At first glance, the subject "performance requirements" looks simple enough. Almost every book or paper about performance has a few pages about defining performance requirements. Quite often, a performance requirements section can be found in project documentation. But, the more you examine the area of performance requirements, the more questions and issues arise. The goal of this presentation is not to answer all the questions, but rather to discuss and evaluate the many aspects of this topic and to provide different views and examples.

Introduction

- Looks simple from the first look
- Some topics deeply investigated, some hardly touched
- Points to ponder, not tutorial
  - Metrics
  - Defining
  - Verification

Defining performance requirements is an important part of system design and development. If there are no written performance requirements, it means that they exist in the heads of stakeholders, but nobody bothered to write them down and make sure that everybody agrees on them. Exactly what is specified may vary significantly depending on the system and environment, but the final requirements should be quantitative and measurable. Performance requirements are the main input for performance testing (where they are verified) as well as capacity planning and production monitoring.

This presentation is not a tutorial, it is rather a discussion about important points to ponder that are often missed when performance requirements are discussed. First we discuss different performance metrics, then how we can gather performance requirements and finally the issues related to verification of performance requirements.

Agenda

- **<u>Metrics</u>**

- **Defining**

- **Verification**

Traditional Performance Requirements

- **For today's distributed business systems**
- **It depends**
  - Quantitative and measurable
- **Response / processing times**
- **Throughput**
- **All are important**

There are several classes of performance requirements. Most traditional are response times (how fast the system can handle individual requests) and throughput (how many requests the system can handle). All classes are vital: Good throughput with long response times usually is unacceptable, as are good response times with low throughput.

Response Times

- How fast requests are processed
- Depends on context
  - 30 minutes may be excellent for a large batch job
- Depends on workload
  - Conditions should be defined
- Aggregate metrics usually used
  - Average, percentiles, etc.

Response times (in the case of interactive work) or processing times (in the case of batch jobs or scheduled activities) define how fast requests should be processed. Acceptable response times should be defined in each particular case. A time of 30 minutes could be excellent for a big batch job, but absolutely unacceptable for accessing a Web page in a customer portal. Response times depend on workload, so it is necessary to define conditions under which specific response times should be achieved; for example, a single user, average load or peak load.

Significant research has been done to define what the response time should be for interactive systems, mainly from two points of view: what response time is necessary to achieve optimal user's performance (for tasks like entering orders) and what response time is necessary to avoid Web site abandonment (for the Internet). Most researchers agreed that for most interactive applications there is no point in making the response time faster than one to two seconds, and it is helpful to provide an indicator (like a progress bar) if it takes more than eight to 10 seconds.

Response times for each individual transaction vary, so we need to use some aggregate values when specifying performance requirements, such as averages or percentiles (for example, 90 percent of response times are less than X). Maximum / timeout times should be provided also, if necessary.

For batch jobs, remember to specify all schedule-related information, including frequency (how often the job will be run), time window, dependency on other jobs and dependent jobs (and their respective time windows to see how changes in one job may impact others).

Throughput is the rate at which incoming requests are completed. Throughput defines the load on the system and is measured in operations per time period. It may be the number of transactions per second or the number of adjudicated claims per hour.

Defining throughput may be pretty straightforward for a system doing the same type of business operations all the time like processing orders or printing reports. Still additional metrics may be necessary in some cases, such as the number of items in an order or the size of a report. It may be more difficult for systems with complex workloads: The ratio of different types of requests can change with the time and season.

It is also important to trace how throughput varies with time. For example, throughput can be defined for a typical hour, peak hour, and non-peak hour for each particular kind of load. In some cases, it may be helpful to further detail what the load is hour-by-hour.

Number of Users

- **Number of users by itself doesn't define throughput**
  - Without defining what each user is doing and how intensely
  - 500 users running one short query each minute: throughput 30,000 queries per hour
  - 500 users running one short query each hour: throughput 500 queries per hour
  - Same 500 users, 60X difference between loads

The number of users doesn't, by itself, define throughput. Without defining what each user is doing and how intensely (i.e. throughput for one user), the number of users doesn't make much sense as a measure of load. For example, if 500 users are each running one short query each minute, we have throughput of 30,000 queries per hour. If the same 500 users are running the same queries, but only one query per hour, the throughput is 500 queries per hour. So there may be the same 500 users, but a 60X difference between loads (and at least the same difference in hardware requirements for the application – probably more, considering that not many systems achieve linear scalability).

