**SANDIA REPORT**
SAND2003-1897
Unlimited Release
Printed June 2003)

# Using Modeling and Simulation to Analyze Application and Network Performance at the Radioactive Waste and Nuclear Material Disposition Facility

Roy A. Life, Joseph H. Maestas, Dennis Bateman

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,
a Lockheed Martin Company, for the United States Department of Energy's
National Nuclear Security Administration under Contract DE-AC04-94AL85000.



Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.



Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865)576-8401
Facsimile: (865)576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: http://www.doe.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800)553-6847
Facsimile: (703)605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2003-1897
Unlimited Release
Printed June 2003

# Using Modeling and Simulation to Analyze Application and Network Performance at the Radioactive Waste and Nuclear Material Disposition Facility

***Roy A. Life, Joseph H. Maestas, and Dennis Bateman***
Telecommunication Operations Department
Sandia National Laboratories
PO Box 5800
Albuquerque, NM 87185-0812I

***Abstract***

Telecommunication services customers at the Radioactive Waste and Nuclear Material Disposition Facility (RWNMDF) have endured regular service outages that seem to be associated with a custom Microsoft Access Database. In addition, the same customers have noticed periods when application response times are noticeably worse than at others. To the customers, the two events appear to be correlated. Although many network design activities can be accomplished using trial-and-error methods, there are as many, if not more occasions where computerized analysis is necessary to verify the benefits of implementing one design alternative versus another. This is particularly true when network design is performed with application flows and response times in mind. More times than not, it is unclear whether upgrading certain aspects of the network will provide sufficient benefit to justify the corresponding costs, and network modeling tools can be used to help staff make these decisions. This report summarizes our analysis of the situation at the RWNMDF, in which computerized analysis was used to accomplish four objectives: 1) identify the source of the problem; 2) identify areas where improvements make the most sense; 3) evaluate various scenarios ranging from upgrading the network infrastructure, installing an additional fiber trunk as a way to improve local network performance, and re-locating the RWNMDF database onto corporate servers; and 4) demonstrate a methodology for network design using actual application response times to predict, select, and implement the design alternatives that provide the best performance and cost benefits.

**Acknowledgements**
The authors would like to thank the many individuals who provided assistance throughout this project. Of particular note are:

Anne Van Arsdall, Telecommunication Operations Department, for providing editorial assistance.

Dianna Muller, Radioactive Waste Nuclear Material Disposition Department, for her patience, perseverance, and assistance while we obtained information needed to perform the analysis.

Michael Rios, Telecommunication Operations Department, for reviewing this document, for providing the documentation about the network at the Radioactive Waste and Nuclear Material Disposition Facility and for providing funding for the subsequent network upgrades.

Diana Eichert, Cyber Enterprise Management and Process Integration, for providing assistance with SNMP and MRTG setup for collecting link utilization statistics, which we were able to import as background utilization statistics.

Bobby Weaver, Tonopah Test Range Department, for providing information about the wireless infrastructure between Tech Areas 3 and 5 and for providing assistance with setting up the Ravlin encryptors for SNMP-based management.

# Table of Contents



# List of Figures

## 1.0 Introduction

Sandia National Laboratories Telecommunication Operations Department is developing a network modeling and simulation capability for diagnosing network and application performance issues and for evaluating alternatives for major designs of network infrastructure.

Ideally, this analytical work will be performed before deploying any major change in the corporate network. The network at the Radioactive Waste and Nuclear Material Disposition Facility (RWNMDF) is an example of an existing network primarily implemented with legacy Ethernet hardware from the pre-switch era. During intervals of peak utilization, it is beset with poor performance, which underscores the fact that the existing infrastructure will not support growth. The RWNMDF network marks an area at Sandia where a combination of encryption, microwave, and legacy shared Ethernet elements are used to interconnect to the Sandia Restricted Network (SRN). Staff at RWNMDF reside in several buildings and commonly use a centralized Microsoft Access® database application residing in a single server.

In modern networks, clients, servers, application traffic (e.g., mix, type, pattern) and network infrastructure all contribute to the perception of a performance issue. In environments where shared network resources are stressed, each of these components becomes especially crucial in the performance equation. This paper looks at the contributions of each component and identifies those areas where improvements will make the most impact. Improvements in the network can be achieved in a number of ways, and several ideas have been exchanged. Without additional analysis, it would be unclear whether upgrading certain aspects of the RWNMDF network would provide sufficient benefit to justify the corresponding costs; we report the results of such an analysis here. Modeling and simulation tools, techniques, and SNMP-based network monitoring tools will be used to ascertain existing network utilization, identify problem areas, and to evaluate alternative designs. Actual traffic flows in the existing network will be taken into account to ensure that the new design and recommendations offer the users at the Radioactive Waste and Nuclear Material Disposition Facility visible improvements moving forward.

## 2.0 State of the RWNMDF Network Prior to Re-design

The Radioactive Waste and Nuclear Material Disposition Facility comprises several buildings where staff resides. It shares a single common subnet with approximately 166 active registered hosts, of which an estimated 75 are still active today and use the network facility on a regular basis. The current network connecting the facility was first implemented in the early-to-mid nineties, prior to the introduction of switched Ethernet at Sandia. As a result, hosts attach to network via segmented hubs, and in the worst cases, via cascaded hubs. Most hosts attach via the second tier of hubs.

A profile of network usage at the RWNMDF includes corporate email, corporate web access, and corporate data storage facilities. Users at this facility also commonly access an internal database based on the Microsoft Access product, which is essential to the daily business operations at this facility. Users typically log onto the database and remain connected, although they may not be generating transactions. These users commonly report that they are not able to establish a connection to the database; that is, the connections are rejected, and this occurs at least once a day. Repairing the database, which takes thirty minutes to complete, clears the condition and then users are once again able to use the database.

The analytical tools used during this study do not lend themselves, specifically, to trouble shooting database problems. Therefore, it is difficult to say with certainty the cause of database corruption. It is very unlikely that a heavily congested network would contribute to the corruption since the Transmission Control Protocol (TCP) ensures that the application receives data that is free from errors and is in the proper order. It seems more likely that the server or the application itself causes the corruption. Figure 1 represents the network implementation at the RWNMDF before the upgrade. Note that the figure does not reflect the proper number of workstations attached to the network, nor the proper room locations for the hubs, but serves to establish the general idea.

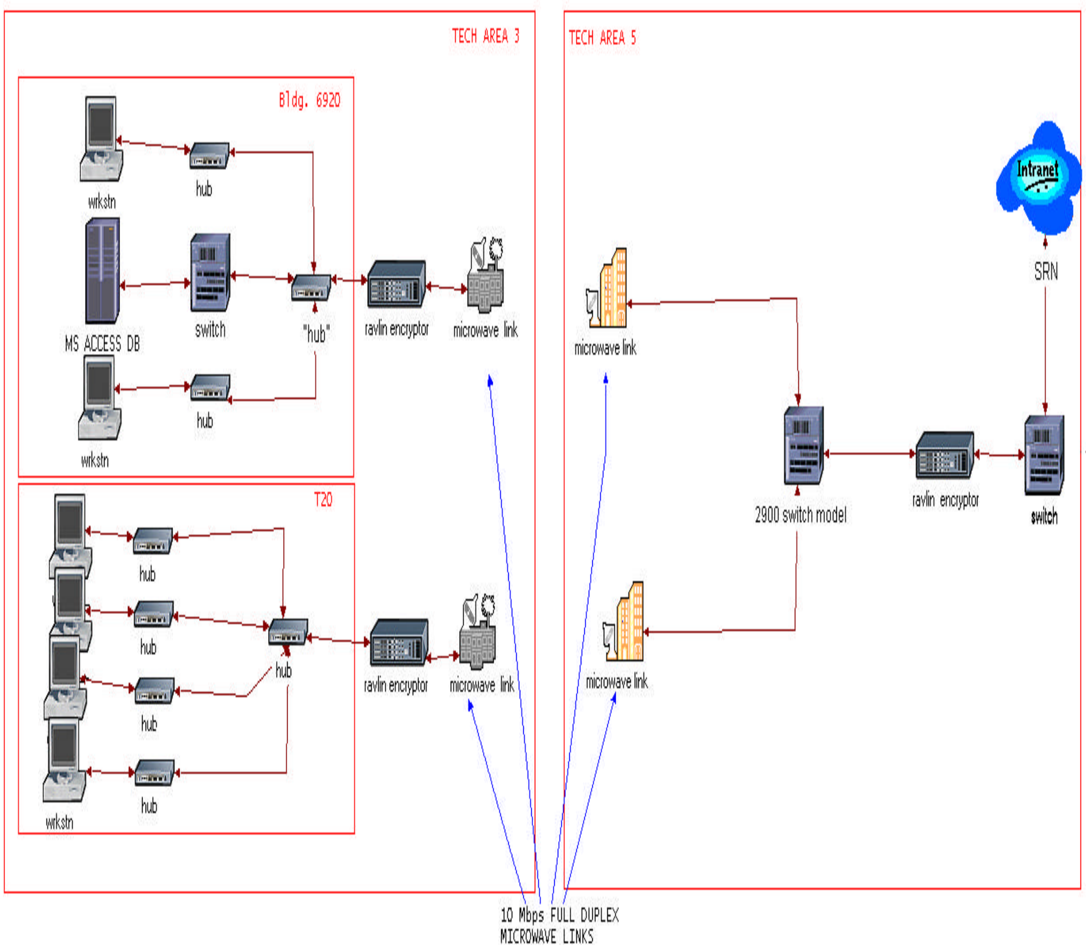


***Figure 1: RWNMDF's Current Network (before upgrade)***

To get a glimpse of the network activity at the RWNMDF, a Network Associates Sniffer®, hereafter referred to as Sniffer, was attached to a hub. In the figure, this element is labeled "hub." Additionally, packet capture agents were installed in the server and in a locally attached client workstation. A user case session between the server and the client was recorded that consisted of the following tasks:

1.) Log in to the database server
2.) Search for an entry, and look at some data
3.) View report that does calculations, query, open a report, print 2 pages, close the report
4.) Look for label, find label, view label, and print label
5.) Add new record to table, save, close the form
6.) Close all forms and log out of database

The packet captures were imported into the Application Characterization Environment software package, ACE®, by OPNET Technologies, where application characterization and analysis is performed.

