

The Node Craftsman Book

An advanced
Node.js tutorial

The Node Craftsman Book

An advanced Node.js tutorial

Manuel Kiessling

This book is for sale at <http://leanpub.com/nodecraftsman>

This version was published on 2015-10-10



Leanpub

This is a [Leanpub](#) book. Leanpub empowers authors and publishers with the Lean Publishing process. [Lean Publishing](#) is the act of publishing an in-progress ebook using lightweight tools and many iterations to get reader feedback, pivot until you have the right book and build traction once you do.

©2013 - 2015 Manuel Kiessling

Tweet This Book!

Please help Manuel Kiessling by spreading the word about this book on [Twitter!](#)

The suggested hashtag for this book is [#nodecraftsman](#).

Find out what other people are saying about the book by clicking on this link to search for this hashtag on Twitter:

<https://twitter.com/search?q=#nodecraftsman>

Also By Manuel Kiessling

The Node Beginner Book

Node入门

El Libro Principiante de Node

Livro do Iniciante em Node

Beginning Mobile App Development with React Native

For Nora, Aaron and Melinda.

Contents

Preface	i
About	i
Status	i
Notes on code formatting	i
Intended audience	i
Part 1: Node.js basics in detail	1
Introduction to Part 1	2
Working with NPM and Packages	3
Test-Driven Node.js Development	5
Object-oriented JavaScript	11
Blueprints versus finger-pointing	11
A classless society	12
Creating objects	12
Object-orientation, prototyping, and inheritance	20
A classless society, revisited	21
Summary	23
Synchronous and Asynchronous operations explained	25
Visualizing the Node.js execution model	25
Blocking and non-blocking operations	29
Using and creating Event Emitters	33
Introduction	33
Creating your own Event Emitter object	37
Summary	47
Optimizing code performance and control flow management using the <i>async</i> library . .	48
Executing expensive asynchronous background tasks in parallel	49
Optimizing code structure with <i>async</i>	55

CONTENTS

Node.js and MySQL	61
Using the node-mysql library	61
A first database application	61
Using node-mysql's Streaming API	65
Making SQL queries secure against attacks	70
Summary	78
Node.js and MongoDB	79
Some MongoDB basics	79
Applying CRUD operations with the low-level mongodb driver	80
Retrieving specific documents using filters	86
More complex update operations	90
Working with indexes	92
Querying collections efficiently	95
Summary	99
Part 2: Building a complete web application with Node.js and AngularJS	100
Introduction	101
The requirements from a user's perspective	101
High level architecture overview	102
Setting up the development environment	104
Milestone 1: A first passing test against the server	106
Milestone 2: The API responds with actual database content	110
Abstracting database access	110
Ensuring a clean slate for test runs	112
Completing the first spec	113
Milestone 3: Setting the stage for a continuous delivery workflow	117
Introducing automatic database migrations	120
Milestone 4: Giving users a frontend	125
Setting up frontend dependencies through bower	125
Serving the frontend through the backend server	125
Adding the frontend code	126
Adding AngularJS view templates	133
Milestone 5: More work on the backend	136
Adding a route for retrieving categories	136
Making the backend recognize different environments	142

CONTENTS

Milestone 6: Completing the backend and finalizing the application	147
Creating the route for adding new keywords	147
Creating the route for updating keywords	150
Creating the route for deleting keywords	155
Summary	160
Epilogue	161

Preface

About

The aim of this book is to help beginning JavaScript programmers who already know how to write basic Node.js applications in mastering JavaScript and Node.js thoroughly.

Status

This book is finished and will only receive updates regarding errors. It was last updated on October 10, 2015.

All code examples have been tested to work with Node.js v4.1.2.

Notes on code formatting

Please note that long lines in code examples may receive a line-break, denoted by the backslash sign `\`. Take care to put those lines into your editor as one line.

Take, for example, this code block:

```
1         dbSession.fetchAll('SELECT id, value, categoryID FROM keyword ORDER BY i\  
2 d', function(err, rows) {
```

The PDF version of this book shows a line break right after the *i* of the *ORDER BY id* part, denoted by a `\` after the *i*.

However, in your editor, you need to make sure that the whole code resides on one line.

Also, note that from what I can see, code blocks can not be copy-pasted from the PDF version of this book. While this might seem like a major annoyance, experience shows that learning to program works way better if one is forced to transcribe code from a book into the editor.

Intended audience

This book will fit best for readers that are familiar with the basic concepts of JavaScript and have already written some basic Node.js applications.

As this book is a sequel to *The Node Beginner Book*, I recommend reading it before starting with this book.

Part 1: Node.js basics in detail

Introduction to Part 1

The goal of this book is to enable you to write Node.js programs of any level of complexity, from simple utilities to complex applications that use several external modules and consist of several layers of code, talking to external systems and serving thousands of users.

In order to do so, you need to learn about all the different aspects of Node.js - the tools, the methodologies, the libraries, the APIs, the best practices - and you need to learn how to put all of that together to create a working whole.

Therefore, I have split this book into two parts: A collection of different basics in Part 1, and a thorough tutorial on how to put these basics together to build a complex application in Part 2.

In Part 1, every chapter stands on its own and isn't directly related to the other chapters. Part 2 is more like a continuous story that starts from scratch and gives you a finished and working application at the end.

Thus, let's start with Part 1 and look at all the different facets of Node.js software development.

Working with NPM and Packages

We already used NPM, the *Node Package Manager*, in order to install a single package and its dependencies for the example project in *The Node Beginner Book*.

However, there is much more to NPM, and a more thorough introduction is in order.

The most useful thing that NPM does isn't installing packages. This could be done manually with slightly more hassle. What's really useful about NPM is that it handles package dependencies. A lot of packages that we need for our own project need other packages themselves. Have a look at <https://npmjs.org/package/request>, for example. It's the overview page for the NPM package *request*. According to its description, it provides a "simplified HTTP request client". But in order to do so, *request* not only uses its own code. It also needs other packages for doing its job. These are listed under "Dependencies": *qs*, *json-stringify-safe*, and others.

Whenever we use the NPM command line tool, *npm*, in order to install a package, NPM not only pulls the package itself, but also its dependencies, and installs those as well.

Using *npm install request* is simply a manual way to implicitly say "my project depends on *request*, please install it for me". However, there is an explicit way of defining dependencies for our own projects, which also allows to have all dependencies resolved and installed automatically.

In order to use this mechanism, we need to create a control file within our project that defines our dependencies. This control file is then used by NPM to resolve and install those dependencies.

This control file must be located at the top level folder of our project, and must be named *package.json*.

This is what a *package.json* file looks like for a project that depends on *request*:

```
{
  "dependencies": {
    "request": ""
  }
}
```

Having this file as part of our code base makes NPM aware of the dependencies of our project without the need to explicitly tell NPM what to install by hand.

We can now use NPM to automatically pull in all dependencies of our project, simply by executing *npm install* within the top level folder of our code base.

In this example, this doesn't look like much of a convenience win compared to manually installing *request*, but once we have more than a handful of dependencies in our projects, it really makes a difference.

The *package.json* file also allows us to “pin” dependencies to certain versions, i.e., we can define a version number for every dependency, which ensures that NPM won’t pull the latest version automatically, but exactly the version of a package we need:

```
{
  "dependencies": {
    "request": "2.27.0"
  }
}
```

In this case, NPM will always pull *request* version 2.27.0 (and the dependencies of this version), even if newer versions are available.

Patterns are possible, too:

```
{
  "dependencies": {
    "request": "2.27.x"
  }
}
```

The *x* is a placeholder for any number. This way, NPM would pull in *request* version 2.27.0 and 2.27.5, but not 2.28.0.

The official documentation at <https://npmjs.org/doc/json.html#dependencies> has more examples of possible dependency definitions.

Please note that the *package.json* file does much more than just defining dependencies. We will dig deeper in the course of this book.

For now, we are prepared to use NPM for resolving the dependencies that arise in our first project - our first test-driven Node.js application.

Test-Driven Node.js Development

The code examples in *The Node Beginner Book* only described a toy project, and we came away with not writing any tests for it. If writing tests is new for you, and you have not yet worked on software in a test-driven manner, then I invite you to follow along and give it a try.

We need to decide on a test framework that we will use to implement our tests. A lack of choice is not an issue in the JavaScript and Node.js world, as there are dozens of frameworks available. Personally, I prefer *Jasmine*, and will therefore use it for my examples.

Jasmine is a framework that follows the philosophy of *behaviour-driven development*, which is kind of a “subculture” within the community of test-driven developers. This topic alone could easily fill its own book, thus I’ll give only a brief introduction.

The idea is to begin development of a new software unit with its specification, followed by its implementation (which, by definition, must satisfy the specification).

Let’s make up a real world example: we order a table from a carpenter. We do so by *specifying* the end result:

“I need a table with a top that is 6 x 3 ft. The height of the top must be adjustable between 2.5 and 4.0 ft. I want to be able to adjust the top’s height without standing up from my chair. I want the table to be black, and cleaning it with a wet cloth should be possible without damaging the material. My budget is \$500.”

Such a specification allows to share a goal between us and the carpenter. We don’t have to care for how exactly the carpenter will achieve this goal. As long as the delivered product fits our specification, both of us can agree that the goal has been reached.

With a test-driven or behaviour-driven approach, this idea is applied to building software. You wouldn’t build a piece of software and then define what it’s supposed to do. You need to know in advance what you expect a unit of software to do. Instead of doing this vaguely and implicitly, a test-driven approach asks you to do the specification exactly and explicitly. Because we work on software, our specification can be software, too: we only need to write functions that check if our unit does what it is expected to do. These check functions are unit tests.

Let’s create a software unit which is covered by tests that describe its expected behaviour. In order to actually drive the creation of the code with the tests, let’s write the tests first. We then have a clearly defined goal: making the tests pass by implementing code that fulfills the expected behaviour, and nothing else.

In order to do so, we create a new Node.js project with two folders in it:

```
src/  
spec/
```

spec is where our test cases go - in Jasmine lingo, these are called “specifications”, hence “spec”. The *spec* folder mirrors the file structure under *src*, i.e., a source file at *src/foo.js* is mirrored by a specification at *spec/fooSpec.js*.

Following the tradition, we will test and implement a “Hello World” code unit. Its expected behaviour is to return a string “Hello Joe!” if called with “Joe” as its first and only parameter. This behaviour can be specified by writing a unit test.

To do so, we create a file *spec/greetSpec.js*, with the following content:

```
1  'use strict';  
2  
3  var greet = require('../src/greet');  
4  
5  describe('greet', function() {  
6  
7      it('should greet the given name', function() {  
8          expect(greet('Joe')).toEqual('Hello Joe!');  
9      });  
10  
11     it('should greet no-one special if no name is given', function() {  
12         expect(greet()).toEqual('Hello world!');  
13     });  
14  
15 });
```

This is a simple, yet complete specification. It is a programmatical description of the behaviour we expect from a yet-to-be-written function named *greet*.

The specification says that if the function *greet* is called with *Joe* as its first and only parameter, the return value of the function call should be the string “Hello Joe!”. If we don’t supply a name, the greeting should be generic.

As you can see, Jasmine specifications have a two-level structure. The top level of this structure is a *describe* block, which consists of one or more *it* blocks.

An *it* block describes a single expected behaviour of a single unit under test, and a *describe* block summarizes one or more blocks of expected behaviours, therefore completely specifying all expected behaviours of a unit.

Let’s illustrate this with a real-world “unit” described by a Jasmine specification:

```
describe('A candle', function() {

  it('should burn when lighted', function() {
    // ...
  });

  it('should grow smaller while burning', function() {
    // ...
  });

  it('should no longer burn when all wax has been burned', function() {
    // ...
  });

  it('should go out when no oxygen is available to it', function() {
    // ...
  });

});
```

As you can see, a Jasmine specification gives us a structure to fully describe how a given unit should behave.

Not only can we *describe* expected behaviour, we can also *verify* it. This can be done by running the test cases in the specification against the actual implementation.

After all, our Jasmine specification is just another piece of JavaScript code which can be executed. The NPM package *jasmine-node* ships with a test case runner which allows us to execute the test case, with the added benefit of a nice progress and result output.

Let's create a *package.json* file that defines *jasmine-node* as a dependency of our application - then we can start running the test cases of our specification.

As described earlier, we need to place the *package.json* file at the topmost folder of our project. Its content should be as follows:

```
{
  "devDependencies": {
    "jasmine-node": ""
  }
}
```

We talked about the *dependencies* section of *package.json* before - but here we declare *jasmine-node* in a *devDependencies* block. The result is basically the same: NPM knows about this dependency

and installs the package and its dependencies for us. However, dev dependencies are not needed to run our application - as the name suggests, they are only needed during development.

NPM allows to skip dev dependencies when deploying applications to a production system - we will get to this later.

In order to have NPM install *jasmine-node*, please run

```
npm install
```

in the top folder of your project.

We are now ready to test our application against its specification.

Of course, our *greet* function cannot fulfill its specification yet, simply because we have not yet implemented it. Let's see how this looks by running the test cases. From the root folder of our new project, execute the following:

```
./node_modules/jasmine-node/bin/jasmine-node spec/greetSpec.js
```

As you can see, Jasmine isn't too happy with the results yet. We refer to a Node module in *src/greet.js*, a file that doesn't even exist, which is why Jasmine bails out before even starting the tests:

```
Exception loading: spec/greetSpec.js  
{ [Error: Cannot find module '../src/greet'] code: 'MODULE_NOT_FOUND' }
```

Well, let's create the module, in file *src/greet.js*:

```
1 'use strict';  
2  
3 var greet = function() {};  
4  
5 module.exports = greet;
```

Now we have a general infrastructure, but of course we do not yet behave as the specification wishes. Let's run the test cases again:

FF

Failures:

1) greet should greet the given name

Message:

TypeError: object is not a function

Stacktrace:

TypeError: object is not a function

at null.<anonymous> (./spec/greetSpec.js:8:12)

2) greet should greet no-one special if no name is given

Message:

TypeError: object is not a function

Stacktrace:

TypeError: object is not a function

at null.<anonymous> (./spec/greetSpec.js:12:12)

Finished in 0.015 seconds

2 tests, 2 assertions, 2 failures, 0 skipped

Jasmine tells us that it executed two test cases that contained a total of two assertions (or expectations), and because these expectations could not be satisfied, the test run ended with two failures.

It's time to satisfy the first expectation of our specification, in file *src/greet.js*:

```
1 'use strict';
2
3 var greet = function(name) {
4   return 'Hello ' + name + '!';
5 };
6
7 module.exports = greet;
```

Another test case run reveals that we are getting closer:

```
.F
```

Failures:

1) greet should greet no-one special if no name is given

Message:

Expected 'Hello undefined!' to equal 'Hello world!'.

Stacktrace:

Error: Expected 'Hello undefined!' to equal 'Hello world!'.

at null.<anonymous> (spec/greetSpec.js:12:21)

Finished in 0.015 seconds

2 tests, 2 assertions, 1 failure, 0 skipped

Our first test case passes - *greet* can now correctly greet people by name. We still need to handle the case where no name was given:

```
1  'use strict';
2
3  var greet = function(name) {
4    if (name === undefined) {
5      name = 'world';
6    }
7    return 'Hello ' + name + '!';
8  };
9
10 module.exports = greet;
```

And that does the job:

```
..
```

Finished in 0.007 seconds

2 tests, 2 assertions, 0 failures, 0 skipped

We have now created a piece of software that behaves according to its specification.

You'll probably agree that our approach to create this unbelievably complex unit of software - the *greet* function - in a test-driven way doesn't prove the greatness of test-driven development in any way. That's not the goal of this chapter. It merely sets the stage for what's to come. We are going to create real, comprehensive software through the course of this book, and this is where the advantages of a test-driven approach can be experienced.

Object-oriented JavaScript

Let's talk about object-orientation and inheritance in JavaScript.

The good news is that it's actually quite simple, but the bad news is that it works completely different than object-orientation in languages like C++, Java, Ruby, Python or PHP, making it not-quite-so simple to understand.

But fear not, we are going to take it step by step.

Blueprints versus finger-pointing

Let's start by looking at how "typical" object-oriented languages actually create objects.

We are going to talk about an object called *myCar*. *myCar* is our bits-and-bytes representation of an incredibly simplified real world car. It could have attributes like *color* and *weight*, and methods like *drive* and *honk*.

In a real application, *myCar* could be used to represent the car in a racing game - but we are going to completely ignore the context of this object, because we will talk about the nature and usage of this object in a more abstract way.

If you would want to use this *myCar* object in, say, Java, you need to define the blueprint of this specific object first - this is what Java and most other object-oriented languages call a *class*.

If you want to create the object *myCar*, you tell Java to "build a new object after the specification that is laid out in the class *Car*".

The newly built object shares certain aspects with its blueprint. If you call the method *honk* on your object, like so:

```
myCar.honk();
```

then the Java VM will go to the class of *myCar* and look up which code it actually needs to execute, which is defined in the *honk* method of class *Car*.

Ok, nothing shockingly new here. Enter JavaScript.

A classless society

JavaScript does not have classes. But as in other languages, we would like to tell the interpreter that it should build our *myCar* object following a certain pattern or schema or blueprint - it would be quite tedious to create every *car* object from scratch, “manually” giving it the attributes and methods it needs every time we build it.

If we were to create 30 *car* objects based on the *Car* class in Java, this object-class relationship provides us with 30 cars that are able to drive and honk without us having to write 30 *drive* and *honk* methods.

How is this achieved in JavaScript? Instead of an object-class relationship, there is an object-object relationship.

Where in Java our *myCar*, asked to *honk*, says “go look at this class over there, which is my *blueprint*, to find the code you need”, JavaScript says “go look at that other object over there, which is my *prototype*, it has the code you are looking for”.

Building objects via an object-object relationship is called Prototype-based programming, versus Class-based programming used in more traditional languages like Java.

Both are perfectly valid implementations of the object-oriented programming paradigm - it’s just two different approaches.

Creating objects

Let’s dive into code a bit, shall we? How could we set up our code in order to allow us to create our *myCar* object, ending up with an object that is a *Car* and can therefore *honk* and *drive*?

Well, in the most simple sense, we can create our object completely from scratch, or *ex nihilo* if you prefer the boaster expression.

It works like this:

```
1 var myCar = {};  
2  
3 myCar.honk = function() {  
4   console.log('honk honk');  
5 };  
6  
7 myCar.drive = function() {  
8   console.log('vrooom...');  
9 };
```

This gives us an object called *myCar* that is able to *honk* and *drive*:

```
myCar.honk(); // outputs "honk honk"  
myCar.drive(); // outputs "vrooom..."
```

However, if we were to create 30 cars this way, we would end up defining the honk and drive behaviour of every single one, something we said we want to avoid.

In real life, if we made a living out of creating, say, pencils, and we don't want to create every pencil individually by hand, then we would consider building a pencil-making machine, and have this machine create the pencils for us.

After all, that's what we implicitly do in a class-based language like Java - by defining a class `Car`, we get the car-maker for free:

```
Car myCar = new Car();
```

will build the `myCar` object for us based on the `Car` blueprint. Using the `new` keyword does all the magic for us.

JavaScript, however, leaves the responsibility of building an object creator to us. Furthermore, it gives us a lot of freedom regarding the way we actually build our objects.

In the most simple case, we can write a function which creates “plain” objects that are exactly like our “ex nihilo” object, and that don't really share any behaviour - they just happen to roll out of the factory with the same behaviour copied onto every single one, if you want so.

Or, we can write a special kind of function that not only creates our objects, but also does some behind-the-scenes magic which links the created objects with their creator. This allows for a true sharing of behaviour: functions that are available on all created objects point to a single implementation. If this function implementation changes after objects have been created, which is possible in JavaScript, the behaviour of all objects sharing the function will change accordingly.

Let's examine all possible ways of creating objects in detail.

Using a simple function to create plain objects

In our first example, we created a plain `myCar` object out of thin air - we can simply wrap the creation code into a function, which gives us a very basic object creator:

```
1 var makeCar = function() {  
2   var newCar = {};  
3   newCar.honk = function() {  
4     console.log('honk honk');  
5   };  
6 };
```

For the sake of brevity, the drive function has been omitted.

We can then use this function to mass-produce cars:

```
1 var makeCar = function() {
2   var newCar = {}
3   newCar.honk = function() {
4     console.log('honk honk');
5   };
6   return newCar;
7 };
8
9 myCar1 = makeCar();
10 myCar2 = makeCar();
11 myCar3 = makeCar();
```

One downside of this approach is efficiency: for every *myCar* object that is created, a new *honk* function is created and attached - creating 1,000 objects means that the JavaScript interpreter has to allocate memory for 1,000 functions, although they all implement the same behaviour. This results in an unnecessarily high memory footprint of the application.

Secondly, this approach deprives us of some interesting opportunities. These *myCar* objects don't share anything - they were built by the same creator function, but are completely independent from each other.