Not So Traditional Performance Requirements

- **For today's distributed business systems**
- **Concurrency**
- **Resources**
- **Scalability**
- **Context**

While some requirements are considered traditional and absolutely necessary for some kind of systems and environments, they are often missed or not elaborated enough for interactive distributed systems.

•Concurrency

•Resources

•Scalability

•Context

Concurrency

- **Number of simultaneous users or threads**
- **Take resources even if doing nothing**
- **Named, <u>active</u>, concurrent users**
- **Trap: "really concurrent" users**
  - **Example: requirement 20 "concurrent" users**
  - **Request takes 10 seconds, throughput 120 requests/minute**
  - **Improve to one sec, same throughput - two "concurrent" users**

Concurrency is the number of simultaneous users or threads. It is important: connected, but inactive users still hold some resources. For example, the requirement may be to support up to 300 active users.

The terminology used to describe the number of users is somewhat vague. Usually three metrics are used:

 • Total or named users. All registered or potential users. This is a metric of data the system works with. It also indicates the upper potential limit of concurrency.

 • Active or concurrent users. Users logged in at a specific moment of time. This one is the real measure of concurrency in the sense it is used here.

 • Really concurrent. Users actually running requests at the same time. While that metric looks appealing and is used quite often, it is almost impossible to measure and rather confusing: the number of "really concurrent" requests depends on the processing time for this request. For example, let's assume that we got a requirement to support up to 20 "concurrent" users. If one request takes 10 sec, 20 "concurrent" requests mean throughput of 120 requests per minute. But here we get an absurd situation that if we improve processing time from 10 to one second and keep the same throughput, we miss our requirement because we have only two "concurrent" users. To support 20 "concurrent" users with a one-second response time, we really need to increase throughput 10 times to 1,200 requests per minute.

It is important to understand what users you are discussing: the difference between each of these three metrics for some systems may be drastic. The number of online users (the number of parallel session) looks like the best metric for concurrency (complementing throughput and response time requirements). Finding the number of concurrent users for a new system can be tricky, but information about real usage of similar systems can help to make the first estimate.

Resources

- CPU, I/O, memory, and network
- Resource Utilization
  - Related to a particular configuration
  - Often generic policies like CPU below 70%
- Relative values (in percents) are not useful if configuration is not given
  - Commercial Off-the-Shelf (COTS) software
  - Virtual environments

The amount of available hardware resources is usually a variable at the beginning of the design process. The main groups of resources are CPU, I/O, memory, and network.

When resource requirements are measured as resource utilization, it is related to a particular hardware configuration. It is a good metric when the hardware the system will run on is known. Often such requirements are a part of a generic policy; for example, that CPU utilization should be below 70 percent. Such requirements won't be very useful if the system deploys on different hardware configurations, and especially for Commercial Off-the-Shelf (COTS) software.

In [Performance07] authors separate performance requirements / goals (which they apply mainly to response times) from performance targets / thresholds (which they apply mainly to resource utilization).

Resources: Absolute Values

- Absolute values
  - # of instructions, I/O per transaction
  - Seen mainly in modeling
  - MIPS in mainframe world
- Importance increases again with the trends of virtualization, cloud computing, and service-oriented architectures
  - VMware: CPU usage in MHz

When specified in absolute values, like the number of instructions to execute or the number of I/O per transaction (as sometimes used, for example, for modeling), it may be considered as a performance metric of the software itself, without binding it to a particular hardware configuration. In the mainframe world, MIPS was often used as a metric for CPU consumption, but I am not aware about such a widely used metric in the distributed systems world.

The importance of resource-related requirements will increase again with the trends of virtualization, cloud computing, and service-oriented architectures. When you depart from the "server(s) per application" model, it becomes difficult to specify requirements as resource utilization, as each application will add only incrementally to resource utilization for each service used. And indeed, looks like in the VMware world 'CPU usage in MHz' or usagemhz is used.

Scalability

- **Ability of the system to meet performance requirements as the demand increases**

- **Increasing # of users, transaction volumes, data sizes, new workloads, etc.**

- **Performance requirements as a function, for example, of load or data and configuration**
  - **No free / ideal scalability**

Scalability is a system's ability to meet the performance requirements as the demand increases (usually by adding hardware). Scalability requirements may include demand projections such as an increasing number of users, transaction volumes, data sizes, or adding new workloads.