Figure 2 reveals the summary of delays for a specific task in the use case. It was generated by the AppDoctor Summary of Delays® application within ACE. The purpose of AppDoctor Summary of Delays is to decompose the total delay into various network and application categories. The size of each rectangle in Figure 2 reveals the relative amount of delay contributed by each component. From this graph, it is clear that delays due to network propagation, transmission delays, and protocol congestion are relatively small as compared to the application components.

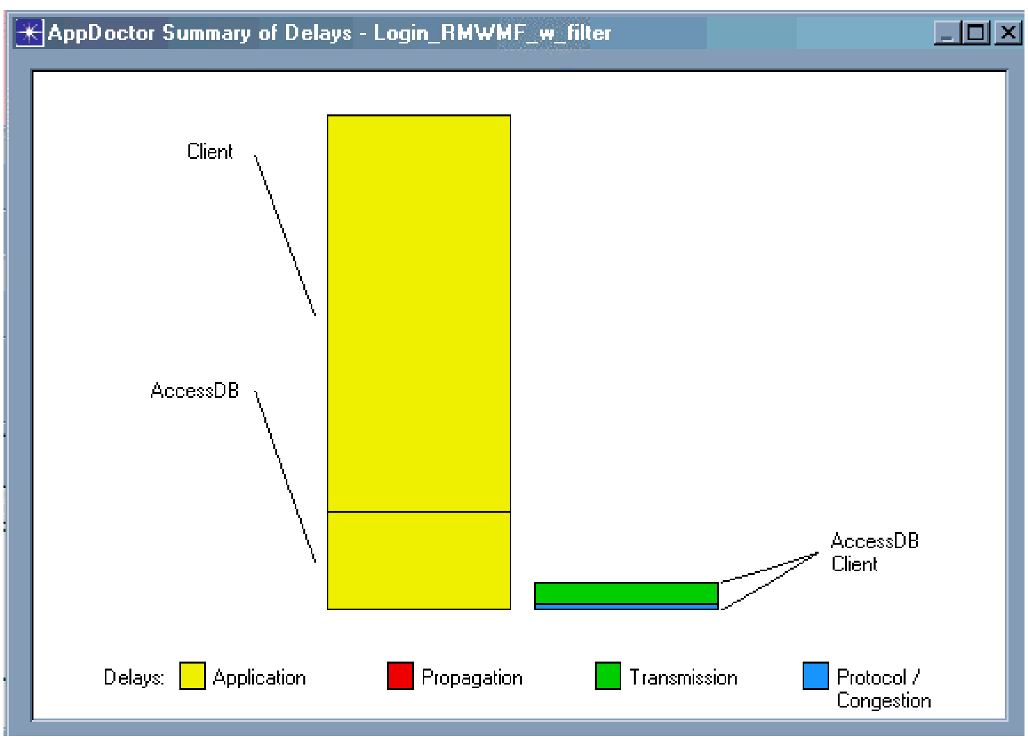


***Figure 2: Summary of Delays for the RWNMDF Network***

Table 1 summarizes the relative amount of delay contributed by each of the components depicted in Figure 1 during the completion of each task. It also identifies the areas where improvements make the most sense. Note that the client and server contribute most of the delay observed in each task; this is where improvements would make the biggest impact.

| All values in percentages (%) | Login | Search | Calculate | Print | Add a form | Logout |
|---|---|---|---|---|---|---|
| Application Response - Client | 76.1 | 86 | 87.2 | 84 | 95.5 | 80.8 |
| Application Response - Server | 18.6 | 7.2 | 7.3 | 6.4 | 2.4 | 14.4 |
| Transmission | 4 | 5.3 | 4.1 | 6.4 | 1.4 | 3.2 |
| Protocol / Congestion | 1.3 | 1.5 | 1.4 | 3.2 | 0.7 | 1.5 |
| Total | 100 | 100 | 100 | 100 | 100 | 100 |

***Table 1: Table of Delays for Database Tasks***

The AppDoctor Statistics window in Figure 3 provides a summary of metrics of application and network performance for the login task. One important metric for this study is application response time. Table 2 summarizes the application response times for each task.

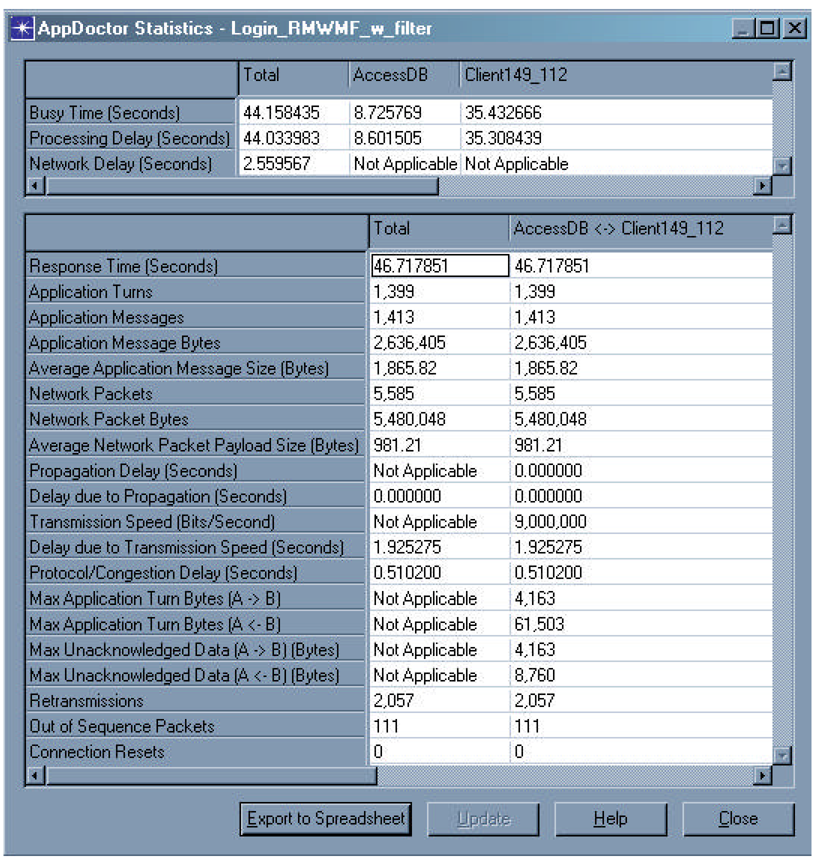


| | Total | AccessDB | Client149_112 |
|---|---|---|---|
| Busy Time (Seconds) | 44.158435 | 8.725769 | 35.432666 |
| Processing Delay (Seconds) | 44.033983 | 8.601505 | 35.308439 |
| Network Delay (Seconds) | 2.559567 | Not Applicable | Not Applicable |

| | Total | AccessDB <-> Client149_112 |
|---|---|---|
| Response Time (Seconds) | 46.717851 | 46.717851 |
| Application Turns | 1,399 | 1,399 |
| Application Messages | 1,413 | 1,413 |
| Application Message Bytes | 2,636,405 | 2,636,405 |
| Average Application Message Size (Bytes) | 1,865.82 | 1,865.82 |
| Network Packets | 5,585 | 5,585 |
| Network Packet Bytes | 5,480,048 | 5,480,048 |
| Average Network Packet Payload Size (Bytes) | 981.21 | 981.21 |
| Propagation Delay (Seconds) | Not Applicable | 0.000000 |
| Delay due to Propagation (Seconds) | 0.000000 | 0.000000 |
| Transmission Speed (Bits/Second) | Not Applicable | 9,000,000 |
| Delay due to Transmission Speed (Seconds) | 1.925275 | 1.925275 |
| Protocol/Congestion Delay (Seconds) | 0.510200 | 0.510200 |
| Max Application Turn Bytes (A -> B) | Not Applicable | 4,163 |
| Max Application Turn Bytes (A <- B) | Not Applicable | 61,503 |
| Max Unacknowledged Data (A -> B) (Bytes) | Not Applicable | 4,163 |
| Max Unacknowledged Data (A <- B) (Bytes) | Not Applicable | 8,760 |
| Retransmissions | 2,057 | 2,057 |
| Out of Sequence Packets | 111 | 111 |
| Connection Resets | 0 | 0 |

***Figure 3: Summary of Statistics of Network and Application Performance***

| Response Time (seconds) | Login | Search | Calculate | Print | Add a Form | Logout |
|---|---|---|---|---|---|---|
| Total | 46.717851 | 70.121963 | 70.711649 | 140.382565 | 105.302111 | 14.172999 |

***Table 2: Summary of Task Application Response Time** (Ref AppDoctor Statistics)*

Each task is further examined to determine the source of the delay. Figure 4 shows an example of the analysis for the login task. Three areas contribute to the processing delay observed for the login task: processing delay, protocol overhead, and retransmissions. Analysis of the other tasks revealed that bottlenecks were also recorded for processing delay, protocol overhead, and retransmissions. The next section discusses each of the areas that were identified as significant contributors to the application response time.

