

PERFORMANCE TESTING

CASE STUDY

Part 1: Understanding the Situation

Ross Collard
rcollard@attglobal.net
Collard & Company

Version 4.9, January 2004

Case Study 1.1
Copyright © 2004 Collard & Company

Understanding the Situation

Introduction to the Case Study

This case study describes a situation in which a performance test is needed. Your assignment is to identify and analyze the issues and develop a test strategy for the situation. This strategy describes how you will predict whether the system's performance is likely to be acceptable when it is placed in live operation.

Organization of this Case Study

The case study is in six parts:

Part 1: Understanding the Situation: In this part (which starts below), you will address the test objectives, justification, scope, technical and business issues, risk factors and overall test approach.

Part 2: Understanding the Situation -- Suggested Answers to Part 1: You will critique suggested answers to the questions in Part 1, and either concur or suggest revisions and improvements to these answers.

Part 3: Reviewing the Proposed Test Strategy: You will review and issue an opinion on a proposed test strategy. This is presented in the form of an executive summary of the strategy and a comprehensive appendix with the supporting details.

Part 4: Reviewing the Consultants' Report: Subject matter experts (the consultants) have rendered an opinion on the adequacy of the proposed test strategy in Part 3. Your job is to determine which of the consultants' findings and conclusions are valid and worth acting on, and what corrective actions to take.

Part 5: Reviewing a "Practical" Test Strategy: In this part, an alternative approach is presented, which was developed based on the consultants' feedback. You will review and compare this alternative to the original proposed approach, and determine which of the two strategies is the most suitable, or what mix of the two you recommend.

Part 6: Developing the Robustness Test Strategy: You will review and critique the proposed strategy for testing the systems' reliability and recoverability. Part 6 is not included in this set of documents.

Understanding the Situation

Learning Objectives

The lessons to be learned from this case study and series of exercises include:

- How to develop a performance test strategy in a typical mixed-technology, mixed-vendor environment with multiple interdependent application systems.
- How to set performance goals and testing objectives.
- How to perform a risk assessment and use it to focus and prioritize the test efforts.
- How to design the test lab, and allow for the differences between the lab and the real-world infrastructure which will be utilized in live operation.
- How to design realistic test work loads.
- How to use automated load testing tools effectively, and minimize opportunities for tool mis-utilization.
- How to decide what to measure and what data to collect.
- How to interpret the harvest of data and form meaningful, trustworthy conclusions.
- Why and how feature testing is fundamentally different from performance testing.

Understanding the Situation

Exercise 1: Determining the Testing Objectives (Allow 60 to 90 minutes for this exercise.)

Instructions

This exercise is the first in a series where you and the other the class participants, working individually and in small teams, analyze performance test issues. The purpose of this first exercise is to understand a typical business situation and consider its performance and robustness test issues.

Read the background to the case study, in the attached Description of the Situation, Part A only. (You do not need to review Part B of the background for this exercise.) Expect that this background reading will take you about 10 to 15 minutes, including a browse through the appendices to Part A. You may say: “Why should I have to read this? The only thing I ever read is TV Guide.” Actually, this reading is important. Testing is context-specific, and we can talk endlessly about test strategy, but there is no substitute for actually getting in there and doing it ourselves. This background reading provides the context – it describes a typical challenging situation you are likely to encounter on the job.

Based on the background (i.e., the description of the situation), answer the questions listed below. Each answer should be no more than a few lines long. In the time available for this exercise you may not be able to get through all the questions. The intention here is not to rush so that all questions are covered, but to think about the issues in developing a performance test strategy. These questions are tough. We do not expect perfect answers, but we would like your best thinking. If you become bogged down on a question, however, it is not worth agonizing over, so move on after a few minutes to the next question.

Questions to Address

- (1.1) What do we want to accomplish with this performance testing project? In the process of determining what we want to accomplish, it is helpful to consider a few other questions:
- o *Why* are we doing this test?
 - o *Who* do we need to satisfy? Who are the vested interests: who have needs and expectations of the test results? These individuals or groups sometimes are called stakeholders or constituencies.
 - o *What* do they want to know?
 - o *What SHOULD* they want to know?
 - o *How* will they use the results of the performance test?

Case Study 1.4

Copyright © 2004 Collard & Company

Understanding the Situation

- (1.2) A set of business objectives are included in the description of the situation. Which of these business objectives can reasonably be addressed in a performance testing project? Which cannot?
- (1.3) Performance goals are stated in the background reading (the Description of the Situation, Part A). Overall, are these performance goals for the system (i) relevant and significant, (ii) realistic (i.e., probably feasible to attain), and (iii) testable or measurable? Has any major goal been omitted?
- (1.4) Performance testing objectives also are stated in the Description of the Situation, Part A. Which of these testing objectives can be directly linked back to one or more specific business objectives? Which cannot? (A performance goal differs from a testing objective: an example of a testing objective is to evaluate whether response time is adequate, whereas the goal is a target for comparison with the measured response time. The goal is the level of response time that the system is expected to meet. It is difficult to evaluate if performance is satisfactory without having reasonably specific goals.)
- (1.5) Are all these testing objectives appropriate, valid and within the scope of the performance testing project?
- (1.6) Are these testing objectives specific enough so that we can evaluate if they have been accomplished? (The specificity of the testing objectives depends on the specificity of the performance goals. If you felt that the stated performance goals are not specific enough, assume for the moment that these goals have been revised and now are realistic and measurable. In other words, answer this question as if the revised goals can be used to evaluate the system's performance.)
- (1.7) What data do we want to collect or measure, in order to evaluate the performance of this system? (In other words, how do we define performance for this system?) Should response time, throughput, availability, error rates or resource utilization be measured, some mix of these, or other characteristics?

Understanding the Situation

Exercise 2: Team Discussion of the Testing Objectives (Allow 45 to 60 minutes for this exercise)

Instructions

Form a team of three to four people for this exercise. Find a comfortable place to gather around and work together. Together with your teammates, compare your answers to the previous questions in Exercise 1. The intention in comparing answers is not necessarily to reach consensus, though that's fine, but to obtain a deeper appreciation of the issues by seeing others' perspectives. Note that we are not looking for polished and detailed answers at this time, just an initial sketch of your thoughts, ready for discussion with the whole class.