From a performance requirements perspective, scalability means that you should specify performance requirements not only for one configuration point, but as a function, for example, of load or data. For example, the requirement may be to support throughput increase from five to 10 transactions per second over next two years with response time degradation not more than 10 percent.

Most scalability requirements I have seen look like "to support throughput increase from five to 10 transactions per second over next two years without response time degradation". Actually, except in a few very special cases, increasing load leads to some response time degradation, so supporting such requirement without any changes in the system is possible only with addition of hardware resources.

It is very difficult to consider performance (and, therefore, performance requirements) without context. It depends, for example, on hardware resources provided, the volume of data operated on, and functionality included in the system. So if any of that information is known, it should be specified in the requirements. Not everything may be specified at the same point: while the volume of data to keep is usually determined by the business and should be documented at the beginning, the hardware configuration may be determined during the design stage.

Agenda

- **Metrics**
- **Defining**
- **Verification**

Defining Requirements

- **Requirements / Architect's vocabulary**
- **Quality Attributes**
  - Part of Nonfunctional Requirements
- **Approaches**
  - Text
  - Quality Attribute Scenarios (SEI)
  - Planguage

Requirement Engineering / Architect's vocabulary is very different from what is used in performance testing / capacity planning. Performance and scalability are examples of Quality Attributes (QA), part of Nonfunctional Requirements (NFR).

In addition to specifying requirements in plain text, there are multiple approaches to formalize documenting of requirements. For example, Quality Attribute Scenarios by The Carnegie Mellon Software Engineering Institute (SEI) or Planguage (Planning Language) introduced by Tom Gilb.

Quality Attribute Scenarios

- QA scenario defines:
  - Source
  - Stimulus
  - Environment
  - Artifact
  - Response
  - Response Measure

QA scenario defines source, stimulus, environment, artifact, response, and response measure [Bass03]. For example, the scenario may be "Users initiate 1,000 transactions per minute stochastically under normal operations, and these transactions are processed with an average latency of two seconds." For this example:

•Source is a collection of users.

•Stimulus is the stochastic initiation of 1,000 transactions per minute.

•Artifact is always the system's services.

•Environment is the system state, normal mode in our example.

•Response is processing the transactions.

•Response measure is the time it takes to process the arriving events (an average latency of two seconds in our example).

Planguage

- Tag: unique identifier
- Gist: brief description
- Scale: unit of measure
- Meter: how to measure
- Minimum / Plan / Stretch/ Wish : levels to attain
- Past / Record / Trend

Planguage (Planning language) was suggested by Tom Gilb and may work better for quantifying quality requirements [Simmons01]. Planguage keywords include:

•Tag: a unique identifier

•Gist: a short description

•Stakeholder: a party materially affected by the requirement

•Scale: the scale of measure used to quantify the statement

•Meter: the process or device used to establish location on a Scale

•Must: the minimum level required to avoid failure

•Plan: the level at which good success can be claimed

•Stretch: a stretch goal if everything goes perfectly

•Wish: a desirable level of achievement that may not be attainable through available means

•Past: an expression of previous results for comparison

•Trend: an historical range or extrapolation of data

•Record: the best-known achievement

It is very interesting that Planguage defines four levels for each requirement: minimum, plan, stretch, and wish.

Goals vs. Requirements

- **Most response times "requirements" are goals**
    - Missing them won't prevent deploying the system
- **For response times, the difference between goals and requirements may be large**
    - For many web applications goals are two-five seconds and requirements somewhere between eight seconds and one minute

One often missed issue, as Scott Barber notes, is goals versus requirements [Barber07]. Most of response time "requirements" (and sometimes other kinds of performance requirements,) are goals, not requirements: something that we want to achieve, but missing them won't necessarily prevent deploying the system.

In many cases, especially for response times, there is a big difference between goals and requirements (the point when stakeholders agree that the system can't go into production with such performance). For many interactive web applications, response time goals are two to five seconds and requirements may be somewhere between eight seconds and one minute.

One approach may be to define both goals and requirements. The problem is that, except when coming from legal or contractual obligation, requirements are very difficult to get. Even if stakeholders define performance requirements, quite often, when it comes to go/no go decisions, it becomes clear that it was not the real requirements, but rather second-tier goals.

