QA to AQ Part Two

Shifting from Quality Assurance to Agile Quality
"Measuring and Monitoring Quality"

Joseph W. Yoder [1] and Rebecca Wirfs-Brock[2]

[1] The Refactory, Inc.,

[2]Wirfs-Brock Associates, Inc.

joe@refactory.com, rebecca@wirfs-brock.com

Abstract. *As organizations transition to agile processes, Quality Assurance (QA) activities and roles need to evolve. Traditionally, QA activities occur late in the process, after the software is fully functioning. As a consequence, QA departments have been "quality gatekeepers" rather than actively engaged in the ongoing development and delivery of quality software. Agile teams incrementally deliver working software. Incremental delivery provides an opportunity to engage in QA activities much earlier, ensuring that both functionality and important system qualities are addressed just in time, rather than too late. Agile teams embrace a "whole team" approach. Even though special skills may be required to perform certain development and Quality Assurance tasks, everyone on the team is focused on the delivery of quality software. This paper is part two of a series of patterns about Agile QA practices and activities. The patterns in this paper are focused primarily on measuring and monitoring system qualities.*

Categories and Subject Descriptors
D.1.5 [**Programming Techniques**]: NEED TO ADD HERE

General Terms
Quality Assurance (QA), Software Developer in Test (SDET), Test Driven Development (TDD), Extreme Programming (XP), Pairing, Agile, Waterfall Methodology, Software Development Lifecycle (SDLC), Continuous Integration (CI), Extract-Transform-Load (ETL) Agile, Patterns, Testing

Keywords
Agile Quality, Quality Assurance, Testing

Introduction

As organizations move to being more agile, it is important that this transition also includes Quality Assurance (QA). Nothing prevents QA from being involved throughout the development process, but generally this has not been the case.

Traditionally, QA only becomes involved late in the development process, just before it is necessary to test and release the final product[1]. This is partly because of a different mindset between QA in traditional software processes and Agile QA. Generally, QA's primary responsibility is to certify the functionality of the application based upon the contract and requirements; usually with black-box tests. Typically, QA groups have worked independently from the software team. However, in agile teams, QA works more closely with the whole team on an ongoing and daily basis.

QA in agile groups are more proactive, working to ensure quality in all facets of the development process. They work closely and coordinate between business, management and developers. Because of this, Agile QA requires additional skills. For example, Agile QA needs to understand both requirements and the code, know how to write their own automated suite cases, and be involved in all aspects of the agile process.

Previously in [YWA] we presented an overview of 22 patterns and wrote patterns for 6 of them and patlets for the rest. A patlet is brief description of a pattern, usually one or two sentences. A list of these patlets is included in the appendix. The patterns written in our previous paper were Integrating Quality into your Agile Process, Agile Quality Scenarios, Quality Acceptance Stories, Fold-Out Qualities, Whole Team and Quality Focused Sprint.

This paper extends that work by writing patterns related to measuring and monitoring qualities: Finding Essential Qualities, Agile Landing Zones, Agree on Quality Targets, Specify Measurable Values for System Qualities, Recalibrate the Landing Zone, and System Quality Dashboards. The patterns are written according to Takashi Iba's Patterns 3.0 pattern languages [IBA] for guiding human action. The most important feature of the Pattern Language 3.0, according to Iba, is that 3.0 pattern readers use patterns to design their own actions. Pattern 3.0 languages are especially useful to those who wish to reflect on their actions and values as they collaborate with others.

Our ultimate goal is to turn all patlets into full-fledged patterns and make a pattern language for action and change in QA.

These patterns are intended for any agile team that wants to focus on important qualities for their systems and better integrating QA into their agile process. These patterns are applicable whether or not you have separate QA teams and roles or extensive agile experience.

[1] "Quality Control" generally refers to inspection activities that occur at the end of a process. "Quality Assurance" or "Total Quality Control" was an alternative to "Quality Control" that recognized that inspection at the end was ineffective and it was more effective to take a more holistic approach building quality into the process from the start. So "Traditionally, QA activities occur late in the process, after the software is fully functioning" is ironic, or perhaps tragic, given that the origin of "Quality Assurance" was about not relying on late-in process activities to produce quality AND to engage everyone in the quality process. The issue is not "Quality Assurance" does activities late so much as too many so-called Quality Assurance departments don't seem to have learned what "Quality Assurance" means. Quoted during an email dialog with one of our reviewers, Jason Yip.