Allow about 10 minutes for each question, to discuss them in your team. (Some of these questions could take 3 weeks each to discuss, but we only have limited time in the classroom.) It is OK if you do not get through all the questions in the time allotted, but move on to the next question if you feel that you are becoming bogged down on any one question. At the end of the exercise, be prepared to discuss and justify your team's answers with the class.

Understanding the Situation

Disclaimers

The organization described in this case study is fictional, and is not intended to represent any real-world organization.

Testing tools from various vendors are mentioned in this case study and its accompanying series of exercises. None of these mentions should be construed as an endorsement of a particular vendor or a recommendation of a particular tool.

A Friendly Warning

The questions in the exercise may seem easy at first reading, but they can be fairly difficult to answer. The questions are hard because they require us to use judgment and think the situation through. They are also hard because the information you are provided in this exercise (as you will see in the next few pages) is not complete and perfect -- just like in the real world.

There is some degree of uncertainty, unresolved issues and information which is open to interpretation in this case study -- again, like the real world. Sometimes the testers get into "analysis paralysis" -- they feel that they cannot do anything at all unless they first know everything about the situation. However, we cannot procrastinate and wait until we have perfect information; the test strategy will never be formed.

Ironically, despite the temptation we sometimes have to want to know everything about a situation, attempting to master all the information can be an overwhelming task. A great deal of information is presented in the next few pages, as background for the exercises, and an important part of the test planners' job is to simply to get a handle on this situation. The test planners have to review the information in the next section, determine what's relevant, and sift out what's not. Just like in a real test project.

Because the questions you have been asked to address are not that easy, there is a temptation to disregard them on a real-life project and "forge ahead" regardless. Though we do not want to become paralyzed in indecision, it is not a good idea to skip these questions. If we do not think them through at the beginning of a testing project, sooner or later we will be forced to return and reconsider them, and perhaps will have to change the direction of the testing project.

Understanding the Situation

Description of the Situation, Part A

A.1. Overview

The following sections provide background information needed for the exercises. These sections are labeled:

- A.2. The Business Background
- A.3. Your Responsibilities
- A.4. Basic Functions of the System
- A.5. The Physical Environment
- A.6. The Database
- A.7. Business Operations and Processes
- A.8. Business Objectives for the System
- A.9. Interfacing Systems
- A.10. Performance Goals
- A.11. Performance Testing Objectives

Appendices

Appendix A: System Work Flows

Appendix B: System Architecture

Later, more information is provided in another section, the Description of the Situation, Part B, but this information is not needed for the first exercise.

A.2. The Business Background

Testing Books (TB) is a book club which specializes in selling testing and quality assurance books. The official company slogan is: “Test Geeks ‘R Us.” The senior managers of the book club describe the business as thriving, growing and profitable. They believe the book club has a core of loyal fans that prefer its services to generalists like Barnes & Noble and Amazon, and they want to build on this success by providing even better service and more competitive prices.

To support this business goal, TB is in the process of building a comprehensive new information system to support its core business operation, which is the ordering and shipping of books to its members. This new system will essentially replace the existing automated and manual systems which the book club uses for ordering and shipping books, and also replace an existing Web site which is not considered to be very effective.

A.3. Your Responsibilities

Your job will be to test the performance of this new system, because it is critical to the success

Case Study 1.8

Copyright © 2004 Collard & Company

Understanding the Situation

of the book club. The scope of this testing project includes measuring and evaluating the system response time, throughput, error rates, resource utilization, scalability and ability to handle peak loads. Congratulations (or condolences). You will report directly to the vice president of information systems for the purposes of this project.

Your immediate assignment is to draft a high-level strategy which describes how you will proceed and the approach you recommend for this testing project. You will be presenting and discussing this performance test strategy with the senior managers of Testing Books *next week*. In this presentation, they will want to know how you will test the performance of the system -- not the details as yet, but your overall approach. The managers expect an insightful, cogent analysis, and they are confident that what you say next week will be well thought through, organized and pertinent. (No pressure here at all!)

A.4. Basic Functions and Work Flow of the System

The new system processes book orders: it provides order entry and fulfillment capabilities. With this system, customers can order books directly from the web site and book club employees also can enter orders through a client/server network. The system also manages the fulfillment of these orders. In all, it will support the following five business activities:

- o Ordering of books from the book club.
- o Picking, packing and shipping of books to members in fulfillment of orders.
- o Answering queries on the status of memberships, orders and shipments.
- o Publication of electronic and printed catalogs which show what books are available to order.
- o Reporting of the information needed by the book club executives to manage the business. This information, which is automatically generated by the system, includes order volumes, trends (which books are selling swiftly or slowly), order backlogs, the turnaround times to fulfill orders, book inventory levels, etc., and it is used by the senior managers who are running the business. While the volume of transactions in this category is likely to be relatively low, it is important to be able to provide timely answers.

Please note that this is not a complete list of the business activities and the system functions which support them, just the major ones. The complete list of functions which the system is expected to provide is listed in the system requirements document, and a related set of use cases describe the services provided and how the functions should work. The system requirements and the use cases are not attached to this exercise, but this overview will provide enough information

Understanding the Situation

for you to do the exercises.

To realistically measure performance, testers need to know how the system works. At this stage, you may have many questions about the functionality and work flow, such as: “How does one become a book club member?”, “Can non-members access all parts of the web site?”, “Is this particular function even on the web?”, and “How does this system tie into other systems?”

Appendix A: System Work Flows, briefly explains the system-user interactions. (Later, Appendix C lists the features and outlines how frequently the features will be utilized, but this appendix is not needed to answer the questions in this exercise.)

