



# Crafting Quality Software

## Gems From The Qafoo Blog

Benjamin Eberlei, Kore Nordmann,  
Manuel Pichler & Tobias Schlitt

Published by Qafoo GmbH

Copyright © 2016 Qafoo GmbH  
Published by Qafoo GmbH  
Version 1.1.1; Published December 15, 2016  
<https://qafoo.com/> | <contact@qafoo.com>

# Contents

<b>1</b>	<b>Introduction</b>	
<b>1.1</b>	<b>About This Book</b>	<b>9</b>
<b>1.2</b>	<b>About Qafoo</b>	<b>10</b>
<b>1.3</b>	<b>The Authors</b>	<b>11</b>
1.3.1	Kore Nordmann .....	11
1.3.2	Tobias Schlitt .....	11
1.3.3	Bejamin Eberlei .....	12
1.3.4	Manuel Pichler .....	12
<b>2</b>	<b>Clean Code</b>	
<b>2.1</b>	<b>Developers Life is a Trade-Off</b>	<b>13</b>
2.1.1	The NoSQL Dilemma .....	13
2.1.2	Overengineered State Machines .....	14
2.1.3	Hack Hack Hack .....	15
2.1.4	Bottom Line .....	16
<b>2.2</b>	<b>Never Use null</b>	<b>17</b>
2.2.1	What it is Used For .....	17
2.2.2	null as Return Value .....	19

2.2.3	Summary . . . . .	20
<b>2.3</b>	<b>Struct classes in PHP</b>	<b>21</b>
2.3.1	Implementation . . . . .	22
2.3.2	Copy on write . . . . .	24
2.3.3	Summary . . . . .	25
<b>3</b>	<b>Object Oriented Design</b>	
<b>3.1</b>	<b>Learn OOD - to unlearn it again</b>	<b>27</b>
3.1.1	Learning OOD the classical way . . . . .	27
3.1.2	OOD in fast pace and agile . . . . .	28
3.1.3	Refactoring is the key . . . . .	29
3.1.4	Learning OOD to unlearn it . . . . .	29
3.1.5	Conclusion / TL;DR. . . . .	31
<b>3.2</b>	<b>Object lifecycle control</b>	<b>32</b>
3.2.1	Why is this bad? . . . . .	32
3.2.2	How can I solve this? . . . . .	33
3.2.3	Conclusion . . . . .	34
<b>3.3</b>	<b>Ducks Do Not Type</b>	<b>35</b>
3.3.1	Duck Typing . . . . .	36
3.3.2	Prototyping . . . . .	36
3.3.3	Using Foreign Code . . . . .	37
3.3.4	Package Visibility . . . . .	38
3.3.5	Conclusion . . . . .	38
<b>3.4</b>	<b>Abstract Classes vs. Interfaces</b>	<b>39</b>
3.4.1	Definitions . . . . .	39
3.4.2	Classes are Types. . . . .	39
3.4.3	interface . . . . .	40
3.4.4	Telling Apart . . . . .	41
3.4.5	Examples & Hints . . . . .	42
3.4.6	tl;dr . . . . .	43

<b>3.5</b>	<b>ContainerAware Considered Harmful</b>	<b>44</b>
3.5.1	Background . . . . .	44
3.5.2	Issues . . . . .	44
3.5.3	Conclusion . . . . .	47
<b>3.6</b>	<b>Code Reuse By Inheritance</b>	<b>48</b>
3.6.1	Inheritance . . . . .	48
3.6.2	Active Record . . . . .	48
3.6.3	A Smaller Example . . . . .	49
3.6.4	The Helper Method . . . . .	51
3.6.5	Testing Private Methods . . . . .	51
3.6.6	Depth Of Inheritance Tree (DIT) . . . . .	52
3.6.7	Summary . . . . .	52
<b>3.7</b>	<b>Utilize Dynamic Dispatch</b>	<b>53</b>
<b>3.8</b>	<b>When to Abstract?</b>	<b>58</b>
3.8.1	Summary . . . . .	60
<b>4</b>	<b>Testing</b>	
<b>4.1</b>	<b>Finding the right Test-Mix</b>	<b>61</b>
4.1.1	The Test-Mix Tradeoff . . . . .	63
4.1.2	Conclusion . . . . .	64
<b>4.2</b>	<b>Practical PHPUnit: Testing XML generation</b>	<b>65</b>
4.2.1	Test case basics . . . . .	68
4.2.2	The naive way . . . . .	69
4.2.3	The lazy way . . . . .	71
4.2.4	Selecting with CSS . . . . .	71
4.2.5	Tag match assertions . . . . .	75
4.2.6	Using XPath . . . . .	77
4.2.7	Conclusion . . . . .	81
<b>4.3</b>	<b>Testing file uploads with PHP</b>	<b>82</b>
<b>4.4</b>	<b>Mocking with Phake</b>	<b>88</b>
4.4.1	Test Doubles Explained . . . . .	88
4.4.2	Benefits of Test Doubles . . . . .	88

4.4.3	Introduction to Phake . . . . .	89
4.4.4	Conclusion . . . . .	91
<b>4.5</b>	<b>Testing Effects of Commands With Phake::capture()</b>	<b>92</b>
<b>4.6</b>	<b>Using Mink in PHPUnit</b>	<b>94</b>
<b>4.7</b>	<b>Introduction To Page Objects</b>	<b>97</b>
4.7.1	Groundwork . . . . .	97
4.7.2	A First Test . . . . .	98
4.7.3	Refactoring The Frontend . . . . .	100
4.7.4	Problems With Page Objects . . . . .	102
4.7.5	Conclusion . . . . .	102
<b>4.8</b>	<b>Database Tests With PHPUnit</b>	<b>103</b>
4.8.1	Removing Data versus Schema Reset . . . . .	103
4.8.2	Point of Data Reset . . . . .	104
4.8.3	Mocking the Database Away . . . . .	106
4.8.4	Conclusion . . . . .	106
<b>4.9</b>	<b>Database Fixture Setup in PHPUnit</b>	<b>107</b>
4.9.1	Dump & Insert Live Data . . . . .	107
4.9.2	Base Data . . . . .	108
4.9.3	Test Data . . . . .	108
4.9.4	Conclusion . . . . .	110
<b>4.10</b>	<b>Using Traits With PHPUnit</b>	<b>111</b>
4.10.1	An Example . . . . .	111
4.10.2	Traits . . . . .	112
4.10.3	Whats The Difference? . . . . .	113
4.10.4	Summary . . . . .	114
<b>4.11</b>	<b>Outside-In Testing and the Adapter and Facade Patterns</b>	<b>115</b>
4.11.1	Conclusion . . . . .	120
<b>4.12</b>	<b>Behavior Driven Development</b>	<b>122</b>
4.12.1	Example . . . . .	122
4.12.2	Behat . . . . .	123
4.12.3	Rationale . . . . .	124
4.12.4	Conclusion . . . . .	124

<b>4.13</b>	<b>Code Coverage with Behat</b>	<b>125</b>
4.13.1	Preparation . . . . .	125
4.13.2	Collecting Code Coverage . . . . .	125
4.13.3	Running Tests . . . . .	126
4.13.4	Conclusion . . . . .	127
<b>4.14</b>	<b>Testing Micro Services</b>	<b>128</b>
<b>5</b>	<b>Refactoring</b>	
<b>5.1</b>	<b>How to Refactor Without Breaking Things</b>	<b>133</b>
5.1.1	Tests . . . . .	133
5.1.2	Baby Steps . . . . .	134
<b>5.2</b>	<b>Testing legacy code</b>	<b>137</b>
<b>5.3</b>	<b>Teaching and Learning Domain-Driven Design without Patterns</b>	<b>142</b>
<b>6</b>	<b>Architecture</b>	
<b>6.1</b>	<b>Why Architecture is Important</b>	<b>147</b>
6.1.1	Summary . . . . .	148
<b>6.2</b>	<b>Scaling Constraints of Languages</b>	<b>149</b>
6.2.1	Why PHP? . . . . .	149
6.2.2	So, Why Not PHP? . . . . .	150
6.2.3	Summary . . . . .	152
<b>6.3</b>	<b>How To Synchronize a Database With Elasticsearch?</b>	<b>153</b>
6.3.1	Transactions . . . . .	153
6.3.2	Changes Feed . . . . .	154
6.3.3	Generating Correct Sequence Numbers . . . . .	156
6.3.4	Storing The Sequence Number . . . . .	156
6.3.5	Conclusion . . . . .	157
<b>6.4</b>	<b>Common Bottlenecks in Performance Tests</b>	<b>158</b>
6.4.1	System Test Setup . . . . .	158
6.4.2	Stack Analysis . . . . .	159
6.4.3	It is Not Always The Database . . . . .	159

6.4.4	Summary . . . . .	160
<b>6.5</b>	<b>Embedding REST Entities</b>	<b>161</b>
6.5.1	Entities . . . . .	161
6.5.2	The Use-Case . . . . .	162
6.5.3	Resource Embedding with HAL . . . . .	162
6.5.4	Better Resource Embedding . . . . .	163
6.5.5	Bottom Line . . . . .	164
<b>7</b>	<b>Workflow</b>	
<b>7.1</b>	<b>Coding in Katas</b>	<b>165</b>
<b>7.2</b>	<b>Why you need infrastructure and deployment automation</b>	<b>168</b>
7.2.1	Can you make a build in one step? . . . . .	168
7.2.2	Do you make daily builds? . . . . .	169
7.2.3	Do you use configuration management tools to automate infrastructure? . . . .	170
7.2.4	Is the development setup documented and automated? . . . . .	170
7.2.5	Can you rollout and rollback deployments in one step? . . . . .	171
7.2.6	Can applications be deployed to a new server setup without changes to the code? 171	
7.2.7	Conclusion . . . . .	172

# 1. Introduction

## 1.1 About This Book

This book is a curated collection of blog posts from the Qafoo Team Blog<sup>1</sup>. Over the time we created many blog posts focussing on the topics of Clean Code, Object Oriented Design, Testing, Refactoring and Software Architecture. To make it easier for you to consume those blog posts we re-structured them and collected them in this book.

Many of the topics in this book are covered in our trainings and you can use this book as a guide after those trainings to read more about topics you particularly care about or topics not covered in your training – since we tailor our trainings to the concrete requirements each customer has.

If you are interested in further information on any of the topics in this book get in contact<sup>2</sup> with us or directly request one of our workshops<sup>3</sup> to be held on-site, just for you.

When reading through this book remember that you read blog posts. For every topic you find a link to the original post where you can leave remarks, the author and the original publishing date. We'd love to hear back from you on our blog.

---

<sup>1</sup><https://qafoo.com/blog>

<sup>2</sup><https://qafoo.com/contact>

<sup>3</sup><https://qafoo.com/services/workshops>

## 1.2 About Qafoo

Qafoo combines deep knowledge with passion for software quality. It is our goal to impart this passion to our customers through consulting, coaching and workshops. This enables our customers to create amazing software with sustainable quality.

All experts from Qafoo do not only have a solid training in their field but also extensive experience from practical application of their knowledge. All of us successfully develop open source and commercial software since years.

If you like to get more information about Qafoo we suggest you to visit the following pages:

- Our Customers<sup>4</sup>
- Possible Workshops<sup>5</sup>
- Presentations at conference<sup>6</sup>

You are, of course, always welcome to get in contact with us<sup>7</sup>.

---

<sup>4</sup><https://qafoo.com/clients>

<sup>5</sup><https://qafoo.com/services/workshops/workshops>

<sup>6</sup><https://qafoo.com/resources/presentations>

<sup>7</sup><https://qafoo.com/contact>

## 1.3 The Authors



The team members from left to right:

### 1.3.1 Kore Nordmann

Kore Nordmann has a university degree in computer science and extraordinary broad experience as a software developer in the professional and open-source PHP projects. Based on this unique mix of theoretical and practical knowledge, he supports teams to take the right path in critical software design and architecture decisions. Kore is one of the founders of Qafoo GmbH where he works as an expert consultant and trainer.

### 1.3.2 Tobias Schlitt

Tobias Schlitt has a degree in computer science and works in professional PHP-based web projects since 1999. As an open source enthusiast, he contributed to various community projects. Tobias is co-founder of Qafoo GmbH which helps development teams to produce high-quality web application in terms of expert consulting and individual training. Tobias main focus is on software architecture, object oriented design and automated testing.

### 1.3.3 Benjamin Eberlei

Benjamin works at Qafoo GmbH where he is teaching and consulting on software-quality, refactoring and architecture. Recently he has also founded the Tideways company, which offers performance monitoring, profiling and exception tracking solutions for PHP. As a project leader of Doctrine he is mentoring the next generation of ORM developers.

### 1.3.4 Manuel Pichler

Manuel Pichler is a software architect with extraordinary practical experience basis. As the creator of some of the most widely used PHP quality assurance tools (PHPMD, PHP\_Depend and Ant Build Commons) he has a vivid passion on static code analysis, software metrics and other quality assurance related technologies. Manuel works as an expert consultant and trainer at Qafoo GmbH, of which he is also a co-founder.

## 2. Clean Code

### 2.1 Developers Life is a Trade-Off

[https://qafoo.com/blog/075\\_developers\\_life\\_trade\\_off.html](https://qafoo.com/blog/075_developers_life_trade_off.html)

*Kore Nordmann at 27. May, 2015*

At Qafoo, we train a lot of people on topics like object oriented software design, automated testing and more. It happens quite often that an attendee asks questions like "Which is the best solution for problem class \$x?", "What is the optimal tool for task \$y" or "There is a new technology \$z, is that the future of web development?". Some are disappointed when I reply "It depends" or "That does not exist", but that's the truth.

There is no silver bullet and one of the most important skills every developer needs to hone is to assess possibilities and to find the best trade-off for the current challenge.

To make that point clear I'm giving three examples from my personal experience, some where it went well and some where it did not.

#### 2.1.1 The NoSQL Dilemma

Choosing a storage system is probably the most common decision to be made in any web project. For one project we already knew that there will be a large amount of data, so we started RnD on recent storage solutions and NoSQL was the big topic then. Cassandra appeared to suite our needs perfectly from its description.

So we gave it a try and prototyped against it, using a thin class layer that hid the database detail.

It was a wise decision to create that layer whereas going with Cassandra from the start was not a very good one. While the database met our expectations, it crashed multiple times with loosing data in our development VMs and none of us had a clue, why, or more importantly, how to fix it. So we revised that decision and eventually went with MySQL and a denormalized schema, where only important attributes became dedicated table fields and the main data structure was serialized JSON in blob. That worked fine for several years.

What did we do? We found the solution which promised to fit our requirements quite well. Luckily, we did not jump blindly on it, but kept it a bit away from the system and evaluated further. Constraints like "knowledge on maintenance" were eventually considered and so we made a trade-off between these requirements that pushed us in contrary directions.

### 2.1.2 Overengineered State Machines

You don't need to jump on such a high-level bandwagon to find examples on where you need to make trade-offs. In one project I created a component to match our business model against a 3rd party one to use their data in our system. While the models were similar, a fundamental difference was how it was determined when data was published and deleted.

That problem appeared perfect for applying a state machine to it to trigger changes on our side when the 3rd party model changed. So I went the extra mile by creating an abstract state machine with `Node` and `Transition` classes. A transition was triggered through an abstract `Condition` and a node fired an `Event` which eventually triggered that change in our model. It was beautiful code in my eyes, with really small classes (1-2 lines executable code). To "ease" configuration and since I expected frequent changes, I made it configurable in XML.

Some time later, a co-worker needed to check if the implemented logic still met new requirements. Digging through that, starting from the XML configuration, running through so many classes to get the complete picture of what it did, took him a large amount of time. It turned out that the logic was fine, but the amount of time spent re-occured some more times after that experience.

A year later, we needed to finally adjust the logic. But the model did not fit the new requirement well, so that we needed to implement many additional classes and the XML configuration became even more complex.

What went wrong? I did not evaluate all available options and pick a good trade-off for the situation. As a minimal solution, I could have hidden the state processing behind a class and code it straight with nested `if` conditions. Maybe that would have been 40 lines of code, but a good bunch of unit tests could have covered that. Maybe a solution in between, abstracting only the states and hard-coding the trigger and transition logic, would have been the optimal trade-off.

So, what's the moral of the story? Whenever you take a software design decision in your project, there is a ton of possible solutions. Just picking a random, interesting, clean, ... one is most probably not the right choice. Instead, you need to check your actual constraints and then find the best trade-off between the possibilities. Which one that is can vary greatly.

### 2.1.3 Hack Hack Hack

There are times in a project, where the option of just hacking a problem solution into the core of the system looks viable. Be warned explicitly: only do that with great care and foresight. In many environments that can lead to stacked up technical debt and possibly result in a really hard to maintain project.

In that project we had an import workflow which performed time-costly analysis on the imported data. It worked well as the incoming chunks of data were small. But then a party registered which provided huge chunks of data and the system became unresponsive. Not only that their imports were delayed, but it affected all other involved parties, too. On the other hand, that big-chunk-party did not even need all the analysis we performed on their data.

A clean solution would have been to create a configuration flag for the kinds of analysis. Another one would have introduced sharding for the import process, that was already in the backlog with low priority. Both solutions would have required too much time, we needed to fix the actual problem immediately. So we hacked a condition into a prominent place in the core, skipping the expensive analysis for the big-chunk-party.

I had an eye on that code quite for a long time, being prepared to refactor it

whenever viable. But there was no need. The code worked fine, nobody needed to touch it again and the problem was solved. Eventually, the expensive analysis was not necessary at all anymore, so we removed it and with it the ugly hack.

What did we actually do? We made a trade-off. That one was very much in one of the possible directions, but because we knew of the potential problems and were prepared to revise the decision, we could take the risk.

#### 2.1.4 Bottom Line

One of the most important tasks of a developer is to make trade-offs. They occur wherever you look in your every day life. It is a highly important step to realize and accept this. And it is important to hone that skill. You need to open your mind for new technology and techniques, learn and try them wherever you can. But then you need to step back, analyze the current situation and then find the best trade-off between all possible approaches. In addition, you need to be able to reflect your decisions and be prepared to revise them. The plus side: by doing so, you will surely learn something for all the upcoming trade-offs.

## 2.2 Never Use null

[https://qafoo.com/blog/083\\_never\\_use\\_null.html](https://qafoo.com/blog/083_never_use_null.html)

*Kore Nordmann at 3. May, 2016*

When doing code reviews together with our customers we see a pattern regularly which I consider problematic in multiple regards – the usage of `null` as a valid property or return value. We can do better than this.

Let's go into common use cases first and then discuss how we can improve the code to make it more resilient against errors and make it simpler to use. Most issues highlighted are especially problematic when others are using your source code. As long as you are the only user (which hopefully is not the case) those patterns might be fine.

### 2.2.1 What it is Used For

One common use case is setter injection for optional dependencies, like:

```
class SomeLoggingService {  
    private $logger = null;  
  
    public function setLogger(Logger $logger) {  
5         $this->logger = $logger;  
    }  
}
```

The logger will be set in most cases but somebody will forget this when using your service. Now a second developer enters the scene and writes a new method in this class using the property `$logger`. The property is always set in the use cases they test during development, thus they forget to check for `null` – obviously this will be problematic in other circumstances. You rely on methods called in a certain order which is really hard to document. An internal `getLogger()` method constructing a default `null` logger might solve this problem, but it still might not be used because the second developer wasn't aware of this method and just used the property.

In PHP versions < 7 a call to `$this->logger->notice(...)` will result in a `Fatal Error` which is particularly bad since the application can't handle this kind of errors in a sane way. In PHP 7 those errors are finally catchable but still nothing you'd expect in this situation.

What is even worse is *debugging* this kind of initialization. This is often even used together with aggregated objects which are *required* by the aggregating class. (You should not use setter injection for mandatory aggregates, but it is still used this way.) Let's consider the following code now:

```
class SomeService {  
    public function someMethod() {  
        $this->mandatoryAggregate->someOtherMethod(/* ... */);  
    }  
5 }
```

When calling `someMethod()` and the property `$mandatoryAggregate` is not initialized we get a fatal error, as mentioned. Even if we get a backtrace through `XDebug`<sup>1</sup> or change the code to throw an exception and get a backtrace it is still **really hard** to understand why this property is not initialized since the construction of `SomeService` usually happens outside of the current callstack but inside the Dependency Injection Container or during application initialization.

The debugging developer is now left with finding all occurrences where `SomeService` is constructed, check if the `$mandatoryAggregate` is properly initialized and fix it, if not.

### The solution

All mandatory aggregates **must always** be initialized during construction. If you want a slim constructor consider a pattern like the following:

```
class SomeService {  
    public function __construct(Aggregate $aggregate, Logger $logger = null) {  
        $this->aggregate = $aggregate;  
        $this->logger = $logger ? : new Logger\NullLogger();  
    }  
5 }
```

The parameter `$aggregate` now is really mandatory, while the `logger` is optional – but it will still always be initialized. The `Logger\NullLogger` now can be `logger` which just throws all log messages away. This way there is no need to care about checking the logger every time you want to use it.

Use a so called null object if you need a default instance which does nothing. Other examples for this could be a null-mailer (not sending mails) or a null-cache (does not cache). Those null objects are usually really trivial to implement. Even it

---

<sup>1</sup><https://xdebug.org/>

costs time to implement those you'll save a lot time in the long run because you will not run in Fatal Errors and have to debug them.

### 2.2.2 null as Return Value

A similar situation is the usage of `null` as a return value for methods which are documented to return something else. It is still commonly used in error conditions instead of throwing an exception.

It is, again, a lot harder to debug if this occurs in a software you use but you are not entirely familiar with. The `null` return might pass through multiple call layers until it reaches your code which makes debugging that kind of code a journey through layers of foreign and undiscovered code – sometimes this can be fun but almost never what you want to do when in a hurry:

```
class BrokenInnerClass {
    public function innerMethod() {
        // ...
        if ($error) {
            return null;
        }
        // ...
    }
}
5
10
class DispatchingClass {
    public function dispatchingMethod() {
        return $this->brokenInnerClass->innerMethod();
    }
}
15
class MyUsingClass {
    public function myBeautifulMethod() {
        $value = $this->dispatchingClass->dispatchingMethod();
        $value->getSomeProperty(); // Fatal Error
    }
}
20
```

Usually there are even more levels of indirection, of course. We live in the age of frameworks after all.

**The solution**

If a value could not be found do not return `null` but throw an exception – there are even built in exceptions for such cases like the `OutOfBoundsException`, for example.

In the callstack I can see immediately where something fails. In the optimal case the exception message even adds meaning and gives some hints of what I have to fix.

**2.2.3 Summary**

Using `null` can be valid inside of value objects and sometimes you just want to show nothing is there. In most cases `null` should be either replaced by throwing an exception or providing a null object which fulfills the API but does nothing. Those null objects are trivial and fast to develop. The return on investment will be **huge** due to saved debugging hours.

## 2.3 Struct classes in PHP

[https://qafoo.com/blog/016\\_struct\\_classes\\_in\\_php.html](https://qafoo.com/blog/016_struct_classes_in_php.html)

*Kore Nordmann at 24. January, 2011*

PHP arrays are a wonderful tool and one of the reasons I like PHP. Their versatility makes it possible to easily set up proof of concepts (POC), either used as hash maps storing multiple keys, or as lists, stacks, trees or whatever you like.

But once you are past the phase of the initial POC, the excessive usage of arrays and exactly their versatility has some drawbacks: If you see an `array` type hint or return documentation, you know nearly nothing about the data structure. Using arrays as key-value hash maps for storing configuration keys or data sets you also know nearly nothing about the expected contents of the array.

This is no problem during the initial implementation, but can become a problem during maintenance - it might not be trivial to find out what the array contains or is supposed to contain (without dumping it). There are no common ways to document such array structures nor you get auto-completion from common IDEs. If such a hash map is filled with data in different locations in your application it even gets worse. Also, mistyping a key - whether on read or write - creates a serious debugging hell.

In Apache Zeta Components and in several of my own projects we are using - so called - struct classes to solve this issue: The struct classes do not define any methods but just contain documented properties. They just deal as a data container, similar to a hash map.

There are several benefits and one drawback using this approach. The benefits:

- Struct classes are far easier to document
- Your IDE can provide you with correct auto-completion
- Your IDE even knows the type of each child in a struct allowing you to create and process deeply nested structures correctly
- You can be sure which properties a passed struct has - no need to check the availability of each property on access
- Structs can throw exceptions access to non-existent properties

The drawback:

- The structs are objects, which means they are passed by reference. This can be an issue if you are operating on those structs. I will show an example later.

### 2.3.1 Implementation

To see what I am talking about let's take a look at a example base class for structs:

```
<?php
abstract class Struct
{
5   public function __get( $property )
    {
        throw new RuntimeException( 'Trying to get non-existing property ' .
            $property );
    }
10  public function __set( $property , $value )
    {
        throw new RuntimeException( 'Trying to set non-existing property ' .
            $property );
    }
}
```

In a struct base class you can implement `__get()` and `__set()` so they throw an exception if an unknown property is accessed. For me PHPs behavior of silently creating public properties on property write access caused quite some irritations over time. A typo in a property name and your code does strange things. I like to get a warning or (even better) an exception for that. Now, let's take a look at a concrete struct:

```
<?php
class LocationStruct extends Struct
{
5   /**
    * @var string
    */
    public $city;
10  /**
    * @var string
    */
    public $country;
15  public function __construct( $city = null , $country = null )
    {
        $this->city    = $city;
        $this->country = $country;
    }
}
```

```

20 }
    }

```

The `LocationStruct` has two documented, public properties. Each one, of course, could be a struct again. If the `LocationStruct` is used as a type hint somewhere in your application or library you now know exactly what data is expected and can create it comfortably, supported by your favorite IDE. The definition of a constructor is really helpful to easily create new struct instances.

### Extending the base struct

There are some sensible extension you probably want to use for the base struct: As mentioned before the structs are passed by reference, which is not always what you want. You therefore probably want to implement `__clone()` in a sensible way, generically for all your structs:

```

<?php
abstract class Struct
{
    // ...

    public function __clone()
    {
        foreach ( $this as $property => $value )
        {
            if ( is_object( $value ) )
            {
                $this->$property = clone $value;
            }
        }
    }
}

```

Another functionality you might want to implement, and a good use case of late static binding<sup>2</sup> (LSB) in PHP 5.3, is the `__set_state()` method, so you can export your struct using `var_export()` just like arrays:

```

<?php
abstract class Struct
{

```

<sup>2</sup>[https://qa.fo/book-language\\_oop5\\_late-static-bindings](https://qa.fo/book-language_oop5_late-static-bindings)

```

5 // ...
  public static function __set_state( array $properties )
  {
    $struct = new static ();
10     foreach ( $properties as $property => $value )
    {
        $this->$property = $value;
    }
15     return $struct;
  }
}

```

If you are using `__set_state()` to ex- and import structs in your application, this is a good reason to define sensible default values for all constructor arguments.

### 2.3.2 Copy on write

As mentioned before, one problem with this usage of struct classes is that they are always passed by reference. It is not entirely obvious why this would be a problem, but it already caught me some times, so here is a example.

In the Graph component from the Apache Zeta Components<sup>3</sup> we, for example, use a struct class to represent coordinates (`ezcGraphCoordinate`). Obviously there are quite some calculations to perform when rendering (beautiful) charts.

Now imagine you want to draw a set of circles at increasing offsets:

```

$offset = new ezcGraphCoordinate( 42, 23 );
for ( $i = 0; $i < $shapeCount; ++$i )
{
5   $driver->drawCircle( $coordinate, 10 );

   $offset->x += 15;
}

```

The `drawCircle()` method now might perform additional calculation on the passed coordinate, for example, because the currently used driver does not use the center point, but the top left edge of the circle as a drawing offset. In this case the method might internally modify the coordinate and thus the offset in the shown loop would

<sup>3</sup><http://zetac.org/Graph>

also be modified. Hopefully you got tests for this in place and therefor add a `$offset = clone $offset` in the `drawCircle()` method. This hit me **very** seldomly until now, but it might be an issue you should be aware of when using struct classes.

### 2.3.3 Summary

Even requiring slightly more work when writing software, the benefit of struct classes during the maintenance phase of projects makes them a true winner - in my personal opinion.

For POCs I tend to still use arrays for structs, but once the software reaches production quality I tend to convert array structs into struct classes since some time in the software I write / maintain.

In C#, for example, such struct classes are a language element and differ from common object exactly in the copy-on-write vs. pass-by-reference behaviour mentioned in this post. I would love to see that in PHP but my knowledge of the Zend Engine is limited and maybe I should bribe a more experienced PHP internals developer...

#### Final note

There are other ways to implement struct classes, like using a properties array instead of public properties, which enable you to perform type checks on property write access. Those might be discussed in another blog post but would exceed the purpose of this blog post.



