safety system is comprised of different organizations' systems. The MITRE analysis focused on the architecture and design as it currently exists. During the assessment, the team noted that the HEC and other departments are in the process of making improvements or have identified upgrades that are needed but these planned changes are

not included in this assessment. Figure 3-2 highlights the current end-to-end systems that contribute or interface to the public safety system.

Figure 3-2. [redacted]

The key systems comprising the City of Houston Public Safety System are identified in Tables 3-1, 3-2, and 3-3.

Table 3-1. [redacted]

[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]

While the CAD system is the core system used by the call takers and dispatchers, there are many other systems (applications) required to be available and functioning in an optimal manner in order to efficiently and effectively enable communications between the call takers, dispatchers, and HPD/HFD response personnel when responding to 911 emergency incidents.

Other key components comprising the City of Houston's public safety system are shown in Table 3-2:

Table 3-2. [redacted]

[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]

The above list highlights those device components also required to be operational in order for all data to be successfully transmitted to and from the CAD system. In the event that any one of the above device components is not properly functioning, back-up procedures are activated in order for HPD/HFD emergency response personnel to continue responding to 911 emergency incidents.

[redacted] Table 3-3:

Table 3-3. [redacted]

[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]

The Tables depict the various groups involved in the end-to-end delivery of the HEC IT portfolio. Figure 3-1 illustrates the complexity of the operations and support for the public safety system. Root cause analysis of a perceived system problem may require multiple organizations to become involved in order to validate and verify that their particular scope of supported component is not the root cause of the problem or issue being experienced. Finally, outages to the system may be prolonged due to differences in service levels from the various groups identified below. While some groups and organizations provide 7x24 support for their components, other groups are only responsible for delivering support during regular business hours Monday through Friday.

As is the case with the device components and subsystems, public safety system performance may become degraded or unavailable to all or portions of the public safety system users when any of the above systems are not properly functioning. While back-up/contingency processes and procedures are instantly activated in order to eliminate disruptions to 911 emergency operations, unavailability of any of these components may have a performance impact to the CAD system.

The architecture was evaluated for possible single points of failure. This analysis primarily focused on the CAD, RMS, MSS, SAN, and network and did not address the possible failure of the radio, EAS, or PBX systems. For purposes of this assessment, single point failure analysis explored the impact of the majority of users being denied access to computer resources necessary to carrying out their mission. Based on the design of the HEC infrastructure, key systems such as CAD and RMS have been redundantly maintained [redacted] In addition, alternative systems to support voice via radio interface can function without the need for CAD or RMS and still permit the dispatch of resources (while somewhat less efficient than if CAD were available). Therefore, if CAD were to shut down, the ability to use Orbacom in conjunction with the call taking front end still permits the emergency dispatch function to occur. The following is a list of potential single points of failure that should be further studied and remedied:

- [redacted]
- Data Interfacing to MDT and RMS-HPD system from/to CAD.
- SAN Architecture.
- Integrated Database – Integrating CAD and RMS-Fire.
- [redacted]

3.2.1 [redacted]

3.2.2 RMS-HPD System From/To CAD

The next single point of failure addresses loss of services within 61 Riesner even when network services are available. While the CAD at HEC has a redundant design for both Fire/EMS and HPD, HEC only supports a redundantly configured RMS for Fire/EMS. A corresponding RMS-HPD exists at 61 Riesner that is not redundantly configured – implying a server failure (RMS-HPD) will result in down time with no failover option (as exists for the RMS-Fire/EMS at HEC).

The team was concerned that a past CAD outage resulted in problems at RMS-HPD [redacted] causing CAD to "hang" due to large queue backlogs that were unanswered. This problem has been resolved through upgrades to memory and processor functions.

3.2.3 Storage Area Network Architecture

The primary disk storage supporting the [redacted] was designed to support a high degree of redundancy from a RAID design as well as provide mirror imaging to another system (connected via a SAN) that could be used for recovery in the event of catastrophic failure of the primary RAID system. However, the implementation of the storage environment does not provide for real-time dynamic failover (it was neither specified in any requirements nor tested during acceptance). The intent of the components that are used in the design is primarily for disaster recovery. [redacted] In the event the primary RAID fails, the secondary system (which is basically a mirror image of the primary) is "resurrected" and becomes the primary system (until manually switched back to the backup role). This resurrection process is manual and can take many minutes depending on the disk configuration. [redacted] Network storage systems do exist that can provide real time dynamic failover but the current implementation does not provide that level of support.

3.2.4 Integrated Database – Integrating CAD and RMS-Fire

The database failover capability is primarily oriented at ensuring that when the primary system fails the system fails over to the backup system. This capability supports redundancy but does not adequately address other critical areas that need to be planned for to prevent database failures from causing outages. As noted in the review of the outages, one or more of the longer outages related to human error associated with the CAD/RMS databases. While disk mirroring can protect against catastrophic equipment failure, human error that erases the database is instantly "mirrored" to the backup disks. At that point, reconstruction of some type will be necessary to have an operational system. The best protection against this type of failure is to prevent it through better process management (configuration control practices). Other database improved performance measures can include:

- Database management tools exist that can create checkpoint rollbacks for certain types of transactions.
- Database replication with properly configured delays to remedy human error.

- Other database tools that can analyze the impact of a change prior to that change. Further analysis would need to be conducted to determine the best tools and techniques to mitigate this type of potential failure.

[redacted]

The team also noted that the current database version used in the system is Oracle 8. Oracle 8 is not fully supported by Oracle Corporation. Several security vulnerabilities have been identified with Oracle 8 and patches are not being provided to remedy these vulnerabilities. Oracle 9i can improve the reliability, security, maintainability, and performance of the system. For example, Oracle Real Application Cluster (RAC) is provided with Oracle 9i and it has distinct advantages over Oracle Parallel Server (OPS) used by Oracle 8.

RAC, introduced with Oracle 9i, is an advanced version of OPS with many additional self-tuning, management, and data warehousing features.

Oracle 9i introduced many new features to help the database administrator such as the ability to change database configuration "on the fly," enhanced availability, automatic performance and configuration tuning, and enhanced manageability. Given the nature of how past database administrator activities have led to system outages, the additional functionality reduces the risk of making a wrong decision that impacts the overall operations.

3.3 Systems Operations and Support

The public safety operations includes call takers from HEC, and dispatchers from the Houston Fire Department and the Houston Police Department. The primary roles are:

- Neutral 911 Call Taker
- HEC Call Taker (Fire/EMS)
- HFD Dispatcher
- HEC Call Taker (Police)
- HPD Dispatcher
- Supervisor HEC Call takers (Fire/EMS)
- Supervisor HEC Call takers (Police)
- Supervisor HPD Dispatch
- Supervisor HFD Dispatch

Appendix B provides a description of the call takers and dispatchers functions and how they operate the system.

4 Performance Analysis

This section analyzes the performance of the public safety data system. The analysis includes investigation of past records, interview results, and recently captured data traffic to determine whether the system performs in accordance with the scope of services and if it meets normal service level requirements. The analysis also assesses on other factors which provide past and current measurements of success performance. Specifically, this section discusses the following:

- Analysis of system outages that occurred from the period September 2003 through December 2004.
- Estimates of operational and inherent availability.
- Workload statistics.
- Scope of services performance analysis.
- Reliability assessment.
- Network configuration analysis.
- System performance monitoring.

4.1 Analysis of Outages and Errors

The scope of incidents considered in this report is based on the "HEC Outage Status" Excel spreadsheet and the "CAD System Availability From September 23, 2003, through December 16, 2004." The outages documented in the spreadsheet caused system-wide downtimes. Downtime is defined as a period of time when the system was unavailable to the call takers and dispatchers).[3] In all but two outage incidents,[4] the system was completely at the down state, and all users had to use some other means to get their jobs done. Isolated problems are identified in another report called the Software Incident Report Tracking (SIRT) and are analyzed separately.

The total assessment period for the HEC availability covers from September 23, 2003, 04:00, when the live operation of the upgraded system commenced, till January 31, 2005, 23:59. The upgraded system was accepted on January 2, 2004. In Section 4.2, two sets of availability calculations are provided, one for the total assessment period starting from the live operation commencement date, and the other for the shorter assessment period starting from the acceptance date.

---

[3] See Section J of Scope of Services: *CAD & RMS Acceptance Test Plans*, Page 11.

[4] For two outages CAD was able to operate partially. But, during these two incidents, either new logons could not be established or new emergency events could not be recorded. For incident # B10, all systems eventually had to be shut down.

Seventeen outages have occurred since the system went live. Table 4-1 shows a short summary of these outages. Ten of them occurred before system acceptance in a period less than 3 ½ months, and are labeled B1 through B10. After the acceptance, the frequency of outages has been significantly reduced, with only seven outages occurring over a period of almost 12 months, but their recovery times were generally longer. These outages are labeled A1 through A7. Each downtime period of an outage consisted of corrective downtime, preventive downtime, and/or delay time (for lack of logistic or administrative support). The last two outages were scheduled repairs and hence considered as preventive downtimes.

Table 4-1. HEC System Outages

Incident #	Date	Start Time	Total Time (hours)	Corrective downtime (hours)	Preventive downtime (hours)	Delay time (hours)	Problem	Cause	Device (Location)
(System went live and the acceptance test period started on 23 Sep 2003. This period had 10 outages: B1 – B10.)									
B1	9/24/2003	?	0.23	0.23			Incompatible software upgrade (for MDC sign-on) and human error	The down time was caused by an attempt by Northrop Grumman to install a software upgrade relating to MDC unit sign-on	CAD
B2	9/30/2003	16:00	0.08	0.08			Software bug		CAD
B3	10/2/2003	21:50	0.28	0.28			External interface and human error	Root cause was a network problem at 61 Riesner. The network problem was diagnosed to	CAD and SNA gateway @ Riesner

Incident #	Date	Start Time	Total Time (hours)	Corrective downtime (hours)	Preventive downtime (hours)	Delay time (hours)	Problem	Cause	Device (Location)
								be caused by the backup SNA gateway computer	
B4	10/8/2003	21:58	0.45	0.45			Software bug (system deadlock)		CAD
B5	11/5/2003	12:15	0.25	0.25			Software and client-server communication	Root cause was a network problem at 61 Riesner	CAD and workstation @ Riesner
B6	11/7/2003	17:21	0.12	0.12			Software and client-server communication	Root cause was a network problem at 61 Riesner	CAD and workstation @ Riesner
B7	11/10/2003		0.62	0.62			Software bug (archive logging)		CAD
B8	11/16/2003	22:08	0.25	0.25			Hardware failure (memory module)		CAD
B9	11/28/2003	8:30	4.38	4.38			Software bug (database lock) and procedure error		RMS and CAD
B10	12/3/2003	14:05	0.98	0.98			Hardware failure (RMS memory module) and software bug (database lock		RMS and CAD

Incident #	Date	Start Time	Total Time (hours)	Corrective downtime (hours)	Preventive downtime (hours)	Delay time (hours)	Problem	Cause	Device (Location)
(System was accepted on 1/2/2004)									
A1	4/10/2004	0:30	3.18	2.18		1.00	Database configuration mistake and human error (Northrop Grumman DBA)	Problems with expansion of the data table	CAD
A2	4/25/2004	16:26	0.90	0.90			Software bug (memory leak)		CAD
A3	5/10/2004	15:10	12.00	12.00			Human error (system admin to backup database)	Programmer issued a command at the Operating System (UNIX TRU64) level that caused the problem	CAD
A4	8/8/2004	12:10	5.00	5.00			Hardware failure (SAN disk array controller) and human error (HP)	Error on the HP technician part on loading a previous version of the firmware	SAN
A5	12/1/2004	7:30	8.00	8.00			Hardware failure (SAN disk array controller) and human error (HP)	Bad disk controller	SAN

Incident #	Date	Start Time	Total Time (hours)	Corrective downtime (hours)	Preventive downtime (hours)	Delay time (hours)	Problem	Cause	Device (Location)
A6	12/7/2004	?	2.75		2.75		Hardware replacement (SAN disk array controller)		SAN
A7	12/14/2004	?	2.42		2.42		Hardware replacement (SAN CPU and cache modules)		SAN

The counts of incidents by problem type are listed in Table 4-2. Table 4-2 shows that some incidents had multiple types of problems.

Table 4-2. Counts of Incidents by Problem Type

Problem Type	Count
Software	9
Human or procedure	7
Hardware	6
Interfaces (workstation-server communications or networking)	3
Database configuration	1

A more detailed classification of incident types can be found in Table 4-3.

Table 4-3. Classification of Incident Types

Incident #	Downtime (Hour)	CAD				RMS (H/W)	SAN (H/W)	SNA Gateway (H/W)	Workstation-Server communication	Human or procedure error
		Software	Hardware	DB Config	Admin					
(System went live and the acceptance test period started on 23 Sep 2003.)										
B1	0.23	x								Northrop Grumman
B2	0.08	x								
B3	0.28							x		Northrop Grumman
B4	0.45	x								
B5	0.25	x							x	
B6	0.12	x							x	
B7	0.62	x								
B8	0.25		x							
B9	4.38	x				x				
B10	0.98	x				x				
(System was accepted on 1/2/2004)										
A1	3.18			x						Northrop Grumman
A2	0.90	x								
A3	12.00				x					Northrop Grumman
A4	5.00						x			HP

Incident #	Downtime (Hour)	CAD				RMS (H/W)	SAN (H/W)	SNA Gateway (H/W)	Workstation-Server communication	Human or procedure error
		Software	Hardware	DB Config	Admin					
A5	8.00						x			HP
A6	2.75						x			
A7	2.42						x			

Incident B1 (9/24/2003). The outage was caused by an incompatible software upgrade and is not likely to occur again if configuration requirements are carefully processed.

Incident B2 (9/30/2003). The outage was a software bug in an analysis program that is not critical to call processing and dispatching. Portions of the program were temporarily disabled and are not likely to cause future outage.

Incident B3 (10/2/2003). The outage originated from a bad network card at the backup SNA gateway. This is regarded as a single point of failure. Unless fault isolation is considered, whether in the architecture or at the application level, this kind of outage may happen again.

Incident B4 (10/8/2003). The outage was caused by a system deadlock for database transactions. This was fixed with a code change.

Incidents B5 (11/5/2003) and B6 (11/7/2003) are the same kind of outage. CAD had more than 800,000 TCP packets pending transmission/retransmission from CAD to a remote workstation at 61 Riesner. This large amount of communications backlog caused CAD to go down. The resolution was to limit the amount of data that could be requested at one time from each workstation. Users needing large amounts of data would have to do queries outside of CAD; e.g., using SQL on database server. This problem should not occur again, but the root cause of CAD ability to operate when large communications backlog happens may still be a problem. A better understanding of capacity limits will help develop fault detection and performance monitoring capabilities.

Incident B7 (11/10/2003). The outage was caused by an archive logging process error. This problem should not occur again if the correct procedures are followed.

Incident B8 (11/16/2003). This was the only CAD hardware (memory module) failure. Reoccurrence is dependent on the hardware reliability.

Incidents B9 (11/28/2003) and B10 (12/3/2003). Both outages had the same symptom: incomplete transactions between RMS and CAD or the failure of RMS to report completed transactions caused the integrated database locked. Manual unlock was done by support

contractors. Transaction process functions were examined and reengineered by Northrop Grumman in conjunction with Oracle. The root cause, bad memory modules in RMS, was identified, and all memory modules in RMS were replaced by HP. Reoccurrence of the problem is dependent on the hardware reliability.

Incident A1 (4/10/2004). The outage was caused by insufficiently allocated space in database, and it was compounded by an inexperienced DBA on site. Database space was expanded and a more experienced DBA is on site. This kind of problem is unlikely to happen again.

Incident A2 (4/25/2004). The outage was caused by a software bug (memory leak) in the CAD application. This problem was fixed.

Incident A3 (5/10/2004). Improper system administration (database backup) caused system outage for 12 hours. The contract system administrator has been replaced. This kind of problem is unlikely to happen again.

Incidents A4 (8/8/2004) and A5 (12/1/2004) were both SAN hardware problems, causing downtime 5 and 8 hours, respectively. This signified single-point-of-failure in the system architecture.

The last two outages on 12/7/2004 and 12/14/2004 were both preventive maintenance.

Table 4-3 also shows that the primary CAD (CADB), as the central component interfacing many devices, was vulnerable and thus its unavailability status caused some of the outages. Some incidents did not start from CAD directly, but they still caused CAD to also be unavailable. The system design should isolate CAD from being impacted by failures in other systems.

In addition to the outages, two other categories of problem resolutions were identified. This included problems identified as minor and new requirements. These 61 problems were documented in an SIRT (Software Incident Report Tracking) list and a change order list covering the period from October 31, 2003 to December 17, 2004. The SIRT list provides a description of each problem, estimated completion date, and resolution status. None of these problems were serious enough to cause system downtime on the primary system. Most of them require only system patches, documentation, or demonstration, while a few may need additional design. A summary of SIRT problem types is shown in Table 4-4.

4.2 System Availability Calculations

The scope of service covering the Northrop Grumman agreement specifies a requirement of 99.9% system availability for the CAD and RMS systems. Hardware failures are excluded from Northrop Grumman's availability calculations. MITRE recommends that all major systems meet or exceed 99.99 system availability. MITRE independently assessed the system availability based on universally accepted definition.

Table 4-4. Summary of SIRT Problems

Problem Type	Count
Data entry/recording/display	24
Communication or data transmission	10
Address/location verification	5
Application error	5
Database configuration or management	3
System startup or switchover	3
Data edit check	2
GUI bug	2
Additional data required	2
Documentation	2
Data error	1
Erroneous messages	1
OS update	1

System availability is defined as a system (consisting of hardware and software) is operating at any point in time, when subject to a sequence of "up" and "down" cycles. It addresses the question of "How likely will the system be available in a working condition when it is needed?" In this analysis, availability was evaluated by two standard measurements, operational availability and inherent availability. The availability of the overall system will be discussed first, followed by the computation for CAD, RMS, and SANs. There are two sets of availability calculations based on two alternative views of the starting point of the system life cycle: (1) starting from the system go-live date September 23, 2003; (2) starting from January 3, 2004. All availability calculation results are summarized in Table 4-5. The upper limits of availability for the 95% confidence level are shown in Table 4-6. The purpose is to provide an objective basis for setting reasonable expectations of the system availability. The percentages of uptime for individual months and days are presented in Tables 4-7 and 4-8. Some relevant concepts and definitions can be found in Appendix C.

4.2.1 Availability of the Overall System

The operational availability[5] of the overall system starting from go-live is:

[5] This is similar to the availability defined in Section J of Scope of Services: *CAD & RMS Acceptance Test Plans*, Page 9. But the downtime considered in this report is plain and general: whenever the system is not operational, caused by either hardware or software failure, users are experiencing downtime.

A_o = Total Uptime / Assessment Period = 1- Total Downtime / Assessment Period = 0.9965

The total downtime includes all corrective repair times, preventive maintenance times, and delay times caused by administrative and logistics processes.

The inherent availability of the overall system is:

A_i = MTBF / (MTBF + MTTR) = 0.9970, where MTBF is Mean-Time-Between-Failure and MTTR is Mean-Time-To-Repair.

