An Improved Approach to Product Quality Through Testing

Margaret A. Hillis
Gail W. Minkus
Craig L. DeCaluwe
Michael N. Meyers
Edward T. Schafer

Due to the crucial importance of the 1B processor project, AT&T Network Systems (NS) implemented extensive testing programs to ensure that the product delivered was of the highest quality and met customers' needs. In addition to normal customer acceptance testing, the AT&T Network Services Division (NSD) introduced an innovative new strategy called the accelerated life-cycle test program (ALTP) to facilitate the identification and repair of defects. NSD used a 4ESS™ switch having 100-percent reversible traffic to test the retrofit process and to ensure that customer service would not be adversely affected by the introduction of the 1B into the AT&T network. This paper describes the final product integration and verification processes performed by NS, along with the special testing programs established by NSD customers.

Introduction

The 1B program represented a "brain transplant" for the AT&T network.[1] Thus, the 1B processor underwent one of the most stringent testing programs ever devised for an AT&T product. Three separate, independent test procedures were established, and Figure 1 shows the timeline for the completion of these procedures.

The first phase, known as the accelerated life-cycle test program (ALTP), stressed fault recovery and diagnostics. As a test, technicians purposely introduced over 1,000 defects into the 1B and then assessed its ability to recover. The ALTP also provided the opportunity for hands-on training, allowing field technicians to gain practical experience in the operation, repair, and maintenance of the processor. This first phase also provided the AT&T Network Services Division (NSD) with much valuable information about the length of time needed to detect and repair faults in the field.

The second phase of the 1B testing program consists of system integration and performance-verification (SI/PV), carried out in addition to and parallel with the ALTP. SI/PV is an extremely rigorous testing routine having more than 3,000 different scenarios, and the combinationof the two eliminated test-procedure hand-offs among different groups. This allowed each tester to become expert in at least one functional area, which resulted in the detection of even minute defects before product release.

The network risk avoidance program (NRAP) is the third testing phase shown on the timeline. A 4ESS™ switch in the AT&T network—having 100-percent reversible traffic and handling limited, "live" network activity—was carefully preconditioned and retrofitted with the 1B processor. The NRAP allowed NSD to conduct acceptance testing and study the 1B in a live network environment and, at the same time, eliminated the possibility of adversely affecting customer service.

In addition to the development of an extremely robust three-phase testing procedure, a valuable partnership was formed between NSD customers and the SI/PV group. It was established during the planning phase and continued through the implementation and results-reporting activities. NSD found that this partnership resulted in the following three test-program benefits:

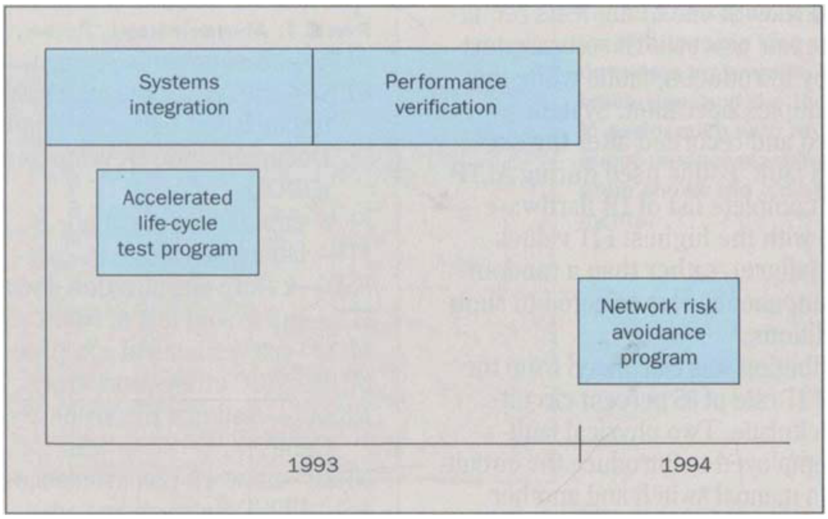

Figure 1. This timeline shows the completion schedule for the three separate, independent test procedures established for the 1B processor development program. The first phase is called the accelerated life-cycle test program (ALTP). The second phase is known as system integration and performance verification (SI/PV). The third testing phase consists of the network risk avoidance program (NRAP).

- Establishment of many new test strategies;
- Reduced interval between fault detection and resolution; and
- Enhanced reporting of testing progress.

NSD—recognizing the value of customer and design-organization partnering—is continuing the practice in developing and testing other AT&T products.