Finding Essential Qualities

"The ability to simplify means to eliminate the unnecessary so that the necessary may speak." —Hans Hofmann

Quite often system essential qualities are overlooked or simplified until late in the development process. This can cause delays due to extensive refactoring and rework of the software design in order to correct quality flaws. To avoid extensive rework it is important that agile teams identify important qualities and make those qualities visible to the team in a timely manner.

How can agile teams understand important qualities for an evolving system?

Not focusing on important qualities early enough can cause significant problems, delays and rework. Correcting functional flaws can be time-consuming. Remedying performance or scalability deficiencies can require significant changes and modifications to the system's architecture.

If important system qualities were identified and addressed during earlier sprints, significant architectural verification could be performed earlier, preventing significant disruptions or delays caused by architectural flaws.

Focusing on non-functional requirements can sometimes distract from important functional requirements outlined by the product owner.

On one hand it would be good if you could develop and automate tests for all qualities. It would be great to test system qualities daily. However, for example, testing something like usability could be expensive and not practical to test so frequently. Trying to find the right balance between how often the tests are ran versus the costs of performing the tests can sometimes b a difficult balancing act.

Therefore, have team meetings at specific times with important stakeholders to brainstorm important qualities that need to be considered for the system. This can be done via a quality attribute workshop where you decide what on critical qualities, and make sure they are visible to team. These workshops should include core members such as the product owner, developers, quality assurance, and the customer.

During a quality workshop, simple collaborative techniques can be used to identify and characterize qualities. A quality chart can be put on a whiteboard that team members use to not specific quality concerns. People can identify a concern and write it on a sticky note that is associated with a specific quality (such as performance or reliability). The team can vote on what they consider most important and urgent and then write quality scenarios for those essential qualities. Teams can use quality sheets or templates to record the quality scenarios.

Workshops can be held at opportune times. Clearly, early on it is important to identify the most important qualities. After formulating a product or project roadmap, which outlines the major features and the order that they could be delivered, potential architectural risks can be identified. Based on these risks, quality-specific concerns can be identified and tied to roadmap features. A rough timeline of when specific qualities and architectural capabilities need to be delivered that enable specific features can also be sketched out. Also, during release planning is another good time to look at qualities. Team members can identify specific qualities that need to be delivered in the release and key activities needed to achieve them.

Agile Landing Zones

"Skate where the puck's going, not where it's been"—Wayne Gretzsky

On a complex project or product, you need to be aware of important system qualities that will make your project a success. You don't want these essential success criteria to get lost in with the myriad of other system requirements.

You also want to make design tradeoffs as you implement your system. Almost always these tradeoffs have architectural implications, so your definition of success needs to be somewhat flexible—you may have to compromise on one design goal in order to achieve another. **How can you understand and monitor those important system qualities that need to be addressed in a way that encourages you to make thoughtful design tradeoffs?**

It is important to identify essential qualities early so that they can be prioritized and contribute to your definition of done. You also need some flexibility in defining what's "good enough". So, you don't want hard and fast (inflexible) acceptance criteria values for all system qualities.

For some system qualities, there isn't one specific number you are aiming for, but you know what is minimally acceptable. For other qualities, you may have specific targets, but you are willing to compromise on them in order to achieve other system quality objectives. You want flexibility in achieving some of quality requirements and overall accountability.

Therefore, define and use an agile landing zone. A landing zone [Gilb] is a set of criteria used to monitor and characterize the "releasability" of a product. An *agile* landing zone is one where not all criteria are fixed at the beginning. The criteria and values of an agile landing zone take shape over the lifetime of a project. Landing zone criteria are similar to release criteria, except they provide for tolerances in acceptable values. There isn't one number you are aiming for; you have a range of values for each attribute you are targeting. This gives you some flexibility in defining what's "good enough", allowing you to make tradeoffs.

Alternatively, where you have the least flexibility in your requirements, you might simply want to define acceptance criteria with specific values that must be met. Only use landing zones for those attributes that have some degree of flexibility in their outcome.

The table below is an example of a landing zone for a loan processing system (all the values have been concocted, for simplicity's sake; any relation to landing zones for real projects is coincidental). Each row represents a task that needs to be performed using the loan processing system.

The values represent the actual time it takes to complete a business task using the system. Tasks may or may not be initiated by users. Some are triggered by incoming data or by a change in the state of a loan (adjusting a loan's interest rate or assigning a loan processor). Others involve adding configuration data, writing additional code and deploying changes to production.