Also known as Intrinsic Availability, the Inherent Availability A_i does not consider delay times and preventive maintenance times.

4.2.2 Availability of CAD/RMS

The CAD/RMS availability is derived from incidents caused by problems with CAD/RMS. Outages caused by other components of the system (e.g., SAN failures) are not included.

The calculation for the operational availability of CAD/RMS includes all outages except the last four that were caused by SAN problems. The operational availability of CAD/RMS since go-live is A_o = 0.9980

In computing the inherent availability of CAD/RMS, incident A3 is not included, since it was initiated by a system administration error that subsequently caused CAD to go down. The inherent availability of CAD/RMS since go-live is A_i = 0.9991

4.2.3 Availability Since Acceptance

If the system life cycle is considered to start from the day after the system acceptance date, as opposed to the system go-live date, then the start time of the assessment period is shifted to January 3, 2004, and the first 10 items in Table 4.1 are not counted against the availability calculation.

The operational availability of the overall system since acceptance is A_o = 0.9964

The operational availability after the acceptance is slightly worse than the operational availability previously calculated for the entire period since the go-live date. The analysis of the outages prior to the acceptance show that even though they were more frequent but were also much shorter (less than an hour), than those that occurred after the acceptance period. One explanation for the apparent difference in the recovery time is that both the system developer and the technical staff might have been more expeditious for problem resolution during the Acceptance Testing phase.

After the system was accepted, the system formally moved from testing to maintenance. The maintenance and support might be less agile than in the testing period. The records show certain degree of failure to meet contingency, which was also compounded by the deficiency in the skill set of the contractors. As a matter of fact, the majority of failures after the acceptance were either caused directly or aggravated by human errors. Based on the interviews, the MITRE team

believes the maintenance team is now more experienced. It is reasonable to believe that the worst time is over; it is also a fair expectation to see reduced downtime in future outages.

There were only five outages after the acceptance. (A6 and A7 are outages for preventive maintenance.) The inherent availability of the overall system since acceptance is $A_i = 0.9970$

It is a coincidence that the inherent availability of the overall system after the acceptance has exactly the same four digits as the inherent availability before the acceptance.

Next, to calculate the availability for CAD/RMS, after the acceptance date, the three incidents (A1 to A3) are used to determine the operational availability calculation and the two incidents (A1 to A2) are used for the inherent availability calculation. Therefore, the estimated availability values for CAD/RMS after acceptance are $A_o = 0.9983$ and $A_i = 0.9997$

The CAD/RMS availability after acceptance has improved largely because all of the outages except one before the acceptance involved CAD/RMS, whereas after the acceptance, only less than half were related to CAD/RMS.

4.2.4 Result Summary of System Availability

A summary of all availability numbers computed in Sections 4.2.1 through 4.2.3 is presented in Table 4-5. As mentioned earlier, these statistical estimates are meant to provide a forward-looking view of the likelihood that the system will be available at any point in time. The percentages of uptime for individual months and days are shown in Tables 4-6 and 4-7.

Table 4-5. Results of System Availability With Different Assessment Periods

	Suppose the system life cycle started from the go-live date (23 Sep 2003)	
	Operational Availability	Inherent Availability
Overall system	0.9965	0.9970
CAD/RMS	0.9980	0.9991
	Suppose the system life cycle started after the acceptance date (3 Jan 2004)	
	Operational Availability	Inherent Availability
Overall system	0.9964	0.9970
CAD/RMS	0.9983	0.9997

These estimates indicate that inherently the CAD/RMS system looks promising for keeping up with the required 99.9% availability level, while the overall system may not achieve the same level of performance. Other parts of the overall system other than CAD/RMS have negatively impacted the overall availability. In order to provide uninterrupted services to end-users, a highly-available CAD/RMS system by itself is not enough, since the past incidents have already shown that it is susceptible to failures of other parts. Thus, it is recommended that efforts be focused on raising the availability of other parts of the overall system, in particular the SANs, and in making CAD/RMS more resilient to failures passed from these interfaces.

4.2.5 Confidence Level and Confidence Limit for Availability Estimates

Various expectations or industry norms for system availability may exist. This analysis calculates availability based entirely on empirical data associated with true events. Furthermore, standard statistical methods can also use the same set of empirical data to calculate the upper limit for system availability given a desirable confidence level.

Details of confidence limit calculation are provided in Appendix C. The results for a 95% confidence level are shown in Table 4-6, where A_o and A_i denote Operational Availability and Inherent Availability, respectively.

Table 4-6. Confidence Limits of Availability With 95% Confidence Level

	Assuming the system life cycle started from the go-live date (23 Sep 2003), the 95% confidence limits are:	
Overall system	$A_o < 0.9983$	$A_i < 0.9984$
CAD/RMS	$A_o < 0.9995$	$A_i < 0.9995$
	Assuming the system life cycle started after the acceptance date (3 Jan 2004), the 95% confidence limits are:	
Overall system	$A_o < 0.9984$	$A_i < 0.9990$
CAD/RMS	$A_o < 0.9998$	$A_i < 0.9999$

The entry "Overall system $A_o < 0.9983$" means the following: If no major improvement is to be made, we can predict with 95% confidence the operational availability of the HEC overall system will be less than 0.9983. That means the overall HEC system will suffer at least 14.8

hours of total operational downtime (both planned and unplanned) per year. Other entries have similar meaning.

These confidence limits indicate that it would not be realistic to expect the availability of the overall system to reach the 0.999 level. The CAD/RMS should achieve higher availability calculations but will probably not reach the recommended 99.99.

4.2.6 Monthly and Daily Availability

The concept of monthly and daily availability has been used by some organizations for checking against service level agreement. It no longer serves as an indication of the probability that the system is in a working condition but reports the percentage of system uptime during a calendar month/day. Dividing the continuous system operation into months and days will inevitably change the calculated results[6].

All monthly operational availability numbers after the acceptance date (3 January 2004) are shown in Table 4-7.

Table 4-7. Monthly Availability (Percentage of Uptime) After Acceptance

Month	Overall System	CAD/RMS	SAN
2004-Jan	100	100	100
2004-Feb	100	100	100
2004-Mar	100	100	100
2004-Apr	99.43	99.43	100
2004-May	98.39	98.39	100
2004-Jun	100	100	100
2004-Jul	100	100	100
2004-Aug	99.33	100	99.33
2004-Sep	100	100	100
2004-Oct	100	100	100
2004-Nov	100	100	100
2004-Dec	98.23	100	98.23
2005-Jan	100	100	100

[6] This is similar to the availability defined in Section J of Scope of Services: CAD & RMS Acceptance Test Plans, Plan 9. But the downtime considered in this report is plain and general: whenever the system is not operational, caused by either hardware or software failure, users are experiencing downtime.

After acceptance, all daily operational availability numbers are 100%, except for the following days in Table 4-8:

Table 4-8. Daily Availability (Percentage of Uptime) After Acceptance

Day	Overall System	CAD/RMS	SAN
2004-04-10	86.74	86.74	100
2004-04-25	96.25	96.25	100
2004-05-10	50.00	50.00	100
2004-08-08	79.17	100	79.17
2004-12-01	66.67	100	66.67
2004-12-07	88.54	100	88.54
2004-12-14	89.93	100	89.93

4.2.7 System Availability Enhancement

MITRE assessed methods to improve system availability given the less than desirable availability results of the overall system. In addition, while the CAD/RMS is close to meeting the 99.9 availability requirements, MITRE recommends an availability of 99.99. Thus, two methods to increase system availability were assessed. They included:

- Increasing reliability by acquiring more reliable components and also make service delivery more reliable. This method increases the MTBF.
- Increasing maintainability by performing repairs and maintenance work more efficiently and effectively. This method shortens MTTR.

In general, improving MTTR has better leverage than improving MTBF for increasing the system availability. The assumptions and formulas for this analysis are contained in Appendix C. Figure 4-1 shows the progressively estimated MTBF for the system calculated after each incident cycle. This chart indicates that the MTBF is getting better (longer) but is not yet reaching a steady state, thus, implying that the integrated HEC system has not passed the so-called "infant mortality" stage. As long as the MTBF continues to get longer, then the system reliability will continue to improve. A steady state will be achieved as the system matures.

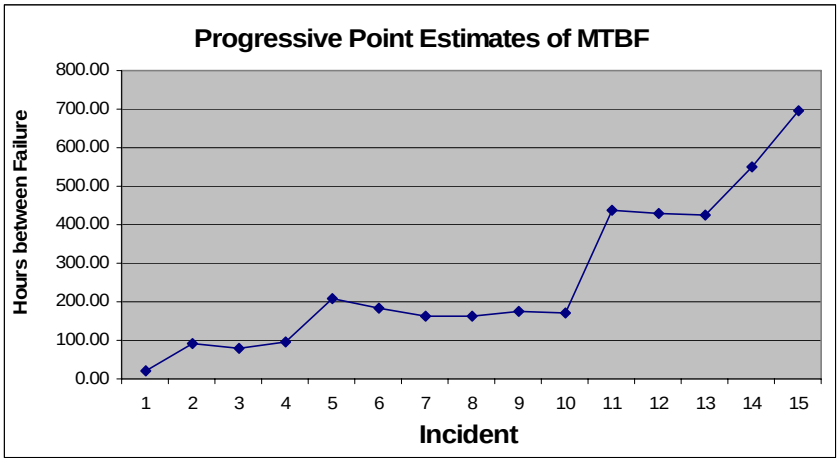

Figure 4-1. Progressive Point Estimates of MTBF

Based on the fact that the MTBF (or the failure rate) of the various subsystems has not reached a steady state, the frequency of the outages will generally be reduced even without major upgrades to the system. However, this statement is not true when assessing the system's MTTR trends. For the short run, the most urgent need is to improve the maintainability, such as having an adequate and better-trained support staff. A required strategy is to amend all identified weak points in the system and infrastructure, particularly, the single points of failure identified in Section 3. This will improve the overall system availability by:

- Reducing the vulnerability to system-wide failure and consequently increase the uptime.
- Providing automatic switchover to the backup and thus effectively reduce the downtime.

The HEC incident report shows that 6 out of 15 system failures since the system went live involved human or procedural errors, and they account for 78% (32.08 hours / 35.73 hours) of the total unexpected outage downtime. The outage history indicates that, not only did the maintenance staff from Northrop Grumman and HP make mistakes; they also seemed to complicate the problems further resulting in a very slow recovery of the system. Some problems leading to human errors may have been intrinsically difficult. Nevertheless, improvements in the skill level of the maintenance staff may help to increase the MTTR and, thus, improve the availability of the system.

4.3 Workload and Performance Assessment

The MITRE team observed the operations of the CAD system during several on-site visits and interviewed management staff, call takers, and call dispatchers to discuss system performance. During these discussions, concerns were raised about the performance of the system during busy time periods and when upgrades were made to the system. In order to attempt to determine whether the workload of the system impacted system and performance, call volume statistics and

data were gathered. This data identified call volume statistics handled by the various departments for the period of January 2004 to December 2004. A secondary analysis was to try to determine if the demand level pointed out a probable cause for some of the system outages that occurred at HEC.

System performance can be affected by the amount of demand using the system. There are three major types of demands that require system resources and may contribute to component wear out and cause degradation in system performance. These three types are:

(1) 911 (including 10-digit calls) call volume statistics.

(2) Call takers and dispatchers use of the CAD systems.

(3) Police and Fire/EMS units and stations that have to respond to the dispatching assignments and make information queries.

Among those three types of demands, only the call volume statistics has data available for each hour during the assessment period. A more useful demand data will be the staffing level records (how many call takers and dispatchers are connecting to the system at each hour), which corresponds to the second type of demands. A series of charts of call volume data are presented in Appendix D. As an example, Figure 4-2 shows the chart for November 2003 (prior to the acceptance). The purpose of the analysis is to determine whether there is any correlation between call volume and outage occurrences. Each chart covers a one month interval within the assessment period. The call volume value includes all calls for Fire, EMS, and Police events. Each data point is the call volume for the corresponding hour. Each triangle on the chart indicates the start time of one of the seventeen outages since the system went live.

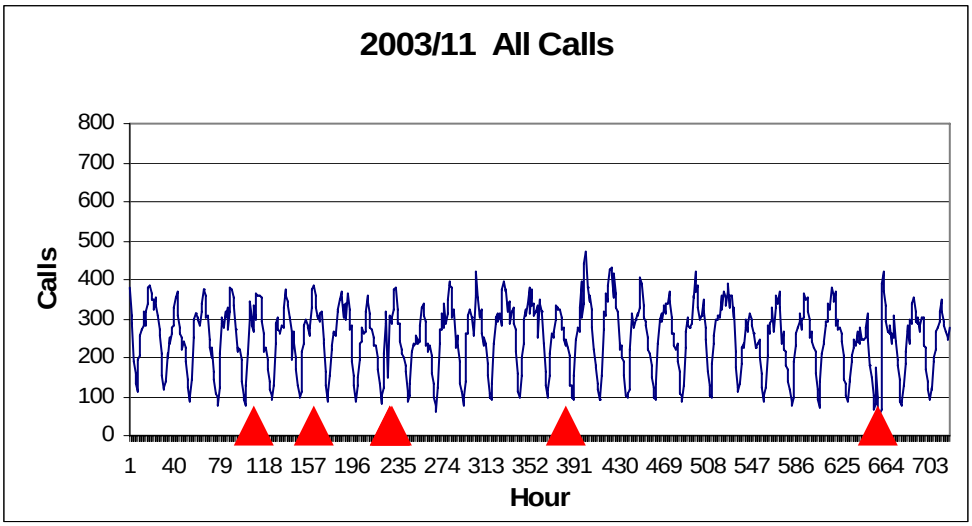

Figure 4-2. Call Volume Statistics and Outages

The analysis of the data from the charts concludes that none of the outages correlate with a spike or surge of the total 911 call volume. Using separated department-wise[7] call volume data, there is still no evidence of correlation between system outages and the volume of either Police calls or Fire/EMS calls. Although a careful analyst may sometimes look further into the fluctuation and variation of the raw data to extract hidden patterns, observations obtained from the plain call volume data deemed further analyses unnecessary. MITRE concludes that the call volume data by itself, does not show impact on the outages.

The MITRE team attempted to gather performance statistical data at various levels to make a clear determination or root cause analysis of system performance.

Figure 4-3 identifies our approach to the end-to-end performance analysis effort. By taking this approach, the team planned to gather performance data at each layer of the subsystems, and then correlate the data to make an accurate assessment of potential system performance issues.

The HP Systems Insight Manager was not adequate for the performance analysis because it was not completely configured. The team was able to use UNIX level command scripts to gather performance data on the CAD and RMS servers in order to conduct a performance analysis for these two subsystems. Figure 4-4 summarizes CPU utilization for a specific snapshot period.

[7]Individual charts for Police calls and Fire/EMS calls are not included in this report, but they have been inspected and led to the same conclusion.

Performance Analysis Scope	Performance Analysis Toolkit	Results
Application Layer (CAD, RMS)	• Transaction Response Tool	Tool turned off; Monthly data from Altaris Command Stats
Database Layer	• Oracle Enterprise Manager	Tool not available; Data from DB Data Status Reports
Hardware / Operating System Layer	• Unix Perf; Tools • HP System Insight Manager	New intall of Insight Mgr; Used Unix tools to snapshot perf.
Network Management Layer	• OPNET NetDoctor	Perf. Data captured and analyzed
Telecom Infrastructure	• Network Sniffer	Perf. Data captured and analyzed

Figure 4-3. Performance Analysis Layers

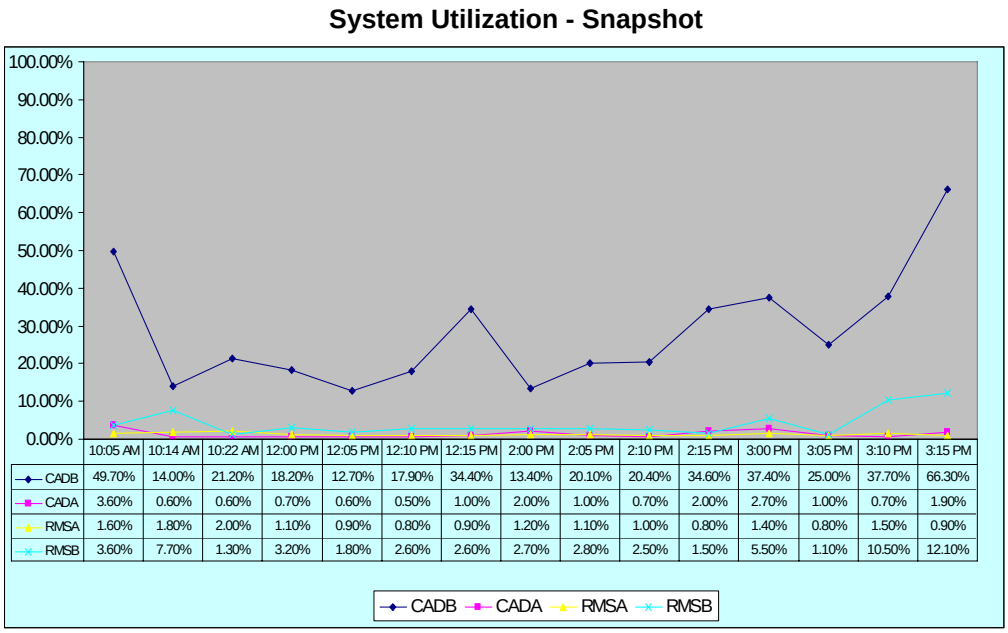

	10:05 AM	10:14 AM	10:22 AM	12:00 PM	12:05 PM	12:10 PM	12:15 PM	2:00 PM	2:05 PM	2:10 PM	2:15 PM	3:00 PM	3:05 PM	3:10 PM	3:15 PM
CADB	49.70%	14.00%	21.20%	18.20%	12.70%	17.90%	34.40%	13.40%	20.10%	20.40%	34.60%	37.40%	25.00%	37.70%	66.30%
CADA	3.60%	0.60%	0.60%	0.70%	0.60%	0.50%	1.00%	2.00%	1.00%	0.70%	2.00%	2.70%	1.00%	0.70%	1.90%
RMSA	1.60%	1.80%	2.00%	1.10%	0.90%	0.80%	0.90%	1.20%	1.10%	1.00%	0.80%	1.40%	0.80%	1.50%	0.90%
RMSB	3.60%	7.70%	1.30%	3.20%	1.80%	2.60%	2.60%	2.70%	2.80%	2.50%	1.50%	5.50%	1.10%	10.50%	12.10%

Figure 4-4. System Utilization – Snapshot

As evidenced in the above Figure, for this specific snapshot, the CADB systems' (primary production environment) CPU utilization increased to over 48% in the early morning (approximately 10:00 am), returned to an average range between 13%-34% for the mid-day period, then increased to over 66% utilization towards the latter part of the day (approximately 3:15 pm). Unfortunately, because no other performance data is available for either the Application Layer or the Database Layer during this period, we are unable to perform any root cause of the increases in the CPU utilization.