This paper describes the specifics of all three testing phases (ALTP, SI/PV, and NRAP) and the details of the partnership between the NSD and SI/PV groups. The discussion highlights the innovations used during the 1B testing processes and points out some strategies that could be used on other projects.

Accelerated Life-Cycle Test Program

As previously mentioned, the first phase of testing is the ALTP, which stresses fault recovery and diagnostics. A discussion of its elements follows.

Strategy. The ALTP is a test program developed by NSD and NS to provide 1B-processor software and procedural-reliability metrics. Failure in time (FIT) data is used to identify the most likely hardware faults. The hardware is then purposely degraded at an accelerated rate to measure system responses to expected failures. The data obtained is normalized to achieve the expected performance during deployment.

In addition, technicians from the NSD On-Site Work Force, Technology Control Centers (TCCs), and the National Electronic Systems Assistance Center (NESAC) participate in the ALTP. The NSD field organization provides technicians with high-level exposure to the 1B, develops subject-matter experts during 1B deployment, and uses experienced testers to assist in identifying any 1B weaknesses.

Process. The ALTP was performed at the NSD Integrated Test Network (ITN) facility in Warrenville, Illinois. ALTP results provided a measurement of software quality for the 1B fault-recovery, diagnostic, and trouble-location programs, as well as verifying the associated product maintenance documentation.

The NSD 1B project team set an objective to implement 1,000 ALTP tests. This objective was based on four tests per shift per day, with one of the tests requiring use of the product repair process. 1B software testing was accomplished by introducing faults while the 4ESS switch was in full-duplex operation. System responses were analyzed and recorded after the successful introduction of a fault. Faults used during ALTP were selected from the complete list of 1B hardware components. Hardware with the highest FIT values (greatest probability of failure), rather than a random selection from all 1B components, was selected to simulate expected field conditions.

ALTP fault distribution was calculated from the 1B-processor hardware FIT rate of 85 percent circuit packs and 15 percent backplane. Two physical fault-insertion devices were employed to introduce the circuit-pack defects—one with a manual switch and another with an automatic time-delay trigger. The faults were applied directly to the component under test. Backplane faults were split between pins and cables. Pin faults were introduced on the backplane in circuit packs having the highest FIT rates. A clip lead was used to simulate a short circuit to ground. Cable faults were inserted by removing the designated cable while all units were in service, simulating an "open" cable.

Benefits. To take full advantage of the 1B evaluation interval, NSD field technicians were employed to conduct the ALTP. As a result, they gained valuable expertise in the repair and maintenance of the 1B processor by working directly with the system and its fault-locating tools. The technicians used diagnostic routines, task-oriented procedures, and equipment drawings to isolate, identify, and repair defective system components.

Through their analysis of the test results, the NSD team identified areas for improvement in the hardware, software, training procedures, and documentation.

In addition to carrying out the ALTP, field technicians were given an opportunity to perform multiple 1B retrofits. They also implemented the procedure necessary to return to the 1A configuration, which is part of the 1B-system contingency plan. Feedback from each test installation was used to improve the retrofit process and provide the following benefits:

- *Additional hands-on training.* Each field crew member attended a hands-on Essential System Maintenance course for the 1B processor before taking part in the ALTP. Subsequent ALTP participation enhanced this training by allowing both technicians and supervisors to carry out repair procedures, increasing their comfort and familiarity with specific field activities. During the 1B retrofits completed during each ALTP session, faults were purposely introduced to allow the field crew to detect them and experience the potential failures in the retrofit procedures.
- *Design improvements.* System deficiencies—identified for the 1B design team through the implementation of the testing plan—were reviewed and analyzed. This resulted in reliability improvements to both the hardware and software, which included the correction of such critical faults as "silent death"—the failure of a system to initiate appropriate alarms.

During the ALTP, concerns were raised regarding the reliability and integrity of the 1B master control complex (MCC) and MCC utility processor (MUP). The MCC provides control and display functions for the technicians, whether they happen to be working at a switch location or remote site, such as a TCC. During the course of testing, many MCC and MUP changes were introduced into the MUP firmware loads. As test-

Panel 1. Abbreviations, Acronyms, and Terms

ALTP—accelerated life-cycle test program
CENT—the AT&T Customer Education and Training organization (formerly known as the Documentation Development Organization [DDO])
FFA—first field application
FIT—failure in time
FSD—feature specification document
ITN—integrated test network
MCC—master control complex
MUP—MCC utility processor
NESAC—National Electronic Systems Assistance Centers
NRAP—network risk avoidance program
NS—AT&T Network Systems
NSD—AT&T Network Services Division
SI/PV—system integration and performance verification
TCC—Technology Control Center