It's really like with real cars from a real car factory: They all look the same, but once they leave the assembly line, they are totally independent. If the manufacturer should decide that pushing the horn on already produced cars should result in a different type of honk, all cars would have to be returned to the factory and modified.

In the virtual universe of JavaScript, we are not bound to such limits. By creating objects in a more sophisticated way, we are able to magically change the behaviour of all created objects at once.

Using a constructor function to create objects

In JavaScript, the entities that create objects with shared behaviour are functions which are called in a special way. These special functions are called *constructors*.

Let's create a constructor for cars. We are going to call this function *Car*, with a capital C, which is common practice to indicate that this function is a constructor.



In a way, this makes the constructor function a class, because it does some of the things a class (with a constructor method) does in a traditional OOP language. However, the approach is not identical, which is why constructor functions are often called *pseudo-classes* in JavaScript. I will simply call them classes or constructor functions.

Because we are going to encounter two new concepts that are both necessary for shared object behaviour to work, we are going to approach the final solution in two steps.

Step one is to recreate the previous solution (where a common function spilled out independent car objects), but this time using a constructor:

```
1 var Car = function() {
2   this.honk = function() {
3     console.log('honk honk');
4   };
5 };
```

When this function is called using the *new* keyword, like so:

```
var myCar = new Car();
```

it implicitly returns a newly created object with the *honk* function attached.

Using *this* and *new* makes the explicit creation and return of the new object unnecessary - it is created and returned “behind the scenes” (i.e., the *new* keyword is what creates the new, “invisible” object, and secretly passes it to the *Car* function as its *this* variable).

You can think of the mechanism at work a bit like in this pseudo-code:

```
1 // Pseudo-code, for illustration only!
2
3 var Car = function(this) {
4   this.honk = function() {
5     console.log('honk honk');
6   };
7   return this;
8 };
9
10 var newObject = {};
11 var myCar = Car(newObject);
```

As said, this is more or less like our previous solution - we don't have to create every car object manually, but we still cannot modify the *honk* behaviour only once and have this change reflected in all created cars.

But we laid the first cornerstone for it. By using a constructor, all objects received a special property that links them to their constructor:


```
1 var Car = function() {
2   this.honk = function() {
3     console.log('honk honk');
4   };
5 };
6
7 var myCar1 = new Car();
8 var myCar2 = new Car();
9
10 console.log(myCar1.constructor); // outputs [Function: Car]
11 console.log(myCar2.constructor); // outputs [Function: Car]
```

All created *myCars* are linked to the *Car* constructor. This is what actually makes them a *class* of related objects, and not just a bunch of objects that happen to have similar names and identical functions.

Now we have finally reached the moment to get back to the mysterious prototype we talked about in the introduction.

Using prototyping to efficiently share behaviour between objects

As stated there, while in class-based programming the class is the place to put functions that all objects will share, in prototype-based programming, the place to put these functions is the object which acts as the prototype for our objects at hand.

But where is the object that is the prototype of our *myCar* objects - we didn't create one!

It has been implicitly created for us, and is assigned to the *Car.prototype* property (in case you wondered, JavaScript functions are objects too, and they therefore have properties).

Here is the key to sharing functions between objects: Whenever we call a function on an object, the JavaScript interpreter tries to find that function within the queried object. But if it doesn't find the function within the object itself, it asks the object for the pointer to its prototype, then goes to the prototype, and asks for the function there. If it is found, it is then executed.

This means that we can create *myCar* objects without any functions, create the *honk* function in their prototype, and end up having *myCar* objects that know how to honk - because everytime the interpreter tries to execute the *honk* function on one of the *myCar* objects, it will be redirected to the prototype, and execute the *honk* function which is defined there.

Here is how this setup can be achieved:

```
1  var Car = function() {};  
2  
3  Car.prototype.honk = function() {  
4    console.log('honk honk');  
5  };  
6  
7  var myCar1 = new Car();  
8  var myCar2 = new Car();  
9  
10 myCar1.honk(); // executes Car.prototype.honk() and outputs "honk honk"  
11 myCar2.honk(); // executes Car.prototype.honk() and outputs "honk honk"
```

Our constructor is now empty, because for our very simple cars, no additional setup is necessary.

Because both *myCars* are created through this constructor, their prototype points to *Car.prototype* - executing *myCar1.honk()* and *myCar2.honk()* always results in *Car.prototype.honk()* being executed.

Let's see what this enables us to do. In JavaScript, objects can be changed at runtime. This holds true for prototypes, too. Which is why we can change the *honk* behaviour of all our cars even after they have been created:

```
1  var Car = function() {};  
2  
3  Car.prototype.honk = function() {  
4    console.log('honk honk');  
5  };  
6  
7  var myCar1 = new Car();  
8  var myCar2 = new Car();  
9  
10 myCar1.honk(); // executes Car.prototype.honk() and outputs "honk honk"  
11 myCar2.honk(); // executes Car.prototype.honk() and outputs "honk honk"  
12  
13 Car.prototype.honk = function() {  
14   console.log('meep meep');  
15 };  
16  
17 myCar1.honk(); // executes Car.prototype.honk() and outputs "meep meep"  
18 myCar2.honk(); // executes Car.prototype.honk() and outputs "meep meep"
```

Of course, we can also add additional functions at runtime:

```
1  var Car = function() {};  
2  
3  Car.prototype.honk = function() {  
4    console.log('honk honk');  
5  };  
6  
7  var myCar1 = new Car();  
8  var myCar2 = new Car();  
9  
10 Car.prototype.drive = function() {  
11   console.log('vrooom...');  
12 };  
13  
14 myCar1.drive(); // executes Car.prototype.drive() and outputs "vrooom..."  
15 myCar2.drive(); // executes Car.prototype.drive() and outputs "vrooom..."
```

But we could even decide to treat only one of our cars differently:

```
1  var Car = function() {};  
2  
3  Car.prototype.honk = function() {  
4    console.log('honk honk');  
5  };  
6  
7  var myCar1 = new Car();  
8  var myCar2 = new Car();  
9  
10 myCar1.honk(); // executes Car.prototype.honk() and outputs "honk honk"  
11 myCar2.honk(); // executes Car.prototype.honk() and outputs "honk honk"  
12  
13 myCar2.honk = function() {  
14   console.log('meep meep');  
15 };  
16  
17 myCar1.honk(); // executes Car.prototype.honk() and outputs "honk honk"  
18 myCar2.honk(); // executes myCar2.honk() and outputs "meep meep"
```

It's important to understand what happens behind the scenes in this example. As we have seen, when calling a function on an object, the interpreter follows a certain path to find the actual location of that function.

While for *myCar1*, there still is no *honk* function within that object itself, that no longer holds true for *myCar2*. When the interpreter calls *myCar2.honk()*, there now is a function within *myCar2* itself.

Therefore, the interpreter no longer follows the path to the prototype of *myCar2*, and executes the function within *myCar2* instead.

That's one of the major differences to class-based programming: while objects are relatively "rigid" e.g. in Java, where the structure of an object cannot be changed at runtime, in JavaScript, the prototype-based approach links objects of a certain class more loosely together, which allows to change the structure of objects at any time.

Also, note how sharing functions through the constructor's prototype is way more efficient than creating objects that all carry their own functions, even if they are identical. As previously stated, the engine doesn't know that these functions are meant to be identical, and it has to allocate memory for every function in every object. This is no longer true when sharing functions through a common prototype - the function in question is placed in memory exactly once, and no matter how many *myCar* objects we create, they don't carry the function themselves, they only refer to their constructor, in whose prototype the function is found.

To give you an idea of what this difference can mean, here is a very simple comparison. The first example creates 1,000,000 objects that all have the function directly attached to them:

```
1 var C = function() {
2   this.f = function(foo) {
3     console.log(foo);
4   };
5 };
6
7 var a = [];
8 for (var i = 0; i < 1000000; i++) {
9   a.push(new C());
10 }
```

In Google Chrome, this results in a heap snapshot size of 328 MB. Here is the same example, but now the function is shared through the constructor's prototype:

```
1 var C = function() {};
2
3 C.prototype.f = function(foo) {
4   console.log(foo);
5 };
6
7 var a = [];
8 for (var i = 0; i < 1000000; i++) {
9   a.push(new C());
10 }
```

This time, the size of the heap snapshot is only 17 MB, i.e., only about 5% of the non-efficient solution.

Object-orientation, prototyping, and inheritance


So far, we haven't talked about inheritance in JavaScript, so let's do this now.

It's useful to share behaviour within a certain class of objects, but there are cases where we would like to share behaviour between different, but similar classes of objects.

Imagine our virtual world not only had cars, but also bikes. Both drive, but where a car has a horn, a bike has a bell.

Being able to drive makes both objects vehicles, but not sharing the *honk* and *ring* behaviour distinguishes them.

We could illustrate their shared and local behaviour as well as their relationship to each other as follows:



Designing this relationship in a class-based language like Java is straightforward: We would define a class *Vehicle* with a method *drive*, and two classes *Car* and *Bike* which both extend the *Vehicle* class, and implement a *honk* and a *ring* method, respectively.

This would make the *car* as well as *bike* objects inherit the *drive* behaviour through the inheritance of their classes.

How does this work in JavaScript, where we don't have classes, but prototypes?

Let's look at an example first, and then dissect it. To keep the code short for now, let's only start with a car that inherits from a vehicle:

```

1  var Vehicle = function() {};
2
3  Vehicle.prototype.drive = function() {
4    console.log('vroom...');
5  };
6
7
8  var Car = function() {};
9
10 Car.prototype = new Vehicle();
  
```

```
11
12 Car.prototype.honk = function() {
13   console.log('honk honk');
14 };
15
16
17 var myCar = new Car();
18
19 myCar.honk(); // outputs "honk honk"
20 myCar.drive(); // outputs "vrooom..."
```

In JavaScript, inheritance runs through a chain of prototypes.

The prototype of the *Car* constructor is set to a newly created vehicle object, which establishes the link structure that allows the interpreter to look for methods in “parent” objects.

The prototype of the *Vehicle* constructor has a function *drive*. Here is what happens when the *myCar* object is asked to *drive()*:

- The interpreter looks for a *drive* method within the *myCar* object, which does not exist
- The interpreter then asks the *myCar* object for its prototype, which is the prototype of its constructor *Car*
- When looking at *Car.prototype*, the interpreter sees a *vehicle* object which has a function *honk* attached, but no *drive* function
- Thus, the interpreter now asks this *vehicle* object for its prototype, which is the prototype of its constructor *Vehicle*
- When looking at *Vehicle.prototype*, the interpreter sees an object which has a *drive* function attached - the interpreter now knows which code implements the *myCar.drive()* behaviour, and executes it

A classless society, revisited

We just learned how to emulate the traditional OOP inheritance mechanism. But it's important to note that in JavaScript, that is only one valid approach to create objects that are related to each other.

It was Douglas Crockford who came up with another clever solution, which allows objects to inherit from each other directly. It's a native part of JavaScript by now - it's the *Object.create()* function, and it works like this:

```
1 Object.create = function(o) {
2   var F = function() {};
3   F.prototype = o;
4   return new F();
5 };
```

We learned enough now to understand what's going on. Let's analyze an example:

```
1 var vehicle = {};
2 vehicle.drive = function () {
3   console.log('vroom...');
4 };
5
6 var car = Object.create(vehicle);
7 car.honk = function() {
8   console.log('honk honk');
9 };
10
11 var myCar = Object.create(car);
12
13 myCar.honk(); // outputs "honk honk"
14 myCar.drive(); // outputs "vroom..."
```

While being more concise and expressive, this code achieves exactly the same behaviour, without the need to write dedicated constructors and attaching functions to their prototype. As you can see, *Object.create()* handles both behind the scenes, on the fly. A temporary constructor is created, its prototype is set to the object that serves as the role model for our new object, and a new object is created from this setup.

Conceptually, this is really the same as in the previous example where we defined that *Car.prototype* shall be a *new Vehicle()*;

But wait! We created the functions *drive* and *honk* within our objects, not on their prototypes - that's memory-inefficient!

Well, in this case, it's actually not. Let's see why:

```
1  var vehicle = {};  
2  vehicle.drive = function () {  
3    console.log('vrooom...');  
4  };  
5  
6  var car = Object.create(vehicle);  
7  car.honk = function() {  
8    console.log('honk honk');  
9  };  
10  
11 var myVehicle = Object.create(vehicle);  
12 var myCar1 = Object.create(car);  
13 var myCar2 = Object.create(car);  
14  
15 myCar1.honk(); // outputs "honk honk"  
16 myCar2.honk(); // outputs "honk honk"  
17  
18 myVehicle.drive(); // outputs "vrooom..."  
19 myCar1.drive(); // outputs "vrooom..."  
20 myCar2.drive(); // outputs "vrooom..."
```

We have now created a total of 5 objects, but how often do the *honk* and *drive* methods exist in memory? Well, how often have they been defined? Just once - and therefore, this solution is basically as efficient as the one where we built the inheritance manually. Let's look at the numbers:

```
1  var c = {};  
2  c.f = function(foo) {  
3    console.log(foo);  
4  };  
5  
6  var a = [];  
7  for (var i = 0; i < 1000000; i++) {  
8    a.push(Object.create(c));  
9  }
```

Turns out, it's not *exactly* identical - we end up with a heap snapshot size of 40 MB, thus there seems to be some overhead involved. However, in exchange for much better code, this is probably more than worth it.

Summary

By now, it's probably clear what the main difference between classical OOP languages and JavaScript is, conceptually: While classical languages like Java provide *one* way to manage object creation and

behaviour sharing (through classes and inheritance), and this way is enforced by the language and “baked in”, JavaScript starts at a slightly lower level and provides building blocks that allow us to create several different mechanisms for this.

Whether you decide to use these building blocks to recreate the traditional class-based pattern, let your objects inherit from each other directly, with the concept of classes getting in the way, or if you don't use the object-oriented paradigm at all and just solve the problem at hand with pure functional code: JavaScript gives you the freedom to choose the best methodology for any situation.

Synchronous and Asynchronous operations explained


Visualizing the Node.js execution model

For the chapters that follow it's important to fully understand what it means, conceptually, that a Node.js application has synchronous and asynchronous operations, and how both operations interact with each other.

Let's try to build this understanding step by step.

The first concept that we need to understand is that of the Node.js event loop. The event loop is the execution model of a running Node.js application.

We can visualize this model as a row of loops:



I've drawn boxes because circles look really clumsy in ascii art. So, these here look like rectangles, but please imagine them as circles - circles with an arrow, which means that one circle represents one iteration through the event loop.

Another visualization could be the following pseudo-code:

```
while (I still have stuff to do) {
  do stuff;
}
```

Conceptually, at the very core of it, it's really that simple: Node.js starts, loads our application, and then loops until there is nothing left to do - at which point our application terminates.

What kind of stuff is happening inside one loop iteration? Let's first look at a very simple example, a Hello World application like this:

```
console.log('Hello');  
console.log('World');
```

This is the visualization of the execution of this application in our ascii art:

```
+----> -----+  
|               |  
| Write         |  
| 'Hello'       |  
| to the screen |  
|               |  
| Write         |  
| 'World'       |  
| to the screen |  
|               |  
+-----+-----+
```

Yep, that's it: Just one iteration, and then, exit the application.


The things we asked our application to do - writing text to the screen, and then writing another text to the screen, using *console.log* - are *synchronous* operations. They both happen within the same (and in this case, only) iteration through the event loop.

Let's look at the model when we bring asynchronous operations into the game, like this:

```
console.log('Hello');  
  
setTimeout(function() {  
  console.log('World');  
}, 1000);
```

This still prints *Hello* and then *World* to the screen, but the second text is printed with a delay of 1000 ms.

setTimeout, you may have guessed it, is an *asynchronous* operation. We pass the code to be executed in the body of an anonymous function - the so-called *callback function*. Why do we do so? The visualization helps to understand why:



This, again at the very core of it, is what calling an asynchronous function does: It starts an operation *outside the event loop*. Conceptually, Node.js starts the asynchronous operation and makes a mental note that when this operations triggers an event, the anonymous function that was passed to the operation needs to be called.

Hence, the *event loop*: as long as asynchronous operations are ongoing, the Node.js process loops, waiting for events from those operations. As soon as no more asynchronous operations are ongoing, the looping stops and our application terminates.



Note that the visualization isn't detailed enough to show that Node.js checks for outside events *between* loop iterations.

Just to be clear: callback functions don't need to be anonymous inline functions:

```

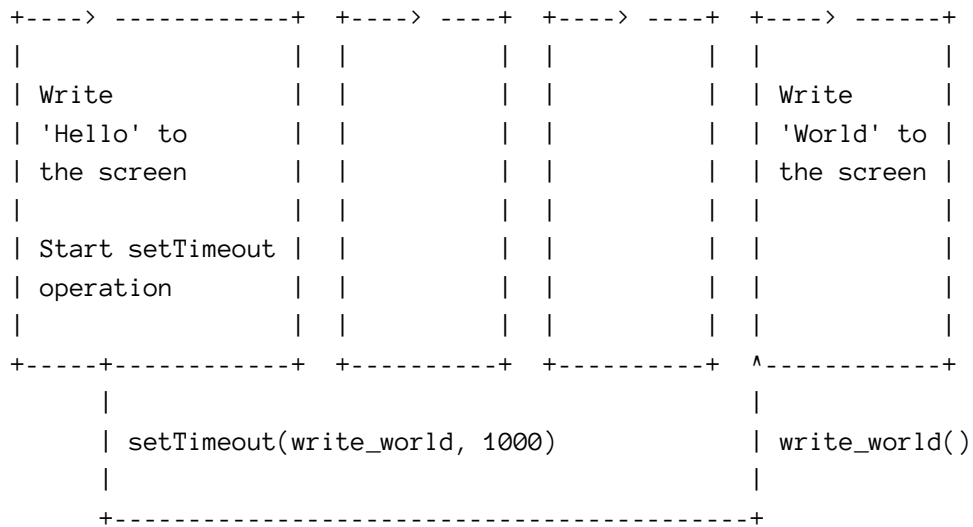
var write_world = function() {
  console.log('World');
};

console.log('Hello');

setTimeout(write_world, 1000);

```

results in:



As time goes by...

It's just that more often than not, declaring a named function and passing it as the callback isn't worth the hassle, because very often we need the functionality described in the function only once.

Let's look at a slightly more interesting asynchronous operation, the *fs.stat* call - it starts an IO operation that looks at the file system information for a given path, that is, stuff like the inode number and the size etc.:


```

var fs = require('fs');

fs.stat('/etc/passwd', function(err, stats) {
  console.dir(stats);
});

```

The visualization:




```
console.log('Hello World');
```

results in a blocking operation - while the Node.js process writes the text to the screen, this is the *only* thing that it does. There is only one single piece of JavaScript code that can be executed within the event loop at any given time.

The reason this doesn't result in a problem in the case of a `console.log()` is that it is an extremely cheap operation. In comparison, even the most simple IO operations are way more expensive. Whenever the network or harddrives (yes, including SSDs) are involved, expect things to be *incredibly much slower* compared to operations where only the CPU and RAM are involved, like `var a = 2 * 21`. Just how much slower? The following table gives you a good idea - it shows actual times for different kind of computer operations, and shows how they relate to each other compared to how time spans that human beings experience relate to each other:

1 CPU cycle	0.3 ns	1 s
Level 1 cache access	0.9 ns	3 s
Level 2 cache access	2.8 ns	9 s
Level 3 cache access	12.9 ns	43 s
Main memory access	120 ns	6 min
Solid-state disk I/O	50-150 μ s	2-6 days
Rotational disk I/O	1-10 ms	1-12 months
Internet: SF to NYC	40 ms	4 years
Internet: SF to UK	81 ms	8 years
Internet: SF to Australia	183 ms	19 years
OS virtualization reboot	4 s	423 years
SCSI command time-out	30 s	3000 years
Hardware virtualization reboot	40 s	4000 years
Physical system reboot	5 m	32 millenia

So, the difference between setting a variable in your code and reading even a tiny file from disk is like the difference between preparing a sandwich and going on vacation for a week. You can prepare a lot of sandwiches during one week of vacation.

And that's the sole reason why all those Node.js functions that result in IO operations also happen to work asynchronously: it's because the event loop needs being kept free of any long-running operations during which the Node.js process would practically stall, which would result in unresponsive applications.