## 3. Object Oriented Design

### 3.1 Learn OOD - to unlearn it again

[https://qafoo.com/blog/064\\_learn\\_ood\\_to\\_unlearn\\_it.html](https://qafoo.com/blog/064_learn_ood_to_unlearn_it.html)

Tobias Schlitt at 11. February, 2014

One topic we regularly teach in workshops for our customers is *object oriented design* (OOD), i.e. the art of crafting classes/interfaces in a way that the result is an easy-to-understand, maintainable and flexible code base. With the agile focus on shipping working software, some might consider the skill of OOD less important. One popular argument is that quick reaction to change is more important than designing objects and their interaction carefully. I personally look at it the exact other way around. This blog post summarizes why you need to learn OOD first, in order to avoid it to some degree again.

#### 3.1.1 Learning OOD the classical way

Back in the days where the waterfall project methodology was most wide spread, object oriented design was a primary skill for software developers: After analyzing the tremendous functional specification, lots of class diagrams were created, itemizing each and every object in order to provide a full way navigation for the actual implementors.

Developers needed to understand interaction between objects in perfect detail, modeling business processes to their long tail and infrastructure alike. Applying a

design pattern wherever possible in a UML diagram was considered the holy grail in many software development projects back then.

Junior developers those days were left to do the actual implementation of the upfront design, occasionally leaving a small component unspecified to allow them to try out the skills they learned from watching the seniors' design.

One of the results of this way to handle projects was a massive risk to over-engineer the system, taking away the flexibility for spontaneous change requests and leading to a long time to market. These are two of the reasons that led to the agile movement, from a technical perspective - leaving the even more important social aspects aside here.

### 3.1.2 OOD in fast pace and agile

With the rise of agile, the way we handle projects changed drastically: Instead of creating a full system design upfront, teams start with a small set of features, rolling these out to the user as fast as possible and iteratively refining the software together with the customer. Shipping working software is one of the values in the agile manifesto, while clean code is not.

While not directly connected, agility and time to market is also where PHP and JavaScript have their strengths as a development platform. Both languages do not enforce a compelling clean code structure, but leave you the option to implement a hack in order build prototypes quickly. This is both their biggest strength and weakness at the same time.

As can be seen in many projects, working at the fast feature pace often led to *big ball of mud* (BBOM) code bases, which reduced the agility of the project day by day, eventually leading to the stage where changes are costly and developers demanded yet another re-write to free themselves of legacy hell. This is what can be described as *under-engineering*.

Luckily, the area of automatic testing and clean code rose alike with the demand for agility, restraining us from just hacking away, in the best case. And especially unit testing requires you to slice your code carefully into fine-grained objects, which in turn requires a deep understanding of OOD.

### 3.1.3 Refactoring is the key

The solution to the mismatch of fast change and clean code is the art of refactoring. Instead of creating a full blown object design upfront, we nowadays let the code flow for some time until we realize that there is a need for cleanup. We then hopefully take the time to rip the created code apart carefully and to restructure it into a nicely crafted object structure, without moving into the direction of over-engineering.

But being able to refactor requires again a deep knowledge of how structures can be crafted in a clean and maintainable way, how objects should interact to stay flexible. In addition, a particularly good test coverage is needed to enable that process of changing the underlying structure of a system without breaking its business processes (refactoring).

### 3.1.4 Learning OOD to unlearn it

Today's software projects require us to react flexible to emerging changes. This can be achieved by letting the object design of a system emerge instead of attempting to craft it upfront. Automatic testing and continuous refactoring are the key methodologies to master this balance act. Both topics require a deep knowledge about how objects relate to each other, about how to craft them in a sensible way and how to slice the code to extract dependencies. But this is essentially what we call *object oriented design*.

This comprehension leads to what we experienced with many development teams during the past years: in order to enable a team to follow the described development path, knowledge about clean code and various topics object oriented design (like SOLID and design patterns) is an essential basis. Otherwise, the developers will not be able to test responsibly and to perform refactorings successfully.

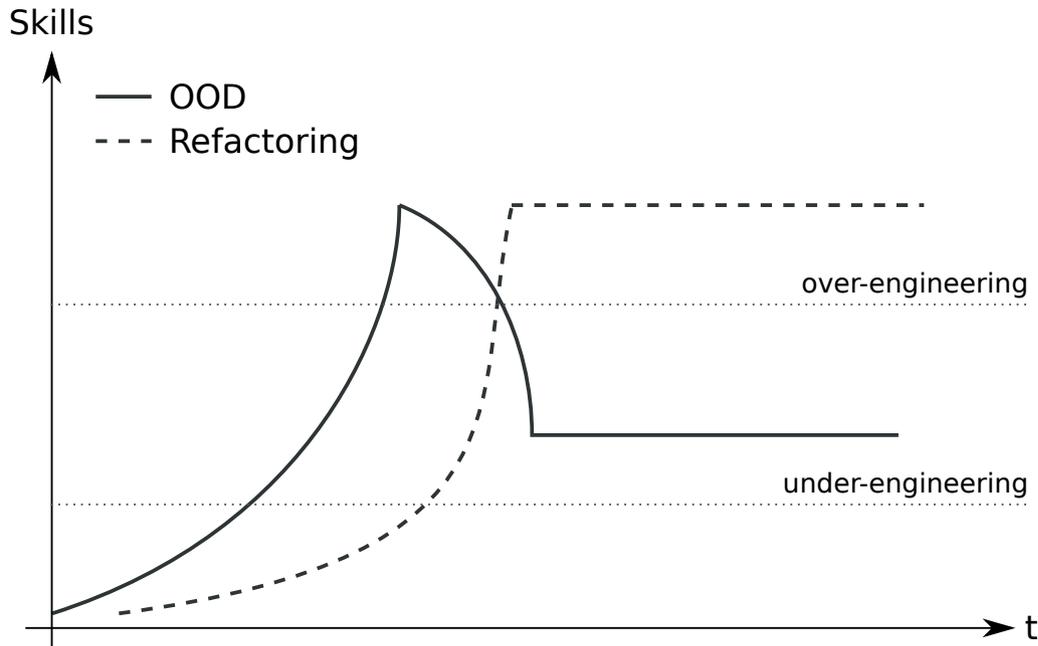
Honing OOD skills requires a pretty good deal of training and practice. Experimenting with object relations, trying to get it right from scratch and iterating learning cycles are absolutely necessary. But once a developer moves in this direction, evidence shows that he automatically begins to over-engineer his code. It is natural and important to undergo this phase of learning, because it is the only way to build up experience and finally acquire a intuitive understanding of how objects relate to each other.

Once this level of perception is reached, the developer needs to restrain himself

again. He needs to *unlearn* parts of what he learned before, in order to let the code flow in a controlled manner, without planning each and every detail and without creating a large object oriented design upfront.

"Unlearn" is of course not the ideal term here, because forgetting the knowledge would be a disaster. Furthermore, refactoring the code requires exactly that know-how and experience. Assessing which code can safely be let slide until refactoring is required and encapsulating that code behind a thin and safe abstraction is essential in order to not end up in a BBOM. Especially for being able to apply a sensible test mix (which is another complex topic in itself) the latter point is highly important.

So there is no way around that learning process, which is illustrated in the following graphic:



Of course, undergoing this process takes time, and failures are required to learn properly. Therefore, junior developers need proper guidance from good seniors - on technical, social and mentoring levels - to go through that process on the job.

Pair programming, coding dojos, code retreats and code review can support this process and speed it up to some degree.

### 3.1.5 Conclusion / TL;DR

In order to master agility, fast feature pace and high quality code at the same time, every developer needs to learn the art of refactoring and automated testing. This in turn requires a high skill level in object oriented design, which can only be reached through a good deal of training and practice. Undergoing a phase of over-engineering during this learning process is natural and even desired to enable the developer to deliver his full potential.

Qafoo experts can help you with a sound training concept<sup>1</sup> to guide your developers safely on their individual way, giving them a well designed mixture of theoretical knowledge, participation and practice. We can furthermore provide sustainable mentoring for your team, providing them with expert advice and constructive feedback, enabling them to take the shortest possible path to becoming a professional.

---

<sup>1</sup><https://qafoo.com/services/training.html>

## 3.2 Object lifecycle control

[https://qafoo.com/blog/020\\_object\\_lifecycle\\_control.html](https://qafoo.com/blog/020_object_lifecycle_control.html)

*Kore Nordmann at 5. April, 2011*

From time to time I trip over APIs, which do not allow me to control the object lifecycle of the used objects myself. This is a bad thing, because it breaks with several concepts of object oriented programming, and forces you back into the dark ages of class oriented programming<sup>2</sup>.

The problem I am talking about is that the API expects a class name instead of an instance (object). The PHP stream wrapper API, for example, let you register a class for a certain schema:

```
stream_wrapper_register( 'virtual', 'myVirtualFileSystemHandler' );
```

In the example above a new object of the class `myVirtualFileSystemHandler` would be created for each file system operation, by the internal stream wrapper implementation.

Such class based plugin mechanisms are not uncommon, but do have several drawbacks I want to outline here - and also provide you with solutions to this very problem.

### 3.2.1 Why is this bad?

The main problem with all class-name based APIs is, that it is not possible to inject further dependencies into the objects resulting from the class name dependency.

Let's examine this in further detail: When we only pass a class name for dependency injection, there are two ways the component in question is using our class:

1. It only calls static methods (this is bad<sup>3</sup>).
2. It creates the object internally, itself.

If the object is created internally, the user of the component is not able to inject additional dependencies into the object just created.

Depending on the API we are interfacing with, the object to be created could for example need a database connection or something alike. The user of the API now only has one choice: Introduce a global state, so that the class can access

<sup>2</sup>[https://qa.foo/book-0103\\_static\\_considered\\_harmful](https://qa.foo/book-0103_static_considered_harmful)

<sup>3</sup>[https://qa.foo/book-0103\\_static\\_considered\\_harmful](https://qa.foo/book-0103_static_considered_harmful)

this global state to fetch a database connection. Since the object is not provided explicitly with the dependency, there is no way but fetching it from some globally known place.

The global state can either be a static property, singleton, registry or something alike - all are global states, and introduce various problems, which are discussed elsewhere<sup>4</sup>.

To summarize: Class based APIs force the user to create a global state in his / her application, if he / she wants to do something non-trivial with your API. It breaks dependency inversion.

### 3.2.2 How can I solve this?

In most cases it should be sufficient to let the user pass an object instead of a class name. If you call methods on that object which require data generated by your code, you can pass this data as an argument to the called method.

Here we get to something different - a sometimes difficult decision: When should I pass something as an argument to a method, and when should I pass dependencies to the constructor of an object? A good rule of thumb is: If an object requires something to exist and cannot live without it, then use the constructor. For everything else use method arguments.

This should solve most issues, but for example the PHP stream API does not just use one object, but wants to create new objects for all kind of use cases. Everytime you access a stream URL somehow, a new object needs to be created for that path. (The constructor is only called sometimes, btw. For example a `stat()` call on URL does not trigger the constructor, but you get a new object.)

There we have the situation that a parameter is vital to the object (the URL). The stream wrapper API needs to construct multiple objects and cannot cope with a single instance we can inject upfront. Therefore you currently specify a class name and the stream-wrapper implementation creates needed instances of this class for you. If you now want to maintain connections (HTTP, FTP) or maybe emulate a file system in memory (PHP variables), you need to introduce a global state, because there is no way to inject any state into those internally constructed objects.

---

<sup>4</sup>[https://qa.fo/book-0103\\_static\\_considered\\_harmful](https://qa.fo/book-0103_static_considered_harmful)

The better way for the stream-wrapper-API would be to require an **instance** of a factory to be registered for some URL schema. The API can then request a new object for a given path from the factory. While the factory needs to follow a certain interface, the implementation is left entirely to the user, who can decide to pass additional other dependencies to the created object - like a connection handler, or a data storage.

The user can even decide to re-use objects, if this is really desired. A typical use for this would be some kind of *identity management* inside of the factory: Reusing the same object for the same path. (Note that this might not make sense in case of the stream wrapper.)

### 3.2.3 Conclusion

From this elaboration you can learn one very important rule for your API design: *Do not create class based APIs, since they force everybody using your API to create a global state, sooner or later.* Allowing to pass an object or a factory keeps the user of your API in control of the lifecycle of his objects. He / she can act far more flexible then.

### 3.3 Ducks Do Not Type

[https://qafoo.com/blog/050\\_ducks\\_do\\_not\\_type.html](https://qafoo.com/blog/050_ducks_do_not_type.html)

Kore Nordmann at 11. July, 2013

Even in ecosystems which generally follow a high standard of code quality I tend to find public methods in classes which do not originate from an interface or an abstract class. I think that this is a really bad habit for several reasons I will explain below. To give this some context, let's start with a simple example:

```
abstract class HttpClient
{
    public function request($method, $path, $body = null, array $headers = array()
    );
}
5
class MyHttpClient extends HttpClient
{
    public function request($method, $path, $body = null, array $headers = array()
    )
    {
        // ...
    }
    public function setConfiguration($key, $value)
    {
        // ...
    }
}
10
15
```

We have an abstract base type for HTTP clients (which could also be an interface (See: *Abstract Classes vs. Interfaces*)) and one implementation. Implementations then sometimes tend to define additional public methods, like `setConfiguration()` in this case, for various different reasons. I want to explain why **I consider this a code smell**.

Semantically, the main point behind abstract classes and interfaces is to build an abstraction other components can implement against. In practice, abstract classes are also often used for code-re-use, but this is also a bad practice and might be a topic of a different blog post.

To stay with the example above, any other component in the application, which wants to use an HTTP client, defines a dependency on the abstract type `HttpClient`. The actual implementation used would then depend on the application configura-

tion. This is the basic idea of *Dependency Inversion* (S.O.L.I.D.), or to say it with the words of Martin Fowler:

- A. High-level modules should not depend on low level modules. Both should depend on abstractions.
- B. Abstractions should not depend upon details. **Details should depend upon abstractions.**

Abstractions can be both: interface definitions or abstract classes.

### 3.3.1 Duck Typing

Then there is *Duck Typing*. Let's start the explanation with the opposite of Duck Typing: Java.

If you pass something to any method in Java, you define a type hint for the value you are expecting. If the type hint is on `HttpClient`, the compiler verifies that you only use methods / properties defined by that interface. You cannot (should not) use the concrete implementation as something else. It does not matter which concrete instance is passed to the method.

When starting to develop with Java I thought this is horribly annoying. And I still think it is – more on that later.

Duck Typing on the other hand means, that you can use anything passed to your method as anything. You usually would check if a method exists and just call it right away. You do not enforce a base type or a formal interface definition.

In the example above this would mean, that anything which can act as a HTTP client would implement a method `request()` and you just go and use those classes as a HTTP client. The underlying problem here is, of course, that entirely different concepts also might define a `request()` method, which then might break a lot.

### 3.3.2 Prototyping

When developing prototypes (not talking about prototype based object orientation, even though it also gets interesting there) or Proof Of Concepts the concept of Duck Typing is extremely important. And this is one of the reasons I usually prefer PHP over Java.

In PHP you can pass stuff around as you like and use the full power of duck typing. Want to see if a couple of components work together well and implement your business requirement? Want to play with a new library? Want to hack a feature seconds before a release date? Just do it.

**BUT:** If you are developing infrastructure components, Open Source libraries or any code other developers depend on: Please design and use your abstraction as if a small Java Compiler sits on your shoulder and pours hot coffee over you anytime you misuse an object.

It is the only way to make code extensible by others. A fact, even advocates of Duck Typing are aware of is that you must know your software really well, if you are using Duck Typing. If synonyms occur in method / property names, you can easily get code behaving in really strange ways. Also you cannot easily identify your ducks, since there is no clearly documented concept behind it. Wikipedia<sup>5</sup> says:

In essence, the problem is that, if it walks like a duck and quacks like a duck, it could be a dragon doing a duck impersonation. One may not always want to let dragons into a pond, even if they can impersonate a duck.

### 3.3.3 Using Foreign Code

The *Open Closed Principle* (S.O.L.I.D.) wants us to change the behavior of our application without modifying code. This requires code to define abstractions which we can replace in the application configuration.

When using type hints on abstractions in your library you make the user of your code assume that she / he can just implement a new class based on this abstract concept and be done. If your code then calls additional methods it will just break. If it breaks immediately one can fix the code, but often enough it happens that the code will only break in obscure error conditions. And since a call to an undefined method even triggers a fatal error this will kill PHP immediately.

And often enough I want to replace such implementations / dependencies in foreign libraries. The most common thing is, that I want your library to use my

<sup>5</sup>[https://en.wikipedia.org/wiki/Duck\\_Typing](https://en.wikipedia.org/wiki/Duck_Typing)

Logger through an adapter, or anything similar. If you then call undocumented methods on that logger (in a special error condition)...

### 3.3.4 Package Visibility

There is one case where public methods are OK, even if they are not defined by a base concept: PHP misses package visibility. It occurs seldom "enough" (to me), but sometimes you want "internal" access to methods on a package level, even if this might not be part of the external API. Until we have package visibility in PHP, from my point of view, it is OK to define additional public methods. But please flag them clearly, to make sure everyone knows it is not part of the public API, like this:

```
class MyHttpClient extends HttpClient
{
    /**
     * @private
     */
    public function setCookie($key, $value)
    {
        // ...
    }
}
```

### 3.3.5 Conclusion

It is fine to employ duck typing. At the end of the day PHP is popular because the language does not impose any limits on the way you want to work with it.

But if you start with class based object orientation, please do it right. Mixing the concepts of class based object orientation and other stuff like Duck Typing messes with everybody's head. It *will* break code in obscure situations which might be really hard to cover by tests.

Abstractions are documentation. The best you can provide for developers. They are not easy to find, but you should really stay with them, once defined. Interfaces can, by the way, be used to extend them over time.

### 3.4 Abstract Classes vs. Interfaces

[https://qafoo.com/blog/026\\_abstract\\_classes\\_vs\\_interfaces.html](https://qafoo.com/blog/026_abstract_classes_vs_interfaces.html)

Kore Nordmann at 2. October, 2012

Features of object oriented languages are often use from a purely technical perspective, without respect to their actual semantics. This is fine as long as it works for you, but might lead to problems in the long run. In this article I discuss the semantical differences between abstract classes and interfaces. I also outline why following the semantics of those language constructs can lead to better code.

*Disclaimer:* This is of course kind of artificial. If using the language constructs in other ways works for you, it's fine. It can make communication, extension of code and maintenance of code harder, though.

#### 3.4.1 Definitions

First, we need to differ between `interface` and `interface`. When typeset in mono space I mean the language construct -- otherwise I talk about the public methods of any object / class. The second one often is also called "public interface", "class interface" or even "class signature".

#### 3.4.2 Classes are Types

A class denotes a **type**.

We know that objects consist of internal state and methods to operate on this state. The interactions between objects are implemented by calling methods on other, related objects. Those methods might, or might not, modify the internal state of the called object.

Classes, as they define the blueprints for those objects, define the type of objects, thus we know how to use an object. When talking about types there are some natural implications, which apply especially when extending from another class, which means to create a sub-type. One is the Liskov Substitution Principle<sup>6</sup>:

“Functions that use pointers or references to base classes must be able to use objects of derived classes without knowing it.”

---

<sup>6</sup><http://butunclebob.com/ArticleS.UncleBob.PrinciplesOfOod>

This basically translates to: "Do not mess up the interface". It is required, so you can use a sub-type (derived class) without any special knowledge, just by reading the documentation / interface of the parent class. Imagine a `Cache` class, from which you extend a `CacheFilesystem`, `CacheMemcached` and maybe a `CacheNull`. You do not want to implement any kind of special handling in every place where an object of any of those derived classes is used in your application. All of them should just behave like a `Cache`. Keep this in mind.

If we can use any sub-type of a type just as the parent type defines, we talk about **Subtype Polymorphism**. Something you definitely want to achieve. Abstract classes are naturally used when you need such a parent type, which requires some specialization, like the `Cache` class mentioned above. I'll get back to why it makes no sense to extract an `interface` from this. And yes, you are allowed to define abstract classes, even if you don't intend to implement any code already and don't have any properties to define.

### 3.4.3 interface

An `interface` describes an aspect of the interface of a type.

PHP does not support multiple inheritance, but it allows you to implement any number of interfaces, right? The reason for this is, that an `interface` annotates usage aspects on the interface of a type. The Interface Segregation Principle<sup>7</sup> also claims:

"Clients should not be forced to depend upon interfaces that they do not use."

This means that a type using another type should only depend on an interface, which is actually used. Depending on a god class, for example, is the exact opposite. You would not know which methods are actually called from looking at the injection signature. But this also implies that the interfaces should be designed small and task-specific.

To stay in the context of the example mentioned above, a proper interface would be `Cacheable` to annotate on certain classes that its instances can be cached. The natural counterpart of a `Cache`, if you want to handle objects. This is a minimal usage

<sup>7</sup><http://butunclebob.com/ArticleS.UncleBob.PrinciplesOfOod>

aspect of a class and other classes requiring this `interface` will most likely, and may only, use the methods required by this `interface`. Implementors of this interface would also not be forced to implement meaningless methods.

#### 3.4.4 Telling Apart

Those definitions are nice, and fluffy -- but what happens if you define an `interface` instead of using an abstract class? (This is at least what I see most.)

- You invite others to create god classes

Interfaces do not allow to already provide code or define properties. Their only "benefit" over abstract classes is that you can implement multiple of them in a class / type.

So why would you define a `Cache` interface. Obviously you do not want that some class implements the `Cache` logic while also being a `Logger` and a `ORM` -- even all might access the same storage.

While this, in theory, provides more "flexibility" for the future, every usage of this flexibility would hurt your software design -- a lot.

- You overload interfaces

A common developer pattern is to define a type, and then extract one single `interface` from the class interface. Such interfaces do not define a single usage aspect any more. If someone wants to implement one aspect of this `interface` and two or three aspects from some other "interfaces" it would result in one *really big* class with dozens of empty methods. Or, even worse, dozens of implemented methods, which are never used. This would obviously be wrong.

But... ====

There is this one popular quote:

“Code against interfaces, not implementations.”

"Interface" is not typeset in mono space. Look it up!

To phrase it better, like it is done in the Dependency Inversion Principle<sup>8</sup>: “**Depend on abstractions, not on concretions.**”. "Interfaces" in the above quote may very well be abstract classes. Don't forget about those.

<sup>8</sup><http://butunclebob.com/ArticleS.UncleBob.PrinciplesOfOod>

### 3.4.5 Examples & Hints

Classic examples for proper interfaces would be:

- Cacheable
- Serializeable

I have two mnemonics for people I discuss Object Oriented Design with:

- interface names should end with `able` OR `ing`.

This is obviously not always true. But ending on `able` is a strong indicator that such an interface just annotates one usage aspect.

- interfaces make sense to be implemented in entirely different types.

An interface usually makes sense to be implemented in types, which otherwise have barely anything in common. A prime example again would be `Cacheable` or `Countable`, which even is a proper interface defined in the SPL.

To come up with a "real-world" example: Let's consider an interface `Drinkable`. You could implement that one on cups, on the sea and maybe even on humans, if there are vampires around. Otherwise humans, seas and cups probably do not have that much in common.

Examples for proper abstract classes could be:

- Log(ger)
- Cache

Again I have a mnemonic to help you with the decision:

- An implemented abstract class can stand on its own.

If you implement all remaining abstract methods of an abstract class, you get a proper member in your world. It does not require any additional methods to be usable by others. It probably will require some dependencies (like a storage), but most other objects will happily just call the methods provided by the abstract class.

Final hint:

An interface **must not** define a constructor. **Never**. The same is true for most abstract classes. By defining a constructor you predefine and limit what dependencies may be injected in implementations. It is very likely to change for different implementations. But this is actually a topic which deserves its very own blog post.

### 3.4.6 tl;dr

An `interface` annotates a usage aspect of a type / class. Do not just extract interfaces from classes. Use abstract classes, when defining a base type.

### 3.5 ContainerAware Considered Harmful

[https://qafoo.com/blog/057\\_containeraware\\_considered\\_harmful.html](https://qafoo.com/blog/057_containeraware_considered_harmful.html)

Tobias Schlitt at 7. October, 2013

A while ago I tweeted<sup>9</sup>

ContainerAware is the new Singleton.

While many people agreed by retweeting and faving. I feel the need to elaborate some more on this statement and save the explanation for the future.

*TL;DR:* No class of your application (except for factories) should know about the Dependency Injection Container (DIC).

#### 3.5.1 Background

The `ContainerAware` interface (actually `ContainerAwareInterface`, `ContainerAware` is a basic implementation of it) is part of the `Symfony2`<sup>10</sup> API<sup>11</sup>, but a similar concept is known from many other frameworks and many applications rely on it. It defines only the one method `setContainer()`, which allows to inject the DIC into an object so that it can directly retrieve services from it.

It is most common to have controllers implement this interface. In fact, the `Symfony2` base controllers<sup>12</sup> do so and there is a base class for shell commands<sup>13</sup> that does it.

#### 3.5.2 Issues

Accessing the DIC in your classes can seriously harm maintainability and code re-usability in the long run.

---

<sup>9</sup><https://twitter.com/tobySen/status/378780141826355200>

<sup>10</sup><http://symfony.com>

<sup>11</sup><https://qa.foo/book-ContainerAware>

<sup>12</sup><https://qa.foo/book-Controller>

<sup>13</sup><https://qa.foo/book-ContainerAwareCommand>

### Reduced Testability

Unit testing is the most common automatic test method in the PHP world. If you have the container injected into your objects, this becomes much harder. There are three strategies to approach testing a class that gets the container injected:

1. Mock the container and make it return the mocked service mocks,
2. Use the container in your test cases and make it return the mocks
3. Mock the subject of test to override the `get()` method that is commonly used to access services inside the class.

The first solution actually requires you to create a container mock that basically does the same thing as the real container. Except for mocking overhead you don't win much with this. So, if you are already in the situation that your classes have a dependency to the DI container, you're better off with version 2.

If you choose the second variant, you're formerly not writing a unit test, but an integration test instead. This is not a problem by default and integration tests are important, especially if you use an external framework for development. However, if you intend to write unit tests, you simply don't in this case, since another class (the DIC) is put under test.

The third variant is a simple no-go. You should never mock methods of your subject, because you cannot ensure that the mocked version of the code mimics what really happens. Your test cases lose a big amount of the safety that they should actually provide you with. Whenever you feel the need to mock a method of the test subject, that is a clear sign for the need to refactor (so-called *code smell*).

### Hidden Dependencies

If you go for a real dependency injection approach (my favorite is constructor injection), you can easily see when a class does too much, just by looking at its dependencies. Think about a constructor like this:

```
<?php
5 class UserProfileController
  {
    // ...

    public function __construct(
        UserRepository $userRepository,
        Search $search,
```

```
10     MailGateway $mailGateway,  
11     Router $router,  
12     HttpClient $httpClient,  
13     PermissionService $permissions  
14 )  
15 {  
16     // ...  
17 }  
18  
19 // ...  
20 }
```

You can easily spot that there is something fishy. The class receives quite many dependencies and these do not fit very well together. Some are from the infrastructure layer, others seem to be business services and even others correlate with the MVC structure. If you see this constructor, it's eminent that there is need for refactoring. You will note that at the latest when you add another dependency to it.

Now compare that with the following code:

```
class UserProfileController extends ContainerAware  
{  
    // ...  
5     public function __construct()  
6     {  
7         // ...  
8     }  
9  
10    // ...  
11 }
```

There is no way to see what the actual dependencies of the class are. In order to see that, you need to go through the code and look for places where the Dependency Injection Container is used to retrieve objects. In real life, nobody will ever do that, especially if this is hidden behind `$this->get('serviceName')` calls.

### Feature Sneak-In

This issue is a direct result of the previous one: If you can, you eventually will. And by this I mean access whatever service you might think you just need.

With access to the DIC a developer has the freedom to just grab any object and use it. This might be convenient to implement hacks, if you don't know where else

to put a certain functionality when being in a rush. But on the other side of the coin, there is nothing which forces you to clean up the mess.

A common result is, that more and more business logic is collected in the controllers, making them become huge transaction scripts, which leads to code duplication and hidden business rules.

### 3.5.3 Conclusion

Making your classes aware of the DIC is not an option. Even if you really feel you need to, just don't do it. To be on the safe side I even recommend to make your controllers services<sup>14</sup>.

If you develop Symfony2 based applications and run into the issue of having to inject too many framework services<sup>15</sup> into your controllers, my co-worker Benjamin has a great post about that in his personal blog.