For example, a minimally acceptable time for "Adding a new loan agreement" is two weeks; the target is to enable the user enter all the information for a new loan agreement and have it configured into the system within 24 hours. "Adding a new loan product" is targeted for two weeks because several activities and actions that need to be completed in order to support processing a new kind of loan.

Attribute	Minimum	Target	Outstanding
Adding new loan agreement	2 weeks	24 hours	12 hours
Add new loan product	3 weeks	2 weeks	1 week
Adjust loan	4 days	2 days	1 day
Access loan risk	1 day	6 hours	10 minutes
Assign loan processor	1 month	1 week	1 day

Each row in the landing zone represents a measurable requirement (see Specify Measurable Values for System Qualities) which has a range of acceptable values labeled Minimum, Target, and Outstanding. The goal is to have each requirement within this range at the end of development. Inside the range is the desired value, labeled *Target*. *Minimum*, *Target*, and *Outstanding* are relative to your budget and timeframe.

If you have more than a few attributes, it can be helpful to organize your landing zone according to category: e.g. cost, performance, data quality, reliability, usability, etc.). The table below organizes landing zone criteria according to the specific system quality, or non-functional requirement, being measured.

Performance	Throughput (transactions per day)	50,000	70,000	90,000
	Average transaction time	2 seconds	1 second	<1 second
	…			
Data Quality	Inter system data consistency (percent critical data attributes consistent)	95%	97%	98%
	Data Accuracy	97%	99%	>99%
	…			

A landing zone helps you focus on a few critical things to monitor (instead of hundreds). Your goal should be to only include key criteria that are critical to your project's success. If you do, it will be easier to see a bigger picture and make sense of it: when one attribute is edging below its minimum, what is happening with the others? Are they trending below minimum, too? If so, you have a big problem with achieving your overall product goals. No, and you have a landing zone which allows you to achieve a successful product/system launch even if every requirement isn't exactly on target.

Expect the criteria in an agile landing zone to shift and be adjusted over time. Initially, you may define those parts of your landing zone that you expect to achieve over the next few months, leaving the rest of the landing zone purposefully sketchy. What initially appeared to

be achievable or reasonable targets may change in light of new facts or market changes. No one wants to deliver yesterday's product to today's market. Landing zones, like release criteria can and do change.

For example, you may have worked hard to meet some early achieved landing zone targets, only to find out that your early decisions had negative consequences on future work. You may have created some technical debt that either needs to be paid off in order to achieve your next targets. Given time or budget constraints, you may decide to recalibrate your landing zone (and set expectations lower).

Specify Measurable Values for System Qualities

"Every line is the perfect length if you don't measure it."— Marty Rubin

To know whether a desired quality has been achieved it has to be measured. The description of the quality and the specific aspect you are trying to measure can't be vague or fuzzy. **How can you decide on what values you expect for a quality and how to measure them?**

For system qualities, like performance or throughput, this may be relatively easy: performance can be measured by profiling system performance for a particular scenario, perhaps repeatedly to obtain an overall average. Other qualities, like reliability, may require a complicated set of measures made over a period of time. Some qualities, like usability, at first glance may appear entirely subjective and as a consequence impossibly difficult to measure. High-level quality attributes may need to be decomposed into smaller ones that are measured and aggregated.

Some qualities are difficult or costly to measure. Complex qualities can take quite a bit of time or effort to measure. Balancing the effort required to make a measurement with the information it yields can be difficult.

Therefore, define an appropriate way to measure a quality and to describe it with only as much accuracy and precision as you need. This involves defining or finding an appropriate way to measure (the meter) and describing accurately the values you expect (the scale) [Gilb].

There are three types of scales of measure: natural, constructed, or proxy. A natural scale is one that is obviously associated with a specific quality and is usually the easiest to agree upon. Examples are elapsed time to perform a system operation in milliseconds or the number of page hits in 24 hours. A constructed scale is one that is built specifically to measure a quality, for example, a 7-point user satisfaction scale. Finally, a proxy scale is an indirect measure of quality, for example, projecting system throughput by using sample data and running transaction scripts. You would select a proxy scale if it would be too costly or time-consuming to measure a quality directly.