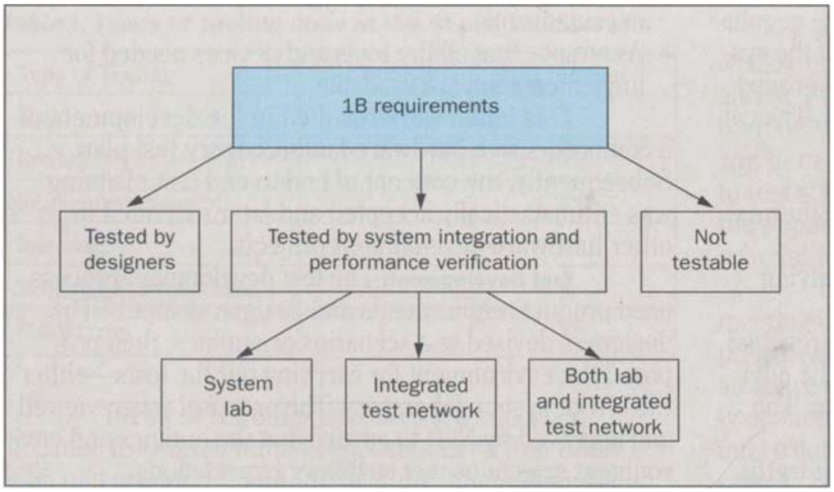

Figure 2. An up-front resolution of requirements verification was first achieved during the 1B development project. The process has proven invaluable, and the technique was used again, in partnership with NSD customers, on subsequent product planning efforts. The block diagram shows the categorization of 1B requirements.

ing progressed, the performance of the MCC-to-MUP interface improved. The 1B project team, however, continued to monitor and evaluate the performance of both the MCC and MUP closely until all issues were resolved.

- *Documentation and training improvements.* Throughout the testing interval, the AT&T Documentation Development Organization (now known as the Customer Education and Training organization [CENT]) provided on-site support. In addition, CENT representatives obtained first-hand feedback about the manuals, instruction sheets, and other printed material produced for the 1B project. They recorded comments from the field groups during the debriefings following each ALTP session, noting any areas for improvement.

Results. In the course of the four-month duration of the ALTP, 1,021 tests were completed, which averages 5.3 tests per shift. Field technicians repaired circuit packs during 171 of these tests using instructions provided in the trouble-clearing documentation. The mean time to repair was 61 minutes, and the outstanding stability of the 1B hardware allowed for near continuous testing. Only one extended failure occurred out of all 1,021 faults introduced into the system.

Initial coverage results were very good in most areas. The ALTP did uncover the need to enhance the diagnostic coverage. The 1B design team put forth a concerted effort in resolving all problem areas identified by the ALTP. Final results of the test procedures indicated that the feature specification document (FSD) coverage requirements were achieved in all areas.

The ALTP provided an excellent means both to train NSD field technicians and to provide design engineers with valuable product information from the perspective of end-users—for example, feedback regarding MCC reliability.

System Integration and Performance Verification

This section highlights the SI/PV process, discusses the innovations used to test the 1B, and focuses on the important partnerships forged with NSD.

Single Team Approach. The 4ESS-switch product realization process defines the SI and PV processes.[2] A single team was formed to execute both SI and PV to reduce resource costs, increase tester expertise, and improve testing quality. This single-team approach was first tried on a smaller project to ensure its viability. Throughout the planning, test development and execution, and product evaluation stages, the SI/PV team established partnerships with its customers and suppliers to improve the quality of the 1B product and the 4ESS-switch system.

Planning. The 1B SI/PV team planning effort began early in the project, when SI/PV representatives participated in FSD reviews. The representatives' role was to review the FSD requirements for clarity and testability. Using the NSD Network Test Plan as a guide, they also began consulting with customers to understand the specific types of tests and test environments that were required. Subsequently, the SI/PV team proposed and negotiated a test plan and strategy. As part of the negotiation process, the team coordinated the efforts of the system-design, systems-engineering (SE), and NSD groups to categorize the testability of each requirement. The categories were established as follows:

- Testable by the SI/PV team;
- Best verified by product design (for example, thermal requirements); and
- Not testable (for example, requirements specifying that certain designs are extensible).