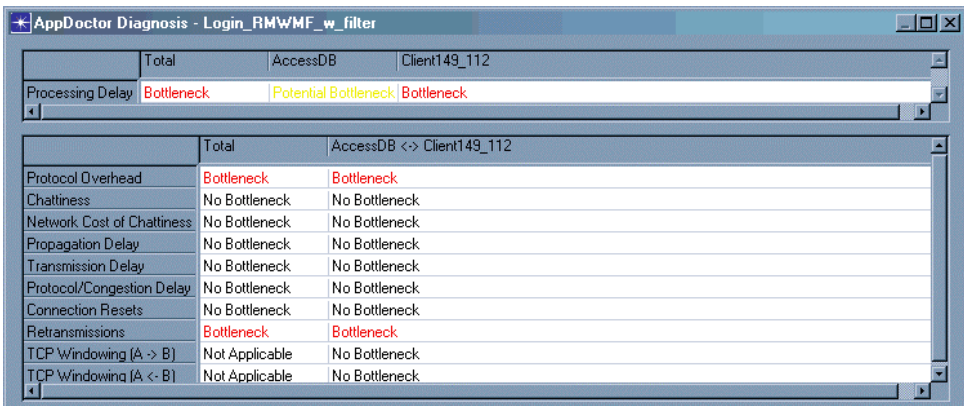


***Figure 4: Application Diagnosis of the RWNMDF Network***

**Processing Delay**

Excessive processing delays, indicated by the "Bottleneck" condition are observed at the client. The processing delays at the server are considered potential bottlenecks. Processing delays are due to file I/O, CPU processing, memory access, etc. They are considered bottlenecks when the percentage of delay exceeds thirty percent (30%) of the total application response time and "potential bottlenecks" when within 60% of this threshold. Processing delays at the server impact all users, whereas processing delays at the client usually only impact the user at the client location. The relatively larger delays observed at the client should not be over emphasized since user think-time is included in the delay calculation. The processing delay observed at the server includes the load imposed by 13 additional users accessing the database at the same time the packets were being captured for this study.

If the server at the RWNMDF were also being used to host other applications, this would also potentially impact the response time of the database application. We did not verify if this was the case. Additionally, there may be an internal problem with the database itself that is contributing to the processing delay at the server. The recurring need to repair the database may be indicative of this problem.

**Protocol Overhead**
Figure 4 also reveals that protocol overhead is a bottleneck. Bottlenecks due to protocol overhead indicate that at least 30% of the data being exchanged between the client and the database server is protocol overhead (protocol headers, TCP acknowledgements, etc.). An application that sends fewer, larger application messages between server and client will utilize resources more efficiently than an application that generates many more messages to send the same amount of data. The greater the percentage of protocol overhead in a data exchange, more time must be spent processing that overhead, which intuitively degrades application performance. This situation is not necessarily a problem in some applications. For example, in an application where timely and consistent feedback is provided to the client, such as would occur between a thin client and its terminal server, this behavior is both normal and expected. However, it would be a problem in an application whose purpose is to move large data sets between a client and its server. The users at the RWNMDF would have to decide which is the case for this application. If the data sets at the RWNMDF are large, then this problem is not one that can be solved by the user or by the database administrator. It is a problem with the commercial off-the-shelf application and one that the users will have to endure unless the application is replaced with something else.

**Retransmissions**
Figure 4 also reveals bottleneck conditions due to TCP retransmissions, which are indicative of corrupted Ethernet frames and is generally the result of one or more problems with the network infrastructure. The hubs in the network are the cause for the retransmissions. Hubs operate in half-duplex mode; that is, only a single host can transmit at any given time. All other hosts attached to a hub wait their turn to transmit, which is a random hit-and-miss event. Conversely, usually only a single host receives information at a time. Since a hub broadcasts all frames to all ports, essentially all hosts can potentially receive the information. Only the hosts that are configured to operate in promiscuous mode will receive all the information. However, by default, a receiver only receives frames that are destined to it, determined by the media access control (MAC) address on the host network interface card (NIC).

A MAC address in the frame that is not its own is blocked. Since only a single host can transmit at any given time through the hub, frames will collide as two or more hosts attempt to transmit. When this occurs, each host detects the collision, then either attempt to retransmit the frame immediately if signal energy is not detected on the bus or backs off for a random amount of time (exponential back off) before attempting to transmit again. All hosts attached to a single segmented hub are part of the same collision domain. When the physical dimension of the domain exceeds the physical limit, collisions are not detected within the required time period, which further degrades the ability to successfully transmit packets [1]. This is the fundamental reason why hubs should not be cascaded. The physical size of the collision domain can be erroneously expanded beyond the ability to successfully detect collisions within the required time period. Care must be exercised when hubs are cascaded. Clearly, with an increase in the utilization of the shared network, comes an increase in the number of frame collisions.

For the RWNMDF, Figure 5 and Figure 6 demonstrate the dependencies between these events. That is, during the busy hour intervals the number of collision errors increase.[Ó]

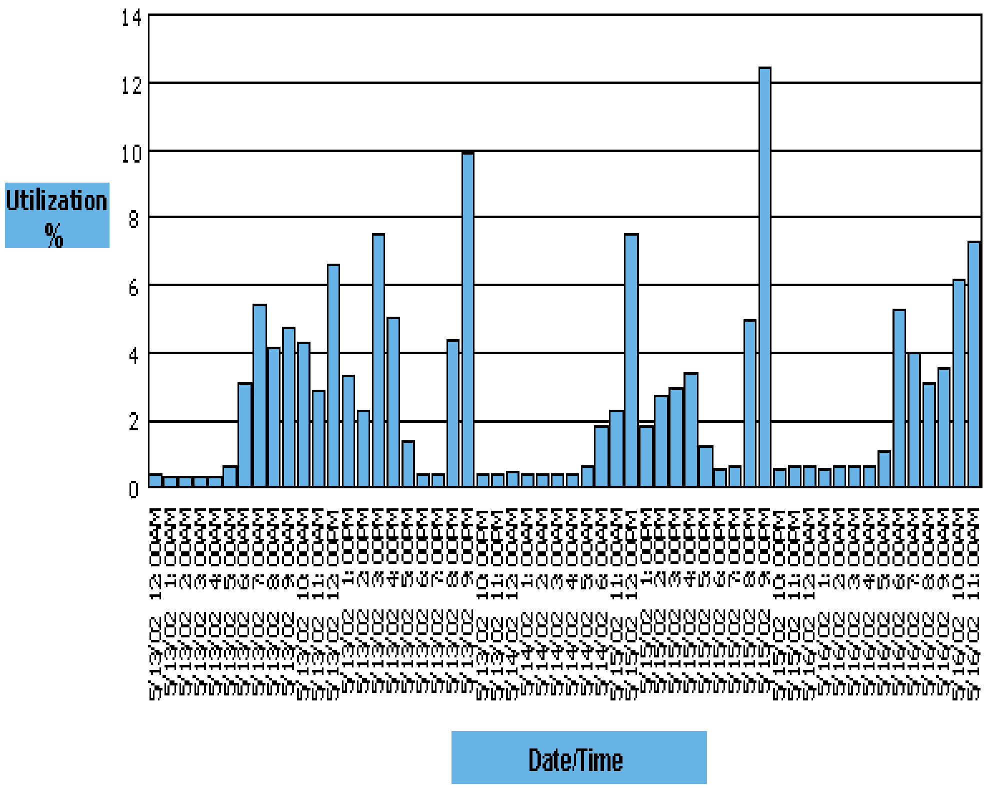


*Figure 5: 4-day Utilization of the RWNMDF Network*

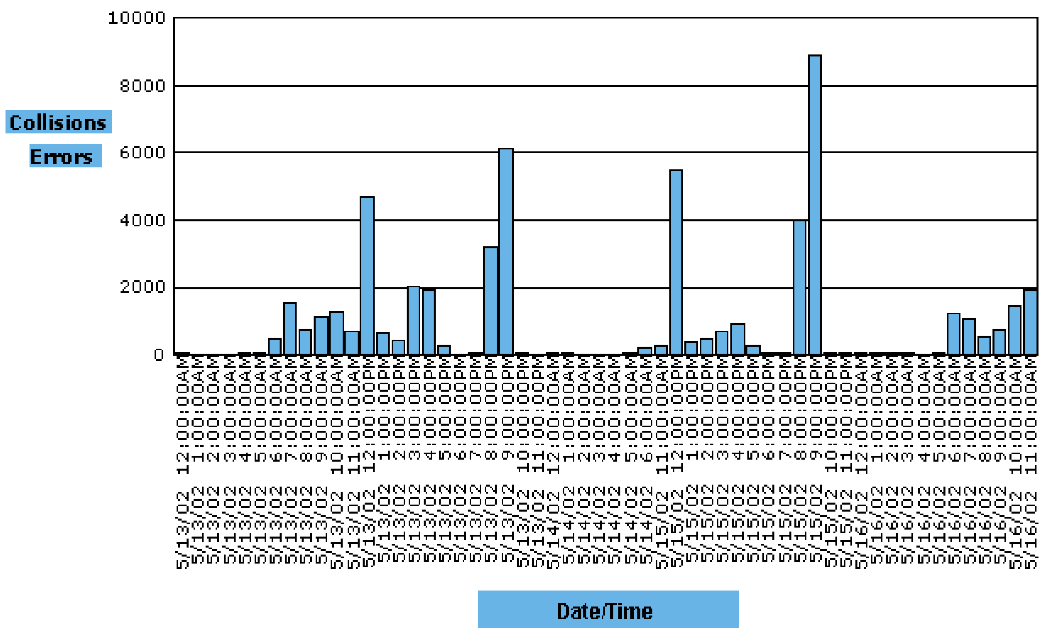


*Figure 6: 4-day Collision Error Count for the RWNMDF Network*

[Ó] This data was periodically captured by the Network Associates Sniffer and subsequently summarized in the Network Associate's Network Performance Orchestrator Visualizer (NPO Visualizer)® database.

Furthermore, TCP uses a combination of sequence numbers, acknowledgements, and (or) timers to track whether a data segment is received, is received without error, and to determine whether to re-transmit the same segment or transmit the next segment in the sequence. Application response time rapidly degrades when TCP has to re-transmit a data segment because of an expired timer or because a segment was corrupted by the collision of another host trying to send data simultaneously. This coupled with the already inefficient packet exchange described earlier, acutely impacts application response times at this facility. It is not a stretch of the imagination to understand why this network will not scale with increased network usage.

