Adjusting Performance Testing to the World of Agile

It looks like agile methodologies are somewhat struggling with performance testing. Theoretically it should be a piece of cake: every iteration you have a working system and know exactly where you stand with the system's performance. Unfortunately, it doesn't always work this way in practice. Performance related activities tend to slip toward the end of the project. Another issue is that agile methods are oriented toward breaking projects into small tasks, which is quite difficult to do with performance: performance-related activities usually span the whole project.

From one side we have a challenge to add performance testing into Continuous Integration making it a part of DevOps processes. From another side, approaching performance testing formally, with rigid, step-by-step approach and narrow specialization often leads to missing performance problems or to prolonged agony of performance troubleshooting. With small extra efforts, making the process more exploratory, efficiency of performance testing increases significantly and these extra efforts usually pay off multi-fold even before the end of performance testing.

This paper discusses how performance testing may be adjusted to the modern industry trends and become an integral part of agile processes. While there are too many open questions there, it looks like such adjusting is not an option anymore.

After a short introduction into the subject, two major aspects would be discussed: performance testing in agile projects and doing performance testing the agile way.

Performance Testing

- Here used as an umbrella term including load, stress, concurrency, etc. testing
- A project itself
  - We can apply software development methodologies to this project
- Waterfall approach in most cases
  - Often pre-production performance validation

Madrid, 10th & 11th June 2015

Most large corporations have performance testing / engineering groups today, performance testing becomes a must step to get the system into production. Still, in most cases, it is pre-production performance validation only.

Even if it is only a short pre-production performance validation, performance testing is a project itself with multiple phases of requirement gathering, test design, test implementation, test execution, and result analysis. So in most cases software development methodologies could, with some adjustments, be applicable to performance testing.

Waterfall Approach to Performance Testing

- Flowing steadily downwards through the phases
  - Get the system ready
  - Develop scripts requested
  - Run scripts in the requested combinations
  - Compare with the provided requirements
  - If requirements missed, throw back to development

Madrid, 10th & 11th June 2015

The waterfall approach in software development is a sequential process in which development is seen as flowing steadily downwards (like a waterfall) through the phases of requirements analysis, design, implementation, testing, integration, and maintenance [Waterfall]. Being a step on the project plan, performance testing is usually scheduled in the waterfall way, when you need to finish one step to start next. Typical steps could be, for example:

-Get the system ready
-Develop scripts requested (sometimes offshore)
-Run scripts in the requested combinations
-Compare with the requirements provided
-Allow some percentage of errors according to the requirements
-Involve the development team if requirements are missed

Manifesto for Agile Software Development

We are uncovering better ways of developing software by doing it and helping others do it.

Through this work we have come to value:
- *Individuals and interactions over processes and tools*
- *Working software over comprehensive documentation*
- *Customer collaboration over contract negotiation*
- *Responding to change over following a plan*

That is, while there is value in the items on the right, we value the items on the left more.

Madrid, 10th & 11th June 2015

The word agile here doesn't refer to any specific development process or methodology. It rather refers to the original definition of this word as stated in "Manifesto for Agile Software Development" [Manifesto01]:

*We are uncovering better ways of developing software by doing it and helping others do it.*

Through this work we have come to value:

Individuals and interactions over processes and tools
Working software over comprehensive documentation
Customer collaboration over contract negotiation
Responding to change over following a plan

That is, while there is value in the items on the right, we value the items on the left more.
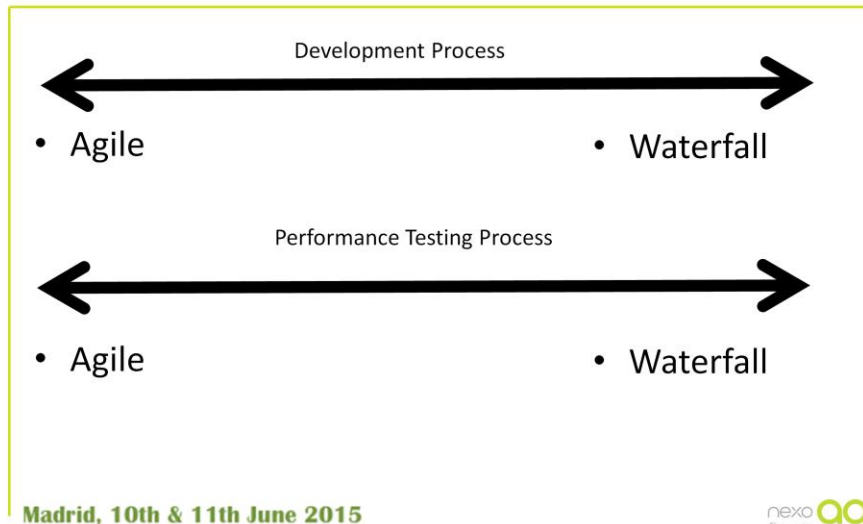
We have two different dimensions here. The first is the process used for the whole project. The second is the approach used to performance testing as a project (independently from the process used for the whole project).