A.5. The Physical Environment

The new system will run on a multi-tier server architecture which will be shared by other applications needed for the book club’s business operations, such as the billing system. The system environment includes the database, the web site, wireless capabilities, voice call center facilities and high-volume printers. Appendix B: System Architecture describes the technical infrastructure of the operating environment, in terms of the network topology, the existing equipment which can be re-used and incorporated into the new infrastructure, the various types of servers and the support software. You do not need expertise as a system architect, system administrator or network engineer to read this appendix.

A.6. The Database

The system has a centralized database which contains data about books in inventory, members, orders, shipments and bills. A new database is currently being converted from other existing databases of book club members, orders and books. The senior managers have identified the database conversion as a complex, high-risk activity, and they are concerned that several cycles of database tuning may be needed before the system performance goals are met.

When it has been converted from existing files, the database will contain approximately 100,000 active members. (Active members are those who have notified the club of their desire to receive a monthly catalog of books, or who have ordered a book within the last year.) The database will also contain another 200,000 inactive people who have ordered at least one book in the past, but not recently, or who have expressed an interest in the book club. The data on these other, inactive people is used in marketing campaigns. People are deleted from the database after three years with no activity.

A.7. Business Operations and Processes

The book club business is organized into five main groups: (a) senior management, (b) the customer service group, (c) the catalog publishing group, (d) the warehouse distribution group, and (e) the information systems group. Each group has an assigned set of responsibilities, and

Understanding the Situation

there is little or no cross-over of tasks among the five groups.

The *customer service group* works directly with the book club members. Members can access the club's Web site or telephone (speaking to a book club employee), in order to place orders, query order status, make complaints, and change information about their memberships. The Web and phone orders and member data changes are processed while the member is on-line or is on the phone. The customer service group is intended to have approximately 100 personal computers (PCs) for its employees.

The *catalog publishing group* chooses the books for the club to offer, collects book reviews and prepares a monthly catalog of available books, which is either printed and mailed or is distributed via the Internet to the members. (These monthly catalogs do not contain the full list of books available for sale, just the most topical ones.) The catalog publishing group is intended to have 25 PCs to support their work. The entire staff of the catalog publishing group is located at the remote satellite office. (All the other business groups are located in the headquarters building.)

The *warehouse group* distributes books to the club members, either from a member order or as the default monthly selection. A copy of the current monthly book selection is sent to each active member, unless he or she explicitly informs the club that he does not want this book. Stocks of high-demand titles are maintained at the warehouse for filling customer orders. When the supply of any title becomes depleted the warehouse issues a request to replenish the stock to the publisher. A book shipment to a member can originate only from the warehouse -- other users cannot authorize shipments. Personnel at the warehouse pack and ship approximately 2,000 books to members in a typical day, or about 45,000 books a month. The warehouse will have 50 PCs.

The *senior managers* make ad-hoc queries and receive on-line status reports and graphs which help them to manage the business. They also expect to receive regularly printed-out monitoring reports from the new system. There are 25 PCs planned for this group's use, including support staff such as the administrative assistants to the managers. The *information systems group* supports and maintains the computer systems and will be expected to support the new system. This group will have 25 PCs including a test lab.

A.8. The Business Objectives for the System

The overall business goals of the senior managers are to (a) grow revenues, (b) improve profitability, and (c) increase member satisfaction. They have agreed to fund the new system, on the understanding that it will facilitate meeting these goals.

The specific business objectives for the new system are as follows:

Understanding the Situation

- o Support the operations of the book club, by providing the order entry, fulfillment and related features. (These features are described in more detail in the system requirements document.)
- o Improve productivity in the customer service and warehouse distribution groups by 25%.
- o Distribute 95% of all books within 15 working days from the date of an order and at the least available shipping cost.
- o Distribute books with 98% accuracy, i.e., only 2% of the books or less should be returned because the wrong book was shipped to a member or because the member's address was incorrect.
- o Answer 95% of customer telephone calls (e.g., queries about order status), within 3 minutes.
- o Support the projected growth in membership for the next 5 years. Today's system configuration does not have to support the load predicted for 5 years in the future, but the system must be upgradeable to meet the projected demand. The projected growth is described later in this document.

A.9. Interfacing Systems

In addition to the book ordering system, the plan is to migrate other existing application systems to the same technical environment (the client/server network), and run them in this environment also. The main systems which share the same resources are:

- o Member Billing: generates invoices for books ordered from the book club by members, and tracks shipping and payment status of these orders.
- o Publisher Ordering: generates orders for books ordered by the book club from publishers, and tracks their delivery status.
- o Communications: handles internal and external e-mail messages.

A.10. Performance Goals

In order to fulfill the business objectives, the system must meet these performance goals:

- (a) Response times must be satisfactory when the system is operating under a typical realistic load.

Understanding the Situation

- (b) The system must operate correctly when accessed simultaneously by multiple users.
- (c) The system must be able to handle heavy loads.
- (d) Performance and reliability levels must be maintained over an extended period of use (24x365 operation).
- (e) The system availability (uptime) must be adequate in live operation.
- (f) The entire system must be tuned optimally in order to efficiently utilize the computing resources.
- (g) The system must degrade gracefully, not fail catastrophically, when it is pushed up to, and beyond, its planned maximum capacity.
- (h) The system must be scalable, so it can be upgraded in the future to accommodate the projected growth.

A.11. Performance Testing Objectives

The purpose of this performance testing project, as outlined by the senior managers, is to assess the adequacy of the system's performance in live operation. Each of the following testing objectives focuses on whether a particular goal is likely to be met:

- (a) Are response times satisfactory when the system is operating under a realistic load?
 - Response time goals have already been specified in service level agreements (SLAs) which have been negotiated between the IS group and the business units. These service level agreements are described in the Description of the Situation, Part B. As mentioned earlier, you do not need to review Part B in order to answer the questions in Exercises 1 and 2.
 - The guideline for response time consistency has been set as plus or minus 25% of the average. That is, if the average response time for a transaction is expected to be 2 seconds, then 90% or more of the measured response times should fall within the range of 1.5 and 2.5 seconds.
- (b) Does the system operate correctly when accessed simultaneously by multiple users? Problems such as features interfering with each other, database optimization, resource contention and transaction priorities may need to be considered here.
- (c) How well does the system handle heavy loads? We cannot assume that the performance degradation from adding additional users or performing extra work is linear: a significant increase in response time may occur when only a few users are added or the workload

Understanding the Situation

increases by an increment.