The development of a new network for this facility is the next step. A model of the existing network was developed to represent the baseline. Alternative designs for the new network are compared to this baseline to obtain a measure of improved performance.

## 3.0 Creating and Validating the Baseline Model

A baseline model of the existing network was created. Since comprehensive network baseline tools currently do not exist at Sandia, the baseline model was created by collecting information from various sources including conversations with networking staff, existing network drawings, configuration files from pertinent network elements, and network management tools including HP OpenView NNM and the Multi Router Traffic Grapher.[Ñ]

Model efficiency is an important aspect to consider during model development. The key is to represent only as much detail of the system as necessary to obtain accurate results. With this in mind, certain liberties were taken to simplify the model for this facility. Consider Figure 7 while referring to Figure 1.

[Ñ] Work is currently underway to create a continuously valid unified view of the end-to-end network. For more information, please refer to the following URL: http://sass2391/modeling/Goals_&_Objectives/Proposals/2003/VNES_Proposal.doc

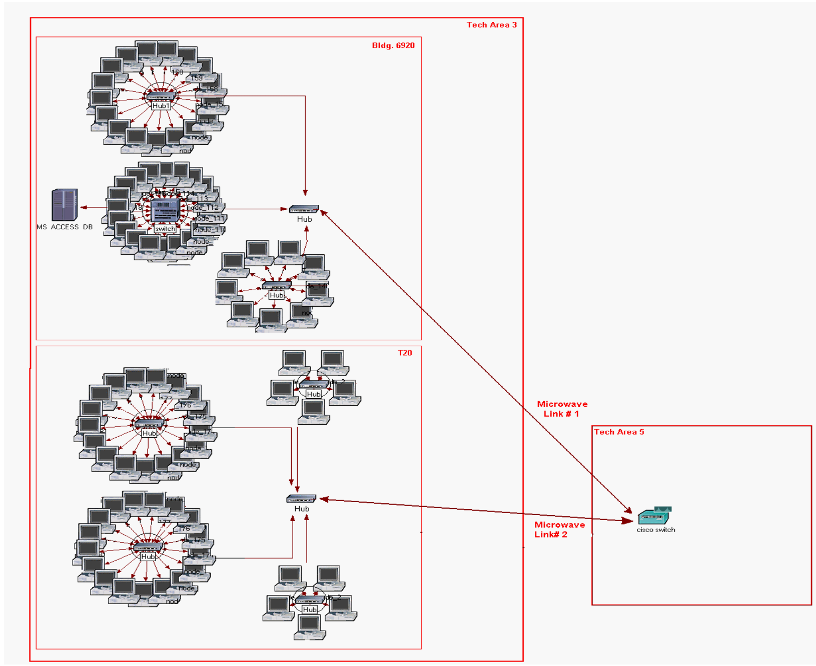


***Figure 7: Model of the Existing RWNMDF Network***

The full duplex microwave links between Tech Area 3 and Tech Area 5 were represented with full duplex 10 Mbps Ethernet links in the model. Transport protocol behavior is unaffected when these types of substitution are made. Additionally, the Ravlin encryptors in Tech Area 3 and 5 and the 2900 series switch in Tech Area 5 were eliminated from the model because the limiting factors, like capacity constraints, exist in the 10 Mbps half duplex hubs and links elsewhere in the network. Latency in the encryptors and in the switch was lumped together and accounted for in each of the 10 Mbps links replacing the microwave links. These substitutions simplify the model and preserve the accuracy of the results. Additionally, the modeled microwave links were loaded with actual background utilization data collected from the live network using the Multi Router Traffic Grapher (MRTG) software package to capture realistic loading effects in the links.[Ó] Five-minute average link utilization levels between the switch in Area 5 and Tech Area 1 were measured with MRTG. Rather than loading this link with background utilization data from MRTG, an equivalent amount of traffic load was simulated by including an email server in Tech Area 1 and multiple email clients in the RWNMDF network. The email server in Tech Area 1 is not shown in Figure 7. The database server at the RWNMDF was represented with a standard model of a Sun Ultra 10 333 MHz server. The use case for the database application was used to create a model of the application. This was accomplished by simply importing the ACE files (created during analysis of the application tasks) into the modeling environment. User data is stripped away, but accounted for, during this process.

---

[Ó] MRTG is currently being phased out and replaced with COS Statscout®. Statscout file formats are currently unsupported in the OPNET Technologies Modeling and Simulation tool suite. Customization will be required to use this data in the future.

Comparing the application response times for each task calculated in ACE and summarized in Table 2 against the simulation results of the baseline network validates the model. Figure 8 demonstrates equivalence between the baseline network model and the actual RWNMDF network for the login task. Similar results were obtained for the other application tasks.

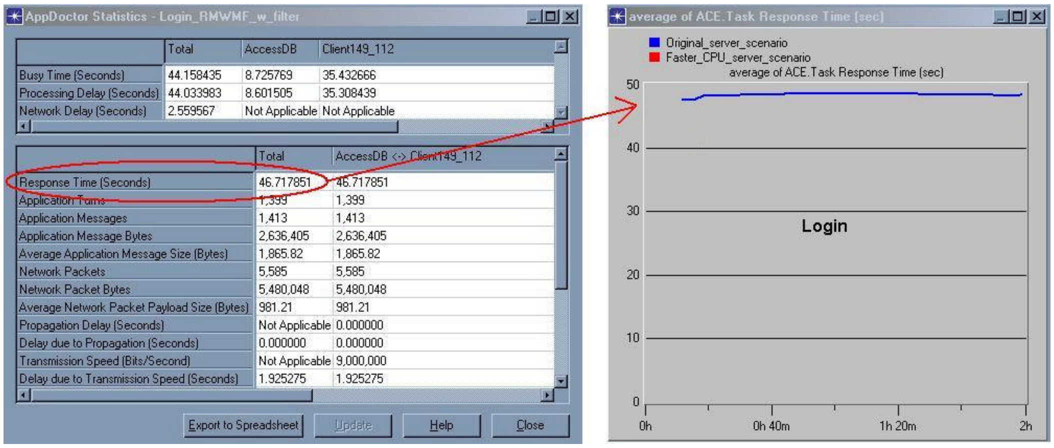

***Figure 8: Validating the Baseline Model***

## 4.0 Modeling the Alternative Network

Several scenarios are evaluated for the new design, keeping in mind performance benefits and implementation cost. Each of the scenarios is discussed and evaluated in the next section. In the first scenario, the processing speed of the server is increased by a factor of four and the network is unchanged. In the second scenario, the cascaded hubs are replaced with Ethernet switches, a fundamental necessity to modernize the RWNMDF network. The database server and its location are maintained as they were in the baseline model. In the third scenario, a 100 Mbps fiber trunk interconnects the two core switches in Tech Area 3. The 802.1Q trunking protocol is enabled on this link and one of the microwave links is removed. Alternatively, the microwave link may also be placed into standby via the Spanning Tree Protocol if link redundancy and diversity redundancy are important; that is, if a fiber cut in Area 3 is likely. The forth scenario evaluates the benefits if the functionality of the database server is moved onto a corporate server in Tech Area 1, while preserving the other changes identified in the second and third scenarios. In this scenario, the microwave link utilization is scrutinized, since suddenly all traffic previously traversing two microwave links is consolidated to traverse the remaining link. The MS Access database would be replaced with an Oracle database or with Microsoft SQL Server [2]. A fifth scenario maintains the corporate database idea, but reinstates the two microwave links and removes the fiber trunk. This scenario is considered another possibility because the installation cost of the fiber trunk may be prohibitive, which would require that the two microwave links remain in place.

## 5.0 Analysis and Recommendations

The Radioactive Waste and Nuclear Material Disposition Facility (RWNMDF) is an example of an existing network where hubs were used as the primary means to access SRN resources. Network utilization is mostly light, 4 – 10 percent, but periodically ramps high during short intervals of time. Light use of the Access Database does not contribute significantly to the traffic load across the network, but may still be high enough to slow the processing time of the server. The database server at the RWNMDF may also be hosting additional applications that were unaccounted for in this study. We did not verify if this was the case. Because of the light usage of the network, network delay was not a contributing factor to poor application response time. However, application response time significantly increases when network load increases.

### *5.1 First Scenario: Upgrading the Database Server*

Since processing delay in the server is a potential bottleneck (see Figure 4), we would expect moderate improvements in the task application response times if processing delays in the server were reduced. Figure 9 shows that in some cases, task application response times can be moderately improved by using a server that is supposed to be better equipped to handle the load. In the model, we simply increased the processing speed in the server by a factor of four. Improvements in response times range from a decent 8 seconds, approximately, for the print task to nearly no improvement in response time for the add-a-form task. A server upgrade is not recommended unless the existing server is in dire need of an upgrade. A better estimate to predict the scalability of the server could have been performed using detailed server models [3, 4, 5, 6, 7]. However, detailed server modeling is outside the scope of this project and would be better served if it were performed in a separate study or reserved for studies that involve large expenditures for enterprise-class servers.