First we discuss performance testing in software development projects using agile methodologies.

Agile Development

- Agile development should be rather a trivial case for performance testing
  - You have a working system each iteration to test early by definition.
  - You need a perf engineer for the whole project
    - Savings come from detecting problems early
- You need to adjust requirements for implemented functionality
  - Additional functionality will impact performance

Madrid, 10th & 11th June 2015

Practical agile development is struggling with performance in general. Theoretically it should be a piece of cake: every iteration you have a working system and know exactly where you stand with the system's performance. You shouldn't wait until the end of the waterfall process to figure out where you are – on every iteration you can track your performance against requirements and see the progress (making adjustments on what is already implemented and what is not yet). Clearly it is supposed to make the whole performance engineering process much more straightforward and solve the first problem of traditional approach that the system should be ready for performance testing (so it usually happened very late in the process).

The Main Issue on the Agile Side

- It doesn't [always] work this way in practice
- That is why you have "Hardening Iterations", "Technical Debt" and similar notions
- Same old problem: functionality gets priority over performance

Unfortunately, it looks like it doesn't always work this way in practice. So such notions as "hardening iterations" and "technical debt" get introduced. Although it is probably the same old problem: functionality gets priority over performance (which is somewhat explainable: you first need some functionality before you can talk about its performance). So performance related activities slip toward the end of the project and the chance to implement a proper performance engineering process built around performance requirements is missed.

The Main Issue on
the Testing Side

expo:QA '15

- Performance Engineering teams don't scale well
  - Even assuming that they are competent and effective
- Increased volume exposes the problem
  - Each iteration
- Remedies: automation, making performance everyone's job

Madrid, 10th & 11th June 2015

nexo QA
Events

The fundamental issue is, as I see it, that performance engineering teams don't scale well, even assuming that they are competent and effective. At least not in their traditional form. They work well in traditional corporate environments where they check products for performance before release, but they face challenges as soon as we start to expand the scope of performance engineering (early involvement, more products/configurations/scenarios, etc.). And agile projects, where we need to test the product each iteration or build, expose the problem through an increased volume of work to do.

Just to avoid misunderstandings, I am a strong supporter of having performance teams and I believe that it is the best approach to building performance culture. Performance is a special area and performance specialists should have an opportunity to work together to grow professionally. The details of organizational structure may vary, but a center of performance expertise should exist. Only thing I am saying here is that while the approach works fine in traditional environments, it needs major changes in organization, tools, and skills when the scope of performance engineering should be extended (as in the case of agile projects).

Actually remedies are well known: automation, making performance everyone jobs (full immersion), etc. [Crispin09, Barber11] However they are not wide-spread yet.

Automation (continuous performance testing) means here not only using tools (in performance testing we almost always use tools), but automating the whole process including setting up environment, running tests, and reporting / analyzing results.

Historically performance testing automation was almost non-existent (at least in traditional environments). Performance testing automation is much more difficult than, for example, functional testing automation. Setups are much more complicated. A list of possible issues is long. Results are complex (not just pass/fail). It is not easy to compare two result sets. So it is definitely much more difficult and would probably require much more human intervention, but it isn't impossible.

And, of course, changing interfaces is a major challenge. Especially when recording is used to create scripts and it is difficult to predict if product changes break scripts.