- (d) Will performance and reliability levels be maintained over an extended period of use? Insidious problems such as memory leaks will not reveal themselves in short run tests or by testing with a small number of users. Such problems usually lead to performance degradation and eventually to system failures.
- (e) Will the system availability (uptime) be adequate in live operation? Apart from the planned downtime for maintenance, is there reasonable confidence that the unplanned downtime will be held to acceptable levels?
- (f) Is the entire system tuned optimally? Successful tuning requires the empirical testing of server and client configurations using carefully controlled workloads.
 - Are upgrades to dual-processor or quad-processor servers required to achieve the performance goals?
 - On the other hand, an optimally tuned system may allow the book club to increase the number of users without costly hardware upgrades. (Experience has shown us that very few systems are initially tuned for optimum performance.)
- (g) Does the system degrade gracefully, or fail catastrophically, when it is pushed up to, and beyond, its planned maximum capacity?
- (h) Is the system scalable: can it be upgraded in the future to accommodate the projected growth over the next 3 to 5 years, without major software re-writes or a major re-structuring or conversion of the database?

Understanding the Situation

Appendices to Part A

These appendices provide supplementary details about the book club's situation. While the details are worth reviewing, it is not necessary to analyze them in depth in order to answer the questions. Appendix A outlines the main interactions among the system and its users. Appendix B describes the technical environment in which the system will operate.

Appendix A: System Work Flows

People can order books either through the web site, or by phone or mail request. Using the web site, interested parties can search for books by topic, author, etc., query the availability and price of a book, and query the status of an order in progress. If they choose, people can enroll as book club members via the web site, phone or mail. A person does not have to become a member to order a book, but members get special privileges such as early notification of sales and occasional discounts. There are no membership fees. Book club members receive monthly catalogs and promotions, either electronically or in printed form.

An order can contain books from various publishers, with different quantities of each book ordered. Orders can be modified or cancelled at any time up until the day of shipment. If a book is not in stock, the system informs the person (or the internal book club employee) and asks if he or she wants to place a backorder.

A database record is created for every new order, and this record is updated to track the progression of the order through to fulfillment. The ordering system maintains records on the book inventory as well as on orders. When an order is entered, the system generates directives to the warehouse staff to pick, pack and ship the order, as well as printing the paperwork needed to ship the order. The system provides the capability to query and update the inventory of books on hand, and automatically decrements the inventory as books are shipped. Ordering books from publishers and book distributors is not part of this system.

The ordering system itself does not generate bills or process payments, but triggers these actions by the separate billing system. People usually pay by credit card, and members with acceptable credit history can be billed for payment within 30 days. To collect payment, the ordering system sends a transaction to the separate billing system. In the event of returned books, the billing system issues refunds but these do not flow through the ordering system. The system also provides the capability for the catalog publishing group to update book information such as descriptions and reviews, and to compose and publish the catalogs.

Understanding the Situation

Appendix B: System Architecture

a. Architecture Overview

The system architecture is based on these principles: (a) the business drives the services, and the services drive the technology; (b) agility to respond to change is a fundamental business requirement, and the architecture is expected to be in flux on an on-going basis; (c) the system is as platform-independent as possible; (d) the system is loosely coupled, in order to minimize bug propagation (so that a code change on one server, for example, does not necessitate code changes on other servers); (e) interfaces will be mainstream and standards-based (i.e., only reasonably mature, non-proprietary interfaces will be utilized); (f) and the system is designed for robustness (e.g., fault-tolerant with fail-over capability, load balancing, redundancy and reserve capacity).

The system design utilizes many small, dedicated-function servers. These include database, web, wireless, application, print, voice telephone, network and remote location servers. The design provides flexibility and agility; improves reliability through high server redundancy, and facilitates troubleshooting by simplifying problem isolation and helping to pinpoint the location of a problem to the nearest server. The remainder of this section describes the planned initial system configuration, which has been derived from a business model (in the form of a set of use cases which describe the services provided). Later in the Description of the Situation, Part B, an operational profile outlines how frequently the services will be utilized. (While you are welcome to peek, you do not need to review Part B in order to do Exercises 1 and 2.)

b. Internal Network Topology

While book club members and visitors can access the system only through the web site, the internal part of system used by the book club employees will operate on a client/server network. This network will include approximately 250 users, with 225 in fixed locations with one workstation per user, and 25 wireless users with hand-held devices.

Of the 225 fixed-location workstations, 200 will be physically located at the book club headquarters, which is an office building in Los Angeles, and will be on a local area network (LAN). The other 25 fixed-location workstations will be located at a satellite office, which is several miles away and which will be connected to the headquarters building by a wide area network (WAN). The 25 wireless devices will be used in the warehouse, which is attached to the Los Angeles headquarters building.

The client workstations will be supported by approximately 16 servers. Following the system architecture principles expressed earlier, each server will be dedicated to a specialized use and coupled to another redundant, load-sharing server for the same specialized use. It is anticipated that the server configuration will include two database servers, two application servers, two Web servers, two voice telephone CTI (computer telephony) servers, two held-held device servers, two

Understanding the Situation

print servers, two network servers, and two servers remotely located at the satellite office. This server configuration may possibly change as the system design is refined and finalized, or as necessary to balance and tune the system, or as and when the system load characteristics change.

The servers in the cluster are intended to fail-over, that is, to back each other up as necessary without a loss of functionality. In the event that one server is disabled, the client/server network should remain operational, even if it provides a degraded (slower) level of performance to the users until full operation can be restored.

Initially, each server is intended to have a single processor. However, the flexibility has been left in the budget to upgrade selected servers to dual-processor or even quad-processor machines, if necessary, depending on the results of the performance testing and the performance tuning and debugging.

c. The Existing Infrastructure

The parts of existing infrastructure which can be re-used, namely the equipment and facilities which are already in place and support the existing application systems, will be integrated into this new environment, as described above. The equipment which will be carried over to the new operating environment has already been included and counted in the inventory of equipment for the new system, which is described in this section.