Northrop Grumman provided summarized data corresponding to the command response time statistics in task category "Check response time of CAD commands," see Section 4.6. For each of the 464 commands, the data shows the total number of invocations in each month and the number of elapsed seconds averaged over each month. Many commands were not used in every month. About 80% of commands were infrequently used. (Less than 7200 invocations per month. If there were averagely 10 CAD users per hour, this would imply that such a command was used averagely less than 1 time per user.) Table 4-9 compares the January 2005 measurement results with the CAD response time requirements for the acceptance test defined in Attachment 27 Section J (Acceptance Test Plan for Altaris CAD Implementation). Northrop Grumman provided measurement results for 18 out of 27 command types. For each of these 18 command types, the response time averaged over the entire month compared very well with the corresponding threshold value. However, the two sets of numbers do not represent the same time scale. The requirements defined for the acceptance testing were meant for evaluating a peak load hour of at

least 433 events (Page B-6 of Test Plan), while the Jan 2005 results were averaged out over the entire month of January 2005 and did not correspond to a particular hour with peak load condition. The comparison results are shown here only for indicating that, on a monthly average basis, the CAD response time average performance fared very well in January 2005.

MITRE recommends the entire performance analysis toolkit referenced above be activated and an extended performance data capturing period be established in order for the appropriate level of performance data correlation and root cause analysis be conducted.

The process for the creation, delivery, verification of the Geofile has a major impact on the performance of the public safety system. During interviews with HEC, Northrop Grumman, call takers and dispatchers, the Geofile accuracy was a major discussion item. The team was unable to determine if the concerns were technical or training issues. For example, it was noted during the interviews with call taker and dispatcher personnel that location information sometimes has to be keyed in different ways before the system recognized the street address and provided the proper system and agency recommendation. These concerns were also discussed with the Graphical Information System (GEO) Team lead. Based on the discussions, the MITRE team believes that these issues could be caused by a number of reasons, many of which cannot be solved through technology solutions. Many of the issues require improvements in processes and communications to be assessed. The process improvements include the use of written documentation to record and forward the problems to the Graphical Information System Team. Next, the recommendations in Section 5 for improved processes should also help to identify and resolve these issues.

4.4 Scope of Services and Performance

MITRE analyzed the performance of the system to determine where performance may not meet requirements in the Northrop Grumman scope of services. Three areas of noncompliance and possible noncompliance with the scope of services were identified.

First, the CAD and RMS did not meet the required monthly 99.9% availability. The Scope of Service requirement number 14.1.2 requires Northrop Grumman to maintain the required system availability of 99.9% for the CAD upgrade and RMS systems. The requirement states that the City shall record the system downtime on a monthly basis. If the system availability is not met, Northrop Grumman is required to submit a report that documents the event and to detail a plan of action to prevent a recurrence in the future. As identified in Section 4.2, the CAD system has not met the monthly requirements of 99.9% availability.

Second, the system performance and monitoring is not being provided. The Scope of Service requirement number 3.3.17 requests CAD reporting statistics, including transaction response times, be provided for any time/data range required. Northrop Grumman has several tools that are

Table 4-9. Measurement Results

Command Type		**Response Time Requirement**	**Jan 2005 Result** (monthly average seconds)
Call-Taker Commands			
Event Entry Form Call-up		1-second or less	-
Location Verification		1-second or less	.26
Access Geo Location Information		1-second or less	.07
Access Location Information		1-second or less	.30
Access Location History		1-second or less	.08
Add Event		3-seconds or less	.35
Dispatch Commands			
Display Event		1-second or less	.25
Unit Suggestion – by Geographic District		2-seconds or less	.26
Dispatch Single Unit		1-second or less	.22
Assist – Single Unit		1-second or less	.63
On-Scene – Single Unit		1-second or less	.20
Change Location – Single Unit		1-second or less	.11
Change Unit Status		1-second or less	.20
Clear Unit and Close Event		3-seconds or less	.15
Mapping Commands			
Center on Location		4-seconds or less	-
Center on Unit		4-seconds or less	-
Pan Left – Predefined Increment		4-seconds or less	-
Pan Right – Predefined Increment		4-seconds or less	-
Zoom-In – Predefined Increment		4-seconds or less	-
Zoom-Out – Predefined Increment		4-seconds or less	-
CAD Inquires – Local Database			
Event History – Single Event Number		2-seconds or less	.12
Unit History – Single Unit		2-seconds or less	.26
Unit Summary – Single Area		3-seconds or less	-
Recent Event History – Last 20 Events		2-seconds or less	.24
Event Query – by Key Field		2-seconds or less	-
Administrative Messages			
Send Message		2-seconds or less	.11
Retrieve Message		2-seconds or less	.17

gathering and collecting this data. However, this information has either not been properly requested or is not being properly provided to the HEC.

Third, the team was unable to determine whether the transaction performance requirements applied to the operational system and whether they were being met. The Scope of Service requirement number 15.10.3 identified requirements for the response time for a list of user command types. This section applies to the acceptance test period and the team could not determine if the requirements were applied after system acceptance. Northrop Grumman performed and successfully passed a modified version of these requirements in[8] the Response Time/Load Test in a simulated environment before the acceptance testing in a live operation environment. Since the system acceptance, there has not been any systematic reporting of such statistics, even though users have expressed concerns about perceived slow response times on occasions.

4.5 Reliability

4.5.1 End of Life

Figure 4-5 shows equipment that is near the end of its life that should be replaced. The systems that are close to or past end of life means that the systems are or will become obsolete within a year.

This figure incorporates all of the major public safety data and radio systems. While the majority of the assessment of data in this section focuses on the HEC systems, this chart emphasizes the importance of evaluating performance parameters for the whole system.

These systems can still function and meet existing operational requirements but will become difficult to maintain. Therefore, key critical components may be reaching their end of life (EOL) which increases the risk to the resiliency and reliability of the overall public safety system.

In addition to concerns with certain systems approaching their end of life, the analysis also showed that two major components used by these systems are becoming obsolete, Oracle database and network equipment.

As of December 31, 2004, Oracle discontinued its support of Oracle 8i products. Users can still get support for Oracle 8i products by signing a contract with Oracle. However, such contracts do not include any potential patches that may be needed to resolve software problems that are found. Currently, most of the user community is using Oracle 9i, with a few users moving to Oracle 10g, because of some functionality that they may need that is not

---

[8] In the PRC Response to the City of Houston Revised Scope of Services for Houston Public Safety Dispatch Consolidation, PRC identified the Acceptance Test Plan, Attachment 27, Section J as the list of the response values that would be the basis for test. Pages B-6 to B-8 of the document titled: Acceptance Test Plan for Altaris® CAD Implementation, August 2001, identify the actual operations and the corresponding test values.

found in Oracle 9i. Oracle 9i is quite stable at this time, and therefore it would be a good option to migrate to that version.

[redacted]

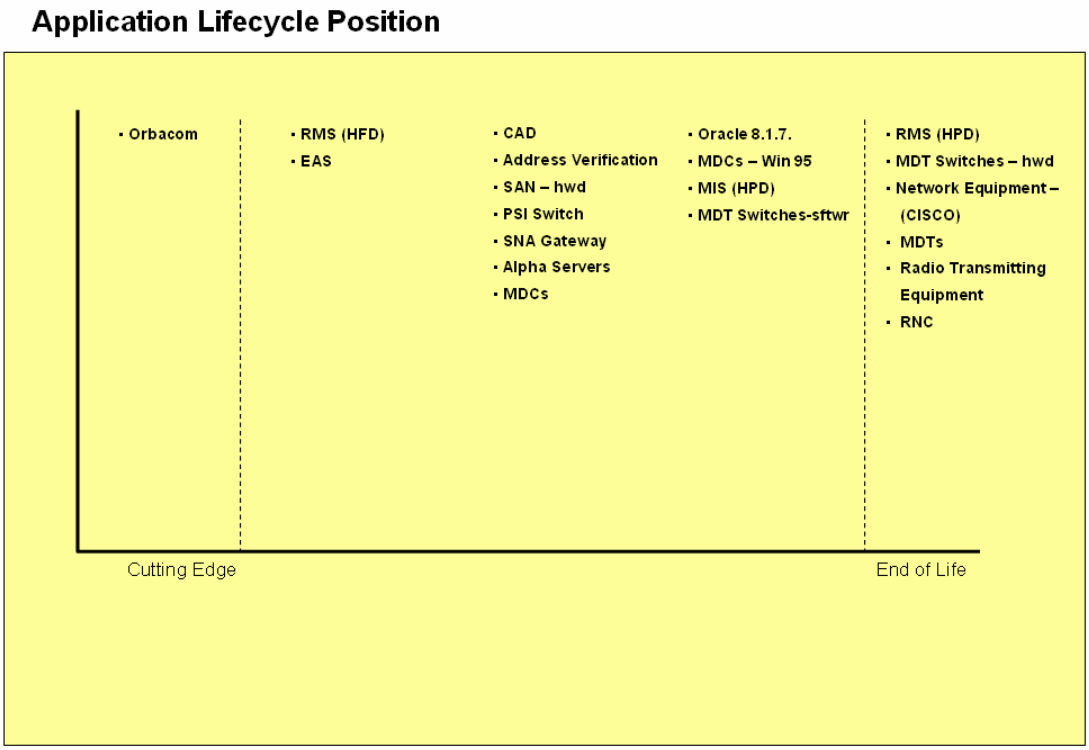

Figure 4-5. Public Safety System Life-Cycle

[redacted]

[redacted]

[redacted]

[redacted]

[redacted]

[redacted]

4.5.2 Network Configuration Analysis

MITRE conducted an analysis of the network router configuration to determine potential problems in the configuration of the network. OPNET's NetDoctor (Version 11.0) tool was used to examine the suite of routers and switches used at the HEC. The results of this analysis are provided in the report titled "Net Doctor Report", February 4, 2005. The network topology is shown in Figure 4-6. The nodes in the figure represent workstations and servers that are attached to the network devices and are used to establish traffic demand for routing analysis. While real traffic data was collected, that data was not used for this simulation since the primary goal of this

assessment was the router configuration. There were a total of 426 different tests run on both the individual (IOS) routers and (CatOS) switches as well as the total network as a whole (the latter involving simulating traffic on the network to determine the viability of the routing configuration between routers and switches).

[redacted]

Figure 4-6. [redacted]

Of the 426 test that were run, no major errors were detected [redacted]

[redacted]

4.6 System Performance Monitoring

The team obtained information from interviewed employees of HEC, ITD, HFD, HPD, Greater Harris County 9-1-1 Emergency Network, and Northrop Grumman to discuss how the performance of the systems were monitored. [redacted] The team primarily focused on the monitoring capability for the HEC portions of the public safety system and the ITD networks. A similar analysis needs to occur for the HPD and HFD systems.

During interviews with ITD, MITRE was informed that the ITD may set up a Network Operations Center (NOC) to provide centralized network management functions for several organizations including HEC. The idea is still at early inception stage and the scope focuses only on detecting network outage alarms and providing centralized problem resolution in a timely manner. The NOC will also report uptime and downtime statistics. In other words, the function of the planned NOC is mainly to react to incidents when they occur. [redacted]

The Scope of Service identifies requirements for system performance monitoring. During interviews with HEC IT staff and Northrop Grumman, MITRE attempted to gain an understanding of what tools were in place and how they were used for monitoring and reporting system performance.

The interviews showed that the system is not being adequately monitored and reported. Basically, the HEC IT staff and Northrop Grumman have a process for system monitoring and reporting on an as needed basis. This process is put in place whenever an outage or significant lack of performance occurs. The process that is in place during normal operations is not clear.

The following paragraphs provide a brief description of possible network monitoring tools that could be used. Some of these capabilities exist in Northrop Grumman's existing configuration but are proprietary. The City of Houston should determine if these proprietary tools could provide the monitoring capabilities needed or if commercial tools are needed.

Performance monitoring and reporting tools fall into three broad categories that are applicable to the system:

- Client monitoring
- Network and server monitoring
- Application-level monitoring

The client monitoring tools gathers metrics about the end-user experience, such as response time for specific interactions in the application. It may be useful for measuring CAD Transaction Response, as defined in Section 15.10.3 of the Scope of Services. [redacted] However, these applications are data-intensive and should only be used as a tool to occasionally gauge the system load to aid in the decision on allocating resources.

Network and server monitoring tools monitor the performance of system infrastructure, connection status, and assist in error detection. They usually use SNMP and RMON agents with real-time event filtering for fault alerting and problem resolution. For non-SNMP equipment, a protocol mediation solution or a proxy agent can mediate standard alarm outputs from various types of equipment to SNMP. They can collect statistics and report throughput, uptime, data link utilization, CPU usage, packet loss, packet latency, etc. Some platforms may also be able to check on connections involved with any given application and provide information about the host server as well. Basic SNMP statistics collection, storage, exception reporting ("Top-N" lists, etc.) and historical trend graphing are built in to most of the major network monitoring platforms, and there are a number of commercial products focused specifically on performance. Concord Communications is one of those with the broadest coverage and largest customer base. HP OpenView series of solutions also provide comprehensive monitoring capabilities. There are also open source tools (MRTG is widely used). HEC is currently considering adopting a network performance monitoring tool.

It must be noted that generic tools used for monitoring network performance usually are not capable of detecting application-level problems. A better approach for application-level monitoring is application instrumentation, which involves writing specific code within an application to check key transaction performance indicators, such as message queue length, waiting time, and completeness of transaction. It may also report other measurements; e.g., response time, database connectivity, system load, etc. An application-monitoring tool may help avoid some of the incidents;, e.g., B4, B5, B6, B9, B10, and A2 from happening again.

A caveat application instrumentation is an invasive method that requires modifying the original application and may be too resource-intensive. The Application Response Measurement (ARM) and Application Instrumentation and Control (AIC) technical standards and APIs have been published by an industry consortium for some years and adopted by a number of leading providers of performance monitoring tools. The system managers can monitor transactions by using simple function calls embedded in the application code. An agent captures these calls and sends them to an ARM or AIC reporting application, such as the IBM Tivoli Management Environment platform. This popular platform is by far one of the best solutions in the industry to enable end-to-end management of all elements in a multi-vendor environment, from the network, to computers, to applications and databases, and to business management of IT services. Northrop Grumman has stated that the ARM tech suite is installed by default. The tool can monitor CPU performance by user and additional functions are available through all-on licenses.

Computer Associates' Unicenter monitoring environment also has a built-in functionality that can provide application-level monitoring. The AIC standard was based on Unicenter's TNG management software. The Insight Integration for CA Unicenter can integrate HP hardware monitoring and event notification functions into Unicenter. Together they provide an integrated platform for managing and monitoring systems and business applications. As stated at the HP web site:

> "This comprehensive, scalable solution builds upon the core elements of Insight Management to complement and extend Unicenter, and to maximize existing IT investments. Network environments that implement Unicenter as the preferred enterprise management platform can use the Insight Integration to help streamline administration and increase systems availability."

Since HEC has already chosen HP System Insight Manager as part of the server management strategy, the combination of Unicenter and Insight Manager seems a suitable candidate and is recommended.

There is a 3-page document called HEC System Monitoring Policy, Draft V1, with the following purpose:

> "The purpose of the Security Monitoring Policy is to ensure that Information Resource security controls are in place, are effective, and are not being bypassed. One of the benefits of security monitoring is the early identification of wrongdoing or new security vulnerabilities. This early identification can help to block the wrongdoing or vulnerability before harm can be done, or at least to minimize the potential impact. Other benefits include Audit Compliance, Service Level Monitoring, Performance Measuring, Limiting Liability, and Capacity Planning."

This policy focuses on security monitoring. It requests that automated tools be used to perform real-time monitoring of Internet traffic, email traffic, LAN traffic, and operating system security parameters. It also requests that some logs and records "be checked for signs of wrongdoings and vulnerability exploitation at a frequency determined by risk." Although Performance Measuring is listed as one of other benefits, there is no specific procedure or data requirement defined.

The policy also states:

> "Currently, all security monitoring is conducted [redacted]

The Altaris® CAD System Manager's Guide prepared for HEC by Northrop Grumman describes eight categories of daily and weekly tasks for monitoring the system performance:

[redacted]. Some of these commands can also be organized in scripts associated with UNIX cron processes to facilitate scheduled monitoring. However, the current guides and practices do not show the employment of automated problem alerting. Without on-site 7x24 support, the collected information cannot be interpreted in time either. It is agreeable that all these tasks are useful; without them, other incidents may have occurred. Nevertheless, some problems slipped through and caused substantial outages. To reduce the hit or miss situations, the following strategy should be considered:

- Perform more frequent checking.[9]
- Provide adequate analysis and interpretation of the collected information, e.g., trending and correlation with historical events. Some (but not all) application-level problems may be reflected upon system performance data. To discover this causal relationship requires collecting and interpreting data collected over a long period of time.

---

[9]There has been a concern that more frequent monitoring might cause degradation of the system performance. The concern is legitimate; however, it should not deter adequate monitoring. The HEC system computing power and network bandwidth are designed to handle above the current workload. It is unlikely that they cannot accommodate the resource requirements for monitoring. On the other hand, a proper planning for capacity should take monitoring requirements into consideration.

- Use automated monitoring tools with comprehensive coverage and robust performance. The automation is not just for automating data collection; it should also do automated event identification and filtering using the relationships obtained from the analysis above.

The Altaris CAD System Manager's Guide also provides a comprehensive list of troubleshooting steps for problem identification and resolution after a problem has occurred.

4.7 Security

[redacted]

5 Process Analysis

In addition to assessing the technical performance of the system, MITRE evaluated the engineering process that were or should be in place to support the overall end-to-end system performance. This analysis is applicable to the overall public safety data and radio systems. As explained in this section, the HEC processes are primarily analyzed but the outcome and recommendations apply to all City of Houston departments public safety data and radio systems. Some of the processes assessed included the following:

- Configuration management to monitor and control the system baseline.
- Risk management to identify areas that may negatively impact system performance, raise these levels to appropriate management level, and plan budget.
- Change management to prepare staff and employees for new methods of doing business and systems operations.
- Requirements management to ensure users have method to identify requirements; system is designed and tested to satisfy requirements.

Our analysis showed that from an end-to-end perspective, the City of Houston engineering processes are fragmented and very few processes are documented. Informal processes do exist in some areas but the effectiveness of these undocumented processes is very hard to evaluate. This section identifies the informal critical processes that are in place at the HEC and then suggest two critical processes that need to be established, risk management and configuration management. This section also discusses two other important items that need to be addressed: training and testing. The criticality of establishing processes can first be shown through the analysis of the outages that occurred.

5.1 Outage Cause Summary

A summarization of the type of problems resulting in the system outages and relevant processes and practices are presented in Table 5-1. This table shows MITRE's analysis of whether the lack of formal processes and practices may have contributed to the outage. The "Relevant Processes" column identifies applicable processes while the "Relevant Practices" column shows applicable practices. The notations in the "Relevant Processes" and "Relevant Practices" columns in Table 5-1 indicate that these processes and controls are relevant to the type of problem encountered. They do not indicate that there was a deficiency in the indicated process or practice, but indicate focus on that process or practice could prevent or minimize similar problems in the future.