---

<sup>14</sup><https://qa.fo/book-service>

<sup>15</sup>[https://qa.fo/book-extending\\_symfony2\\_controller\\_utilities](https://qa.fo/book-extending_symfony2_controller_utilities)

## 3.6 Code Reuse By Inheritance

[https://qafoo.com/blog/063\\_code\\_reuse\\_by\\_inheritance.html](https://qafoo.com/blog/063_code_reuse_by_inheritance.html)

Kore Nordmann at 20. January, 2014

### 3.6.1 Inheritance

To me, inheritance has two properties:

- Defining an *is-a* relationship
- Making it possible to share code between classes by extending from a common base class

The *is-a* relationship between classes is one thing I like to use inheritance for. So, to me, a concrete `Cache` - for example a `MemcacheCache` implementation - extending from an abstract `Cache` base class is a very sane thing to do. Implementing an interface means, in my opinion, adding an additional usage aspect / feature to a class, while the base class defines the *type* of a class. Mind: The base class, most likely, only defines abstract methods and does not provide any implementation in this case. I discussed this in more detail in a dedicated blog post<sup>16</sup>, which is why I skip a detailed explanation now.

The other thing you can use inheritance for is to share code between multiple classes. A prime example for this is `Active Record`<sup>17</sup>, where a concrete record class (for example a `BlogPost`) extends from a base class which adds the database access logic.

I personally think that the latter is one of the worst things you can do in object oriented design. Bear with me for a moment and let me try to explain why I think that.

### 3.6.2 Active Record

With `Active Record`, the problem is fairly obvious: The base class usually implements the storage logic, while the extending class is supposed to implement the business logic and maybe additional table query logic.

---

<sup>16</sup>[https://qafoo.com/blog/026\\_abstract\\_classes\\_vs\\_interfaces](https://qafoo.com/blog/026_abstract_classes_vs_interfaces)

<sup>17</sup>[https://en.wikipedia.org/wiki/Active\\_record](https://en.wikipedia.org/wiki/Active_record)

The business logic (or domain logic) is the most important part of your application – the stuff you earn your money with. You probably want to test this logic since you might be broke if it fails too hard.

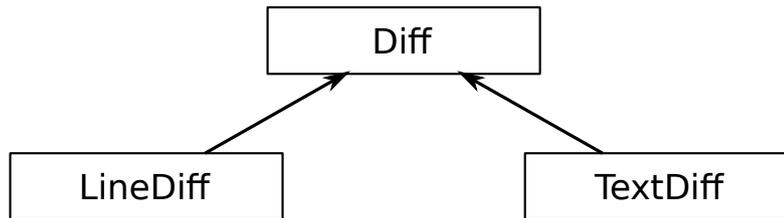
Unit-testing a class which uses logic from a base class, for example accessing the database, is a lot of work. The best way is usually to mock all methods (from the parent class) which access the database and then run the tests on the mocked subject while correctly simulating the return values from your database. This is damn tedious.

Not wanting to go into too much detail here, but testing an Active Record class as-is is often even worse since tests which hit your database are usually damn slow, harder to set up and harder to keep atomic. But most importantly, unit-tests should fail for only one reason, and also testing a database access layer will likely make it a lot harder to locate the exact reason for unit test failures.

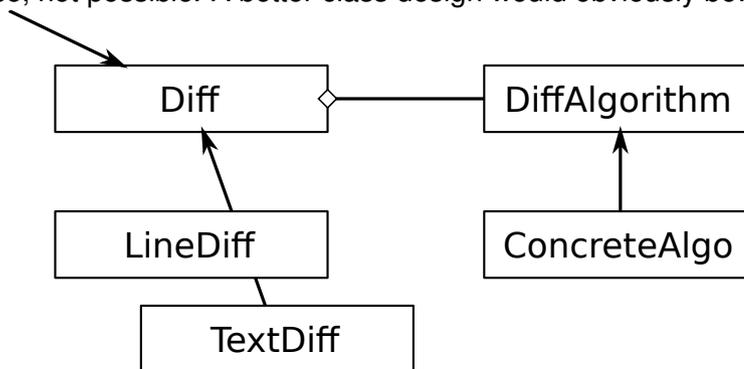
### 3.6.3 A Smaller Example

One of my preferred examples when describing this to others is a mistake in a component design I made not too many years ago. I wanted to write a little *diff* component, which makes it possible to evaluate the line-wise diffs in source code, but also paragraph and word-wise diffs for wiki articles and blog posts.

When writing the component, I started implementing the most common diff algorithm on top of "tokens", which should allow me to handle all the cases mentioned before. Once the algorithm worked, I derived from the class, implementing the different diff flavours:



That worked fine for some time, but once the diff component was integrated into PHPUnit, Sebastian Bergmann noticed problems with large texts and we realized that there are better diff algorithms for particular large texts. The `LineDiff` and `TextDiff` classes, implementing their respective tokenizing rules, are still fine, but we wanted to replace the algorithm in the base class depending on the diff use case. This is, of course, not possible. A better class design would obviously be:



We can now replace the used diff algorithm, test it directly based on simple token streams, and every concrete `Diff` implementation (`LineDiff`, `TextDiff`, ...) will still work with all the diff algorithm implementations you can come up with. A downside,

of course, is that creating the Diff object gets slightly more complex, but this is usually handled by a Dependency Injection Container, anyway.

### 3.6.4 The Helper Method

The points where Code-Reuse-By-Inheritance still keeps creeping into my code from time to time are small helper methods. For example for complex data structures it is so easy to just define a method in the data structure which calculates some kind of hash, to apply simple transformations or to validate some values. The next thing you notice is that you will move those methods up in the extension hierarchy or copy it around (maybe slightly modified).

Why not make it a (single) public method in some Hasher, Validator or Visitor class? You might be surprised how much easier stuff gets to test and that you might even be able to re-use that code in even more places. I guess the *fear of (many) classes* applies here again, which I consider void. But this is another blog post.

Of course it is valid to have such helper methods during prototyping, but if you start using such code in multiple places, start testing it or start to put it into production you should refactor it to follow the mentioned concerns.

In most cases, those helper methods are also a different concern which you start mixing into your class. Even the concern seems closely related, it probably does not hurt to move it somewhere else and make it explicit.

It may happen, though, that people start to over-engineer given these constraints. But every developer walks on a very fine line between missing abstraction and over-engineering all the time, anyway. A good base for a decision could be: Will it ever make sense to somehow re-use this piece of code or may it be possible that someone wants to replace this implementation in a similar use case or during testing?

### 3.6.5 Testing Private Methods

The urge to test private or protected methods is, in my opinion, a code smell which should directly lead to refactoring.

Once you separate your concerns and move code you share by inheritance into its own classes, every non-trivial code will be contained in public methods which are easy to test. And more importantly: easier to understand.

### 3.6.6 Depth Of Inheritance Tree (DIT)

Then there is the "Depth of Inheritance Tree" (DIT) metric with a common boundary value of 5, while the counting even stops at component borders. To me, the maximum value for this metric should be considered 2. Except for some struct classes<sup>18</sup> / value objects, there is, in my opinion, no reason for more than one level of extension of a class. If you use inheritance just for defining the `_type_` of classes, you will never extend more than once. If you are tempted to do that, use aggregation instead and you are probably fine in 99% of all cases.

If you have valid use cases for an inheritance hierarchy greater than two, please share those with me.

### 3.6.7 Summary

To me, by now, Code-Reuse-By-Inheritance is a clear code smell. Every time I am tempted to do this or find this in existing code, I will refactor out the code into a separate class as soon as possible.

---

<sup>18</sup>[https://qafoo.com/blog/016\\_struct\\_classes\\_in\\_php](https://qafoo.com/blog/016_struct_classes_in_php)

### 3.7 Utilize Dynamic Dispatch

[https://qafoo.com/blog/072\\_utilize\\_dynamic\\_dispatch.html](https://qafoo.com/blog/072_utilize_dynamic_dispatch.html)

Tobias Schlitt at 16. October, 2014

A while ago I replied to the tweet<sup>19</sup> by @ramsey<sup>20</sup>

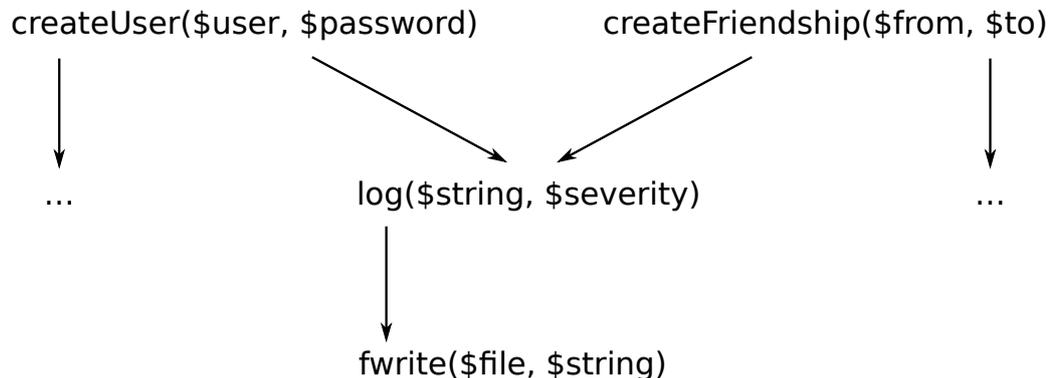
Traits are a nice way to provide common functionality to unrelated classes without using static methods on a global Util class.

with<sup>21</sup>

Which makes them exactly as evil as static access. Funktionalität you dispatch to becomes irreplaceable destroying a fundament of OO: Dynamic dispatch.

I want to use this blog post to illustrate the concept of *dynamic dispatch* which I use a lot recently to motivate creating clean OO structures in my trainings. In my experience, this helps people to understand why we want to write code in this way. After that I will show why traits are bad in this direction.

Dynamic dispatch is actually an implementation detail of object oriented programming languages. To explain it I need to take a little leap back in time and illustrate the code flow in a classic procedural program:



<sup>19</sup><https://twitter.com/ramsey/status/509028110465908737>

<sup>20</sup><https://twitter.com/ramsey>

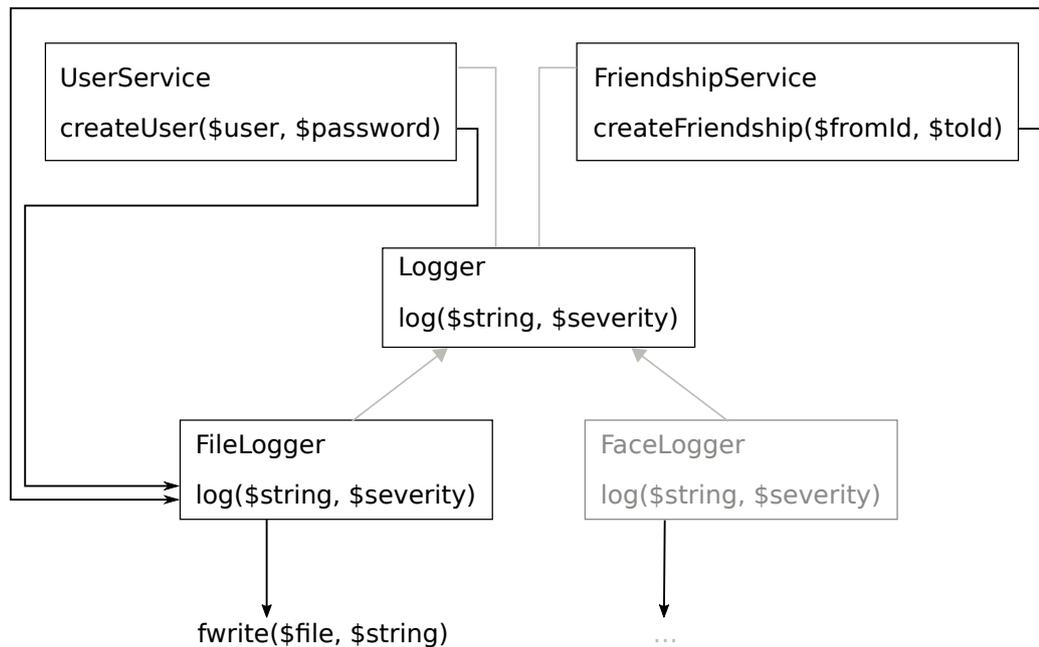
<sup>21</sup><https://twitter.com/tobySen/status/509040311122022400>

The graphic shows two procedures which call a shared `log()` procedure, the arrows visualize the execution flow of the program. As can be seen, a procedural program classically works top down. In order to solve a complex task, a procedure dispatches to other procedures where each fulfills a smaller fraction of the task. It is important to notice that the procedures to be executed are exactly known at compile time of the program (simplified). This is what is called a *static dispatch*.

Imagine you now want to log all friendship related operations to the new *Facelog* server, because it generates incredibly useful insights for you. What are your options to implement this change?

You can a) change the `log()` procedure itself, i.e. patch it. But this is of course no real option here, because `createUser()` would also be affected by this change and most probably many other modules that use `log()`. Option b) is to touch the `createFriendship()` procedure (and any other friendship related modules) and change them to use another procedure instead of `log()`, e.g. `facelog()`. Following this approach is viable, but also means quite some work and potential for errors.

In the object oriented (OO) paradigm the situation is different:



The most obvious difference here is that there are two types of errors: Black arrows visualize the actual code flow again, while gray ones indicate object (actually class) dependencies. The `UserService` depends on `Logger`, `FileLogger` extends `Logger` and so on.

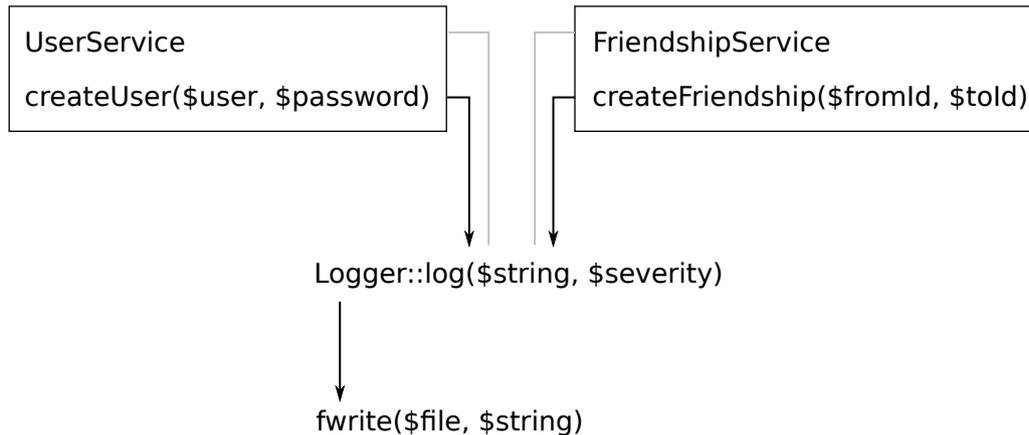
And this exactly is the crucial point: At compile time, an object oriented program typically only knows about the class dependencies. At this stage it is unclear, how the actual execution flow will be at run-time. It depends upon which particular object is given to the dependant, which might be influenced by configuration, user-input and so on. On a technical level, the programming environment decides out of a set of available method implementations which one is the correct one to use. This is what we call *dynamic dispatch*.

So how could you perform the required change in this environment? Of course you can just give a different object to `FriendshipService`, which provides the API defined by the `Logger` type but talks to `FaceLogger` instead of writing to a file. The dynamic dispatch will take care. The nice thing: You neither need to touch the `FriendshipService`

nor do you need to fear undesired side-effects on `UserService` and others by touching the `FileLogger`.

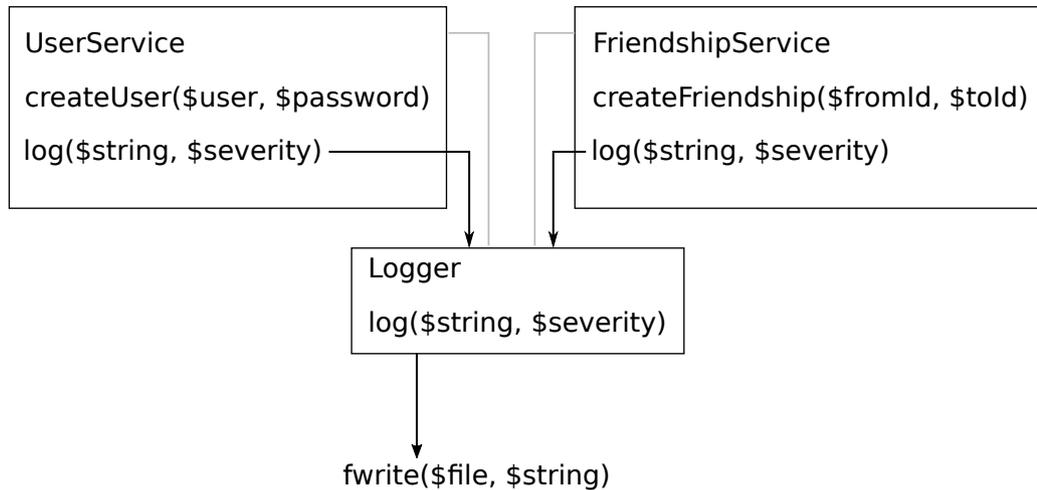
Now, that is dynamic dispatch: Your execution environment decides at run-time, which method is actually to be called, based on the actual object reference given for a type dependency. There is no need to recompile the program to achieve a change in this dispatching, a different user input or configuration can suffice.

Now take a look at a static method call as shown in the following graphic:



What do you realize? Right, a static method call is actually a procedural call. Indeed, the dispatch is static in this case: There is no way for you to replace this call fine grained, except for the two procedural approaches explained further up. Whenever you introduce a static method call, you throw away the powers that dynamic dispatch unveils to you.

Now look at the next graphic that shows a trait implementation for the `Logger`:



Using a trait copies a piece of code into your class, making it a part itself. How would you realize the desired change now? True, there is now way to utilize dynamic dispatch here, but only the procedural approaches work. *A trait is static dispatch, too.*

**Disclaimers:**

1. There is of course much more in the background of this very narrow view on object oriented mechanics, for example subtype polymorphism.
2. I'm aware that there are concepts in procedural languages to implement a dynamic dispatch, too, for example function pointers. However, for illustration purposes, it is helpful to take these out of scope here.

### 3.8 When to Abstract?

[https://qafoo.com/blog/084\\_when\\_to\\_abstract.html](https://qafoo.com/blog/084_when_to_abstract.html)

*Kore Nordmann at 17. May, 2016*

One of the most difficult challenges in a developers life is finding the "right" abstraction, or at least the best one given the current circumstances. The core problem is that abstraction is a bet on the future development of the software and we know that future is volatile. The circumstances will change, so will the view on the best abstraction change.

But there is another dimension which influences this decision: What kind of software are you developing? When it comes to decisions about Object Oriented Design we separate between three different types of projects:

- Internal Project

It is just used by you and your co-workers. You are working in a small team. When you are changing an API you can adapt all code using this API easily – especially if you are using a Monorepo<sup>22</sup> approach.

- Library

There is a common entry point for your library, an interface which everybody is using. This could be everything from a popular Open Source library to a company internal library used in multiple teams or projects. Changing the API is already a no-go because it was develop to fulfill the use-case of the library. You cannot see who is using this API and what code will break. The only chance is a new major release for each change to this API.

- Adaptable Product

There are generic software solutions out there which will be adapted by agencies and developers like Wordpress, Magento or similar products. Users of those products will use **any** API in the source code of those products to change the behaviour of the system. You cannot change any API without breaking someones application.

So how does this influence the decision about abstracting your code? In the first case (Internal Project) you can delay abstracting code pretty far while in the last case (Adaptable Product) you should get it right from the very beginning –

---

<sup>22</sup><https://qa.foo/book-monorepos>

since you shouldn't change anything afterwards.

But, in an Adaptable Product and in Libraries your use case is usually well defined, thus you know what the interface looks and you can think about it, analyze it and come up with a sane abstraction. One could even say: You define the use case with your API, thus your abstraction is fine by definition. Your library and product will be used in other use cases as well, which you did not consider when writing the software, but it usually less of an issue to adapt to that later.

Authors of those libraries and software products are often well known and talk about their code in public. Also this is the code you, as a developer, will see when learning and trying to improve.

But especially in the area of Internal Projects the circumstances change often and fast. You will have no idea what the next requirement of the project stakeholders will be. So, let's be blatant: Do not abstract! Do not use interfaces or abstract classes.

OK, sometimes you still should use them. But at least wait until the requirements get obvious. I, by now, suggest a process like the following:

1. Implement the current requirement in a concrete class
2. Adapt the concrete class to the ever changing requirements
3. If a different implementation for the same use case is required create an abstraction, but not earlier

Example: If you are supposed to integrate with some newsletter provider just implement the code in a concrete class. When starting, I assure you, not all final requirements are provided. There will be some additional requirements even before the first release of this small part of your application. If you came up with an abstraction first you'll have to adapt the abstraction **and** the concrete class. Also both will be more complicated since you already thought about making this somehow generic. Usually with requirements which will **not** show up because the other stakeholders will think about different things then you do.

Only when the second and third newsletter provider shows up and you are supposed to connect them with the same classes create an abstraction. Since we are a lot later in the project you are now better aware of all requirements and have a better understanding for the actual use-cases and differences between the endpoints.

### 3.8.1 Summary

Trying to follow "best practices" from library development in other projects will lead to wrong abstractions and increased workload. Do not abstract before you are sure about the actual requirements and before there is a need to do so. Concrete classes without interfaces are fine in Internal Projects, while they are not in Libraries or Adaptable Products.

## 4. Testing

### 4.1 Finding the right Test-Mix

[https://qafoo.com/blog/055\\_finding\\_the\\_right\\_test\\_mix.html](https://qafoo.com/blog/055_finding_the_right_test_mix.html)

*Benjamin Eberlei at 19. August, 2013*

The topic of Test Driven Development (TDD) and unit-testing usually creates heated discussions among developers. We at Qafoo are not an exception and the topic how to apply TDD in our daily work regularly leads to heated discussions. One reason for this is that both TDD and unit-testing combined are by some people seen dogmatically as the only way to do software development.

This blog post tries to summarize our current state of discussion on TDD, software testing and finding the right test-mix. It is written from my own perspective and all errors in description are mine, not my colleagues'. I also try to review the topic in context of the recent discussions of others.

Let's start with Uncle Bob, who is the primary voice of TDD as the only approach to software development. He recently wrote about this in his article "The startup trap<sup>1</sup>" spurring several discussions (HackerNews<sup>2</sup>, Storify<sup>3</sup>, Greg Youngs Blog<sup>4</sup>).

---

<sup>1</sup><https://qa.foo/book-TheStartUpTrap>

<sup>2</sup><http://martinfowler.com/bliki/TestPyramid.html>

<sup>3</sup><https://qa.foo/book-tdd-and-startups-uncle-bob-vs-nate-et-al>

<sup>4</sup><https://qa.foo/book-tdd-and-startups-uncle-bob-vs-nate-et-al>

However, he also acknowledges in a blog post the day later<sup>5</sup>, that he does not apply TDD dogmatically to every aspect of software development. Notably, he doesn't write tests for:

- state of objects, for example getters/setters and attributes themselves.
- GUIs, because they require lots of tinkering and small adjustments to configuration variables.
- Third-party libraries
- while experimenting with code (what I would call periods of "Spikes").

Most importantly, check how Uncle Bob never mentions the word unit-test in both blog posts. This is because TDD is about software design, not about testing your individual units (objects, functions, ...) or 100% code coverage. We think, TDD can be applied on any level of the testing pyramid<sup>6</sup>: the acceptance-, integration- and unit-test level. Kent Beck describes this in his book "Extreme Programming explained".

Why do I mention this? We have made the experience first hand that starting a new project with TDD and mostly unit-tests can actually slow down development considerably. This happens because during the development of new projects and features the requirements change frequently. This conflicts with tests locking the code to the current functionality through high test-coverage. Johann Peter Hartmann discussed this topic with Toby<sup>7</sup> in an interview on our blog some weeks ago.

Unit-tests are necessary to stabilize your code, but when you know the requirements are not stable yet, then having too many unit-tests can be a burden as well. It is very difficult to write unit-tests that are immune to change, especially if you don't know what changes will happen. Combined with a dynamic language like PHP that has poor automatic refactoring support and you will suddenly see yourself wasting lots of time adjusting unit-tests to changing requirements.

There are several learnings from this that we try to teach our customers in our trainings:

1. Don't focus on unit-tests exclusively. A good test-mix is required and uses acceptance-, integration- and unit-tests. A "good" ratio for software is some-

---

<sup>5</sup><https://qa.fo/book-ThePragmaticsOfTDD>

<sup>6</sup><http://martinfowler.com/bliki/TestPyramid.html>

<sup>7</sup>[https://qafoo.com/blog/051\\_testing\\_sweet\\_spot.html](https://qafoo.com/blog/051_testing_sweet_spot.html)

thing like 70% unit, 20% integration and 10% acceptance tests.

2. During development-spikes, with frequent requirement changes and uncertainty, it can be better to drive the design using acceptance- and integration-tests instead of unit-tests. Unit-test only those parts of the application that are mission critical, highly reused parts of the infrastructure (high cohesion) or stable already.
3. Once the software stabilizes, you should refactor tests from the acceptance- and integration-levels towards the unit-test-level. This makes the tests run much faster, less prone to failure due to side-effects, and allows you to write much more tests for different input combinations and edge-cases. Failures in unit-tests are also much easier to analyze than failures in acceptance tests.
4. Apply risk management to different components of your software: Having a lot of unit-tests is only important for those parts of your software that have a high business value. Features that are less important don't need the same test-coverage as those features that generate the majority of the business value. A few acceptance-tests might already be enough for those less important components.
5. Learning how to write tests that don't break on every occasion. This is beyond the scope of this blog-post.

#### 4.1.1 The Test-Mix Tradeoff

It is important to state that we don't advocate to stop testing. Instead, we are putting forth the notion of TDD decoupled from unit-testing and depending on business value (and risk) instead. This is a tradeoff by reducing the time for refactoring tests and increasing the time to run tests as well as the risk of failure due to uncovered code.

This approach to TDD is not a silver bullet, though: If your business doesn't allow for periods of stabilization or you wait too long before stabilizing, then you will end up with an inverted test-pyramid of many slow acceptance tests and just a few unit-tests or even worse, with no tests at all.

We found that our approach is closely related to the discussion on "Spike and

Stabilize" in Liz Keogh's blog<sup>8</sup> (including the comments), with the difference that we suggest using at least acceptance-tests during spikes. Her list of bullet points on applying TDD is a systematic checklist for the pragmatic choices Uncle Bob had in his blog-post.

#### 4.1.2 Conclusion

TDD is about design and not about unit-testing and 100% coverage. Using acceptance- and integration-tests is a valid approach for TDD and serves well during periods of spikes and frequent requirement changes. This decision trades a slower test suite and less stability for more flexibility to adjust the code. Neglecting to introduce unit-tests during code stabilization however might lead your code base to rot in the long run. Only a decent unit-test-coverage provides a developer with security to change code on all levels of the application.

---

<sup>8</sup><https://qa.fo/book-beyond-test-driven-development>

## 4.2 Practical PHPUnit: Testing XML generation

[https://qafoo.com/blog/007\\_practical\\_phpunit\\_testing\\_xml\\_generation.html](https://qafoo.com/blog/007_practical_phpunit_testing_xml_generation.html)  
Tobias Schlitt at 17. September, 2010

Testing classes which generate XML can be a cumbersome work. At least, if you don't know the right tricks to make your life easier. In this article, I will throw some light upon different approaches and show you, how XML generation can be tested quite easily using XPath.

### Disclaimer

The shown test code can even be more optimized by e.g. using a data provider in some places, but this is not in scope of this article. In addition, I do not obey to the one-assertion-per-test principle, since I disagree with it. Please feel free to discuss both in the comments anyway.  
;)

For better illustration, I use a real-world example we had at a customers project. Of course I anonymized and simplified the code for this blog entry. In the named project, XML was generated in order to communicate with a small web service. One part of the XML structure consisted of person data, which was encapsulated in a simple class:

```
class qaPerson
{
    const GENDER_MALE = 1;
    const GENDER_FEMALE = 2;

    protected $firstName;
    protected $lastName;
    protected $gender;
    protected $dateOfBirth;

    public function __construct(
        $lastName,
        $firstName
    ) {
        $this->setLastName( $lastName );
        $this->setFirstName( $firstName );
    }
}
```

```

20   public function getLastName()
    {
        return $this->lastName;
    }

25   public function getFirstName()
    {
        return $this->firstName;
    }

30   // ...
}

```

As you can see, there is nothing special in it: Some properties and corresponding getters. Of course there are also some setters, but I kept them out here, because they are not significant.