See The Whole Picture

- For example, the requirement is 10 seconds

- We got 15 seconds for peak load

- But what if
    - Only on busiest day of the year
    - All other days it will be below 10 seconds
    - It is CPU-constrained and may be fixed by additional hardware

In addition, multiple performance metrics only together provide the full picture. For example, you may have a 10-second requirement and you get 15-second response time under the full load. But what if you know that this full load is the high load on the busiest day of year, that response times for the maximal load for other days are below 10 second, and you see that it is CPU-constrained and may be fixed by a hardware upgrade? Real response time requirements are so environment- and business-dependent that for many applications it may be problematic to force people to make hard decisions in advance for each possible combination of circumstances. One approach may be to specify goals (making sure that they make sense) and only then, if they are not met, make the decision what to do with all the information available.

What Metrics to Use?

- Average
- Max
- Percentiles (X% below Y sec)
- Median
- Typical
- etc.

Another question is how to specify response time requirements or goals. Individual transaction response times vary, so aggregate values should be used. For example, such metrics as average, maximum, different kinds of percentiles, or median. The problem is that whatever aggregate value you use, you lose some information.

The Issue

- SLA (Service Level Agreement)
  - "99.5% of all transactions should have a response time less than five seconds"
- What happens with the rest 0.5%?
  - All 6-7 seconds
  - All failed/timeout
- Add different types of transactions, different input data, different user locations, etc.

Percentiles are more typical in SLAs (Service Level Agreements). For example, "99.5 percent of all transactions should have a response time less than five seconds". While that may be sufficient for most systems, it doesn't answer all questions. What happens with the remaining 0.5 percent? Do these 0.5 percent of transactions finish in six to seven seconds or do all of them timeout? You may need to specify a combination of requirements: for example, 80 percent below four seconds, 99.5 percent below six seconds, and 99.9 percent below 15 seconds (especially if we know that the difference in performance is defined by distribution of underlying data). Other examples may be average four seconds and maximal 12 seconds, or average four seconds and 99 percent below 10 seconds.

As Adrian Cockcroft noted about observability [Cockcroft00], in addition to losing information when you aggregate, there are different viewpoints for performance data that need to be provided for different audiences. You need different metrics for management, engineering, operations, and quality assurance. For operations and management percentiles may work best. If you do performance tuning and want to compare two different runs, average may be a better metric to see the trend. For design and development you may need to provide more detailed metrics; for example, if the order processing time depends on the number of items in the order, it may be separate response time metrics for one to two, three to 10, 10 to 50, and more than 50 items.

Moreover, the tools providing performance information for different audiences are usually different, present information in a different way, and may measure different things. For example, load testing tools and active monitoring tools provide metrics for the used synthetic workload that may differ significantly from the actual production load. This becomes a real issue if you want to implement some kind of process, such as Six Sigma, to keep performance under control throughout the whole system lifecycle.

Metrics to Use

- **Combination of percentile and availability metric works in many cases**
  - 97% below 5 sec, less than 1% failed/timeout
- **An example of another approach:**
  - Apdex (Application Performance Index)
  - Objective user satisfaction metric
  - A number between 0 and 1
  - 0 no users satisfied, 1 all users satisfied

Things get more complicated when there are many different types of transactions, but a combination of percentile-based performance and availability metrics usually works in production for most interactive systems. While more sophisticated metrics may be necessary for some systems, in most cases they make the process overcomplicated and results difficult to analyze.

There are efforts to make an objective user satisfaction metric. For example, Apdex - Application Performance Index. Apdex is a single metric of user satisfaction with the performance of enterprise applications. The Apdex metric is a number between 0 and 1, where 0 means that no users were satisfied, and 1 means all users were satisfied. The approach introduces three groups of users: satisfied, tolerating, and frustrated. Two major parameters are introduced: threshold response times between satisfied and tolerating users T, and between tolerating and frustrated users F [Apdex, Sevcik08]. There probably is a relationship between T and the response time goal, and between F and the response time requirement. However, while Apdex is a good metric for management and operations, it may be too high-level for engineering.

Where do performance requirements come from?

- Business

- Usability

- Technology

If we look at performance requirements from another point of view, we can classify them into business, usability, and technological requirements.

Business requirements come directly from the business and may be captured very early in the project lifecycle, before design starts. For example, "a customer representative should enter 20 requests per hour and the system should support up to 1000 customer representatives". If translated into more technical terms, requests should be processed in five minutes on average, throughput would be up to 20,000 requests per hour, and there could be up to 1,000 parallel sessions.