The parts of the existing infrastructure which cannot be re-used, such as older clients and servers, will be discarded after the new environment is operational.

There are no hand-held devices or wireless systems currently being used by the book club.

d. The Database Servers

Each of the two database servers will contain a complete copy of the database, and each will mirror the database on the other database server. In other words, all updates will be made concurrently to both copies of the databases, and these copies will be periodically and automatically monitored for consistency by the database management system (DBMS). The database servers will use RAID (redundant disk) technology, so that each database server in itself has redundant copies of the data.

A complete back-up copy of the database will be taken once every night. During this back-up process, which is expected to be completed within a period of 30 minutes, user response times may be slow. Few users are expected at the time, probably 3.00 am, so that slow responses are a minor issue. No guarantees have been made to the user community as to the worst-case performance of the system in this situation. Nevertheless, since the managers want to know if the user response times are likely to be extremely slow during the database back-up, measuring

Understanding the Situation

the user response times during the back-up is within the scope of this project.

e. The Application Servers

A full set of the application systems will reside on each of the two application servers, so that both servers will have the capability to support all the system features and process all transactions.

Load balancing software will be employed to automatically share (balance) the load among the servers: it will route transactions between the two application servers, to route query transactions between the two database servers, and also to route print requests between the two print servers, in order to optimize performance. This load balancing capability is expected to be physically located prior to the application and database servers, and since the system design includes no front-end proxy servers it thus will reside on the web servers... It is anticipated that no extra add-on load balancing tool will be utilized; the balancing will be done by the built-in capabilities of the server operating system.

f. The Print Servers

The print servers are connected to eight high speed laser printers, which will be shared by several users, and which are located in the work areas close to their main intended users. These printers are used to print in high volume the paper copies of the book club catalog. Although the catalog is available on-line, printed copies are available on request. Approximately 5,000 copies of the 300-page catalog are printed and distributed per month. The printers are also used for a variety of internal business reports.

In addition, dedicated local printers will be directly connected to individual personal computers and normally will not be shared across the network. Measuring the print times on these local printers has been declared to be outside the scope of this test project. Though the performance testers are not interested in the level of service provided, they will need to consider the impact, if any, of the localized printing on overall network performance. If network bandwidth is tight enough for local printing to be a problem, even with large print runs, then the system probably will have other bandwidth-related bottlenecks too. Testing the local print response is relatively simple and could provide early feedback on the existence of bandwidth problem areas.

g. The Web Servers

The Web servers will be connected to routers, which connect to the Internet through multiple T1 lines. Each T1 line has a rated throughput of 1.544 megabits per second (Mbps), or sufficient capacity to accommodate up to 25 simultaneous sessions using 56 kbps modems, assuming no data compression technology is used. The book club expects that the Web servers will be accessed from a variety of browsers on a variety of remote platforms. There will be firewalls on

Understanding the Situation

the Web servers which may slow their performance. Since not all the security software tools and controls have as yet been determined, the impact of the security controls on the system performance is not as yet fully known.

The Web servers will have their own databases, physically located on these servers. The Web databases will contain redundant copies of the same data and will be used to store frequently downloaded data, such as FAQs, and to help isolate the Web site visitors from the internal systems of the book club.

The Web server databases will not contain all the data needed to answer queries or process orders. The Web servers will need to access the database servers for some data which is not available on the Web servers. (Although the Web servers are physically connected to the database servers, some of the confidential company data on the internal databases is intended to be essentially invisible to Web visitors for security reasons -- in other words, this data cannot be accessed by the Web servers.)

h. The Voice Telephone Servers

The voice telephone servers will be used to coordinate incoming calls with customers' profiles, so that the person answering a call has ready access to the caller's account status. For example, a customer's data will be retrieved automatically, based on the telephone number of an incoming call, and displayed on the user's screen as he or she answers the call.

The voice telephone service will be supported by call center software which resides on the voice response servers. The voice service is handled by the same set of T1 lines which handle the Web service.

i. The Wireless Servers

The wireless servers will be used to communicate with the hand-held devices in the warehouse. These devices will help direct the order fulfillment, specifically the picking, packing and shipping of books.

In the future, the hand-held devices will be shared with other application systems in addition to the order processing system. The main other application is expected to be inventory management, where the hand-held devices will be used to take and report inventory counts and report out-of-stock conditions directly from the warehouse floor.

The inventory management system is not expected to be implemented until at least 9 months after the order processing system goes live. When the inventory management system is fully operational, it is expected to account for approximately 25% of the total wireless message traffic. (The order processing system will account for the other 75%.)

Case Study 1.19

Copyright © 2003 Collard & Company

Understanding the Situation

j. The Remote Location Servers

Two servers will be situated at the satellite office. These servers will support a local area network for the client workstations in the satellite office, and will connect the satellite office to headquarters by a wide area network.

The business group which is located at the satellite office, and thus is the main group supported by these servers, is the catalog publishing group. The role of this group is described in a later section of this document.

k. Network Technologies

Internally within the headquarters building, the fixed-location clients and servers will be connected by a fast Ethernet local area network (LAN), using the 100Base-T standard and rated at 100 Mbps. If the performance measurements indicate that the LAN is a major bottleneck, the speed of the LAN could possibly be increased to a gigabit per second (using gigabit Ethernet). However, an allowance for this contingency has not been included in the project budget.

The Web and CTI (telephone) servers will be connected to a set of T1 lines, which can be shared by both the Web and voice telephone traffic. These lines have a rated capacity of 1.544 megabits per second for each line. The performance of the protocol conversions from external networks, such as SONET, to internal Ethernet LAN, is assumed not to be a significant issue and will not be measured as part of this project.

Internally within the satellite office building, the clients and servers will be connected by a traditional Ethernet LAN, rated at 100 Mbps.

The wide area network (WAN) connecting the headquarters and the satellite office will utilize one or more dedicated T1 lines.

l. Software

The decision has been made to use Windows XP on both the servers and on the clients.