Automation: Considerations

- You need know system well enough to make meaningful automation
- If system is new, overheads are too high
  - So almost no automation in traditional environments
- If the same system is tested again and again
  - It makes sense to invest in setting up automation
- Automated interfaces should be stable enough
  - APIs are usually more stable on early stages

Madrid, 10th & 11th June 2015

However, the cost of performance testing automation is high. You need to know system well enough to make meaningful automation. Automation for a new system doesn't make much sense - overheads are too high. So there was almost no automation in traditional environment [with testing in the end with a record/playback tool]. When you test the system once in a while before a next major release, chances to re-use your artifacts are low.
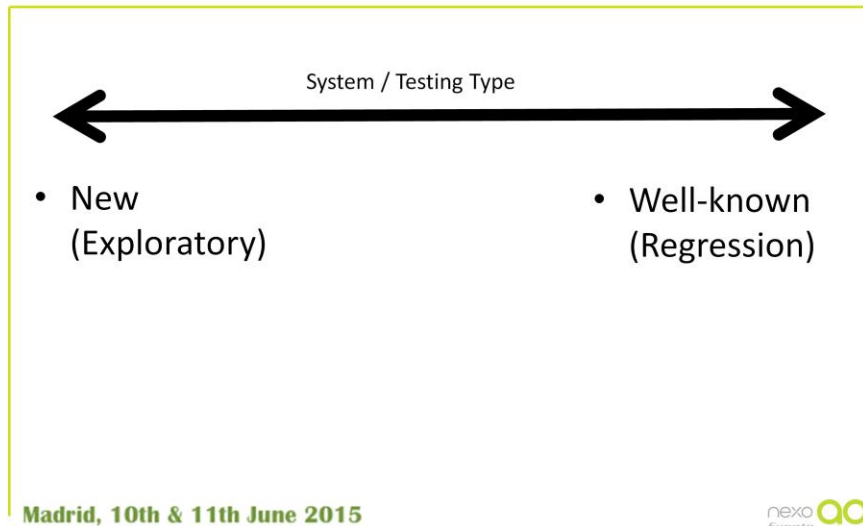
It is opposite when the same system is tested again and again (as it should be in agile projects). It makes sense to invest in setting up automation. It rarely happened in traditional environments – even if you test each release, they are far apart and the difference between the releases prevents re-using the artifacts (especially with recorded scripts – APIs, for example, is usually more stable). So demand for automation was rather low and tool vendors didn't pay much attention to it. Well, the situation is changing – we may see more automation-related features in load testing tools soon.

I am not saying that automation would replace performance testing as we know it. Performance testing of new systems is agile and exploratory in itself and can't be replaced by automation (well, at least not in the foreseen future). Automation would complement it – together with additional input from development offloading performance engineers from routine tasks not requiring sophisticated research and analysis.

Actually it is similar to functional testing where you need to find a trade-off between automated regression testing and exploratory testing – with the difference that you use tools in performance testing in any case and setting up continuous performance

One More Dimensions to Consider

System / Testing Type

- New (Exploratory)
- Well-known (Regression)

Madrid, 10th & 11th June 2015

So we have one more dimension to consider: what kind of the system we have and what kind of performance testing we need. For a new system, when we don't know much about it, we need an agile / exploratory kind of testing (one definition of exploratory testing is: *Exploratory testing is simultaneous learning, test design, and test execution* [Bach03] - that is exactly what we need for a new system). For a well-known system, when changes are rather small, we don't need much exploration and automation may be very effective.

Until recently, while there were some vendors claiming that their load testing tools better fit agile processes, it usually meant that the tool is a little easier to handle (and, unfortunately, often just because there is not much functionality in it). While ease of use is an important factor, it is not specific to agile and should be accompanied by all necessary functionality.

What makes agile projects really different is need of running large number of tests repeatedly – resulting in need for tools to support performance testing automation. Unfortunately, even if a tool has something that may be used for automation, like starting by a command line with parameters, it is difficult to find out.

So far I read about few full-scale implementation of continuous performance testing (for example, [Harding11] or [OpTier12]). Probably we don't see it more often because we don't have much infrastructure for that kind of automation and performance testers may be not the best people to create complex integrated solutions from dozens of not-related pieces (as those who implemented it did). When we get more automation support in tools, we will probably see it more often.

Tool Support: Recent Developments

- Recently agile support became the main theme
  - A lot of new developments
- Integration with Continuous Integration Servers
  - Several tools announced integration recently
- Cloud integration
- Support of newest technologies

Madrid, 10th & 11th June 2015

Recently agile support became the main theme in load testing tools, a lot of new developments. Several tools recently announced integration with Continuous Integration Servers (such as Jenkins or Hudson).

Cloud integration, including automated deployment to public clouds (almost all major load testing tools) and private cloud automation (Oracle Testing as a Service - TaaS).

Support of newest technologies (such as WebSocket or SPDY by Neoload or ADF in Oracle Load Testing).

