


The TIOBE Quality Indicator

a pragmatic way of measuring code quality



Document ID: TIOBE-20120718.1
Version: 2.2 (authorized)
Date: 10-Mar-2014
Author: Paul.Jansen@tiobe.com



Table of Contents

1 Introduction	2
2 Software Quality Attributes	2
3 Software Metrics	3
4 Mapping Software Metrics to Quality Attributes	3
5 Measuring and Judging Metric Values	7
6 TIOBE Compliance Factor	11
7 TIOBE Quality Indicator	12
7.1 Scope.....	12
8 Conclusions	12
Appendix A References	13
Appendix B Reviewers	13

1 Introduction

The proverb "the proof of the pudding is in the eating" applies perfectly to software quality. Because only after a software product has been shipped, the true quality of a software product reveals itself. Software quality is determined by

- the number of defects found after release
- the severity of these defects
- the effort needed to solve these defects

More than 30 years ago, software engineer Barry Boehm already observed that the costs of repairing defects increase exponentially if they are found later on in the software development process [1]. So if it is possible to have a way to measure the software quality of a system before release, it will potentially save a lot of money.

The goal of this document is to define such a software quality measurement system based on a pragmatic approach. The focus is on code quality (as opposed to e.g. quality of requirements or the architecture). The defined approach is based on more than 10 years of experience in this field and the analysis of more than 200 million lines of industrial production software code that are checked each day.

2 Software Quality Attributes

There is an ISO definition of software quality, called ISO 25010 [2]. This standard defines 8 main quality factors and a lot of subattributes. The 8 main quality factors are:

- **Functional suitability.** The degree to which the product provides functions that meet stated and implied needs when the product is used under specified conditions.
- **Reliability.** The degree to which a system or component performs specified functions under specified conditions for a specified period of time.
- **Performance efficiency.** The performance relative to the amount of resources used under stated conditions.
- **Operability.** The degree to which the product has attributes that enable it to be understood, learned, used and attractive to the user, when used under specified conditions.
- **Security.** The degree of protection of information and data so that unauthorized persons or systems cannot read or modify them and authorized persons or systems are not denied access to them.
- **Compatibility.** The degree to which two or more systems or components can exchange information



- and/or perform their required functions while sharing the same hardware or software environment.
- **Maintainability.** The degree of effectiveness and efficiency with which the product can be modified.
- **Transferability.** The degree to which a system or component can be effectively and efficiently transferred from one hardware, software or other operational or usage environment to another.

The ISO 25010 standard helps as a starting point to determine quality in an early stage. It has 2 main drawbacks, however:

- The standard does not specify how to measure quality attributes. Some of the quality attributes even seem unfit for objective measurement. Take "Operability" for instance, with subattributes such as "Attractiveness" and "Ease of Use". How to measure this and what is the unit of measurement?
- Most of the quality attributes defined have different meanings in different contexts. So even if it is possible to measure a quality attribute, it is impossible to define clear objective criteria for what is considered good or bad. "Performance efficiency" is a good example of such a quality attribute. For some software systems a response within 1 second is sufficient, whereas others demand a response within 1 millisecond.

3 Software Metrics

We could try to define measurement systems for the ISO quality attributes in a top-down scientific way. However, this is too ambitious.

Can we measure anything at all? Yes, but we need to take a more pragmatic approach. Lots of metrics are applied to software code nowadays, but unfortunately there is insufficient proof (yet) whether these metrics contribute to better code. Examples of such metrics are cyclomatic complexity [3], code duplication [4] and all kinds of code coverage [5]. These metrics are approximations of some of the quality attributes of the ISO 25010 standard.

To obtain a systematic way of measuring and qualifying these measurements, the 8 most commonly used software code quality metrics in industry today have been selected that can be measured in an automated way. These are:

1. Code coverage
2. Abstract interpretation [6]
3. Cyclomatic complexity
4. Compiler warnings
5. Coding standards [7]
6. Code duplication
7. Fan out [8]
8. Dead code [9]

4 Mapping Software Metrics to Quality Attributes

We define the metrics of the previous section and map them on the quality attributes of the ISO 25010 standard.

1. **Code coverage.** Before software engineers hand over their code to the next stage in the software development cycle, they usually perform unit tests. These are small automated tests that check a particular part of a program such as a single function. The actual results of these automated tests are compared to the expected results. Unit tests are a powerful way to check whether a program behaves like it is designed to behave at the lowest level. The code coverage metric indicates how many lines of code or executable branches in the code have been touched during the unit test runs.