In order to generate an XML representation of the person data, the common visitor pattern<sup>9</sup> was used. The `qaPersonVisitor` takes a `DOMElement` in its constructor, which deals as the root for the XML serialization process. The method `visitPerson()` receives a `qaPerson` object and generates the corresponding XML elements:

```

class qaPersonVisitor
{
    protected $document;
    protected $root;
5   protected $currentElement;

    public function __construct( DOMElement $root )
    {
10      $this->root          = $root;
        $this->currentElement = $root;
        $this->document     = $root->ownerDocument;
    }

    public function visitPerson( qaPerson $person )
15   {
        $this->currentElement = $this->currentElement->appendChild(
            $this->document->createElement( 'Person' )
        );

20      $this->currentElement->appendChild(
            $this->document->createElement(
                'LastName',

```

<sup>9</sup>[https://qa.foo/book-Visitor\\_pattern](https://qa.foo/book-Visitor_pattern)

```
        $person->getLastName()
    )
25    );
    $this->currentElement->appendChild(
        $this->document->createElement(
            'FirstName',
            $person->getFirstName()
30        )
    );

    if ( null !== ( $gender = $person->getGender() ) )
    {
35        $this->currentElement->appendChild(
            $this->document->createElement(
                'Gender', $gender
            )
40        );
    }
    if ( null !== ( $dateOfBirth = $person->getDateOfBirth() ) )
    {
45        $this->currentElement->appendChild(
            $this->document->createElement(
                'DateOfBirth', $dateOfBirth->format( 'Y-m-d' )
            )
50        );
    }
}
```

So, the code does no magic at all and it is pretty straight forward, as a visitor should be. So, how would you attempt to write a unit test for the `visitPerson()` method?

In following, I first present two basics approaches to this task, which I've seen and wrote myself many times. After that, I show you two more decent approaches, which are shipped with PHPUnit, based on CSS selectors and a custom tag matching data structure. Finally, I will show how we solved the task in the customer project with a very slick and clean XPath based approach.

With each presented approach I will give a short discussion on what the pros and cons are. Please feel invited to tell us about additional approaches you might know and to discuss your impressions on the described approaches.

You can find the complete code examples for testing XML generation<sup>10</sup> in our Github repository for blog code examples.

<sup>10</sup>[https://qa.fo/book-phpunit\\_xml](https://qa.fo/book-phpunit_xml)

### 4.2.1 Test case basics

In order to test the XML generation, the `qaPersonVisitor` class has the precondition of receiving a `DOMElement` object in its constructor. In addition to that, the `visitPerson()` method always needs a valid instance of `qaPerson` to be serialized. Therefore, the test class `qaPersonVisitorTest` makes use of PHPUnit's `setUp()` method, to set up the DOM basics, and implements a custom method to return a fixture for the `qaPerson` class:

```
class qaPersonVisitorTest extends PHPUnit_Framework_TestCase
{
    protected $domDocument;
    protected $rootElement;

    5     protected function setUp()
        {
            $this->domDocument = new DOMDocument( '1.0', 'utf-8' );
            $this->rootElement = $this->domDocument->appendChild(
    10             $this->domDocument->createElement(
                'root'
            )
        );
    }

    15     protected function getDomDocument()
        {
            return $this->domDocument;
        }

    20     protected function getDomRootElement()
        {
            return $this->rootElement;
        }

    25     protected function getPersonFixture()
        {
            $person = new qaPerson(
    30             'My Last Name',
                'Some First Name'
            );
            $person->setGender( qaPerson::GENDER_FEMALE );
            $person->setDateOfBirth(
    35             new DateTime( '2000-01-01 00:00:00+00:00' )
            );

            return $person;
        }
}
```

```
40 | // ...  
    | }
```

This class is used as the basis for all of the test methods described in following.

### 4.2.2 The naive way

The naive way of testing if XML is generated correctly is to manually rebuild the expected DOM structure in the test case and assert that this equals the generated one:

```
public function testVisitPersonNaive ()  
{  
    $person = $this->getPersonFixture ();  
  
    $expDom = new DOMDocument( '1.0', 'utf-8' );  
    $expRoot = $expDom->appendChild (  
        $expDom->createElement( 'root' )  
    );  
    $expPersonElem = $expRoot->appendChild (  
        $expDom->createElement( 'Person' )  
    );  
    $expPersonElem->appendChild (  
        $expDom->createElement (  
            'LastName',  
            $person->getLastName ()  
        )  
    );  
    $expPersonElem->appendChild (  
        $expDom->createElement (  
            'FirstName',  
            $person->getFirstName ()  
        )  
    );  
    $expPersonElem->appendChild (  
        $expDom->createElement (  
            'Gender',  
            $person->getGender ()  
        )  
    );  
    $expPersonElem->appendChild (  
        $expDom->createElement (  
            'DateOfBirth',  
            $person->getDateOfBirth ()->format( 'Y-m-d' )  
        )  
    )  
}
```

```
);  
  
$visitor = new qaPersonVisitor( $this->getDomRootElement() );  
$visitor->visitPerson( $person );  
39  
  
$this->assertEquals(  
    $expDom,  
    $this->getDomDocument()  
);  
44 }
```

Before I inspect this technique a bit deeper, it is important to note that one thing: In a real life environment this test case should not be the only one for the `visitPerson()` method. Since the method contains two conditions, you would at least need four test cases to fully test its functionality: One for each combination of the conditions.

Comparing the code of the test method to the tested `visitPerson()` method, you can see that they have quite some code in common. In fact, their code is equal except for the two conditions. There are two flaws hidden in this observation:

First, there is the risk that developers get lazy and just copy & paste between test and actual code. Why should you re-write code that already exists in either of the methods? Even if no real C&P happens, a developer still reminds how he wrote one of the methods and will do in-mind copying when developing the counterpart.

But why is this actually dangerous? Whenever a human tends to duplicate code instead of writing new code for testing, logic errors are likely to be evident in both versions. This is pretty much like reading your own writings: You read what you were intended to write, not what you actually wrote. So a goal of each test should be to reformulate the expectations regarding the tested method in an independent way.

The second problem with the shown method is, that it involves quite some code. 43 lines is not what one would call a short and simple test method. Even if the code is still simple, there is a danger that people will not have the patients to read the code again when touching the method. Again, if you need to writing a lot of code twice, it is pretty seductive to just copy & paste it.

### 4.2.3 The lazy way

One of the flaws of the naive testing method, where the XML structure is rebuilt by hand for comparison, is the amount of code to write for testing purposes. In case of testing generated XML structures, it is easy to ship around this by storing the expected XML to disc and compare this to the structure generated by the `visitPerson()` method:

```
public function testVisitPersonCompareFile ()
{
    $person = $this->getPersonFixture ();

    5     $visitor = new qaPersonVisitor( $this->getDomRootElement () );
        $visitor->visitPerson( $person );

    $this->assertXmlStringEqualsXmlFile (
    10     'data/' . __CLASS__ . '_' . __FUNCTION__ . '.xml',
        $this->getDomDocument ()->saveXml ()
    );
}
```

The first thing to be noted here is the shortness of the code: 11 lines compared to 43 in the naive test method. This is the shortest test method presented in the article, which is quite appealing. But this is only half the truth, since the expected XML structure has to be written in addition. This way of testing XML also involves a serious danger: Why would you write XML by hand, if you are about to write (or have already implemented) a method for this? So chances are good that you switch to the dark side and just store the generated XML and inspect it manually. A no-go for serious testing.

Of course, you can force yourself to write the comparison XML by hand. If you manage to do so, this way of testing XML is very convenient. You can even automate the assertions some more and therefore shorten your code. However, I suggest that you better take a look at the rest of the proposed test proceedings, just to make sure you don't lose self-control once in a while. Remember when projects become hounded, will you then still write the expected XML by hand?

### 4.2.4 Selecting with CSS

The third way of testing XML involves some not so well known functionality of PHPUnit: Assertions which allow you to select parts of the XML structure by CSS

selectors. Since I expect you did not use the corresponding `assertSelectCount()` and `assertSelectEquals()` methods before, I'll explain the code in more detail in following. If you already know these methods, feel free to skip directly to the discussion part, which follows after the next section.

### The code

```
public function testVisitPersonSelectCSS()
{
    $person = $this->getPersonFixture();

    5    $visitor = new qaPersonVisitor( $this->getDomRootElement() );
    $visitor->visitPerson( $person );

    $this->assertSelectCount(
    10     'Person',
        1,
        $this->getDomDocument(),
        'Invalid number of Person elements',
        false
    );

    15    $this->assertSelectEquals(
        'Person > FirstName',
        $person->getFirstName(),
        1,
    20     $this->getDomDocument(),
        'Invalid content of FirstName element',
        false
    );

    25    $this->assertSelectEquals(
        'Person > LastName',
        $person->getLastName(),
        1,
    30     $this->getDomDocument(),
        'Invalid content of LastName element',
        false
    );

    35    $this->assertSelectEquals(
        'Person > Gender',
        $person->getGender(),
        1,
        $this->getDomDocument(),
        'Invalid content of Gender element',
        false
    );
};
```

```
40     $this->assertSelectEquals(
        'Person > DateOfBirth',
        $person->getDateOfBirth()->format( 'Y-m-d' ),
        1,
        $this->getDomDocument(),
45     'Invalid content of DateOfBirth element',
        false
    );
}
```

There are two assert methods called in this code, which I suspect many people did not know about: The `assertSelectCount()` method selects a set of XML nodes through a CSS selector and ensures that the selected number of nodes equals a defined count. The method `assertSelectEquals()` also asserts the number of selected nodes, but in addition makes an assertion of their contents.

But lets look at the usage of these methods step by step: The first assertion in the example ensures that exactly one `<Person>` element is generated. The first parameter to the `assertSelectCount()` method defines the CSS selector to be matched. Note that there is of course not the full CSS 3 specification<sup>11</sup> supported by the parser implemented in PHPUnit, but the implemented sub set works quite nicely for testing purposes.

The second parameter defines the number of XML nodes to be found by this selector and the third one gives the `DOMDocument` to be tested. The fourth parameter provides the typical optional failure message, which is printed by PHPUnit if the test fails. The fifth and last parameter indicates if the given `DOMDocument` contains a HTML structure (default) or some custom XML dialect.

Note that there were some flaws in the PHPUnit test code, regarding the CSS assertions on XML instead of HTML documents, which have been fixed recently. So, in order to use this way of testing right now, you should use a recent PHPUnit checkout from Github<sup>12</sup>.

The other used assertion, `assertSelectEquals()`, works quite similar: The first parameter is again the CSS selector. In the first call to this method, this selects the `<FirstName>` elements below the `<Person>` element. Remember, that the previous assertion already ensured that there is only 1 `<Person>` element. The second parameter

<sup>11</sup><http://www.w3.org/TR/css3-roadmap/>

<sup>12</sup><http://github.com/sebastianbergmann/phpunit/>

contains the expected string content of the selected XML elements. In the first use of the assertion, this is the persons last name. Parameter three determines how many elements are expected to be matched by the assertion. The last three parameters are the same as for `assertSelectCount()`: The `DOMDocument`, a message and the indicator for XML instead of HTML.

There exists a third method in the round dance of CSS selector based assertions in PHPUnit: `assertSelectRegExp()` allows you to provide a regular expression for the expected content instead of just a plain string. This is quite useful for e.g. generated lists of elements, which have a common text structure, but slightly different actual content.

### Discussion

After the code has been explained, let's return to the actual topic: How well is this method of testing suited?

First of all, the potential flaw of copying and pasting code from or to the `visitPerson()` method to the test is not evident. By making use of CSS selector matches, you make sure that you write your test code completely independent from the code you're testing. This is a very good thing, because you need to re-think what you actually expect and re-formulate these thoughts in a completely different way. This fact makes the CSS selector way of testing much more robust than the previously shown test strategies.

However, this way of testing is not flawless, too. Still the number of code lines to write is quite large: 47 lines is even more code than generating the expected XML structure by hand in the test method. In addition to that, the typical developer dealing with XML might not be familiar with the way of selecting nodes through CSS selectors, although this technique became quite common through e.g. jQuery<sup>13</sup> in the JavaScript world. For someone not familiar with this, your test code is harder to read than any of the previously shown methods.

Finally, the code involved for the assertion in the depth of PHPUnit is quite large. It's not a simple equality or type assertion, but involves parsing the CSS selector and performing the corresponding selection of DOM nodes. Involving that much code raises the potential of code flaws in the test code itself, so it is hard to guarantee that your test actually asserts what you are looking for.

---

<sup>13</sup><http://jquery.com/>

### 4.2.5 Tag match assertions

Build on the same background code base, PHPUnit also offers a custom way of matching XML structures through the `assertTag()` method. In fact, if you use one of the CSS selecting assertions, your selectors will be parsed into the structure used by this assertion. Again, if you already know the `assertTag()` method, feel free to skip the following subsection, and go directly to the discussion part.

#### The code

```
public function testVisitPersonTag ()
{
    $person = $this->getPersonFixture ();

    5    $visitor = new qaPersonVisitor ( $this->getDomRootElement () );
    $visitor->visitPerson ( $person );

    $this->assertTag (
        10    array (
            'tag' => 'Person' ,
            'child' => array (
                'tag' => 'LastName' ,
                'content' => $person->getLastName () ,
            ) ,
        15    ) ,
        $this->getDomDocument () ,
        'Incorrect LastName tag' ,
        false
    );

    20    $this->assertTag (
        array (
            'tag' => 'Person' ,
            'child' => array (
                25    'tag' => 'FirstName' ,
                'content' => $person->getFirstName () ,
            ) ,
        ) ,
        $this->getDomDocument () ,
        30    'Incorrect FirstName tag' ,
        false
    );

    35    $this->assertTag (
        array (
            'tag' => 'Person' ,
```

```

        'child' => array(
            'tag'      => 'Gender',
            'content' => (string) $person->getGender(),
        ),
    ),
    $this->getDomDocument(),
    'Incorrect Gender tag',
    false
);

$this->assertTag(
    array(
        'tag' => 'Person',
        'child' => array(
            'tag'      => 'DateOfBirth',
            'content' => $person->getDateOfBirth()->format( 'Y-m-d' ),
        ),
    ),
    $this->getDomDocument(),
    'Incorrect DateOfBirth tag',
    false
);
}

```

A tag assertion is defined by an array, potentially reflecting structure and contents of XML elements at the same time. In the shown example, only two of the possible definitions are used: Definition of a tag and one of its children and definition of the contents of this child.

The first call to `assertTag()` in the above example asserts that there exists a tag `<Person>` which has a child `<LastName>` with the expected content. In order to define this, the given array contains the key `tag` assigned to the name of the tag to look for (`Person` in this case). The second array key `child` defines an assertion on a child of this tag, again using an array structure: In this second level array, again a specific tag is required (`LastName`). However, in this case, not the structure of this tag is defined, but the `content`, indicated by the corresponding key.

There are various array keys which allow you to define different XML constructs. Assertions are possible on attributes, the parent of a tag, the number of children and some more. You can find an overview on the possibilities in the PHPUnit docs on `assertTag()`<sup>14</sup>, as well as in the API docs of the method.

<sup>14</sup><https://qa.foo/book-api>

### Discussion

As with the assertion through CSS selectors shown in the previous section, the `assertTag()` method allows you to formulate your expectations in a completely different way than the code to be tested is written. This is, as explained, a good thing.

A fundamental drawback of `assertTag()` is the syntax of the tag matches. I'm not aware of any construct in PHP (or any other language I know) which uses a similar way of defining how an XML structure should look like. Therefore, it might be really hard for other developers to read your test cases.

Furthermore, while the possibilities with this way of assertions are quite large, it can still not cover all edge cases and you might be forced to choose an alternative way of testing for these. The amount of code for testing is also quite high: With 58 lines, this is the longest test method presented in this article.

#### 4.2.6 Using XPath

A very decent, but not really documented, way for testing the generation of XML content is using XPath. For this purpose, we have implemented a very simple custom assertion method in the named customer project, which is presented below. After that, I show two ways of how this method can practically be used in our visitor example. Finally, the pros and cons of using XPath are discussed.

I won't go into details on XPath itself in this article, so if you are not familiar with this standard language for addressing XML nodes, please refer to one or more of the following resources:

- The extensive XPath introduction article<sup>15</sup> on my private blog
- The W3C specs of XPath 1.0<sup>16</sup>
- The PHP documentation on DOMXPath<sup>17</sup>

Please note that `DOMXPath` does only support XPath version 1.0 and does not support XPath 2.0.

---

<sup>15</sup>[http://schlitt.info/opensource/blog/0704\\_xpath.html](http://schlitt.info/opensource/blog/0704_xpath.html)

<sup>16</sup><http://www.w3.org/TR/xpath/>

<sup>17</sup><http://php.net/domxpath>

### Custom assertion

The custom assertion we used is quite simple and implemented in a new base test case class, so it can be re-used for different XML generation tests:

```

abstract class qaXmlTestCase extends PHPUnit_Framework_TestCase
{
    protected abstract function getDomDocument();

5    protected function assertXPathMatch( $expected, $xpath,
        $message = null, DOMNode $context = null )
    {
        $dom = $this->getDomDocument();
        $xpathObj = new DOMXPath( $dom );

10        $context = $context === null
            ? $dom->documentElement
            : $context;

15        $res = $xpathObj->evaluate( $xpath, $context );

        $this->assertEquals(
            $expected,
            $res,
            $message
20        );
    }
}

```

The abstract `getDomDocument()` method is required in order to instantiate the `DOMXPath` object properly. The actual assertion method `assertXPathMatch()` receives two mandatory parameters: The result expected to be produced by the XPath query and the query itself. Furthermore, it optionally receives the typical PHPUnit message to be displayed if the assertion failed (third parameter). The last parameter can optionally be a `DOMNode` that deals as the context for the XPath query. If this parameter is left out, the query is meant to have the root element as its context.

### Extensive XPath

The usage of this custom assertion is quite simple, too. Below you find a quite extensive code example which asserts the correctness of the generated XML structure step by step:

```

public function testVisitPersonXPathExtensive()
{

```

```
    $person = $this->getPersonFixture();  
5    $visitor = new qaPersonVisitor( $this->getDomRootElement() );  
    $visitor->visitPerson( $person );  
  
    $this->assertXPathMatch(  
10        1,  
        'count(/root/Person)',  
        'Incorrect number of Person elements.'  
    );  
  
    $this->assertXPathMatch(  
15        $person->getLastName(),  
        'string(/root/Person/LastName)',  
        'Incorrect or missing LastName element.'  
    );  
    $this->assertXPathMatch(  
20        $person->getFirstName(),  
        'string(/root/Person/FirstName)',  
        'Incorrect or missing FirstName element.'  
    );  
    $this->assertXPathMatch(  
25        $person->getGender(),  
        'string(/root/Person/Gender)',  
        'Incorrect or missing Gender element.'  
    );  
    $this->assertXPathMatch(  
30        $person->getDateOfBirth()->format( 'Y-m-d' ),  
        'string(/root/Person/DateOfBirth)',  
        'Incorrect or missing DateOfBirth element.'  
    );  
}
```

The first assertion ensures that exactly one `<Person>` element was actually generated. The XPath function `count()` counts the number of nodes in the selected node set. The query inside this function looks from the root of the document for all `<root>` elements and selects all of their children which are `<Person>` elements.

The subsequent assertions take care of each child element of the `<Person>` and assert the correctness of their content. The XPath function `string()` converts the nodes in a selected nodes set into a single string, by simple concatenating their text content. So inherently, the assertion ensures that only a single such element (with content) exists.

You can make use any of the XPath 1.0 function here, to make the XPath query

return something else than a node set, i.e. a DOMNodeList in PHP. If you don't want to use XPath functions or have an assertion which is not possible with them, you can still write an additional assertion method performing your custom operations on the XPath match.

### Short XPath

While the last example was quite long, in terms of code lines, a similar assertion could have been achieved with some more complex XPath in fewer lines of code:

```
public function testVisitPersonXPathShort()
{
    $person = $this->getPersonFixture();

    $visitor = new qaPersonVisitor( $this->getDomRootElement() );
    $visitor->visitPerson( $person );

    $this->assertXPathMatch(
        1,
        sprintf( 'count(/root/Person[
            . 'LastName = "%s" and '
            . 'FirstName = "%s" and '
            . 'Gender = "%s" and '
            . 'DateOfBirth = "%s"])' ,
        15         $person->getLastName(),
                $person->getFirstName(),
                $person->getGender(),
                $person->getDateOfBirth()->format( 'Y-m-d' )
        ),
        20         'Mismatching XPath.'
    );
}
```

The example uses a single XPath expression to count only those <Person> elements with correct content, i.e. a <LastName> element with correct text content and a <FirstName> element with correct content and so on. For this purpose, the XPath matches the <Person> element and puts conditions on it, which are called predicates in XPath terminology.

It has to be noted that this single expression is not as accurate as the multiple ones from the previous example were: For example, it could happen here, that two <Person> elements are generated of which only one matches the given conditions.

**Discussion**

So, let's review all the pros and cons I talked about for the previously presented testing methods:

With XPath you need to re-formulate the expectations to your XML structure and contents explicitly and you cannot do copy & paste between test code and actual code. Furthermore, you can hardly auto-generate an XPath from the generated XML, so you are actually forced to think yourself and do write your test code on your own.

The amount of code to write for your test cases is rather flexible, depending on your custom needs for the test case. The two examples above have 33 and 21 lines of code, which is not few but also not much.

As with the CSS selector and tag match examples, XPath queries are quite flexible in what you want to expect from the generated XML. In fact, XPath is even more powerful here. In addition to that, XPath has the clear benefit of being a standardized language which many developers, who deal with XML, do know well.

**4.2.7 Conclusion**

In this article, I've shown different ways of how you can test XML generation on basis of PHPUnit. In my eyes the most effective, easy and sensible way is the XPath solution shown in the last section. What you actually use in your projects depends much on the amount of XML generation to test and the circumstances, but the XPath way should always be a reasonable choice.

So, how do you test XML generation? Do you use similar techniques than the ones presented here? Did you learn something useful from this blog entry? Please leave a comment!

### 4.3 Testing file uploads with PHP

[https://qafoo.com/blog/013\\_testing\\_file\\_uploads\\_with\\_php.html](https://qafoo.com/blog/013_testing_file_uploads_with_php.html)

Manuel Pichler at 9. December, 2010

A question I am asked on a regular basis is, how you can test a file upload with PHP. In this blog post, I take a precise look at this topic and show you how to test your file uploads from within your standard testing environment, using widely unknown testing framework for PHP.

Let me begin with the testing framework that we all know and use in our daily work, PHPUnit<sup>18</sup>. With this testing framework you have the opportunity to test almost every part of your software. This also applies to file uploads, as long as the application under test realizes uploads only via the PHP's magic `$_FILES` variable. In this case you can easily prepare an appropriately crafted `$_FILES` array in the test's `setUp()` method, that can be accessed by the software under test:

```
// ...
protected function setUp()
{
    parent::setUp();

    5     $_FILES = array(
        'file_valid' => array(
            'name' => 'foo.txt',
            'tmp_name' => '/tmp/php42up23',
            'type' => 'text/plain',
            'size' => 42,
            'error' => 0
        )
    );
    15 }
// ...
```

But fortunately, in most cases this is not quite so simple because the software to be tested utilizes the safer PHP functions `is_uploaded_file()` and `move_uploaded_file()`. And in this case the manipulation of the `$_FILES` array does not work, because both functions operate on another level, so that you cannot manipulate the input data within userland code:

```
class UploadExample
{
```

<sup>18</sup><http://phpunit.de>

```
5   protected $dest;
   public function __construct($dest)
   {
       $this->dest = rtrim($dest, '/') . '/';
   }
   public function handle($name)
   {
10      if (is_uploaded_file($_FILES[$name]['tmp_name'])) {
           move_uploaded_file(
               $_FILES[$name]['tmp_name'],
               $this->dest . $_FILES[$name]['name']
           );
15      }
   }
}
```

Of course, you can still test this part of the application with a heavyweight framework like Selenium<sup>19</sup>. Which, however, brings a number of disadvantages: You must prepare and integrate a variety of other infrastructure components. Beginning with a webserver, the Selenium server, a graphical user interface with a browser and other infrastructure + bootstrap code that is required for a working version of the software. All this increases the required effort to execute a single test and the execution time of the test itself. This carries the danger that false positives arise, caused as a side effects from the complex test infrastructure.

An alternative is the *PHP Testing Framework* or in short *PHPT*, which the core developers use to test PHP<sup>20</sup> itself and that is used by various other PHP related projects like PEAR<sup>21</sup> and PECL<sup>22</sup>. I would describe PHPT as a lightweight testing framework, that is easy to learn due to its simple syntax and a very expressive description format. In this article I will only take a look at a limited set of the PHPT syntax. A complete documentation of PHPT can be found on the PHP Quality Assurance<sup>23</sup> website. The following example shows the three base elements of a minimal PHPT test case.

```
# A description of the test itself:
```

---

<sup>19</sup><http://seleniumhq.org>

<sup>20</sup><http://php.net>

<sup>21</sup><http://pear.php.net>

<sup>22</sup><http://pecl.php.net>

<sup>23</sup>[http://qa.php.net/phpt\\_details.php](http://qa.php.net/phpt_details.php)

```

---TEST---
Example test case

    # The code under test(CUT):

---FILE---
<?php
3 var_dump(strpos('Manuel Pichler', 'P'));
var_dump(strpos('Manuel Pichler', 'Z'));

    # And the test expectations:

1 ---EXPECT---
int(7)
bool(false)

```

But the coolest thing is that, without knowing it, almost everyone has PHPT already installed, because every PEAR installation already contains a PHPT test-runner. This test-runner can be called by the following command:

```

~ $ pear run-tests example-test-case.phpt
2 Running 1 tests
PASS Example test case[example-test-case.phpt]
TOTAL TIME: 00:00
1 PASSED TESTS
0 SKIPPED TESTS

```

After this brief introduction into PHPT, let's come back to the original problem: How to test a file-upload with PHP? Here exactly lies the great strength of PHPT compared to other testing frameworks. With PHPT we can simulate almost every state the application under test could run into. For example, it allows you to alter all `php.ini` settings, to disable internal functions, or to configure an `open_basedir` restriction for a test. With PHPT it is also possible to manipulate all data that is passed into PHP process, including the various super global variables like `$_FILES`, `$_POST`, `$_SERVER`, `$_ENV`. These changes occurs on a different abstraction level than you can get in PHPUnit, so that even internal functions like `is_uploaded_file()` and `move_uploaded_file()` operate on the manipulated test data.

In order to test a file-upload one brick is still missing, namely how to simulate an upload with PHPT? For this you need the `--POST_RAW--` element, where — as the name implies — a raw HTTP Post message can be specified. The easiest way to get appropriate test data is probably to record a real file-upload with a tool like

Wireshark<sup>24</sup> and copy the relevant data into the PHPT test. The following listing shows such a recording:

```

Content-Type: multipart/form-data; boundary=-----WebKitFormBoundaryfywL8UCjFtqUBTQn
-----WebKitFormBoundaryfywL8UCjFtqUBTQn
Content-Disposition: form-data; name="file"; filename="example.txt"
5 Content-Type: text/plain

Qafoo provides quality assurance support and consulting

-----WebKitFormBoundaryfywL8UCjFtqUBTQn
10 Content-Disposition: form-data; name="submit"

Upload
-----WebKitFormBoundaryfywL8UCjFtqUBTQn--

```

Now you have all information together to test the upload of a file with PHPT. So here is the actual test's description:

```

--TEST--
2 Example test emulating a file upload

```

Then you add the Wireshark recording with the file upload in the `--POST_RAW--` element of the PHPT-file:

```

--POST_RAW--
Content-Type: multipart/form-data; boundary=-----WebKitFormBoundaryfywL8UCjFtqUBTQn
-----WebKitFormBoundaryfywL8UCjFtqUBTQn
5 Content-Disposition: form-data; name="file"; filename="example.txt"
Content-Type: text/plain

Qafoo provides quality assurance support and consulting

-----WebKitFormBoundaryfywL8UCjFtqUBTQn
10 Content-Disposition: form-data; name="submit"

Upload
-----WebKitFormBoundaryfywL8UCjFtqUBTQn--

```

Now you just need a little bit glue code to execute the upload component and the actual test assertions:

```

1 --FILE--

```

<sup>24</sup><http://wireshark.org>

```

<?php
require __DIR__ . '/UploadExample.php';

$upload = new UploadExample('/tmp');
$upload->handle('file');

var_dump(file_exists('/tmp/example.rst'));
?>
--EXPECT--
11 bool(true)

```

And that's it. You have a complete PHPT test. Of course, you verify that it works by calling the `run-tests` command:

```

~ $ pear run-tests upload-example.phpt
Running 1 tests
PASS Example test emulating a file upload[upload-example.phpt]
TOTAL TIME: 00:00
5 1 PASSED TESTS
0 SKIPPED TESTS

```

The test runs and the task has been completed successfully.