For each requirement, the SI/PV team reached an agreement with NSD customers regarding the environments in which verification would take place. The following three alternatives were established:

- The laboratory environment (in which a high traffic load can be generated);
- The ITN environment (in which connections with other switches can be maintained); or
- Both environments.

This up-front resolution of the requirements verification issue was first achieved during the 1B project. The process has proven invaluable, and the technique was used again—in partnership with NSD customers—on subsequent product-planning efforts. Figure 2 illustrates the categorization of the 1B requirements.

The SI/PV team worked closely with both the hardware- and software-development groups to plan the test program. In addition, one tester interned with the 1B development integration team to understand the testing program more clearly. This helped prevent skipped tests and ensured that none was redundantly conducted. Subsequently, this tester used the knowledge gained and contacts established through the internship to direct trouble reports quickly to the appropriate designer.

The goal of end-to-end test planning was also established, and it was carried out during *hardware fault recovery*. This process specifies the sequence of actions a switch must follow to detect, isolate, report, and recover from a hardware defect. Hardware fault recovery was identified as a key factor during test planning. Thus, the SI/PV, design, and NSD customer groups—working together—established the following three guidelines to ensure its implementation:

- Identification of all parameters and environments requiring testing;
- Determination of where and by whom testing is best executed; and
- Assurance that all the tools and devices needed for implementation are available.

This initial work resulted in the development of a comprehensive hardware-fault-recovery test plan. Subsequently, the concept of end-to-end test planning was enthusiastically accepted and later extended to other hardware-development projects.

Test Development. The test development process used product requirements and designs as input.[3] The designers devised test scenarios or outlines, then proposed the environment for carrying out the tests—either the ITN or system laboratory. The proposal was reviewed and approved by NSD to ensure that the outlines and environment selections met customer expectations.

Customers, designers, and the other SI/PV team members reviewed and approved the scenarios for accuracy and completeness. Each outline included references to requirements that the particular scenario verified. When the outlines were stored on the electronic test storage system, special tools were used to determine the requirements coverage for the scenarios selected for each phase of test implementation. This provided an additional check on the test program's coverage.

Test Implementation. The process of carrying out testing occurred during both the SI and PV intervals. Table I summarizes the types of testing done at each interval. SI/PV test implementation began with SI, in which the strategy was to implement tests in a "soak" environment. A 4ESS-switch soak means operating the switch, with traffic, in a field system profile for one week. During that time, the first and second eight-hour shifts were dedicated to SI testing and observing the overall 4ESS-switch performance. On the third shift, the system "ran" traffic without tester activity.

The 4ESS-switch print-outs for the third shift are analyzed and a stability index calculated. The stability index is a number between zero and 100 that characterizes how well the system is operating. Thus, a perfectly running system would score 100. Tracking these stability indices over time allowed both team members and NSD to see how hardware and software quality were improving.

During the PV interval, most testing moved to the ITN, which provided the environment. The focus was on the following four processes: requirements verification, retrofit, feature regression, and interswitch testing.

Table I. Types of testing done in the SI and PV intervals

Type of Testing	SI Interval	PV Interval
Soak	X	X
Retrofit		X
Feature regression		X
Interswitch		X
Requirements verification	X	X
Stability run	X	X

Flawless retrofit procedures are important because they often influence a customer's first impression of a product. Using the field documentation, the SI/PV team tested the retrofit process frequently in the ITN. To understand the usefulness of these procedures, the team commissioned different testers to implement them. A retrofit testing expert was also present to help detect problems. Furthermore, the team ensured that a switch traffic load existed during the retrofits—which was the case for all the ITN tests—so the impact on calls could be clearly understood and evaluated.

During the PV interval, feature-regression tests were implemented. These tests check whether old features still operate in the new system. The testers who verified the same features on the 1A processor sampled their tests on the 4ESS-switch system with the 1B. They did so because the features were implemented the same way on both the 1B and 1A processors. The testers also ensured that interfaces between the 4ESS switch, the network databases, and adjuncts were the same whether the 4ESS switch was 1A or 1B based. During times when no testers were scheduled on the 1B processor in the ITN, the SI/PV team ran a call load through the system as a further verification and check on switch stability.

While the primary focus during the PV interval was on ITN testing, the SI/PV team still carried out many tests in the system lab. Soak testing continued, including stability runs and index tracking. Call-load capacity testing was completed during soaks to ensure that the 4ESS switch with the 1B installed met the capacity requirements specified in the FSD. Some anomaly tests were conducted with the system at full call-load capacity. These tests showed how the system—already under stress—would react to the loss of a crucial component.