The lower the coverage, the lower the quality of the performed unit tests. Code coverage is an indicator of both "Functional Suitability" and "Reliability".

A simple example of the output of a code coverage tool is shown for the C# code below. Every line that is coloured "green" is touched during at least one of the tests, whereas "red" lines are not touched by any test.

```
25: if (
26:     element.ElementType == ElementType.Class &&
27:     element.Declaration.Name.EndsWith("Class")
28: )
29: {
30:     addViolation(element, element.Declaration.Name, "Class");
31: }
32: else if (
33:     element.ElementType == ElementType.Struct &&
34:     element.Declaration.Name.EndsWith("Struct")
35: )
36: {
37:     addViolation(element, element.Declaration.Name, "Struct");
38: }
39: return true;
```

The output of the code coverage tool shows that all lines in this code sample are covered by (unit) tests, except for line 37.

- Abstract Interpretation.** A fairly new technology is to detect possible reliability issues in software programs by running *abstract interpretation* tools, also known as *deep flow analysis* tools. These tools are capable of automatically detecting all kinds of programming errors related to the control flow of a program. Examples are null pointer dereferences, buffer overflows and unclosed database connections. The advantage of these tools is that they generate their results without actually running the programs. This is done by calculating all possible paths through a program in an efficient way. Errors found by abstract interpretation are severe programming errors that may result in crashes. This metric is mapped to the "Reliability" attribute.

A simple example of an abstract interpretation issue is shown in the Java code below.

```
159: public Order getOrder() {
160:     // Only return orders with a valid date
161:     if (orderDate.isValid()) {
162:         return order;
163:     } else {
164:         return null;
165:     }
166: }
...
227: public List<Order> getOrderPackages() {
228:     return getOrder().getCorrespondingOrderPackages(company);
229: }
```

Abstract interpretation tools will flag a possible null pointer dereference at line 228, because the function "getOrder" can return null in case the order has no valid date. If this situation occurs an exception will be thrown, possibly resulting in program abortion.

- Cyclomatic complexity.** One of the oldest software metrics is cyclomatic complexity. Cyclomatic



complexity counts the number of independent paths through a program. For instance, each "if" statement adds one extra code path. The higher the cyclomatic complexity the harder it is to understand a program. Moreover, the more paths there are, the more test cases need to be written to achieve a decent code coverage. The average cyclomatic complexity per function is an indicator that enables comparisons of complexity between programs. It is part of the "Maintainability" attribute.

The C# code below shows a simple example of how cyclomatic complexity is calculated.

```
123: public int getValue(int param1)
124: {
125:     int value = 0;
126:     if (param1 == 0)
127:     {
128:         value = 4;
129:     }
130:     else
131:     {
132:         value = 0;
133:     }
134:     return value;
135: }
```

The cyclomatic complexity of the function "getValue" at line 123 is 2 (one path through "then" and one through "else").

- 4. Compiler warnings.** In order to execute a software program on a computer it first must be compiled or interpreted. Compilers/interpreters generate errors and warnings. Errors must be fixed otherwise the program cannot run. Warnings on the other hand do not necessarily need to be solved. However, some compiler warnings indicate serious program flaws. Leaving these unresolved has probably impact on the "Reliability" of the code. Apart from this, most compilers also warn about portability issues. So this metric can also mapped to "Transferability" in most cases.

A simple example of a compiler warning is shown in the C code below.

```
31: int func(int i) {
32:     if (i = 0) {
33:         return -1;
34:     }
...
58: }
```

Most compilers will complain about the assignment in the if condition at line 32 (probably a comparison was meant instead).

- 5. Coding standards.** Software maintenance is one of the most time consuming tasks of software engineers. One of the reasons for this is that it is hard to understand the intention of program code long after it has been written, especially if it has been updated a lot of times. A way to reduce the costs of software maintenance is to introduce a coding standard. A coding standard is a set of rules that engineers should follow. These coding rules are about known language pitfalls, code constructions to avoid, but also about naming conventions and program layout. Since coding standards usually contain many different rules they can be mapped to most quality attributes. Most rules concern "Maintainability" and "Reliability", but there are also rules available for "Transferability", "Performance Efficiency" and "Security".