Performance Challenges in Agile Projects

- Performance-related activities usually span the whole project
- Specifying performance requirements
  - Constraints
    - As user stories should represent finite manageable tasks
  - User stories
    - Separating of initial and ongoing compliance

Madrid, 10th & 11th June 2015

Another issue here is that agile methods are oriented toward breaking projects into small tasks, which is quite difficult to do with performance (and many other non-functional requirements) – performance-related activities usually span the whole project.

There is no standard approach to specifying performance requirements in agile methods. Mostly it is suggested to present them as user stories or as constraints. The difference between user stories and constraints approaches is not in performance requirements per se, but how to address them during the development process. The point of the constraint approach is that user stories should represent finite manageable tasks, while performance-related activities can't be handled as such because they usually span multiple components and iterations. Those who suggest to use user stories address that concern in another way – for example, separating cost of initial compliance and cost of ongoing compliance [Hazrati11].

Now we will discuss another agile aspect of performance testing : approaching performance testing projects in the agile way (independently from what approach is used in the overall software development project).

Agile Performance Testing

- Finding new opportunities inside existing processes using collaboration, iterative and adaptive processes
- Separate from performance testing in agile development, which was discussed earlier
- Most good performance engineers already use similar approaches
  - Although waterfall-like plan is usually presented to management

Madrid, 10th & 11th June 2015

Even if the project scope is limited to pre-production performance testing, approaching testing with an agile, iterative approach you meet your goals faster and more efficiently and, of course, learn more about the system along the way. After we prepare a script for testing (or however the workload is generated), we can run one, a few, and many users (how many depends on the system), analyze results (including resources utilization), and try to sort out any errors. The source of errors can be quite different – script error, functional error, or a direct consequence of a performance bottleneck. It doesn't make much sense to add load until you figure out what is going on. Even with a single script you can find many problems and, at least partially, tune the system. Running scripts separately also allows you to see how much resources are used by each type of load and make some kind of system's "model".

Using the "waterfall" approach doesn't change the nature of performance testing; it just means that you probably do a lot of extra work and still come back to the same point, performance tuning and troubleshooting, much later in the cycle. Not to mention that large tests using multiple use cases are usually a bad point to start performance tuning and troubleshooting - symptoms you see may be a cumulative effect of multiple issues.

Using an agile, iterative approach doesn't mean that you need to re-define the software development process, but rather find new opportunities inside existing processes. I believe that most good performance engineers are already doing performance testing in an agile way but just presenting it as "waterfall" to management (some kind of guerilla tactic). In most cases you need to present a waterfall-like plan to management, and then you are free to do whatever is necessary to properly test the system inside the scheduled timeframe and scope. If

Testing Early

- The main problem in waterfall development
- Nobody argues against, but rarely happens in practice
  - In agile supposed to happened automatically
- Require different, more agile approach
- Any early involvement would be beneficial
  - Even asking a few key questions can save a lot of time and resources later

Madrid, 10th & 11th June 2015

I have never read or heard anybody argue against testing early. Nevertheless, it still rarely happens in practice. Usually some project-specific reasons are stated like tight schedule or budget preventing such activities (if somebody thought about them at all).

The Software Performance Engineering (SPE) approach to the development of software systems to meet performance requirements has long been advocated by Dr. Connie Smith and Dr. Lloyd Williams (see, for example, [Smith02].). While their methodology doesn't focus on testing initiatives, it cannot be successfully implemented without some preliminary testing and data collection to determine both model inputs and parameters as well as to validate model results. Whether you are considering a full-blown SPE or guerrilla–style "back-of-the-envelope" approach, you still need to obtain baseline measurements on which to build your calculations. Early performance testing at any level of detail can be very valuable at this point.

Mentality Change

- Making performance everyone's job
- Late record/playback performance testing -> Early Performance Engineering
- System-level requirements -> Component-level requirements
- Record/playback approach -> Programming to generate load/create stubs
- "Black Box" -> "Grey Box"

Madrid, 10th & 11th June 2015

While early performance engineering is definitely the best approach (at least for product development) and has long been advocated, it is still far from commonplace. The main problem here is that mentality should be changed from a simplistic "record/playback" performance testing occurring late in the product life-cycle to a more robust true performance engineering approach starting early in the product life-cycle. You need to translate "business functions" performed by the end user into component/unit-level usage, and end-user requirements into component/unit-level requirements, etc. You need to go from the record/playback approach to utilizing programming skills to generate the workload and create stubs to isolate the component from other parts of the system. You need to go from "black box" performance testing to "grey box".