Just to be clear: we can easily stall the event loop even if no IO operations are involved and we only use cheap synchronous functions - if only there are enough of them within one loop iteration. Take the following code for example:

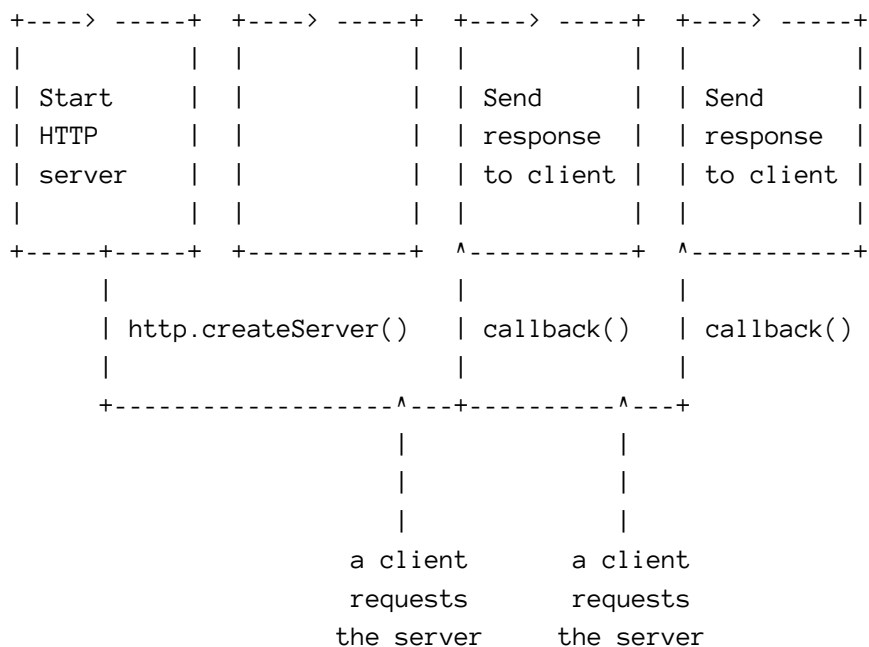
```
var http = require('http');

http.createServer(function(request, response) {
  console.log('Handling HTTP request');
  response.writeHead(200, {'Content-Type': 'text/html'});
  response.write('<b>Hello World</b>');
  response.end();
}).listen(8080);

var a;
for (var i=0; i < 10000000000; i += 1) {
  a = i;
}

console.log('For loop has finished');
```

It's our minimalistic web server again, plus a for loop with 10,000,000,000 iterations. Our event loop visualization basically looks the same:



But here is the problem: The HTTP server will start listening in the background as an asynchronous operation, but we will stay within the first loop operation as long as the *for* loop is running. And although we only have a very cheap operation happening within the for loop, it happens 10,000,000,000 times, and on my machine, this takes around 20 seconds.

When you start the server application and then open your browser at `http://localhost:8080/`, you won't get an answer right away. External events, like our HTTP request, are only handled between one loop iteration and the next; however, the first loop iteration takes 20 seconds because of our *for* loop, and only then Node.js switches to the next iteration and has a chance to handle our request by calling the HTTP server callback.



As you will see, the application will print *For loop has finished*, and right after that, it will answer the HTTP request and print *Handling HTTP request*. This demonstrates how external events from asynchronous operations are handled at the beginning of a new event loop iteration.

You've probably heard time and again how Node.js isn't suited for writing applications with CPU intensive tasks - as we can see, the reason for this is the event loop model.

From this, we can distill the two most important rules for writing responsive Node.js applications:

- Handle IO-intensive operations through asynchronous operations
- Keep your own code (that is, everything that happens synchronously within event loop iterations) as lean as possible



This leaves the question: what are sensible solutions if you *have* to do expensive CPU-bound operations within your JavaScript code? As we will learn in later chapters, we can mitigate the problem that Node.js itself simply isn't particularly suited for these kinds of operations.

Using and creating Event Emitters

Introduction

By now, you are probably more than familiar with this:

```
someFunction(function(err) {  
  if (!err) {  
    console.log('Hey, looks like someFunction has finished and called me.');  } else {  
    console.log('Oh no, something terrible happened!');  }  
});
```

We call a function, *someFunction* in this case, which does something asynchronously in the background, and calls the anonymous function we passed in (the *callback*) once it has finished, passing an *Error* object if something went wrong, or *null* if all went fine. That's the standard callback pattern, and you will encounter it regularly when writing Node.js code.

For simple use cases, where we call a function that finishes its task some time in the future, either successfully or with an error, this pattern is just fine.

For more complex use cases, this pattern does not scale well. There are cases where we call a function which results in multiple events happening over time, and also, different types of events might happen.

One example is reading the contents of a file through a *ReadStream*. The details of handling files are discussed in a later chapter, but we will use this example to illustrate event emitters.

This snippet demonstrates how to read data from a large file using a *ReadStream*:

```
'use strict';  
  
var fs = require('fs');  
  
fs.createReadStream('/path/to/large/file');
```

When reading data from a file, two different things can happen: We either receive content, or we reach the end of the file. Both cases *could* be handled using the callback pattern - for example, by using one callback with two parameters, one with the data and another one that is *false* as long as

the end of the file has not been reached, and *true* once it has been reached. Or we could provide two separate callback functions, one that is called when content is retrieved and one that is called when the end of the file has been reached.

But there is a more elegant way. Instead of working with classical callbacks, *createReadStream* allows us to use an *Event Emitter*. That's a special object which can be used to attach callback functions to different events. Using it looks like this:

```
1  'use strict';
2
3  var fs = require('fs');
4
5  var stream = fs.createReadStream('/path/to/large/file');
6
7  stream.on('data', function(data) {
8    console.log('Received data: ' + data);
9  });
10
11 stream.on('end', function() {
12   console.log('End of file has been reached');
13 });
```

Here is what happens in detail:

- On line 5, we create a read stream that will start to retrieve the contents of file */path/to/large/file*. The call to *fs.createReadStream* does not take a function argument to use it as a callback. Instead, it returns an object, which we assign as the value of the variable *stream*.
- On line 7, we attach a callback to one type of events the *ReadStream* emits: *data* events
- On line 11, we attach another callback to another type of event the *ReadStream* emits: the *end* event

The object that is returned by *fs.createReadStream* is an *Event Emitter*. These objects allow us to attach different callbacks to different events while keeping our code readable and sane.

A *ReadStream* retrieves the contents of a file in chunks, which is more efficient than to load the whole data of potentially huge files into memory at once in one long, blocking operation.

Because of this, the *data* event will be emitted multiple times, depending on the size of the file. The callback that is attached to this event will therefore be called multiple times.

When all content has been retrieved, the *end* event is fired once, and no other events will be fired from then on. The *end* event callback is therefore the right place to do whatever we want to do after we retrieved the complete file content. In practice, this would look like this:

```
1  'use strict';
2
3  var fs = require('fs');
4
5  var stream = fs.createReadStream('/path/to/large/file');
6
7  var content = '';
8
9  stream.on('data', function(data) {
10     content = content + data;
11 });
12
13 stream.on('end', function() {
14     console.log('File content has been retrieved: ' + content);
15 });
```



It doesn't make too much sense to efficiently read a large file's content in chunks, only to assign the whole data to a variable and therefore using the memory anyways. In a real application, we would read a file in chunks to, for example, send every chunk to a web client that is downloading the file via HTTP. We will talk about this in more detail in a later chapter.

The *Event Emitter* pattern itself is simple and clearly defined: a function returns an event emitter object, and using this object's *on* method, callbacks can be attached to events.

However, there is no strict rule regarding the events themselves: an event *type*, like *data* or *end*, is just a string, and the author of an event emitter can define any name she wants. Also, it's not defined what arguments are passed to the callbacks that are triggered through an event - the author of the event emitter should define this through some kind of documentation.

There is one recurring pattern, at least for the internal Node.js modules, and that is the *error* event: Most event emitters emit an event called "error" whenever an error occurs, and if we don't listen to this event, the event emitter will raise an exception.

You can easily test this by running the above code: as long as you don't happen to have a file at */path/to/large/file*, Node.js will bail out with this message:

```
events.js:72
    throw er; // Unhandled 'error' event
      ^
Error: ENOENT, open '/path/to/large/file'
```

But if you attach a callback to the *error* event, you can handle the error yourself:

```
1  'use strict';
2
3  var fs = require('fs');
4
5  var stream = fs.createReadStream('/path/to/large/file');
6
7  var content = '';
8
9  stream.on('error', function(err) {
10     console.log('Sad panda: ' + err);
11 });
12
13 stream.on('data', function(data) {
14     content = content + data;
15 });
16
17 stream.on('end', function() {
18     console.log('File content has been retrieved: ' + content);
19 });
```

Instead of using *on*, we can also attach a callback to an event using *once*. Callbacks that are attached this way will be called the first time that the event occurs, but will then be removed from the list of event listeners and not be called again:

```
stream.once('data', function(data) {
    console.log('I have received the first chunk of data');
});
```

Also, it's possible to detach an attached callback manually. This only works with named callback functions:

```
var callback = function(data) {
    console.log('I have received a chunk of data: ' + data);
}

stream.on('data', callback);

stream.removeListener('data', callback);
```

And last but not least, you can remove all attached listeners for a certain event:

```
stream.removeAllListeners('data');
```

Creating your own Event Emitter object

We can create event emitters ourselves. This is even supported by Node.js by inheriting from the built-in `events.EventEmitter` class. But let's first implement a simple event emitter from scratch, because this explains the pattern in all its details.

For this, we are going to create a module whose purpose is to regularly watch for changes in the size of a file. Once implemented, it can be used like this:

```
'use strict';

watcher = new FilesizeWatcher('/path/to/file');

watcher.on('error', function(err) {
  console.log('Error watching file:', err);
});

watcher.on('grew', function(gain) {
  console.log('File grew by', gain, 'bytes');
});

watcher.on('shrank', function(loss) {
  console.log('File shrank by', loss, 'bytes');
});

watcher.stop();
```

As you can see, the module consists of a class `FilesizeWatcher` which can be instantiated with a file path and returns an event emitter. We can listen to three different events from this emitter: `error`, `grew`, and `shrank`.

Let's start by writing a spec that describes how we expect to use our event emitter. To do so, create a new project directory, and add the following `package.json`:

```
{
  "devDependencies": {
    "jasmine-node": ""
  }
}
```

Afterwards, run `npm install`.

Now create a file `FilesizeWatcherSpec.js`, with the following content:

```
1  'use strict';
2
3  var FilesizeWatcher = require('./FilesizeWatcher');
4  var exec = require('child_process').exec;
5
6  describe('FilesizeWatcher', function() {
7
8    var watcher;
9
10   afterEach(function() {
11     watcher.stop();
12   });
13
14   it('should fire a "grew" event when the file grew in size', function(done) {
15
16     var path = '/var/tmp/filesizewatcher.test';
17     exec('rm -f ' + path + ' ; touch ' + path, function() {
18       watcher = new FilesizeWatcher(path);
19
20       watcher.on('grew', function(gain) {
21         expect(gain).toBe(5);
22         done();
23       });
24
25       exec('echo "test" > ' + path, function() {});
26
27     });
28
29   });
30
31   it('should fire a "shrank" event when the file grew in size', function(done) {
32
33     var path = '/var/tmp/filesizewatcher.test';
34     exec('rm -f ' + path + ' ; echo "test" > ' + path, function() {
35       watcher = new FilesizeWatcher(path);
36
37       watcher.on('shrank', function(loss) {
38         expect(loss).toBe(3);
39         done();
40       });
41
42       exec('echo "a" > ' + path, function() {});
```

```
43
44     });
45
46 });
47
48 it('should fire "error" if path does not start with a slash', function(done) {
49
50     var path = 'var/tmp/filesizewatcher.test';
51     watcher = new FilesizeWatcher(path);
52
53     watcher.on('error', function(err) {
54         expect(err).toBe('Path does not start with a slash');
55         done();
56     });
57
58 });
59
60 });
```



Because this is just an example application, we will not create a *spec* and a *src* directory, but instead just put both the specification file and the implementation file in the top folder of our project.

Before we look at the specification itself in detail, let's discuss the *done()* call we see in each of the *it* blocks. The *done* function is a callback that is passed to the function parameter of an *it* block by Jasmine.

This pattern is used when testing asynchronous operations. Our emitter emits events asynchronously, and Jasmine cannot know by itself when events will fire. It needs our help by being told “now the asynchronous operation I expected to occur did actually occur” - and this is done by triggering the callback.

Now to the specification itself. The first expectation is that when we write “test” into our testfile, the *grew* event is fired, telling us that the file gained 5 bytes in size.



Note how we use the *exec* function from the *child_process* module to manipulate our test file through shell commands within the specification.

Next, we specify the behaviour that is expected if the monitored test file shrinks in size: the *shrank* event must fire and report how many bytes the file lost.

At last, we specify that if we ask the watcher to monitor a file path that doesn't start with a slash, an error event must be emitted.



I'm creating a very simplified version of a file size watcher here for the sake of brevity - for a realworld implementation, more sophisticated checks would make sense.

We will create two different implementations which both fulfill this specification.

First, we will create a version of the file size watcher where we manage the event listener and event emitting logic completely ourselves. This way, we experience first hand how the event emitter pattern works.

Afterwards, we will implement a second version where we make use of existing Node.js functionality in order to implement the event emitter pattern without the need to reinvent the wheel.

The following shows a possible implementation of the first version, where we take care of the event listener callbacks ourselves:

```
1  'use strict';
2
3  var fs = require('fs');
4
5  var FilesizeWatcher = function(path) {
6    var self = this;
7
8    self.callbacks = {};
9
10   if (/^\w/.test(path) === false) {
11     self.callbacks['error']('Path does not start with a slash');
12     return;
13   }
14
15   fs.stat(path, function(err, stats) {
16     self.lastfilesize = stats.size;
17   });
18
19   self.interval = setInterval(
20     function() {
21       fs.stat(path, function(err, stats) {
22         if (stats.size > self.lastfilesize) {
23           self.callbacks['grew'](stats.size - self.lastfilesize);
24           self.lastfilesize = stats.size;
25         }
26         if (stats.size < self.lastfilesize) {
27           self.callbacks['shrank'](self.lastfilesize - stats.size);
28           self.lastfilesize = stats.size;
29         }

```

```
30     }, 1000);
31   });
32 };
33
34 FilesizeWatcher.prototype.on = function(eventType, callback) {
35   this.callbacks[eventType] = callback;
36 };
37
38 FilesizeWatcher.prototype.stop = function() {
39   clearInterval(this.interval);
40 };
41
42 module.exports = FilesizeWatcher;
```

Let's discuss this code.

On line 3, we load the *fs* module - we need its *stat* function to asynchronously retrieve file information.

On line 5 we start to build a constructor function for *FilesizeWatcher* objects. They are created by passing a path to watch as a parameter.

On line 6, we assign the object instance variable to a local *self* variable - this way we can access our instantiated object within callback functions, where *this* would point to another object.

We then create the *self.callbacks* object - we are going to use this as an associative array where we will store the callback to each event.

Next, on line 10, we check if the given path starts with a slash using a regular expression - if it doesn't, we trigger the callback associated with the *error* event.

If the check succeeds, we start an initial *stat* operation in order to store the file size of the given path - we need this base value in order to recognize future changes in file size.

The actual watch logic starts on line 19. We set up a 1-second interval where we call *stat* on every interval iteration and compare the current file size with the last known file size.

Line 22 handles the case where the file grew in size, calling the event handler callback associated with the *grew* event; line 26 handles the opposite case. In both cases, the new file size is saved.

Event handlers can be registered using the *FilesizeWatcher.on* method which is defined on line 34. In our implementation, all it does is to store the callback under the event name in our *callbacks* object.

Finally, line 38 defines the *stop* method which cancels the interval we set up in the constructor function.

Let's see if this implementation works by running `./node_modules/jasmine-node/bin/jasmine-node ./FilesizeWatcherSpec.js`:

```
..F
```

Failures:

1) FilesizeWatcher should fire "error" if the path does not start with a slash

Message:

```
TypeError: Object #<Object> has no method 'error'
```

Stacktrace:

```
TypeError: Object #<Object> has no method 'error'
```

```
at new FilesizeWatcher (FilesizeWatcher.js:11:28)
```

```
at null.<anonymous> (FilesizeWatcherSpec.js:51:15)
```

```
at null.<anonymous> (...mine-node/lib/jasmine-node/async-callback.js:45:37)
```

Finished in 0.087 seconds

3 tests, 3 assertions, 1 failure, 0 skipped

Well... nearly. The core functionality works: the *grew* and *shrank* events are fired as expected.

But the file path check makes problems. According to the message, the problem arises on line 11:

```
self.callbacks['error']('Path does not start with a slash');
```

The error message says that the *self.callbacks* object doesn't have a method called *error*. But in our specification, we defined a callback for this event, just as we did for the other events, on line 53 of *FilesizeWatcherSpec.js*:


```
watcher.on('error', function(err) {
```

The problem here is quite subtle and can only be understood if one remembers how the execution model of Node.js works. Let's recall the chapter on synchronous and asynchronous operations. Our specification code looks like this:

```
watcher = new FilesizeWatcher(path);
```

```
watcher.on('error', function(err) {  
  expect(err).toBe('Path does not start with a slash');  
  done();  
});
```

According to the visualization technique we used to explain the inner workings of the event loop model, this is what actually happens:



Now it becomes clear that the order in which things are happening is simply wrong: We first call the error callback, and then we attach it. This can't work.

When creating the *watcher* object, we call the *FilesizeWatcher* constructor function. Up to the point where it checks the path and calls the *error* event handler callback, everything happens synchronously:

```

var FilesizeWatcher = function(path) {
  var self = this;

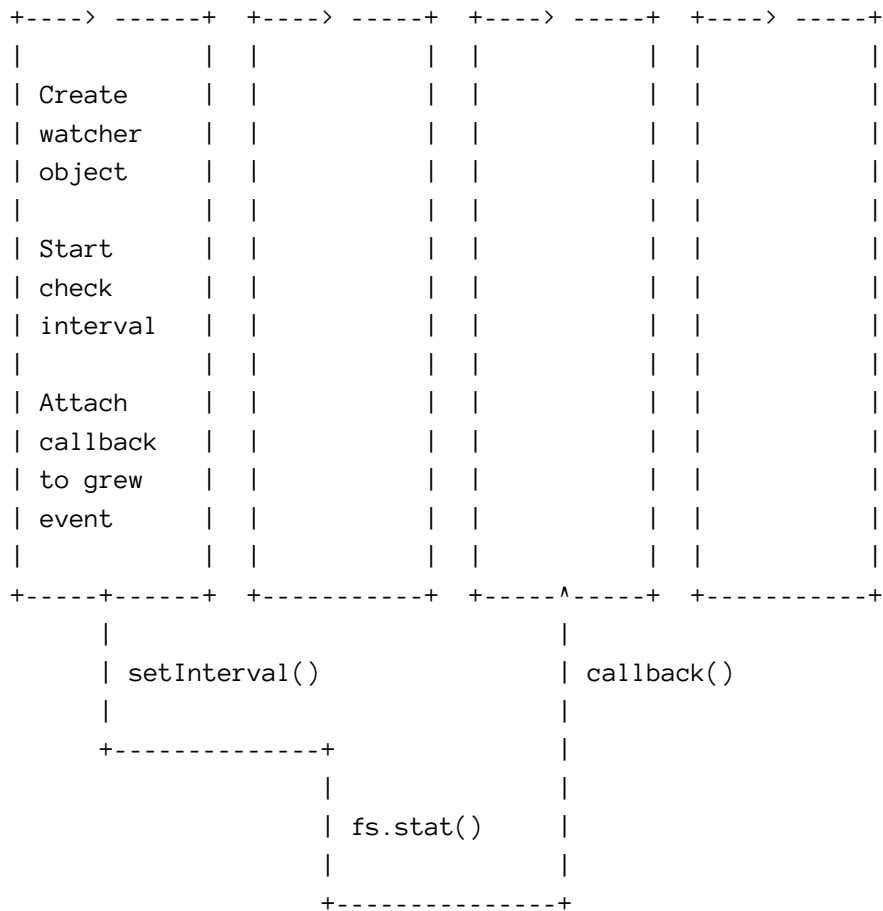
  self.callbacks = {};

  if (/^\./.test(path) === false) {
    self.callbacks['error']('Path does not start with a slash');
    return;
  }
}

```

The constructor function returns and Node.js execution continues with the main code, all the while we are still in the first event loop iteration. Only now - afterwards - we attach the event handler callback for the *error* event; but the constructor already tried to call it!

Why doesn't this problem arise with the *grew* and *shrank* events? Here is the visualization:



This shows how our main code in the first event loop starts the asynchronous *setInterval* operation which in turn starts an asynchronous *stat* operation. Only several event loop iterations later, when the callback for the *grew* event has long been attached, is this callback triggered. No problem arises.

How can we solve the problem for our *error* event? The solution is to put the triggering of the not-yet-attached callback into the future, into the next event loop iteration; this way, the callback will be triggered in the second event loop iteration *after* it has been attached in the first event loop iteration.