An example of a coding standard violation is shown below.

```
31: int abs(int i) {
32:     int result;
33:
34:     if (i < 0) {
35:         result = -i;
36:         goto end;
37:     }
38:     result = i;
39: end:
40:     return result;
41: }
```

Any C coding standard will complain about the goto statement used at line 36. It is considered bad practice to use goto statements.

- Code duplication.** Sometimes, it is very tempting for a software engineer to copy some piece of code and make some small modifications to it instead of generalizing functionality. The drawback of code duplication is that if one part of the code must be changed for whatever reason (solving a bug or adding missing functionality), it is very likely that the other parts ought to be changed as well. But who is to notice? If nobody does, code duplication will lead to rework in the long term. This has a negative effect on "Maintainability".
- Fan out.** Software programs are structured in terms of *modules* or *components*. These modules and components "use" each other. The fan out metric indicates how many different modules are used by a certain module. If modules need a lot of other modules to function correctly (high fan out), there is a high interdependency between modules, which makes code less modifiable. Hence, fan out is mapped to the "Maintainability" ISO attribute.

An example of a high fan out is shown in the Java code below.

```
1: package com.tiobe.plugins.eclipse.analyzer;
2:
3: import java.io.IOException;
4: import java.util.Map;
5:
6: import org.apache.commons.exec.CommandLine;
7: import org.apache.commons.exec.DefaultExecutor;
8: import org.apache.commons.exec.ExecuteException;
9: import org.apache.commons.exec.ExecuteResultHandler;
10: import org.apache.commons.exec.ExecuteWatchdog;
11: import org.apache.commons.exec.Executor;
12: import org.apache.commons.exec.PumpStreamHandler;
13: import org.apache.commons.exec.environment.EnvironmentUtils;
14: import org.apache.commons.io.output.NullOutputStream;
15: import org.eclipse.core.resources.IProject;
16: import org.eclipse.core.resources.IResource;
17:
18: import com.tiobe.plugins.eclipse.console.ITICSConsole;
19: import com.tiobe.plugins.eclipse.console.TICSConsole;
20: import com.tiobe.plugins.eclipse.util.EclipseUtils;
21:
22: public class TICSanalyzer implements ITICSAnalyzer {
```



In this article we have adopted the simple definition of fan out to measure the number of `import` statements. Hence, the fan out of the Java file above is 16.

- Dead code.** If requirements change, some programming code might become obsolete. Due to time pressure, lack of rigorous maintenance or even just ignorance, such dead code might linger in software archives for years. The drawback of dead code is that it consumes unnecessary maintenance effort. Removing dead code increases "Maintainability".

An example of dead code is given in the following C code.

```
318: int calculateindex(unsigned int i) {
319:     if (i < 0) {
320:         return 0;
321:     } else {
322:         return i;
323:     }
324: }
```

The code at line 320 is never executed because unsigned integers can never be less than 0.

5 Measuring and Judging Metric Values

This section defines how the 8 metrics of the previous section are measured. It also specifies how the obtained metric values are to be judged on a scale between 0 and 100 (called the *score*). The formulas that are used to calculate the scores for metrics have been determined empirically, based on analyzing the more than 200 million lines of code that are checked by TIOBE Software each day.

If one of the metrics is not measured at all for whatever reason, the score for this metric is 0.

There are 6 different categories distinguished based on the normative system. These are similar to the European Union energy labels [10]. See the table below.

Category	Name	TQI Score
	Outstanding	$\geq 90\%$
	Good	$\geq 80\%$
	Fairly Good	$\geq 70\%$
	Moderate	$\geq 50\%$
	Weak	$\geq 40\%$
	Poor	$< 40\%$

The category "Moderate" is a bit larger than the other categories to create a Gaussian-like distribution.

The 8 metrics are measured and valued in the following way.

- Code coverage.** Code coverage is measured by taking the average of the available decision, branch and statement coverage figures. At least one of these three coverage types should be available. This is based on research done by Steve Cornett [11]. The following formula is applied to value code coverage:



$$\text{score} = \min(0.75 * \text{test_coverage} + 32.5, 100)$$

The definition of the code coverage score is based on the fact that a code coverage above 90% is perfect (score = 100). Improving the code coverage if it is above 90% is not worthwhile. On the other hand, if the code coverage is above 10%, i.e. if at least some tests are performed, the score should be in category E (score = 40). Between 40 and 100, the score is evenly distributed.