Table 5-1. Processes and Practices Relevant to the Outage Problems

Legend: D - Directly control; I - Indirectly control

Time Line		Relevant Processes								Relevant Practices								
		Requirements Management	Configuration Management	Deficiency Reporting	Risk Management	Test & Evaluation	Continuity Management	Release Management	Continuous Process Improvement	Integrating Vision	Acceptance Testing	Involving Stakeholders	Documentation	Usability	on-going support	Maintainabity	Checklists	Training
Sep-03																		
	9/24/2003	D	I		I								I	I	D	I	D	D
	9/30/2003	D									I				D	I		D
Oct-03																		
	10/2/2003				I		I						D				D	D
	10/8/2003				I	I	I											
Nov-03																		
	11/5/2003				I		I											
	11/7/2003				I		I											
	11/10/2003					D												
	11/16/2003				I		I											
	11/28/2003				I		D										D	D
Dec-03																		
	12/3/2003				I		D										D	D
End of Acceptance Testing																		
Jan-04																		
Feb-04																		
Mar-04																		
Apr-04																		
	4/10/2004				I		I								D			D
	4/25/2004				I		I								D			D
May-04																		
	5/10/2004				I		I									I	D	D
Jun-04																		
Jul-04																		
Aug-04																		
	8/8/2004		D		I		I								D	D	D	I
Sep-04																		
Oct-04																		
Nov-04																		
Dec-04																		
	12/1/2004		D	I	I		I											D
	12/7/2004																	
	12/14/2004				I	I	I											

A look at the outage data (Table 4-2, 4-3, and 5-1) shows that many of the problems encountered in the outages are the result of insufficient training or human error. In fact, as explained in Section 4.2, the majority of outage time can be contributed to insufficient training or human error.

Only two of the outages were the result of requirements. In both cases, human error caused the wrong requirement to be provided. Thus, MITRE concludes that the requirement management process used has been sufficient. However, there were two incidents where the lack of adequate configuration control had a major impact. In a couple of these instances, the problem causing an outage was repeated later. This reoccurrence indicates that the problem resolution process may not be adequate since the actual cause of the first outage was not resolved sufficiently to prevent it from reoccurring a second time.

Table 5-1 shows that institution of risk management and continuity management would be beneficial in the likelihood of decreasing the outages occurring. A risk management process could have identified potential problems before they became outages. A risk management process will make the City of Houston more proactive in assessing and improving the performance of the overall system. The purpose of a risk management process is to identify risks and prioritize them so that limited funds can be spent where they will have the most beneficial impact on the program.

Table 5-1 also shows that a continuity management process could have a beneficial role in most of the outages. The continuity management process is not only responsible for the contingency plan in the event of an outage, but also evaluates other changes that could minimize the chance of an outage in the event of the type of problems seen. MITRE noted that the City of Houston has an effective process for continuing operations during system outages. The "paper" process adequately maintains the degree of call taking and dispatching operations needed for emergency services. The chart was primarily focused on continuity of operations for the CAD system and should not be misunderstood to indicate there is no continuity management.

5.2 Critical Processes

MITRE assessed and evaluated processes applicable to all departments and those within the HEC. The major process or strategic planning activity that relates to the scope of this assessment was the City of Houston Executive Order "Policy to Direct and Monitor Technology Efforts." As stated, the purpose of the Executive Order is to establish and communicate the City's technology strategic direction. The Technology Steering Committee, or a committee of equivalent scope, could provide a forum to oversee and pursue the recommendations made in this report.

The team was unable to find documented engineering processes that were applicable to and being used by all departments. Thus, MITRE recommends that the City of Houston identify and adopt policies and procedures that implement critical engineering processes such as problem resolution, risk management and configuration management. Next, the MITRE team reviewed HEC policies and interviewed HEC staff to identify processes in place at the HEC. Table 5-2 contains a list of the process, the applicable policy and general status information.

Table 5-2. HEC Processes

Process	Applicable Policy	HEC Process Status
Risk Management		No formal process
Requirements Development & Management		Uses a system of functional design and detailed work packages to manage requirements.
Configuration Management		Informal process in place for change control.
Enterprise Integration		No formal process.
Integrated Testing		Internal testing process. Individuals test requirements using a loose test plan that tests all functionality.
Capacity Management		No formal process.
Continuity Management		No formal process.
Incident Management	Software Incident Report Tracking (SIRT) Form	On a "by exception" basis.
Change Management	HEC Change Management Policy draft	As implemented, a change control process not change management.
Change Control		Informal exists.
Release Management		No formal process.
Service Management	Maintenance Agreement	No formal process.
Problem Resolution		Informal process exists.
System Enhancement		Informal process exists.

5.2.1 Problem Resolution Process

Figure 5-1 shows one of the critical support processes that currently exist is the problem resolution process. This process is utilized when a system issue arises. HEC's Problem Resolution Process, as shown in Figure 5-1, provides the correct steps in collecting the problem, sending the problem to the correct people for authorization and resolution and requesting resolution confirmation from the problem originator. However, in a couple of outages, the problem was not correctly identified or resolved; thus, the problem occurred again and caused an additional outage. Assuming this process is consistently followed, there may be difficulties in reproducing the problem which make it difficult to know if the problem has or has not been correctly resolved. HEC also has a method of documenting problems for their systems in a SIRT List and a Change Order List. These lists provide the following information: applicable agency that submitted the problem, the agency priority of the problem, description of the problem, estimated completion date, and status. The lists provide an excellent summary for tracking problems and a similar tracking capability should be implemented city-wide for all of the public safety system problems identified. This list could also support the risk management process by providing input on the major issues that need to be resolved and identified.

Additionally, the problem process resolution appears to be followed by all stakeholders when system outages and system performance degradations are experienced. The breakdown occurs when non-critical system issues and problems are identified, especially for issues with functionality of the system. There is currently no clear, documented escalation process for functionality issues. The escalation procedures for non-critical system issues, as well as the communication feedback on resolutions or decisions to end-users raises such functionality issues.

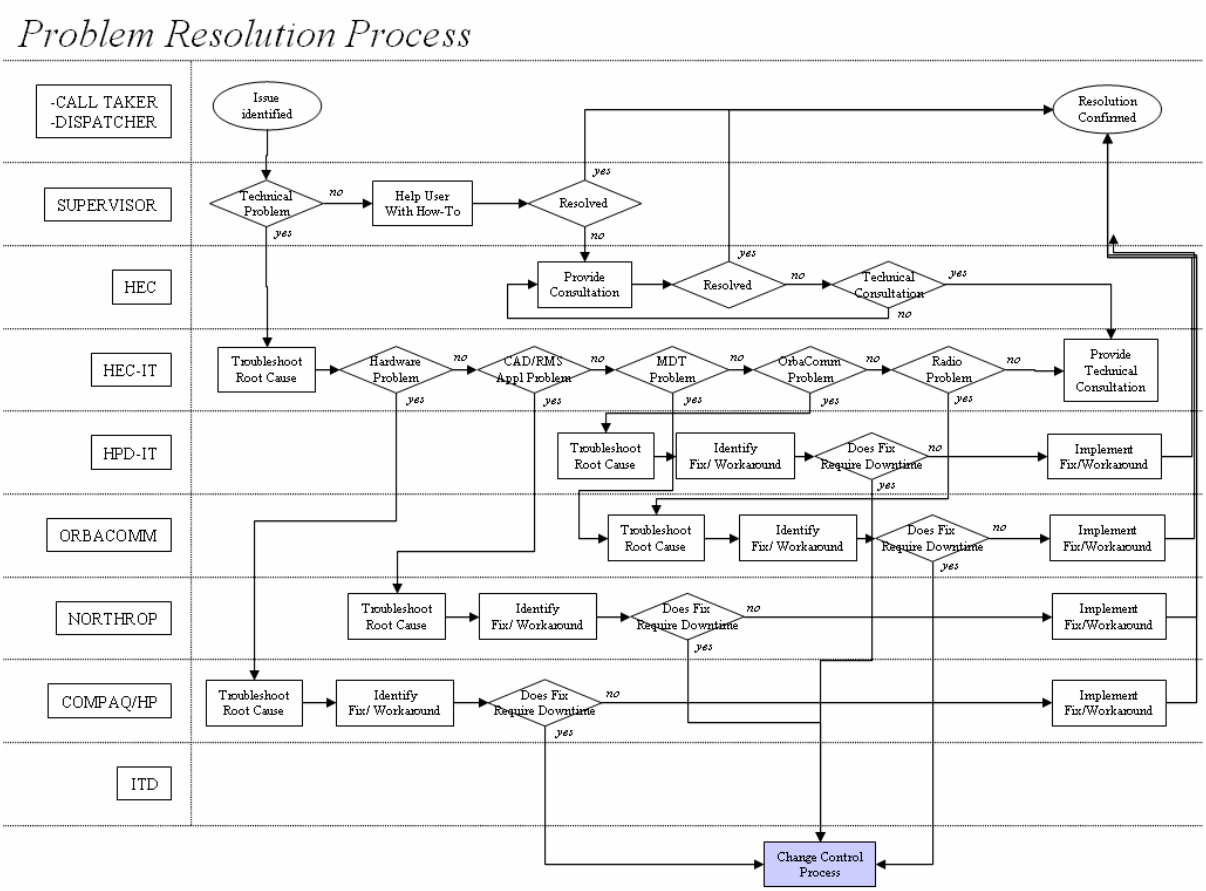

Figure 5-1. Problem Resolution Process

5.2.2 System Enhancement Process

The system enhancement process shown in Figure 5-2 is used to address issues that have been identified as changes to the original or existing system functionality or architecture that the City of Houston accepted as part of the acceptance sign-off between Northrop Grumman and the City of Houston.

All system changes are considered enhancements and therefore must undergo a process of review to determine the following: whether the change is needed, impacts of the changes to the existing functionality and architecture, prioritization of requested enhancements/changes, funding for the enhancements/ changes, and the expected turnaround for the vendor to deliver agreed/accepted enhancements.

The HEC System Enhancement Process, is a good process. The diagram shows that the proper steps, correct people, and validation are included. However, the actual turnaround time from when the enhancements are approved by HEC and the time it takes for the vendor to deliver the

agreed to enhancements is not in alignment with customer expectations and appears not to be in accordance with mutually agreed to timelines established at the beginning of this process.

MITRE recommends that HEC and Northrop Grumman establish an enhancement/release task team to clear out the backlog of changes and enhancements in existence for quite some time now. These changes/enhancements have been reviewed, designed, approved, and scheduled for development but no enhancement deliverables have been provided.

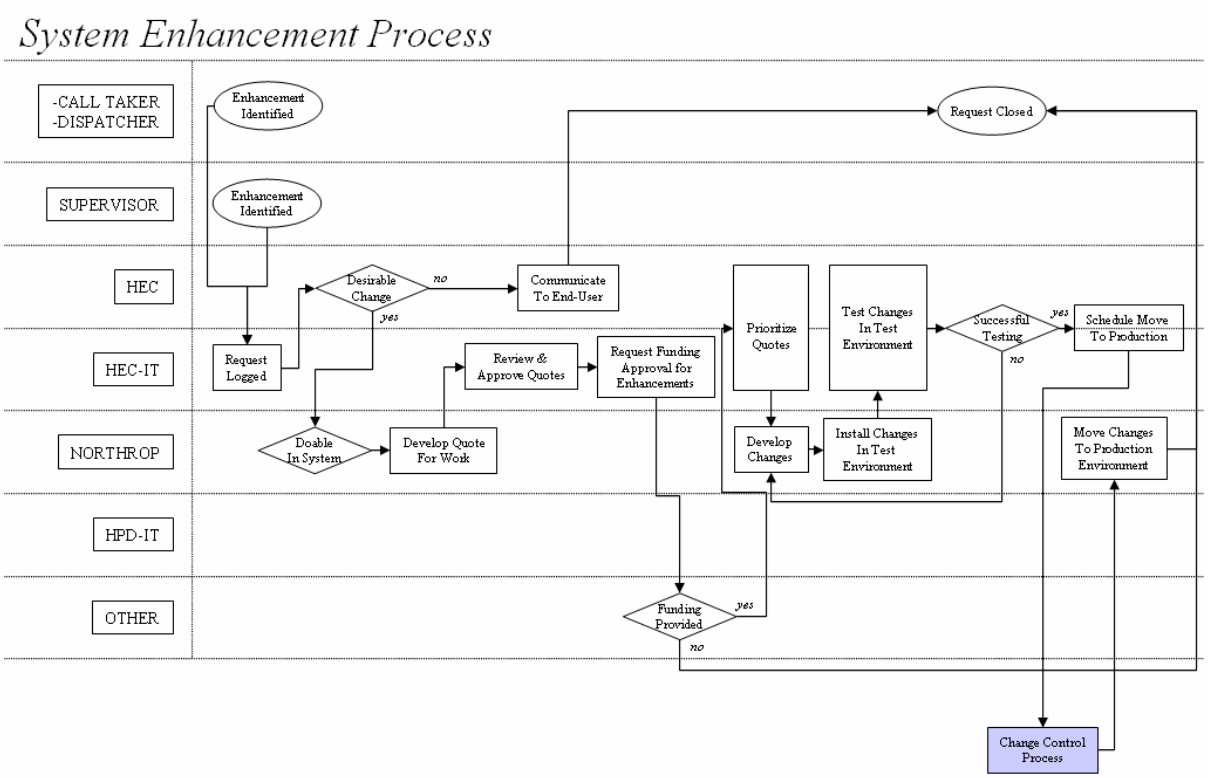

Figure 5-2. System Enhancement Process

5.2.3 Change Control Process

Figure 5-3 illustrates the existing Change Control process in place at HEC for changes. The source of changes can either be enhancements or problems/issues that are affecting the system performance.

The HEC Change Control Process is not complete. It does not provide details on the following:

- Approval process
- Review board
- Organization roles
- Tracking of changes
- Configuration control

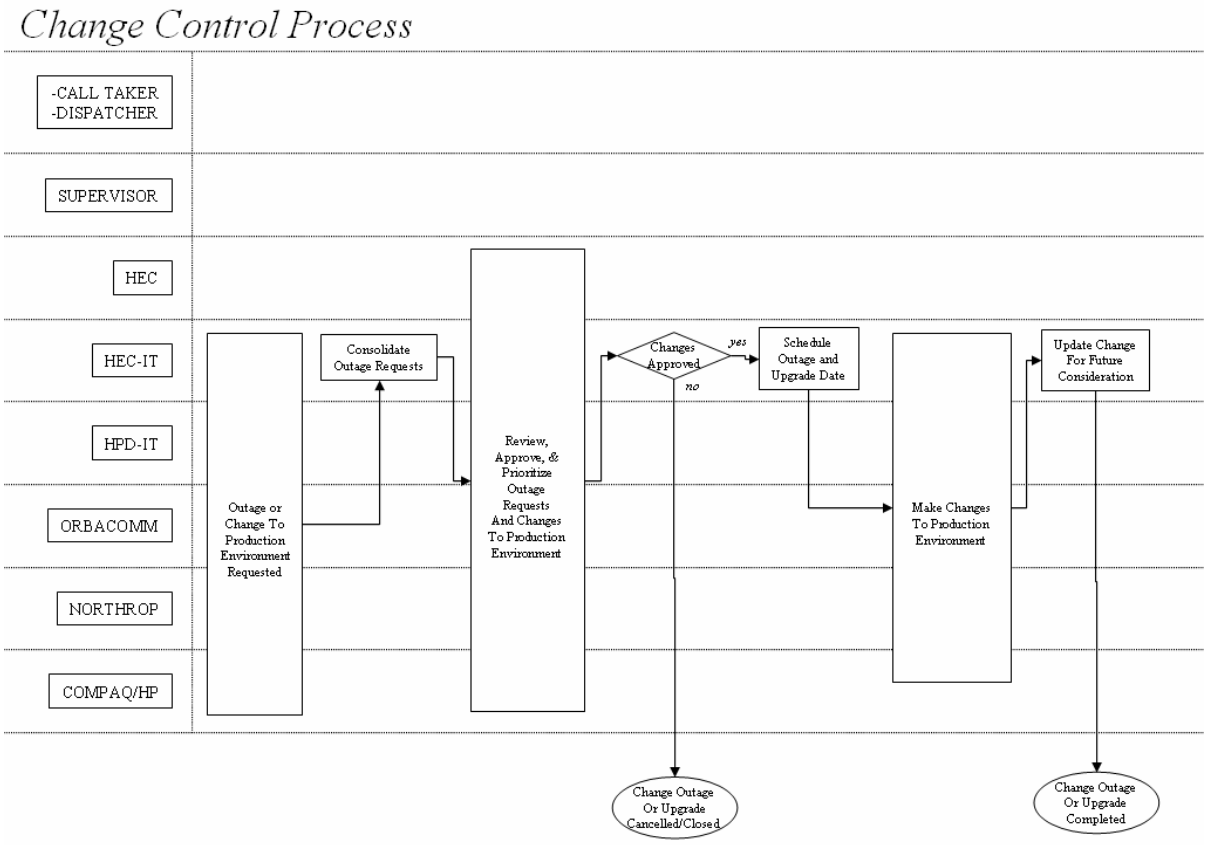

Figure 5-3. Change Control Process

While the change control process is understood by all IT support stakeholders, no evidence of documented procedures and change control communications were identified. The lack of documentation makes it difficult to conduct after action reviews (AARs) on executed changes. Furthermore, documentation on successfully testing as well as back-up procedures for proposed

changes was also not found increasing the risk of prolonging change outages in case of unsuccessful changes, as well as mitigating risks of proposed changes.

5.2.4 Recommended Engineering Processes

The MITRE analysis of the existing and recommended processes evaluated several engineering processes. The team recognized the benefit of each of the processes but realized it would be unrealistic to implement all of them at this time. The critical processes that are lacking and are recommended include risk management and configuration management. An effective risk management process focuses on the risks with the highest probability of occurring and greatest impact if they do occur. For example, the single point of failure and high priority SIRTs may be risks identified by the City of Houston. This process will help the City of Houston become more proactive rather than reactionary. To be effective, the risk identification must be solicited from all departments, contractors, and the Greater Harris County 9-1-1 Emergency Network and it must be visible to all team members so that risks are seen from all points of view. To avoid an unwieldy number of risks from being tracked, the risk identifier simply "proposes" risks. A risk review board, with representatives from HFD, HPD, HEC and ITD should be used to review a proposed risk and then either accepts or rejects the risk. After risk identification, the risk should be assessed for their impact on cost, schedule, technical performance or other impacts such as regulatory, security, or political. The risks should then be prioritized based on their probability of occurring and the consequences if the risk were to occur. The risk review board, senior leadership and budget authorizers decide which risks should be given resources for mitigation.

Next, risk managers should be assigned for the risks given resources. Risk managers should develop mitigation plans, determine how they will know if a mitigation is successful and develop contingency plans in the event a risk mitigation is not successful. They then track the status of the mitigate on plan and close the risk when appropriate. This risk management process results in the avoidance or minimization of the impact of consequences of risks with the smallest expenditure of funds.