Since adding necessary precision and accuracy can be difficult, especially for usability qualities, let's illustrate how to improve upon the extremely vague statement, "the system must be easy to use"

A first attempt adds more precision by identifying a specific task and what "easy" means for that task:

> *Eighty percent of novice users should successfully place an order for a single item in under 3 minutes without assistance.*

We can add more details; we're not only want to qualify the speed of placing on order but also whether online help is an aid or a hindrance:

> *Eighty percent of novice users should successfully place an order for a single item in under 3 minutes only using online help for assistance.*

There are two key ideas about measuring "easy to use." First, there is a scale which constrains the possible values of what we are measuring: Time required for a novice to complete a 1-item order using only online help for assistance. Second, there is a meter, which defines how we are going to make our measurement. Since we don't want to only measure

one user and extrapolate to all users, we may decide on average the times obtained for 100 users during testing.

It's best to find a natural scale. People usually won't to argue about it being "good." If you can't find a natural scale, look next for a proxy. You may need to decompose what you are trying to measure into smaller parts and try again. For example, *"Adding a New Loan Agreement"* involves several sub-steps, each requiring time to perform. And you may need several different scenarios to specify expected values under different circumstances.

Finally, you may need to incorporate qualifiers to make things specific when you need the precision. For example, it isn't just any old user's response we're trying to measure, it is:

Time required for a novice to complete a 1-item order using only online help for assistance.

A meter can be an agreed upon way to provide a measurement. To find a meter, look at the scale. If no obvious meter comes to mind, ask others for their experiences or look for "off-the-shelf" tools that come with reasonable meters.

Once you've found a meter, check that:

- Stakeholders agree it is adequate,
- There isn't a more cost effective meter, and
- You can test it on the system, ideally, before it is deployed.

Agree on Quality Targets

"An agreement cannot be the result of an imposition." —Nestor Kirchner

There are several areas where you need to define specific quality-related targets. You may have targets for performance, usability, internationalization, reliability or other non-functional qualities that broadly apply to several user stories or across a number features. Or, you may have a specific system quality that you want to focus on and improve. How much improvement to strive for may be open to debate.

However, if you've done something similar in the past, the quality criteria to choose and their acceptable values may be obvious. At other times it can more of a challenge to reach agreement. **How can you reach consensus when defining quality acceptance criteria?**

Diverse stakeholders have different interests, backgrounds and expectations. Not everyone may be equally informed. Some may hold contradictory opinions. Yet, in order to work towards a specific quality-related objective, everyone needs to buy in and work towards common measurable targets.

Quality requirements priorities are often influenced by the effort to implement them and the effort needed to perform the benchmark.

Therefore, work towards informed consensus on quality-related targets. If you have diverse stakeholders with varying opinions, you may decide to give each stakeholder group a voice in identifying several qualities that are particularly relevant to them. These can be added to landing zone criteria that you've already established.

For first time landing zone builders, you might want to choose someone who knows about the product to take a first cut at establishing landing zone criteria. A business architect, product owner, or lead engineer might prepare a "proposed landing zone" containing reasonable values for quality criteria and values that are questioned, challenged, and then reviewed by a small group. For a landing zone, minimum, target and acceptable values should be agreed upon by the group.

When you are coming up with acceptable values for quality scenarios, you might also use a similar approach. Some informed individual might make a rough cut at "proposed" values that are to be achieved. But a group of informed experts might refine initial values.

Discussions should be to the point, collaborative, and non-confrontational. Someone might propose a set of values based on historical trends and extrapolation. Or a software architect might propose values based on prototyping results or benchmark data. Or the team might declare a design spike to investigate reasonable and possible values. The group might end up agreeing to adjust numbers because the prototyping or design spike evidence was compelling. To effectively set quality value targets, the group should have mutual respect, trust and transparency, and no hidden agendas.

For example, on one program, the chief business architect made the initial cut of quality criteria and their initial values in the landing zone. He was a former techno geek who knew his technical limits. He had deep business knowledge, product vision and a sense about where to be precise and where there should be a lot of flexibility in the landing zone values. Consequently some criteria were very precise. Since they were in the business of processing a lot of transactions, they knew were they needed to improve based on projected increases in

transaction volumes. The transaction throughput target for one business process was based on extrapolations from the existing implementation, taking into account the new architecture and system deployment capabilities. The minimum acceptable value was better than the current implementation (because why else would they be building a new system), but target and outstanding values were based on extrapolations of current capabilities. Other landing zone criteria related to maintainability were only generally categorized as requiring either a patch, a new system release, or online update support. The definitions for what was a patch, a release or an online update were nailed down so that there was no ambiguity in what was meant.

Possible ways of coming up with values include averaging informed individuals' estimates, using an existing system as a baseline, extrapolating values from similar scenarios, or benchmarking working code. Sometimes it may be necessary to create a spike solution to obtain estimates.