Unit Performance Testing

- Any part of the system
- Not a standard practice
- Do not wait the system is assembled
- Test cases are simpler, fewer variables
- Many systems are monolithic
  - Test-Driven Development may be an answer
- Third-party components

Madrid, 10th & 11th June 2015

One rarely discussed aspect of early performance testing is unit performance testing. The unit here maybe any part of the system like a component, service, or device. It is not a standard practice, but it should be. As we get later in the development cycle, it is more costly and difficult to make changes. Why should we wait until the whole system is assembled to start performance testing? We don't wait in functional testing, why should we in performance? The pre-deployment performance test is an analogue of system or integration tests, but usually it is conducted without any "unit testing" of performance.

The main obstacle here is that many systems are pretty monolithic; if there are parts - they don't make much sense separately. But there may be significant advantages to test-driven development. If you can decompose the system into components in such way that you may test them separately for performance, then you will only need to fix integration issues when you put the system together. Another problem is that large corporations use a lot of third-party products where the system appears as a "black box" and is not easily understood making it more difficult to test effectively.

During unit testing different variables such as load, the amount of data, security, etc. can be reviewed to determine their impact on performance. In most cases, test cases are simpler and tests are shorter in unit performance testing. There are typically fewer tests with limited scope; e.g., fewer number of variable combinations than we have in a full stress and performance test.

Another kind of early performance testing is infrastructure benchmarking; the hardware and software infrastructure is also a component of the system.

We shouldn't underestimate the power of the single-user performance test. If the performance of the system for a single user isn't good, it won't be any better for multiple users. Single-user testing is conducted throughout application development life-cycle, during functional testing and user acceptance testing; gathering performance data during these stages can be extremely helpful. The single-user performance alone may facilitate the detection of some performance issues earlier. Single-user performance can provide a good indication of what business functions and application code needs to be investigated further. Additionally, between single-user tests and load tests there are also functional multi-user tests as described in Karen Johnson's article [Johnson08]. A well-constructed test with a few users can also help to identify a lot of problems which may be very difficult to diagnose during load testing.

I believe that the title of the Andy Grove book "Only the Paranoid Survive" [Grove96] relates even better to performance engineers than it does to executives. I can imagine an executive who isn't paranoid, but I can't imagine a good performance engineer without this trait. And it is applicable to the entire testing effort from the scenarios you consider, to the scripts you create, and the results you report.

Be a Performance Test Architect

- Gather all requirements and project them onto the system architecture
  - Business requirements are not the "Holy Scripture"
  - Iterative process, evaluate and validate
  - Scrutinize workload
  - Project requirements onto the system architecture
    - What components are involved

Madrid, 10th & 11th June 2015

There are two large set of questions requiring architect-type expertise:

Gathering and validation of all requirements (first of all, workload definition) and projecting them onto the system architecture.

Too many testers consider all detailed information that they obtain from the business people (i.e., workload descriptions, scenarios, use cases, etc.) as the "holy script". But business people know the business, and they rarely know anything about performance engineering. So obtaining requirements is an iterative process and every requirement submitted should be evaluated and, if possible, validated [Podelko11]. Sometimes performance requirements are based on reliable data, sometimes they are just a pure guess, and it is important to understand how reliable they are.

The load the system should handle should be carefully scrutinized; the workload is an input to testing, while response times are output. You may decide if response times are acceptable even after the test – but you should define workload before.

The gathered requirements should be projected onto the system architecture – it is important to understand if included test cases add value testing a different set of functionality or different components of the system. From another side, it is important to make sure that we have test cases for every component (or, if we don't, we know why).

Be a Performance Test Architect

- Make sure that the system is properly configured and results are meaningful
  - Difference between environments
  - Data used
  - User access
- Make the test environment as close to the production as possible
- Performance testing isn't an exact science

Madrid, 10th & 11th June 2015

Making sure that the system under test is properly configured and the results obtained may be used (or at least projected) for the production system. Environment and setup-related considerations can have a dramatic effect. Here are a few:
•What data are used? Is it real production data, artificially generated data, or just a few random records? Does the volume of data match the volume forecasted for production? If not, what is the difference?
•How are users defined? Do you have an account set with the proper security rights for each virtual user or do you plan to re-use a single administrator id?
•What are the differences between the production and the test environment? If your test system is just a subset of your production - can you simulate the entire load or just a portion of that load?

It is important to get the test environment as close as possible to the production environment, but some differences may still remain. Even if we executed the test in the production environment with the actual production data, it would only represent one point in time, other conditions and factors would also need to be considered. In real life the workload is always random, changing each moment, including actions that nobody could even guess.