Fault Reviews. System faults were reviewed weekly during design team meetings. Both the NSD and SI/PV teams participated in 1B development reviews to remain informed about documented defects and to monitor trouble resolution timetables as an aid to test scheduling. This procedure provided testers the opportunity to request fault remedies to help meet scheduled intervals.

The design community encouraged the weekly meetings because such reviews assisted in establishing priorities for correcting identified faults. NSD assumed an active role in the meetings and retained final authority on fault dispositions. Thus, NSD must concur with a final trouble-report resolution for it to be considered closed. This concurrence provided the design team with a customer perspective and helped ensure that important end-user issues were resolved.

Evaluation. The test-implementation phase was considered complete when the SI/PV team met its exit criteria as specified by NSD customers. The team documented its findings in a report that included initial and final test results, disposition of all faults found, certification letters signed by the testers, and a recommendation for deployment. The report also included the results of all requirements verification completed by product design. Thus, the final SI/PV report provided NSD with an overview of the testing program, which was established to verify each testable 1B requirement.

Results and Future Direction. The success of the SI/PV testing program is best measured by the output metrics associated with the PV process. These metrics include the test-implementation rate, fault-finding rate, percent of valid faults reported, and overall customer satisfaction as measured by a PV report card. PV-process customers include internal design and field-support groups. Results indicated that SI/PV testing met or exceeded customer process requirements for all metrics. Higher-than-normal productivity can be attributed, in part, to the expertise developed by the testers, who tested the same 1B parameters during both the SI and PV intervals. The extraordinary level of quality assurance can be accredited to that expertise, as well as to the fault-screening process.

The marks on the PV report card for the 1B were extremely high. In almost every case, both internal and external customers graded the PV process as exceeding expectations. According to customer comments compiled by the SI/PV team, the success of the testing process resulted from effective communication and consistent follow-through. Early in test-process planning, the SI/PV team provided customers with a clear explanation of what to expect, contributing to the high report-card marks.

The SI/PV team strategy will be adopted for future hardware development projects. In addition, requirements verification during a project's planning phase and the tracing of requirements to specific tests have both become standard. End-to-end test planning has been used in a number of new platforms. Plans have also been established to provide customers with SI/PV design-team letters verifying requirements in platform reports. Some of the lessons learned from 1B testing are already being implemented on other hardware development projects, which will result in even more robust SI/PV testing processes in the future.

Network Risk Avoidance Program

NSD established a test plan, known as the NRAP, for the first field application (FFA) switch, providing additional assurance that the 1B processor would perform as designed. An FFA site was selected that met the NSD need to ensure that customers would not experience any service interruptions.

The FFA site required the capabilities of increasing or decreasing switch traffic as needed. The NRAP not only tested the product; it also tested both the process and support systems required for 1B-processor retrofit into the AT&T network. This plan allowed NSD to perform testing in a highly controlled environment. The 1B project team and the AT&T Network Operations Center developed a plan to perform progressive testing and training on the FFA switch. Maintenance technicians from throughout the country were sent to the switch site, and they conducted multiple retrofits with the support of the TCCs.

Retrofit Testing. The progressive test plan required that all traffic be removed from the 4ESS switch, and the initial retrofit was performed while monitoring the switch's service indicators. Following the analysis of a successful retrofit, both the TCC and NESAC performed multiple practice retrofits for two weeks. In addition, field technicians were brought in to practice and refine their tasks associated with the retrofit process.

Once the practice retrofit testing was completed, the 1A processor was placed back into service and a light traffic load applied to the switch. The switch was then retrofitted with the 1B processor one final time to assess its effect on live traffic. All service indicators were monitored and assessed confirming the belief that the 1B was ready for deployment.

Once convinced that the retrofit process could be performed reliably, NSD focused attention on the in-service operation of the 1B processor. The FFA switch selected for this testing was chosen because its traffic was fully backed up and could be rerouted quickly to ensure uninterrupted customer service. During the first month of service, the AT&T Network Operations Center in Basking Ridge, New Jersey continued to increase switch traffic while monitoring service quality. Meanwhile NSD, in cooperation with the TCC and NESAC, monitored the performance of the 1B processor's hardware and software. Each element performed flawlessly, confirming results of previous testing.

The increase in the 1B processor's real time, which resulted in additional call-processing capacity, was part of the expected performance improvement as specified in the FSD. The Network Operations Center compared the 1A and 1B processors' real time and confirmed that the expected gains were achieved. At this point, all involved in the 1B project were confident that the processor was ready for deployment.