To reach consensus on specific quality scenario targets, you may need someone to play the role of facilitator. The facilitator should know enough about the program or product to be constructive, but they need not be the "authority" or "expert". That person should be good at gaining consensus and get the best from individuals who may have strongly held opinions and disagreements. Ideally, a facilitator knows enough about the product to offer constructive observations and has the ability to lead a small group forward in defining acceptable criteria and values. It can be more effective to have an informed facilitator to guide quality target definitions, than a dispassionate, uniformed one.

Recalibrate the Landing Zone

"*Test fast, fail fast, adjust fast."*— Tom Peters

Initially, you defined a set of landing zone criteria that you expected to achieve over a few iterations. You left the rest of your landing zone purposefully sketchy. As you've implemented new functionality, you have continued to add new landing zone criteria while monitoring the values of existing ones. **How can you continue to evolve your landing zones and keep them update to date?**

As you continue with development, it can become harder to keep criteria within their landing zone target values. Solutions that achieve newly identified landing zone criteria may impact your ability to maintain other values.

What initially appeared to be achievable targets may change in light of new information and your current implementation.

Although the product owner or client will see these targets as important, they maybe see these as a lower priority over other things that need to be done and may not understand the implication to the overall system.

Budget and Time constraints can limit the effort you are able to devote to achieving important quality constraints.

Therefore, rather than simply throw up your hands in defeat, revisit your landing zone criteria and reset expectations. Some values may not be appropriate, given what you know now.

It is natural for the criteria in an agile landing zone and their values to shift and be adjusted over time. What initially appeared to be reasonable targets may change in light of new facts or market changes. No one wants to deliver yesterday's product to today's market. Previous implementation decisions can affect or limit your ability to achieve newly identified criteria. For example, deciding to focus on meeting specific performance targets may have impacted usability or flexibility criteria.

Landing zones, like release criteria can and do change. In fact changing acceptable values for your landing criteria is not always a bad thing to do, especially if you are reacting to the current situation and making thoughtful design tradeoffs.

For example, you may have created some technical debt that needs to be paid off in order to achieve some landing zone performance targets. Given time and budget constraints, you decide not to invest in design rework for the current release. It is more important to deliver working functionality on time than make it fast. So, you opt to recalibrate your landing zone (setting acceptable performance criteria lower). You've made a conscious decision to redefine what is acceptable.

You might also recalibrate/readjust landing zone criteria upwards based on new information/system capabilities/technologies. For example, with experimentation you find that by tweaking cache and buffer sizes, you can increase throughput for an important data translation (ETL) process. Rather than simply move into the "outstanding" range, you also adjust the minimum acceptable value upwards and note that cache and buffer tuning should be considered for any time critical ETL process.

System Quality Dashboards

"The dashboard needs to deliver data in a timely fashion, and that timeliness is dictated by which process is being represented in the dashboard" —Keith Gile
http://thinkexist.com/quotes/keith_gile

Typically, agile software development focuses on features and functionality first before paying attention to other important system aspects such as architecture and critical qualities. On agile projects you hear statements like, "Make it work, make it right, then optimize it". Most agile practices push to develop important functional requirements as outlined by the product owner, which are prioritized on the work backlog. As the system evolves the team begins to better understand what system qualities are important and how to better measure them. As the system evolves, keeping track of these qualities becomes increasingly important.

How can agile teams provide a means to make this information accessible and visible to the team?

❖ ❖ ❖

Creating tools and dashboards takes time and often there are limited resources and people dedicated to building QA tools. There is often not enough time committed to consider important quality implications of one's design and what is important to measure and monitor.

It can be difficult to know what qualities are important to monitor. As more and more qualities are built into the system, certain ones are important to keep a watch on while others once validated and made testable are good enough.

Creating a lot of tools and dashboards can seem like a pointless luxury compared to making sure the system is meeting the requirements well enough to ship.

Certain qualities such as performance and reliability, if not regularly tracked can become very difficult to improve late in the development process. Although originally the system might meet quality constraints, as the system evolved sometimes qualities become invisible and as a consequence aren't maintained over time.

❖ ❖ ❖

Therefore, create dashboards to test and validate important qualities where needed. These dashboards can overlap with operational dashboards for production systems. As important system qualities are outlined and included in the backlog, note which ones should be monitored and where tools can be created to measure the system as it evolves. Dashboards can show real-time results, for example, performance of running processes, or display quality values measured during check-in or system build by quality tests.