According to this formula the mapping to the code quality categories is as follows.

Code Coverage	TQI Score	Category
>= 76.7%	>= 90%	
>= 63.3%	>= 80%	
>= 50%	>= 70%	
>= 23.3%	>= 50%	
>= 10%	>= 40%	
< 10%	< 40%	

Table 1: Code Coverage Scores

2. **Abstract Interpretation.** Abstract interpretation results are measured by taking all errors found by the abstract interpreter. The resulting set of errors is mapped to a scale between 0% and 100% via the TIOBE compliance factor definition [12]. The following formula is applied to the compliance factor to get the scores for abstract interpretation.

$$\text{score} = \max(\text{compliance_factor}(\text{abstract_interpretation_violations}) * 2 - 100, 0)$$

Abstract interpretation errors are considered to be important, so its perfect score (score = 100) means there are no abstract interpretation errors at all. A compliance less than 70% indicates that there are lots of such errors, thus having a poor score (score = 40). The score is distributed evenly.

According to this formula the mapping to the code quality categories is as follows.

Compliance Factor	TQI Score	Category
>= 95%	>= 90%	
>= 90%	>= 80%	
>= 85%	>= 70%	
>= 75%	>= 50%	
>= 70%	>= 40%	
< 70%	< 40%	

Table 2: Abstract Interpretation Scores

3. **Cyclomatic complexity.** The definition of cyclomatic complexity has been given by McCabe [13]. This definition is also used in this article. The average cyclomatic complexity per function is mapped on a normative scale by using the formula

$$\text{score} = \min(\max(140 - 20 * \text{cyclomatic_complexity}, 0), 100)$$

The definition of the cyclomatic complexity score is based on the fact that an average complexity of

less than 3 is considered to be good, and an average complexity of more than 5 is very bad. The average cyclomatic complexity of more than 200 million lines of industrial software code as checked by TIOBE is 2.97.

According to this formula the mapping to the code quality categories is as follows.

Cyclomatic Complexity	TQI Score	Category
<= 2.5	>= 90%	
<= 3	>= 80%	
<= 3.5	>= 70%	
<= 4.5	>= 50%	
<= 5	>= 40%	
> 5	< 40%	

Table 3: Cyclomatic Complexity Scores

- 4. Compiler warnings.** Compiler warnings are measured by running the compiler used at the highest possible warning level. If more than one compiler is used (e.g. because code is generated for multiple platforms), the warnings of all compilers are combined. Since different compilers check for different compiler warnings, it is not sufficient to use the number of compiler warnings as input for the score. Hence, the set of compiler warnings should be normalized, based on the number of different checks a compiler performs and the severity of these checks. TIOBE uses its compliance factor for this [12], which is a figure between 0 (no compliance) and 100 (complete compliance, i.e. no compiler warnings). A brief summary of the way the TIOBE compliance factor is calculated is given in the next section.

Once the compliance factor is known, the following formula is applied to determine the score for compiler warnings:

$$\text{score} = \max(100 - 50 * \log_{10}(101 - \text{compliance_factor}(\text{compiler_warnings})), 0)$$

This rather complex formula is based on the observation that most compilers have lots of different warnings and most of these warnings don't occur in the software. Hence, the compliance will be high most of the cases. That's why a logarithmic function is used.

According to this formula the mapping to the code quality categories is as follows.

Compliance Factor	TQI Score	Category
>= 99.42%	>= 90%	
>= 98.49%	>= 80%	
>= 97.02%	>= 70%	
>= 91.00%	>= 50%	
>= 83.22%	>= 40%	
< 83.22%	< 40%	

Table 4: Compiler Warning Scores

- 5. Coding standards.** It is important to make sure that as many coding standard rules as possible are automated by code checkers. For this metric only automated rules are taken into account. It is



assumed that the coding rules have been categorized in severity levels. The resulting set of coding rule violations is mapped to a scale between 0 and 100 via the TIOBE compliance factor definition [12]. Coding standards are mapped on a normative scale by using the formula

$$\text{score} = \text{compliance_factor}(\text{coding_standard_violations})$$


Compliance Factor	TQI Score	Category
$\geq 90\%$	$\geq 90\%$	
$\geq 80\%$	$\geq 80\%$	
$\geq 70\%$	$\geq 70\%$	
$\geq 50\%$	$\geq 50\%$	
$\geq 40\%$	$\geq 40\%$	
$< 40\%$	$< 40\%$	

Table 5: Coding Standard Scores

6. **Code duplication.** This metric is calculated by counting the number of semantically equivalent chains of 100 tokens (default for most tools). The total number of lines of code that contains a chain is taken and expressed as a percentage of the total size of the system. Code duplication is mapped on a normative scale by using the formula

$$\text{score} = \min(-20 * \log_{10}(\text{code_duplication}) + 60, 100)$$

According to this formula the mapping to the code quality categories is as follows.