Performance testing isn't exact science. It is a way to decrease the risk, not to eliminate it completely. Results are as meaningful as the test and environment you created. Usually performance testing has small functional coverage, no emulation of unexpected events, etc. Both the environment and the data are often scaled down. All these factors confound the straightforward approach to performance testing – which states that we simply test X users simulating test cases A and B. This way we leave aside a lot of questions, for example: How many users the system can handle?

Scripting Process

- Record/playback is easy for plain HTML
  - Not many such systems anymore
- Software Development Project
  - Learn as you script
  - Correlate and parameterize
  - Validate
    - Lack of error messages is not the proof
    - Test different input data

Madrid, 10th & 11th June 2015

There are very few systems today that are stateless with static content using plain HTML, the kind of systems that lend themselves to a simplistic "record/playback" approach. In most cases there are many stumbling blocks in your way to create a proper workload. Starting from the approach you use to create the workload – the traditional "record/playback" approach just doesn't work in many cases [Podelko12]. If it is the first time you see the system there is absolutely no guarantee that you can quickly record and playback scripts to create the workload, if at all.

Creating performance testing scripts and other objects is, in essence, a software development project. Sometimes automatic script generation from recording is mistakenly interpreted as the whole process of script creation, but it is only the beginning. Automatic generation provides ready scripts in very simple cases; in most non-trivial cases it is just a first step. You need to correlate (get dynamic variables from the server) and parameterize (use different data for different users) scripts. These are operations prone to errors because we make changes directly in the communication stream. Every mistake is very dangerous because such mistakes usually can not happen in the real world where the users works with the system through a user interface or API calls.

After the script is created it should be evaluated for a single user, multiple users, and with different data. You should not assume that the system works correctly when the script was executed without errors. A very important part of load testing is workload validation. We should be sure that the applied workload is doing what it is supposed to do and that all errors are caught and logged. It can be done directly by analyzing server responses or, in cases when this is impossible, indirectly. It can be done, for example, by analyzing the application log or database for the existence of particular

My group specializes in performance testing of financial analysis products. A few of scripting challenges exist for almost every product. Nothing exceptional – you should resolve them if you have some experience and would be attentive enough, but time after time we are called to save performance testing projects ("nothing works" or "the system doesn't perform") to find out that there are serious problems with scripts and scenarios which make test results meaningless. Even very experienced testers stumble, but problems could be avoided if more time were spent analyzing what is going on. Let's consider an example – probably typical for challenges you can face with modern Web-based applications.

```
web_custom_request("XMLDataGrid.asp_7","URL={URL}/Data/XML
    DataGrid.asp?Action=EXECUTE&TaskID=1024&RowStart=1&ColSt
    art=2&RowEnd=1&ColEnd=2&SelType=0&Format=JavaScript",
    LAST);
web_custom_request("XMLDataGrid.asp_8","URL={URL}/Data/XML
    DataGrid.asp?Action=GETCONSOLSTATUS", LAST);
web_custom_request("XMLDataGrid.asp_9","URL={URL}/Data/XML
    DataGrid.asp?Action=GETCONSOLSTATUS", LAST);
web_custom_request("XMLDataGrid.asp_9","URL={URL}/Data/XML
    DataGrid.asp?Action=GETCONSOLSTATUS", LAST);
```

Some operations, like financial consolidation, can take a long time. The client starts the operation on the server and then waits until it will finish, a progress bar is shown in meanwhile. When recorded, the script looks like (in LoadRunner pseudo-code):

web_custom_request("XMLDataGrid.asp_7",
"URL={URL}/Data/XMLDataGrid.asp?Action=EXECUTE&TaskID=1024&RowStart=1&Co
lStart=2&RowEnd=1&ColEnd=2&SelType=0&Format=JavaScript", LAST);

web_custom_request("XMLDataGrid.asp_8",
"URL={URL}/Data/XMLDataGrid.asp?Action=GETCONSOLSTATUS", LAST);

web_custom_request("XMLDataGrid.asp_9",
"URL={URL}/Data/XMLDataGrid.asp?Action=GETCONSOLSTATUS", LAST);

web_custom_request("XMLDataGrid.asp_9",
"URL={URL}/Data/XMLDataGrid.asp?Action=GETCONSOLSTATUS", LAST);