The database management software used by the book club will be Oracle 9i.

The Web servers will run Apache software to service Internet traffic.

m. Physical Installation

A contractor is in the process of installing and wiring the computers and network equipment

Case Study 1.20

Copyright © 2003 Collard & Company

Understanding the Situation

needed for the new system.

This contractor will test and confirm that this equipment has been installed and works, and also will confirm that the support software, such as Windows 2000 and device drivers, has been installed correctly on the equipment.

Understanding the Situation

Exercise 3: Calculating the Work Load (Allow 20 to 30 minutes for this exercise)

Instructions

Answer these questions, based on the following set of definitions and assumptions and on the data presented in the case study.

Questions

- (3.1) How many web site visits or sessions are expected per month?
- (3.2) How many web site visits or sessions are expected in a typical hour?
- (3.3) How many web site visits or sessions are expected in a peak-demand hour?
- (3.4) What is the expected peak number of concurrent visitors in a typical week?
- (3.5) What is the expected peak number of concurrent visitors in a typical month?
- (3.6) How many hits will happen to the home page in a typical hour?
- (3.7) How many hits will happen to the home page in the peak hour in a typical month?
- (3.8) What other significant assumptions besides the ones listed below, if any, have you made in this exercise?

Definitions of Terms

Work Load: The mix of demands placed on the application or system under test. Depending on the situation, these demands may be database queries, events, transactions, web page download requests, messages, and so on.

Hit: A single event, activity or transaction; typically a request to download a web page.

Visit or Session: The series of individual events or activities (usually the externally visible interactions between one user and the system), which occur from the time when the user logs on to the system until the same user logs off.

Concurrency: The occurrence of two or more events or activities at the same instant or during the same brief time interval. If two events are not simultaneous, to be considered concurrent they should occur closely enough in time to have a non-trivial chance of competing for the same resources, or of interacting and possibly interfering with each other.

Case Study 1.22

Copyright © 2003 Collard & Company

Understanding the Situation

Concurrent Users (or Visitors): The total number of overlapping users who are actually using the system or who have active sessions at a specific instant in time. These users are all connected to the system, may be in different operating states (ready, waiting for response, running), and may or may not be performing the same tasks. The system typically must dedicate system resources to manage the session of each connected user.

Hourly Users (or Visits or Sessions): The total number of users who have an open active session during the span of an hour.

Session Length or Duration: The total elapsed time that a single user is using the system during a single session or visit.

Assumptions

Use the following assumptions to calculate the work load. Please document any new assumptions you make.

The 100,000 active users (book club members) will visit the web site, on average twice per month. Another 15,000 causal visitors (non-members) will access the site per month.

A site visit, or session, lasts an average of 15 minutes.

While session lengths can be indefinitely long, 95% of sessions lengths fall in the range from 3 minutes to 45 minutes.

Since calculating the number of overlapping sessions -- what many people call concurrent users -- is complex, assume that this approximation is sufficiently accurate: the number of concurrent users is equal to the number of hourly users multiplied by the average session length (in hours).

Each visit or session accesses and downloads the home page once on average.

The book store is open every day of the month (30 days on average), and the volume of traffic is approximately the same from day to day.

This site is visited mostly by U.S. residents, and primarily for personal use, not business use. The hours of reasonable use are 10 am EST thru 12 pm PST (or 17 hours a day). Not many people use the system at 3 am.

The peak hour traffic is 20% of the total day's traffic.

The expected peak number of concurrent users in a typical week is expected to be 4 times the

Understanding the Situation

number in the daily peak hour.

The expected peak number of concurrent users in a typical month is expected to be 8 times the daily peak.

Understanding the Situation

Exercise 4: Outlining the Test Environment (Allow 30 to 45 minutes for this exercise)

Instructions

The purpose of this exercise is to identify and analyze issues which tend to be important in planning and organizing a performance and robustness test project.

Read the background to the case study, in the attached Description of the Situation, Part B. (Expect that this background reading will take you about 10 to 15 minutes.) Review your earlier answer to question (1.7): “What data do we want to collect or measure, in order to evaluate the performance of this system?” Revise your answer to this question as appropriate. Then answer the questions below.

Questions to Address

- (4.1) Where – at what points or locations -- will the measurements be taken? (In other words, at what points internally within the system or externally to the system do we want to gather data or monitor the system’s behavior? What performance data needs to be captured at which locations? We do not need a precise answer yet, just a general sense of where to look)
- (4.2) What loads or mixes of demands should we place on this system, while measuring its performance, such as an average load, a peak load or an overload?
- (4.3) What types of load testing should we utilize in our strategy, such as spike and bounce, hot spot or duration testing?
- (4.4) How do we generate and drive these loads?
- (4.5) What environment (test equipment and facilities) do we need for the performance testing?
- (4.6) What kinds of tools are needed for this performance testing, if any? (At first glance this question may seem to be the same as question (10) above. The former question focuses on the core tools for load generation. By contrast, this question (12) addresses the broader picture, including all the other major tools that the performance testers need.)
- (4.7) What skills are required for this performance test? What kinds of people should be on the performance test team?
- (4.8) How should the data (the test results) be analyzed and interpreted? How do we determine, based on the measurements, whether the performance objectives of the system

Case Study 1.25

Copyright © 2003 Collard & Company

Understanding the Situation

are likely to be met in real-world operation of the system?

- (4.9) What are the likely difficulties involved with this performance test project? What are the risks, if any, that the performance test will not deliver trustworthy or usable results?
- (4.10) What additional information, if any, do you need to know before you can develop an adequate performance test strategy?
- (4.11) What assumptions have you made about this situation? Do these assumptions need to be validated, and if so, how can we validate them?
- (4.12) Some of the senior managers are skeptical about the performance testing project, and believe in the “fat server” approach instead. How would you respond to the suggestion that we skip the performance testing, save the time and money needed for this effort, and simply “beef up” the servers, clients and communications channels as and when necessary in live operation?
- (4.13) Do we need an initial impact assessment (IIA) for this project? An IIA assesses the need for a performance test based on limited information about the situation.