Code Duplication	TQI Score	Category
$\leq 0.03\%$	$\geq 90\%$	
$\leq 0.10\%$	$\geq 80\%$	
$\leq 0.32\%$	$\geq 70\%$	
$\leq 3.16\%$	$\geq 50\%$	
$\leq 10.00\%$	$\geq 40\%$	
$> 10.00\%$	$< 40\%$	

Table 6: Code Duplication Scores

7. **Fan out.** The fan out is measured by counting the average number of imports per module. This measurement is language dependent. For C and C++ the number of `include` statements is used, for Java the number of `import` statements. Wild cards in Java `import` statements and `using` statements in C# in general appear to be difficult because these statements import several modules at once. That is why we choose to count these statements as 5. The average fan out of a software system is mapped on a normative scale by using the formula

$$\text{score} = \min(\max(120 - 5 * \text{fan_out}, 0), 100)$$

According to this formula the mapping to the code quality categories is as follows.



Fan Out	TQI Score	Category
<= 6	>= 90%	
<= 8	>= 80%	
<= 10	>= 70%	
<= 14	>= 50%	
<= 16	>= 40%	
> 16	< 40%	

Table 7: Fan Out Scores

8. **Dead code.** Dead code is simply expressed as the percentage of code in LOCs that is not reachable via any program execution. The dead code rate is mapped on a normative scale by using the formula

$$\text{score} = \max((100 - 2 * \text{dead_code}), 0)$$

According to this formula the mapping to the code quality categories is as follows.

Dead Code	TQI Score	Category
<= 5%	>= 90%	
<= 10%	>= 80%	
<= 15%	>= 70%	
<= 25%	>= 50%	
<= 30%	>= 40%	
> 30%	< 40%	

Table 8: Dead Code Scores

6 TIOBE Compliance Factor

Three of the eight metrics of the code quality system in this article use the TIOBE compliance factor to express compliance to a set of rules. In this section, the TIOBE compliance factor is explained in more detail. Its mathematical definition can be found in [12], where it is known as the *TIOBE confidence factor*.

Basically, the TIOBE compliance factor expresses how much code complies to a certain set of rules. This could be a set of compiler warnings or a set of coding standard rules. This is needed because the number of compiler warnings or coding standard violations in itself doesn't say anything about code quality. If there are 3,000 compiler warnings left in your code is that all right or plain wrong?

There are various parameters that play a role in mapping the number of warnings/violations to a number between 0 and 100 in a sensible way. These are:

- Total number of distinct checks/rules measured. If your compiler only measures one kind of check, the number of compiler warnings will usually be lower if compared to a compiler that measures hundreds of different checks.
- Severity level of these checks/rules. Compiler warnings that indicate programming errors should be weighed as more important than compiler warnings that identify minor issues.
- Size of the software system. The larger a software system, the higher the chances to have more



warnings/violations.

- The lines of code that could be checked. If only half of your code can be compiled, then the number of compiler warnings (not compiler errors) will be lower if compared to a situation in which all code could be compiled and checked.

All these parameters are taken into account by the TIOBE compliance factor. Its definition is:

TIOBE Compliance Factor = Percentage of code checked * Weighted warnings

So if only 20% of the code could be checked, the TIOBE compliance factor will never exceed 20%. A *weighted warning* is calculated by dividing the number of warnings of a specific kind by the total physical lines of code (warning density) and the number of rules for the severity level of the kind of compiler warning/coding rule. Then we divide by 4 to the power of its severity level. For instance, a compiler warning of level 3 is considered 16 times less important than a compiler warning of level 1.

7 TIOBE Quality Indicator

The 8 code quality metrics defined in this article all help to get a complete picture of the code quality before release. However, not all code quality metrics are equally important. For instance, a low code coverage has much more impact on quality than a high dead code rate. This section defines how the 8 metrics are combined into one overall code quality figure, called the TIOBE Quality Indicator (TQI).