What each request is doing is defined by the *?Action=* part. The number of *GETCONSOLSTATUS* requests recorded depends on the processing time. In the example above it was recorded three times; it means that the consolidation was done by the moment the third *GETCONSOLSTATUS* request was sent to the server. If playback this script, it will work in the following way: the script submits the consolidation in the *EXECUTE* request and then calls *GETCONSOLSTATUS* three times. If we have a timer around these requests, the response time will be almost instantaneous. While in reality the consolidation may take many minutes or even an hour (yes, this is a good example when sometimes people may be happy having one

```
web_custom_request("XMLDataGrid.asp_7","URL={URL}/Data/XML
    DataGrid.asp?Action=EXECUTE&TaskID=1024&RowStart=1&ColS
    tart=2&RowEnd=1&ColEnd=2&SelType=0&Format=JavaScript",
    LAST);
do {
    sleep(3000);
    web_reg_find("Text=1","SaveCount=abc_count",LAST);
    web_custom_request("XMLDataGrid.asp_8","URL={URL}/Data/X
    MLDataGrid.asp?Action=GETCONSOLSTATUS",    LAST);
    } while (strcmp(lr_eval_string("{abc_count}"),"1")==0);
```

Consolidation scripts require creating an explicit loop around GETCONSOLSTATUS to catch the end of consolidation:

```
web_custom_request("XMLDataGrid.asp_7",
"URL={URL}/Data/XMLDataGrid.asp?Action=EXECUTE&TaskID=1024&RowStart=1&Co
lStart=2&RowEnd=1&ColEnd=2&SelType=0&Format=JavaScript", LAST);

do {
            sleep(3000);
            web_reg_find("Text=1", "SaveCount=abc_count", LAST);
            web_custom_request("XMLDataGrid.asp_8",
            "URL={URL}/Data/XMLDataGrid.asp?Action=GETCONSOLSTATUS",
            LAST);
    } while (strcmp(lr_eval_string("{abc_count}"),"1")==0);
```

Here the loop simulates the internal logic of the system sending *GETCONSOLSTATUS* requests each 3 sec until the consolidation is done. Without such loop the script just checks the status and finishes the iteration while the consolidation continues for a long time after that.

Usually, when people are talking about performance testing, they do not separate it from tuning, diagnostics, or capacity planning. "Pure" performance testing is possible only in rare cases when the system and all optimal settings are well known. Some tuning activities are usually necessary at the beginning of the testing to be sure that the system is properly tuned and the results are meaningful. In most cases, if a performance problem is found, it should be diagnosed further up to the point when it is clear how to handle it. Generally speaking, 'performance testing', 'tuning', 'diagnostics', and 'capacity planning' are quite different processes and excluding any of them from the test plan (if they are assumed) will make it unrealistic from the beginning.

Both performance tuning and troubleshooting are iterative processes where you make the change, run the test, analyze the results, and repeat the process based on the findings. The advantage of performance testing is that you apply the same synthetic load, so you can accurately quantify the impact of the change that was made. That makes it much simpler to find problems during performance testing than wait until they happen in production where workload is changing all the time.

Still, even in the test environment, tuning and performance troubleshooting are quite sophisticated diagnostic processes usually requiring close collaboration between a performance engineer running tests and developers and/or system administrators making changes. In most cases it is impossible to predict how many test iterations would be necessary. Sometimes it makes sense to create a shorter and simpler test still exposing the problem under investigation. Running a complex, "real-life" test on each tuning or troubleshooting iteration can make the whole process very long as well as make the problem less evident due to different effects the problem may have on different workloads.

Building a Model

- Significantly increases the value of performance testing
  - One more way to validate tests and help to identify problems
  - Answers questions about sizing the system
- Doesn't need to be a formal model
  - May be just simple observations as to how much resources are needed by each component

Madrid, 10th & 11th June 2015

Creating a model of the system under test is very important and significantly increases the value of performance testing. First, it is one more way to validate test correctness and help to identify problems with the system – if you see deviations from the expected behavior it may mean issues with the system or issues with the way you create workload (see good examples in [Gunther06]). Second, it allows answering questions about sizing and capacity planning of the system.

It doesn't need to be a formal model created by a sophisticated modeling tool (at least for the purpose of performance testing – if any formal model is required, it is a separate activity). It may be just simple observations how much resources on each component of the system are needed for the specific workload. For example, the workload A creates significant CPU utilization on the server X while the server Y is hardly touched. This means that if increase the workload A the lack of CPU resources on the server X will create a bottleneck. As you run more and more complex tests, you verify results you get against your "model", your understanding how the system behaves – and if they don't match, you need to figure out what is wrong.