Important considerations that need to be answered when considering dashboards are: When should I build a tool that creates a dashboard? When should I buy one? If a tool provides all that is needed and a relatively easy way to make measurements, then use the tool. Another consideration is how easy the tool is to set up or if it "works out of the box"[ref]. Sometimes tools that provide powerful means of measurement may be costly. So you have to make a tradeoff between purchasing a powerful tool or using an open source dashboard that may not be as powerful. Another consideration is how well the dashboard integrates with your development environment.

Additionally, whether you buy or build a tool, consider what quality aspects should be shown and what frequency that they are measured. Does the dashboard perform quality-related tests when initiated by a user, showing results of tests executed during build or integration, or is in

monitoring the production system? How frequently should contents of a dashboard be updated in order to be useful and what happens when measured values fall below minimally acceptable criteria? Some dashboard tools allow you to configure alerts and notifications when measured values cross a threshold.

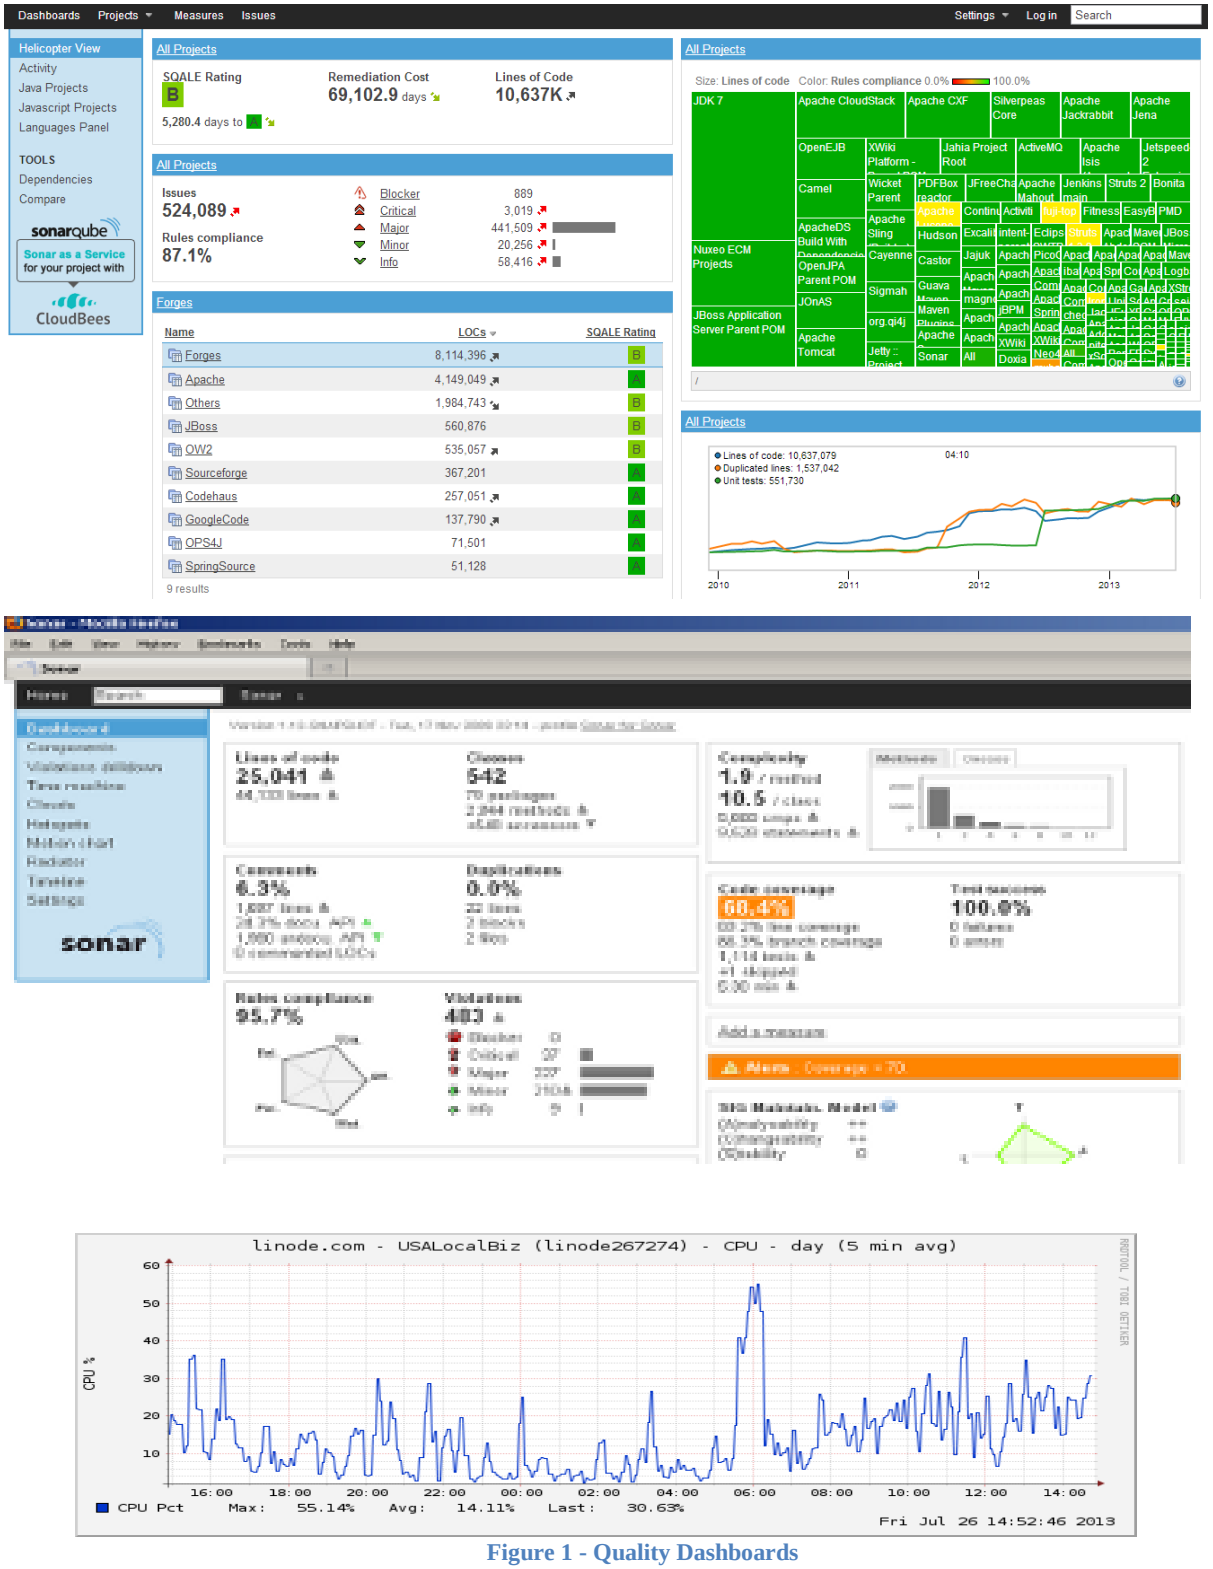

Figure 1 - Quality Dashboards

Summary

This paper outlined a subset of important patterns when shifting from Quality Assurance (QA) to Agile Quality (AQ). The complete set includes both ways of incorporating QA into the agile process and agile techniques for describing, measuring, adjusting, and