Maybe you wonder, how you can integrate these PHPT tests into an existing test infrastructure? A fairly simple solution exists: A relatively unknown feature of PHPUnit<sup>25</sup> is the built-in support for PHPT tests. This provides the great benefit that you must not worry about the integration, your tests must only inherit from the appropriate test classes:

- `PHPUnit_Extensions_PhptTestCase`
- `PHPUnit_Extensions_PhptTestSuite`

The `PHPUnit_Extensions_PhptTestCase` class can be used to reference a single PHPT test file, which is then executed by the PHPUnit testing framework. It has to be noted that the absolute path to the PHPT file must be specified:

```

<?php
require_once 'PHPUnit/Extensions/PhptTestCase.php';

5 class UploadExampleTest extends PHPUnit_Extensions_PhptTestCase
{
    public function __construct()
    {
        parent::__construct(__DIR__ . '/upload-example.phpt');
    }
}

```

<sup>25</sup><http://phpunit.de>

```
10 }  
    }
```

Alternatively, you can use the `PHPUnit_Extensions_PhptTestSuite` class, that takes a directory as its first constructor argument and then searches for all `*.phpt` files within this directory:

```
<?php  
require_once 'PHPUnit/Extensions/PhptTestSuite.php';  
  
5 class UploadExampleTestSuite extends PHPUnit_Extensions_PhptTestSuite  
  {  
    public function __construct()  
    {  
        parent::__construct(__DIR__);  
    }  
10 }
```

Using this powerful combination it should be easy for you to integrate file upload tests into your existing test infrastructure.

As an attachment to this article you can find a more complex example that tests a file-upload within the context for the Zend-Framework. The great advantage of PHPT tests is that such a test can be executed without the complicated path through a webserver and the HTTP protocol.

You can find the full code examples<sup>26</sup> from this blog post in our github repository.

---

<sup>26</sup>[https://qa.fo/book-testing\\_file\\_uploads](https://qa.fo/book-testing_file_uploads)

## 4.4 Mocking with Phake

[https://qafoo.com/blog/037\\_phake.html](https://qafoo.com/blog/037_phake.html)

*Benjamin Eberlei at 13. March, 2013*

Update (14.3.2013): Introduced Test Double wording instead of using mock objects for everything to more cleanly describe the different concepts.

The use of Test Doubles is an important skill to learn when using Test Driven Development (TDD). Test Doubles allow you to replace dependencies of an object with lookalikes, much like crash test dummies are used during automobile safety tests so humans aren't harmed.

### 4.4.1 Test Doubles Explained

Creating a test double in PHP involves creating a new class that extends the public API of the original class with empty methods. You can safely use all methods of a test double and they will do nothing, rather than calling the original code of the original class. There are two ways to create these test doubles: You can write them yourself or use one of the many existing libraries.

I would strongly recommend to use any of the existing libraries that can simplify and automate this task for you. Technically they work using code-generation at run-time.

To allow interactions with test doubles there are three ways to configure them in any library:

- Add expectations of the arguments passed to a method (Verification)
- Add results that are returned from a method-call to the mock object (Stubbing)
- Delegate calls to the original code (Spying)

Test doubles using the first approach are called Mock Objects. Objects of the second type are called Stubs, of the third type Spies.

### 4.4.2 Benefits of Test Doubles

There are many reasons why test doubles are useful:

- Allow units (objects) to be tested in isolation of their dependencies. This is done by replacing the dependencies with test doubles.

- Allow verification of behavior between objects. In contrast to assertions that can only verify the state of objects in isolation. This makes them very useful with relation to Behavior Driven Development (BDD) inside your unit-tests.
- Test Doubles are useful to test-drive new interfaces based on required behavior without caring for the implementation at the moment.

That means testdoubles are invaluable to move from state-based object-oriented programming to a behavioral approach based on sending messages between objects.

#### 4.4.3 Introduction to Phake

Using Test Doubles in PHPUnit tests means using the built-in MockObjects library<sup>27</sup> for quite some years. In the last years two contenders emerged that can be used as optional dependencies in PHPUnit:

- Mockery by Padraic Brady<sup>28</sup>
- Phake by Mike Lively<sup>29</sup>

Both can be installed using Composer and integrated into your projects very easily.

This blog post introduces Phake, because it works quite differently than both PHPUnit Mock Objects and Mockery:

1. In the well known four phase test<sup>30</sup> with Setup, Exercise, Verify, Teardown both PHPUnit mocks and Mockery require expectations to be part of the "setup" phase.

This is unfortunate, because mock expectations are much more related to the "Verify" phase instead. This happens in PHPUnit and Mockery, because they don't explicitly differentiate between methods that are mocked (verification) or stubbed (returning results). Phake introduces a differentiation by requiring different APIs to be used for configuration.

2. Instead of using strings for method names and builder methods for arguments, Phake let's you prototype the method-call in actual PHP code. This simplifies

---

<sup>27</sup><https://github.com/sebastianbergmann/phpunit-mock-objects>

<sup>28</sup><https://github.com/padraic/mockery>

<sup>29</sup><https://github.com/mlively/Phake>

<sup>30</sup><http://xunitpatterns.com/Four%20Phase%20Test.html>

the mock object configuration considerably and requires much less typing.

Let's see an example containing both mock and stub objects in one test for loading the weather data for a given location.

```

<?php
class LoaderTest extends \PHPUnit_Framework_TestCase
{
5   public function testGetWeather()
    {
        // 1. setup
        $logger = \Phake::mock('Qafoo\Weather\Logger');
        $functional = \Phake::mock('Qafoo\Weather\Service');

10        // Stubbing with Phake::when()
        \Phake::when($functional)->getWeatherForLocation()->thenReturn(
            new Struct\Weather('Fair', 23, 0, 'NW')
        );

15        $loader = new Loader($functional, $logger);

        // 2. exercise
        $locatedWeather = $loader->getWeatherForLocation(
20            new Struct\Location('Berlin', 'Germany')
        );

        // 3. verify
        $this->assertInstanceOf(
25            'Qafoo\Weather\Struct\LocatedWeather',
            $locatedWeather
        );

        // Verification with Phake::verify()
30        \Phake::verify($logger)->log('Fetched weather for Berlin Germany.');
```

Using the `Phake::when()` call and passing a mock object you can prototype what a method call should return for your code to show a desired behavior. See how we can just call `->getWeatherForLocation()` as a prototype for how the stub behaves instead of `PHPUnits ->method('getWeatherForLocation')` or `Mockerys ->shouldReceive('getWeatherForLocation')`.

Using `thenReturn` specifies the return value of this method call and completes the description of how the stub works. If you want to return different values on con-

secutive calls, just chain multiple `thenReturn` calls. You can use `thenThrow` to throw exceptions.

Verification is done with `Phake::verify()` after the tested code was actually exercised. We again prototype what method calls and which arguments we want to verify, in this case `->log('Fetched weather for Berlin Germany')`; on the logger mock.

This is a very simple stubbing and verification example with Phake. Comparable to PHPUnit and Mockery, we could add more complex expectations:

- Using argument matchers to verify the structure of arguments used in method calls.
- Checking multiple invocations such as exactly n-times, at least N-times or others.
- Verify no interaction with a mock happened at all.
- Verify no further interaction happened than the ones already verified.

#### 4.4.4 Conclusion

Phake is a great mocking library and can be easily integrated into PHPUnit. Its new approach to prototype mocks and stubs and the separation between stubbing and verification phases is very refreshing and easy to use.

If you want to go into more detail and learn about Phake, you should check out the extensive documentation<sup>31</sup>.

---

<sup>31</sup><http://phake.digitalsandwich.com/docs/html/>

## 4.5 Testing Effects of Commands With Phake::capture()

[https://qafoo.com/blog/078\\_phake\\_capture.html](https://qafoo.com/blog/078_phake_capture.html)

*Benjamin Eberlei at 8. March, 2016*

Today I want to share a simple trick for the excellent Mocking library Phake<sup>32</sup> (I wrote about it before (*See: Mocking with Phake*)) when testing state on APIs that don't return values.

Testing becomes harder when you are using command / query separation in your code and service operations don't return values anymore. If your command method creates objects and passes them down the stack, then you usually want to make assertions on the nature of the changes.

Take the following example that creates a new `Order` object:

```
<?php
class CheckoutHandler
{
    5     private $orderRepository;
        private $productRepository;

        public function __construct($orderRepository, $productRepository)
        {
            10     $this->orderRepository = $orderRepository;
                $this->productRepository = $productRepository;
        }

        public function checkout(Checkout $command)
        {
            15     $order = new Order();
                $order->setAddress($command->address);

                foreach ($command->productIds as $id => $amount) {
                    20     $product = $this->productRepository->find($id);
                        $order->addItem($product, $amount);
                }

                $this->orderRepository->save($order);
            }
        }
    25 }
```

A "usual" PHPUnit test for this class can only make a single assertion that the

---

<sup>32</sup><http://phake.readthedocs.org/en/2.1/>

`OrderRepository` is called with an `Order` object. But we might want know if a product was correctly assigned.

With `Phake::capture($value)` we can assign the argument passed to `OrderRepository#save($order)` to a variable that is available inside the Unit-Test, ready to run assertions on.

```
<?php
class CheckoutHandlerTest extends \PHPUnit_Framework_TestCase
{
    5     public function testCheckout()
        {
            $orderRepository = \Phake::mock(OrderRepository::class);
            $productRepository = \Phake::mock(ProductRepository::class);

            10     $product = new Product();
            \Phake::when($productRepository)->find(42)->thenReturn($product);

            $handler = new CheckoutHandler($orderRepository, $productRepository);
            $handler->checkout(new Checkout([
            15         'productIds' => [42 => 1],
                'address' => new Address(),
            ]));

            \Phake::verify($orderRepository)->save(\Phake::capture($order));

            20     $this->assertEquals(1, count($order->getProducts()));
            $this->assertSame($product, $order->getProducts()[0]);
        }
}
```

See after the `\Phake::capture($order)` call, the `$order` variable contains the argument that was passed to the `OrderRepository` from your code.

This arguably reaches into the tested class quite a bit, but when you use Command / Query separation and London-Style TDD the only way to observe behaviour and state is mocking. I still think Phake is the best mocking library for PHP and the `capture` method is another good argument for it.

## 4.6 Using Mink in PHPUnit

[https://qafoo.com/blog/081\\_phpunit\\_mink\\_functional\\_tests.html](https://qafoo.com/blog/081_phpunit_mink_functional_tests.html)

*Benjamin Eberlei at 5. April, 2016*

Another day for a short PHPUnit trick. If you want to use PHPUnit to control a browser for functional or acceptance tests, then you can easily do this using the Mink library. Mink is well known from the Behat community to facilitate Behaviour-Driven Development (BDD), but it is a standalone library that can be used with PHPUnit just as easily.

This is more flexible than using dedicated browser abstractions such as Selenium directly from PHPUnit, because you can switch between different implementations or even run tests with multiple implementations using the same code base.

To start install Mink into your PHP project using Composer:

```
$ composer require behat/mink behat/mink-goutte-driver --dev
```

This will install Mink and the Guzzle/Goutte based Driver to crawl your site using a very simple cURL based browser abstraction.

Lets start using it for a simple PHPUnit test that verifies Wikipedia Search:

```
<?php
class WikipediaTest extends \PHPUnit_Framework_TestCase
4 {
    public function testSearch()
    {
        $baseUrl = 'https://en.wikipedia.org/wiki/Main_Page';
        $driver = new \Behat\Mink\Driver\GoutteDriver();
        $session = new \Behat\Mink\Session($driver);
        $session->start();
        $session->visit($baseUrl);
        $page = $session->getPage();
        $page->fillField('search', 'PHP');
        $page->pressButton('searchButton');
        $content = $session->getPage()->getContent();
        $this->assertContains('PHP: Hypertext Preprocessor', $content);
        $this->assertContains('Rasmus Lerdorf', $content);
    }
}
```

Setting up the Driver and Session over and over again can become quite complicated, lets introduce a reusable trait:

```
2 <?php
3
4 trait MinkSetup
5 {
6     private $minkBaseUrl;
7
8     private $minkSession;
9
10    /**
11     * @before
12     */
13    public function setupMinkSession()
14    {
15        $this->minkBaseUrl = isset($_SERVER['MINK_BASE_URL'])
16            ? $_SERVER['MINK_BASE_URL']
17            : 'http://localhost:8000';
18
19        $driver = new \Behat\Mink\Driver\GoutteDriver();
20        $this->minkSession= new \Behat\Mink\Session($driver);
21        $this->minkSession->start();
22    }
23
24    public function getCurrentPage()
25    {
26        return $this->minkSession->getPage();
27    }
28
29    public function getCurrentPageContent()
30    {
31        return $this->getCurrentPage()->getContent();
32    }
33
34    public function visit($url)
35    {
36        $this->minkSession->visit($this->minkBaseUrl . $url);
37    }
38 }
```

The `@before` annotation is relatively new, it makes sure that the annotated method is called during each test cases setup phase, whenever we use the `MinkSetup` trait in a test class.

This allows us to write the actual test in a much simpler way:

```
<?php
```

```

class WikipediaTest extends \PHPUnit_Framework_TestCase
{
    use MinkSetup;

    public function testSearch()
    {
        $this->visit('/');

        $page = $this->getCurrentPage();
        $page->fillField('search', 'PHP');
        $page->pressButton('searchButton');

        $content = $this->getCurrentPageContent();

        $this->assertContains('PHP: Hypertext Preprocessor', $content);
        $this->assertContains('Rasmus Lerdorf', $content);
    }
}

```

If you followed the `MinkSetup` implementation, you saw the `MINK_BASE_URL` environment variable. We can configure this from the `phpunit.xml` configuration:

```

<?xml version="1.0" ?>
<phpunit bootstrap="vendor/autoload.php">
    <php>
        <env name="MINK_BASE_URL">http://en.wikipedia.org/wiki</env>
    </php>
</phpunit>

```

You can improve this by adding more helper methods onto the `MinkSetup` trait, for example by closely following the possibilities that Mink provides inside of Behat (See `MinkContext`<sup>33</sup>).

<sup>33</sup><https://qa.fo/book-MinkContext>

## 4.7 Introduction To Page Objects

[https://qafoo.com/blog/089\\_introduction\\_to\\_page\\_objects.html](https://qafoo.com/blog/089_introduction_to_page_objects.html)

Manuel Pichler at 6. September, 2016

A while ago we wrote about writing acceptance tests (end-to-end tests) with Mink and PHPUnit (*See: Using Mink in PHPUnit*). While this is a great set of tools for various applications such tests tend to be susceptible to changes in the frontend. And the way they break is often hard to debug, too. Today I will introduce you to Page Objects<sup>34</sup> which can solve some of these problems.

The basic idea behind a Page Object is that you get an object oriented representation of your website. The Page Objects maps the HTML (or JSON) to an object oriented structure you can interact with and assert on. This is more initial work than writing tests with PHPUnit and Mink directly, but it can be worth the effort. I will introduce you to Page Objects by writing some simple tests for Tideways – our application performance monitoring platform<sup>35</sup>.

### 4.7.1 Groundwork

We will again use the awesome Mink<sup>36</sup> to simulate browsers and make it easy to interact with a website. Thus we are actually re-using the `FeatureTest` base class from the *Using Mink in PHPUnit* (*See: Using Mink in PHPUnit*) blog post. We have set up a repository<sup>37</sup> where you can take a full look at the working example and maybe even try it out yourself.

You'll need some tools to set this up – in the mentioned repository it is sufficient to execute `composer install`. Setting it up in a new project you'd execute something like:

```
| composer require --dev phpunit/phpunit behat/mink behat/mink-goutte-driver
```

The `FeatureTest` base class<sup>38</sup> handles the basic mink interaction and has already been discussed in the last blog post so that we can skip it here.

---

<sup>34</sup><http://martinfowler.com/bliki/PageObject.html>

<sup>35</sup><https://tideways.io>

<sup>36</sup><http://mink.behat.org/en/latest/>

<sup>37</sup><https://github.com/QafooLabs/PageObjects>

<sup>38</sup><https://qa.foo/book-FeatureTest>

### 4.7.2 A First Test

As mentioned we want to test Tideways<sup>39</sup> and Tideways requires you to login. Thus we start with a simple login test:

```

class LoginTest extends FeatureTest
{
    public function testLogInWithWrongPassword()
    {
        $page = (new Page\Login($this->session))->visit(Page\Login::PATH);

        $page->setUser(getenv('USER'));
        $page->setPassword('wrongPassword');
        $newPage = $page->login();

        $this->assertInstanceOf(Page\Login::class, $newPage);
    }

    // ...
}

```

This test already uses a page object by instantiating the class `Page\Login`. And by using this one it makes the test very easy to read. You instantiate the page, `visit()` it and then interact with it in an obvious way. We set username and password, and then call `login()`. Since we set a wrong password we expect to stay on the login page.

This already is the nice thing with page objects. The test are readable and this is something we want to optimize for, right?

One the other hand the logic must be implemented in the Page Object. By implementing it in a Page Object it is re-usable in other tests as you will see later. So let's take a look at this simple Page Object:

```

use Qafoo\Page;

class Login extends Page
{
    const PATH = '/login';

    public function setUser($user)
    {
        $this->find('input[name="_username"]')->setValue($user);
    }
}

```

<sup>39</sup><https://tideways.io>

```
15     public function setPassword($password)
        {
            $this->find('input[name="_password"]')->setValue($password);
        }

        public function login()
        {
            $this->find('input[name="_submit"]')->press();
20         return $this->createFromDocument();
        }
    }
}
```

Since we use Mink and implement some logic in the `Qafoo\Page` base class this still does not look that complex. What you should note is the fact that the method `setUser()` (and alike) hide the interaction with the DOM tree. If the name of those form fields change you'll have to change it in one single location. The methods `find()` and `visitPath()` can be found in the `Page` base class<sup>40</sup> and just abstract Mink a little bit and provide better error messages if something fails.

The `login()` method will execute a HTTP request to some page. If the login failed we will be redirected back to the login page (like in the test above), otherwise we expect to be redirected to the dashboard:

```
5     public function testSuccessfulLogin()
        {
            $page = (new Page\Login($this->session))->visit(Page\Login::PATH);

            $page->setUser(getenv('USER'));
            $page->setPassword(getenv('PASSWORD'));
            $newPage = $page->login();

10         $this->assertInstanceOf(Page\Dashboard::class, $newPage);
        }
}
```

We expect the user name and password to be set as environment variables since there are no public logins for Tideways. If you want to run the tests yourself, just create an account and provide them like mentioned in the README<sup>41</sup>.

There is one "magic" method left in the page object shown before – the method `createFromDocument()`. The method maps the path of the last request back to `Page`

---

<sup>40</sup><https://qa.fo/book-Page>

<sup>41</sup><https://github.com/QafooLabs/PageObjects>

Object. Something like the router in about every framework would do, but we map to a Page Object instead of a controller. This method will get more complex for complex routes but it helps us to make assertions on the resulting page.

### 4.7.3 Refactoring The Frontend

We recently migrated the dashboard from being plain HTML rendered using Twig templates on the server into a React.js component. What happens to our page objects in this case? Let's take a look at our dashboard tests first:

```
class DashboardTest extends FeatureTest
{
    use Helper\User;

    // ...

    public function testHasDemoOrganization()
    {
        $this->login();

        $page = (new Page\Dashboard($this->session))->visit(Page\Dashboard::PATH);

        $organizations = $page->getOrganizations();
        $this->assertArrayHasKey('demo', $organizations);
        return $organizations['demo'];
    }

    // ...
}
```

This test again makes assertions on a page object – now `Page\Dashboard` which can be instantiated after logging in successfully. The test itself does not reveal in any way if we are asserting on HTML or some JSON data. It is simple and asserts that we find the `demo` organization in the current users account (which you might need to enable<sup>42</sup>).

So let's take a look at the dashboard Page Object, where the magic happens:

```
class Dashboard extends Page
{
    const PATH = '/dashboard';
```

<sup>42</sup><https://app.tideways.io/settings/profile>

```

5   public function getOrganizations ()
    {
        $dataElement = $this->find ('[data-dashboard]');
        $dataUrl = $dataElement->getAttribute ('data-dashboard');

10        $data = json_decode($this->visitPath($dataUrl)->getContent());
        \PHPUnit_Framework_Assert::assertNotNull($data, "Failed to parse JSON
            response");

        $organizations = array ();
        foreach ($data->organizations as $organization) {
15            $organizations[$organization->name] = new Dashboard\Organization(
                $organization, $data->applications);
        }

        return $organizations;
    }
20 }

```

We are currently migrating (step by step) from jQuery modules to React.js components and are still using data attributes to trigger loading React.js components in the UI. Instead of asserting on the HTML, what we would have done when still rendering the dashboard on the server side, we check for such a data attribute and load the data from the server. For each organization found on the page we then return another object which represents a partial (organization) on the page.

Using this object oriented abstraction of the page allows us to transparently switch between plain HTML rendering and React components while the test will look just like before. The only thing changed is the page object, but this one can still power many tests which can make the effort worth it. On top of those Organization partials we can then execute additional assertions:

```

/**
 * @depends testHasDemoOrganization
 */
public function testMonthlyRequestLimitReached (Page\Dashboard\Organization
    $organization)
5   {
        $this->assertFalse ($organization->getMonthlyRequestLimitReached ());
    }

/**
10  * @depends testHasDemoOrganization
    */
public function testHasDemoApplications (Page\Dashboard\Organization $organization)

```

```
15 {  
    $this->assertCount(3, $organization->getApplications());  
}  
// ...
```

#### 4.7.4 Problems With Page Objects

As you probably can guess providing a full abstraction for your frontend will take some time to write. Those page objects can also get slightly more complex, so that you might even feel like testing them at some point.

Since end-to-end tests also will always be slow-ish (compared to unit tests) we advise to only write those tests for critical parts of your application. The tests will execute full HTTP requests which take time – and nobody runs a test suite which takes multiple hours to execute.

Also remember that, like with any integration test, you probably need some means to setup and reset the environment the tests run on. In this simple example we run the test on the live system and assume that the user has the demo organization enabled. In a real-world scenario you'd boot up your application (provision a vagrant box or docker container), reset the database, mock some services, prime the database and run the tests against such a reproducible environment. This takes more effort to implement, again.

While the tests are immune to UI changes this way (as long as the same data is still available) they are not immune to workflow changes. If your team adds another step to a checkout process, for example, the corresponding Page Object tests will still fail and you'll have to adapt them.

#### 4.7.5 Conclusion

Page Objects can be a good approach to write mid-term stable end-to-end tests even for complex applications. By investing more time in your tests you can get very readable tests which are easy to adapt to most UI changes.

## 4.8 Database Tests With PHPUnit

[https://qafoo.com/blog/090\\_database\\_tests\\_with\\_phpunit.html](https://qafoo.com/blog/090_database_tests_with_phpunit.html)

*Tobias Schlitt at 4. October, 2016*

Most of us do not use PHPUnit solely for Unit Tests but we also write Integration or Acceptance Tests with PHPUnit. One very common question then is how to interact with the database correctly in those tests. Let me show you the different options and their trade offs...

There are multiple aspects of database tests where our decision has impact on test atomicity and test runtime. All decisions boil down to: More test atomicity leads to longer test runs, and we can buy test speed by omitting test atomicity. One might be tempted to immediately favour test atomicity but fast tests are a feature which is often underappreciated. If your tests are slow (more then 30 seconds or even taking minutes) developers will not run those tests before each commit any more – they feel it just takes them too much time. Your tests will still be run on your Continuous Integration server but even when your tests fail it is annoying to debug them, because you have to wait so long until the failing test is reached.

### 4.8.1 Removing Data versus Schema Reset

You have a schema of your database laying around somewhere, right? It should even be somewhere in your source code repository to be able to initialize a full build of your application. We prefer using something like DBDeploy<sup>43</sup> to maintain and document changes to our schema but the solution does not really matter. With a schema you can drop the entire database to clean it up before a test and re-apply the schema. This works well with MySQL or SQLite but takes a lot of time with systems like Oracle.

The other idea is to reset the tables your test modified. Since this does not do any schema changes but just removes data it is usually faster. The most common implementation is to remove data manually in a `tearDown()` method in all tables you know the test modified. The problem now is that your application will change. This means that the tables data is written into will also change. At some point you'll forget removing some data in some table and this will cause a side effect on another test

---

<sup>43</sup><http://dbdeploy.com/>

– a horribly hard to debug side effect. And this only works for new data, it will not revert changed data.

Some people are adding something like a `changed` column to each table which automatically receives the time of the last change to a row. If you are doing this you can execute something like `DELETE FROM $table WHERE changed >= $testStartTime` on each table. But then again, this only works for data added in a test and would remove all rows which were changed in a test.

### Summary

Resetting the full schema is the cleanest approach, but also takes the most time. Resetting a selected number of tables is faster, but also more cumbersome and error-prone. If you have some kind of change management you might be able to use that. Especially for large and complex schemas (where you would prefer a cleaner approach) resetting the full schema often takes far too much time.

## 4.8.2 Point of Data Reset

Now we know different ways to reset the data in our database, but *when* should we do the reset. There are basically three options:

- Before each test
- Before each test class
- Once before the whole test run

### Before Each Test

The `setUp()` method in `PHPUnit_Framework_TestCase` is the right point to run initialisations before each test and we can initialize the database here if our test wants to access the database and clean it up later in the `tearDown()` method. This way we preserve the atomicity of our tests. But since a single cleanup can take up to some seconds this is only feasible for simple and small schemas.

### Before Each Test Class

PHPUnit offers the two (static) methods `setUpBeforeClass()` and `tearDownAfterClass()` which are run once before any test in the test case class is executed and respectively after all tests in the test case class have been executed.

These methods allow us to reset or setup the database once for entire test case. If you do something like this you must remember that the tests in this test case depend on each other. The first test might create some row and the second test then accesses this row and uses it for something. I usually add `@depends` annotations on the tests to make this extra obvious. Also the second test will not work without the first one anyways. You can see an example for such a test case in our TimePlanner demo project<sup>44</sup>.

The initialisation of the database happens in the mentioned methods in a base class for our integration tests<sup>45</sup>. The methods look slightly more complicated because the tests are run against three different databases (MySQL, SQLite and CouchDB) and all are resetting the entire schema using Doctrine.

With this method you cannot easily reorder your tests in a test case anymore, but you can still run your tests cases in any order, since they are still supposed to be independent. If you want to debug a single test it still means that you have to run all tests which it depends on before it – which can be annoying.

#### **Before the whole test run**

PHPUnit allows you to specify a bootstrap file, either on the CLI or in the `phpunit.xml`. Nowadays this is mostly used to just specify the `vendor/autoload.php` from composer, but you can do more in this file. For example you can initialize your database schema. You can also do this in some build step which is always executed before each test run or somewhere similar.

This allows you to initialize the database once before all tests. This approach is supposed to be the fastest because it minimizes database interaction and thus IO. But there are some drawbacks:

- If you don't specify the order of your test cases manually the order they are picked up by PHPUnit might changed depending on your OS or file system. This can lead to very weird and hard to debug test failures.
- State is leaked between all your tests across your complete test suite. This again can lead to situations which are very hard to debug.
- There will be situations where you can debug one test only when running the

---

<sup>44</sup><https://qa.fo/book-PublicHolidayGatewayTest>

<sup>45</sup><https://qa.fo/book-IntegrationTest>

entire test suite because it depends on the state of some random test you might not even be aware of.

### Summary

Except for very small and simple projects it is usually best to initialize the database before each test class. This seems to be the best compromise between test stability and speed for most projects – your mileage might vary.

#### 4.8.3 Mocking the Database Away

Another entirely different option which comes to mind is mocking the entire interaction with your database. But such complex mocks are often error prone and a lot of effort to implement and maintain. If your database interaction is very simple it might work for you. In any project with non trivial queries (aggregations, joins, ...) you probably do not want to walk this path.

In a MySQL project you can think about using a SQLite in-memory database, but at some point you'll discover differences between these two database management systems. And this could mean production bugs which will not be discovered during testing.

#### 4.8.4 Conclusion

In general you should try to reduce access to global state (like a database) as much as possible in your tests. But at some point you'll want to test the interaction with your database – in this case you must decide for a way which works best for your project. In most projects the best way for us seems to reset the full schema before each test case which interacts with the database.

## 4.9 Database Fixture Setup in PHPUnit

[https://qafoo.com/blog/091\\_database\\_fixture\\_setup\\_in\\_phpunit.html](https://qafoo.com/blog/091_database_fixture_setup_in_phpunit.html)

*Tobias Schlitt at 15. November, 2016*

We already discussed when and how to reset a database between tests (*See: Database Tests With PHPUnit*). But on top of that you often need some basic or more complex data to run the tests against. We discuss different variants and their respective trade-offs in this post...

### 4.9.1 Dump & Insert Live Data

The simplest approach - we still see this in the wild - for schema and test data management is to dump and replay live data for testing purposes. The benefit here is that you only have to adapt data in one place when the schema changes and it is really easy to set up. But there are a couple of drawbacks to this approach:

- **Large live data sets**

The live data usually is "large" – at least a couple of hundred MB (while database are usually not considered large before reaching TB). It takes a lot of time to reset a database even with such a dump. You do not want to spend such an amount of time before each test, test suite or even before each test run.