Moving code execution into the next event loop iteration is simple thanks to the *process.nextTick* method. Here is how the relevant part in file *FilesizeWatcher.js* needs to be rewritten:

```

if (/^\//.test(path) === false) {
  process.nextTick(function() {
    self.callbacks['error']('Path does not start with a slash');
  });
  return;
}

```

With this, we put the callback triggering for the *error* event listener into a callback function that will be executed by *process.nextTick* on the next iteration of the event loop:

```

+----> -----+      +----> -----+ +----> -----+ +----> -----+
|              |      |              | |              | |              |
| Create       |      |              | |              | |              |
| watcher     |      |              | |              | |              |
| object      |      |              | |              | |              |
|             |      |              | |              | |              |
| Check path  |      |              | |              | |              |
|             |      |              | |              | |              |
| Put         |      |              | |              | |              |
| callback    |      |              | |              | |              |
| into next   |      |              | |              | |              |
| iteration   |      |              | |              | |              |
|             |      |              | |              | |              |
| Attach     |      |              | |              | |              |
| callback    |      |              | |              | |              |
| to error    |      |              | |              | |              |
| event       |      |              | |              | |              |
|             |      |              | |              | |              |
+-----+-----+      +-----^-----+ +-----+-----+ +-----+-----+
|              |      |              | |              | |              |
| process.nextTick() | callback() |              | |              |
|              |      |              | |              | |              |
+-----+-----+

```

Thus we ensure that it is called *after* it has been defined.

As mentioned, Node.js supports us in creating event emitters, sparing us from writing the callback handling code for our own emitters.

We can use this support by inheriting from an existing class *EventEmitter* in module *events*.

Because the expected behaviour of our *MyEmitter* doesn't change from an outside perspective, we will leave the spec unchanged, and only change our implementation from using a self-made solution to utilizing the existing Node.js class:

```

1  'use strict';
2
3  var fs = require('fs');
4  var util = require('util');
5  var EventEmitter = require('events').EventEmitter;
6
7  var FilesizeWatcher = function(path) {
8    var self = this;
9

```

```
10  if (/^\//.test(path) === false) {
11    process.nextTick(function() {
12      self.emit('error', 'Path does not start with a slash');
13    });
14    return;
15  }
16
17  fs.stat(path, function(err, stats) {
18    self.lastfilesize = stats.size;
19  });
20
21  self.interval = setInterval(
22    function() {
23      fs.stat(path, function(err, stats) {
24        if (stats.size > self.lastfilesize) {
25          self.emit('grew', stats.size - self.lastfilesize);
26          self.lastfilesize = stats.size;
27        }
28        if (stats.size < self.lastfilesize) {
29          self.emit('shrank', self.lastfilesize - stats.size);
30          self.lastfilesize = stats.size;
31        }
32      }, 1000);
33    });
34 };
35
36 util.inherits(FilesizeWatcher, EventEmitter);
37
38 FilesizeWatcher.prototype.stop = function() {
39   clearInterval(this.interval);
40 };
41
42 module.exports = FilesizeWatcher;
```

The interesting part here is line 36. We use a helper function that ships with Node.js, *util.inherits*. We use it to make our *FilesizeWatcher* class a descendant of *events.EventEmitter*, which also ships with Node.js. This class implements all the logic that is needed to act as an event emitter, and our *FilesizeWatcher* class inherits this logic. This way, our own code now concentrates on implementing its specific logic, without having to care about the details of registering callbacks and checking if these can be called etc.

According to our tests, *FilesizeWatcher* still behaves as it did before:

...

Finished in 0.052 seconds

3 tests, 3 assertions, 0 failures, 0 skipped



Instead of using *util.inherits*, we can also create the inheritance by placing a new *events.EventEmitter* object on the prototype of *FilesizeWatcher*:

```
FilesizeWatcher.prototype = new EventEmitter();
```

Although *util.inherits* adds some extra properties to our constructor, the result is the same, behaviour-wise. Give it a try if you like and see if the tests still pass.

Summary

The asynchronous development style using callbacks that is a recurring pattern in JavaScript and Node.js applications can at times result in a code structure that is hard to understand. Applying the event emitter pattern comes in handy in situations where more than one callback event must be handled, resulting in a cleaner code structure that is easier to follow and reason about. The pattern is so useful that we should use it for our own objects and functions where it makes sense, something that's made easy by utilizing the existing *util.inherits* class.

Optimizing code performance and control flow management using the *async* library

Writing or using functions or methods in your code that are executed one after the other gets you a long way in your applications.

Sometimes, those functions are simple synchronous steps:

```
console.log('Starting calculation...');
var result = 5 + 4;
console.log('The result is', result);
```

Often, callbacks are used when operations are executed in the background and jump back into our code's control flow asynchronously at a later point in time:

```
console.log('Starting calculation...');
startExpensiveCalculation(5, 4, function(err, result) {
  if (!err) {
    console.log('The result is', result);
  }
});
```

If those asynchronous background operations bring forth a more complex callback behaviour, they might be implemented as an event emitter:

```
console.log('Starting calculation...');

calculation = Calculator.start(5, 4);

calculation.on('error', function(err) {
  console.log('An error occurred:', err);
});

calculation.on('interim result', function(result) {
  console.log('Received interim result:', result);
});
```

```
});  
  
calculation.on('done', function(result) {  
  console.log('Received final result:', result);  
});
```

Handling expensive operations asynchronously in the background, especially if they are IO-bound, is an important key to making Node.js applications perform efficiently - reading a large file or writing a lot of records to a database will always be a costly procedure, but handling it asynchronously at least ensures that the other parts of our application won't be blocked while that procedure is going on.

Nevertheless, there is often potential for optimization within our own code and its control flow.

Executing expensive asynchronous background tasks in parallel

Let's consider an example where our application queries several different remote web services, presenting the retrieved data on the console.

We are not going to query real remote web services, instead we will write a very simple Node.js HTTP server that will serve as a dummy web service. Our web server doesn't really do anything significant, and therefore we will make it react to requests a bit slower than necessary, in order to simulate a real web service that has a certain workload - as you will see, this makes it easier for us to show the performance gain in our own code optimizations.

Please create a new project folder and create a file *server.js* with the following content:

```
1  'use strict';  
2  
3  var http = require('http');  
4  var url = require('url');  
5  var querystring = require('querystring');  
6  
7  http.createServer(function(request, response) {  
8  
9    var pathname = url.parse(request.url).pathname;  
10   var query = url.parse(request.url).query;  
11   var id = querystring.parse(query)['id'];  
12  
13   var result = {  
14     'pathname': pathname,
```

```
15     'id': id,
16     'value': Math.floor(Math.random() * 100)
17   };
18
19   setTimeout(function() {
20     response.writeHead(200, {"Content-Type": "application/json"});
21     response.end(JSON.stringify(result));
22   }, 2000 + Math.floor(Math.random() * 1000));
23
24 }).listen(
25   8080,
26   function() {
27     console.log('Echo Server listening on port 8080');
28   }
29 );
```

This gives us a very simple “echo” server - if we request the URL `http://localhost:8080/getUser?id=4`, we receive `{"pathname": "/getUser", "id": "4", "value": 67}` as the response. This is good enough to give us the simulation of a remote webservice API to play around with.

But alas, it’s a slow webservice! Someone didn’t do his optimization homework, and we now have to deal with an API where every response takes between 2 to 3 seconds (this is simulated with the *setTimeout* construct on lines 19-22).

This allows to show how different request patterns will result in different runtime characteristics.

We will now write a Node.js client for this webserver. This client will make two consecutive requests to the server, and print the output for both requests on the command line:

```
1  'use strict';
2
3  var request = require('request');
4
5  request.get(
6    'http://localhost:8080/getUserName?id=1234',
7    function(err, res, body) {
8      var result = JSON.parse(body);
9      var name = result.value;
10
11     request.get(
12       'http://localhost:8080/getUserStatus?id=1234',
13       function(err, res, body) {
14         var result = JSON.parse(body);
15         var status = result.value;
```

```
16
17     console.log('The status of user', name, 'is', status);
18   });
19
20 });
```



Save this in a file called *client.js*. Also, you need to install the *request* module via `npm install request`.

This is probably the most straight-forward solution. We start the first request, wait for it to finish, print the result, then start the second request, wait for it to finish, and print that result, too.

How long does that take? Let's see:

```
$ time node client.js
The status of user 62 is 68
```

```
real 0m5.810s
user 0m0.172s
sys  0m0.033s
```



Don't forget to start the server via `node server.js` beforehand!

Not surprisingly, it takes around 5-6 seconds, because we only start the second request after the first request has been completed, and each request takes around 2-3 seconds.

We can't do anything about the terribly slow remote webservice, but our own code isn't exactly optimal either. Our two requests don't inherently depend on each other, and yet, we are executing them serially.

Of course starting these requests in parallel is simple, because both are asynchronous operations:

```
1  'use strict';
2
3  var request = require('request');
4
5  var name, status;
6
7  request.get(
8    'http://localhost:8080/getUserName?id=1234',
9    function(err, res, body) {
10     var result = JSON.parse(body);
11     name = result.value;
12   });
13
14  request.get(
15    'http://localhost:8080/getUserStatus?id=1234',
16    function(err, res, body) {
17     var result = JSON.parse(body);
18     status = result.value;
19   });
20
21  console.log('The status of user', name, 'is', status);
```

No, wait, sorry! That's not going to work - *console.log* will execute within the first event loop iteration while the request callbacks are triggered in later iterations. Mh, how about...

```
1  'use strict';
2
3  var request = require('request');
4
5  var name, status;
6
7  request.get(
8    'http://localhost:8080/getUserName?id=1234',
9    function(err, res, body) {
10     var result = JSON.parse(body);
11     name = result.value;
12   });
13
14  request.get(
15    'http://localhost:8080/getUserStatus?id=1234',
16    function(err, res, body) {
17     var result = JSON.parse(body);
```

```
18     status = result.value;
19
20     console.log('The status of user', name, 'is', status);
21 });
```

No, that's not good either: We *start* both request in parallel, but we have no guarantee that they will finish at the same time. We risk printing

```
1 The status of user undefined is 75
```

if the second request finishes earlier than the first. Well, looks like we need some additional code to synchronize our finished calls. How about this:

```
1 'use strict';
2
3 var request = require('request');
4
5 var name, status;
6 var firstHasFinished, secondHasFinished = false;
7
8 request.get(
9   'http://localhost:8080/getUserName?id=1234',
10  function(err, res, body) {
11    var result = JSON.parse(body);
12    name = result.value;
13    markFinished('first');
14  });
15
16 request.get(
17   'http://localhost:8080/getUserStatus?id=1234',
18  function(err, res, body) {
19    var result = JSON.parse(body);
20    status = result.value;
21    markFinished('second');
22  });
23
24 function markFinished(step) {
25   if (step == 'first') {
26     firstHasFinished = true;
27   }
28
29   if (step == 'second') {
```

```
30     secondHasFinished = true;
31   }
32
33   if (firstHasFinished && secondHasFinished) {
34     console.log('The status of user', name, 'is', status);
35   }
36 }
```

Seriously now - that's not even funny! What if we need to synchronize dozens or hundreds of operations? We could use an array where we store the state of each operation... no, this whole thing doesn't feel right.

async to the rescue, I say!

async is a clever little module that makes managing complex control flows in our code a breeze.

After installing the module via `npm install async`, we can write our client like this:

```
1  'use strict';
2
3  var request = require('request');
4  var async = require('async');
5
6  var name, status;
7
8  var getUsername = function(callback) {
9    request.get(
10     'http://localhost:8080/getUserName?id=1234',
11     function(err, res, body) {
12       var result = JSON.parse(body);
13       callback(err, result.value);
14     });
15 };
16
17 var getUserStatus = function(callback) {
18   request.get(
19     'http://localhost:8080/getUserStatus?id=1234',
20     function (err, res, body) {
21       var result = JSON.parse(body);
22       callback(err, result.value);
23     });
24 };
25
26 async.parallel([getUsername, getUserStatus], function(err, results) {
```

```
27   console.log('The status of user', results[0], 'is', results[1]);  
28   });
```

Let's analyze what we are doing here.

On line 4, we load the *async* library. We then wrap our requests into named functions. These functions will be called with a *callback* parameter. Inside our functions, we trigger this callback when our operation has finished - in this case, when the requests have been answered.

We call the callbacks with two parameters: an *error* object (which is *null* if no errors occurred), and the result value.

The orchestration happens on lines 26-28. We use the *parallel* method of the *async* object and pass an array of all the functions we want to run in parallel. Additionally, we pass a callback function which expect two parameters, *err* and *results*.

async.parallel will trigger this callback as soon as the slowest of the parallel operations has finished (and called its callback), or as soon as one of the operations triggers its callback with an error.

Let's see what this does to the total runtime of our script:

```
$ time node client.js  
The status of user 95 is 54  
  
real    0m3.176s  
user    0m0.240s  
sys     0m0.044s
```

As one would expect, the total runtime of our own code matches the runtime of one request because both requests are started in parallel and will finish roughly at the same time.

Optimizing code structure with *async*

async offers several other mechanisms for managing the control flow of our code. These are interesting even if our concern isn't performance optimization. Let's investigate them.

For these cases, let's remove the artificial slowness from the API server by removing the *setTimeout* operation on line 19 and 22, making the server respond immediately:


```
1  'use strict';
2
3  var http = require('http');
4  var url = require('url');
5  var querystring = require('querystring');
6
7  http.createServer(function(request, response) {
8
9      var pathname = url.parse(request.url).pathname;
10     var query = url.parse(request.url).query;
11     var id = querystring.parse(query)['id'];
12
13     var result = {
14         'pathname': pathname,
15         'id': id,
16         'value': Math.floor(Math.random() * 100)
17     };
18
19     response.writeHead(200, {"Content-Type": "application/json"});
20     response.end(JSON.stringify(result));
21
22 }).listen(
23     8080,
24     function() {
25         console.log('Echo Server listening on port 8080');
26     }
27 );
```

Sometimes we want to run operations in series. This is of course possible by putting method calls into the callback functions of previous method calls, but the code quickly becomes ugly if you do this with a lot of methods:

```
'use strict';

var request = require('request');

var url = 'http://localhost:8080/';

request.get(url + 'getUserName?id=1234', function(err, res, body) {
    console.log('Name:', JSON.parse(body).value);

    request.get(url + 'getUserStatus?id=1234', function(err, res, body) {
```

```
    console.log('Status:', JSON.parse(body).value);

    request.get(url + 'getUserCountry?id=1234', function(err, res, body) {
        console.log('Country:', JSON.parse(body).value);

        request.get(url + 'getUserAge?id=1234', function(err, res, body) {
            console.log('Age:', JSON.parse(body).value);
        });
    });

});

});
```

This is already starting to look messy, and we haven't even added any notable “business logic” to our code.

Note how our code is intended another level with every method call, creating the so-called “boomerang pattern” that is typical for multi-level nested callback control flows.

We can use *async.series* to achieve the same control flow with much cleaner code:

```
1  'use strict';
2
3  var request = require('request');
4  var async = require('async');
5
6  var url = 'http://localhost:8080/';
7
8  async.series([
9
10     function(callback) {
11         request.get(url + 'getUserName?id=1234', function(err, res, body) {
12             console.log('Name:', JSON.parse(body).value);
13             callback(null);
14         });
15     },
16
17     function(callback) {
18         request.get(url + 'getUserStatus?id=1234', function(err, res, body) {
19             console.log('Status:', JSON.parse(body).value);
20             callback(null);
21         });
22     }
23 ], function(err, results) {
24     // ...
25 });
```

```
21     });
22   },
23
24   function(callback) {
25     request.get(url + 'getUserCountry?id=1234', function(err, res, body) {
26       console.log('Country:', JSON.parse(body).value);
27       callback(null);
28     });
29   },
30
31   function(callback) {
32     request.get(url + 'getUserAge?id=1234', function(err, res, body) {
33       console.log('Age:', JSON.parse(body).value);
34       callback(null);
35     });
36   }
37
38 ]);
```

Just as with *async.parallel*, we can use *async.series* to collect the results of each step and do something with them once all steps have finished. This is again achieved by passing the result of each step to the callback each step triggers, and by providing a callback function to the *async.series* call which will receive an array of all results:

```
1  'use strict';
2
3  var request = require('request');
4  var async = require('async');
5
6  var url = 'http://localhost:8080/';
7
8  async.series([
9
10   function(callback) {
11     request.get(url + 'getName?id=1234', function(err, res, body) {
12       callback(null, 'Name: ' + JSON.parse(body).value);
13     });
14   },
15
16   function(callback) {
17     request.get(url + 'getStatus?id=1234', function(err, res, body) {
18       callback(null, 'Status: ' + JSON.parse(body).value);
```

```
19     });
20   },
21
22   function(callback) {
23     request.get(url + 'getUserCountry?id=1234', function(err, res, body) {
24       callback(null, 'Country: ' + JSON.parse(body).value);
25     });
26   },
27
28   function(callback) {
29     request.get(url + 'getUserAge?id=1234', function(err, res, body) {
30       callback(null, 'Age: ' + JSON.parse(body).value);
31     });
32   }
33 ],
34
35
36 function(err, results) {
37   for (var i=0; i < results.length; i++) {
38     console.log(results[i]);
39   }
40 }
41
42 );
```

In case that one of the series steps passes a non-null value to its callback as the first parameter, the series is immediately stopped, and the final callback is triggered with the results that have been collected to far, and the *err* parameter set to the error value passed by the failing step.

async.waterfall is similar to *async.series*, as it executes all steps in series, but it also enables us to access the results of a previous step in the step that follows:

```
1 'use strict';
2
3 var request = require('request');
4 var async = require('async');
5
6 var url = 'http://localhost:8080/';
7
8 async.waterfall([
9
10   function(callback) {
11     request.get(url + 'getSessionId', function(err, res, body) {
```

```
12     callback(null, JSON.parse(body).value);
13   });
14 },
15
16   function(sId, callback) {
17     request.get(url + 'getUserId?sessionId=' + sId, function(err, res, body) {
18       callback(null, sId, JSON.parse(body).value);
19     });
20   },
21
22   function(sId, uId, callback) {
23     request.get(url + 'getUserName?userId=' + uId, function(err, res, body) {
24       callback(null, JSON.parse(body).value, sId);
25     });
26   }
27 ],
28 ],
29
30 function(err, name, sId) {
31   console.log('Name:', name);
32   console.log('SessionID:', sId);
33 }
34
35 );
```

Note how for each step function, *callback* is passed as the last argument. It follows a list of arguments for each parameter that is passed by the previous function, minus the error argument which each step function always passes as the first parameter to the callback function.

Also note the difference in the final callback: instead of *results*, it too expects a list of result values, passed by the last waterfall step.

async provides several other interesting methods which help us to bring order in our control flow and allows us to orchestrate tasks in an efficient manner. Check out [the *async* README¹](https://github.com/caolan/async/blob/master/README.md) for more details.

¹<https://github.com/caolan/async/blob/master/README.md>

Node.js and MySQL

Using the node-mysql library

Node.js is able to connect to MySQL database servers, but support for this is not built in. We need to make use of a Node.js package that handles the low level details of establishing a connection and talking to the database. Additionally, such a package should enable us to issue statements to the database written in SQL.

The most widely accepted package for this is *node-mysql* by Node.js core contributor Felix Geisendörfer.

In order to use it, let's start by declaring *mysql* as a dependency for our project in our *package.json* file:

```
{
  "dependencies": {
    "mysql": "~2.0.0"
  }
}
```

As always, *npm install* pulls in everything we need to start working with a MySQL database.

For what follows, I assume that you have a working MySQL database up and running on your local machine. You should be able to connect to it locally. You can verify that this is the case by running

```
mysql -h127.0.0.1 -uroot -p
```

on your command line. Provide the database root password if needed - you will also need this password in your Node.js code in order to establish a connection.

A first database application

Once you are up and running, let's start with a first, simple database application. Create a file *index.js* in the directory where you installed the *node-mysql* module, and fill it with the following code:

```
1  'use strict';
2
3  var mysql = require('mysql');
4
5  var connection = mysql.createConnection({
6    host: 'localhost',
7    user: 'root',
8    password: 'root'
9  });
10
11 connection.query(
12   'SELECT "foo" AS first_field, "bar" AS second_field',
13   function (err, results, fields) {
14     console.log(results);
15     connection.end();
16   }
17 );
```

Let's see what exactly happens here. First, we include Felix' external *mysql* library (line 3). We then create a new object dubbed *connection* on line 5.

This object is a representation of our connection with the database. No connection without connecting, which is why we call the *createConnection* function with the parameters necessary to connect and authenticate against the database server: the host name, the username, and the password.

Note that there is a parameter *port* which defaults to 3306 if omitted. You only need to set this parameter if your database server does not run on this default port.

For a list of all possible parameters (a total of 17 are available), see the *node-mysql* documentation at <https://github.com/felixge/node-mysql#connection-options>.

Through the newly created connection object, we can now talk to the database server, which is what we do on line 11. To do so, the *connection* object provides a method named *query*.

We provide two parameters to the method call. The first is an SQL statement, the second is an anonymous function. This function gets called when Node.js receives the answer to our query from the server - a *callback*.