The metrics are combined by weighing them. This is based on empirical evidence. It is important to note that TIOBE has started research to correlate software defects to code quality metrics for the more than 200 million lines of code it measures each day. Once this research has been finished, the weighing will be more solidly founded on statistical data.

The 8 metrics are weighted as follows.

Metric	Weight
Code Coverage	20%
Abstract Interpretation	20%
Cyclomatic Complexity	15%
Compiler Warnings	15%
Coding Standards	10%
Code Duplication	10%
Fan Out	5%
Dead Code Detection	5%

Table 9: Weights of TQI metrics

7.1 Scope

Not all code is subject to the TQI. The following kinds of code are excluded:

- Generated code
- External/third party code
- Test code



7.2 Recommended TQI Levels

What TQI level should one aim for? That depends on your application domain. Different domains have different quality constraints. For instance, for avionics we recommend TQI level A but for administrative software we only require TQI level D. The recommended levels for various application domains can be found in the table below.

Application Domain	TQI Level Goal
Avionics, Defense	 A
Space, Medical, Automotive	 B
Semiconductors, Bioinformatics	 C
Administrative	 D

Table 10: Recommended TQI Levels

8 Conclusions

The TIOBE Quality Indicator (TQI) is a pragmatic way to get an overview of the quality of software code before release or even before system testing. The indicator combines the most well-known code quality metrics by defining how they are measured and how the outcome of the resulting measurements should be judged. Based on this a software system is labelled between A (outstanding quality) and F (poor quality).

Appendix A References

- [1] Boehm, Barry W.; Philip N. Papaccio, "Understanding and Controlling Software Costs", IEEE Transactions on Software Engineering, v. 14, no. 10, October 1988, pp. 1462-1477.
- [2] ISO, "Systems and software engineering – Systems and software Quality Requirements and Evaluation (SquaRE) – System and software quality models", ISO/IEC 25010:2011, 2011, obtainable from http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=35733.
- [3] Wikipedia, "Cyclomatic Complexity", extracted July 2012, obtainable from http://en.wikipedia.org/wiki/Cyclomatic_complexity.
- [4] Wikipedia, "Duplicated Code", extracted July 2012, obtainable from http://en.wikipedia.org/wiki/Duplicate_code.
- [5] Wikipedia, "Code Coverage", extracted July 2012, obtainable from http://en.wikipedia.org/wiki/Code_coverage.
- [6] Wikipedia, "Abstract Interpretation", extracted July 2012, obtainable from http://en.wikipedia.org/wiki/Abstract_interpretation.
- [7] Wikipedia, "Coding Conventions", extracted July 2012, obtainable from http://en.wikipedia.org/wiki/Coding_standard.
- [8] Henry, S.; Kafura, D., "Software Structure Metrics Based on Information Flow", IEEE Transactions on Software Engineering Volume SE-7, Issue 5, September 1981, pp. 510–518.
- [9] Wikipedia, "Dead Code", extracted July 2012, obtainable from http://en.wikipedia.org/wiki/Dead_code.
- [10] Wikipedia, "European Union energy label", extracted July 2012, obtainable from http://en.wikipedia.org/wiki/European_Union_energy_label.
- [11] Cornett, Steve, "Code Coverage Analysis", obtainable from <http://www.bullseye.com/coverage.html>.



[12] Jansen, Paul; Krikhaar, Rene; Dijkstra, Fons, "*Towards a Single Software Quality Metric – The Static Confidence Factor*", 2006, obtainable from <http://www.tiobe.com/content/paperinfo/DefinitionOfConfidenceFactor.html>.

[13] McCabe, Thomas J., "*A Complexity Measure*", IEEE Transactions on Software Engineering Volume SE-2, Issue 4, December 1976, pp. 308–320.

Appendix B Reviewers

The following persons have reviewed this document.

Name	Company
Johan van Beers	Philips Healthcare
Rob Goud	TIOBE Software
Rene van Hees	Thales
Nicolas de Jong	TIOBE Software
Marco Louwerse	Océ Technologies
Randy Marques	Randy Marques Consultancy
Goce Naumoski	ASML
Jan van Nunen	TIOBE Software
Ben Ootjers	Unisys
Dennie Reniers	TIOBE Software
Ben van Rens	Océ Technologies
Bram Stappers	TIOBE Software
Walfried Veldman	ASML