The MITRE team also recommends that the City of Houston implement a city-wide configuration management process. The "Little Book of Configuration Management," November 1998 from the Software Program Managers Network provides an ideal framework for creating a configuration management process.

As defined in the above source, configuration management is the basic project control mechanism that establishes and maintains the integrity of products through the project's life cycle. Configuration Management will support the City of Houston by providing:

- Configuration Identification -- The ability to identify what information has been approved for concurrent use in the project, who owns the information, how the information was approved for CM control, and the latest approved release.
- Configuration Control -- The configuration control process and procedures designating the level of control through which each work product must pass (for

example, author control, project-level control, acquirer control); identifying the persons or groups with authority to authorize changes and to make changes at each level (for example, the programmer/analyst, the software lead, the project manager, the acquirer); and the steps to be followed to obtain required authorization for changes, to process change requests, to track changes, to distribute changes, and to maintain past versions. Change control provides the mechanism to build software systems for tests that have a known configuration and can be exactly reproduced.

- Status Accounting -- Formalized recording and reporting of the established configuration documents, the status of proposed changes, and the status of the implementation of approved changes. Status record information provides an accessible and current record of the status of each controlled piece of information that is planned to be used, the content of each release from CM, and who has checked out or is working on a piece of information that the test organization plans on accessing through CM.
- Reviews and Audits -- Frequent evaluation of the content, baseline integrity, and release integrity of all controlled products to ensure they conform to their configuration documents.

5.3 Training

The MITRE analysis included a review of training documents and interviews with the staff to gain an understanding of the past and current training. Two types of training were identified as important to preparing staff to operate and sustain the system. The first is operator or user training for the CAD system and the other components used in the performance of call taking or dispatching. The second is the training of the HEC and support staff to support the monitoring, maintenance, management, and utilization of the system.

The two methods use different models for the training relationship. The user training side of the HEC is performed by training staff that is part of HEC (call taking) or assigned to HEC duties (dispatching). The relationship between HEC staff and system providers in this area is a "train-the-trainers" model. The training from Northrop Grumman and third-party training providers to the HEC IT and support staff uses a direct training model, as it assumes that the trainees will be doing the work themselves.

The assessment focused on identifying future training issues for the CAD user community and support staff at the HEC. As such, it identified the steps that can be taken to ensure that the training needs of the CAD system users will be met as the system evolves and changes. In turn, this will include recommendations that will address processes used in managing changes in the CAD system.

5.3.1 IT/Support Staff Training

Initial requirements for training of the support staff are detailed in the Scope of Services Section E, Training Plan, June 13, 2001. The discussions with Northrop Grumman personnel and HEC

staff suggested that these training requirements had evolved in the context of a set of assumptions about the future responsibility of the HEC staff for expanding the role of HEC to support the maintenance and control of the CAD system.

The current assumption can be described as "HEC manages and monitors, while Northrop Grumman corrects (SIRT) and extends (change order)." The training plan has evolved to reflect this assumption.

Certain training regarding the setup, initialization, and startup of the core systems of the HEC were completed as originally planned. The staff that took these courses are still in place, so the consequences of turnover in this area is not being addressed.

The original plan for training encompassed a "Programmer Training Course," and a collection of third-party training on [redacted] administration and management.[10] The lack of change access to the source code undercut the original intent of the programming course. The compression in the startup of the HEC system prevented the achievement of most of the third-party training. As a result, two types of training programs were developed.

The programming course originally focused on modification of the interface code to the user and to various external systems. At the request of the students in the session, the instruction method was converted to a more detailed architectural review of the system structure. The intent was to enable the monitoring and management role for the HEC support staff.

The second response was in the area of third-party training. A change order to the contract was processed to convert the money assigned to the listed third-party training into a pool that could be drawn from after startup to meet the training needs for the support of the system. This change has enabled the training to be shifted in time. Some of the resources have been used in exactly the same arena as originally planned, such as [redacted] administration and management. As HEC and related staff still have responsibility for the development of database applications, these resources have been used for training in Oracle related areas.

The approach that has evolved here appears to meet the current intent of the HEC management regarding support for system maintenance and management. Eventually, issues such as a response to turnover and refresher training will need to be addressed.

5.3.2 User Training and Training Processes

A key component in the assessment of training needs at the HEC derives from the different roles that the HEC encompasses. A number of distinct communities exist within the mixture of call taking and dispatching, Police and Fire/EMS. The nature of performance requirements laid upon members of each community differ, and the norms for handling issues differ by profession.

The assessment was based on interviews conducted during the period December 15-17, 2004, plus a review of documents provided by HEC, HPD, HFD, and the City of Houston staff.

---

[10] See pp. 13-21, Section E, Training Plan, June 13, 2001.

5.3.2.1 Issues

Historically, both the police call takers and dispatchers were employees of the HPD and worked in an HPD facility. This proximity to the officers undoubtedly helped to inform and communicate to the call takers and dispatchers what the primary issues were for the dispatch role, from the officer perspective. The police call takers are now employees of HEC, and the dispatchers, while still employees of the HPD, are now resident and trained at the HEC facility. The decreased direct exposure to the HPD environment carries with it an implicit training loss.

The HEC has recently completed [February 4, 2005] its first full training class of new 911/Police call takers. The training period exceeded six weeks, and includes supervised floor time beyond that time period. Preliminary comparison of the training approaches and schedules used by HEC and related dispatch organizations place the Houston effort in the middle to high range of call taker and dispatch training efforts across the United States.[11] Police domain language and expertise has been addressed in the first training session, and the issue is being watched.

Similar arguments and concerns apply to the Fire/EMS call takers. On the EMS part, the additional training to support delivery of pre-arrival instructions will provide a degree of exposure to the language and concerns of EMS. A similar effort will be needed to address fire issues.

Training issues can engage other agencies more directly. For example, some of the discussions with HPD personnel made it clear that HPD training issues outside of the HEC have specific impacts on police dispatcher workload. For example, HPD officer training on the full range of capabilities of the MDT system would provide the opportunity to take some of the load off HPD dispatchers working within the HEC facility.

---

[11] See, for example:
Washington State Criminal Justice Training Commission
CJTC Telecommunicator Program Website
http://www.cjtc.state.wa.us/telecom/index.htm

State of New Jersey
Office of Emergency Telecommunications Services
http://www.state.nj.us/911/
http://www.state.nj.us/911/trainingregs.html

Dispatch Monthly Magazine Training Resources
http://www.911dispatch.com/train_file/training_menu.html
http://www.911dispatch.com/train_file/train_survey.html

Illinois' Public Safety Telecommunicator Training & Standards
http://www.911dispatch.com/train_file/illinois_training.pdf

The HPD dispatcher can provide a number of query services to various databases for the officer in the field. Examples of these services include queries regarding status of driver's licenses, automobile registration, outstanding warrants, etc. Some of the functions of the CAD/RMS and query systems can be performed by the officer directly from the MDT in the police patrol vehicle.

Anecdotal evidence was provided to the effect that the HPD had not received formal training for officers on the use of the MDT system for approximately eight years. No MDT training was provided at the academy, instead, use of the MDT was one focus of on-the-job training. The result was a wide variance in the individual officer familiarity with the capabilities and methods for use of the MDT. In turn, this leads to officer requests for dispatcher performance of tasks that could be done by the officer in the vehicle, increasing dispatcher and system workload.

Currently, as of December 16, 2004, an in-service training program on the MDT capabilities had been developed by [redacted] HPD. According to him, 56 in-service training sessions had been scheduled to bring MDT training to the entire force, and 28 of them had been completed at that time. [redacted]

5.3.2.2 Processes

Preparing for the future engenders two types of process consideration for training staff. The first is enabling the feedback that supports the continuous improvement in the quality of the training process. This addresses both improvement in training methods and content of material. The second is ensuring that the training considerations of future system changes are addressed before the system changes are implemented.

In the interviews conducted at HEC, there is considerable evidence that call takers and dispatchers are working to learn the system, and to learn ways of making it work. The HFD dispatchers tend to create personal "bibles" of techniques for achieving specific goals through the system. These documents or other collections of information and experience are a potential source of new material for training. The experience of each individual as they move from the classroom to the real world of the operations floor is a new pair of eyes looking at the potential of tuning the training. The experience of new and experienced officers can provide input to improve the training process. Processes to obtain and "mine" these ad hoc tools for information to improve training should be developed.

The planned change process of the HEC needs to explicitly engage the training community for a number of reasons. As changes are proposed, the impact to the normal training process needs to be evaluated. The costs of training to meet the new environment, in both schedule and resource terms, needs to be part of the change management process.

Examples exist of how this type of issue has been implicitly addressed and missed were provided during the interviews. Some loops are closed coincidentally because of staff fulfilling multiple

roles. For example, the members of the HPD Dispatch Training Unit are also the representatives to the Change Board Meetings, so the issue of training impact of proposed changes is automatically considered. Our current understanding is that HFD does have representation on the Change Board, so a presence has been established with the responsibility for assessing training consequences on the Fire/EMS side of the HEC. This is one mechanism to achieve this goal.

The key issue here is that current written policies of the HEC do not guarantee consideration of training issues. Membership in the change management committee for the affected training communities is one mechanism for achieving that consideration. Other mechanisms, based on process requirements, can also address training consequences of change. While voting membership in a Change Committee may not be required, guarantees of awareness and a forum for noting training consequences and needs are essential.

5.4 Testing

MITRE reviewed the implementation plan, acceptance testing documents, Go-Live documents, and other material provided by the HEC to determine the processes used for testing. MITRE also had extensive conversation with HEC IT and Northrop Grumman to discuss this topic. The analysis shows that the testing process for the pre-acceptance period was much more exhaustive and complete than the post-acceptance testing. In addition to the specified Scope of Services testing, the Go-Live testing provided an excellent measure to evaluate the systems readiness for operations. [redacted]

- [redacted].
- [redacted].
- [redacted].

Northrop Grumman's internal processes ensure that testing of software changes and new releases verify that the software operates correctly and that the functional change meets requirements.

Given the criticality of the entire system, some form of periodic failover testing is needed to ensure that the site is ready to accommodate different contingencies. This testing should include all major systems including SAN, router, database, and communication link. Eventually this failover test concept should encompass a full business continuity plan that includes disaster recovery. With the current A and B sides to the servers and SAN storage along with a test system (that could be further expanded), a concept for a three sided environment could be architected that could provide better testing, training, integration, and failover analysis. The "third" side could also evolve toward a disaster recovery system that could be eventually remotely located and continue to function not only as a hot spare but to support additional training, conduct improved testing, and better integration of future capabilities.

MITRE recommends that the HEC acquire a full test suite of equipment and software. At a minimum, the full test suite should include the same architecture configuration of the production system including CAD and RMS equipment and software as well as external equipment such as MDTs. The configuration should be used for the following:

- Functional and regression testing of major system upgrades.
- Functional and regression testing of maintenance software releases.
- Load testing of current configuration and to address possible growth.
- Interoperability testing of current and future changes to legacy and external systems.
- General troubleshooting and analysis.

6 Recommendations

The end-to-end performance of the City of Houston public safety system can be improved through the incorporation of short term technical solutions and long term strategic activities. Many of the recommendations are similar in overall scope to major goals and priorities identified in the "City of Houston Technology Investment Plan, Fiscal Years 2005 – 2009." The MITRE assessment independently identified solutions that can improve the performance of the existing public safety system. This section will identify the solutions and activities that should be taken by the City of Houston.

The first actions that should be taken are those that are needed to reduce the occurrences of major outages that have been experienced in the past. These actions include:

- Establish responsibility for end-to-end system management and integration.
- Eliminate single points of failure and establish effective automatic fail over.
- Increase system maintenance scope and time periods to provide a tiered 7x24 support team (technicians and public safety system help desk).
- Enhance HEC system performance monitoring and analysis.
- Enhance security [redacted].
- Document current processes and incorporate formal configuration management and risk management processes.

MITRE recommends that the City of Houston appoint or identify at least two positions with overall responsibility for the end-to-end system management and integration of the public safety system. The first position would be responsible for the performance of all of the systems (i.e., network, radio, voice and computer) that support the full operations of public safety from call taking to emergency response as shown in Figure 6-1. The position would be responsible for resolving system integration issues, budget preparation, technical staffing, contact monitoring and direction, and other related management responsibilities. The second position would establish a lead program engineer to provide technical support to the management position. The lead program engineer would be responsible for resolving technical issues, overseeing system testing and performance monitoring, establishing and implementing engineering processes, and providing technical advise as necessary. These positions should not be established to replace or supplement current roles served by HEC, HPD, and HFD staff. On the contrary, they would support the Director of Public Safety in his responsibilities to oversee the operations of the public safety system. This recommendation requires increased personnel budget costs for the manager and engineer and is a recurring cost.

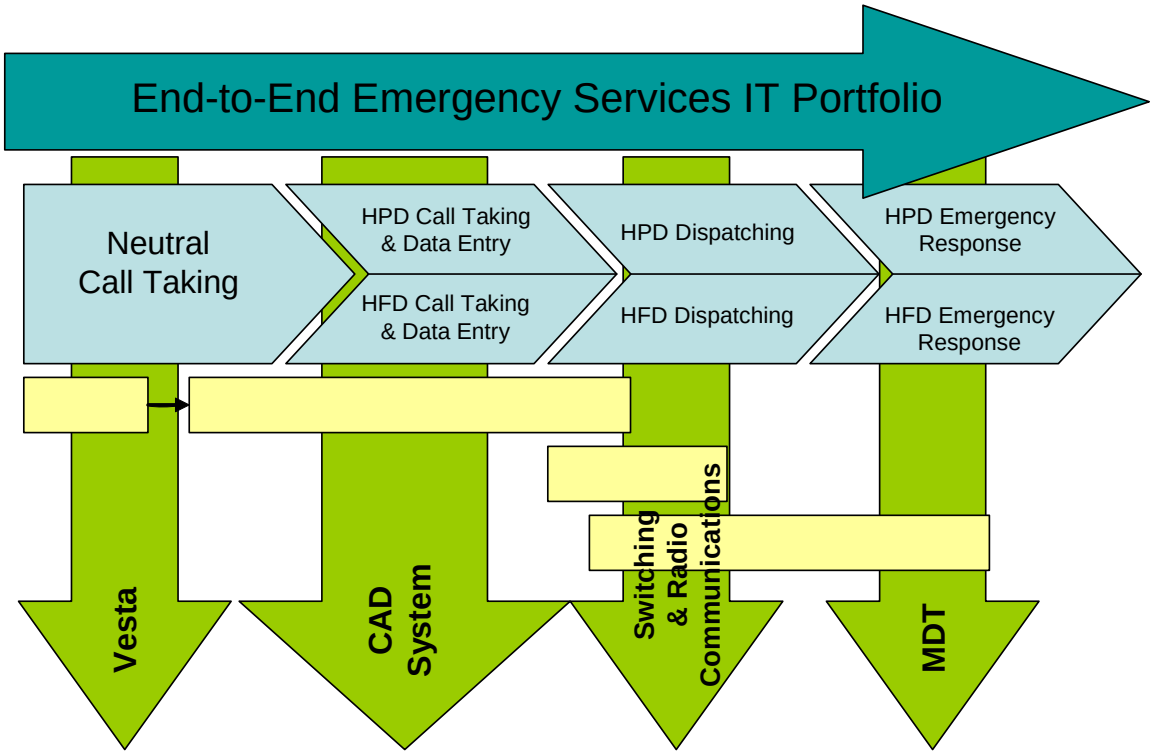

Figure 6-1. End-to-End Portfolio

MITRE recommends that the City of Houston immediately eliminate major single points of failure that could render the public safety system unavailable to HEC, HPD, and HFD users. The SANs and the integrated database should be upgraded [redacted] as discussed in Section 3. Both of these components have caused major outages in the past and a failure in them could cause repeat occurrences. Some of the fixes may involve technical changes while others may incorporate new processes or procedures. This recommendation impacts one-time equipment cost during the year of purchase.

MITRE recommends that the City of Houston expand the maintenance contract to expand the equipment warranty coverage, help desk support, and 7x24 service. The City of Houston currently has basic support service and preventative maintenance under their current agreement with Northrop Grumman. The City of Houston should consider exercising the option to add corrective maintenance offered by Northrop Grumman or to obtain an equivalent service. This option would help to potentially resolve issues while the new system is still going through its early stages of operations. The estimated cost to add the expanded maintenance coverage above what the City of Houston currently has is approximately $550,000 annually.

MITRE recommends that the City of Houston enhance its system performance monitoring and analysis capability. The network monitoring tools discussed in Section 3 should be procured and operated to help monitor system performance and to support isolation of technical problems before they become major problems. In addition, the City should request the monthly and other reports required to be gathered by Northrop Grumman according to the Scope of Service and Maintenance Agreement to help analyze overall system performance. This recommendation requires a contract review with Northrop Grumman to determine if the tools provided by Northrop Grumman can meet the performance monitoring and reporting requirements. If not, then additional tools will need to be purchased.

[redacted]

The existing engineering processes should be documented and approved by the appropriate manager. In addition, the City of Houston should incorporate configuration management and risk management processes that can be applied to all departments. This recommendation requires development of formal policies and procedures. It impacts all departments current operations and requires personnel resources from all of them.

The next recommendations are intended to identify how the general system performance can be improved. They include:

- Measure and monitor the system's end-to-end availability.
- Develop end-to-end performance monitoring and analysis.
- Replace obsolete equipment and software. HEC should establish tighter control and tracking of equipment and software expected life through a formal configuration management process. At a minimum, the equipment identified as end-of-life in this report should be replaced.
- Enhance testing capabilities and processes.
- Identify and measure user and system performance statistics.

The current system availability requirements only apply to the CAD and RMS applications. Other failures can occur and cause long outages without consequence or requirements for immediate resolution. MITRE recommends that the City of Houston define and measure system availability to include all hardware, applications, software, communications systems and interfaces. The availability numbers should be based on the criticality of the system or function to the effective operations of the call takers and dispatchers. This recommendation

impacts current contracts and may require significant changes. It also impacts the budget for recurring warranty services and new equipment.

Performance monitoring and analysis are needed at all levels and for all parts of the system. MITRE recommends that each department have the capability to monitor and analyze the portions of the system that they primarily maintain and operate. This information should be shared to provide a City-wide view of system performance that can be analyzed and shared with all departments. This recommendation impacts current department operations and sustainment. It requires a significant financial investment and change in some staff roles.

All equipment and software should be tracked and monitored to identify possible end of life or obsolescence. This information should be documented and incorporated within the strategic planning and budgeting for new systems. The cost of this upgrade is for equipment that needs to be upgraded or replaced.

Northrop Grumman and HEC's testing role in incorporating new changes and upgrades need to be formally documented to ensure that complete testing occurs. The testing process needs to include the following tests: functional, regression, loading, and interoperability. In addition, the City of Houston should acquire and maintain a complete test configuration of the system. This test configuration should include all of the CAD, RMS, MSS, SANS equipment and software, as well as external systems, where feasible. This configuration will provide full testing capabilities and may also be an additional back-up system to the existing equipment. [redacted] This recommendation impacts recurring personnel costs for new staff role. It also impacts one-time equipment costs.