This is once again the classical asynchronous pattern of JavaScript applications that makes sense whenever our code triggers an operation whose execution could potentially block the program flow if executed synchronously: Imagine the database needing 5 seconds before it has compiled the answer to our query - that's not unlikely for more complex queries. If we wouldn't wait for the database's answer asynchronously, then our whole application would have to wait for it. No other code could be executed for these 5 seconds, rendering our whole application unusable during this time.

Line 12 shows which parameters the *query* method passes to our callback function: *err*, *results*, and *fields*.

For now, let's concentrate on *results*, because that's where the fun is.

What kind of thing is *results*? It's simply an array of objects, with each object representing one row of the result set and each attribute of such an object representing one field of the according row.

Let's run our application and examine its output to see how this looks. Please run `node index.js`. The result should look like this:

```
[ { first_field: 'foo', second_field: 'bar' } ]
```

Let's add a legend:

```
|----- An array -----|
|                               |
| |----- with an object -----| | |
| |                               | |
| | |----- that has an attribute -----| |
| | |                               | |
| | | with a value -----| |
| | |                               | |
| | ----- ---| |
| | { first_field: 'foo', second_field: 'bar' } | |
| | ----- ---| |
| | |                               | |
| | | with a value -----| |
| | |                               | |
| | |----- that has a field -----| |
| | |                               | |
| |----- with a row -----| |
| |                               | |
|----- A result set -----|
```

(Note that we didn't query for any real data from any real database table, but relational databases are really stubborn and treat everything as a set of rows of fields, even if we only query for some text strings mapped to made up fields like we just did.)

According to this structure we can of course access the value of any field in any row directly, by accessing the right attribute of the right object:


```
console.log(results[0].first_field); // Outputs 'foo';
```

Let's examine the other parameters that are passed to our anonymous callback function.

err is either an error object, if a problem occurred, or *null* if everything went well. Let's modify our code a bit to see how this works:

```
1  'use strict';
2
3  var mysql = require('mysql');
4
5  var connection = mysql.createConnection({
6    host: 'localhost',
7    user: 'root',
8    password: 'root'
9  });
10
11 connection.query(
12   'SLECT "foo" AS first_field, "bar" AS second_field',
13   function (err, results, fields) {
14     console.log(err);
15     console.log(results);
16     connection.end();
17   }
18 );
```

Note how I changed the SQL statement in line 12 and added logging of the *err* parameter on line 14. The altered SQL statement is invalid and results in the database answering the query with an error. *node-mysql* translates this error into an error object which is then passed as the first parameter to our callback function.

If you run the new code, a string representation of the error object is written to the console:

```
{ [Error: ER_PARSE_ERROR: You have an error in your SQL syntax; check
  the manual that corresponds to your MySQL server version for the
  right syntax to use near 'SLECT "foo" AS first_field, "bar" AS
  second_field' at line 1]
  code: 'ER_PARSE_ERROR',
  errno: 1064,
  sqlState: '42000',
  index: 0 }
```

The fact that the *err* parameter is *null* if no error occurred makes error handling in our code very simple:

```
1  'use strict';
2
3  var mysql = require('mysql');
4
5  var connection = mysql.createConnection({
6    host: 'localhost',
7    user: 'root',
8    password: 'root'
9  });
10
11 connection.query(
12   'SELECT "foo" AS first_field, "bar" AS second_field',
13   function (err, results, fields) {
14     if (err) {
15       console.log('A database error occurred!');
16     } else {
17       console.log(results);
18     }
19     connection.end();
20   }
21 );
```

That's a very typical pattern in a lot of Node.js applications: Call a library function while providing a callback, and check if the first parameter passed to this callback is *null* in order to decide if an error occurred, continuing otherwise.

fields returns an array with one object for every field that was requested in the query - each object contains meta information about the according field. It isn't of any significant use for us here. Feel free to *console.log()* the *fields* parameter if you like.

Using node-mysql's Streaming API

Next, we will query some real data from existing tables. For this, we first need to create a database, add a table to it, and insert some rows. Let's do this through a Node.js application. Please create a file *create.js* with the following content:

```
1  'use strict';
2
3  var mysql = require('mysql');
4
5  var connection = mysql.createConnection({
6    host: 'localhost',
7    user: 'root',
8    password: 'root'
9  });
10
11 connection.query('CREATE DATABASE node', function(err) {
12   if (err) {
13     console.log('Could not create database "node".');
14   }
15 });
16
17 connection.query('USE node', function(err) {
18   if (err) {
19     console.log('Could not switch to database "node".');
20   }
21 });
22
23 connection.query('CREATE TABLE test ' +
24   '(id INT(11) AUTO_INCREMENT, ' +
25   ' content VARCHAR(255), ' +
26   ' PRIMARY KEY(id))',
27   function(err) {
28     if (err) {
29       console.log('Could not create table "test".');
30     }
31   }
32 );
33
34 connection.query('INSERT INTO test (content) VALUES ("Hello")');
35 connection.query('INSERT INTO test (content) VALUES ("World")');
36
37 connection.end();
```

This program is meant to be executed just once in order to set up a new database with some example data. Please run it now.

Let's rewrite our *index.js* file and try to query for this newly added data:

```
1  'use strict';
2
3  var mysql = require('mysql');
4
5  var connection = mysql.createConnection({
6    host: 'localhost',
7    user: 'root',
8    password: 'root',
9    database: 'node'
10 });
11
12 connection.query(
13   'SELECT id, content FROM test',
14   function (err, results, fields) {
15     if (err) {
16       console.log('A database error occurred!');
17     } else {
18       console.log(results);
19     }
20     connection.end();
21   }
22 );
```

Please note how on line 9 I have added an additional connection parameter named *database*. This way, a default database is already selected for subsequent queries on this connection, and doesn't have to be selected through a *USE database* query.

If we execute the script, the result should look like this:

```
[ { id: 1, content: 'Hello' }, { id: 2, content: 'World' } ]
```

As one would expect, we now receive an array with two objects, because the resultset consists of two rows.

This gives us the opportunity to discuss an even better way to handle resultsets from queries.

I have explained how the classical callback pattern should be used on the *query* call because otherwise expensive database operations could block the whole Node.js application.

However, there is another source of potentially blocking behaviour: large resultsets.

Imagine that our test database consists of 1,000,000 rows instead of just two, and we would like to output the content of the whole resultset to the console like this:

```
connection.query(
  'SELECT id, content FROM test',
  function (err, results, fields) {
    for (var i = 0; i < results.length; i++) {
      console.log('Content of id ' + results[i].id +
                  ' is ' + results[i].content);
    }
  }
);
```

Using the callback pattern ensures that our Node.js application won't block while the database retrieves these 1,000,000 rows. However, once we receive the resultset with these rows in our application - that is, within the callback - and we start operating on them, executing these operations also will take a long time.

During this time, our application will block, too. Remember that in Node.js, only one piece of our JavaScript code is executed at any given time. If we iterate over these 1,000,000 array entries and do something with each entry, and this takes, say, 10 seconds, then our application cannot do something else during these 10 seconds.

The solution: instead of using the all-or-nothing callback pattern, we can utilize *node-mysql's* streaming API.

Instead of providing a callback function to the *query* call, we only provide the SQL statement, but we use the return value of the *query* function: a *Query* object.

This object emits several events that happen during the lifetime of the query: *error*, *field*, *row* and *end*.

Let's concentrate on *row* and see how a more efficient version of our 1,000,000 row output application would look like:

```
var query = connection.query('SELECT id, content FROM test');

query.on('row', function(row) {
  console.log('Content of id ' + row.id + ' is ' + row.content);
});
```

Now we have a piece of our code running 1,000,000 times, once for each row. While this doesn't make our application faster in the narrower sense (it will still take the assumed 10 seconds until all rows have been handled), it does make it way more efficient. Our application now only blocks during the handling of one row, immediately returning control to the event loop of our application, allowing to execute other code in between.

Of course, this isn't critical for an application that is used by only one person on one shell console. But it is *very* critical if such a database operation takes place within a multi-user client-server application:

Imagine you use a website to retrieve a small amount of data from a database - you wouldn't want to wait 10 seconds for your page to load only because another user requested a page that retrieved 1,000,000 rows for him from the database.

Let's rewrite our *index.js* script once again to demonstrate usage of the complete streaming API of *node-mysql*:

```
1  'use strict';
2
3  var mysql = require('mysql');
4
5  var connection = mysql.createConnection({
6    host: 'localhost',
7    user: 'root',
8    password: 'root',
9    database: 'node'
10 });
11
12 var query = connection.query('SELECT id, content FROM test');
13
14 query.on('error', function(err) {
15   console.log('A database error occurred:');
16   console.log(err);
17 });
18
19 query.on('fields', function(fields) {
20   console.log('Received fields information.');
```

```
21 });
```

```
22
```

```
23 query.on('result', function(result) {
```

```
24   console.log('Received result:');
```

```
25   console.log(result);
```

```
26 });
```

```
27
```

```
28 query.on('end', function() {
```

```
29   console.log('Query execution has finished.');
```

```
30   connection.end();
```

```
31 });
```

It's pretty much self explanatory. Note how the order of the *on* blocks doesn't say anything about when (and if) they are executed - it's up to *node-mysql* to decide when each block (or, to be more precise: the anonymous callback function associated with each block) will be called.

When executing this script, this is what the output should look like:

```
Received fields information.  
Received result:  
{ id: 1, content: 'Hello' }  
Received result:  
{ id: 2, content: 'World' }  
Query execution has finished.
```

Making SQL queries secure against attacks

Let's now go full circle and create a simple web application that allows to insert data into our table and also reads and displays the data that was entered.

We need to start a web server with two routes (one for displaying data, one for taking user input), and we need to pass user input to the database and database results to the webpage. Here is the application in one go:

```
1  'use strict';  
2  
3  var mysql      = require('mysql'),  
4      http       = require('http'),  
5      url        = require('url'),  
6      querystring = require('querystring');  
7  
8  
9  // Start a web server on port 8888. Requests go to function handleRequest  
10  
11 http.createServer(handleRequest).listen(8888);  
12  
13  
14 // Function that handles http requests  
15  
16 function handleRequest(request, response) {  
17  
18   // Page HTML as one big string, with placeholder "DBCCONTENT" for data from  
19   // the database  
20   var pageContent = '<html>' +  
21                   '<head>' +  
22                   '<meta http-equiv="Content-Type" ' +  
23                   'content="text/html; charset=UTF-8" />' +  
24                   '</head>' +  
25                   '<body>' +  
26                   '<form action="/add" method="post">' +
```

```
27         '<input type="text" name="content">' +
28         '<input type="submit" value="Add content" />' +
29         '</form>' +
30         '<div>' +
31         '<strong>Content in database:</strong>' +
32         '<pre>' +
33         'DBCCONTENT' +
34         '</pre>' +
35         '</div>' +
36         '<form action="/" method="get">' +
37         '<input type="text" name="q">' +
38         '<input type="submit" value="Filter content" />' +
39         '</form>' +
40         '</body>' +
41         '</html>';
42
43     // Parsing the requested URL path in order to distinguish between
44     // the / page and the /add route
45     var pathname = url.parse(request.url).pathname;
46
47     // User wants to add content to the database (POST request to /add)
48     if (pathname == '/add') {
49         var requestBody = '';
50         var postParameters;
51         request.on('data', function (data) {
52             requestBody += data;
53         });
54         request.on('end', function() {
55             postParameters = querystring.parse(requestBody);
56             // The content to be added is in POST parameter "content"
57             addContentToDatabase(postParameters.content, function() {
58                 // Redirect back to homepage when the database has finished adding
59                 // the new content to the database
60                 response.writeHead(302, {'Location': '/'});
61                 response.end();
62             });
63         });
64
65     // User wants to read data from the database (GET request to /)
66     } else {
67         // The text to use for filtering is in GET parameter "q"
68         var filter = querystring.parse(url.parse(request.url).query).q;
```



```
69     getContentsFromDatabase(filter, function(contents) {
70         response.writeHead(200, {'Content-Type': 'text/html'});
71         // Poor man's templating system: Replace "DBCCONTENT" in page HTML with
72         // the actual content we received from the database
73         response.write(pageContent.replace('DBCCONTENT', contents));
74         response.end();
75     });
76 }
77 }
78
79
80 // Function that is called by the code that handles the / route and
81 // retrieves contents from the database, applying a LIKE filter if one
82 // was supplied
83
84 function getContentsFromDatabase(filter, callback) {
85     var connection = mysql.createConnection({
86         host: 'localhost',
87         user: 'root',
88         password: 'root',
89         database: 'node'
90     });
91     var query;
92     var resultsAsString = '';
93
94     if (filter) {
95         query = connection.query('SELECT id, content FROM test ' +
96                                 'WHERE content LIKE "' + filter + '%"');
97     } else {
98         query = connection.query('SELECT id, content FROM test');
99     }
100
101     query.on('error', function(err) {
102         console.log('A database error occurred:');
103         console.log(err);
104     });
105
106     // With every result, build the string that is later replaced into
107     // the HTML of the homepage
108     query.on('result', function(result) {
109         resultsAsString += 'id: ' + result.id;
110         resultsAsString += ', content: ' + result.content;
```

```
111     resultsAsString += '\n';
112   });
113
114   // When we have worked through all results, we call the callback
115   // with our completed string
116   query.on('end', function(result) {
117     connection.end();
118     callback(resultsAsString);
119   });
120 }
121
122
123 // Function that is called by the code that handles the /add route
124 // and inserts the supplied string as a new content entry
125
126 function addContentToDatabase(content, callback) {
127   var connection = mysql.createConnection({
128     host: 'localhost',
129     user: 'root',
130     password: 'root',
131     database: 'node',
132   });
133
134   connection.query('INSERT INTO test (content) ' +
135     'VALUES (' + content + ')',
136     function(err) {
137       if (err) {
138         console.log('Could not insert content "' + content +
139           '" into database.');
```

Simply start the application by running `node index.js`, open up its homepage in your browser by visiting <http://localhost:8888/>, and play around with it.

So, that's quite a bit of code here. Let's walk through the important parts. If the user requests the homepage, here is what happens:

- The HTTP server that is started on line 11 receives the requests and calls the *handleRequest* function that is declared on line 16.

- On line 45 the variable *pathname* is declared and initialized with the requested path, which is / in this case, by parsing the request URL.
- The code continues on line 66, where the request for / is handled.
- The handling starts by parsing the URL for the value of the GET parameter *q* which is later used to filter the database content.
- On line 69 things start to get interesting. We call *getContentsFromDatabase* with two parameters: The *filter* string we just parsed from the URL, and an anonymous callback function.
- Execution continues on line 84, starting with the creation of a new database connection.
- Then, on lines 94 to 99, the query to retrieve the contents from the database is started.
- Each retrieved row triggers the execution of the *result* event handler on line 108, where we extend the string *resultsAsString* with the contents of each row.
- Execution jumps to line 116 once all rows have been retrieved. Here, in the *end* event handler, we close the connection to the database and then call the callback we passed into the *getContentsFromDatabase* function. This returns control back to the request handler, in line 70.
- There, we start responding to the request, first by sending a status 200 header, followed by the page content - with the database content replaced into the HTML at the proper place on line 74.

With this, the request has been handled and responded to.

When the user uses the *Filter content* form, the same procedure is executed, with the only difference that now the GET parameter *q* has a value which is then used to filter the database results using the SQL *LIKE* clause.

When using the *Add content* button, a different execution path is taken:

- The HTTP server that is started on line 11 receives the requests and calls the *handleRequest* function that is declared on line 16.
- On line 45 the variable *pathname* is declared and initialized with the requested path, which is */add* in this case, by parsing the request URL.
- The code continues on line 48, where the request for */add* is handled.
- Then, on line 51, the body of the request (the POST data) is retrieved step-by-step, in the same way that large resultsets from a database can be retrieved when using the streaming API. With every step, a variable is extended which finally holds the complete request body.
- Once the request has been fully received, execution continues on line 54 when the *end* event is triggered for the request.
- We then parse the request body in order to translate the POST data into a JavaScript object which allows us to easily access the value of POST parameters, as is done on line 57 where we call the *addContentToDatabase* function with the string that is to be added to the database as its first parameter, followed by a callback function that we expect to be triggered once the data has been inserted.

- With this, execution continues on line 126 where we start by connecting to the database, followed by an *INSERT* query on line 134. The query function triggers our anonymous callback defined on line 135 when the query execution has finished, which in turn triggers the callback that was passed in from the request handler function.
- This leads to execution continuing on line 60, where we redirect the user's browser back to the homepage.

So, nice application, does what it's supposed to do. And has a huge problem.

Why? Because the way we treat the incoming data that is provided by the user puts our database at risk. Our application is vulnerable to so-called SQL injection attacks. Let's demonstrate how.

To do so, let's create another table within our database *node*, one that's supposed to carry sensitive information that's not meant for any user's eyes. We will just pretend that our nice little content-viewer app is part of a bigger application that somewhere also stores user passwords.

We will call this table *passwords*, and you need to create this within the *node* database on your local server, either using the MySQL command line, a MySQL GUI client you happen to use, or, if you like, you can simply rewrite the *create.js* script we used earlier, like this:

```
1  'use strict';
2
3  var mysql = require('mysql');
4
5  var connection = mysql.createConnection({
6    host: 'localhost',
7    user: 'root',
8    password: 'root'
9  });
10
11 connection.query('USE node', function(err) {
12   if (err) {
13     console.log('Could not switch to database "node".');
14   }
15 });
16
17 connection.query('CREATE TABLE passwords ' +
18                 '(id INT(11) AUTO_INCREMENT, ' +
19                 ' password VARCHAR(255), ' +
20                 ' PRIMARY KEY(id))',
21   function(err) {
22     if (err) {
23       console.log('Could not create table "passwords".');
24     }
25   }
26 });
```

```
25     }
26   );
27
28   connection.query('INSERT INTO passwords (password) VALUES ("secret")');
29   connection.query('INSERT INTO passwords (password) VALUES ("dont_tell")');
30
31   connection.end();
```

Execute *create.js* through Node.js in order to create the new table.

Now, here is the scenario: We have a table *test* that contains public information, and a table *passwords* that contains sensitive information. We have written an application that operates only on the *test* table, thus allowing public access to the public information in that table, but not allowing access to any sensitive information like the passwords in the *passwords* table.

And yet, the passwords can be retrieved through the application. Here's how to do that:

- Start the application with *node index.js*
- Visit <http://localhost:8888> in your browser
- Enter the following text into the second text input field and hit *Filter content*

```
%" UNION SELECT id, password AS content FROM passwords WHERE password LIKE "
```

Note that you need to enter the text *exactly* as shown above, starting with a percentage sign and ending with a quotation mark.

You will receive a list of data from the test table (if any content has been added before), but at the end of this list, you will also see the two passwords we have added to the *passwords* table, like this:

```
id: 1, content: foo bar
id: 1, content: secret
id: 2, content: dont_tell
```

Whoops. How is that possible? It's because a set of SQL commands has been injected into our application's SQL, although only data - the string to filter for - was meant to be injected. Hence an SQL injection attack.

If you look at the critical piece of our code, what happens becomes clear. On lines 95 and 96, our SQL query is constructed:

```
query = connection.query('SELECT id, content FROM test ' +  
                          'WHERE content LIKE "' + filter + '%"');
```

If the content of the *filter* variable is “foo”, this is what the final SQL looks like:

```
SELECT id, content FROM test WHERE content LIKE "foo%"
```

But if the malicious filter text is entered into the web page, we end up with the following SQL query:

```
SELECT id, content FROM test WHERE content LIKE "%" UNION  
SELECT id, password AS content FROM passwords WHERE password LIKE "%"
```

(I’ve added a line break in order to stay within the boundaries of the text page.)

The malicious filter text simply closes the LIKE query of the original SQL and continues with an additional SELECT on the *passwords* table, combining the results of both SELECTs into a single (and identically structured) result set using the UNION clause. Another LIKE clause is added at the end of the malicious code, which makes sure that the whole SQL is syntactically correct.

This is a made up example for demo purposes, but that’s how actual attacks look like. A real attack might need a lot more effort - e.g., the attacker needs to find out the structure of the passwords table, but supporting information can be gathered by querying meta data from the database; if an attacker finds a part of the application that allows to extend harmless SQL queries into more potent ones, finding sensitive information is only a question of time and patience.

What can we, as application developers, do against such attacks? There is no 100% secure system, as the saying goes, however, there are a lot of ways to raise the bar, making attacks more difficult.

In our scenario, a very effective counter-measure is to use placeholders with value assignment when building SQL statements, instead of hand-crafting the final query by concatenating strings.

This way, we can guarantee that any part of your query that is supposed to be treated as a value is indeed interpreted by *node-mysql* as a value, escaping it in a way that doesn’t allow the value to end up at the database server as a command.

To achieve this, we need a way to create a query that can unambiguously differentiate between the command parts of a query and the value parts.