validating important system qualities. This paper wrote a few of these that focuses on measuring and monitoring qualities as patterns. Ultimately it is the authors' plan to write all of the patlets as patterns and weave them into a 3.0 pattern language for evolving from Quality Assurance to an Agile Quality mindset.

Acknowledgements

We thank our shepherd Eduardo Guerra for his valuable comments and feedback during the PLoP 2014 shepherding process. We also thank our 2014 PLoP Writers Workshop Group, xxx, yyy, and zzz, for their valuable comments.

Note to Shepherd and Workshop Reviewers

Our references have not been completed yet as we are working on these for the final paper. We will go through all our references and update them. The following section is a place holder for when we complete this task.

References

[GILB] Gilb, Tom. 2005. *Competitive Engineering: A Handbook For Systems Engineering, Requirements Engineering, and Software Engineering Using Planguage*. Butterworth-Heinemann.

[FY98] Foote B., and Yoder, J. 1998. *The Selfish Class*. Proceedings of the 3rd Conference on Pattern Languages of Programs (Monticello, Illinois). PLoP '96, Technical Report #WUCS-97-07, September 1996, Department of Computer Science, Washington University. Pattern Languages of Program Design 3 edited by Robert Martin, Dirk Riehle, and Frank Buschmann, Addison-Wesley, 1998

[IBA] Iba, T. 2011. "Pattern Language 3.0 Methodological Advances in Sharing Design Knowledge," International Conference on Collaborative Innovation Networks 2011 (COINs2011).