The current system has limited system and user performance requirements that need to be met throughout the life of the system. The City of Houston needs to define minimum performance requirements that need to be met and monitored. These requirements should identify the critical functions and performance times that must be sustained throughout the system life. Both the average and threshold performance parameters should be specified. This recommendation impacts contractual agreements. It requires new change order specifications or contract modifications to identify performance requirements.

The last recommendations are those that are needed to support the life-cycle of the system through operations and sustainment.

- Determine appropriate Contractor and City of Houston system operations and sustainment model.
- Develop end-to-end public safety strategic plan, architecture and roadmaps.
- Incorporate disaster recovery system and processes.
- Decrease application customization.

The current agreement and operations do not clearly specify the roles of contractors and the City of Houston to support an end-to-end system. This lack of understanding contributes to operations and sustainment performance issues. As a part of its overall strategic planning, the City of Houston needs to determine whether it will primarily use an outsourcing model for operations and sustainment or whether these skills will be developed in-house. Next, the City of Houston will need to specify with the contractors, the various roles and approach for achieving the model selected.

The City of Houston has developed strategic goals and plans for its information technology consolidation. MITRE recommends that a separate effort be focused on developing strategic goals, planning, and budgeting for the public safety system. This end-to-end approach should include the current system as well as all systems that rely on it.

MITRE recommends that the City of Houston develop a disaster recovery capability for the current data system. This capability should include the minimum equipment necessary for call takers and dispatchers to continue emergency services operations in the event the HEC facility equipment is not available.

The current system software has a large percentage of customized code. This customization matches the current operations of HFD, EMS and HPD. Over time, the degree of customization will affect the systems long-term performance and sustainment. Based on information that is widely accepted in the industry, most CAD systems are replaced every ten years. However, the software is usually updated periodically during this period as new software releases are made available. For those systems that are highly customized, agencies budget out year monies to port existing customizations to the vendor's latest software releases. Because of the additional out year costs, (integration services) associated with upgrading these one of a kind systems to the vendor's current software release(s), smaller locals and agencies tend to install the vendors base software offering from day one. Larger locals and agencies that have customized software and can afford these integration costs budget accordingly.

On the hardware side, users usually change out their hardware (servers, disk storage, routers/switches, and workstations) every three to five years depending on new technological breakthroughs, vendor discontinuation and support, or new configurations or new software capabilities that require newer hardware to function properly.

MITRE recommends that the City of Houston assess where customization can be decreased or eliminated. This assessment will evaluate the two primary functions that the customization supports. These functions include: fire and police tall taking and dispatch operations and (2) monitoring legacy systems.

.

Appendix A Referenced Documents

1. Contract to PRC Public Sector, Inc., to Implement Consolidated Dispatching at the Houston Emergency Center, Request for Council Action, August 20, 2001
2. Houston Emergency Center Technology Management Plan, Arthur Andersen, May 15, 2002
3. Strategic Technical Plan For the City of Houston, Friday, July 20, 2001, Arthur Andersen
4. Policy to Direct and Monitor Technology Efforts, Executive Order No. 1-44 Revised, November 25, 2002
5. Agreement for the Purchase of Equipment and Licenses of the Police and Fire CAD System and Fire RMS/MIS Systems Software and Hardware and Provision of Services and Maintenance
   Attachments:
   1. Scope of Services
   2. Police MIS Database
   3. ANI/ALI Message Format
   4. Existing Police CAD Interfaces
   5. Racial Profiling Data Collection
   6. Workcard Data Collection
   9. Response Levels
   11. Emergency Alerting System
   12. TDH Trauma Reporting
   13. Quick Dispatch Requirements
   14. Resource Recommendation Procedure
   15. Highrise Documentation
   16. HAZMAT Documentation
   19. RMS Issues
   22. RMS EMS Field Layout
   23. Existing Keyboard Layouts
   24. Disposition Processing
   25. Personnel Accountability Review

26. HFD Running Schedule Regeneration Documentation
27. PSI Deliverables and Pricing
    a. PRC's Response to the City of Houston's Revised Scope of Services for Houston Public Safety Dispatch System
    b. CAD & RMS Acceptance Test Plans
    c. Acceptance Test Plan for Altaris® Fire RMS Implementation
    d. Cost Proposal
    e. Bill of Materials
    f. Compaq Value-Added Implementation Services
    g. Preliminary Project Schedule
    h. Training Plan
    i. Documentation and Manual Requirements
    j. Electrical Specifications
    k. Oracle Licensing Requirements
    l. Network Requirements
28. Payment Schedule
29. Maintenance Agreement

6. Approval of Change Order #1 to Contract With PRC Public Sector, Inc,. for Implementation of Houston Emergency Center Consolidated Dispatching System, Request for Council Action, July 16, 2003
7. City Of Houston Technology Investment Plan V1, Fiscal Years 2005 – 2009, Draft, December 10, 2004
8. HEC ITS FY04 Budget (1820)
9. Houston Emergency Center, Dress Rehearsal #2, August 12, 2003
10. Houston Emergency Center, Dress Rehearsal #3, September 12, 2003
11. HEC Go-Live Timeline, September 22, 2003
12. Police Call Processing Time Reports, September 2003 – November 2004
13. Monthly Report, Response Time by District, Priority One – Priority Three, January 2003 – November 2003

14. Monthly Report, Response Time by District, Priority One – Three, January 2004 – November 2004
15. Emergency Communications Division, Call Data 2003
16. Emergency Communications Division, Call Data 2004
17. Number of Calls by Priority, Year: January – December 2003
18. Number of Calls by Priority, Year: January – November 2004
19. Houston Emergency Center, 2004/2003 Police Call Volume
20. Houston Emergency Center, Total Call Volume, Comparison 2004 vs. 2003
21. Houston Emergency Center 2004/2003 Fire/EMS Call Volume
22. Houston Emergency Center, Call Processing Times, January – November 200
23. Application Performance, Table Names:iApplicationStat, December 15 – 16, 2004
24. Application Performance, Table Names:mApplicationStat, January – December 2004
25. EMS Incident/Response/Patient Summary, January 1 – September 23, 2003
26. EMS Incident/Response/Patient Summary, January – November 2004
27. Houston Fire Department Summary Report for Fire, January 1 – September 23, 2003
28. Houston Fire Department Summary Report for Fire, January – November 2004
29. Houston Fire Department Call Processing Time Report, January 1 – September 23, 2003
30. Houston Fire Department Call Processing Time Report, January – November 2004
31. CAD/RMS Outages, September 2003 – December 2004
32. CAD Call Taker Reference Guide
33. CAD Call Taker Practice Guide
34. 40-Hour CAD Call Taker Training Course
35. CAD Course Testing
36. Dress Rehearsals – Dress Rehearsal Summary
37. [redacted]
38. [redacted]
39. [redacted]
40. Additional Training Roll Call Training Issues

41. Houston Police Call Takers Attending Training, March – May 2003
42. Policies and Procedures

    HEC Acceptable Use Policy, Draft

    HEC Antivirus Update and Configuration Procedure, Draft

    HEC Change Management Policy, Draft V1

    HEC Data Backup and Recovery Policy, Draft V1

    HEC Desktop/Laptop Build Procedures, Draft V1

    HEC Intrusion Detection Standard

    HEC IT Department Security Policy, Draft V1

    HEC Network Access Policy, Draft V1

    HEC Password Standards Policy, Draft V1

    HEC Patch Management Procedure, Draft V1

    HEC Router/Switch Security Procedure, Draft V1

    HEC Systems Monitoring Policy, Draft V1

    HEC Third Party Equipment Standard, Draft V1

    HEC [redacted]

    HEC VPN Access Procedure, Draft V1

    HEC Windows 2000 Security Update Procedure, Draft V1
43. Security Assessment Report Findings and Recommendations, Strategic Network Consulting, July 26, 2004 and Supporting Material
44. Budgets 2003 – 2005
45. CAD Appendix F: PRC TCP/IP Protocol Specification
46. CAD Appendix G: Houston CAD Message Content Document
47. Functional Design Control Number, 001 – 071
48. City of Houston Houston Emergency Center Computer Aided Dispatch Server Configuration Information, October 22, 2004
49. City of Houston Houston Emergency Center Message Switching System Installation and Configuration, September 2, 2004
50. Altaris® CAD System Manager's Guide Prepared For Houston Emergency Center, February 2, 2004 – Draft

51. City of Houston 911 Production Server System Platform, Houston, Texas, Pioneer Technical Documentation, April 30, 2002
52. HEC Project Sign-in Sheets
    a. Functional Design Session, November 13, 2001
    b. Functional Design Review, January 7, 2002
    c. Functional Acceptance Test – CAD, February 3, 2003
53. Policy to Direct and Monitor Technology Efforts, November 25, 2002
54. Commission on State Emergency Communications, Best Practices for Basic 911 System Training, Training Manual
55. HEC Status 3/25/03, 4/3/03, 4/10/03, 8/17/03, 9/29/03
56. Altaris Status, 9/11/03
57. Houston Altaris® CAD Call taker Train-the-Trainer Schedule
58. Alartis® Computer Aided Dispatch System and Records Management System Project Implementation Plan, December 14, 2001
59. Memorandum of Understanding – CAD Functional Acceptance Testing
60. HEC FSD Evaluation Exceptions Identified 3/29/2002
61. CAD Failover Load Test Report, July 15-16, 2003, Performance Certification
62. HEC Polices and Procedures, January 20, 2005
63. Altaris® CAD Programmer Training Materials
64. Altaris® CAD Call Taker and Dispatch Training Manuals
65. Altaris® Cad Initial System Configuration
66. Altaris® CADLIVE, INTLIVE, MISLIVE Data Dictionary
67. Altaris® CAD and MSS As Built Documentation
68. SIRT List, All Items
69. Change Order List, All Items
70. Altaris® Computer Aided Dispatch System and Records Management System Project Implementation Plan, December 14, 2001
71. Altaris® CAD Command Statistic Report for 2005, January 12, 2005

Appendix B Operations of Call Takers and Dispatchers

Figure B-1 shows the operation of call takers and dispatchers. The Neutral 911 call takers are an initial entry point to the system. They classify a call as going to Police or Fire/EMS, or refer it to another agency. They transfer the caller to either a Police or Fire/EMS call taker, referred to as a "warm-transfer." Combined events, those requiring both Fire/EMS and Police response, are transferred to Fire/EMS call takers. Neutral 911 call takers do not interact with the technical CAD system, but they do use the VESTA call management system.

HPD and HFD call takers are the interface to the public requesting services. They obtain, organize, and enter the information that is the basis for making resource decisions. They define the call type and priority, "coding" the call. Some aspects of the call taking requirements are explicitly incorporated into the CAD information entry system under the Special Instructions (SINS) feature, but there are marked differences between Police and Fire/EMS usage of that tool.

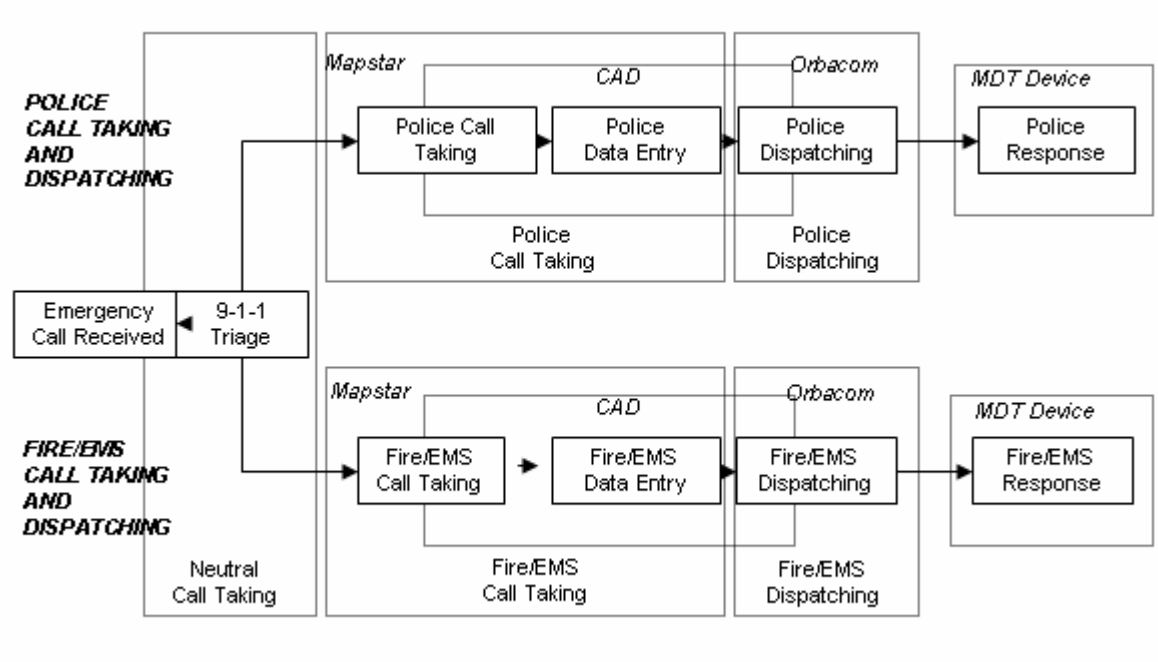

Figure B-1. Operations of Call Takers and Dispatchers

Upon completion of the basic entry of an event, the information is then passed to a dispatcher to assign, monitor, and manage a response. At this point the police call taker may terminate the call, but the Fire/EMS call takers may have responsibility for the delivery of "pre-arrival instructions," the coaching of the caller to take medical action prior to arrival of a medical unit. The dispatch operations between the HFD and HPD have important distinctions. The diversity of response possibilities is larger on the Fire/EMS side, choosing among types of equipment and possible combinations (engines, ladder, tower, BLS ambulances, ALS ambulances, Paramedic vehicles, command officers). There is automatic support by the CAD system for selection of asset combinations. The continuous service delivery from the dispatcher is limited, with no direct involvement in safety issues.

The variety of police dispatch choices are typically much more limited in terms of type of response, although some specialized unit selection is occasionally involved. Practically, if not officially defined, the police dispatchers do provide some degree of load management for the officers in the field, making certain that the load on the officer is not driven by a simple "closest officer" algorithm that might overload one officer. There is a very important continuous service connection from the officer to the dispatcher that is unique to the police side of dispatch.

The analysis of the operations against the initial system design showed major differences and expectations. The roll out of the new CAD system was expected to have minimal disruption to the police call taking and police dispatching processing. The expectation by police dispatchers was that the system would be modified to fit their existing police dispatch processes and that departmental policies and procedures would not be affected. This is consistent with the terms of the acquisition of the new system as an upgrade to the existing police CAD system. In contrast to this view, the Fire/EMS participated in the new CAD system project with the expectation that the implementation of the new CAD system would result in changes to their business processing, but like the police dispatching staff, no expectations existed for changes in the departmental reporting structure or impacts to their existing departmental policies and procedures.

The newly formed HEC organization took on the role of integrating call taking and dispatching business processes in anticipation that a single call taking and call dispatching process and procedure would be followed when utilizing the new CAD system. Furthermore, the HEC organization began to work towards standardizing staff policies and procedures as they saw their role as a "service organization" to the Houston Police and Fire Departments with overall responsibility and accountability for Houston's Emergency Services' call taking and dispatching functions.

Appendix C System Availability Concepts and Calculations

Concepts and Definitions

A system has recurrent up periods (operating) and down periods (in maintenance/repair) in its life cycle. MTBF (Mean Time Between Failure) and MTTR (Mean Time To Repair) are two widely used statistics in availability theory to measure how frequent failure incidents are likely to occur and how fast a repair can be done.

Availability can be evaluated by two standard measurements:

- Operational Availability = Total Uptime/Assessment Period = 1 - Total Downtime/ Assessment Period, where the computation takes into consideration all corrective repair times, preventive maintenance times, and administrative and logistics delay times. This is assessed from end-users' perspective: whenever the system cannot be used due to either planned or unplanned events, the system is viewed as unavailable.
- Inherent Availability = MTBF/(MTBF + MTTR), where the computation excludes preventive maintenance times and administrative and logistics delay times. Also known as intrinsic availability, this measurement based on only failure outages that required corrective repairs, is basically reflecting the system reliability and the ability to recover from failures. MTBF is estimated by total assessment period divided by the number of outages. MTTR is estimated by total repair time divided by the number of repairs.

More details on relevant concepts and definitions can be found in:

U.S. Department of Defense Handbook 3235.1-H "*Test & Evaluation of System Reliability, Availability, and Maintainability*", 1982.

U.S. Department of Defense Handbook MIL-HDBK-338B "*Electronic Reliability Design Handbook,*" 1998.

Calculations of System Availability

If the system life cycle is considered to start from the first day when the system went live for conducting the acceptance test in a live operation environment, then the assessment start time was September 23, 2004 04:00:00, and the data of incidents B1 through B10 and A1 through A7 should all be considered for the availability calculation. If the system life cycle is considered to start from the system acceptance date, then the assessment start time was September 23, 2004, 04:00:00, and only the data of incidents A1 through A7 should be considered.

The parameters and calculations for operational availability are shown in Table C-1.

Table C-1. Operational Availability Calculations

		Assessment Period (hour)	Incidents considered	Total Downtime (hour)	A_o = 1 - TD/AP
System life cycle started from go-live date	Overall System	11924.00	B1 – B10 A1 – A7	41.90	**0.9965**
	CAD/RMS	11924.00	B1 – B10 A1 – A3	23.73	**0.9980**
System life cycle started from acceptance date	Overall System	9480.00	A1 – A7	34.25	**0.9964**
	CAD/RMS	9480.00	A1 – A3	16.08	**0.9983**

When calculating inherent availability, the last two incidents (A6 and A7) classified as preventive maintenance are not counted.

The parameters and calculations for inherent availability are shown in Table C-2.

Table C-2. Inherent Availability Calculations

		MTBF (hour)	Incidents considered	MTTR (hour)	A_i = MTBF/ (MTBF+MTTR)
System life cycle started from go-live date	Overall System	794.93	B1 – B10 A1 – A5	2.38	**0.9970**
	CAD/RMS	993.67	B1 – B10 A1 – A2	0.89	**0.9991**
System life cycle started from acceptance date	Overall System	1896.00	A1 – A5	5.62	**0.9970**
	CAD/RMS	4740.00	A1 – A2	1.54	**0.9997**

Calculations of Confidence Limit for System Availability

Assume the times between failure and the repair times all have exponential distributions. For the inherent availability A_i, it can be shown that the $(1-\alpha)$ one-sided confidence interval is given by:

$$P\left(A_i \geq \frac{\theta}{\theta + \phi \times F_{1-\alpha,2n,2n}}\right) = 1 - \alpha\,,$$

where n is the number of failures, θ is the estimated MTBF, ϕ is the estimated MTTR, and $F_{1-\alpha,2n,2n}$ is the F-statistic such that $P\left(Z \leq F_{1-\alpha,2n,2n}\right) = 1 - \alpha$ for any random variable Z with an F-distribution.