If agile methodology is implemented properly, it should be much easier and less risky to do performance testing as far as the system should be available for testing at the end of each iteration.

However, if you are involved from the beginning and testing the system each iteration as it gets delivered, you probably would need more resources comparing to pre-production performance testing. Payoff is in decreasing performance-related risks and finding problems early (which, unfortunately, is quite difficult to quantify comparing with the cost of resources you need to invest).

Automation and making performance everyone's job is the way to handle work increase and make it more manageable.

Takeaways – Agile Testing

- Small extra efforts, making the process more agile, increase efficiency significantly – and usually pay off multi-fold
  - Get involved early
  - Be paranoid about workload generation / system setup
  - Expect multiple tuning and troubleshooting iterations
  - Build a "model"

Madrid, 10th & 11th June 2015

Even if the project scope is limited to pre-production performance testing, approaching testing with an agile, iterative approach you meet your goals faster and more efficiently and, of course, learn more about the system along the way. After we prepare a script for testing (or however the workload is generated), we can run one, a few, and many users (how many depends on the system), analyze results (including resources utilization), and try to sort out any errors. The source of errors can be quite different – script error, functional error, or a direct consequence of a performance bottleneck. It doesn't make much sense to add load until you figure out what is going on. Even with a single script you can find many problems and, at least partially, tune the system. Running scripts separately also allows you to see how much resources are used by each type of load and make some kind of system's "model".

Using the "waterfall" approach doesn't change the nature of performance testing; it just means that you probably do a lot of extra work and still come back to the same point, performance tuning and troubleshooting, much later in the cycle. Not to mention that large tests using multiple use cases are usually a bad point to start performance tuning and troubleshooting - symptoms you see may be a cumulative effect of multiple issues.

Using an agile, iterative approach doesn't mean that you need to re-define the software development process, but rather find new opportunities inside existing processes. I believe that most good performance engineers are already doing performance testing in an agile way but just presenting it as "waterfall" to management (some kind of guerrilla tactic). In most cases you need to present a waterfall-like plan to management, and then you are free to do whatever is necessary to properly test the system inside the scheduled timeframe and scope. If

Questions?

expo:QA '15

Alexander Podelko

alex.podelko@oracle.com
@apodelko

*References may be found in the slide notes
and at www.alexanderpodelko.com*

Madrid, 10th & 11th June 2015

nexo QA
Events

[Agile] Agile Software Development, Wikipedia.
http://en.wikipedia.org/wiki/Agile_software_development

[AgileLoad12] Agile Performance Testing process, AgileLoad whitepaper, 2012.
http://www.agileload.com/agileload//blog/2012/10/22/agile-performance-testing-process---whitepaper

[Bach03] Bach, J. Exploratory Testing Explained, 2003.
http://people.eecs.ku.edu/~saiedian/teaching/Fa07/814/Resources/exploratory-testing.pdf

[Barber07] Barber, S. An Explanation of Performance Testing on an Agile Team, 2007.
http://www.logigear.com/newsletter-2007/320-an-explanation-of-performance-testing-on-an-agile-team-part-1-of-2.html http://www.logigear.com/newsletter-2007/321-an-explanation-of-performance-testing-on-an-agile-team-part-2-of-2.html

[Barber11] Barber, S. Performance Testing in the Agile Enterprise, STP, 2011.
http://www.slideshare.net/rsbarber/agile-enterprise

[Buksh11] Buksh, J. Performance By Design – an Agile Approach, 2011.
http://www.perftesting.co.uk/performance-by-design-an-agile-approach/2011/11/18/

[Capacitas13] Improving the Value of Your Agile Performance Testing (Part 1), Capacitas's blog, 2013.
http://capacitas.wordpress.com/2013/05/20/improving-the-value-of-your-agile-performance-testing-part-1/

[Crispin09] Crispin L., Gregory, J. Agile Testing: A Practical Guide for Testers and Agile Teams. Pearson Education, 2009.

[Dobson07] Dobson, J. Agile Performance Testing, Agile, 2007
http://agile2007.agilealliance.org/downloads/proceedings/075_Performance%20Testing%20on%20an%20Agile%20Project_617.pdf

[Harding11] Harding A. Continuous Integration - A Performance Engineer's Tale, 2011.
http://www.slideshare.net/sthair/continuous-integration-a-performance-engineers-journey