- **Boundary value test cases**

There is often a need for special test cases, like strange characters or other type of boundary values in your tests. Those data sets might not exist in the live data so you end up creating this data in your tests on top of the data dump anyways.

- **Data privacy**

Depending on your use case and business your developers should not have access to all live data. There might be sensitive information which should be locked down on your servers. In this case a live SQL dump is no option. Especially in Germany we might be required by law to lock certain data away from certain people.

- **Changing data on the live system**

It is obvious that data on the live system changes over time. This can make it

hard to author tests that are stable and reproducible. In addition, tests might become less meaningful if you need to craft them in a way that they can cope with changing live data.

### Modified Live Data Set

To avoid the problems mentioned above a next step usually is a modified SQL file, which is smaller and does contain sensible test data, like boundary values or stable data without any sensitive information for reproducible tests.

The problem which arises now is that you have to adapt two files when you change the data structure. And the schema and properties of the data *will* divert over time – no matter how careful you are. In the end this approach is always hard to maintain, so what can be done?

First we suggest you implement some kind of sensible schema management like DBDeploy<sup>46</sup>, or even Doctrine Migrations if this works for your use-case. In this post we want to focus on the data / fixture management, though.

### 4.9.2 Base Data

Most applications require a set of base data, like default or admin users, some groups and permissions or something similar. This will be same in all installations (production, staging, development, testing, ...). Depending on your schema management solution you will be able to insert this data during schema initialization. With DBDeploy you can just add some `INSERT` statements with Doctrine you could use Doctrine data fixtures<sup>47</sup> (or even the Symfony bundle<sup>48</sup>).

### 4.9.3 Test Data

The more important thing is the test data. Inserting sensible test data is tightly coupled to your test database resetting strategy (*See: Database Tests With PHPUnit*) which we discussed earlier.

---

<sup>46</sup><http://dbdeploy.com/>

<sup>47</sup><https://github.com/doctrine/data-fixtures>

<sup>48</sup><https://qa.fo/book-index>

**Before Each Test**

When you reset your database before each test you also want to insert the test data right inside the test. This is very obvious, makes tests easy to read and understand. In theory this is clearly the best solution. But as mentioned in the referenced blog post this will cost a lot of time to execute and likely be to slow for any non-trivial application.

**Before Each Test Class**

When you reset the database before each test case you can create the data throughout the test case. An simple common CRUD example could be the following tests:

1. Create a new data set
2. Load data set
3. Fail loading a non-existent data set
4. Load listing (with one item)
5. Delete data set
6. Fail deleting non-existent data set

Those tests depend on each other which should be indicated by using `@depends` annotations like in this example from our demo project<sup>49</sup>.

This approach is still very clean and from reading the test case you can understand the full context. Another developer will still be able to understand what is going on. This is a lot harder when the data is inserted in another place, since you'll always have to look in multiple places to get the full image. And new developers might not yet know all the places to look at. Tests which provide all the context you need to know in one file are very helpful for everybody.

This strategy will get more complex if you have dependent data – like tests for user permissions, which also require users and groups to exist. You could either use custom helpers or tools like Alice<sup>50</sup> for this.

**Before the whole test run**

If you decided to only reset the database once before all tests are run you usually need a more complete data fixture. You will want to fill all tables with some basic

---

<sup>49</sup><https://qa.fo/book-PublicHolidayGatewayTest>

<sup>50</sup><https://github.com/nelmio/alice>

data you can work with. Especially in such a case tools like Alice<sup>51</sup> are very useful. By relying on Faker<sup>52</sup> you even get sensible looking data without too much custom work.

#### 4.9.4 Conclusion

The way you initialise your database fixtures depends on your test schema management. We suggest to reset schemas at the begin of each test case and creating the data right in the test case. This proved to be a good compromise between speed, test (case) atomicity and test readability. Tools like Alice<sup>53</sup> and Faker<sup>54</sup> can ease the process of creating data a lot.

---

<sup>51</sup><https://github.com/nelmio/alice>

<sup>52</sup><https://github.com/fzaninotto/Faker>

<sup>53</sup><https://github.com/nelmio/alice>

<sup>54</sup><https://github.com/fzaninotto/Faker>

## 4.10 Using Traits With PHPUnit

[https://qafoo.com/blog/092\\_using\\_traits\\_with\\_phpunit.html](https://qafoo.com/blog/092_using_traits_with_phpunit.html)

Kore Nordmann at 29. November, 2016

As we already wrote that "Code Reuse By Inheritance" (*See: Code Reuse By Inheritance*) has lots of problems and we consider it a code smell. You should always aim to use Dependency Injection, most likely Constructor Injection. But with test cases in PHPUnit we cannot do this because we have no control about how and when our test cases are created. There are a similar problem in other frameworks, like we discussed in "Object Lifecycle Control" (*See: Object lifecycle control*). We also blogged about traits as a Code Smell (*See: Utilize Dynamic Dispatch*), but let me show and explain why they might be fine to use in your test cases.

So in PHPUnit it is common to reuse code by inheritance. Even by now we often create some base class providing common functionality. This works fine and is OK if this really defines an "is a"-relationship. Let's continue with an example from our Page Object test repository<sup>55</sup>. The patterns described here are usually not required for Unit Tests (which you should always also write) but mostly for integration or functional tests.

### 4.10.1 An Example

The example repository<sup>56</sup> implements functional tests for a website using the Page Object pattern (*See: Introduction To Page Objects*). This means that the tests access a website and assert on its contents, try to fill forms, submit them and click links. Such tests are a useful part of your test mix (*See: Finding the right Test-Mix*), but should never be your only tests.

Let's see the options for code reuse we employ in this little test project – and the reasoning behind it. We start with a simple test case:

```
class LoginTest extends FeatureTest
{
    public function testLoginWithWrongPassword()
    {
        $page = (new Page\Login($this->session))->visit(Page\Login::PATH);
    }
}
```

<sup>55</sup><https://github.com/QafooLabs/PageObjects>

<sup>56</sup><https://github.com/QafooLabs/PageObjects>

```

    $page->setUser ( getenv ( 'USER' ) );
    $page->setPassword ( 'wrongPassword' );
    $newPage = $page->login ();
10
    $this->assertInstanceOf ( Page\Login::class , $newPage );
}
// ...
15 }

```

This test case extends from a base class `FeatureTest` to re-use its functionality. The base class uses PHPUnit's `setUp()` method to setup and start the mink test driver which will act as a browser to access the website. And it provides a default `tearDown()` method to reset the browser state again.

In this case the inheritance defines a clear "is a" relationship with the `FeatureTest` and it overwrites default methods in the PHPUnit stack which should be overwritten in any feature test. The `FeatureTest` again extends an `IntegrationTest` which is empty in this example but normally would provide access to the application stack to reset a database, access random services or something similar. We just do not need anything like this in this little test project. Since functional tests can be considered a superset of integration tests this is fine again and provides common functionality which belongs to all tests of this type.

### 4.10.2 Traits

Let's take a look at a slightly more complex test case now:

```

class DashboardTest extends FeatureTest
{
    use Helper\User;
5
    // ...

    public function testHasDemoOrganization ()
    {
10
        $this->login ();

        $page = ( new Page\Dashboard ( $this->session ) )->visit ( Page\Dashboard::PATH );

        $organizations = $page->getOrganizations ();
        $this->assertArrayHasKey ( 'demo', $organizations );

```

```
15 |         return $organizations['demo'];
    |     }
    |     // ...
    | }
```

In this case we use the `Helper\User` trait to include some functionality – it provides the `login()` method which is used in the test `testHasDemoOrganization()`. Not every feature test might need this aspect and in "normal" software you would provide such helpers through constructor injection. But since we do not have any control on the test case creation we include the code using a trait.

The trait enables code reuse – we can use it in any test case which requires login. The trait extracts this concern and we do not clutter every test case requiring login with this kind of code.

### 4.10.3 Whats The Difference?

The trait helps us in this example, the code looks clean, so you might want to ask: **Why would traits ever be considered a code smell?**

One of the most important reasons is that in a test case there probably won't be a reason to change a dependency without adapting the code (Open Closed Principle). In other words: There is no reason for dynamic dispatch.

A trait establishes a dependency to another class which is defined by the name of the trait (instead of an instance of some class which could be a subtype). There is no easy way to change the actually used implementation from the outside. If you include a `LoggerTrait` there is no way to change the used `LoggerTrait`, during tests or when the requirements change, without changing code. Traits establish a `static` dependency which is hard to mock and hard to replace during runtime or configuration.

But we will never mock our test cases, right? And if the use cases change we will change the test cases. This can happen a lot as compared to unit tests.

Especially (Open Source) libraries and extensible software commonly has the requirement that people should be able to change the behaviour without changing the code. Most likely because they do not have direct access to the code or it would have side effects to other usages of the code. But nobody uses your test cases in such a way, thus you are "allowed" to sin in here – at least a little bit.

And there are no other options. This, generally, can be another reason to use traits. Traits are often an option when refactoring legacy software to temporarily use common code before we can migrate to sensible dependency injection. Being a code smell they even help knowing about places which are still not done.

#### 4.10.4 Summary

In tests traits can be a great tool to reuse common code, while we still consider traits a code smell in almost every other case.

## 4.11 Outside-In Testing and the Adapter and Facade Patterns

[https://qafoo.com/blog/087\\_outside\\_in\\_testing\\_adapter\\_pattern.html](https://qafoo.com/blog/087_outside_in_testing_adapter_pattern.html)

Benjamin Eberlei at 5. July, 2016

We at Qafoo are big fans of outside-in testing as described in the book "Growing Object-Oriented Software, Guided by Tests"<sup>57</sup> (Steve Freeman and Nat Pryce). As part of our workshops on Test-Driven Development we explain to our customers<sup>58</sup> how testing from the outside-in can help find the right test-mix.

The technique puts a focus on test-driven-development, but instead of the traditional approach starts at the acceptance test level. The first test for a feature is an acceptance test and only then the feature is implemented from the outside classes first (UI and controllers), towards the inner classes (model, infrastructure). Mocks are used to describe roles of collaborators when starting to write tests for the outside classes. When a class is tested, the roles described by interfaces (See: *Abstract Classes vs. Interfaces*) are implemented and the testing cycle starts again with mocks for the collaborators.

Outside-In testing leads to interfaces that are written from what is useful for the client object using them, in contrast to objects that are composed of collaborators that already exist. Because at some point we have to interact with objects that exist already, we will need three techniques to link those newly created interfaces/roles to existing code in our project:

1. The adapter pattern, mostly for third-party code
2. The facade pattern, mostly to structure your own code into layers
3. Continuous refactoring of all the interfaces and implementations

This blog post will focus on the facade adapter pattern as one of the most important ingredient to testable code.

Even if very well tested, APIs of third party libraries or frameworks are usually not suited for projects using any form of automated testing, either because using them directly requires setting up external resources (I/O) or mocking them is complex, maybe even impossible.

### warning

---

<sup>57</sup><http://www.growing-object-oriented-software.com/>

<sup>58</sup>[https://qafoo.com/blog/051\\_testing\\_sweet\\_spot.html](https://qafoo.com/blog/051_testing_sweet_spot.html)

Side Fact: This even affects libraries that put a focus on testing, such as Doctrine2, Symfony2 or Zend Framework. The reason is that libraries often provide static APIs, fluent APIs, facades with too many public methods or complex object interactions with law-of-demeter violations.

Take a common use-case, importing data from a remote source into your own database. In a Symfony2 project with Doctrine and Guzzle as supporting libraries, the feature could easily be implemented as a console command, using only the third-party code and some glue code of our own:

```
<?php
namespace Acme\ProductBundle\Command;

class ImportProductCommand extends ContainerAwareCommand
5 {
    protected function configure() { /* omitted */ }

    protected function execute(InputInterface $input, OutputInterface $output)
10 {
        $client = $this->getContainer()->get('guzzle.http_client');
        $entityManager = $this->getContainer()->get('doctrine.orm.
            default_entity_manager');

        $request = $client->get('http://remote.source/products.xml');
        $response = $request->send();
15
        $products = $this->parseResponse($response);

        foreach ($products as $product) {
            $entityManager->persist($product);
20        }

        $entityManager->flush();
    }

    protected function parseResponse($response) { /** omitted */ }
25 }
}
```

Looks simple, but in the real world, you can safely assume there is quite some complexity in `parseResponse` and possibly even more logic inside the loop over all the `$products`.

In this scenario, the code is completely untestable, just by combining the APIs of Guzzle, Doctrine and Symfony.

Lets take a different approach, starting with the highest level description of the feature as a test written in Gherkin for Behat:

```

Scenario: Import Products
  Given a remote service with products:
    | name | description          | price
    | A    | Nice and shiny    | 100
5   | B    | Rusty, but cheap    | 10
  When I import the products
  Then I should see product "A" in my product listing
  Then I should see product "B" in my product listing

```

This will be our acceptance (or end-to-end) test that will use as many external systems and I/O as possible. We might need to mock data from the remote product service, but when possible we should use the real data. This test will require the UI to be done as well as the necessary database and remote system APIs and therefore will fail until we have implemented all the code necessary. So lets focus on the code directly with our first unit-test and lets imagine how we want the code to be used:

```

2 <?php
  namespace Catalog;

  class ImportProductTest extends \PHPUnit_Framework_TestCase
  {
7     public function testImportTwoProducts ()
        {
            $remoteCatalog = \Phake::mock('Catalog\RemoteCatalog');
            $productGateway = \Phake::mock('Catalog\ProductGateway');

12         $productA = $this->createSampleProduct();
            $productB = $this->createSampleProduct();

            \Phake::when($remoteCatalog)->fetch()->thenReturn(array($productA,
                $productB));

17         $importer = new ProductImporter($productGateway);
            $importer->import($remoteCatalog);

            \Phake::verify($productGateway)->store($productA);
            \Phake::verify($productGateway)->store($productB);
22     }
  }

```

When we fetch 2 products from the \$remoteCatalog then those products should

be passed to the `$productGateway`, our storage system. To describe the use-case in this simple way, we have abstracted `RemoteCatalog` and `ProductGateway`, which will act as facades used in our `ProductImporter`. The implementation is very simple:

```
<?php
namespace Catalog;

5 class ProductImporter
{
    /**
     * @var ProductGateway
     */
10    private $productGateway;

    public function import(RemoteCatalog $remoteCatalog)
    {
15        $products = $remoteCatalog->fetch();

        foreach ($products as $product) {
            $this->productGateway->store($product);
        }
20    }
}
```

We haven't seen the Guzzle and Doctrine code here again, instead we have written code that is entirely written in our business domain and uses concepts from this domain.

Lets move to the implementation of the `RemoteCatalog` starting with a test:

```
<?php
namespace Catalog\RemoteCatalog;

5 class HttpCatalogTest extends \PHPUnit_Framework_TestCase
{
    public function testGetAndParseData()
    {
10        $client = \Phake::mock('Catalog\Adapter\HttpClient');
        $parser = \Phake::mock('Catalog\RemoteCatalog\Parser');
        $url = 'http://remote.local';

        $productA = $this->createSampleProduct();

15        \Phake::when($client)->get($url)->thenReturn('<xml>');
        \Phake::when($parser)->parse('<xml>')->thenReturn(array($productA));
    }
}
```

```
20     $catalog = new HttpCatalog($url, $client, $parser);
        $data = $catalog->fetch();
    }
    $this->assertSame(array($productA), $data);
}
```

This test again perfectly clear explains how fetching a catalog with HTTP should work on a technical level. Notice how we choose a very simple API for the HTTP client, one that is fetching and \$url and retrieving the body as string. We don't need more.

```
2 <?php
  namespace Catalog\RemoteCatalog;

  use Catalog\RemoteCatalog;

  class HttpCatalog implements RemoteCatalog
  7  {
      private $url;
      private $client;
      private $parser;

      public function fetch()
      12  {
          $body = $this->client->fetch($this->url);
          return $this->parser->parse($body);
      }
  17 }
```

The HttpClient is a real adapter for us now, we want this very simple API and we are going to use Guzzle to implement it:

```
3 <?php
  namespace Catalog\Adapter\Guzzle;

  use Guzzle\Http\Client;

  class GuzzleHttpClient implements HttpClient
  8  {
      private $client;

      public function __construct(Client $client)
      {
          $this->client = $client;
      }
  13 }
```

```
18 public function fetch($url)
    {
        $request = $this->client->get($url);
        $response = $request->send();

        return $response->getBody(true);
    }
}
```

Implementing the Guzzle client we have reached a "leaf" of our object graph. It is important to see how only the Guzzle adapter actually uses Guzzle code. A complete solution would also require to handle the Guzzle Exceptions, but that is only a trivial task to introduce as well.

You can continue with this example and implement the `Parser` and the `ProductGateway` objects. The technique stays the same: Think in terms of what you want the API to look from the outside and invent collaborators that help you think about the problem. Then implement them until you get to the "leafs" of the object graph. Only the leafs should actually contain code to third party software.

Starting with new requirements of the system we will probably be able to reuse some of the code: The `HttpClient` interface is very useful in a good number of use-cases. The `HttpCatalog` can be used with specific `Parser` implementations to import many different product catalogs that have an HTTP interface. And if your customer uses some other protocols like FTP, BitTorrent or anything else then we are good as well. The `ProductGateway` will probably be used in various places to give us access to finding and storing products.

#### 4.11.1 Conclusion

Specification and testing of the requirements from the outside requires you to think about what you need first and not on what you have and how you might combine this. This helps to design simple and reusable objects.

But outside-in testing has additional benefits: When you don't test everything in your system with unit-tests (which is not uncommon) for any of the various reasons (time, prototyping, not core domain, ..) then you still have at least one acceptance test that verifies the complete feature and maybe some unit-tests for the tricky implementation details.

Overall we think making yourself familiar with outside-in testing is beneficial to designing testable and maintainable applications.

## 4.12 Behavior Driven Development

[https://qafoo.com/blog/036\\_behavior\\_driven\\_development.html](https://qafoo.com/blog/036_behavior_driven_development.html)

*Tobias Schlitt at 8. March, 2013*

While unit, integration and system tests - especially combined with the methodology of Test Driven Development (TDD) - are great ways to push the technical correctness of an application forward, they miss out one important aspect: the customer. None of these methods verify that developers actually implement what the customer desires. Behavior Driven Development<sup>59</sup> (BDD) can help to bridge this gap.

The methodology of BDD is actually derived from TDD. Instead of writing a test for a single code unit upfront, an acceptance test for a certain behavior is formulated. More important, the test case is formulated in the ubiquitous language of the project, allowing non-technical persons to read it as the specification of the desired feature.

In the following I want to outline the methodology of BDD from a developer's perspective.

### 4.12.1 Example

So much for theory. In practice I prefer to explain things on the basis of examples. So here is an acceptance test:

```
Feature: The existing CFPs should be listed
  As a visitor I want to get a listing of all existing CFPs.

  Scenario: If there are no CFPs, I see an empty list.
    Given there are no CFPs
    When I view "CFP Listing"
    Then I see 0 CFPs listed
```

The snippet above shows an acceptance test formulated in a language that is based upon the Gherkin framework for Domain Specific Languages (DSLs). The first lines are narrative, documenting the feature to be tested. They do not fulfil any technical purpose, think of it as if it was a class level doc block.

What follows is the definition of a specific scenario within the feature. A scenario is basically a user story, as you probably recognize it from your agile development mode. Again, the introductory sentence is documentation and not processed.

<sup>59</sup>[https://en.wikipedia.org/wiki/Behavior-driven\\_development](https://en.wikipedia.org/wiki/Behavior-driven_development)

The block following afterwards is the actual test case. These three sentences are on the one hand human readable. On the other hand, they are structured so simple and generic that they can be processed by a computer with very little effort. The first word on every line is a keyword from the Gherkin language, while *Given* indicates a precondition, the *When* part indicates the test stimulus and *Then* formulates an expectation.

The wonderful thing about this specification is that your customer can actually read and verify it. You can write down what you, as a developer, understood how your customer wants the system to behave. You can send the specification to her/him and ask her/him if this is exactly what she/he desires. If she/he agrees, you can start implementing right away and verify your progress against the specification, as it is executable through a BDD test tool.

Possibly, if you have a technically skilled customer and train him quite a bit, he will probably even be able to adjust scenarios or even write some on his own.

### 4.12.2 Behat

The Gherkin example from above is almost useless without a proper tool to execute it. Of course, there is no tool that can do it right away, because no software can understand the project specific sentences. However, frameworks exist that offer you a basis. For the PHP world, the toolkit of choice is called Behat<sup>60</sup>.

Behat provides you with the basic infrastructure for BDD and enables you to easily work with custom sentences. Specifically, you create a so-called *FeatureContext* for your test cases that contain a method for every sentence in your projects ubiquitous language. For example:

```
class ListingFeatureContext extends BehatContext
{
    // ...

    /**
     * @Given /^there are no CFPs$/
     */
    public function thereAreNoCfps ()
    {
        $this->cleanupDatabase ();
    }
}
```

<sup>60</sup><http://behat.org/>

```
}  
}
```

The extract from the `ListingFeatureContext` above shows the method that reacts to the sentence `Given there are no CFPs`. Using an annotation, Behat connects the method to the sentence. Whenever it discovers that sentence in a test scenario, it will execute the method.

### 4.12.3 Rationale

Of course, the examples shown above are only very rudimentary, missing e.g. variables and other advanced features. However, they should have explained what BDD is all about: Communication. Especially in teams which follow the Domain Driven Design (DDD) approach, a ubiquitous language for the project domain is already practiced. Toolkits such as Behat provide you with the environment to express expectations in this language and make these executable.

The two essential ideas behind this are a) to ease communication with the client and b) to bridge the gap between (important!) tests for technical correctness and business expectations.

### 4.12.4 Conclusion

BDD is an interesting approach that can work especially well for projects that build on extensive business logic and such that follow Domain Driven Design. Besides that, Behat can become a tool of choice for acceptance tests also in respect to its integration with Mink<sup>61</sup> and Symfony2<sup>62</sup>.

What are your experiences with BDD and with Behat?

---

<sup>61</sup><http://extensions.behat.org/mink/>

<sup>62</sup><http://extensions.behat.org/symfony2/>

## 4.13 Code Coverage with Behat

[https://qafoo.com/blog/040\\_code\\_coverage\\_with\\_behat.html](https://qafoo.com/blog/040_code_coverage_with_behat.html)

*Manuel Pichler at 3. April, 2013*

There is generally no point in having code coverage for Behat test cases because of their nature: The purpose of an acceptance test is to assert a certain behavior of an application, not to technically test a piece of code. Therefore, there is no point in checking for uncovered code pieces in order to write a Behat test for it.

That said, there is still a scenario where you want to peek at code coverage of Behat tests: When creating them as wide-coverage tests before starting to refactor legacy code. Behat in combination with Mink provides you with a great tool for such tests.

Before you can start with refactoring legacy code you need tests to ensure that you don't break working functionality. Web acceptance tests on basis of Behat and Mink are a great tool to realize these. But how can you detect if the code you are about to refactor is touched by at least one test? Code coverage can be of assistance there.

### 4.13.1 Preparation

Since Behat does not ship with code coverage (for very good reason), you need some hand work to get that done, but not much. In order to get started, you need to install the `PHP_CodeCoverage` library<sup>63</sup> and `phpcov`<sup>64</sup>, most probably via PEAR using:

```
$ pear config-set auto_discover 1
$ pear install -a pear.phpunit.de/phpcov
```

### 4.13.2 Collecting Code Coverage

Since the Behat tests stimulate your application through external calls, it is not possible to generate code coverage right from the test code. Instead, you need to trigger the code coverage collection from your application code:

<sup>63</sup><https://github.com/sebastianbergmann/php-code-coverage>

<sup>64</sup><https://github.com/sebastianbergmann/phpcov>

```
<?php
3 // ... forbid production access here ...

$calculateCoverage = file_exists("/tmp/generate-behat-coverage");

8 if ($calculateCoverage) {
    require 'PHP/CodeCoverage/Autoload.php';

    $filter = new PHP_CodeCoverage_Filter();
    $filter->addDirectoryToBlacklist(__DIR__ . "/../vendor");
    $filter->addDirectoryToWhitelist(__DIR__ . "/../src");
13
    $coverage = new PHP_CodeCoverage(null, $filter);
    $coverage->start('Behat Test');
}

18 // ... run your application here ...

if ($calculateCoverage) {
    $coverage->stop();

23
    $writer = new PHP_CodeCoverage_Report_PHP;
    $writer->process($coverage, __DIR__ . "../log/behat-coverage/" . microtime(
        true) . ".cov");
}
```

At first the code detects if code coverage information should be gathered by checking if the file `/tmp/generate-behat-coverage` exists. You can touch and delete that one manually or from your test setup.

The next code block loads and initializes the code coverage collection, creates a filter for 3rd party code and starts the code coverage collection. After that, the comment indicates to run your application, which might e.g. be a Symfony2 kernel handling call.

The final lines write the code coverage information into a file for further processing. It will create a dedicated file for each request, where these files then need to be merged later.

### 4.13.3 Running Tests

With the shown code in place, you can trigger a Behat test run with code coverage using the following Ant code, for example:

```
5 <target name="behat-coverage" depends="clean, initialize">
  <delete dir="${commons:logsdirectory}/behat-coverage" />
  <mkdir dir="${commons:logsdirectory}/behat-coverage" />

  <touch file="/tmp/generate-behat-coverage" />
  <antcall target="behat" />
  <delete file="/tmp/generate-behat-coverage" />

  <exec executable="phpcov" failonerror="false" dir="${basedir}">
  <arg value="--merge" />
  <arg value="--html" />
  <arg value="${commons:logsdirectory}/../coverage/behat" />
  <arg value="${commons:logsdirectory}/behat-coverage" />
  </exec>
15 </target>
```

The Ant target first cleans up code coverage from previous runs. It then touches the file that indicates to the application to run code coverage, executes Behat and removes the trigger file again. Then the `phpcov` utility is executed to merge the results of all requests into a single coverage report and generate HTML from it.

#### 4.13.4 Conclusion

Code coverage is completely out of scope for acceptance tests. However, if you abuse Behat to create wide-coverage tests before refactoring, it might be of help to you to see what is still missing before you start hacking.

## 4.14 Testing Micro Services

[https://qafoo.com/blog/071\\_testing\\_micro\\_services.html](https://qafoo.com/blog/071_testing_micro_services.html)

Tobias Schlitt at 16. September, 2014

I recently had a short exchange with Ole Michaelis on Twitter about how to end-to-end test micro services<sup>65</sup>. Since I didn't have time to make my whole case, Ole suggested that I blog about this, which I'm happily doing now.

The idea behind micro service architecture was discussed by Martin Fowler<sup>66</sup> in a nice to read blog post, so I won't jump on that topic in detail here.

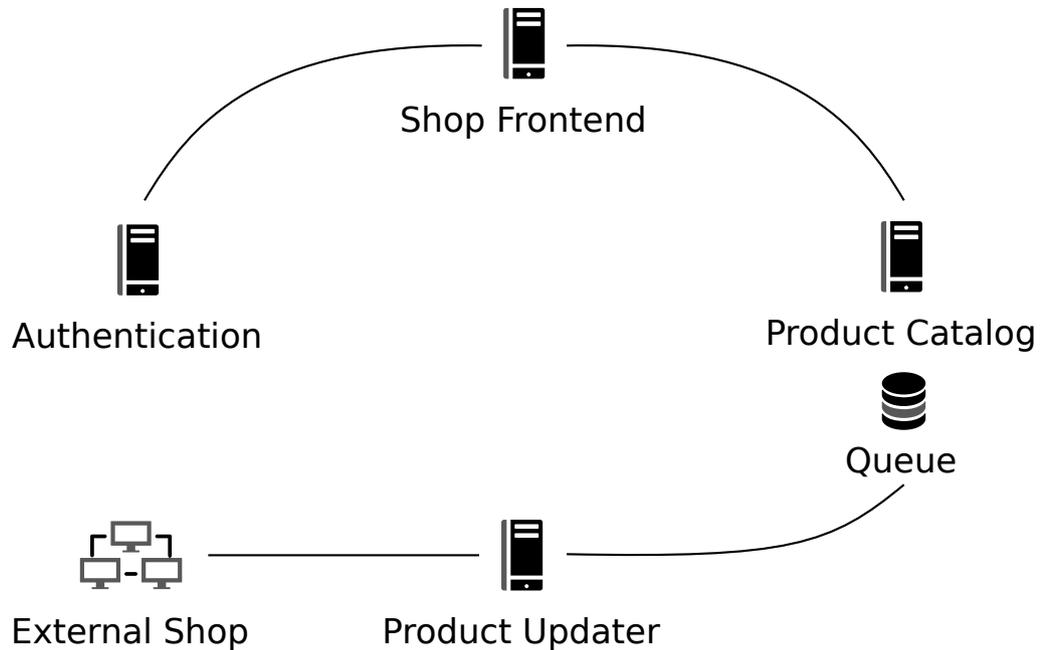
At Qafoo we run a flavor of what would today be called a micro service architecture for Bepado<sup>67</sup>. I sketched a simplified fraction of that below. While the frontend is roughly a classical web application, the other shown components are modelled as independently deployable, slim web services. The *Product Updater* service is responsible for fetching product updates from remote shops. It validates the incoming data and queues the product for indexing. Indexing is then performed in a different service, the *Product Catalog*.