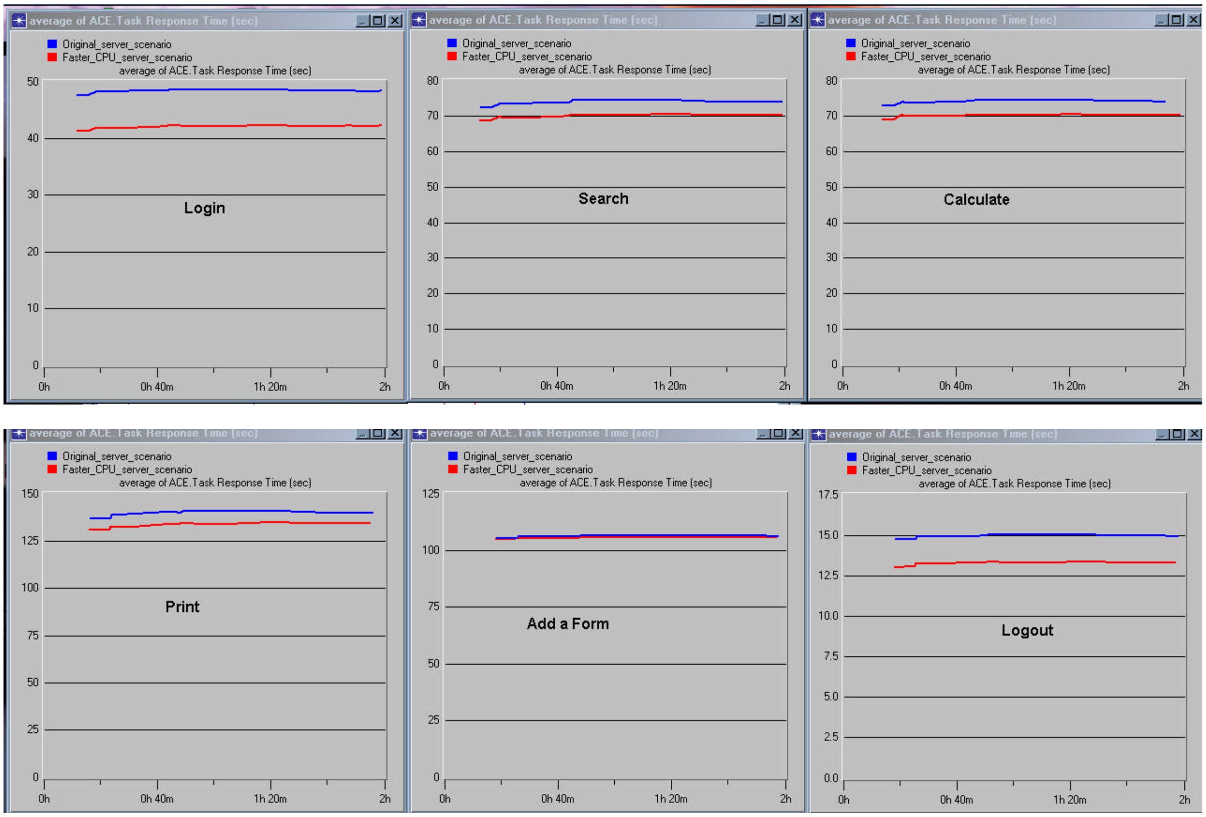


*Figure 9: Response Times Slightly Improve with a Faster Server*

## *5.2 Second Scenario: Upgrading from Hubs to Switches*

In Figure 10, the hubs are replaced with switches. A comparison of the baseline scenario (network of hubs) with the switch upgrade reveals benefits of replacing hubs with switches when network load is high (see Figure 11). In terms of link utilization for the two scenarios, a significant change in microwave link utilization is also visible. The links are better utilized with the switched network. Recall that collisions are common in a network of hubs. In a switched network, they are non-existent, thus more traffic passes through the links (see Figures 12 and 13).

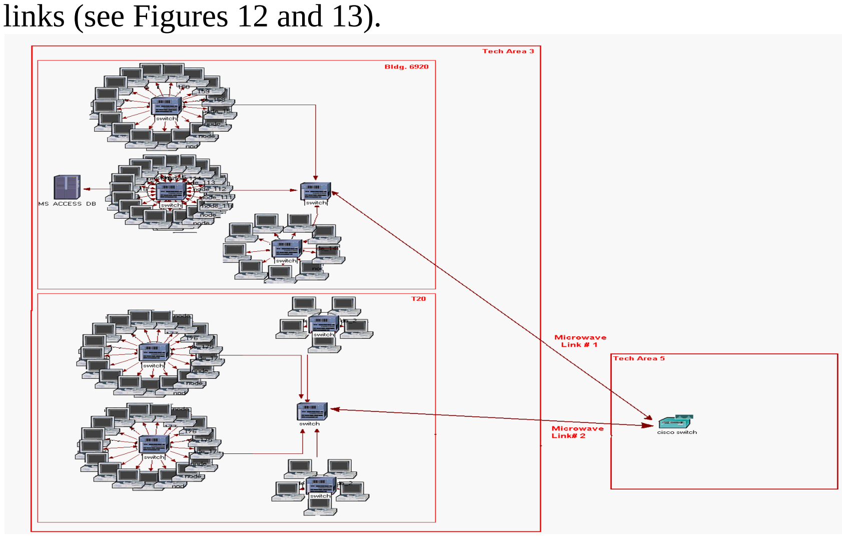


*Figure 10: Replace Concatenated Hubs with Switches*

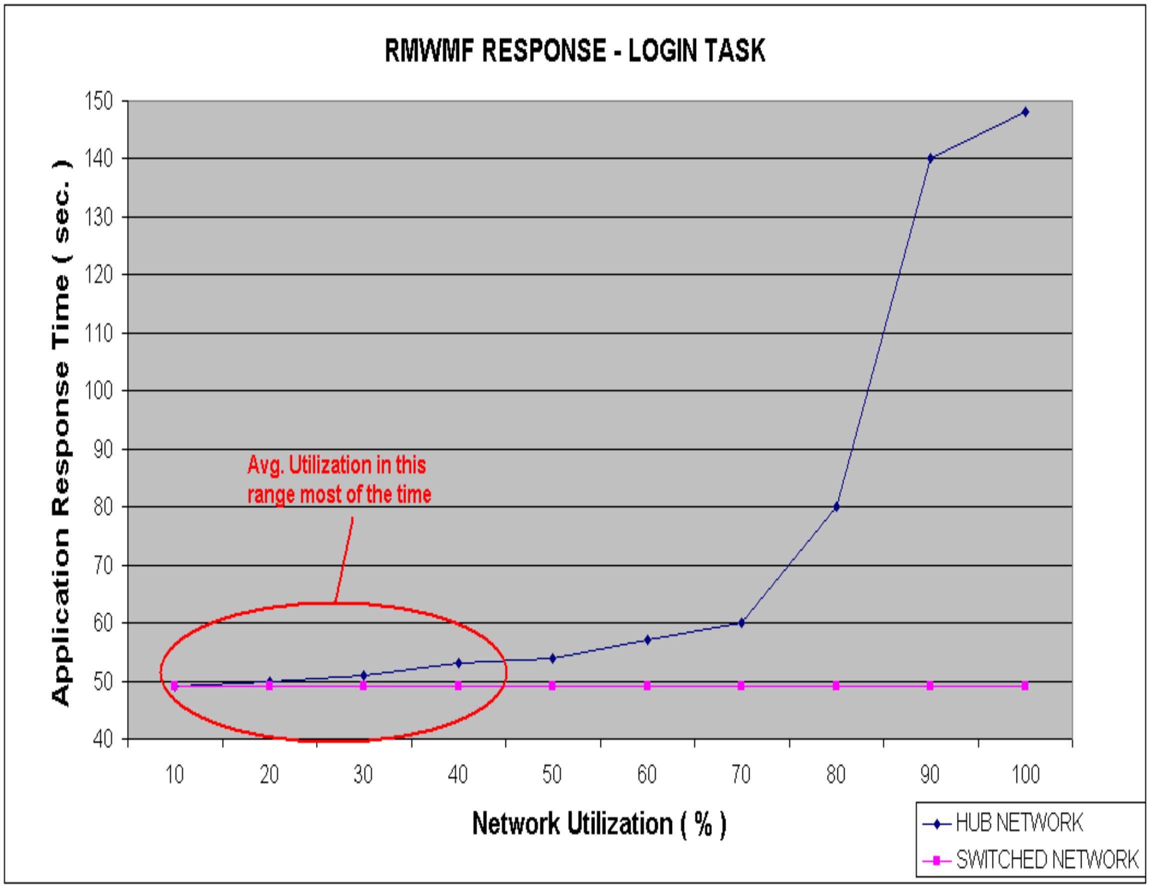


*Figure 11: Performance benefit not realized until network utilization is high*

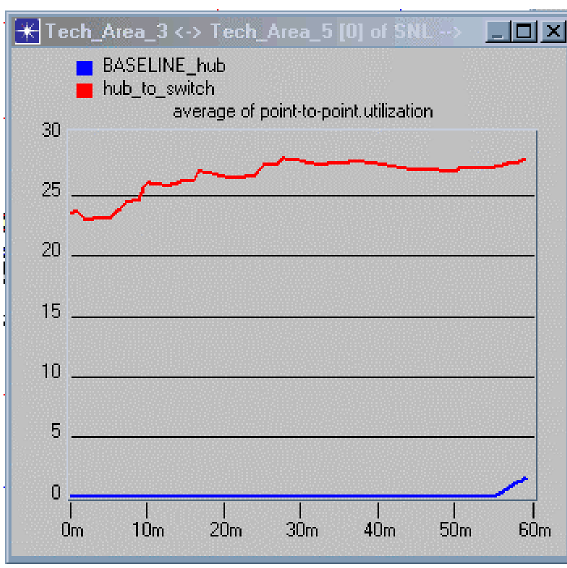


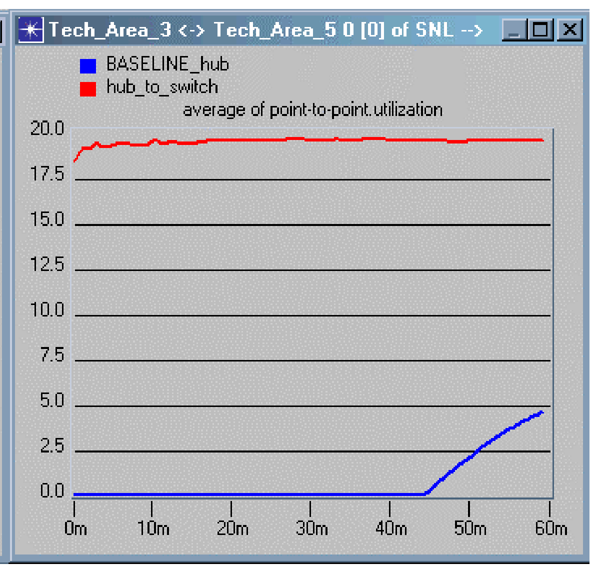


*Figure 12: Link Utilization for Link #1*

*Figure 13: Link Utilization for Link #2*

### 5.3 Third Scenario: Fiber Trunk Between First-tier Switches

The third scenario evaluates whether performance would improve if a 100-Mbps fiber trunk were installed between the two first-tier switches in Tech Area 3 (see Figure 14). This scenario is particularly important because installing a fiber trunk had been proposed to interconnect the two switches (replacing one of the microwave links) providing a possible benefit of significantly improving local network performance. It would be important to determine whether the expenditure for the fiber would provide significant benefits to warrant the cost of the fiber; the installation is potentially expensive. At the same time the fiber trunk is installed, one of the microwave links is disabled or placed in standby. Dual microwave links are unnecessary unless path redundancy and diversity are required. In the event of a fiber cut, the redundant path, i.e., the microwave link, would recover within 20 seconds. Thus, after 20 seconds of downtime, users could resume normal operations.