Using the outage data for the overall system since the go-live date, the confidence limit and corresponding confidence level for the inherent availability A_i are computed and tabulated in Table C-3. The confidence limit can be interpreted as the availability target, and the confidence level indicates the possibility for reaching that target.

Table C-3. Confidence Limit for Inherent Availability A_i of the Overall System Since Go-Live Date

$A_i \geq$ Confidence Limit	Confidence Level
0.9991	0.05%
0.9989	0.50%
0.9984	5.00%
0.9970	50.00%
0.9952	90.00%
0.9945	95.00%

The first row reads: Probability $(A_i \geq 0.9991) = 0.0005$. This means there is extremely low confidence (0.0005) that the Inherent Availability of HEC could reach 0.999. Usually, analysis of a reliability model for the architecture can help identify which components would contribute the most to the overall system unavailability.

The probability expression $\Pr(A_i \geq a) = p$ is equivalent to $\Pr(A_i < a) = 1 - p$. Thus, the third row indicates we are 95% confident that the inherent availability of HEC is lower than 0.9984. That means we can predict with 95% confidence that, if nothing is to be improved, the overall HEC system downtime will be at least 14 hours per year.

For the Operational Availability A_o, there is no simple close form expression for representing the confidence level of A_o. Monte Carlo simulation was used for obtaining the approximated confidence limits and confidence levels, which are shown in Table C-4.

Table C-4. Confidence Limit for Operational Availability A_o of the Overall System Since Go-Live Date

$A_i \geq$ Confidence Limit	Confidence Level
0.9983	5%
0.9977	25%
0.9970	50%
0.9962	75%
0.9949	95%

These results for A_o are very similar to that for A_i.

The same methods are used to calculate the confidence limits for the availability of CAD/RMS alone. The same calculations are repeated for the assessment period that started after the acceptance. These results are discussed in Section 4.2.5.

Analysis of the Tradeoffs Between Reliability and Maintainability

For the 0.999 availability, reducing just one hour in repair time will be as effective as adding 41 days of uptime between two failures. (Whether this approach is more economical in the long run will be subject to further tradeoff analysis, taking into account of an additional set of criteria including finance, support goals, and other relevant factors.) Improving MTTR has better leverage than improving MTBF for increasing the availability value.

Figure C-1 shows the estimated MTBF calculated after each incident cycle after the system go-live date. For example, the third data point is calculated as follows: dividing the total elapsed time until the end of the third incident by three. This chart indicates that the MTBF is getting better (longer) but is not yet reaching a steady state, implying that the integrated public safety data system has not passed the so-called "infant mortality" stage.

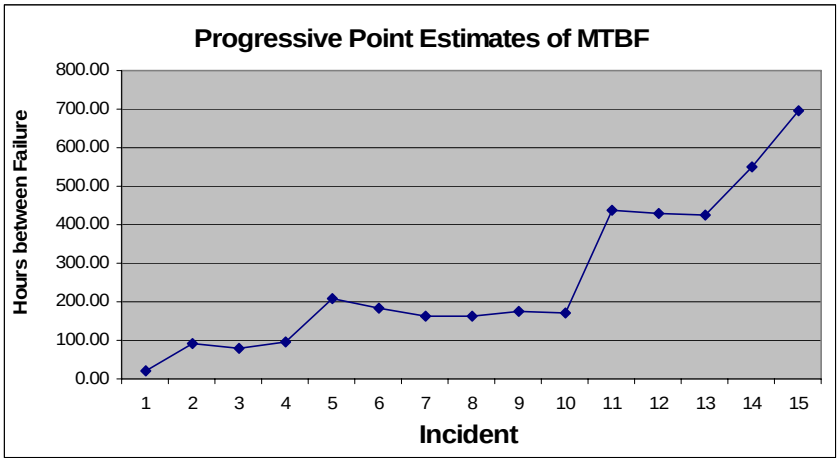

Figure C-1. Progressive Point Estimates of MTBF

Not reaching a steady state is not a bad sign. On the contrary, a bad sign would be when the MTBF has already come to a steady state but is stuck with an undesirable MTBF, such as 700 hours between failures (i.e., approximately one failure per month), for the rest of the system life cycle before the next major upgrade is acquired.

For assessing how frequent failures would occur, it is more accustomed to calculating the failure rate, which is defined as the reciprocal of MTBF, i.e., the number of failures per unit time. Figure C-2 displays the estimated failure rate calculated after each incident cycle. Again, it can be seen that the failure rate has not yet reached a steady state and it is expected to improve further.

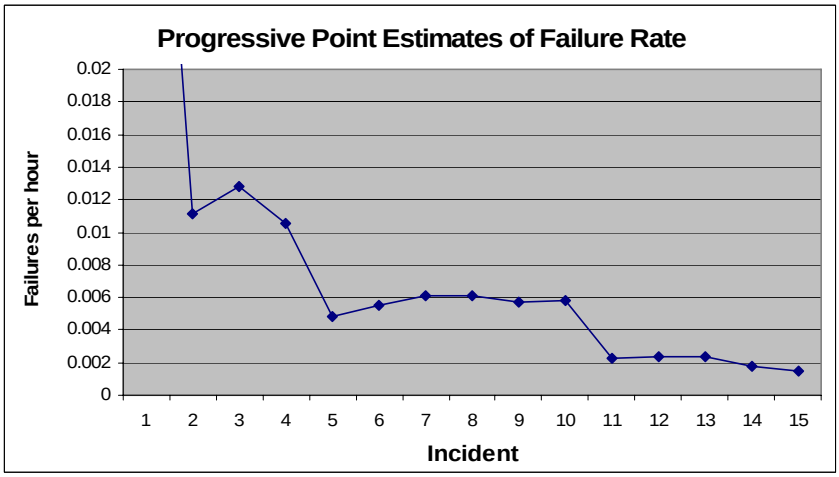

Figure C-2. Progressive Point Estimates of Failure Rate

The DoD Military Handbook MIL-HDBK-338B for Electronic Reliability Design uses the Software Reliability Curve in Figure C-3 to represent the reliability in the life cycle of a typical software system. Clearly, the HEC system has not yet reached the end of Period A on this curve.

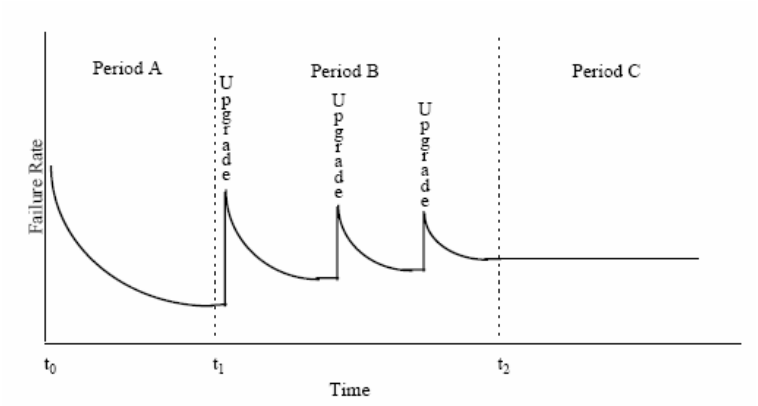

Figure C-3. Software Reliability Curve (from MIL-HDBK-338B)

To improve system and component reliability will require rework of system architecture and integration, which in turn will require a significant amount of resources and time to accomplish. Nevertheless, the single point of-failure identified in Section 4 should be eliminated, other components, and systems can be replaced over time.

Appendix D HEC Call Volume Statistics

Each chart below covers a one month interval within the assessment period from September 2003 to December 2004. The call volume value includes all calls for Fire, EMS, and Police events. Each data point is the call volume within the corresponding hour. Each triangle overlain on the chart indicates the start time of an unplanned outage. The two outages of scheduled maintenance in December 2004 are not shown. The first and the last months have data only for partial months.

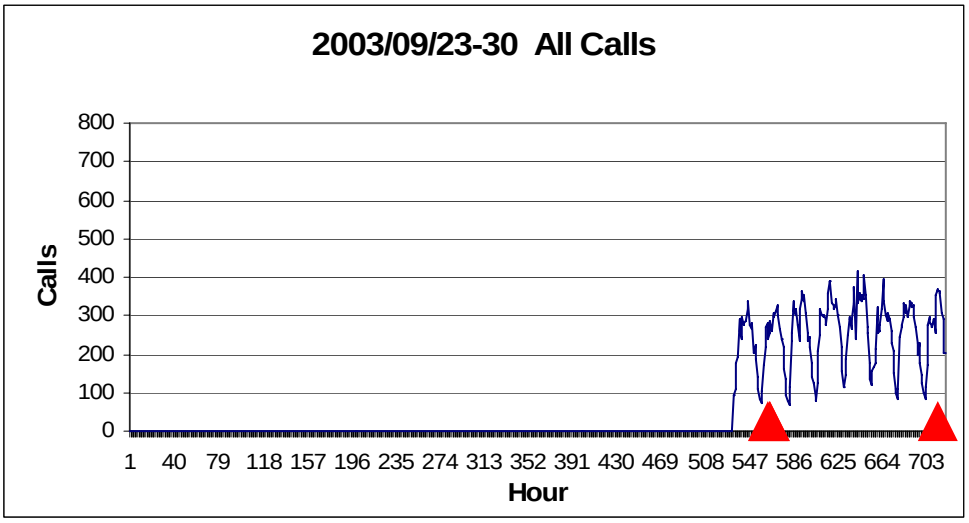

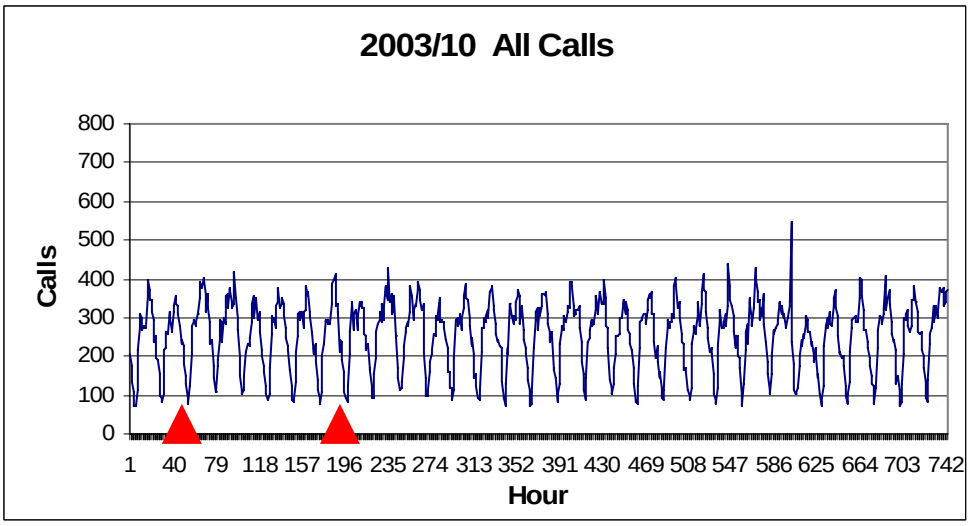

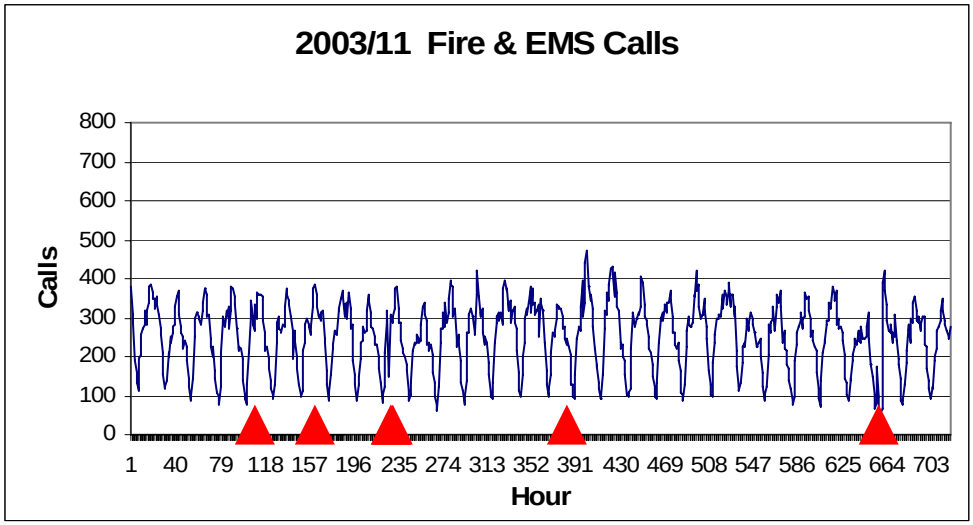
2003/11 Fire & EMS Calls
Calls
800
700
600
500
400
300
200
100
0
1 40 79 118 157 196 235 274 313 352 391 430 469 508 547 586 625 664 703
Hour

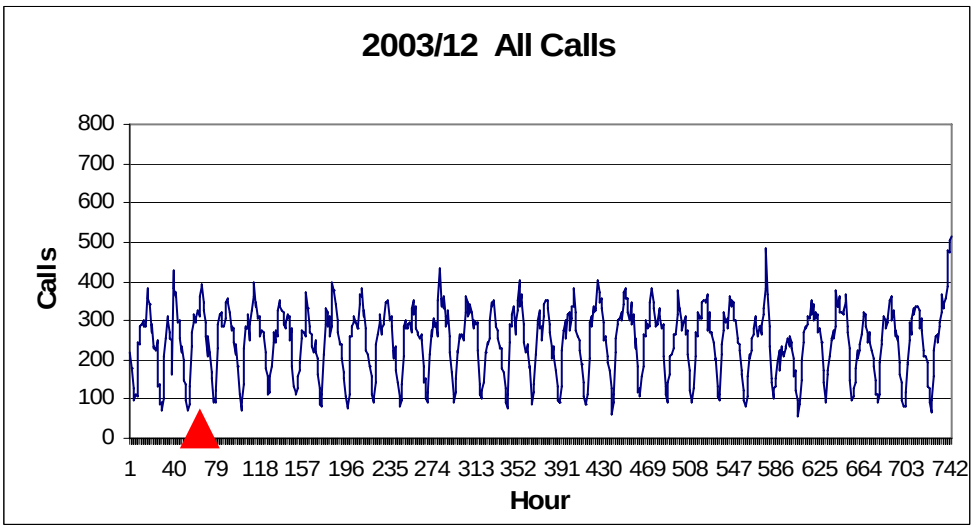
2003/12 All Calls
Calls
800
700
600
500
400
300
200
100
0
1 40 79 118 157 196 235 274 313 352 391 430 469 508 547 586 625 664 703 742
Hour

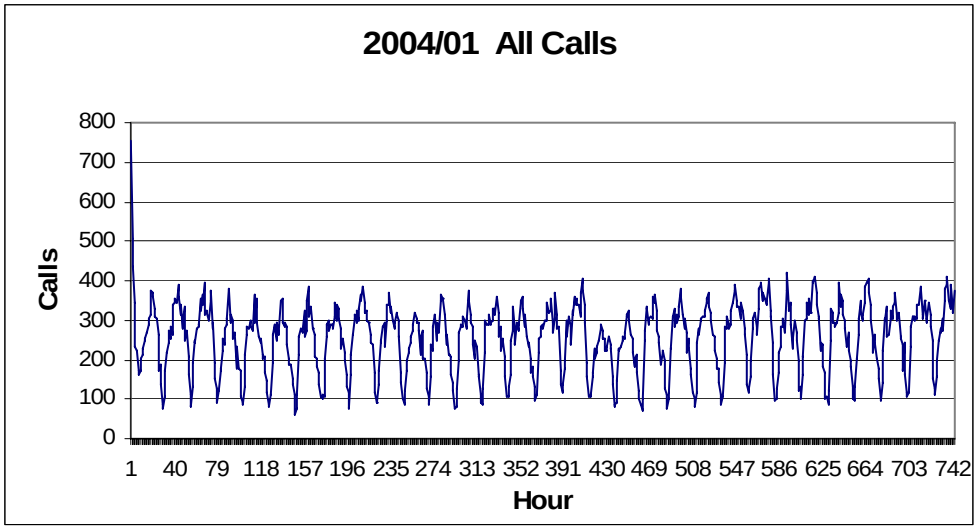
2004/01 All Calls
Calls
800
700
600
500
400
300
200
100
0
1 40 79 118 157 196 235 274 313 352 391 430 469 508 547 586 625 664 703 742
Hour

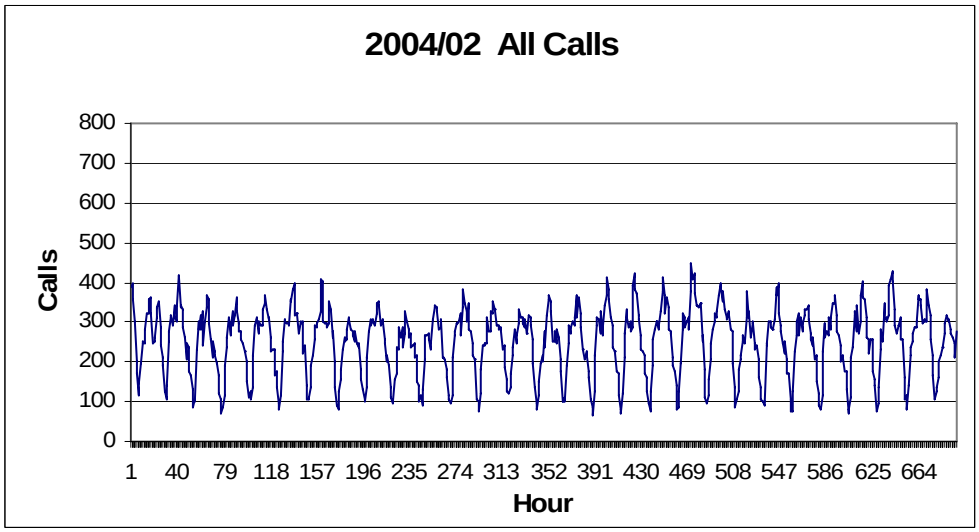
2004/02 All Calls
Calls
800
700
600
500
400
300
200
100
0
1 40 79 118 157 196 235 274 313 352 391 430 469 508 547 586 625 664
Hour