---

<sup>65</sup><https://twitter.com/CodeStars/status/501336599419187200>

<sup>66</sup><http://martinfowler.com/articles/microservices.html>

<sup>67</sup><http://bepado.com>



While this shows only a small part of the full architecture, typical problems with system-testing such environments become evident. Sure, you can unit test classes inside each micro service and apply classical system tests to each service in isolation. But, how would you perform end-to-end system tests for the full application stack?

Trying to apply the classical process of "arrange, act, assert" did not work because of two reasons: a) there is a good portion of asynchronicity (queuing) between services which prevents from reasoning about the point in time when a certain state of the system can be asserted. b) setting up the fixture for the full system to trigger a certain business behavior in isolation would be quite complex task.

A fix for a) could be to mock out the queues and have synchronous implementations for the test setup. Still, the accumulated service communication overhead (HTTP) makes timing hard. Again, you could mock out the service layer and replace it with something that avoids API roundtrips for testing. But this would mean that

you need to maintain two more implementations of critical infrastructure for testing purposes only and it would make your system test rather unreliable, because you moved a large bit away from the production setup.

A similar argument applies to topic b): Maintaining multiple huge fixtures for different test scenarios through a stack of independent components is a big effort. Micro fixtures for each service easily become inconsistent.

For these reasons, we moved away from the classical testing approach for our end-to-end tests and instead focused on metrics.

The system already collects metrics from each service for production monitoring. These metrics include statistical information like "number of updates received per shop", various failover scenarios e.g. "an update received from a shop was invalid, but could be corrected to maintain a consistent system state", errors, thrown exceptions and more. For system testing we actually added some more metrics which turned out to be viable in production, too.

In order to perform a system test, we install all micro services into a single VM. This allows the developer to easily watch logs and metrics in a central place. A very simple base fixture is deployed that includes some test-users. In addition to that, we maintain very few small components that emulate external systems with entry points to our application. These "emulation daemons" feed the system with randomized but production-like data and actions. One example is the *shop emulator*, which provides product updates (valid and invalid ones) for multiple shops.

As a developer I can now test if the whole application stack still works in normal parameters by monitoring this system and asserting the plausibility of metrics. This helps us to ensure that larger changes, and especially refactorings across service borders, do not destroy the core business processes of the system.

While this is of course not a completely automated test methodology, it serves us quite well in combination with unit and acceptance tests on the service level: The amount of time spent for executing the tests is acceptable, while the effort for maintaining the test infrastructure is low.

What are your experiences with testing micro service architectures? Do you have some interesting approaches to share? Please feel invited to leave a comment!

P.S. I still wonder if it would be possible to raise the degree of automation further,

for example by applying statistical outlier detection to the metrics in comparison to earlier runs. But as long as the current process works fine for us, I'd most probably not invest much time into research in this direction.



## 5. Refactoring

### 5.1 How to Refactor Without Breaking Things

[https://qafoo.com/blog/085\\_how\\_to\\_refactor\\_without\\_breaking.html](https://qafoo.com/blog/085_how_to_refactor_without_breaking.html)

*Tobias Schlitt at 1. June, 2016*

Refactoring means to change the structure of your code without changing its behavior. It is an essential part of everyday programming and should become knee-jerk for your whole development team. Refactoring is very helpful to cleanup feature spikes, revise earlier decisions and keep a maintainable codebase in the long run. In a perfect project world - with extensive automated tests of various types - this is just a matter of getting used to. But there are only very few such projects. So getting into proper refactoring is much harder. This article will show you important tips to master this challenge with your team.

From our experience in various (legacy) projects successful refactoring depends on the following points:

1. Tests
2. Baby steps

#### 5.1.1 Tests

Tests help you to ensure that the behavior of your application does not break while restructuring the code. But in many cases you will want to apply techniques of refactoring just to make your code more testable and to come to a stage where

writing unit tests gets cheap. There the dog seems to chase its own tail.

We found out that high-level functional tests can deal as a good basis to get started with refactoring. Even very old legacy applications can usually be tested through the browser using Mink with PHPUnit (*See: Using Mink in PHPUnit*) (or Mink with Behat). Which of both solutions you choose depends on your project team: If you are familiar with PHPUnit and don't plan of involving non-technical people in testing later, PHPUnit + Mink is a solid choice for you.

Before you start writing tests you need to setup at least a rudimentary automation system that can reset your development installation (most likely the database). The goal must not be to get a fully fledged infrastructure automation (*See: Why you need infrastructure and deployment automation*) (which is of course still desirable) but to get a predictable, reproducible starting state for your tests. Maybe you just hack up a shell script or use PHPUnits `setUpBeforeClass()` to apply a big bunch of SQL.

Then you start writing tests through the front-end for the parts of your code that you want to touch first, e.g. a really bad controller action (which code to refactor first will be part of another article).

Make sure to concentrate on the essentials: Keep the current behavior working. Don't care too much about good test code. You can throw these tests away after you finished your refactoring or just keep a few of them and clean them up later. As usual this is a matter of trade-off (*See: Developers Life is a Trade-Off*). What you want to achieve here is an automated version of the click-and-play tests you'd do manually in the browser to verify things still work.

Code coverage can be of good help here to see if you have already enough tests to be safe. We have a blog post on using code coverage with Behat (*See: Code Coverage with Behat*) for exactly this purpose. The same technique can be applied to running PHPUnit with Mink. But beware: the goal is not `$someHighPercent` code coverage! The goal is to give you a good feeling for working with the underlying code. Once you have reached that state, stop writing tests and focus on the actual refactoring again.

### 5.1.2 Baby Steps

When you start with restructuring your code, do yourself a favor and don't be too ambitious. The smaller your steps are, the easier it gets. Ideally, you will only apply

a single refactoring step (e.g. *extract method* or even *rename variable*) at once, then run your tests and commit.

We know that this is hard to get through in the first place, especially when you did not do much refactoring before. But reminding yourself over and over again to go very small steps into a better direction is really helpful for multiple reasons:

1. It reduces the risk of breaking something. The human brain can only cope with a limited amount of complexity. The larger the change is, the more things you need to keep in mind. This raises the chance of messing things up and waisting time.
2. If you messed up the current step (for example by changing behavior or realizing that your change did not lead to a good result) large changes make it harder for you to just reset and restart. You will think about the time you already invested and will probably go on trying to fix the state. However, this typically makes it worse. Reset to HEAD and restart the refactoring step should be the way to go instead.
3. While you might have a big picture in mind where your refactoring should lead, this might not be the best goal. Maybe there are better solutions you did not think about in the first place. Doing baby steps will keep the door open for correcting your path at any time.
4. Chances that you will get through a large refactoring without being disturbed are low. There is always an emergency fix to be applied, a very important meeting to be joined, good coffee to be drunken or just Facebook that will require you to stop. Getting back into your working stack later will be hard and committing a non-working state should be a no-go. With baby steps, you can just cancel the current step or finish it within seconds and leave safely.

Long story short: Do yourself the favor and get used to baby steps. This will sometimes even result in more ugly intermediate steps. Get over it, things will eventually be better!

As a side note: People often ask us, if committing each and every baby step won't lead to polluting your version history. If you feel that way, rather go for squashing your commits later than doing larger steps with each commit.

What are your tips for successful refactoring? Leave us a comment!

The current issue of the German PHP Magazin<sup>1</sup> also has a slightly more extensive article on this topic by us.

---

<sup>1</sup><https://qa.fo/book-php-magazin-4-16-244272>

## 5.2 Testing legacy code

[https://qafoo.com/blog/003\\_testing\\_legacy\\_code.html](https://qafoo.com/blog/003_testing_legacy_code.html)

Manuel Pichler at 30. July, 2010

Today we know about the benefits of Test Driven Development and normally also start new projects using TDD. Most projects from the last couple of years integrated this method in their daily development process, which often results in in good code coverage results of 90% and above. But what about all the other old projects, you still manage in your daily work?

Those projects already exist for several years, doing their jobs and are adding everyday value for their users. They contain of thousands or hundreds of thousands lines of code, but most or all of them are untested. Estimating the time for modifications and maintenance is nearly impossible and risky, because side effects are always possible. So what can we do with such code bases?



Option one: Ee can try to write tests for the whole application, afterwards. But for this task a developer could spend all his working years without doing anything else then writing tests. So, this isn't really a realistic solution.

Option two: We can start to write tests only for newly added features. At a first glance this seems to be a good compromise. You don't have to invest too much time, and all newly added features are tested. But this solution has one major drawback. What's with all that legacy code? The code that makes the maintenance a nightmare and estimating a gambling game. It is still untested.

This brings us to the third possible solution: We write tests for newly added features and also we add tests for the existing code once we have to change it. With this approach, we have the advantage of both new and old source code will be tested over time, without spending too much time and money for legacy code tests.

Following approach three, chances are good that we get an application stable within a few month, because normally changes to legacy code are not spread throughout the whole application. This means that parts of an application will never change a) because no one uses them or b) because they are so stable so that no one has to change them. While other parts of the application will change regularly a) because there are so many bugs in that part of the source code or b) because there are continuous requirement changes and improvements.

If you would like to get professional support<sup>2</sup> for PHP\_ChangeCoverage or quality in general don't hesitate to contact us<sup>3</sup>.

Sounds good, right? But using a normal coverage report, you may never get an appealing percentage of 70%, 80% or more. The report will stay red and the covered lines will stuck around 25% to 30%. This condition can be really frustrating for all participants, because they work hard on their tests, but there is no positive and motivating feedback. And here comes PHP\_ChangeCoverage<sup>4</sup> into play.

PHP\_ChangeCoverage combines the coverage data collected by PHPUnit<sup>5</sup> and the commit history of your version control system and generates a new coverage report, which only reflects those parts of the application that have been changed within a specified time range. You can get a first version of PHP\_ChangeCoverage from its github repository<sup>6</sup>. :

```
~ $ git clone git://github.com/manuelpichler/php-change-coverage.git phpccov
~ $ cd phpccov
```

---

<sup>2</sup><https://qafoo.com/services/support.html>

<sup>3</sup><https://qafoo.com/contact.html>

<sup>4</sup><http://github.com/manuelpichler/php-change-coverage>

<sup>5</sup><http://phpunit.de>

<sup>6</sup><http://github.com/manuelpichler/php-change-coverage>

In order to run `PHP_CodeCoverage` two additional dependencies must be installed. The first one is `PHPUnit` and the second one is `PHP_CodeCoverage`<sup>7</sup>. Both can be installed through `PHPUnit's` PEAR channel. :

```
~ $ sudo pear channel-discover pear.phpunit.de
~ $ sudo pear channel-discover components.ez.no
3 ~ $ sudo pear install --alldeps phpunit/PHPUnit
~ $ sudo pear install --alldeps phpunit/PHP_CodeCoverage-beta
```



Now that all required programs are installed we can start a small example session. Therefore you can checkout out the small example application which can be found under `docs/example`. :

```
~ $ cd docs/example
~ $ cd example
~ $ svn co file:/// 'pwd' /svnrepo checkout
```

Now let's generate the code coverage report for the project the first time: :

```
2 ~ $ ../../phpccov --coverage-html coverage \
    checkout/test/PHP/ChangeCoverage/ExampleTest.php
```

As you can see, the `PHP_CodeCoverage` command line interface accepts the same arguments as the used `PHPUnit` version accepts. And without any special parameters both tools have nearly the same behavior. If we look onto the generated code coverage report (*Coverage for partial tested code*), we can see that the source has a line coverage of round about 60% and method coverage of 50%. If

<sup>7</sup><http://github.com/sebastianbergmann/php-code-coverage>



- CVS<sup>10</sup>
- Git<sup>11</sup>
- Mercurial<sup>12</sup>
- Subversion<sup>13</sup>

If PHP\_ChangeCoverage cannot detect one of these version control systems it will fallback to a simple file based implementation that uses the last modification time to collect only changed files.

Additionally PHP\_ChangeCoverage uses PHPUnit<sup>14</sup> to execute the test cases and PHP\_CodeCoverage<sup>15</sup> to generate the different report formats.

The project is licensed under New BSD license<sup>16</sup> and is available on Github<sup>17</sup> for forking.

---

<sup>10</sup><http://www.nongnu.org/cvs/>

<sup>11</sup><http://git-scm.com/>

<sup>12</sup><http://mercurial.selenic.com/>

<sup>13</sup><http://subversion.apache.org/>

<sup>14</sup><http://github.com/sebastianbergmann/phpunit>

<sup>15</sup><http://github.com/sebastianbergmann/php-code-coverage>

<sup>16</sup><http://www.opensource.org/licenses/bsd-license.html>

<sup>17</sup><http://github.com/manuelpichler/php-change-coverage>

### 5.3 Teaching and Learning Domain-Driven Design without Patterns

[https://qafoo.com/blog/077\\_teaching\\_learning\\_domain\\_driven\\_design.html](https://qafoo.com/blog/077_teaching_learning_domain_driven_design.html)

*Benjamin Eberlei at 26. February, 2016*

When development teams start to use Domain-Driven Design (DDD) in work projects, by focussing on patterns such as Entity, Repository, Services and Value Objects, then there is trouble ahead. In my experience with teams using DDD the pattern-thinking can be very distracting and dangerous. An approach without patterns can bring the essentials across much better. Especially, less experienced developers have a hard time to follow DDD by pattern because the "Why and How are we doing this" is not easy to teach.

For this post, I am assuming you are starting to use DDD as a developer team without real backing of a domain expert – which seems to be the widespread default. Don't forget to remind yourself that you are missing a critical part, the domain expert to help you get the critical nuances of your code right.

In my experience, when you apply very strict DDD in this scenario, the domain expert will change or specify the requirements at some point and leave you helpless with the wrong model. This can be a huge setback and cause a lot of friction and code changes. CRUD (Anemic) models are much more resilient here, because you can easily adapt or work around all the already existing business logic.

So lets say you are using a mostly CRUD-centric style to application development and still want to benefit from some of the achievements of DDD, that is - you would like business rules to be:

1. using the language of the domain.
2. enforced and translated into code as strictly as possible.
3. independent of technical details and frameworks.

The technical DDD patterns are just one means to those ends. They are not an end on their own, just applying the patterns doesn't make any sense.

We can try to bring those properties to live by relentlessly applying the SOLID principles or the four rules of simple design<sup>18</sup>.

Take the following (extremely simplified) code example performing some money calculations and then storing the result in a database table. It is written from a

---

<sup>18</sup><http://martinfowler.com/bliki/BeckDesignRules.html>

CRUD perspective using the transaction script pattern to implement the "discount" use-case:

```
1 <?php
  namespace MyShop\Domain;

  class OrderService
  {
6      /**
       * @param int $orderId
       * @param float $percent (between 0 and 1)
       */
11     public function discount($orderId, $percent)
        {
            if ($percent < 0 || $percent > 1) {
                throw new \InvalidArgumentException("Percent must be float between 0
                    and 1");
            }

16         $orderTotal = $this->selectOrderTotal($orderId);
            $orderTotal = (1 - $percent) * $orderTotal;

            $sql = 'UPDATE order SET total = ? WHERE id = ?';
            $this->connection->executeQuery($sql, [$orderTotal, $orderId]);
21     }
  }
```

Please ignore the consistency problems here – order items and taxes are simply ignored in this simple example.

The business rules in this snippet cannot live up to the DDD practices:

- Business rules are not strictly enforced for Percent and Money values, they are implemented only when they are needed, operating on primitive types like floats and integers.
- No independence of technical details, the business rules about discount handling, percent and money calculations are mixed with SQL statements.

Without DDD patterns, we start with applying SOLID principles. The single responsibility principle teaches us, that no class should have more than one responsibility, but here `OrderService` knows how to validate percentage values and how to multiply money with percentages.

So lets refactor and introduce a class `Percent` and a class `Money`:

```
<?php
namespace MyShop\Domain;
```

```
5 class Percent
{
    private $value;

    public function __construct($value)
    {
10         if ($value < 0 || $value > 1) {
                throw new \InvalidArgumentException("Percent must be float between 0
                    and 1");
            }
            $this->value = $value;
        }

15     public function asFloat()
        {
            return $this->value;
        }
20 }

class Money
{
    private $value;

25     public function __construct($value)
        {
            $this->value = $value;
        }

30     public function reduceBy(Percent $percent)
        {
            return new Money($this->value * (1 - $percent->asFloat()));
        }

35     public function asFloat()
        {
            return $this->value;
        }
40 }

class OrderService
{
45     public function discount($orderId, Percent $percent)
        {
            $orderTotal = $this->selectOrderTotal($orderId);
            $orderTotal = $orderTotal->reduceBy($percent);
        }
}
```

```
50     $sql = 'UPDATE order SET total = ? WHERE id = ?';  
        $this->connection->executeQuery($sql, [$orderTotal->asFloat(), $orderId]);  
    }  
}
```

Much better! We have separated the responsibilities. Please ignore the fact Percent and Money should probably not use floats internally, there is much more to improve in the code example, which is outside of this blog post.

But we are not done with using SOLID as an improvement method to refactor towards DDD principled code. The dependency inversion principle states that high level code should never depend on low level code, but our high level business logic depends on low level database code here.

Let's introduce an interface that abstracts the data access:

```
<?php  
3 namespace MyShop\Domain;  
  
interface OrderTotalGateway  
{  
8     public function fetch($orderId);  
     public function update($orderId, Money $total);  
}  
  
class OrderService  
13 {  
     public function discount($orderId, Percent $percent)  
     {  
         $orderTotal = $this->totalGateway->fetch($orderId);  
         $orderTotal = $orderTotal->reduceBy($percent);  
18         $this->totalGateway->update($orderId, $orderTotal);  
     }  
}
```

Now we can implement this interface outside the MyShop\Domain namespace, where only high level (business logic) is located, not database code.

```
namespace MyShop\Implementations\Doctrine;  
  
use MyShop\Domain\Money;  
use MyShop\Domain\OrderTotalGateway;  
5  
class DbalOrderTotalGateway implements OrderTotalGateway  
{
```

```
public function fetch($orderId) { // ... }  
10 public function update($orderId, Money $total)  
    {  
        $sql = 'UPDATE order SET total = ? WHERE id = ?';  
        $this->connection->executeQuery($sql, [$total->asFloat(), $orderId]);  
    }  
15 }
```

Now only high-level code is located in the `OrderService` and details such as the money being stored as a float in the database has been pushed down into a low level implementation.

Without ever knowing about DDD patterns we have introduced Value Objects (`Percent` and `Money`) and a class loosely resembling Repository (`OrderTotalGateway`). The business logic is independent from technical details and enforced in a much stricter way than before using encapsulation.

The similar result can be achieved by applying the four rules of simple design, in this case "No duplication" and "Reveals intent" can help achieving the same result without knowing about the SOLID principles.

Another approach to achieve yet the same result is to view this problem from an anti-pattern or code-smell perspective: The original `OrderService` class exhibits a severe case of feature-envy, drawing too much logic onto itself. One reason for this is the primitive obsession using floats for `Money` and `Percent`. Getting rid of this problem also leads to introduction of the `Money` and `Percent` classes.

Don't make the mistake to implement DDD by thinking by the blue books building blocks and technical patterns too much. Using a refactoring based approach combined with SOLID principles or the Four Rules of Simple Design lets you move towards DDD in much smaller and sustainable way. If you want to implement technical DDD practices in your team, start with SOLID and Refactoring instead. This is much easier to teach and learn. Gradually showing team members how the code improves and how simple these small steps goes a long way to convince the team to start using more challenging DDD patterns.

## 6. Architecture

### 6.1 Why Architecture is Important

[https://qafoo.com/blog/079\\_why\\_architecture\\_matters.html](https://qafoo.com/blog/079_why_architecture_matters.html)

*Kore Nordmann at 22. March, 2016*

We experience that the system architectures of our customers grow more and more complex. This is either because of scaling requirements or because developers like to try out new technologies like implementing Microservices in other languages than PHP (Node.js, Go, ...) or different storage technologies (MongoDB, CouchDB, Redis, ...). Depending on the application it is also common to introduce dedicated search services (Elasticsearch, Solr, ...), queueing systems (Beanstalk, RabbitMQ, ZeroMQ, ...) or cache systems (Redis, Memcache, ...).

Often there are very valid reasons to do this but there is also an important problem: You are creating a distributed system and they are really hard to get right & operate. Every system spread across multiple nodes in a network is a distributed system. A system consisting of a MySQL server and a PHP application server is already distributed, but this is a well known problem for most teams. Architecture decisions start to get critical once the data is distributed across multiple systems. Why is this the case?

One of the things which are hardest to repair in existing systems are inconsistencies of your data. Repairing this often even requires manual checks and sanitization which, depending on the amount of data, can take up really large amounts of time.

There are even studies <sup>1</sup> pointing out the costs of bugs in your architecture. If they are discovered late, when the system is already in production, then fixing these bugs can amplify the costs hundredfold. This is why we suggest to investigate and analyze your architecture before distributing your data and be careful doing so.

What are the main points you should check when designing a system architecture for a new project, during scaling an existing project or when introducing new storage technologies (search, storage, cache, ...)? There are a couple of questions we can ask ourselves:

- How can the consistency of data be ensured across multiple systems?
- How do we verify that the chosen systems fulfil their requirements?
- What are the technical and operational risks of newly introduced systems?
- How will the system handle latencies and failures of nodes?
- Is the overall application resilient against single node failures or how can this be accomplished?

On top of that those decisions should be documented and valued by certain criteria. There are even frameworks for documenting system architecture decisions and risks, which you might want to follow like ATAM<sup>3</sup>. Important assessment points are:

- Consistency and security of data
- Performance (latency, transaction throughput)
- Modifiability (Applicability to new products, future change costs)
- Availability (Hardware failure, software bugs)

### 6.1.1 Summary

When introducing new systems you should be careful especially when you plan to distribute your data across multiple nodes. Technology and architecture decisions should not be made because some topic is hot right now (like Microservices) but you should assess that the chosen system architecture actually fulfills your requirements and will be beneficial. Since there will be no *perfect* architecture for your use case one should always document the respective benefits, drawbacks and reasoning why some kind of architecture was implemented.

<sup>1</sup>"Code Complete (2nd ed.)"<sup>2</sup> by Steve McConnell (ISBN 0735619670)

<sup>3</sup><https://en.wikipedia.org/wiki/Special:BookSources/0735619670>

## 6.2 Scaling Constraints of Languages

[https://qafoo.com/blog/088\\_scaling\\_constraints\\_of\\_languages.html](https://qafoo.com/blog/088_scaling_constraints_of_languages.html)

*Kore Nordmann at 2. August, 2016*

Micro-Services or any set of small services are common again right now. While it can make a lot of sense to use a dedicated service for a well defined problem those services are sometimes used just to play with a different server software. While it is pretty obvious for most that selecting the right database is important the same is true for selecting the right language (virtual machine) for the job.

There are different types of services or server applications where different types of virtual machines (executing the opcodes / bytecode of the compiled source code) make more or less sense. What are the criteria we should base such a decision on and which language should we choose when?

When coming from a PHP background there are a couple of technology stacks you come across regularly:

- PHP (Connected to your webserver through FPM or mod\_php with Apache)
- Node.js (JavaScript)
- Go (different server paradigms possible, most is also true for stuff like Elixir/OTP or Erlang/OTP)
- Java (doing the heavy crunching)

### 6.2.1 Why PHP?

Why did we start using PHP or how did it get so popular after all? One reason is **LCoDC\$SS** as described in Roy T. Fieldings dissertation<sup>4</sup>. The abbreviations which describes the network architectural properties of HTTP / REST stands for:

- Layered (L)  
We are using a protocol which allows layering. You can put a reverse caching proxy or a load balancer in front of your application servers.
- Code on Demand (CoD)  
We are delivering HTML, CSS (& JavaScript) which are interpreted by browsers. The rendering does not happen on the server.

---

<sup>4</sup>Architectural Styles and the Design of Network-based Software Architectures<sup>5</sup> by Roy T. Fielding (2000)

- **Client-Server (CS)**  
The Browser (Client) interacts with a Server.
- **Cached (\$)**  
The result of certain HTTP verbs (`GET`, `HEAD`) can be cached. This speeds up the web especially for static resources (images, ...).
- **Stateless (S)**  
The client always transmits the full state which is required to generate the response by the server. This might include session cookies. If your application is implemented correctly any of your frontend servers (serving the same application) can answer any request – no matter if the same client was connected to a different server in an earlier request.

PHP was built for this. With it's so called shared-nothing architecture it handles stateless HTTP requests perfectly. The PHP engine throws away any state (variables, file pointers, database connections, ...) after returning the response. It is almost impossible to build shared state between multiple requests (it is, of course, but don't do this).

This way PHP does not only use every CPU core on a server but usually it is even trivial to add additional servers and just scale your frontend this way. Only works if you respect the statelessness of HTTP requests, though.

### 6.2.2 So, Why Not PHP?

We just learned that PHP seems perfect for the web – does it serve all use cases perfectly? No.

Besides HTTP browsers also speak other protocols like WebSocket. With WebSockets you do not have the HTTP request response cycle but bi-directional permanent connections. PHP is not built for this.

Let's start with a use case where WebSockets are commonly used: A chat. You want that a message written by some user is immediately passed on to every other user in the same chat. This requires state on the server (who is in the chat) as well as passing messages between multiple connections. The incoming message should be passed on to all connected users.

You can implement this with HTTP (long polling), PHP and some software managing the shared state (database, in-memory storage) but it will be complicated and

doesn't make much sense. This is where technologies like Node.js come into play.

### **Node.js**

With Node.js it is really simple to develop your own server software. Using HTTP, Websockets or anything else is pretty straight forward and you probably already know the language from some frontend development. Using asynchronous IO Node.js allows to do some work while waiting for the network or file system. This allows to answer many requests per second while sharing some state – like for a chat.

On the other hand does Node.js not really support multiple cores in a sensible way. You can fork multiple processes or start multiple Node.js servers but you are loosing the benefit of shared state immediately. And managing this state will get hard and cumbersome again.

What I like to use Node.js for is quickly developing simple servers. For services which require state and can easily be handled by a single CPU core, Node.js might even be ready to serve production ready services.

### **Go**

Developed by Google Go as a language is not as easy to pick up as Node.js when you come from a pure web development background. This is even worse for languages like Elixir or Erlang.

On the other it is fairly easy to write servers with those languages and they even utilize multiple CPU cores and still maintain shared state. If you need to scale a WebSocket-based application you might want to take a look at those.

When a single server is not large enough any more you'll have to come up with intelligent ideas to distribute your service - but this is far beyond "micro" then.

### **Java**

Java is commonly used for complex backend services or to integrate with existing (legacy) stacks in companies. The Java VM actually allows to build both fast applications embracing the shared nothing nature of HTTP, but also servers with a shared state. Since state is so easy to share, Java applications are often not scalable beyond a single server. The Java VM on the other hand scales really well on a single server.

### 6.2.3 Summary

Choosing the language (virtual machine) is not just a matter of taste. Because of different paradigms the languages / virtual machines are build with the decision has architectural consequences. There are also other factors like developer experience, license costs or similar when it comes to choosing the right technology for your team & application.

## 6.3 How To Synchronize a Database With ElasticSearch?

[https://qafoo.com/blog/086\\_how\\_to\\_synchronize\\_a\\_database\\_with\\_elastic\\_search.html](https://qafoo.com/blog/086_how_to_synchronize_a_database_with_elastic_search.html)

*Kore Nordmann at 14. June, 2016*

Since search engines like Apache Solr and ElasticSearch are easy to use and setup more and more applications are using them to index their content and make it searchable by the user. After all the underlying Lucene index engine provides far more powerful features than a plain MySQL full text search or similar solutions. With Apache Solr and ElasticSearch you can enhance the performance and the functionality of your website.

What we often stumble across, though, is the naive approach of synchronizing both data storages. It boils down to:

- Store data in database
- Store data in search index

The problem with this approach is, as in any distributed system, that both the first and the second write **will** fail at some point. When the first write to the database fails usually nothing is indexed. When the second write to the search index fails you have content which exists but the user won't be able to find it in the search index.

Depending on your domain this might not be a problem. But if a price update of a product in an online shop fails you might actually loose money or at least customers will be angry about the inconsistencies on your site. There will be similar problems in other domains.