Description of the Situation, Part B

This section addresses these topics:

- B.1. System Usage
 - B.2. Growth Projections
 - B.3. Service Level Agreements
 - B.4. System Development and Feature Testing Methodology
 - B.5. Automated Test Facilities
 - B.6. Conditions and Constraints
- Appendix C: Feature List and Operational Profile

B.1. System Usage

The transactions processed by the order processing system and other related systems are listed below, together with their expected frequencies of utilization. The transaction volumes are based on actual counts of transactions within the existing system, which have been adjusted for the expected differences between the existing system and the new one. (There is not always a one-to-one relationship between the transactions in the existing system and the new one.) These counts reflect today’s use and do not include the project growth in the future, and do not reflect the anticipated shift from primarily internal ordering (by book club employees) to primarily external (by members and visitors through the web site). The peak hours of demand for the business groups may not all happen at the same time. The following table summarizes the

Case Study 1.26

Copyright © 2003 Collard & Company

Understanding the Situation

demand:

User Group / Feature or Transaction	Frequency of Utilization (Transactions per hour)	
	Normal hour	Peak hour of month
<i>Customer service group</i>		
Web site hits by external users	950	7,750
Demands generated internally	250	1,500
<i>Catalog publishing</i>		
Monthly catalog of books		
Printed catalog	-	5,000
Distributed via the Internet	-	15,000
<i>Warehouse group</i>		
Shipping and bill of lading papers	150	1,000
Picking instructions	250	1,250
<i>Information systems group</i> (not including periodic database back-ups)		
	250	6,000
<i>Senior management group</i>		
Queries and reports	Under 25	1,750

See Appendix C for the supporting details behind this table.

Occurrences of Peaks

Information about when the peak demands occur during a week or month (i.e., at which specific hours), for each feature and type of system use, is not yet available.

B.2. Growth Projections

The system must be scalable or upgradeable to support book club member growth rates of 15% per year, and order volume growth rates of 25% per year, for the next four years.

The growth is expected to be achieved primarily by providing an enhanced Web site for book club members and potential members. For this reason, the mix of books sold over the Internet versus through the traditional channels is expected to change, as follows:

Time Frame	Percentage of Books Sold	
	By Traditional Channels	By Internet
This year	80%	20%
2 years ahead	50%	50%

Case Study 1.27

Copyright © 2003 Collard & Company

Understanding the Situation

4 years ahead

20%

80%

This changing mix of demands and expected shift of work from the telephones to the Web means that the internally generated workload in response to telephone calls is expected to fall over the next 5 years, not rise, even after allowing for the overall growth in business volumes.

B.3. Service Level Agreements (SLAs)

A separate service level agreement has been developed for each of the three operational groups, the customer service group, the catalog publishing group, the warehouse distribution group. Each service level agreement has been signed by the manager of the pertinent business group and by the manager of the information systems group. The adherence to the service level agreements will be monitored and reported to senior management monthly, and the information systems group will be evaluated in part on their ability to satisfy the service level agreements.

The following statements summarize the key points in each of these service level agreements. The agreements state that, for priority 1 features under normal operating conditions (e.g., with an average workload), the internal users' response times should be less than 0.5 seconds, 90% of the time or better, on the client machines. (Priority 1 is the highest priority.) Under peak load, the internal users' response time should be 1 second or less, 90% of the time. This response time goal must be met for each business unit individually, such as the customer service group, the catalog publishing group, and the warehouse distribution group. In other words, it is not sufficient if the average response time across all the business units meets this goal, while the average for any one business does not meet the goal.

For priority 2 features, the internal users' response times should be less than 2 seconds, 90% of the time or better. For features with priority levels 3 or more, no response time guidelines have been established as part of the SLAs.

For the external users who visit the Web site, the requirement is that the book club's Web site is "noticeably faster" than competitors' Web sites. It has been determined that this means that the book club's home page on the Web site can be downloaded in no more than 4 seconds on average, under an average work load, and no more than 10 seconds under peak load. These times are measured from when a user clicks to initiate an action until when the user starts to see a response, e.g., when a page starts to be rendered, not when the page is fully visible.

These response time targets are the averages for all Web site visitors. (A person using a 9.6 Kbps modem will have a slower response than someone with a dedicated T1 line.)

The book club also anticipates that these Web-based response time targets will need to be revised at least once every 6 months, and tightened, as competitors' Web services become progressively faster.

Case Study 1.28

Copyright © 2003 Collard & Company

Understanding the Situation

On-going processes will need to be put in place, to monitor and report the level of compliance with the SLAs during the on-going live operation of the system. Developing these mechanisms to monitor and report on SLA compliance is not part of the scope of the performance test project.

The service level agreements also address other areas besides response time and throughput, such as acceptable levels of errors in processing, system availability, etc., but these areas are outside the scope of the performance testing.

B.4. System Development and Feature Testing Methodology

The new book club order processing system is currently in the process of being developed. Where possible, the developers are re-using existing software components, which have either been purchased from external sources or derived from other internal systems.

Rapid application development (RAD), which is also called the iterative spiral approach, is being used by the developers. In this approach, the feature testing starts early and overlaps development, and components and sub-assemblies of components are tested as they become available, and then re-tested as the components are debugged or enhanced.

There is a feature testing team, which has been organized to test the features of the new order processing system. This team is not responsible for performance, load or stress testing, but its members have been asked to cooperate and provide assistance to the performance test team, as appropriate and without interfering with the feature test project.

The developers and testers are planning to build component-level test drivers, in order to start unit-testing each major software component as soon as it becomes available.

B.5. Automated Test Facilities

This feature testing team is in the process of building a library of automated feature test cases for each of the features listed above. Depending on the criticality, risk and complexity of each feature, the library is expected to contain anywhere from one to ten test cases for each feature.

These test cases are being built and will be executed and maintained by using WinRunner, a tool from Mercury Interactive of Sunnyvale, CA. The feature testing team has purchased five copies of WinRunner, which could be available for performance testing when the feature testers are not using them.