This is done by creating a query where the value parts are not directly included within the query string - instead, their position is marked with a special character, the quotation mark (?), as a placeholder.

This is how our query looks when using this syntax:

```
SELECT id, content FROM test WHERE content LIKE ?
```

Now *node-mysql* knows that whatever is going to be used at the position marked by the `?` placeholder is a value, treating it accordingly.

Of course, when issuing the query, we also need a way to tell what value to use. This is done by mapping the value onto the query-with-placeholder, like so:

```
filter = filter + '%';
query = connection.query('SELECT id, content FROM test WHERE content LIKE ?',
    filter);
```

In our specific case, we need to add the search wildcard, `%`, to the filter parameter. We then construct the query with the placeholder, passing the value for this placeholder as the second parameter. The rest of the code remains unchanged.

We should change the INSERT query that is used for adding content, too, by rewriting line 135 to

```
connection.query('INSERT INTO test (content) VALUES (?)', content,
    function(err) {
    ...
```

This changes the SQL query to placeholder syntax, and adds the *content* variable as the second parameter. The callback function is now the third parameter.

Mapping multiple values onto multiple placeholders is done by providing an array of values:

```
query('SELECT foo FROM bar WHERE id = ? AND name = ?', [42, 'John Doe']);
```

Here, “42” maps to the *id* placeholder, and “John Doe” to the *name* placeholder.

If you restart the application and try to attack it again, you will see that the attack no longer works. Mission accomplished.

Summary

In this chapter we learned about the *node-mysql* module, how it enables us to connect to relational databases, how using asynchronous callbacks and the Streaming API enable us to write efficient applications even when large resultsets need to be handled, and analyzed how database applications can be attacked and, more importantly, how we can protect our applications against these kinds of attacks.

Node.js and MongoDB

Some MongoDB basics

MongoDB is a document-oriented NoSQL database that stores objects in the following hierarchy:

Server

```
\_Database
  \_Collection
    \_Document
    \_Document
    \_Document
    \_ ...
  \_Collection
    \_Document
    \_Document
    \_Document
    \_ ...
```

If you are coming from a relational SQL database background (I do), then it might help to think of collections as tables, and of documents as rows in that table. It's not a perfect analogy, however; the key idea of a document-oriented database like MongoDB is that documents within the same collection do not need to have the same structure, while each row in a table of a relational database does have the same structure.

Here is how actual data in this hierarchy might actually look like:

localhost

```
\_accounting
  \_customers
    \_ { _id: 53c6c2, name: 'Jane Doe' }
    \_ { _id: 9dc42e, name: 'John Doe', age: 24 }
    \_ { _id: 63a76d, name: 'Jim Doe', options: { billtime: 3 } }
  \_invoices
    \_ { _id: 98ef5a, value: 192.87 }
    \_ { _id: f4eb21, value: 192.87, recurring: true }
```


This example represents the case where one has a MongoDB server running on localhost, with two databases *customers* and *invoices*. Each database contains some documents that are similar, but not identical, in structure.

As you can see, I have represented the document objects in JavaScript object notation - this is how we are going to encounter them when working with a MongoDB server from within Node.js.

Every document has an *_id* attribute, which is a random 24-digit hex value - I have shortened them for the illustration.

We will now examine how to create and work with such a structure.

Applying CRUD operations with the low-level mongodb driver

Again, I won't cover installation of the server software. Please refer to [the official installation instructions](http://docs.mongodb.org/manual/installation/)² in order to get a MongoDB server up and running on your machine.

Once this is done, create a new project folder called *mongodb-test*, and in there create an initial *package.json* as follows:

```
{
  "name": "mongodb-test",
  "version": "0.0.1",
  "description": "",
  "dependencies": {
    "mongodb": "^1.4.7"
  },
  "devDependencies": {}
}
```

As always, we pull in the dependencies with

```
npm install
```

Now we can create a first program in file *index.js*:

²<http://docs.mongodb.org/manual/installation/>

```
1  'use strict';
2
3  var MongoClient = require('mongodb').MongoClient;
4
5  MongoClient.connect(
6    'mongodb://127.0.0.1:27017/accounting',
7    function(err, connection) {
8      var collection = connection.collection('customers');
9
10     collection.insert({'name': 'Jane Doe'}, function(err, count) {
11
12       collection.find().toArray(function(err, documents) {
13         console.dir(documents);
14         connection.close();
15       });
16     });
17   });
18
19 });
```

Let's see what this code does. On line 3, we require the *MongoClient* class from the *mongodb* library. We use the *connect* method of this class on line 5 to establish a connection to our MongoDB server, which is listening on *localhost*, port *27017*. Upon connecting, we declare that we want to use the *accounting* database.

We then create an object *collection* on line 8 which represents the *customers* collection within the *accounting* database.

Using the *insert* method of the collection object, we add a new document to the *customers* collection on line 10.

When the insert operation has finished, we query the collection for all existing documents using the *find* method on line 12. Our callback receives an array with the documents, prints them to the screen, and then closes the database connection.

Running this code by executing

```
node index.js
```

will result in output similar to

```
[ { _id: 53c8ba1b517c86d3c9df71e6, name: 'Jane Doe' } ]
```

and running it again will result in output similar to

```
[ { _id: 53c8ba1b517c86d3c9df71e6, name: 'Jane Doe' },  
  { _id: 53c8ba21d13eb8dbc96e2c19, name: 'Jane Doe' } ]
```

As you can see, each run creates a new document in the database. The *name* attribute is set by our code, the *_id* is generated by MongoDB.

Let's rewrite our code a bit in order to update existing documents:

```
1  'use strict';  
2  
3  var MongoClient = require('mongodb').MongoClient;  
4  
5  MongoClient.connect(  
6    'mongodb://127.0.0.1:27017/accounting',  
7    function(err, connection) {  
8      var collection = connection.collection('customers');  
9  
10     collection.update({}, {'$set': {'age': 24}}, function(err, count) {  
11  
12       console.log('Updated', count, 'documents');  
13  
14       collection.find().toArray(function(err, documents) {  
15         console.dir(documents);  
16         connection.close();  
17       });  
18  
19     });  
20  
21   });
```

The change is on line 10. Instead of inserting a new document, we want to update our documents. The first parameter is the filter that defines which documents to match when applying the update. In this case, we pass an empty object, that is, we don't want to match any specific document and instead we want to apply the update to all documents in the collection.

The second parameter defines what to update within each matched document. In this case, we want to set the *age* attribute of the documents to 24. The documents we already stored do not have an *age* attribute yet, but that's not a problem: because MongoDB works schema-less, it will happily add attributes that did not yet exist on a given document.

This is how the content of our collection looks like after running the new code:

Updated 1 documents

```
[ { _id: 53c8ba1b517c86d3c9df71e6, name: 'Jane Doe' },  
  { _id: 53c8ba21d13eb8dbc96e2c19, name: 'Jane Doe', age: 24 } ]
```

Oh, well. Not quite what we expected - we didn't define a filter to the *update* statement because we wanted to match all existing documents, and yet only one of our two documents has been updated with the *age* attribute.

The reason is that this is the default behaviour for MongoDB updates: only the first match for a given filter - even if it matches every document - is updated.

By changing line 10 from

```
collection.update({}, {'$set': {'age': 24}}, function(err, count) {
```

to

```
collection.update(  
  {},  
  {'$set': {'age': 24}},  
  {'multi': true},  
  function(err, count) {
```

(that is, by adding a parameter `{'multi': true}`), we override this behaviour, which then results in MongoDB updating all matching documents.

Let's have a look at how actual filters that match only some of our documents look like. We are going to rewrite our script because we need some more documents to play with:

```
1  'use strict';  
2  
3  var MongoClient = require('mongodb').MongoClient;  
4  
5  MongoClient.connect(  
6    'mongodb://127.0.0.1:27017/accounting',  
7    function (err, connection) {  
8      var collection = connection.collection('customers');  
9  
10     var doFind = function (callback) {  
11       collection.find().toArray(function (err, documents) {  
12         console.dir(documents);  
13         callback();  
14       });  
15     };
```

```
15     };
16
17     var doInsert = function (i) {
18         if (i < 20) {
19             var value = Math.floor(Math.random() * 10);
20             collection.insert(
21                 {'n': '#' + i, 'v': value},
22                 function (err, count) {
23                     doInsert(i + 1);
24                 });
25         } else {
26             console.log();
27             console.log('Inserted', i, 'documents:');
28             doFind(function () {
29                 doUpdate();
30             });
31         }
32     };
33
34     var doUpdate = function () {
35         collection.update(
36             {'v': {'$gt': 5}},
37             {'$set': {'valuable': true}},
38             {'multi': true},
39             function (err, count) {
40                 console.log();
41                 console.log('Updated', count, 'documents:');
42                 doFind(function () {
43                     collection.remove({}, function () {
44                         connection.close();
45                     });
46                 });
47             });
48     };
49
50     doInsert(0);
51
52 });
```

As you can see, we created three functions, *doFind*, *doInsert* and *doUpdate*. The script starts by calling *doInsert*, which is a recursive function that inserts 20 documents into the collection, and does so in a serial fashion by recursively calling itself (on line 23) every time the previous insert operation

has finished. Each documents has an *n* attribute with the serial number of its insert operation, and a *v* value, which is a random number between 0 and 9.

As soon as the limit of 20 inserts has been reached, the function calls *doFind* (on line 28) which prints all existing documents to the screen. We pass *doFind* a callback function which calls *doUpdate*. This function adds the field *valuable* to some, but not all, documents.

Whether a document is updated or not depends on it's *v* value: using the *\$gt* operator (on line 36), we filter the documents that are affected by the update to those whose value is *greater than 5*.

After the update operation has finished, *doFind* is called once again, printing the results of the update to the screen. We then delete all documents using *remove* (on line 43), again using a filter that matches any document - however, *remove* doesn't need a *multi* parameter in order to work on all matched documents; it does so by default.

Here is a typical result for a run of this script:

Inserted 20 documents:

```
[ { _id: 53d60cc676371927029f95bd, n: '#0', v: 5 },
  { _id: 53d60cc676371927029f95be, n: '#1', v: 8 },
  { _id: 53d60cc676371927029f95bf, n: '#2', v: 0 },
  { _id: 53d60cc676371927029f95c0, n: '#3', v: 4 },
  { _id: 53d60cc676371927029f95c1, n: '#4', v: 6 },
  { _id: 53d60cc676371927029f95c2, n: '#5', v: 2 },
  { _id: 53d60cc676371927029f95c3, n: '#6', v: 4 },
  { _id: 53d60cc676371927029f95c4, n: '#7', v: 1 },
  { _id: 53d60cc676371927029f95c5, n: '#8', v: 0 },
  { _id: 53d60cc676371927029f95c6, n: '#9', v: 7 },
  { _id: 53d60cc676371927029f95c7, n: '#10', v: 5 },
  { _id: 53d60cc676371927029f95c8, n: '#11', v: 8 },
  { _id: 53d60cc676371927029f95c9, n: '#12', v: 2 },
  { _id: 53d60cc676371927029f95ca, n: '#13', v: 8 },
  { _id: 53d60cc676371927029f95cb, n: '#14', v: 1 },
  { _id: 53d60cc676371927029f95cc, n: '#15', v: 0 },
  { _id: 53d60cc676371927029f95cd, n: '#16', v: 1 },
  { _id: 53d60cc676371927029f95ce, n: '#17', v: 4 },
  { _id: 53d60cc676371927029f95cf, n: '#18', v: 4 },
  { _id: 53d60cc676371927029f95d0, n: '#19', v: 6 } ]
```

Updated 6 documents:

```
[ { _id: 53d60cc676371927029f95bd, n: '#0', v: 5 },
  { _id: 53d60cc676371927029f95bf, n: '#2', v: 0 },
  { _id: 53d60cc676371927029f95c0, n: '#3', v: 4 },
  { _id: 53d60cc676371927029f95c2, n: '#5', v: 2 },
  { _id: 53d60cc676371927029f95c3, n: '#6', v: 4 },
```

```
{ _id: 53d60cc676371927029f95c4, n: '#7', v: 1 },
{ _id: 53d60cc676371927029f95c5, n: '#8', v: 0 },
{ _id: 53d60cc676371927029f95c7, n: '#10', v: 5 },
{ _id: 53d60cc676371927029f95c9, n: '#12', v: 2 },
{ _id: 53d60cc676371927029f95cb, n: '#14', v: 1 },
{ _id: 53d60cc676371927029f95cc, n: '#15', v: 0 },
{ _id: 53d60cc676371927029f95cd, n: '#16', v: 1 },
{ _id: 53d60cc676371927029f95ce, n: '#17', v: 4 },
{ _id: 53d60cc676371927029f95cf, n: '#18', v: 4 },
{ _id: 53d60cc676371927029f95be, n: '#1', v: 8, valuable: true },
{ _id: 53d60cc676371927029f95c1, n: '#4', v: 6, valuable: true },
{ _id: 53d60cc676371927029f95c6, n: '#9', v: 7, valuable: true },
{ _id: 53d60cc676371927029f95c8, n: '#11', v: 8, valuable: true },
{ _id: 53d60cc676371927029f95ca, n: '#13', v: 8, valuable: true },
{ _id: 53d60cc676371927029f95d0, n: '#19', v: 6, valuable: true } ]
```

As you can see, 20 documents with a random *v* value are inserted. Afterwards, all documents whose value is greater than 5 have been updated to include the attribute *valuable* set to *true* - a total of 6 documents in this case.

Retrieving specific documents using filters

We have not yet performed any truly interesting queries against our collection. Up until now, all we did was to retrieve all documents contained in our collection by performing *collection.find()* without any filter. More complex filters are available of course.

Let's rewrite the existing code in order to retrieve only specific documents in our *doFind* function. First, a very simple example - we are going to retrieve all documents where the *v* attribute has a value of 5. To do so, we need to change line 11 from

```
collection.find().toArray(function (err, documents) {
```

to

```
collection.find({'v': 5}).toArray(function (err, documents) {
```

which results in only the matching documents being printed:

Inserted 20 documents:

```
[ { _id: 53d7698d99c6107303ad204c, n: '#8', v: 5 },  
  { _id: 53d7698d99c6107303ad2050, n: '#12', v: 5 } ]
```

Updated 10 documents:

```
[ { _id: 53d7698d99c6107303ad204c, n: '#8', v: 5 },  
  { _id: 53d7698d99c6107303ad2050, n: '#12', v: 5 } ]
```

We can write filters that match more than just one attribute, like this:

```
collection.find(  
  {  
    'v': 6,  
    'valuable': true  
  }  
)  
.toArray(function (err, documents) {
```

Both attributes of each document in the collection must match the given values for a document to be included. In other words, this is an AND query.

For the given collection, the above query doesn't make too much sense, because after the update, every document with a *v* value of 6 has the *valuable* attribute set to *true* anyways. But we might want to filter for all *valuable* documents whose *v* value is less than 8, like so:

```
collection.find(  
  {  
    'v': { '$lt': 8 },  
    'valuable': true  
  }  
)  
.toArray(function (err, documents) {
```

We already encountered the *greater than* operator *\$gt*, and *\$lt* is the opposite operator, *less than*. Other available operators are:

- *\$lte*: less than or equal
- *\$gte*: greater than or equal
- *\$ne*: not equal

Filters can also be combined into an OR query, if, for example, we want to query for all documents whose *v* value is either 5 or 8:


```
collection.find(
  {
    '$or': [
      {'v': 5},
      {'v': 8}
    ]
  }
).toArray(function (err, documents) {
```

And of course it is also possible to combine an AND query and an OR query:

```
collection.find(
  {
    'v': {'$gt': 3},
    '$or': [
      {'n': '#5'},
      {'n': '#10'}
    ]
  }
).toArray(function (err, documents) {
```

This will retrieve all documents that have a *v* value greater than 3 AND whose *n* value is #5 OR #10.

Speaking of the *n* attribute: it holds a string value, which cannot be queried with operators like *\$gt* or *\$lt*, of course. We can, however, apply regular expressions:

```
collection.find(
  {
    'n': /^#1/
  }
).toArray(function (err, documents) {
```

This matches all documents whose *n* value starts with #1.



All the filter variations discussed so far can also be applied to the *update* and *remove* methods of the *collection* object. This way, you can update or remove only specific documents in a given collection.

The *find* method also takes an *options* parameter. Possible options are *limit*, *skip*, and *sort*:

```
collection.find(
  {
    'n': /^#1/
  },
  {
    'limit': 5,
    'skip' : 2,
    'sort' : 'v'
  }
).toArray(function (err, documents) {
```

This one retrieves five documents whose *n* value starts with *#1*, skips the first two matches, and sorts the result set by the *v* value:

Inserted 20 documents:

```
[ { _id: 53d7739bbe04f81c05c13b6f, n: '#19', v: 5 },
  { _id: 53d7739bbe04f81c05c13b66, n: '#10', v: 5 },
  { _id: 53d7739bbe04f81c05c13b6b, n: '#15', v: 6 },
  { _id: 53d7739bbe04f81c05c13b67, n: '#11', v: 7 },
  { _id: 53d7739bbe04f81c05c13b68, n: '#12', v: 7 } ]
```

Updated 9 documents:

```
[ { _id: 53d7739bbe04f81c05c13b6f, n: '#19', v: 5 },
  { _id: 53d7739bbe04f81c05c13b66, n: '#10', v: 5 },
  { _id: 53d7739bbe04f81c05c13b6b, n: '#15', v: 6, valuable: true },
  { _id: 53d7739bbe04f81c05c13b68, n: '#12', v: 7, valuable: true },
  { _id: 53d7739bbe04f81c05c13b67, n: '#11', v: 7, valuable: true } ]
```

We can also sort over multiple fields and in different directions:

```
collection.find(
  {
    'n': /^#1/
  },
  {
    'limit': 5,
    'skip' : 2,
    'sort' : [ ['v', 'asc'], ['n', 'desc'] ]
  }
).toArray(function (err, documents) {
```

This sorts the matching documents first by their *v* value in ascending order, and documents with an identical *v* value are then ordered by their *n* value, in descending order, as can be seen here with the documents whose *v* value is *1*:

Inserted 20 documents:

```
[ { _id: 53d77475c0a6e83b0556d817, n: '#14', v: 1 },
  { _id: 53d77475c0a6e83b0556d815, n: '#12', v: 1 },
  { _id: 53d77475c0a6e83b0556d818, n: '#15', v: 3 },
  { _id: 53d77475c0a6e83b0556d81a, n: '#17', v: 4 },
  { _id: 53d77475c0a6e83b0556d814, n: '#11', v: 6 } ]
```

Updated 7 documents:

```
[ { _id: 53d77475c0a6e83b0556d817, n: '#14', v: 1 },
  { _id: 53d77475c0a6e83b0556d815, n: '#12', v: 1 },
  { _id: 53d77475c0a6e83b0556d818, n: '#15', v: 3 },
  { _id: 53d77475c0a6e83b0556d81a, n: '#17', v: 4 },
  { _id: 53d77475c0a6e83b0556d814, n: '#11', v: 6, valuable: true } ]
```

More complex update operations

We already learned how to add attributes or change the value of attributes in existing documents:

```
collection.update(
  {'n': /^#1/},
  {'$set': {'v': 5}},
  {'multi': true},
  function (err, count) {
    // ...
  }
)
```

This would set the value of the *v* attribute to 5 for all documents whose *n* value starts with #1.

This is probably the most regularly used update operation. Some other ways to update documents are available, however. Numeric attributes can be increased, for example:

```
collection.update(
  {'n': /^#1/},
  {'$inc': {'v': +1}},
  {'multi': true},
  function (err, count) {
    // ...
  }
)
```

There is no *\$dec* operation, but of course, we can increase by -1:

```
collection.update(
  {'n': /^#1/},
  {'$inc': {'v': -1} },
  {'multi': true},
  function (err, count) {
    // ...
  }
)
```

Multiplication is possible, too:

```
collection.update(
  {'n': /^#1/},
  {'$mul': {'v': 2} },
  {'multi': true},
  function (err, count) {
    // ...
  }
)
```

There are several other operators:

- `$rename` Renames a field.
- `$unset` Removes the specified field from a document.
- `$min` Only updates the field if the specified value is less than the existing field value.
- `$max` Only updates the field if the specified value is greater than the existing field value.
- `$currentDate` Sets the value of a field to current date, either as a Date or a Timestamp.

Please refer to [the MongoDB documentation³](#) for a detailed discussion of each operation.