### 6.3.1 Transactions

The second idea of most developers now have is to build some kind of transaction support across both storages. We know transactions well from our Computer Science courses and relational database management systems, thus it is an intuitive decision. With a transaction wrapped around you'll get something like this:

```
5 | try {  
    |     $database->store($document);  
    |     $searchIndex->index($document);  
    | } catch (IndexException $e) {  
    |     $database->remove($document);  
    |     // Inform user that write failed
```

```
| }
```

This works, but it puts the burden of resubmitting the document and handling the failure in your application on the user. Resubmitting a complex document might or might not work in a sane way for your users. If the user edited a long and complex text and all edits are lost because the search index failed to update – your user might not be that happy.

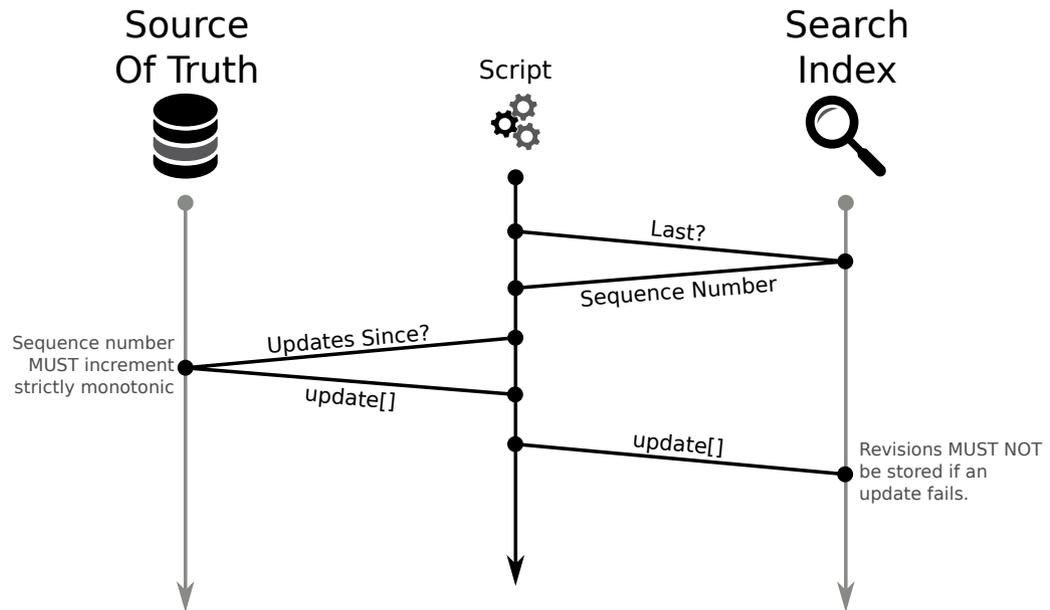
### 6.3.2 Changes Feed

There is a better way to implement this, often used in systems where you have one primary storage and any number of secondary storages. In this scenario your primary storage defines the state of your system – it is commonly called "Source Of Truth". In this example your database knows the definitive truth since everything is first stored there. Any other system, like your search index or caches should be updated based on the state in the database.

What we want to do is passing every change in the database on to the search index. A changes feed can do this. Some databases like CouchDB<sup>6</sup> offer this out of the box for every document, but how can this done with a relational database management system like MySQL? And how can the data be passed on safely to the search index? The process behind this can be illustrated by this simplified sequence diagram:

---

<sup>6</sup><http://guide.couchdb.org/draft/notifications.html>



The idea behind this is that we start asking the target system (search index) what document it already knows. For this we need some kind of sequence number or revision – which **MUST** be sortable & strictly monotonic<sup>7</sup> to make sure that we do not lose any updates. With the last sequence number returned from the target system we can now load all or just a batch of changes from our database. These changes are then pushed to the search index. This sounds simple, right? But there are two questions remaining:

- How to I get MySQL to produce correct sequence numbers?
- How can I store the last sequence number in the search index?

<sup>7</sup>[https://en.wikipedia.org/wiki/Monotonic\\_function](https://en.wikipedia.org/wiki/Monotonic_function)

### 6.3.3 Generating Correct Sequence Numbers

This is the hardest part with the described architectural pattern and a little bit harder to get right than one would assume in the first place. Every sequence number **MUST** be bigger than the last sequence number otherwise changes can be lost.

The easiest way to implement this is an additional table with the following schema:

```
CREATE TABLE changes (  
  3  sequence_number INT NOT NULL AUTO_INCREMENT,  
    document_id INT NOT NULL,  
    PRIMARY KEY (sequence_number)  
);
```

The `document_id` would reference a data set in another table – if you want you can even define a foreign key relation with the data indexed in the search index. Defining the `sequence_number` number as `AUTO_INCREMENT` ensures that MySQL takes care of incrementing the sequence number.

With every change to a document we now also append a row to the `changes` table. Do this inside one transaction. Afterwards we can just query the changes like:

```
SELECT  
  sequence_number, document.*  
FROM changes  
JOIN -- ... --  
5 WHERE sequence_number > :since  
ORDER BY sequence_number ASC  
LIMIT 0, 100;
```

This would request the next 100 changes and join them with the actual data.

This table will grow fast you say? You're right.

But this can be optimized. The only thing we must keep in this table is the latest sequence number of each `document_id` so that we can run a clean import in the correct order. If there are dependencies between your documents this can get a little more complex but can still be solved.

Unnecessary at first, but at some point you might also have to handle the case where the `sequence_number` overflows.

### 6.3.4 Storing The Sequence Number

The sequence number must not be increased in the search index if no document was stored. Otherwise we would lose the document since the next request for new

documents to the database will use the already increased sequence number.

Since systems like ElasticSearch do not support transaction we should store the sequence number associated with the update right in the document. Using a MAX query<sup>8</sup> and an index on the sequence number field we can still fetch the last sequence number from the search index.

Another option would be to store the last sequence number in a special document or somewhere else like the file system. If ElasticSearch now loses some or all documents we will not be aware of it and some documents will again be missing from the search index. Solutions like ElasticSearch tend to only persist their in memory data every few seconds and if the node crashes in the mean time they **will** loose some data. The described pattern ensures the index is brought up to date once the node recovers. Even with a complete data loss the architectural pattern described here will automatically ensure the entire index is rebuilt correctly.

### 6.3.5 Conclusion

The architectural pattern discussed in this blog post is nothing we invented. It is actually used by replication mechanisms of databases or other storage systems. We just showed you an adapted version for the concrete use case of synchronizing contents between a database and ElasticSearch.

The shown pattern is stable and resilient against failures. If ElasticSearch is not available while storing the data it will catch up later. Even longer down times or complete node failures will be automatically resolved. The mechanism is also called *eventual consistency* because of this.

The problem with this approach is that you'll need an extra script or cron job to synchronize both systems. The search index will also always lag slightly behind the database, depending on how often the script is run.

We generally advise to use this pattern to keep multiple system in sync when there is one source of truth. The pattern does not work when data is written to multiple different nodes.

---

<sup>8</sup><https://qa.fo/book-search-aggregations-metrics-max-aggregation>

## 6.4 Common Bottlenecks in Performance Tests

[https://qafoo.com/blog/082\\_common\\_bottlenecks\\_in\\_performance\\_tests.html](https://qafoo.com/blog/082_common_bottlenecks_in_performance_tests.html)

Kore Nordmann at 19. April, 2016

Most developers by now internalized that we should not invest time in optimizations before we know what happens exactly. Or as Donald Knuth wrote:

Programmers waste enormous amounts of time thinking about, or worrying about, the speed of noncritical parts of their programs, and these attempts at efficiency actually have a strong negative impact when debugging and maintenance are considered.<sup>9</sup>

This is true for optimizations in your PHP code but also for optimizations regarding your infrastructure. We should measure before we try to optimize and waste time. When it comes to the assumed performance problems in your system architecture most people guess the root cause will be the database. This might be true but in most projects we put under load it proved to be false.

So, how can we figure out where the problems are located in our stack?

You need to test the performance, as always. For optimizing your PHP scripts you'd use XDebug and KCacheGrind or something like Tideways<sup>11</sup>. You'd trigger a single request and see what is slow in there and optimize that part.

### 6.4.1 System Test Setup

It is slightly more complicated to test your full stack. In the optimal case you simulate the real user behaviour. It is definitely not sufficient to just run `ab` or `siege` against your front page. For an online shop typical user tasks could be:

- Browsing the catalogue (random browser)
- Product searches
- User sign up & login
- Order checkout

We usually discuss the common tasks on a website with the owner of the website. Then we discuss the numbers for each of those task groups which we should

<sup>9</sup>Computing Surveys (PDF)<sup>10</sup>, Vol 6, No 4, December 1974

<sup>11</sup><https://jmeter.apache.org/>

simulate to mimic a certain scenario. With this information a jMeter<sup>12</sup> test can be authored simulating the real behaviour of your users. After a test run you can compare the access logs of the test with common access logs from your application to verify you simulated the right thing – if those logs are available.

### 6.4.2 Stack Analysis

Once you wrote sufficiently generic tests you will be able to simulate larger numbers of users accessing your website – some may call it DDOS'ing<sup>13</sup> your own page. When running these tests you can now watch all metrics for your system closely and you'll be able to identify the performance problems in your stack.

There are couple of tools which might help you here, but the list is far from extensive and depends a lot on the concrete application:

- vmstat watches CPU usage, free memory and other system metrics
- iftop shows if network performance is an issue
- Tideways<sup>14</sup> or XHProf for live analysis of your application servers

On top of that you should definitely watch the error logs on all involved systems.

### 6.4.3 It is Not Always The Database

The root cause for performance problems on websites we put under test weren't rooted in the database for most of our test runs:

- Varnish & Edge Side Includes (ESI)

We tested a large online shop which used ESI so extensively that the application servers running PHP were actually the problem. The used framework had a high bootstrap time so that this was the most urgent performance impediment. The database wasn't even sweating.

- Network File System (NFS) locking issues

Once you put a high load on a server sub systems will behave differently. NFS, for example, tries to implement some locking for a distributed file system. When multiple servers are accessing the same set of files it can stall your

---

<sup>12</sup>[https://qa.fo/book-Denial-of-service\\_attack](https://qa.fo/book-Denial-of-service_attack)

<sup>13</sup><https://tideways.io>

<sup>14</sup><https://qafoo.com/services/workshops/performance.html>

application entirely. Something you will almost never hit during development but in load tests or later in production.

There are even configuration issues which only occur under load and degrade your performance more than any slow query will do.

- Broken server configurations  
In one case a web hoster who claimed to be specialized on cluster setups provided us with a setup where we ran the tests on. The cluster allowed a lot more FPM children to spawn than database connections. Once put under load the MySQL server rejected most of the incoming connections which meant the application failed hard.
- Opcode cache failures  
Wrong Opcode cache (APC, eAccelerator, ...) configuration can even degrade PHP performance. But this is also something you will not notice without putting the system under load. So you'll only notice this when many customers try to access your website or during a load test.

#### 6.4.4 Summary

If you want to ensure the application performs even under load you should simulate this load before going live. The problems will usually not be where you thought they would be. Test, measure and analyze. The time invested in a load test will be a far more efficient investment than random optimizations based on guesswork.

Kore will talk about "Packt Mein Shop das?" ("Will My Shop Manage This?") at Oxid Commons<sup>15</sup> in Freiburg, Germany on 02.06.2016 about load testing.

---

<sup>15</sup><https://qa.fo/book-p261-knuth>

## 6.5 Embedding REST Entities

[https://qafoo.com/blog/048\\_embedding\\_rest\\_entities.html](https://qafoo.com/blog/048_embedding_rest_entities.html)

Tobias Schlitt at 13. June, 2013

During my talk at IPC Spring<sup>16</sup> I showed an approach for embedding entities in REST responses. This methodology might be useful if the standard use-case for your API is to request multiple related entities together. The downside of such an approach is the raised complexity of caching and especially purging logic. In this blog post I will further elaborate on the approach of resource embedding.

### 6.5.1 Entities

As an example I chose the resources *Product* and *Category* from the e-commerce domain. A product entity could be encoded through the media type `application/vnd.com.example.product+xml` as shown as follows:

```

<?xml version="1.0" encoding="UTF-8"?>
<product
  xmlns="urn:com.example.product"
  xmlns:atom="http://www.w3.org/2005/Atom">
5  <atom:link rel="self"
      type="application/vnd.com.example.product+xml"
      href="http://example.com/products/23" />
  <name>Glow Stone Driveway</name>
  <description>Awesome ...</description>
10 <atom:link rel="collection"
      type="application/vnd.com.example.category+xml"
      href="http://example.com/categories/geek_toys" />
  <!-- More links ... -->
</product>

```

A product has a `name` and a `description`. Besides that, the representation provides Atom<sup>17</sup> links as hyper media controls, using relations as recommended by the IANA<sup>18</sup>.

Corresponding, the following example shows an encoded category entity:

```

<?xml version="1.0" encoding="UTF-8"?>
<category

```

<sup>16</sup>[https://qafoo.com/blog/047\\_pragmatic\\_rest\\_bdd\\_ipc.html](https://qafoo.com/blog/047_pragmatic_rest_bdd_ipc.html)

<sup>17</sup><https://tools.ietf.org/html/rfc4287>

<sup>18</sup><https://qa.foo/book-link-relations>

```
  xmlns="urn:com.example.category"
  xmlns:atom="http://www.w3.org/2005/Atom">
5 <atom:link rel="self"
  type="application/vnd.com.example.category+xml"
  href="http://example.com/categories/geek_toys" />
  <name>Geek Toys </name>
  <products>
10   <atom:link rel="item"
    type="application/vnd.com.example.product+xml"
    href="http://example.com/products/23" />
    <!-- ... -->
  </products>
15 <!-- Links: overview, paging, sorting, ... -->
</category>
```

Note that here is a list of product items linked from a category, which could for example be paged, sorted, filtered and so on using URL parameters.

### 6.5.2 The Use-Case

Imagine the major use-case for the API is to retrieve the top 20 products from a category. To retrieve these together with product details, a client needs 21 GET requests on the first go. For each subsequent execution of the use case, the 21 requests must be repeated, although the response might be that the entities did not change.

Alternatively, the server can set expiry times for the entities, so that clients do not need to re-check that frequently. But such time spans are hard to calculate or even guess and can easily lead to stale caches.

So, depending on the requirements for data accuracy and clients, this might easily lead to a huge number of requests.

### 6.5.3 Resource Embedding with HAL

A promising idea to solve the named issue is to deliver resources in an embedded way while making explicit that the delivered data partially belongs to a different resource that can be found at a different URI. Searching the web reveals that there is already an approach for this out there, which is called Hypertext Application Language (HAL)<sup>19</sup>.

<sup>19</sup>[http://stateless.co/hal\\_specification.html](http://stateless.co/hal_specification.html)

Embedding products into a category using a HAL encoding is shown in the following example:

```

5 <?xml version="1.0" encoding="UTF-8"?>
  <resource href="/categories/geek_toys">
    <name>Geek Toys</name>
    <products>
      <resource rel="product" href="/products/23">
        <name>Glow Stone Driveway</name>
        <description>Awesome ...</description>
        <link rel="category" href="/categories/geek_toys" />
      </resource>
10 <!-- ... -->
    </products>
  </resource>

```

HAL encodes any resource using the `<resource>` tag and has its own `<link>` tag. As can be seen, the products which are part of a category are embedded into its representation, but the categories that are assigned to a product are linked.

While this approach appears charming at first glance, it has some flaws:

First of all, HAL is meant to encode all entities in a standardized manner of a `<resource>` tag. That conflicts with one of the fundamental ideas of hyper media: to encode an entity in a semantically meaningful way.

Second, the HAL specification does not provide its own namespace, making it ugly to re-use among other XML languages. Finally, the specification introduces its own `<link>` element instead of re-using existing standards like Atom<sup>20</sup> or XLink<sup>21</sup>.

#### 6.5.4 Better Resource Embedding

Since the idea of HAL is essentially quite nice, I made up my mind to make the approach more standards-compliant and to mitigate the named issues. The basic idea is to re-use the Atom `<link>` elements, as in the original drafts, and embed resources as their children:

```

<?xml version="1.0" encoding="UTF-8"?>
<category xmlns="..." xmlns:atom="..."
  xmlns:p="urn:com.example.product"><!-- ... -->
  <products>

```

<sup>20</sup><https://tools.ietf.org/html/rfc4287>

<sup>21</sup><http://www.w3.org/TR/xlink/>

```
5   <atom:link rel="item" type="..." href="..." >
      <p:product><!-- ... -->
          <p:name>Glow Stone Driveway </p:name>
          <p:description>Awesome ... </p:description>
          <atom:link rel="collection"
10         type="application/vnd.com.example.category+xml"
            href="http://example.com/categories/geek_toys" />
          </p:product>
      </atom:link >
      <!-- ... -->
15 </products>
</category>
```

This example left out some unimportant details, it is basically derived from the original example in this post. However, in addition to the category namespace, the namespace for encoding product entities is also imported using the shortcut `p`. Inside of the `<atom:link>` for a product, the entity representation itself is embedded using this namespace.

Naturally, the mentioned drawbacks of HAL are mitigated: Clear semantics are kept, each media type ships with a decent XML namespace and a standard `<link >` element is used (Atom). Luckily, the Atom specs allow any external element to occur inside of `<atom:link>` so embedding the product entity is perfectly valid.

Furthermore, if clearly documented, that would allow the REST provider to transparently switch on and off embedding of specific resources. A client then must not rely on embedded resources and must be capable of fetching them through their link references. However, if available, it can use the embedded information to cache an embedded resource transparently.

### 6.5.5 Bottom Line

Of course, resource embedding is nothing I would recommend to you in general. It can be considered a *hack* for special use-cases. One of its drawbacks is to weaken the caching capabilities RESTful web APIs provide: The category resource must be re-fetched every time one of the embedded products changes. This of course also affects your application's caching logic, because you need to purge caches accordingly.

So, if you think about making use of embedded resources, make sure you analyzed the use-cases of your API carefully and keep the hack as local as possible.

## 7. Workflow

### 7.1 Coding in Katas

[https://qafoo.com/blog/034\\_coding\\_in\\_katas.html](https://qafoo.com/blog/034_coding_in_katas.html)

*Tobias Schlitt at 18. February, 2013*

In almost any kind of sports you hone your skills by repeating a small piece of practice over and over again. Pretty much the same works for learning to play a musical instrument. The idea of Code Katas applies this simple but effective method of exercise to the world of programming.

A programmer's everyday life mostly consists of a single goal: getting things done. New features are to be shipped, bugs to be found and fixed. While you can certainly learn many new things thereby, this scenario lacks two essential factors for effective learning and self-development: **Time to try and possibility to fail**. Everyday work is typically not the situation where taking a risk and failing by intention is desired.

So, if you want to grow beyond yourself, you need to create an environment for your practice where failure is allowed. Starting a spare time project or even an open source project can help there. But finding a sufficiently clear and defined task to practice your skills is hard. That's where the idea of Code Katas jumps in:

Kata ( or literally: "form") is a Japanese word describing detailed choreographed patterns of movements practiced either solo or in pairs.

-- <https://en.wikipedia.org/wiki/Kata>

A Code Kata is basically a well-defined, small task to be solved by programming. The time to solve a Kata should be reasonable to allow you to get started rapidly and to see first results quickly. However, you usually set yourself a time limit, say 1 hour, and stop precisely when the limit is reached. The actual result of this work is only a secondary concern, same applies to finishing the task. What counts is the actual practice. Pushed even further: You might actually want to throw away the code later.

The important part is to reflect on what you did: How did you approach the problem? Where did your approach lead you? What were the problems? What went well? Ideally, ask another developer for feedback or, if possible, directly work in pairs to learn from each other.

And when you are done: Start over from scratch. Maybe directly after finishing the first iteration or a week later. Do the same Kata again. Try a different approach or try the same and see where it leads you after your reflection. Reflect again and set yourself a goal for the next iteration.

With these characteristics, a Kata gives you a good basis to try out something new, to let off steam and to practice your skills as a programmer. Maybe you could try a purely test driven approach now. Maybe you will think about a complete OO model upfront then. And maybe you could apply Object Calisthenics in another iteration.

One of the most famous Code Katas is the Bowling Game Kata<sup>1</sup> by Robert C. Martin. The task of this Kata is to implement bowling with its simple rules. Not a tough challenge, but something with a connection to reality, clearly structured and well-defined. Uncle Bob used this Kata to visualize the methodology of Test Driven Development.

At Qafoo we tend to use Katas more and more frequently in our trainings for practice. Of course, Katas can neither replace practices derived from the real world nor working on own code. But they offer our training attendees the possibility to perform quick practices in-between theoretical lectures and to reflect on different approaches to certain problems in practical trying.

---

<sup>1</sup>[https://qa.foo/book-ArticleS\\_UncleBob](https://qa.foo/book-ArticleS_UncleBob)

There is a Kata Catalogue<sup>2</sup> where you can find quite some Katas. Tell us about your experience with Coding Katas!

---

<sup>2</sup><http://codingdojo.org/cgi-bin/wiki.pl?KataCatalogue>

## 7.2 Why you need infrastructure and deployment automation

[https://qafoo.com/blog/065\\_infrastructure\\_automation.html](https://qafoo.com/blog/065_infrastructure_automation.html)

*Benjamin Eberlei at 28. February, 2014*

The amount of time wasted on setup, synchronization and deployment of applications is often ignored by teams and managers. Manual server management and application deployment are a huge waste of time and ultimately money. In addition, manually performing these tasks is often prone to error and a big risk for uninterrupted uptime of production.

Software quality does not stop with tests and good CodeSniffer and PHP Mess Detector scores. The deployment and setup is equally important to the quality of an application. From our experience, this is a field where many teams could still achieve huge productivity gains.

As a yardstick for the optimal solution, we can take a look at what the Joel test<sup>3</sup> has to say about deployment and infrastructure and add some of our own additional requirements (some from the OpsReportCard<sup>4</sup>):

- Can you make a build in one step?
- Do you make daily builds?
- Do you use configuration management tools to automate infrastructure?
- Is the development setup documented and automated?
- Can you rollout and rollback deployments in one step?
- Can applications be deployed to a new server setup without changes to the code?

I want to discuss these bullets in short paragraphs and mention why solving these problems allows you to save time and money, if done properly.

### 7.2.1 Can you make a build in one step?

One of the Joel test requirements is building an application in one step. In the PHP world that means that - by executing one command or pushing a button in the CI system - a process is triggered that creates a running version of your project.

What does a complete build include?

---

<sup>3</sup><http://www.joelonsoftware.com/articles/fog0000000043.html>

<sup>4</sup><http://www.opsreportcard.com/>

- Code deployment to the appropriate machines
- Configuration
- Database setup or migrations
- Service location
- Cache Warming

Tools to automate code deployment include Bashscript + rsync<sup>5</sup> + symlinks, Capistrano<sup>6</sup> (with Capifony plugin for Symfony), Ansible<sup>7</sup> or Ant<sup>8</sup>. It is perfectly fine to automate with bash scripts, some automation is always better than none. However, the amount of boilerplate with Bash is often quite big compared to the other solutions that explicitly handle most of the deployment issues already. For example, we regularly implement a combination of Ansible and Ant deployment for our customers.

For database migrations there are tools such as DBDeploy<sup>9</sup>, Liquibase<sup>10</sup> or FlywayDB<sup>11</sup>. These can normally be integrated and automated during deployment via Capistrano, Ansible or Ant. This requires the database schema to be always backwards compatible, something to carefully watch out for.

A compelling argument for introducing a reliable one-step build is that it can be safely triggered by any member of the team and therefore reduces the possibility of errors drastically.

### 7.2.2 Do you make daily builds?

If you don't have a one-step deployment, it is very unlikely that you are able to make daily builds in a cost-effective way. Without daily builds you lose one benefit of PHP, the rapid pace of development.

You lose money because you might be slower with new features than your competitor or because you cannot remove costly bugs from your software fast enough.

---

<sup>5</sup><http://rsync.samba.org/>

<sup>6</sup><http://capistranorb.com/>

<sup>7</sup><http://www.ansible.com/home>

<sup>8</sup><http://ant.apache.org/>

<sup>9</sup><http://dbdeploy.com/>

<sup>10</sup><http://www.liquibase.org/>

<sup>11</sup><http://flywaydb.org>

### 7.2.3 Do you use configuration management tools to automate infrastructure?

When running many different projects, all with their own unique server setup, you can easily see how this will cost you time and money to maintain all these servers. Automating server maintenance helps a lot.

CFengine<sup>12</sup> was the first tool that automated the setup of servers and it was released in 1993. Competitors such as Puppet, Chef and Ansible are much younger, but still old enough to be considered stable and frankly there is no excuse not to use them anymore.

Configuration management tools all solve the problem of consistently automating the setup of new and old servers. The goal is to set up all servers in a similar, consistent way by the use of automation instead of manual SSH'ing, package and configuration management.

Tasks you want to automate with configuration management:

- Deployment of SSH authorized keys
- Security updates of packages
- Management of log files (logrotate, fetching, aggregating)
- PHP, Apache/Nginx, MySQL, NoSQL, ... configuration
- Updates of commonly used software such as Wordpress, TYPO3, Drupal, Shopware, Oxid and their third party extensions.

Recently we recommend using Ansible<sup>13</sup> to our customers, because it is so much easier to understand than other tools.

Puppet<sup>14</sup> is much more complex to understand initially from our own experience and also has a much higher learning curve.

### 7.2.4 Is the development setup documented and automated?

There are two obvious reasons why not automating the setup of development machine is dangerous:

- Bugs and problems of the category "it works on my machine" are much more likely to happen, especially when development and production machines significantly differ (PHP 5.2 vs 5.3 was such a difference).

---

<sup>12</sup><http://cfengine.com/>

<sup>13</sup><http://www.ansible.com/home>

<sup>14</sup><http://puppetlabs.com/>

- You cannot easily introduce new developers to new projects because it always requires a huge amount of time. This contributes to knowledge silos and significantly increases the bus factor of projects.

At a time where php applications use much more services than just the LAMP stack, it is vital that every developer has the same production-like environment to develop on. This can be achieved with virtual machines and tools like Vagrant<sup>15</sup>. Combined with configuration management tools like Puppet<sup>16</sup>, Chef<sup>17</sup> or Ansible<sup>18</sup>, the tools for easy setup of development machines exist and are stable.

A more detailed post from us on Vagrant will be forthcoming in the next weeks.

### 7.2.5 Can you rollout and rollback deployments in one step?

Nothing **always** works and for deployments of applications this is even more true. Not accounting for failure is careless. Failure during deployments is not a black swan event, it's real and probably in the >1% likelihood.

You can reduce the cost of failure by being able to rollback to an old, working version.

You need a deployment system where several versions of your code are available on the production machines and where you can switch between them with a single atomic operation. This also requires asset compilation and database schema migrations that are backwards and forwards compatible.

### 7.2.6 Can applications be deployed to a new server setup without changes to the code?

One common mistake in application development is hardcoding configuration variables. Especially in the early development of greenfield projects, too much configuration got hardcoded into the code. The side effect of this is that you will probably lose a lot of time once the first real production setup should be deployed. Hardcoded configuration variables also often complicate the installation of the applica-

---

<sup>15</sup><http://www.vagrantup.com/>

<sup>16</sup><http://puppetlabs.com/>

<sup>17</sup><http://www.getchef.com/chef/>

<sup>18</sup><http://www.ansible.com/home>

tion on different servers. For continuous integration, staging machines or for multiple tenant setups, however, this is extremely important.

Not properly taking care of application configuration will make deployment much harder.

### 7.2.7 Conclusion

Not automating deployment and infrastructure is expensive and can cause frequent human errors. This article highlighted several deployment problems and discussed tools and procedures to fix them. In the next years, the question of infrastructure automation will become even more important, as automation has positive effects on time to market and allows scaling development teams up. As usual, companies that don't automate will have a hard time competing with companies that do.

If you need help automating your infrastructure and deployment, don't hesitate to contact us<sup>19</sup>. We can help you with the details and problems that commonly occur and we will make sure you get started very quickly.

---

<sup>19</sup><https://qafoo.com/contact.html>

## **Impressum**

Erstellt und vertrieben wird dieses Buch durch:

Qafoo Gesellschaft mit beschränkter Haftung  
Büngelerstr. 64  
46539 Dinslaken

E-Mail: [contact@qafoo.com](mailto:contact@qafoo.com)  
Telefon: +49 (0)209 40501252  
Registergericht: Amtsgericht Duisburg  
Registernummer: HRB 27078

Umsatzsteuer-Identifikationsnummer gem. §27a UStG: DE272316452  
Inhaltlich Verantwortlicher: Kore Nordmann (Anschrift, s.o.)

Vertretungsberechtigte Geschäftsführer:

- Kore Nordmann
- Manuel Pichler
- Tobias Schlitt