Note that the mention of any particular tool in this case study, such as WinRunner, should not be construed to be a recommendation or endorsement of that tool.

Understanding the Situation

B.6. Conditions and Constraints

At this time, the senior managers have not set a deadline or budget limit for this performance testing project. This does not mean that they will be willing to accept an indefinitely long time frame or an indefinitely large budget for this project, but they are waiting until they hear what you propose.

Understanding the Situation

Appendix C: Feature List and Operational Profile

User Group / Feature or Transaction	Frequency of Utilization (Transactions per hour)			Priority (*)
	Normal use	Peak use (**)		
		Week	Month	
<i>Customer service group</i>				
(a) Web site hits by external users (direct external demands)				
Main or home page only	450	1,800	3,600	1
Book query (availability, price)	250	1,000	2,000	1
Book order	50	250	500	2
Credit card authorization	50	250	500	2
Query status of existing order	25	50	75	3
Add new member (sign up)	5	15	50	1
Change membership information	5	15	50	3
Link to another site	125	500	1,000	2

These transactions are all initiated from external users via the Internet.

- (b) Demands generated internally in response to telephone calls (the average duration of these telephone calls is 12 minutes each)

Book query	100	300	600	1
Book order	50	200	400	2
Credit card authorization	35	100	350	2
Status of existing order	10	25	100	2
Add new member	1	5	10	1
Change membership information	1	5	10	3
Delete existing member	0.5	2	5	3
Complaint	0.5	2	25	1

These transactions are all initiated from the personal computers in the customer service group.

(*) Priority of each type of transaction, where 1 is the highest and 4 is the lowest priority.

(**) These peaks represent the expected load during the worst-case hour in a typical week, and in a typical month, in the two columns respectively. The peaks for all these types of transactions are not expected to occur within the same one hour of the week or the month.

Understanding the Situation

User Group / Feature or Transaction	Frequency of Utilization (Transactions per hour)			Priority (*)
	Normal use	Peak use (**)		
		<i>Week</i>	<i>Month</i>	
<i>Catalog publishing group</i>				
Book list				
Add book to list of available books	2	50	100	3
Change book information	1	50	100	3
Delete book	1	10	20	3
Add book review	5	25	50	3
Change book review	1	5	10	3
Monthly catalog of available books				
Printed catalog (***)	0	5,000	5,000	4
Distributed via the Internet (***)	0	15,000	15,000	4

These transactions are all initiated from the personal computers in the catalog publishing group.

(***) These loadings on the system happen only in occasional bursts; normally the transaction traffic is zero.

Warehouse group

Shipping and bill of lading (Shipping papers and invoice)	150	500	1,000	3
Inventory management (****)				
Inventory query (Internal levels of inventory)	10	50	100	2
Publisher query (e.g., review of new book titles)	10	50	100	2
Publisher order	5	25	50	3
Inventory update (Receipt of new book data from publishers)	10	500	500	3

These transactions are all initiated from the personal computers in the warehouse group.

(****) These transactions are not processed by the ordering system, but by other systems which share the same resources.

Understanding the Situation

User Group / Feature or Transaction	Frequency of Utilization (Transactions per hour)			Priority (*)
	Normal use	Peak use (**)		
		Week	Month	
<i>Warehouse group (continued)</i>				
Order fulfillment				
Picking instructions and confirmations	250	750	1,250	2

These transactions are all initiated from the hand-held wireless devices in the warehouse group.

*Information systems group (****)*

System development and maintenance	50	250	500	2
Internal e-mail traffic -- among all departments	150	1,000	5,000	2
System administration				
On-going administration	50	250	500	1
Periodic (e.g., database back-up)	0	50,000	50,000	4

These transactions, except the e-mail, are all initiated from the personal computers in the IS group. The e-mail traffic is distributed equally across all the clients in the network.

Senior management group

Ad-hoc query	0	250	500	1
Daily on-line status report	10	100	250	2
Weekly status summary	5	500	500	4
Monthly financial and operating data (****)	0	0	500	4

These transactions are all initiated from the personal computers in the senior management group.

Other Systems and Features

The ordering system interacts and shares resources with other systems, such as billing, publisher ordering, market data analysis (data mining) and communications (e-mail). The work loads listed above include overhead and background transactions like e-mail between the book club employees, and features which are not part of the ordering system but which are hosted on the same application servers. Other system features have their own application servers, which are not included in the equipment inventory in Appendix B. The work loads for these features are as follows:

Understanding the Situation

User Group / Feature or Transaction	Frequency of Utilization (Transactions per hour)			Priority (*)
	Normal use	Peak use (**)		
		Week	Month	
<i>Billing group</i>				
Generate e-bill	50	250	500	1
Print bill	50	250	500	1
Generate reminder notice	10	25	150	1
Process payment	50	250	500	2
Adjust bill or payment	10	25	150	2
<i>Publisher ordering group</i>				
Order from publisher	10	25	150	1
Receive and store	10	25	150	1
Pay publisher	10	25	150	2
Adjust order or receipt	1	10	50	2
<i>Marketing group</i>				
Data mining query	10	25	250	5

Exercise 5: Team Discussion of the Test Environment (Allow 30 to 45 minutes for this exercise)

Instructions

Together with your teammates, compare your answers to the previous questions. The intention is not necessarily to reach consensus, though that's fine, but to obtain a deeper appreciation of the issues by seeing others' perspectives. Note that we are not looking for polished and detailed answers at this time, just an initial sketch of your key thoughts, ready for discussion with the whole class.

Microsoft and Me

Microsoft has published a similar performance testing case study, for a fictional firm called the Duwamish Book Store. The Web-based business functions supported by the system include point-of-sale, order entry, shipping and receiving, and a book catalog. My case study in this book was published in a substantially complete version well before Microsoft published theirs. The Duwamish Book Store case study is worth reviewing as a counterpoint to this one. Microsoft's case study is simpler and less rich in issues, so it may be worth looking at first, before you get into the complexities and nuances of this case study.