The main trap here is to immediately link business requirements to a specific design, technology, or usability requirements, thus limiting the number of available design choices. If we consider a Web system, for example, it is probably possible to squeeze all the information into a single page or have a sequence of two dozen screens. All information can be saved at once or each page of these two-dozen can be saved separately. We have the same business requirements, but response times per page and the number of pages per hour would be different.

Requirements Elicitation

- *Final* requirements should be quantitative and measurable

- Business people know what the system should do and may provide some information
    - They are not performance experts

- Document real business requirements in the form they are available
    - Then elaborate them into quantitative and measurable

While the final requirements should be quantitative and measurable, it is not an absolute requirement for initial requirements. Scott Barber, for example, advocates that we need to gather qualitative requirements first [Barber07]. While business people know what the system should do and may provide some numeric information, they are not trained in requirement elaboration and system design. If asked to provide quantitative and measurable requirements, they finally provide them based on whatever assumptions they have about system's design and human-computer interaction, but quite often this results in wrong assumptions being documented as business requirements. We need to document real business requirements in the form they are available, and only then elaborate them into quantitative and measurable.

Usability Requirements

- **Many researchers agree that response times should be 2-5 seconds for maximum performance and users lose focus if response times are more than 8-10 seconds**

- **Sometimes linked closely to business requirements**

  - **Make sure that response times are not worse than competitor's**

Usability requirements, mainly related to response times, are based on the basic principles of human-computer interaction. Many researchers agree that users lose focus if response times are more than eight to 10 seconds and that response times should be two to five seconds for maximum productivity. These usability considerations may influence design choices (such as using several Web pages instead of one). In some cases, usability requirements are linked closely to business requirements; for example, make sure that your system's response times are not worse than response times of similar or competitor systems.

Response Times: Review of Research

- In 1968 Robert Miller defined three threshold levels of human attention
- Instantaneous 0.1-0.2 seconds
- Free interaction 1-5 seconds
- Focus on dialog 5-10 seconds

As long ago as 1968, Robert Miller's paper "Response Time in Man-Computer Conversational Transactions" described three threshold levels of human attention [Miller68]. Jakob Nielsen believes that Miller's guidelines are fundamental for human-computer interaction, so they are still valid and not likely to change with whatever technology comes next [Nielsen94]. These three thresholds are:

•Users view response time as instantaneous (0.1-0.2 second)

•Users feel they are interacting freely with the information (1-5 seconds)

•Users are focused on the dialog (5-10 seconds)

Instantaneous Response Time

- Users feel that they directly manipulate User Interface (UI)
- For example, between typing a symbol and its appearance on the screen
- 0.1-0.2 seconds
- Often beyond the reach of application developers
  - System/UI libraries, client-side

Users view response time as instantaneous (0.1-0.2 second): Users feel that they directly manipulate objects in the user interface. For example, the time from the moment the user selects a column in a table until that column highlights or the time between typing a symbol and its appearance on the screen. Robert Miller reported that threshold as 0.1 seconds. According to Peter Bickford 0.2 second forms the mental boundary between events that seem to happen together and those that appear as echoes of each other [Bickford97].

Although it is a quite important threshold, it is often beyond the reach of application developers. That kind of interaction is provided by operating system, browser, or interface libraries, and usually happens on the client side, without interaction with servers (except for dumb terminals, that is rather an exception for business systems today).

Free Interaction

- Notice delay, but "feel" the computer is "working"
- Earlier researchers reported 1-2 sec
    - Simple terminal interface
- For problem solving tasks no performance degradation up to 5 sec
- Depends on the number of elements and repetitiveness of the task

Users feel they are interacting freely with the information (1-5 seconds): They notice the delay, but feel the computer is "working" on the command. The user's flow of thought stays uninterrupted. Robert Miller reported this threshold as one-two seconds [Miller68].

Peter Sevcik identified two key factors impacting this threshold [Sevcik03]: the number of elements viewed and the repetitiveness of the task. The number of elements viewed is, for example, the number of items, fields, or paragraphs the user looks at. The amount of time the user is willing to wait appears to be a function of the perceived complexity of the request. The complexity of the user interface and the number of elements on the screen both impact thresholds. Back in 1960s through 1980s the terminal interface was rather simple and a typical task was data entry, often one element at a time. So most earlier researchers reported that one to two seconds was the threshold to keep maximal productivity. Modern complex user interfaces with many elements may have higher response times without adversely impacting user productivity. Users also interact with applications at a certain pace depending on how repetitive each task is. Some are highly repetitive; others require the user to think and make choices before proceeding to the next screen. The more repetitive the task, the better expected response time.