Appendix

A previous paper on this topic outlines some core patterns when evolving from traditional quality assurance to being agile at quality [ref]. We outlined all that patterns using patlets. A patlet is a brief description of a pattern, usually one or two sentences. Following is an excerpt from that paper outlining the patlets.

Central to successfully using these QA patterns is knowing where quality concerns can fit into your agile process. The following patlet describes those considerations.

Patlet Name	Description
Integrating Quality into your Agile Process	Incorporate QA into your process including a lightweight means for describing and understanding system qualities.

Identifying Qualities

An important but difficult task for software development teams is to identify the important qualities (non-functional requirements) for a system. Quite often system qualities are overlooked or simplified until late in the development process, thus causing time delays due to extensive refactoring and rework of the software design required to correct quality flaws. It is important in agile teams to identify essential qualities and make those qualities visible to the team. The following patlets support identifying the qualities:

Patlet Name	Description
Finding the Qualities	Brainstorm the important qualities that need to be considered.
Agile Quality Scenarios	Create high-level quality scenarios to examine and understand the important qualities of the system.
Quality Stories	Create stories that specifically focus on some measurable quality of the system that must be achieved.
Specify Measureable Values or System Qualities	Specify scale, meter, and values for specific system qualities.
Fold-out Qualities	Define specific quality criteria and attach it to a user story when specific, measurable qualities are required for that specific functionality.
Agile Landing Zone	Define a "landing zone" that defines acceptance criteria values for important system qualities. Unlike traditional "landing zones", an agile landing zone is expected to evolve during product development.
Recalibrate the Landing Zone	Readjust landing zone values based on ongoing measurements and benchmarks.
Agree on Quality Targets	Define landing zone criteria for quality attributes that specify a range of acceptable values: minimally acceptable, target and outstanding. This range allows developers to make tradeoffs to meet overall system quality goals.

Making Qualities Visible

It is important for team members to know important qualities and have them presented so that the team is aware of them. The following patlets outline ways to make qualities visible:

Patlet Name	**Description**
System Quality Dashboard	Define a dashboard that visually integrates and organizes information about the current state of the system's qualities that are being monitored.
System Quality Radiator	Post a display that people can see as they work or walk by that shows information about system qualities and their current status without having to ask anyone a question. This display might show current landing zone values, quality stories on the current sprint or quality measures that the team is focused on.
Qualify the Roadmap	Examine a product feature roadmap to plan for when system qualities should be delivered.
Qualify the Backlog	Create quality scenarios that can be prioritized on a backlog for possible inclusion during sprints.
Quality Chart	Create a chart or listing of the important qualities of the system and make them visible to the team; possibly on the agile board.

Being Agile at Quality

In any complex system, there are many different types of testing and monitoring, specifically when testing for system quality attributes. QA can play an important role in this effort. The role of QA in an Agile Quality team includes: 1) championing the product and the customer/user, 2) specializing in performance, load and other non-functional requirements, 3) focusing quality efforts (make them visible), and 4) assisting with testing and validation of quality attributes. The following patlets support "Becoming Agile at Quality":

Patlet Name	**Description**
Whole Team	Involve QA early on and make QA part of the whole team.
Quality Focused Sprints	Focus on your software's non-functional qualities by devoting a sprint to measuring and improving one or more of your system's qualities.
QA Product Champion	QA works from the start understanding the customer requirements. A QA person will collaborate closely with the Product owner pointing out important Qualities that can be included in the product backlog and also work to make these qualities visible and explicit to team members.
Agile Quality Specialist	QA provides experience to agile teams by outlining and creating specific test strategies for validating and monitoring important system qualities.
Monitoring Qualities	QA specifies ways to monitor and validate system qualities.
Agile QA Tester	QA works closely with developers to define acceptance criteria and tests that validate these, including defining quality scenarios and tests for validating these scenarios.
Spread the Quality Workload	Rebalance quality efforts by involving more than just those who are in QA work on quality-related tasks. Another way to spread the work on quality is to include quality-related tasks throughout the project and not just at the end of the project.
Shadow the Quality Expert	Spread expertise about how to think about system qualities or implement quality-related tests and quality-conscious code by having another person spend time working with someone who is highly skilled and knowledgeable about quality assurance on key tasks.
Pair with a Quality Advocate	Have a developer work directly with quality assurance to complete a quality related task that involves programming.