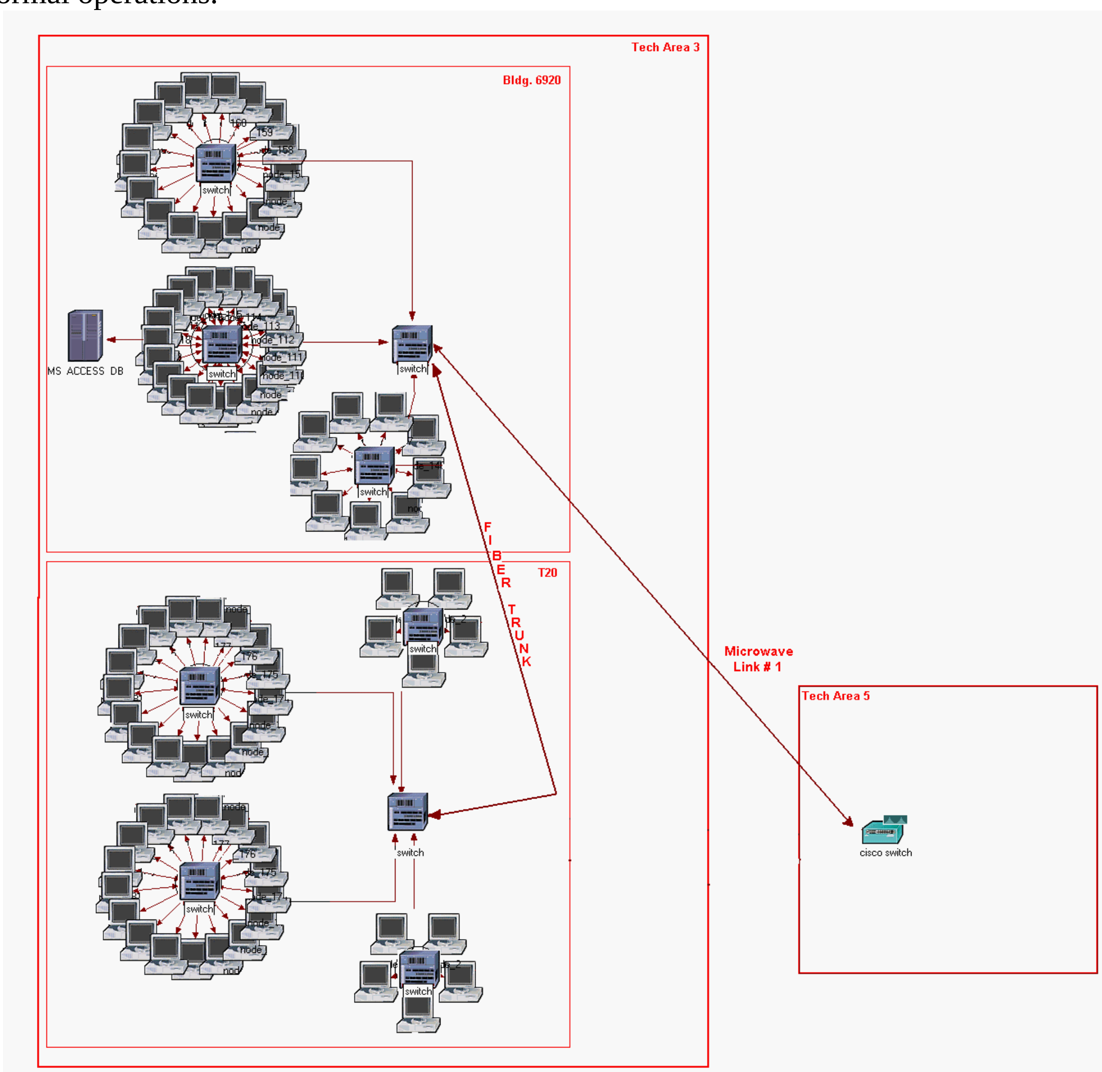


*Figure 14: Interconnect First-tier Switches with 100 Mbps Fiber Trunk*

Figure 15 shows that removal of one of the microwave links, with the fiber trunk in place, minutely improves application response time for the user by an imperceptible one-second. We did not receive a cost estimation for the installation of the fiber, though one was requested. Nevertheless, the fiber would not have been a sensible investment from the standpoint of performance alone. The scale might have been tilted in favor of the fiber from an economics perspective comparing maintenance and reliability costs over the life the fiber versus the microwave link. An economics study was not in the scope of this project.

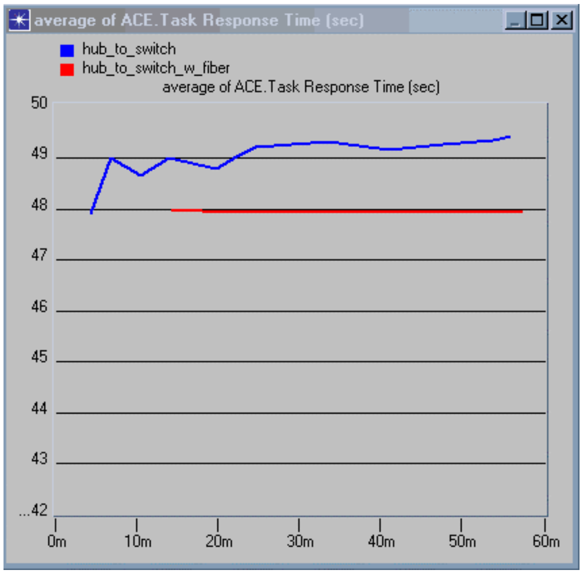


***Figure 15: Application Response Time (Fiber versus Microwave)***

Additionally, link utilization across the remaining microwave link increases moderately by nearly 5 percent as shown in Figure 16. The increase in link utilization is expected because traffic that was bound for the MS Access database server and traffic bound for Area 1 (e-mail, http, etc.) traversed the two microwave links. Now, all traffic bound for Area 1 will traverse a single microwave link. Recall that an e-mail server was used to model the traffic destined to (from) Tech Area 1, but was intentionally omitted from Figure 14 for the sake of simplicity.

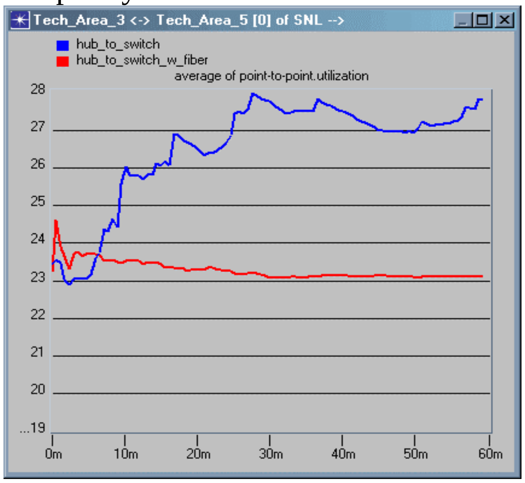


***Figure 16: Link Utilization of Remaining Microwave Link***

## 5.4 Forth Scenario: Move Database to Corporate Servers in Area 1

This scenario looks at the database application response time as if the database were to be moved onto a corporate server in Tech Area 1 (see Figure 17). We assumed that the existing application usage pattern would continue after the database is moved.