The `update` method can be used to perform an operation where documents are updated if the filter being used is able to match existing documents, but a new document is inserted instead if no existing document could be matched. This is called an *upsert*:

³<http://docs.mongodb.org/manual/reference/operator/update-field/>

```
collection.update(  
  {'n': '#20'},  
  {'$set': {'n': '20', 'v': 1} },  
  {'multi': true, 'upsert': true},  
  function (err, count) {  
    // ...  
  }  
)
```

In our initial collection, only entries with *n* from #0 to #19 exist. Therefore, an update call with a filter set to {'n': '#20'} cannot match any document in the collection. Because the *upsert* option is set to *true*, a new document is inserted:

Inserted 20 documents:

```
[ { _id: 53d789d4fff187d607f03b3a, n: '#0', v: 7 },  
  ...  
  { _id: 53d789d4fff187d607f03b4c, n: '#18', v: 8 },  
  { _id: 53d789d4fff187d607f03b4d, n: '#19', v: 5 } ]
```

Updated 1 documents:

```
[ { _id: 53d789d4fff187d607f03b3a, n: '#0', v: 7 },  
  ...  
  { _id: 53d789d4fff187d607f03b4c, n: '#18', v: 8 },  
  { _id: 53d789d4fff187d607f03b4d, n: '#19', v: 5 },  
  { _id: 53d789d44266a0f5acef6f05, n: '#20', v: 1 } ]
```

Working with indexes

Like other databases, MongoDB needs to index data sets if we want queries on large sets to perform well.

In order to experience this firsthand, let's write a script which inserts 200,000 documents instead of just 20:

```
1  'use strict';
2
3  var MongoClient = require('mongodb').MongoClient;
4
5  MongoClient.connect(
6    'mongodb://127.0.0.1:27017/accounting',
7    function (err, connection) {
8      var collection = connection.collection('customers');
9
10     var doInsert = function (i) {
11       if (i < 200000) {
12         var value = Math.floor(Math.random() * 10);
13         collection.insert(
14           {'n': '#' + i, 'v': value},
15           function (err, count) {
16             doInsert(i + 1);
17           });
18       } else {
19         connection.close();
20       }
21     };
22
23     doInsert(0);
24
25   });
```

Call this script *insert.js* and execute it running `node insert.js`.

Now, we time how long it takes MongoDB to return 10 specific rows from this collection, using the following script:

```
1  'use strict';
2
3  var MongoClient = require('mongodb').MongoClient;
4
5  MongoClient.connect(
6    'mongodb://127.0.0.1:27017/accounting',
7    function (err, connection) {
8      var collection = connection.collection('customers');
9
10     collection.find(
11       {'v': {'$gt': 5}},
12       {
```

```
13     'skip': 100000,  
14     'limit': 10,  
15     'sort': 'v'  
16   }  
17   ).toArray(function (err, documents) {  
18     console.dir(documents);  
19     connection.close();  
20   });  
21  
22 });
```

This script queries the database for 10 documents whose *v* attribute is greater than 5, starting at document 100,000, sorting results by their *v* value.

Name this script *query.js* and run it like this:

```
time node query.js
```

This way, you'll get a poor man's performance benchmark. This is the result on my machine:

```
~$ time node query.js
```

```
[ { _id: 53d7955c3ac66b9e0af26745, n: '#36271', v: 6 },  
  { _id: 53d7955c3ac66b9e0af26757, n: '#36289', v: 6 },  
  { _id: 53d7955c3ac66b9e0af26761, n: '#36299', v: 6 },  
  { _id: 53d7955c3ac66b9e0af2677a, n: '#36324', v: 6 },  
  { _id: 53d7955c3ac66b9e0af2677b, n: '#36325', v: 6 },  
  { _id: 53d7955c3ac66b9e0af26792, n: '#36348', v: 6 },  
  { _id: 53d7955c3ac66b9e0af26794, n: '#36350', v: 6 },  
  { _id: 53d7955c3ac66b9e0af267b2, n: '#36380', v: 6 },  
  { _id: 53d7955c3ac66b9e0af267b3, n: '#36381', v: 6 },  
  { _id: 53d7955c3ac66b9e0af267b6, n: '#36384', v: 6 } ]
```

```
real 0m1.078s  
user 0m0.238s  
sys 0m0.040s
```

The script runtime is 1 second on average. Now, write the script that adds an index on the *v* attribute:

```
1 'use strict';
2
3 var MongoClient = require('mongodb').MongoClient;
4
5 MongoClient.connect(
6   'mongodb://127.0.0.1:27017/accounting',
7   function (err, connection) {
8     var collection = connection.collection('customers');
9
10    collection.ensureIndex('v', function(err, indexName) {
11      connection.close();
12    });
13
14  }
15 );
```

Call this one *addIndex.js* and execute it. Then, run the benchmark again. This is my result:

```
~$ time node query.js
```

```
[ { _id: 53d7955c3ac66b9e0af26745, n: '#36271', v: 6 },
  { _id: 53d7955c3ac66b9e0af26757, n: '#36289', v: 6 },
  { _id: 53d7955c3ac66b9e0af26761, n: '#36299', v: 6 },
  { _id: 53d7955c3ac66b9e0af2677a, n: '#36324', v: 6 },
  { _id: 53d7955c3ac66b9e0af2677b, n: '#36325', v: 6 },
  { _id: 53d7955c3ac66b9e0af26792, n: '#36348', v: 6 },
  { _id: 53d7955c3ac66b9e0af26794, n: '#36350', v: 6 },
  { _id: 53d7955c3ac66b9e0af267b2, n: '#36380', v: 6 },
  { _id: 53d7955c3ac66b9e0af267b3, n: '#36381', v: 6 },
  { _id: 53d7955c3ac66b9e0af267b6, n: '#36384', v: 6 } ]
```

```
real 0m0.269s
user 0m0.236s
sys 0m0.033s
```

From 1 second to 0.2 seconds - quite an improvement. This effect becomes more significant the larger the collection is. With a 2,000,000 documents collection, the query takes 5.6 seconds without an index and still only 0.2 seconds with an index on my machine.

Querying collections efficiently

When retrieving large result sets from a MongoDB collection, the same rule that applies to MySQL database result sets also applies here: reading the complete result set into our Node.js process at once

isn't going to be efficient resource-wise. Handling a result array with 2,000,000 entries will eat a lot of memory no matter what. This is what the *toArray* method, which we used until now, does:

```
collection.find().toArray(function (err, documents) {  
  // We now have one large documents array  
});
```

What *collection.find()* returns is a so-called *cursor object*. It can be transformed into an array, which is very convenient, but when handling lots of documents, it's better to handle each one separately using the cursor's *each* method:

```
collection.find().each(function (err, document) {  
  // We retrieve one document with each callback invocation  
});
```

This way, the cursor object will invoke the callback function for each document in the query result set, which means that Node.js only needs to claim memory for one document at a time.

Thus, instead of

```
collection.find().toArray(function (err, documents) {  
  console.dir(documents);  
});
```

we can do

```
collection.find().each(function (err, document) {  
  console.dir(document);  
});
```

if we want to print a whole result set to the screen. The callback will be invoked with the *document* parameter set to *null* if the end of the result set has been reached - therefore, we can rewrite our *doFind-doInsert-doUpdate* script like this:

```
1  'use strict';
2
3  var MongoClient = require('mongodb').MongoClient;
4
5  MongoClient.connect(
6    'mongodb://127.0.0.1:27017/accounting',
7    function (err, connection) {
8      var collection = connection.collection('customers');
9
10     var doFind = function (callback) {
11       collection.find(
12         {},
13         {'sort': '_id'}
14       ).each(function (err, document) {
15         if (document === null) {
16           callback();
17         } else {
18           console.dir(document);
19         }
20       });
21     };
22
23     var doInsert = function (i) {
24       if (i < 20) {
25         var value = Math.floor(Math.random() * 10);
26         collection.insert(
27           {'n': '#' + i, 'v': value},
28           function (err, count) {
29             doInsert(i + 1);
30           });
31       } else {
32         console.log();
33         console.log('Inserted', i, 'documents:');
34         doFind(function () {
35           doUpdate();
36         });
37       }
38     };
39
40     var doUpdate = function () {
41       collection.update(
42         {'n': /^#1/},
```

```
43     {'$mul': {'v': 2} },
44     {'multi': true},
45     function (err, count) {
46         console.log();
47         console.log('Updated', count, 'documents:');
48         doFind(function () {
49             collection.remove({}, function () {
50                 connection.close();
51             });
52         });
53     });
54 };
55
56 doInsert(0);
57
58 });
```

The `mongodb` driver also provides a convenient way to access the records of a result set through a streaming API. The following code excerpt shows a version of the `doFind` function which has been rewritten to make use of this API:

```
var doFind = function (callback) {

    var stream = collection.find(
        {},
        {'sort': '_id'}
    ).stream();

    stream.on('data', function(document) {
        console.dir(document);
    });

    stream.on('close', function() {
        callback();
    });

};
```

As you can see, it's quite simple: The `stream` method of the cursor object returned by `collection.find()` returns a stream object which emits `data` and `close` events, among others.

We can then attach callback functions to these events where we either print the incoming document, or call the provided callback function in order to continue our script.



The stream object returned by `collection.find().stream()` implements the `stream.Readable` interface. See [the corresponding Node.js documentation chapter](#)⁴ for more information.

Summary

In this chapter we learned about the workings of MongoDB and the `mongodb` driver, which enables us to insert, update, retrieve and remove documents from this NoSQL database. We learned about query filters and complex update operations, demonstrated how indexes can significantly accelerate query operations, and showed how to handle large result sets efficiently.

⁴http://nodejs.org/api/stream.html#stream_class_stream_readable

Part 2: Building a complete web application with Node.js and AngularJS

Introduction

I want to admit it right away: Looking at the amount of code, this part is probably more about AngularJS than Node.js - but I think it's important to see how Node.js can be used as one part in a complete application. Also, I'm a huge fan of AngularJS and am firmly convinced that knowing Node.js and AngularJS in combination is extremely valuable.

Nevertheless, besides teaching you a bit about building AngularJS frontends, the chapter does teach you a lot about Node.js, too. We will learn how to create a full-fledged RESTful Webservice API, and probably even more important, we will learn how to do so in a fully test-driven manner using end-to-end test specifications.

Additionally, we will talk about using a database abstraction layer, and we will learn how to integrate database migrations into our development and testing workflow. All in all, we will learn a lot about creating Node.js applications for the real world.



The code of the final application is also available at <https://github.com/manuelkiessling/nodecraftsman/tree/master/keyword-wrangler>

The requirements from a user's perspective

What we are going to build is an application which allows its users to comfortably manage a set of structured data from the world of food and cooking.

The data basically is a collection of keywords like *Aubergine*, *Knife*, or *Paul Bocuse*, and these keywords are mapped into categories, like *Vegetable*, *Utility*, or *Famous chef*.

What's needed now is an application where users can browse and edit this data: create new keywords, delete keywords, change the value of existing keywords, and change the category of keywords. Thus, at the heart of it we need a CRUD application, with a little bit of extra functionality.

It has already been decided that the users want a web-based solution. The data will be stored in a relational database.

This is how the user interface will look like:

```

+-----+ +-----+
New keyword: |           | | Category  V  Create
+-----+ +-----+

```

Keyword	Category	Ops
Aubergine	Vegetable	Save Delete
Knife	Utility	Save Delete
Paul Bocuse	Star chef	Save Delete

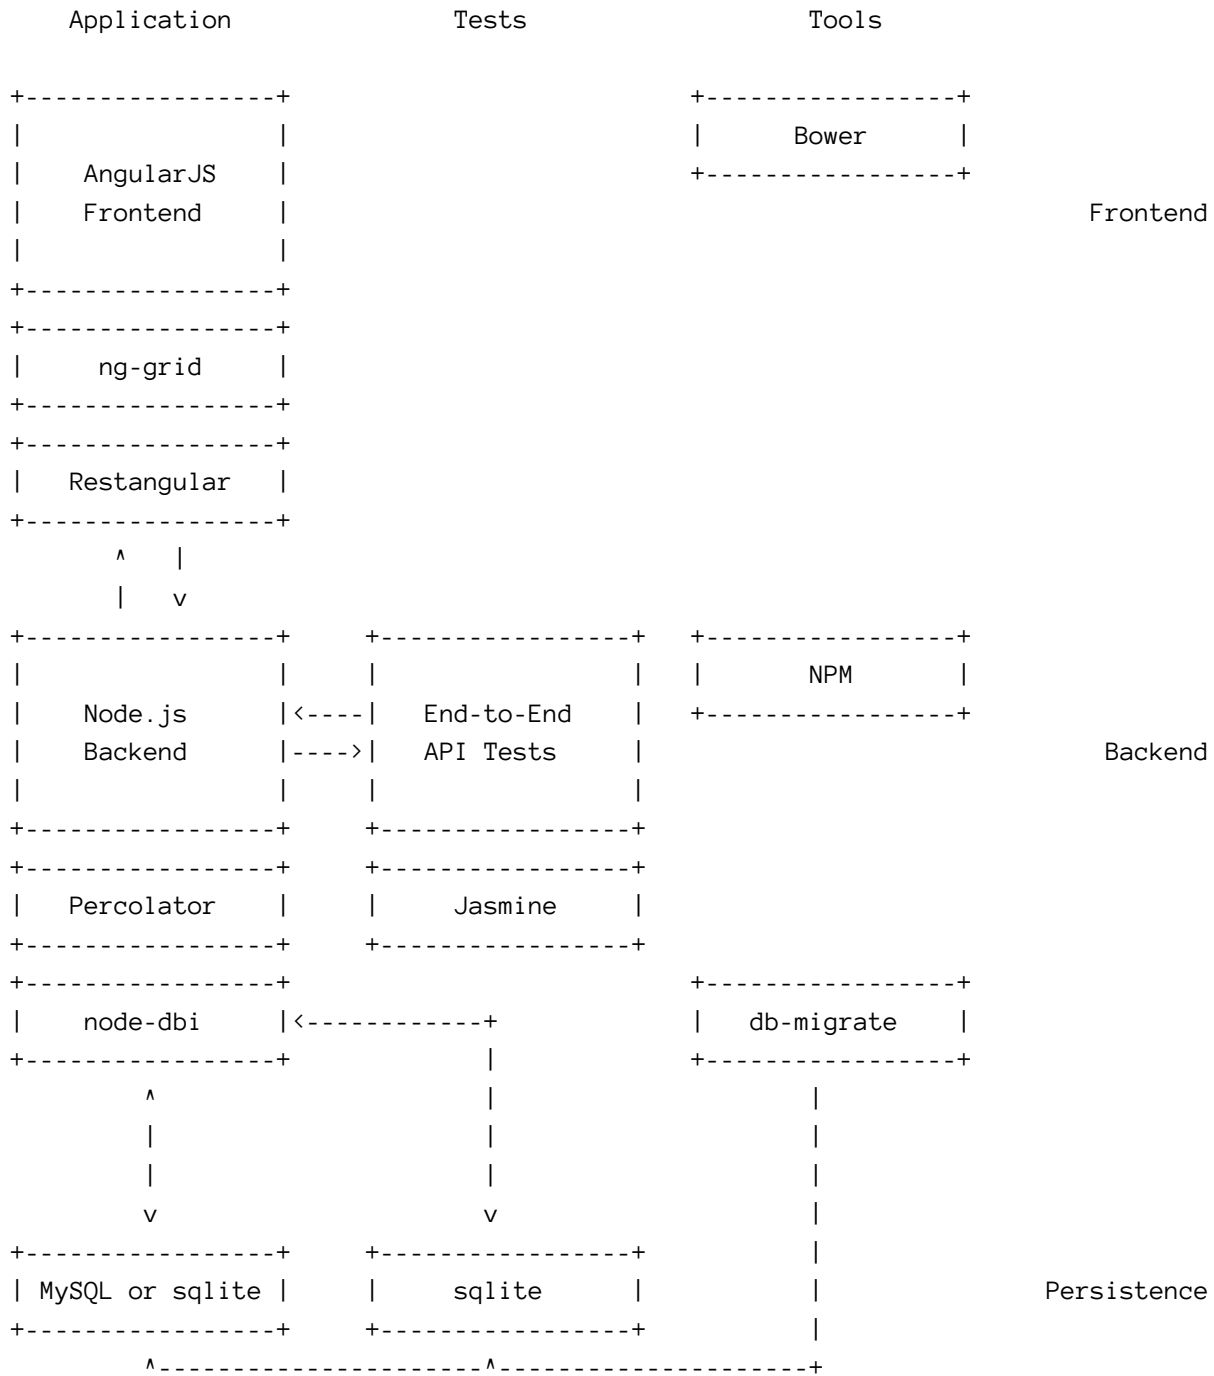
What this depicts is a screen with a table of keywords and the categories they are mapped to. Each row also has an operations column which allows to **Save** and **Delete** a keyword.

Above this table is a form which allows to **Create** a new keyword.

Being the JavaScript aficionados that we are, we are going to create the application as an AngularJS frontend that talks to a RESTful webservice API on the backend written in Node.js.

High level architecture overview

The following is an overview about the most important architectural parts of our application:



We will make use of the AngularJS plugin *Restangular* which will help us manage our webservice calls.

Our backend server will be built using *Percolator*, a library which streamlines the creation of RESTful APIs with Node.js.

We are going to use different database solutions for testing & development (sqlite) and production (MySQL), which is why we will use *node-dbi*, a database access abstraction layer.

In order to ease database schema management in different environments, we will use the *db-migrate* tool that handles the details of versioning the structural changes of our project's database.

Once again *Jasmine* will be our testing framework - this time, however, we are going to create end-to-end tests instead of unit tests.

As always, *NPM* will be our trusty companion when it comes to dependency management on the backend - on the frontend side, *bower* will take that role.

Setting up the development environment

This is going to be our largest application to date, so let's make sure we can work comfortably.

Let's start by creating a project folder - let's call our project *Keyword Wrangler*:

```
$ mkdir keyword-wrangler
```

Next, we initialize a Node.js project in this folder by running `npm init` and answering the upcoming questions as follows:

```
$ npm init
```

This utility will walk you through creating a package.json file.

It only covers the most common items, and tries to guess sane defaults.

See ``npm help json`` for definitive documentation on these fields and exactly what they do.

Use ``npm install <pkg> --save`` afterwards to install a package and save it as a dependency in the package.json file.

Press `^C` at any time to quit.

```
name: (keyword-wrangler)
version: (1.0.0) 0.0.1
description: A tool to manage keywords
entry point: (index.js) src/backend/index.js
test command:
git repository:
keywords:
author:
license: (ISC)
```

About to write to `keyword-wrangler/package.json`:

```
{
  "name": "keyword-wrangler",
  "version": "0.0.1",
  "description": "A tool to manage keywords",
  "main": "src/backend/index.js",
  "scripts": {
    "test": "echo \"Error: no test specified\" && exit 1"
  },
  "author": "",
  "license": "ISC"
}
```

Is this ok? (yes) yes



Make sure to execute the `npm init` command *within* the `keyword-wrangler` folder!

This gives us a sensible `package.json` file to start with. We can now start to pull in all the dependencies we need.

We are going to need the following Node.js packages:

- `percolator` - that's the REST framework we are going to use
- `node-dbi` - an SQL abstraction layer
- `mysql` - The MySQL db driver, used by `node-dbi`
- `sqlite3` - The sqlite3 db driver, used by `node-dbi`
- `db-migrate` - an SQL migrations manager
- `async` - we are going to write some asynchronous orchestration
- `jasmine` - the test framework
- `request` - this enables us to talk to the API we will build from within our tests
- `bower` - the dependency and package manager for the frontend application

Running

```
$ npm install percolator node-dbi mysql db-migrate async bower --save
$ npm install jasmine-node request sqlite3 --save-dev
```

will install the most recent stable versions of these packages and add them as dependencies to the `package.json` file. We only need `Jasmine` and `request` during development, which is why we install them separately with `--save-dev`. This allows to selectively install only the non-dev packages for production (per default, however, `npm` always installs all dependencies).

Milestone 1: A first passing test against the server

We will approach the requirements through a set of milestones. The first milestone we will try to reach is a very first passing end-to-end Jasmine test case against our webservice API.

Let's see what we need to reach that milestone:

- We need to be able to execute Jasmine test cases
- We need a Jasmine spec that requests our http backend
- We need an http backend that responds to the request

In order to verify that we are good to go in regards to executing Jasmine test cases, let's create a subfolder *spec/e2e* in our project directory:

```
mkdir -p spec/e2e
```

Next, we create the file that will hold the end-to-end tests for our yet-to-be-written API server. We name it *apiSpec.js* and store it within the *spec/e2e* folder we just created:

```
1  'use strict';
2
3  describe('The API', function () {
4
5      it('should just work', function () {
6          expect(true).toBe(true);
7      });
8
9  });
```

As long as the test infrastructure itself works correctly, this will of course pass. Let's verify that:

```
$ ./node_modules/.bin/jasmine-node --verbose ./spec/
```

```
The API - 1 ms
  should just work - 1 ms
```

```
Finished in 0.004 seconds
1 test, 1 assertion, 0 failures, 0 skipped
```



`./node_modules/.bin` is a folder with shortcuts to all the executables within `node_modules`. The actual executable in this case is `./node_modules/jasmine-node/bin/jasmine-node`

Ok, great, we can build on that. The goal of our first milestone is to have a test case which successfully requests a server backend and checks for the correct response. Let's rewrite `apiSpec.js` accordingly:

```
1 'use strict';
2
3 var request = require('request');
4
5 describe('The API', function () {
6
7   it('should respond to a GET request at /api/keywords/', function (done) {
8     request.get(
9       {
10        'url': 'http://localhost:8080/api/keywords/',
11        'json': true
12      },
13      function (err, res, body) {
14        expect(res.statusCode).toBe(200);
15        expect(body.foo).toEqual('bar');
16        done();
17      });
18   });
19
20 });
```

If we re-execute the test command, we get - nothing. Not only a failed test run, but literally no output to the console whatsoever. The problem is that there isn't any server yet that could respond to our GET request; however, *Jasmine* remains silent about that.