That is the threshold that gives us response time usability goals for most user-interactive applications. Response times above this threshold degrade productivity. Exact numbers depend on many difficult-to-formalize factors, such as the number and types of elements viewed or repetitiveness of the task, but a goal of three to five seconds is reasonable for most typical business applications.

Users are focused on the dialog (5-10 seconds): They keep their attention on the task. Robert Miller reported that threshold as 10 seconds [Miller68]. Users will probably need to reorient themselves when they return to the task after a delay above this threshold, so productivity suffers.

Peter Bickford investigated user reactions when, after 27 almost instantaneous responses, there was a two-minute wait loop for the 28[th] time for the same operation. It took only 8.5 seconds for half the subjects to either walk out or hit the reboot [Bickford97]. Switching to a watch cursor during the wait delayed the subject's departure for about 20 seconds. An animated watch cursor was good for more than a minute, and a progress bar kept users waiting until the end. Bickford's results were widely used for setting response times requirements for Web applications.

In another research [Bouch00] users were presented with Web pages that had predetermined delays from two to 73 seconds. While performing the task, users rated the latency for each page they accessed as high, average or poor. The team reported the following ratings: good - up to 5 seconds, average - from 6 to 10 seconds, poor - more than 10 seconds.

That is the threshold that gives us response time usability requirements for most user-interactive applications. Response times above this threshold cause users to lose focus and lead to frustration. Exact numbers vary significantly depending on the interface used, but it looks like response times should not be more than eight to 10 seconds in most cases. Still, the threshold shouldn't be applied blindly; in many cases, significantly higher response times may be acceptable when appropriate user interface is implemented to alleviate the problem.

Technological Requirements

- Comes from the chosen design and used technology
  - We call ten web services sequentially to show a page within 3 sec. It translates into requirements of 200-250 ms for each web service
  - resource utilization requirements

The third category, technological requirements, comes from chosen design and used technology. Some of technological requirements may be known from the beginning if some design elements are given, but others are derived from business and usability requirements throughout the design process and depend on the chosen design.

For example, if we need to call ten Web services sequentially to show the Web page with a three-second response time, the sum of response times of each Web service, the time to create the Web page, transfer it through the network and render it in a browser should be below 3 second. That may be translated into response time requirements of 200-250 milliseconds for each Web service. The more we know, the more accurately we can apportion overall response time to Web services.

Another example of technological requirements is resource consumption requirements. For example, CPU and memory utilization should be below 70% for the chosen hardware configuration.

Business requirements should be elaborated during design and development, and merge together with usability and technological requirements into the final performance requirements, which can be verified during testing and monitored in production. The main reason why we separate these categories is to understand where the requirement comes from: is it a fundamental business requirement and the system fails if we miss it or a result of a design decision that may be changed if necessary.

Performance requirements are important input for Software Performance Engineering [Smith02]. During design and development the requirements are further elaborated. For example, the service / stored procedure response time requirements should be determined by its share in the end-to-end performance budget. In this way, the worst possible combination of all required services, middleware and presentation layer overheads will provide the required time. For example, if there is a Web page with 10 drop-down boxes calling 10 separate services, the response time objective for each service may be 0.2 seconds to get three seconds average response time (leaving one second for network, presentation, and rendering).

Determining what specific performance requirements is another large topic that is difficult to formalize. Consider the approach suggested by Peter Sevcik for finding T, the threshold between satisfied and tolerating users. T is the main parameter of the Apdex (Application Performance Index) methodology, providing a single metric of user satisfaction with the performance of enterprise applications. Peter Sevcik defined ten different methods [Sevcik08].

Methods 1-5 to Find T
(by Peter Sevcik)

- Default value (4 sec)

- Empirical data

- User behavior model (# of elements/task repetitiveness)

- Outside references

- Observing users

•Default value (the Apdex methodology suggest 4 sec)

•Empirical data

•User behavior model (number of elements viewed / task repetitiveness)

•Outside references

•Observing the user

Methods 6-10 to Find T (by Peter Sevcik)

- Controlled performance experiment
- Best time multiple
- Find frustration threshold F first and calculate T from F (F=4T in APDEX)
- Interview stakeholders
- Mathematical inflection point

•Controlled performance experiment

•Best time multiple

•Find frustration threshold F first and calculate T from F (the Apdex methodology assumes that F = 4T)

•Interview stakeholders

•Mathematical inflection point.