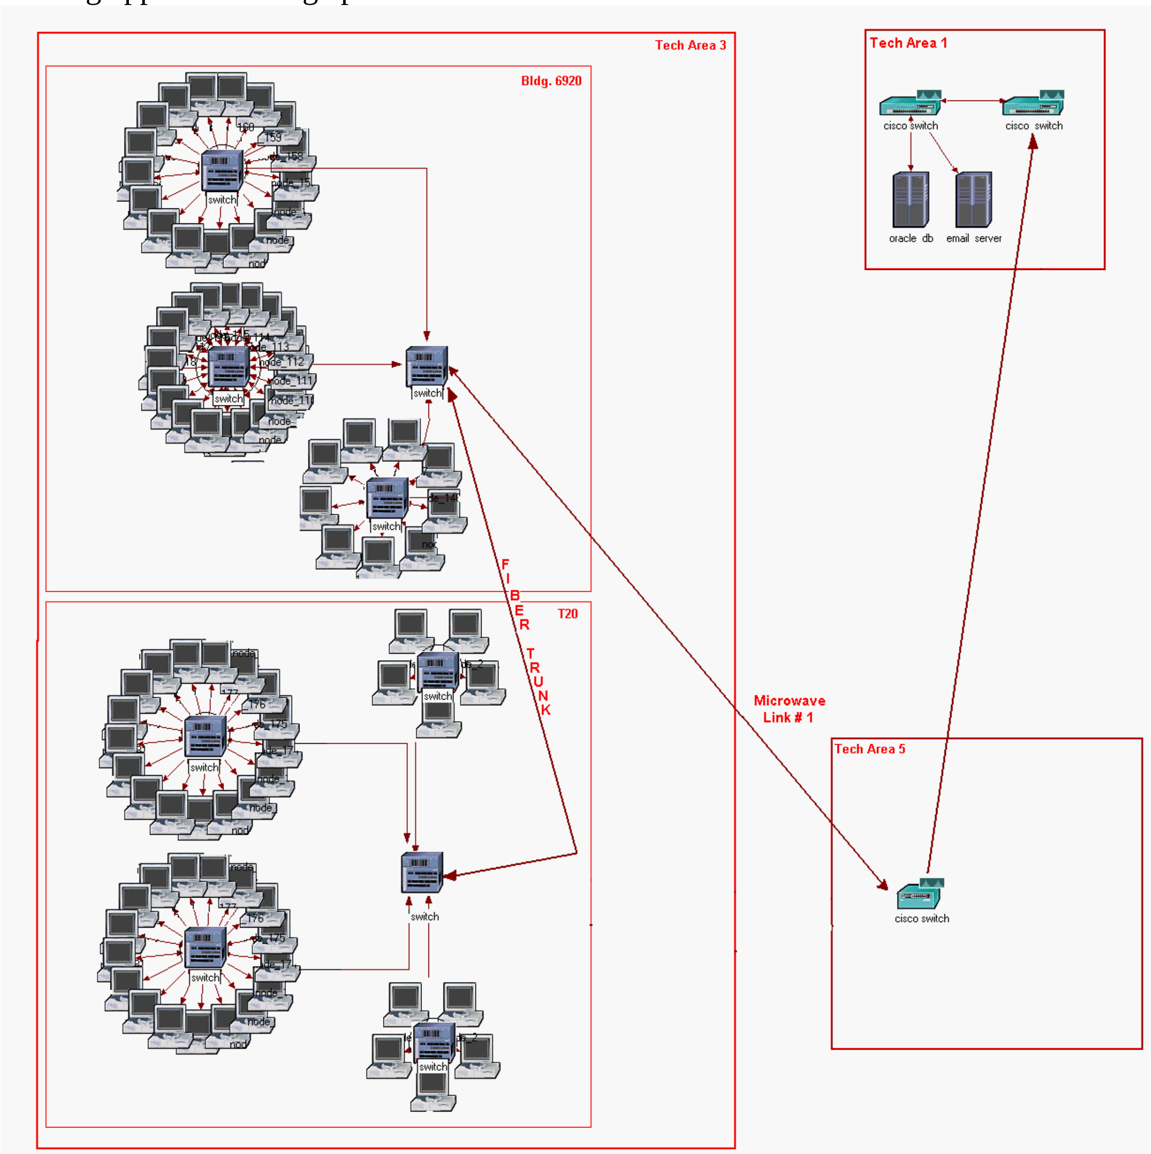


*Figure 17: Move Database to Corporate Servers in Area 1*

Note that the network model was slightly modified to include a database server attached to the distribution switch-router for the corporate servers. Also note that the same server previously located at the RWNMDF is used as the corporate database server. This was done to merely approximate a worst-case performance scenario for users of the RWNMDF database. We would expect that corporate servers are better able to handle corporate-size loads (users of the RWNMDF plus users of any other application that might be hosted on such a server). Realistically, we would expect application response times to be at least as good as if the RWNMDF database server were relocated to Tech Area 1 and its user load were maintained at the same levels. Under these conditions,

Figure 18 shows that the application response time minutely increases (less than 1 second) when the database is moved to Area 1. Additionally, as expected, the utilization of the microwave link increases slightly (see Figure 19).

The benefits of moving the database to corporate servers are not realized in terms of performance or response time, but other factors, probably operating costs, will play a roll in making this decision. Users of the RWNMDF stated that the database needed repair on a near daily basis. Each time repair is required, users of the database are unable to access it for nearly 30 minutes. Time and money spent during the database recovery activity could be recovered, to some degree, by moving the database administration to corporate resources. The downfall, if any, might be less flexibility in extending and changing the functionality of the database to accommodate changes in RWNMDF program requirements. The authors refer the decision makers of the RWNMDF database to the Enterprise Database Administration organization for more information [8].

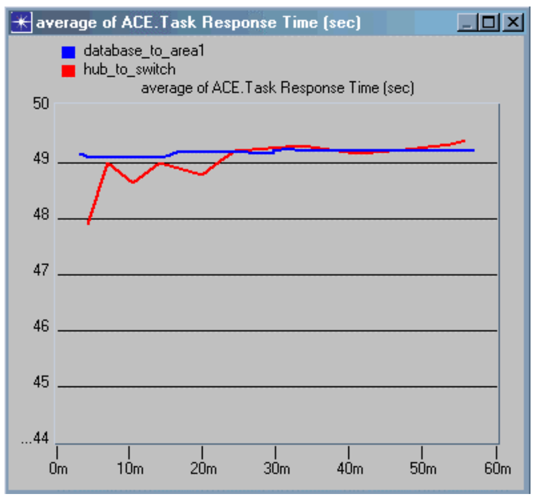


***Figure 18: Application Response Times are nearly the same when the RWNMDF database is moved onto corporate servers***

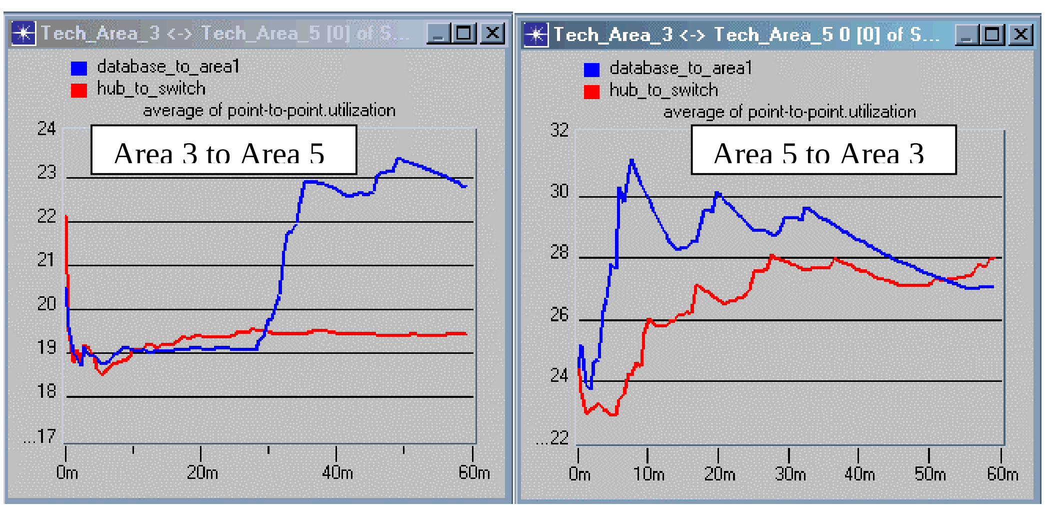


***Figure 19: Microwave Link Utilization (Link #1)***

### 5.5 Fifth Scenario: If the Fiber Link is Too Expensive

This scenario considers the case when both microwave links are maintained as is, in lieu of not installing the fiber trunk. In this scenario, the RWNMDF database is moved onto corporate servers in Tech Area 1 (see Figure 20). For all practical purposes, the application response times for this scenario is similar to the response times in Scenario 2: "Upgrading from Hubs to Switches." Please refer to the response time charts in section 5.2.