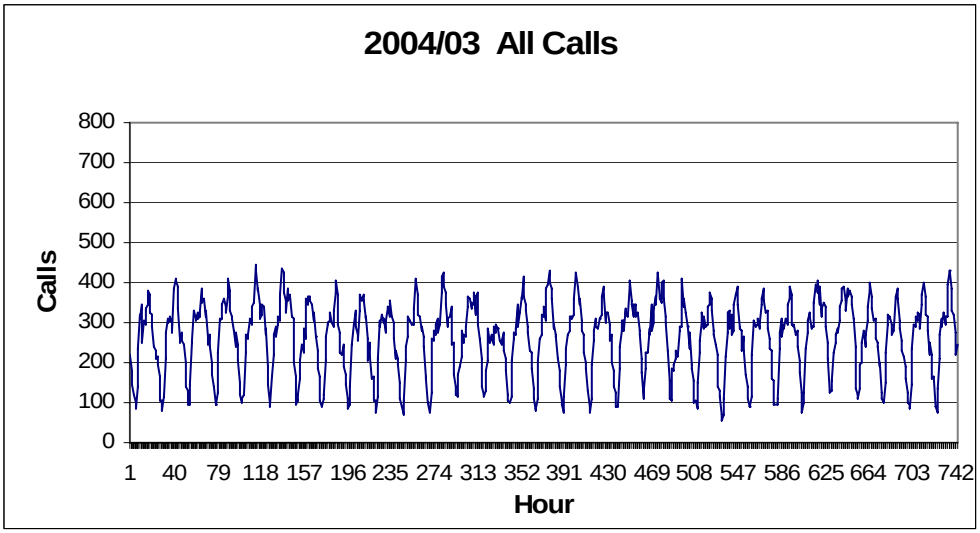
2004/03 All Calls
Calls
800
700
600
500
400
300
200
100
0
1 40 79 118 157 196 235 274 313 352 391 430 469 508 547 586 625 664 703 742
Hour

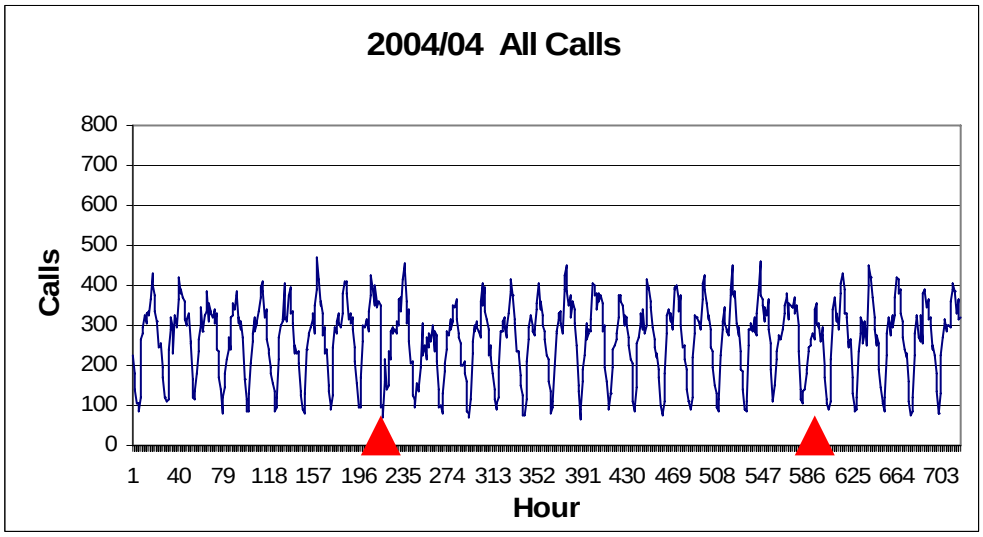
2004/04 All Calls
Calls
800
700
600
500
400
300
200
100
0
1 40 79 118 157 196 235 274 313 352 391 430 469 508 547 586 625 664 703
Hour

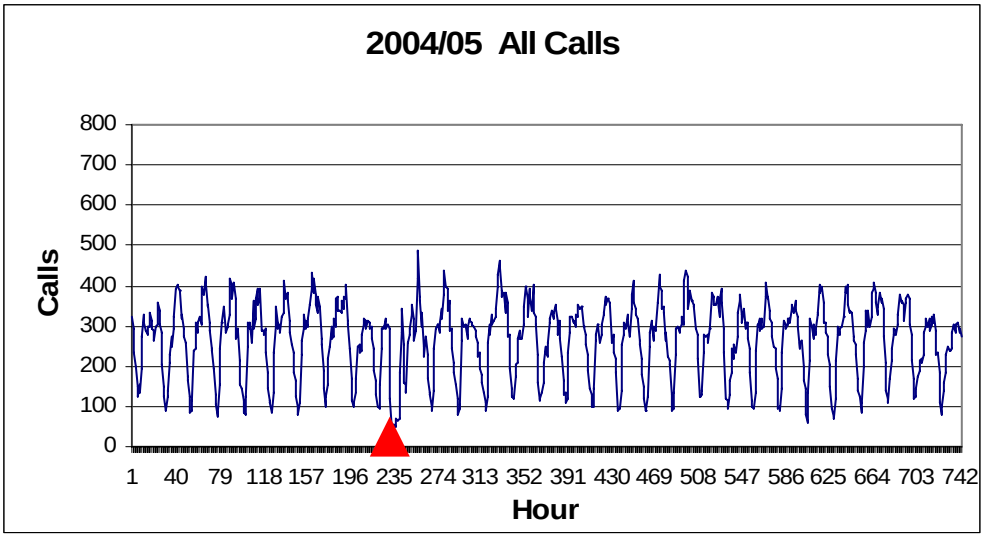
2004/05 All Calls
Calls
800
700
600
500
400
300
200
100
0
1 40 79 118 157 196 235 274 313 352 391 430 469 508 547 586 625 664 703 742
Hour

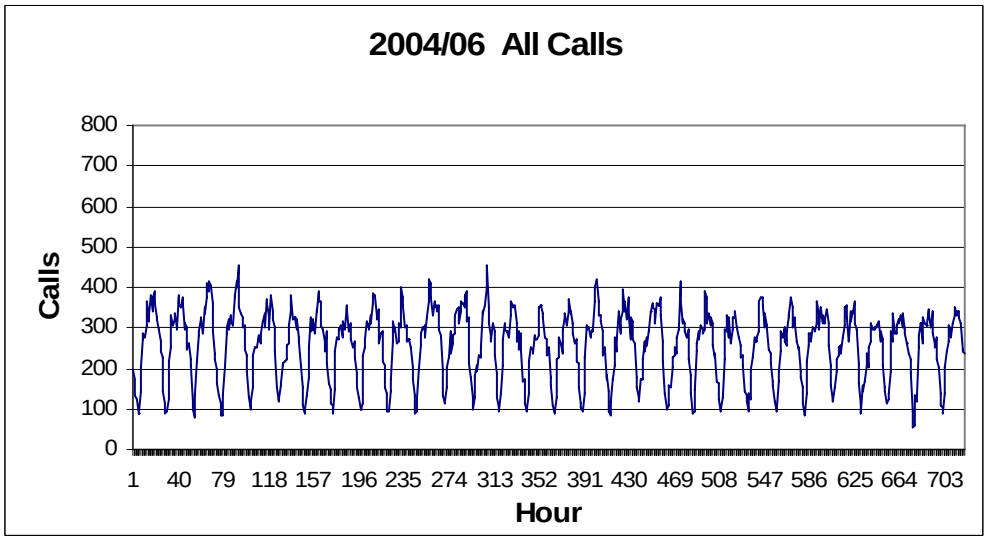
2004/06 All Calls
Calls
800
700
600
500
400
300
200
100
0
1 40 79 118 157 196 235 274 313 352 391 430 469 508 547 586 625 664 703
Hour

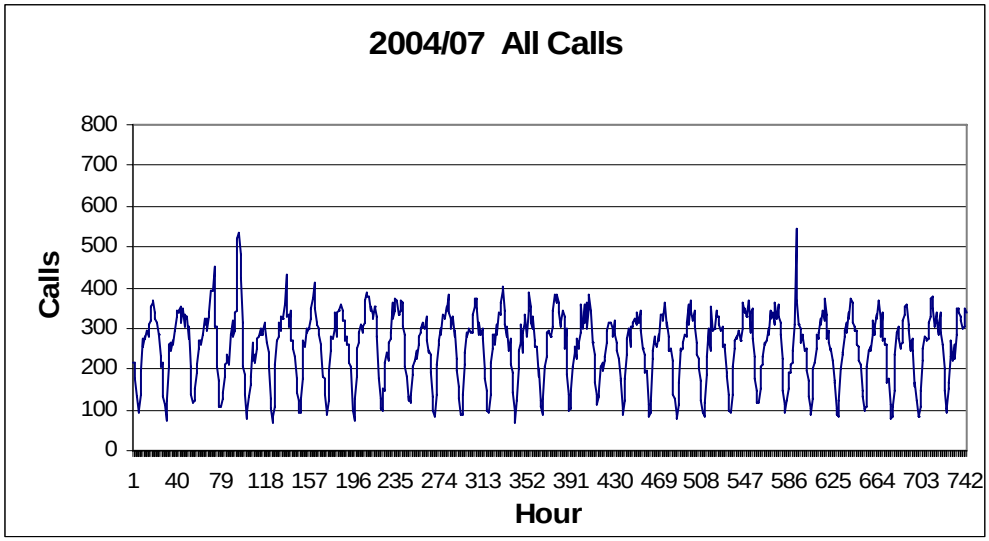
2004/07 All Calls
Calls
800
700
600
500
400
300
200
100
0
1 40 79 118 157 196 235 274 313 352 391 430 469 508 547 586 625 664 703 742
Hour

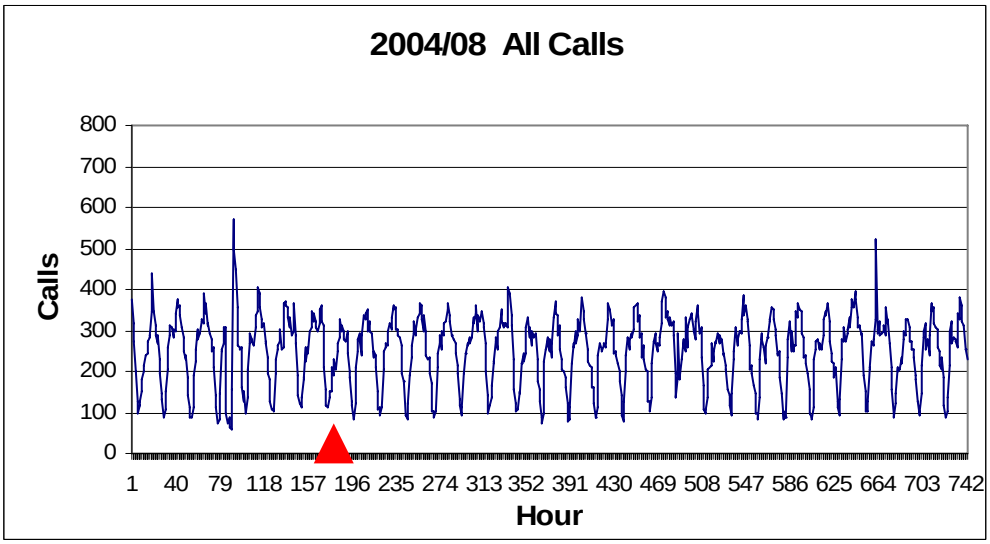
2004/08 All Calls
Calls
800
700
600
500
400
300
200
100
0
1 40 79 118 157 196 235 274 313 352 391 430 469 508 547 586 625 664 703 742
Hour

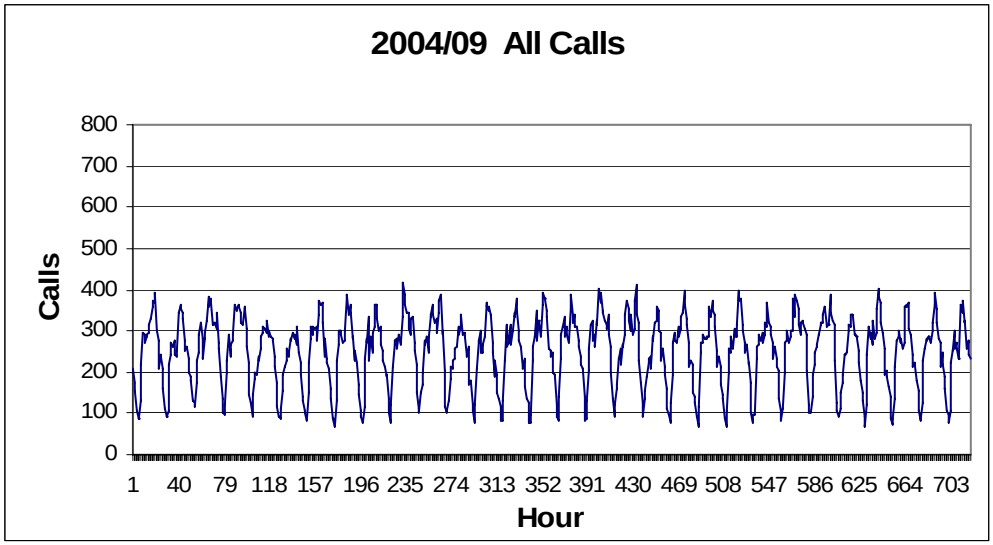
2004/09 All Calls
Calls
800
700
600
500
400
300
200
100
0
1 40 79 118 157 196 235 274 313 352 391 430 469 508 547 586 625 664 703
Hour

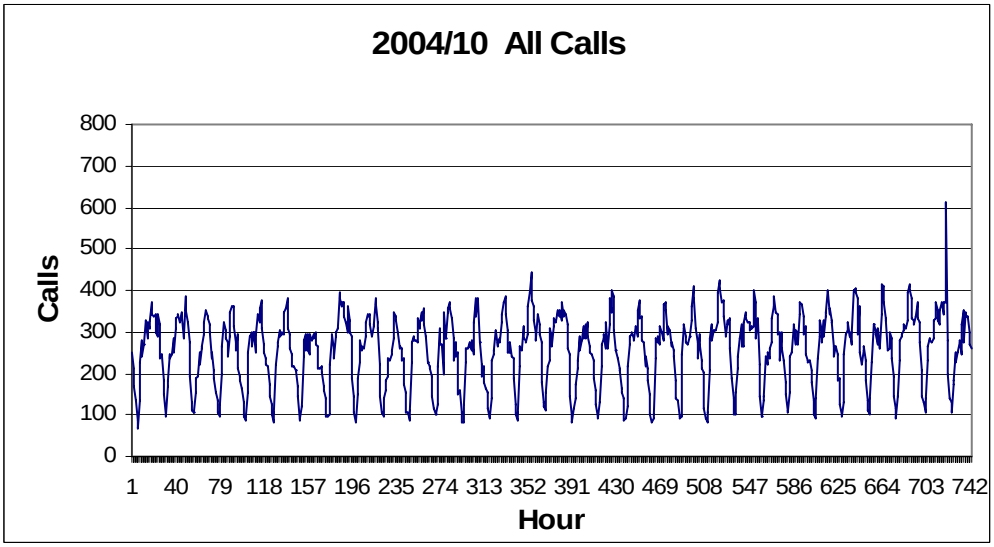
2004/10 All Calls
Calls
800
700
600
500
400
300
200
100
0
1 40 79 118 157 196 235 274 313 352 391 430 469 508 547 586 625 664 703 742
Hour

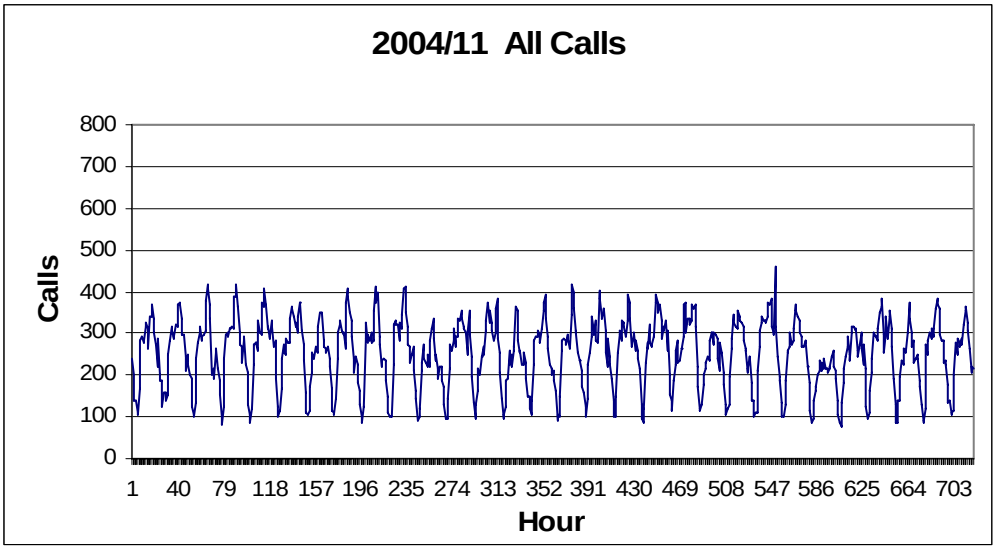

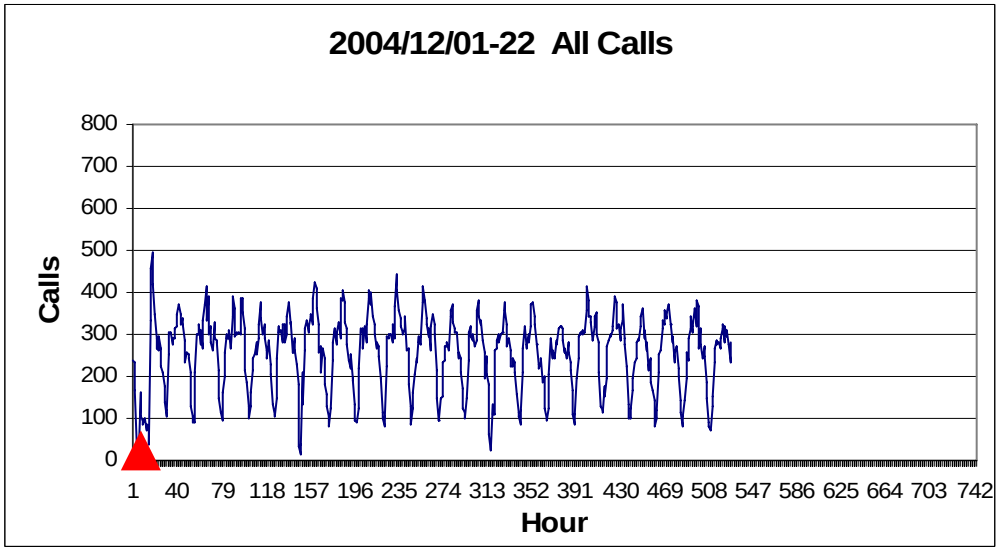

Call volume statistics from 23 September 2003 to 22 December 2004:

Max =	754
Median =	274
Average =	256
Min =	15

Distribution List

Internal

F. G. McLoughlin

M. L. Minter

D. R. Moody

A. G. Moore

R. F. Nesbit

M. S. Rosen

A. M. Shoemaker

T. J. Woodhouse

J. P. L. Woodward

Project

R. Blount, Jr. (10 copies)

S. F. Chang

F. Galdos

S. P. Morrissey

S. Nakamoto

D. P. Shaffer

External

Sharon Counterman, Houston Emergency Center (20 copies)

Barbara Evans, Houston Emergency Center

Roy Morales, Houston Emergency Center

Rene Gonzales, DC Services

Brian Jones, L. Robert Kimball & Associates, Inc.

Bryan Smith, L. Robert Kimball & Associates, Inc.

Rick Taormina, L. Robert Kimball & Associates, Inc.