Each method is discussed in details in [Sevcik08].

Suggested Approach

- So Peter Sevcik suggests to use several of these methods: if all come approximately to the same number it will be T

- A similar approach can be used for performance requirements: use several methods to get the numbers – you get goal/requirement if they are close

  - Investigate / sort out if they differ significantly

The idea is use of several (say, three) of these methods for the same system. If all come to approximately the same number, they give us T. While the approach was developed for production monitoring, there is definitely a strong correlation between T and the response time goal (having all users satisfied sounds a pretty good goal), and between F and the response time requirement. So the approach probably can be used for getting response time requirements with minimal modifications. While some specific assumptions like four seconds for default or the F=4T relationship may be up for argument, the approach itself conveys the important message that there are many ways to determine a specific performance requirement and it would be better to get it from several sources for validation purposes. Depending on your system, you can determine which methods from the above list (or maybe some others) are applicable, calculate the metrics and determine your requirements.

Agenda

- **Metrics**
- **Defining**
- **Verification**

Requirement Verification

- Let's consider the following example
- Response time requirement is 99% below 5 sec
- 99% 3-5 sec, 1% 5-8 sec
  - Looks like a minor performance issue
- 99% 3-5 sec, 1% failed or had strangely high response times (more than 30 sec)
  - Looks like a bug or serious performance issue

Requirement verification presents another subtle issue: how to differentiate performance issues from functional bugs exposed under load. Often, additional investigation is required before you can determine the cause of your observed results. Small anomalies from expected behavior are often signs of bigger problems, and you should at least figure out *why* you get them.

When 99 percent of your response times are three to five seconds (with the requirement of five seconds) and 1 percent of your response times are five to eight seconds it usually is not a problem. But it probably should be investigated if this 1 percent fail or have strangely high response times (for example, more than 30 sec) in an unrestricted, isolated test environment. This is not due to some kind of artificial requirement, but is an indication of an anomaly in system behavior or test configuration. This situation often is analyzed from a requirements point of view, but it shouldn't be, at least until the reasons for that behavior become clear.

Requirement Verification: Performance vs. Bugs

- **Two completely different cases**
  - Performance issue: business decision, cost vs. response time trade off.
  - Bug exposed under load: should be traced down first to make decision

These two situations look similar, but are completely different in nature:

1) The system is missing a requirement, but results are consistent: this is a business decision, such as a cost vs. response time trade off.

2) Results are not consistent (while requirements can even be met): that may indicate a problem, but its scale isn't clear until investigated.

Unfortunately, this view is rarely shared by development teams too eager to finish the project, move it into production, and move on to the next project. Most developers are not very excited by the prospect of debugging code for small memory leaks or hunting for a rare error that is difficult to reproduce. So the development team becomes very creative in finding "explanations". For example, growing memory and periodic long-running transactions in Java are often explained as a garbage collection issue. That is false in most cases. Even in the few cases when it is true, it makes sense to tune garbage collection and prove that the problem went away.

Another typical situation is getting some transactions failed during performance testing. It may still satisfy performance requirements, which, for example, state that 99% of transactions should be below X seconds – and the share of failed transaction is less than 1 percent. While this requirement definitely makes sense in production where we may have network and hardware failures, it is not clear why we get failed transactions during the performance test if it was run in a controlled environment and no system failures were observed. It may be a bug exposed under load or a functional problem for some combination of data.

The equipment is not operating as expected, and therefore there is a danger that it can operate with even wider deviation in this unexpected and not thoroughly understood way. The fact that this danger did not lead to a catastrophe before is no guarantee that it will not the next time, unless it is completely understood.

Dr. Richard Feynman
Roger Commission Report on the
Challenger space shuttle accident

When some transactions fail under load or have very long response times in the controlled environment and we don't know why, we have one or more problems. When we have unknown problems, why not trace it down and fix in the controlled environment? It would be much more difficult in production. What if these few failed transactions are a view page for your largest customer, and you won't be able to create any order for this customer until the problem is fixed? In functional testing, as soon as you find a problem, you usually can figure out how serious it is. This is not the case for performance testing: usually you have no idea what caused the observed symptoms and how serious it is, and quite often the original explanations turn out to be wrong.