1B Versus 1A Performance

After its first year of successful operation, the 1B processor continues its record of outstanding performance. In the AT&T network, it currently operates with less than half the downtime typically experienced with the older 1A processor. Customers outside AT&T also have reported excellent results. In addition, both the anticipated hardware failure rate and associated subscriber service interruptions have been significantly reduced.

Partnership Value. The entire testing process was a true team effort. Abundant trust had to be established among the NS and NSD teams, enabling them to focus on many different factors in delivering a well-coordinated, thorough test plan. The teams had a clear and common goal—to make the 1B processor the best product possible—and their partnership made them much more effective in achieving it.

Conclusion

The close working relationship established between NSD and the SI/PV team was vital to the successful completion of the project. As discussed in this paper—and as exemplified by the high performance of the 1B processor—the robust testing programs developed by these teams contributed greatly to the excellence of the finished product.

Along with providing many valuable improvements, the testing programs increased customer confidence in both the 1B processor's technical capabilities and its readiness for deployment throughout the AT&T switched network. The development and testing phases were enhanced by NSD's involvement. The success of the project is an example of its effectiveness. The daily interfaces established during its completion were a key ingredient in delivering a product having high levels of both performance and quality.

In addition, the hands-on training provided by the ALTP—as well as the close customer and supplier relationships established among the SI/PV team, product design, and NSD—could substantially benefit the sucessful implementation of future hardware-development projects. Many of the processes used in the routines discussed in this paper are now being incorporated into other testing programs. The expectation is that as these ideas and strategies are extended and developed, even more robust testing programs will result.

Acknowledgment

The authors gratefully acknowledge the outstanding dedication and vital contributions of all the NSD 1B and NS SI/PV team members in successfully completing this substantial project.

References

1. T. W. Anderson, P. D. Carestia, J. H. Foster, and M. N. Meyers, "The Evolution of the 4ESS Switch," *AT&T Technical Journal*, Vol. 73, No. 6, November/December 1994, pp. 93-100.
2. Glenford J. Myers, *The Art of Software Testing*, John Wiley and Sons, Inc., New York City, 1979.
3. *Best Current Practices: Development Test Management*, Sofware Quality and Productivity Cabinet, AT&T Bell Laboratories, Murray Hill, New Jersey, January 1991.

(Manuscript approved April 1995)

Margaret A. Hillis *was a technical manager in the 4ESS Customer Support, Testing, and Planning Department at AT&T Network Systems in Naperville, Illinois. During the 1B development project, she was responsible for technical support of the group facilitating processor verification and integration into the 4ESS switch. Ms. Hillis has a B.S. degree in mathematics from St. Norbert College in De Pere, Wisconsin and an M.S. in operations research from the Georgia Institute of Technology in Atlanta. She joined AT&T in 1982.*

Gail W. Minkus *was a technical manager in the 4ESS Customer Support, Testing, and Planning Department at AT&T Network Systems in Naperville, Illinois. At the time the 1B was being developed and tested, she supported the group facilitating processor verification and integration into the 4ESS switch. Ms. Minkus holds a B.S. degree in mathematics and computer science from the University of Illinois in Urbana and an M.S. in engineering–economic systems from Stanford University, Palo Alto, California. She joined AT&T in 1977.*

Craig L. DeCaluwe *was a manager of the 1B Processor Department at the AT&T Network Services Division in Warrenville, Illinois. During the project, he was responsible for ensuring that the 1B processor complied with all requirements as stipulated in the feature specification document, that the installation and deployment processes were flawless, and that all activities were completed on time to meet AT&T network capacity and reliability needs. Mr. DeCaluwe has worked in areas related to the 4ESS switch since 1976. He joined AT&T in 1972.*

Michael N. Meyers *is director of the Network Infrastructure Business Development Department at AT&T Network Systems in Naperville, Illinois. He is responsible for planning all switching, transmission, and operations systems product offerings to meet AT&T affiliates' network infrastructure needs. Mr. Meyers holds a B.A. degree in mathe-*

matics from the University of Illinois in Urbana, as well as an M.S. in computer science and an M.B.A., both from the University of Chicago. He joined AT&T in 1967.

Edward T. Schafer *is district manager of the 1B Processor Department at the AT&T Communications Services Group, Network Services Division (NSD), in Atlanta, Georgia. He is the NSD program manager responsible for the development and deployment of the 1B processor, which includes production, installation, customer acceptance, certification, training, and documentation. He is also responsible for management of project-related information. Mr. Schafer joined AT&T in 1955.*