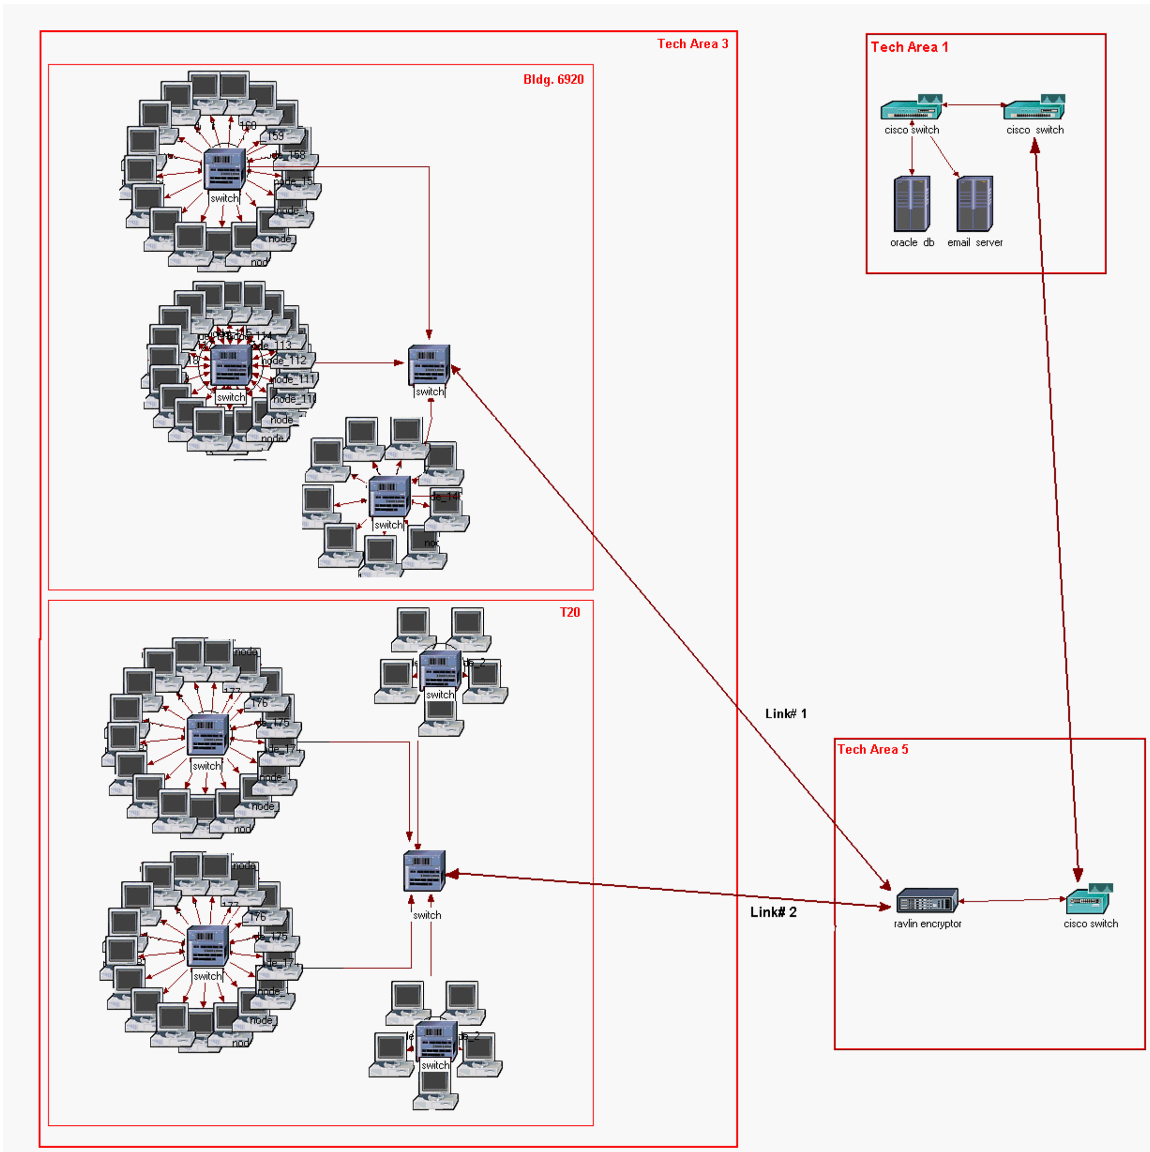


*Figure 20: If Fiber Trunk is too Expensive*

## 6.0 Conclusion

Network performance is not the real issue at this facility most of the time, but it could become an issue in the future. Replacing the network of hubs with Ethernet switches clearly provides benefits to Sandia. Network longevity, increased security, and operational efficiencies are gained by this strategic move. Aside from these fundamental upgrades, the study indicates that additional expenditures in providing fiber connectivity between the first-tier switches in Area 3 do not enhance performance substantially. A one-second improvement in application performance hardly justifies the cost to install this link at this time.

We could not identify with certainty the root cause for the database corruption. The Transmission Control Protocol (TCP) provides reliable and in order delivery of data between a sender and a recipient while assuming very little about the reliability of the underlying communications system. The technique known as positive acknowledgement with retransmission provides an extremely high degree of assurance that data hand-off to the application layer is free of errors and is in the same order as transmitted [9]. Therefore, it is very unlikely that database corruption is caused by poor network performance or reliability. The RWNMDF staff should identify the events that trigger the need to rebuild the database. For example, when does the database get corrupt? Does it occur when a particular record is written or retrieved? Does it occur when a particular user on a particular machine uses the database? Does it occur only when remote users use the database (i.e, connect to the network on a different switch than that which connects to the database server)?

The database server at the RWNMDF and the database application remain an issue to be resolved. The analysis showed that the client contributes 80% to the application response time and the server contributes 14%. The remaining 6 % is contributed by protocol and transmission delay. The apparent alternatives are to replace the existing server with one more suited to this application or to move the database onto corporate servers in Tech Area 1. From purely an application-performance perspective, upgrading the database server to a faster machine should be done only if the existing server is in need of upgrading, since only low-to-moderate improvements in application performance can be realized after the upgrade. The new machine may or may not resolve the database issue. The study showed an imperceptible difference in application performance if the database is moved onto corporate servers in Tech Area 1. In addition, the existing database may also be in need of work (repair or replacement). The large delay component exhibited in the client includes user think time and cannot be neatly separated into the system and human delay components. Cost benefits may be realized by moving the database onto corporate servers.

The Radioactive Waste and Nuclear Material Disposition Facility is one example in which modeling and simulation were used to analyze a problem and to evaluate design alternatives. It was possible to proactively determine where time is spent during a case of application use to identify the most likely areas of response improvement, demonstrating how using today's technology can vastly improve problem resolution without using the trial-and-error techniques of the past. Modeling and simulation not only save the time involved in replacing network hardware in the hopes of fixing a problem, but they also allow network operations staff to efficiently design alternatives to improve network functionality while relying on quantitative analysis rather than anecdotal techniques.

## 7.0 References


1 C. E. Spurgeon, *Ethernet: The Definitive Guide*, O' Reilly & Associates, Inc., California, 2000.

2 W. R. Mertens, DBMS Platform Selection, http://www-irn.sandia.gov/organization/div9000/ctr9500/dpt9523/dbms%20platform%20selection.html.

3 An OPNET Methodology and Case Study, December 19, 2002, Server Capacity Planning Using the Specialized Server Models, http://www.opnet.com.

4 An OPNET Methodology and Case Study, January 6, 2003, Methodology for Organizing Server Assets, http://www.opnet.com.

5 An OPNET Methodology and Case Study, Session 1405 Predicting Server Performance Scalability – Introduction, OPNETWORK 2002 Conference Proceedings, http://www.opnet.com.

6 An OPNET Methodology and Case Study, Session 1416 Predicting Server Performance Scalability - Advanced, OPNETWORK 2002 Conference Proceedings, http://www.opnet.com.

7 An OPNET Methodology and Case Study, Session 1417, Predicting Storage Network Performance, OPNETWORK 2002 Conference Proceedings, http://www.opnet.com.

8 Enterprise Database Administration Organization, http://www-irn.sandia.gov/organization/div9000/ctr9600/dpt9618/dba_home_page.html.

9 D. E. Comer, *Internetworking with TCP/IP Principles, Protocols, and Architectures, Fourth Edition*, Prentice Hall, Inc., New Jersey, 2000.

**Distribution**

| Copies | MS | |
|---|---|---|
| 1 | MS1149 | J. J. Thompson, 3125 |
| 3 | MS1149 | D. L. Muller, 3125 |
| 1 | MS9012 | S. C. Gray, 8949 |
| 1 | MS9012 | R. D. Gay, 8946 |
| 1 | MS9012 | F. H. Blair, 8946 |
| 1 | MS0630 | M. J. Murphy, 9600 |
| 1 | MS0622 | D. S. Rarick, 9620 |
| 1 | MS0662 | C. D. Caton, 9623 |
| 1 | MS0801 | A. L. Hale, 9300 |
| 1 | MS0801 | M. R. Sjulin, 9330 |
| 1 | MS0805 | W. D. Swartz, 9329 |
| 1 | MS0806 | C. R. Jones, 9332 |
| 3 | MS0812 | M. J. Benson, 9334 |
| 1 | MS0812 | R. L. Adams, 9334 |
| 1 | MS0812 | J. L. Akins, 9334 |
| 1 | MS0812 | B. L. Amberg, 9334 |
| 6 | MS0812 | D. B. Bateman, 9334 |
| 1 | MS0812 | L. D. Byers, 9334 |
| 1 | MS0812 | K. R. Carpenter, 9334 |
| 1 | MS0812 | L. S. Chance, 9334 |
| 1 | MS0812 | L. A. Dubes, 9334 |
| 1 | MS0812 | B. Dominguez, 9334 |
| 1 | MS0812 | H. R. Garcia, 9334 |
| 1 | MS0812 | M. D. Gomez, 9334 |
| 1 | MS0812 | M. L. Grassham, 9334 |
| 1 | MS0812 | M. J, Hamill, 9334 |
| 1 | MS0812 | W. E. Hill, 9334 |
| 1 | MS0812 | J. A. Keelin, 9334 |
| 1 | MS0812 | C. M. Keliiaa, 9334 |
| 1 | MS0812 | E. J. Klaus, 9334 |
| 1 | MS0812 | A. B. Leyba-Essary, 9334 |
| 6 | MS0812 | R. A. Life, 9334 |
| 1 | MS0812 | T. Lyon, 9334 |
| 6 | MS0812 | J. H. Maestas, 9334 |
| 1 | MS0812 | P. L. Manke, 9334 |
| 1 | MS0812 | R. A. Mason, 9334 |
| 1 | MS0812 | R. L. Moody, 9334 |
| 1 | MS0812 | J. Pelowitz, 9334 |
| 1 | MS0812 | D. R. Porter, 9334 |
| 1 | MS0812 | M. A. Rios, 9334 |
| 1 | MS0812 | G. Rivera, 9334 |
| 1 | MS0812 | G. F. Rudolfo, 9334 |
| 1 | MS0812 | T. J. Spears, 9334 |

| | | |
|---|---|---|
| 1 | MS0812 | P. M. Torrez, 9334 |
| 1 | MS0812 | B. C. Whittet, 9334 |
| 1 | MS0812 | V. K. Williams, 9334 |
| 1 | MS0806 | L. Stans, 9334 |
| 1 | MS0806 | S. A. Gossage, 9336 |
| 1 | MS0806 | L. F. Tolendino, 9336 |
| 1 | MS0806 | T. D. Tarman, 9336 |
| 1 | MS0806 | J. Wertz, 9336 |
| 1 | MS0806 | E. L. Witzke, 9336 |
| 1 | MS1392 | B. L. Weaver, 15421 |
| 2 | MS0899 | Technical Library, 9616 |
| 1 | MS9018 | Central Technical Files, 8945-1 |
| 1 | MS0612 | Review/Approval Desk, 9612 |