Michael Bolton described this situation concisely [Bolton06]:

As Richard Feynman said in his appendix to the Rogers Commission Report on the Challenger space shuttle accident, when something is not what the design expected, it's a warning that something is wrong. "The equipment is not operating as expected, and therefore there is a danger that it can operate with even wider deviations in this unexpected and not thoroughly understood way. The fact that this danger did not lead to a catastrophe before is no guarantee that it will not the next time, unless it is completely understood." When a system is in an unpredicted state, it's also in an unpredictable state.

We need to specify performance requirements at the beginning of any project for design and development (and, of course, reuse them during performance testing and production monitoring). While performance requirements are often not perfect, forcing stakeholders just to think about performance increases the chances of project success.

What exactly should be specified – goal vs. requirements (or both), average vs. percentile vs. APDEX, etc. – depends on the system and environment. Whatever it is, it should be something quantitative and measurable in the end. Making requirements too complicated may hurt. We need to find meaningful goals / requirements, not invent something just to satisfy a bureaucratic process.

If we define a performance goal as a point of reference, we can use it throughout the whole development cycle and testing process and track our progress from the performance engineering point of view. Tracing this metric in production will give us valuable feedback that can be used for future system releases.

References

[Apdex] Apdex Web site
http://www.apdex.org/

[Barber07] Barber, S. Get performance requirements right - think like a user, Compuware white paper, 2007.
http://www.perftestplus.com/resources/requirements_with_compuware.pdf

[Bass03] Bass L., Clements P., Kazman R. Software Architecture in Practice, Addison-Wesley, 2003.
http://etutorials.org/Programming/Software+architecture+in+practice,+second+edition

[Bickford97] Bickford P. Worth the Wait? Human Interface Online, View Source, 10/1997.
http://web.archive.org/web/20040913083444/http://developer.netscape.com/viewsource/bickford_wait.htm

[Bolton06] Bolton M. More Stress, Less Distress, Better Software, November 2006.
http://www.stickyminds.com/sitewide.asp?ObjectId=11536&Function=edetail&ObjectType=ART

[Bouch00] Bouch, A., Kuchinsky, A. and Bhatti, N. Quality is in the eye of the beholder: Meeting users' requirements for Internet quality of service, CHI 2000, 297-304.
http://research.hp.com/personal/Nina_Bhatti/papers/chi.pdf

[Cockcroft00] Cockcroft, A., Walker B. Capacity Planning for Internet Services. Quick planning techniques for high growth rates. Sun Microsystems, 2000.
http://java.coe.psu.ac.th/SunDocuments/SunBluePrints/caphi.pdf

[Miler68] Miller, R. B. Response time in user-system conversational transactions, In Proceedings of the AFIPS Fall Joint Computer Conference, 33, 1968, 267-277.

[Loosley05] Loosley C. When Is Your WebSite Fast Enough? E-CommerceTimes, 2005.
http://www.ecommercetimes.com/story/46627.html

[Martin86] Martin, G.L. and Corl, K.G. System response time effects on user productivity, Behavior and Information Technology, 5(1), 1986, 3-13.

[Nielsen94] Nielsen J. Response Times: The Three Important Limits, Excerpt from Chapter 5 of Usability Engineering, 1994.
http://www.useit.com/papers/responsetime.html

[Performance07] Performance Testing Guidance for Web Applications, 2007.
http://perftestingguide.codeplex.com/releases/view/6690

[Podelko07] Podelko A. Multiple Dimensions of Performance Requirements, CMG, 2007.
http://www.alexanderpodelko.com/docs/Performance_Requirements_CMG07.pdf

[Sevcik03] Sevcik, P. How Fast Is Fast Enough, Business Communications Review, March 2003, 8–9.
http://www.bcr.com/architecture/network_forecasts%10sevcik/how_fast_is_fast_enough?_20030315225htm

[Sevcik08] Sevcik, P. Using Apdex to Manage Performance, CMG, 2008.
http://www.apdex.org/documents/Session18.0Sevcik.pdf

[Simmons01] Simmons E. Quantifying Quality Requirements Using Planguage, Quality Week, 2001.
http://www.clearspecs.com/downloads/ClearSpecs20V01_Quantifying%20Quality%20Requirements.pdf

[Smith02] Smith C., Williams L. "Performance Solutions", Addison-Wesley, 2002.