The reason is that requesting the missing HTTP server leads to an exception, and we need to explicitly tell *Jasmine* that we want to see exceptions:

```
$ ./node_modules/.bin/jasmine-node --verbose --captureExceptions ./spec/
```

```
TypeError: Cannot read property 'statusCode' of undefined
    at Request._callback (keyword-wrangler/spec/e2e/apiSpec.js:14:19)
    at self.callback (keyword-wrangler/node_modules/request/request.js:373:22)
    at Request.emit (events.js:95:17)
    at Request.onRequestError (keyword-wrangler/node_modules/request/request.js:\
971:8)
    at ClientRequest.emit (events.js:95:17)
    at Socket.socketErrorListener (http.js:1552:9)
    at Socket.emit (events.js:95:17)
    at net.js:441:14
    at process._tickCallback (node.js:442:13)
```

Ok, it's time to create a web server. As said, we will use a library for this, namely *Percolator*. The spec we defined expects a server at port 8080 that responds with `{"foo": "bar"}` when receiving a GET request at `/api/keywords/`. Let's see how to implement this behaviour with *Percolator*.

Put the following into file `src/backend/index.js`:

```
1  'use strict';
2
3  var Percolator = require('percolator').Percolator;
4
5  var port = 8080;
6  var server = Percolator({'port': port});
7
8  server.route('/api/keywords',
9    {
10     GET: function (req, res) {
11       res.object({'foo': 'bar'}).send()
12     }
13   }
14 );
15
16 server.listen(function() {
17   console.log('Server started and listening on port', port);
18 });
```

Then, start this server in a new console window:

Milestone 1: A first passing test against the server

109

```
$ node src/backend/index.js
```

Server started and listening on port 8080

Now, running the test command again should look much better:

```
$ ./node_modules/.bin/jasmine-node --verbose --captureExceptions ./spec/
```

```
The API - 19 ms
```

```
  should respond to a GET request at /api/keywords/ - 18 ms
```

```
Finished in 0.021 seconds
```

```
1 test, 2 assertions, 0 failures, 0 skipped
```

Great, our first milestone has been reached.

Milestone 2: The API responds with actual database content

Let's now look at how to integrate database access into our server application. We defined that within our local development and test environment, we would like to use *sqlite*. Let's start with this.

Abstracting database access

In order to connect to the database, we are going to use an abstraction layer by utilizing *node-dbi*. This allows us to transparently use *MySQL* at a later point in time, without the need to change the code that does database operations.

First, we should change our spec in order to have it expect actual database content. A request to `/api/keywords/` should return a JSON structure that carries an array of objects, where each object represents a keyword with its ID, value, and the ID of its category:

```
[
  {"id": 1, "value": "Aubergine", "categoryID": 1},
  {"id": 2, "value": "Onion", "categoryID": 1},
  {"id": 3, "value": "Knife", "categoryID": 2}
]
```

A data structure like this would be stored in a relational database like this:

Table: keyword

id	value	categoryID
1	Aubergine	1
2	Onion	1
3	Knife	2

Table: category

id	name
1	Vegetable
2	Utility



In this structure, *keyword.categoryID* is a foreign key to *category.id*. For now, however, we will only work with the *keyword* table.

Assuming the above example database content during a test run, our spec expectation needs to look like this (in file *spec/e2e/apiSpec.js*):

```

1  'use strict';
2
3  var request = require('request');
4
5  describe('The API', function () {
6
7    it('should respond to a GET request at /api/keywords/', function (done) {
8      var expected = {
9        "_items": [
10         { 'id': 1, 'value': 'Aubergine', 'categoryID': 1 },
11         { 'id': 2, 'value': 'Onion', 'categoryID': 1 },
12         { 'id': 3, 'value': 'Knife', 'categoryID': 2 }
13       ]
14     };
15
16     request.get(
17       {
18         'url': 'http://localhost:8080/api/keywords/',
19         'json': true
20       },
21       function (err, res, body) {
22         expect(res.statusCode).toBe(200);
23         expect(body).toEqual(expected);
24         done();
25       });
26 });

```



```
27
28 });
```



Note how *expected* is an object with a field *_items*, which contains the actual list of keywords. *Percolator* doesn't allow to return an array of things directly.

There is still some work to do in order to make this succeed.

Ensuring a clean slate for test runs

We are going to make our lives a lot easier if we write some helper code for our tests which ensures that at the start of each test run, the database is empty. Let's put this into *spec/resetDatabase.js*:

```
1  'use strict';
2
3  var async = require('async');
4
5  var resetDatabase = function (dbSession, callback) {
6
7    async.series(
8      [
9
10     function (callback) {
11       dbSession.remove('keyword', '1', function (err) {
12         callback(err)
13       });
14     },
15
16     function (callback) {
17       dbSession.remove('category', '1', function (err) {
18         callback(err)
19       });
20     }
21   ],
22
23   function (err, results) {
24     callback(err);
25   }
26 );
```

```
28
29 };
30
31 module.exports = resetDatabase;
```

As you can see, we make use of a *dbSession* object. We need to create a module where this is set up, in file *src/backend/dbSession.js*:

```
1 'use strict';
2
3 var DBWrapper = require('node-dbi').DBWrapper;
4
5 var dbWrapper = new DBWrapper('sqlite3', {'path': '/var/tmp/keyword-wrangler.test\
6 t.sqlite'});
7 dbWrapper.connect();
8 module.exports = dbWrapper;
```



This is a very simplistic implementation - we will later refine it in order to allow for different database setups depending on the environment.

Let's now create the *sqlite* database file with the *keyword* and *category* table:

```
$ sqlite3 /var/tmp/keyword-wrangler.test.sqlite
```

```
sqlite> CREATE TABLE keyword (id INTEGER PRIMARY KEY, value TEXT, categoryID INT\
EGER);
sqlite> CREATE TABLE category (id INTEGER PRIMARY KEY, name TEXT);
sqlite> .quit
```



Install the *sqlite3* package if it is not available on your system!

Completing the first spec

We now have the boilerplate code and database structures needed to further complete our spec:

```
1  'use strict';
2
3  var request = require('request');
4  var dbSession = require('../src/backend/dbSession.js');
5  var resetDatabase = require('../resetDatabase.js');
6  var async = require('async');
7
8  describe('The API', function () {
9
10   it('should respond to a GET request at /api/keywords/', function (done) {
11     var expected = {
12       "_items": [
13         { 'id': 1, 'value': 'Aubergine', 'categoryID': 1 },
14         { 'id': 2, 'value': 'Onion', 'categoryID': 1 },
15         { 'id': 3, 'value': 'Knife', 'categoryID': 2 }
16       ]
17     };
18
19     async.series(
20       [
21
22         function(callback) {
23           resetDatabase(dbSession, callback);
24         },
25
26         function(callback) {
27           dbSession.insert(
28             'keyword',
29             { 'value': 'Aubergine', 'categoryID': 1 },
30             function(err) { callback(err) });
31         },
32
33         function(callback) {
34           dbSession.insert(
35             'keyword',
36             { 'value': 'Onion', 'categoryID': 1 },
37             function(err) { callback(err) });
38         },
39
40         function(callback) {
41           dbSession.insert(
42             'keyword',
```

```
43         {'value': 'Knife', 'categoryID': 2},
44         function(err) { callback(err) });
45     }
46
47 ],
48
49 function(err, results) {
50     request.get(
51         {
52             'url': 'http://localhost:8080/api/keywords/',
53             'json': true
54         },
55         function (err, res, body) {
56             expect(res.statusCode).toBe(200);
57             expect(body).toEqual(expected);
58             done();
59         }
60     );
61 }
62
63 );
64
65 });
66
67 });
```

With this, our spec is complete - we empty the database, insert our test data, and expect the web service API to respond with the data we just inserted. This is not yet the behaviour of the API - we need to rewrite *src/backend/index.js*:

```
1  'use strict';
2
3  var Percolator = require('percolator').Percolator;
4  var dbSession = require('../src/backend/dbSession.js');
5
6  var port = 8080;
7  var server = Percolator({'port': port, 'autoLink': false});
8
9  server.route('/api/keywords',
10 {
11     GET: function (req, res) {
12         dbSession.fetchAll('SELECT id, value, categoryID FROM keyword ORDER BY id'\
```

```
13 , function(err, rows) {
14     if (err) {
15         console.log(err);
16         res.status.internalServerError(err);
17     } else {
18         res.collection(rows).send();
19     }
20 });
21 }
22 }
23 );
24
25 server.listen(function() {
26     console.log('Server started and listening on port', port);
27 });
```



Don't forget to restart the server after making changes to it!

Our tests now pass again, and thus our second milestone has been reached.

Milestone 3: Setting the stage for a continuous delivery workflow

We are already growing our application driven by tests, however, there are still some inconveniences that hinder us from going fully continuous. If the goal is to put every change we make into production automatically (safeguarded by our tests, of course), then we need to automate the things we still do manually: starting the server and changing the database structure.

Let's tackle the server startup first. It would be great if each test case would start and stop the server process as needed. This is possible by refactoring the server code a bit.

We need to split into two files what is currently handled in *src/backend/index.js*. The goal here is to have the server *logic* in one file, and the server *startup* in another one. This way, our tests can start the server logic up by themselves, and for production, we still have a file we can execute.

To achieve this, let's create a file *src/backend/server.js* with the following content:

```
1  'use strict';
2
3  var Percolator = require('percolator').Percolator;
4  var dbSession = require('../src/backend/dbSession.js');
5
6  var Server = function(port) {
7    var server = Percolator({'port': port, 'autoLink': false});
8
9    server.route('/api/keywords',
10     {
11       GET: function (req, res) {
12         dbSession.fetchAll('SELECT id, value, categoryID FROM keyword ORDER BY id', function(err, rows) {
13           if (err) {
14             console.log(err);
15             res.status.internalServerError(err);
16           } else {
17             res.collection(rows).send();
18           }
19         });
20       });
21     }
22 }
```

```
23   );
24
25   return server;
26 };
27
28 module.exports = { 'Server': Server};
```

This isn't that much different from what we had in *index.js* until now. But instead of configuring the server and starting it up directly, we now merely provide a *Server* function that sets up a *Percolator* server object for us. It's then our business to get that server going - we do so in *src/backend/index.js*:

```
1  'use strict';
2
3  var Server = require('./server.js').Server;
4
5  var server = Server('8080');
6
7  server.listen(function() {
8    console.log('Server started and listening on port', server.options.port);
9  });
```

This way, we can still start the server via `node src/backend/index.js`. However, we can now also manage the server from other parts of our code - for example, from our tests in *spec/e2e/apiSpec.js*.

```
1  'use strict';
2
3  var request = require('request');
4  var dbSession = require('../src/backend/dbSession.js');
5  var Server = require('../src/backend/server.js').Server;
6  var resetDatabase = require('../resetDatabase.js');
7  var async = require('async');
8
9  describe('The API', function () {
10
11    var server;
12
13    beforeEach(function (done) {
14      server = Server('8081');
15      server.listen(function (err) {
16        resetDatabase(dbSession, function() {
17          done(err);
18        });
19      });
```

```
19     });
20   });
21
22   afterEach(function (done) {
23     server.close(function () {
24       resetDatabase(dbSession, function() {
25         done();
26       });
27     });
28   });
29
30   it('should respond to a GET request at /api/keywords/', function (done) {
31     var expected = {
32       "_items": [
33         { 'id': 1, 'value': 'Aubergine', 'categoryID': 1 },
34         { 'id': 2, 'value': 'Onion', 'categoryID': 1 },
35         { 'id': 3, 'value': 'Knife', 'categoryID': 2 }
36       ]
37     };
38
39     async.series(
40       [
41
42         function(callback) {
43           dbSession.insert(
44             'keyword',
45             { 'value': 'Aubergine', 'categoryID': 1 },
46             function(err) { callback(err) });
47         },
48
49         function(callback) {
50           dbSession.insert(
51             'keyword',
52             { 'value': 'Onion', 'categoryID': 1 },
53             function(err) { callback(err) });
54         },
55
56         function(callback) {
57           dbSession.insert(
58             'keyword',
59             { 'value': 'Knife', 'categoryID': 2 },
60             function(err) { callback(err) });
```



```
61     }
62
63   ],
64
65   function(err, results) {
66     if (err) throw(err);
67     request.get(
68       {
69         'url': 'http://localhost:8081/api/keywords/',
70         'json': true
71       },
72       function (err, res, body) {
73         expect(res.statusCode).toBe(200);
74         expect(body).toEqual(expected);
75         done();
76       }
77     );
78   }
79
80 );
81
82 });
83
84 });
```

Now, our specs no longer require an already running server: *before each* test case, the server is started, and *after each* case, it is shut down. We also moved the database reset step into these blocks.

Also, we are now using port *8081* for our tests, versus port *8080* which is used for production (do not forget to change the port on line 14).

With this, you can now stop the running server, and you can execute test runs without the need to start the server.

Introducing automatic database migrations

Let's take the final step in automation by setting up database migrations. A migration is a change to the database structure of our application. E.g., if we change our code and it now expects a new field in one of our tables, the migration would be the code that alters the table structure and adds the missing field. Making these migrations part of our code base allows us to automate this - when the tests or the application is started, all migrations that are not yet applied to the database are applied automatically; we no longer need to ensure manually that our code and our database structures are in sync.

The Node.js module that takes care of this for us is *db-migrate*. We already installed it via *npm*.

This module expects a *database.json* file where it can find the database connection parameters for different environments. An environment is, for example, *test*, or *production*. This way, we can use different database configurations for our different environments. We declared that we would like to use *sqlite* for our tests and *MySQL* in production; the *database.json* file makes this kind of setup simple (and, as we will see, we can use it for *db-migrate* and from our own code).

Here is the *database.json* file we need:

```
1 {
2   "test": {
3     "driver": "sqlite3",
4     "filename": "/var/tmp/keyword-wrangler.test.sqlite"
5   },
6
7   "dev": {
8     "driver": "sqlite3",
9     "filename": "/var/tmp/keyword-wrangler.dev.sqlite"
10  },
11
12  "production": {
13    "driver": "mysql",
14    "host": "localhost",
15    "user": "root",
16    "password": "root",
17    "database": "keyword_wrangler"
18  }
19 }
```

The file needs to reside in the root folder of our project.

We also need a place for our migrations. By default, the migration files are placed in a folder called *migrations*, again in the root folder of our project.

We can now create our first migration; it will programmatically create the *keyword* and *category* tables we already created manually. Migrations are files with JavaScript code, but we don't need to create them manually. *db-migrate* does this for us:

```
$ ./node_modules/.bin/db-migrate create createKeywordAndCategoryTable --env test
```

This gives you a file with a name like *migrations/20150127080251-createKeywordAndCategoryTable.js* and the following content:

```
1 var dbm = require('db-migrate');
2 var type = dbm.dataType;
3
4 exports.up = function(db, callback) {
5
6 };
7
8 exports.down = function(db, callback) {
9
10 };
```

Two functions, one *up*, one *down*. The *up* function carries the code that is to be executed if a migration needs to be applied. The *down* function carries the reverse, i.e., the code that is needed to un-apply or rollback the migration.

Given our database structure, here is how this looks in practice:

```
1 'use strict';
2
3 var dbm = require('db-migrate');
4 var type = dbm.dataType;
5 var async = require('async');
6
7 exports.up = function(db, callback) {
8   async.series(
9     [
10
11     db.createTable.bind(db, 'keyword', {
12       id: { type: 'int', primaryKey: true, autoIncrement: true, notNull: true \
13     },
14     value: { type: 'string', length: '128', notNull: true, unique: true },
15     categoryID: { type: 'int', notNull: true }
16   })),
17
18   db.createTable.bind(db, 'category', {
19     id: { type: 'int', primaryKey: true, autoIncrement: true, notNull: true \
20   },
21     name: { type: 'string', length: '128', notNull: true }
22   })
23
24 ], callback);
25 };
26
```

```
27 exports.down = function(db, callback) {
28   async.series(
29     [
30       db.dropTable.bind(db, 'keyword'),
31       db.dropTable.bind(db, 'category')
32     ], callback);
33   };
```

The *up* step will create our tables, the *down* step will remove them. This is the equivalent of the manual SQL operations we did earlier. And because we already did that manually, we need to remove them before our migrations can work. The most simple way to do this is by deleting the *sqlite* database file:

```
$ rm /var/tmp/keyword-wrangler.test.sqlite
```

We are now able to run our migrations:

```
$ ./node_modules/.bin/db-migrate up --env test
```

```
[INFO] Processed migration 20150127080251-createKeywordAndCategoryTable
[INFO] Done
```

Looks good. This creates the tables, and additionally creates a table named *migrations*, where *db-migrate* keeps track of the migrations already applied:

```
$ sqlite3 /var/tmp/keyword-wrangler.test.sqlite
```

```
sqlite> SELECT * FROM migrations;
```

```
1|20150127080251-createKeywordAndCategoryTable|1422435911741
```

So, from now on we need to combine our migration runs with our test runs; whenever we change the code and its tests in a way that makes them expect a different database structure, we need to create a new migrations file, program the required database changes into them, and can then apply them, followed by a test run:

```
$ ./node_modules/.bin/db-migrate up --env test && \
  ./node_modules/.bin/jasmine-node --verbose --captureExceptions ./spec/
```

The great thing is that now all development steps can be automatized, and we can ensure that code and database are always in sync. Imagine that another developer clones the repository with your code. Given that she has a working Node.js environment, all she needs to do is:

- check out the code
- run `npm install`
- run the migrations

and she is immediately up-to-date with her development environment for your project.

Milestone 3 has been reached!

Milestone 4: Giving users a frontend

We have now reached a point where our backend service deserves a frontend. As said, we want this to be based on AngularJS. Furthermore, we will use the *Restangular* library (which will ease the communication with our backend webservice), and *ngGrid*, which makes building dynamic tables straightforward.

Setting up frontend dependencies through bower

The very first step is to utilize *Bower*, the package manager for frontend JavaScript projects, to pull in the dependencies of our project.

We need to create a directory *src/frontend*, and place a *bower.json* file into it with the following contents:

```
1 {
2   "name": "keyword-wrangler-client",
3   "dependencies": {
4     "angular": "1.3.15",
5     "angular-route": "1.3.15",
6     "angular-grid": "2.0.14",
7     "restangular": "1.4.0",
8     "bootstrap": "3.3.4",
9     "async": "0.9.2"
10  },
11  "analytics": false
12 }
```

Then, while in folder *keyword-wrangler/src/frontend*, run:

```
$ ../../node_modules/.bin/bower install
```

This will install all dependencies into *src/frontend/bower_components*.

Serving the frontend through the backend server

Before we start building our frontend app, let's make sure we can browse to it.

To do so, change line 7 in *src/backend/server.js* from

```
var server = Percolator({'port': port, 'autoLink': false});
```

to

```
var server = Percolator({'port': port, 'autoLink': false, 'staticDir': __dirname\  
+ '/../frontend'});
```

With this change, Percolator will serve everything in *src/frontend* as static files (while still serving the API at the */api* route).

Adding the frontend code

Time to start building the app. Create *src/frontend/index.html* as follows:

```
1 <!DOCTYPE html>  
2 <html ng-app="app">  
3  
4   <head>  
5     <meta charset="UTF-8"/>  
6     <title>Keyword Wrangler</title>  
7  
8     <link rel="stylesheet" href="bower_components/bootstrap/dist/css/bootstr  
9 ap.css">  
10    <link rel="stylesheet" href="bower_components/angular-grid/ng-grid.min.c\  
11 ss">  
12    <style>  
13      .gridStyle {  
14        margin-top: 20px;  
15        height: 800px;  
16      }  
17    </style>  
18  </head>  
19  
20  <body>  
21    <div id="main" class="container">  
22      <div class="span12" ng-view></div>  
23    </div>  
24  
25    <script src="bower_components/async/lib/async.js"></script>  
26    <script src="bower_components/jquery/dist/jquery.js"></script>  
27    <script src="bower_components/lodash/dist/lodash.js"></script>
```

```
28     <script src="bower_components/angular/angular.js"></script>
29     <script src="bower_components/angular-route/angular-route.js"></script>
30     <script src="bower_components/restangular/dist/restangular.js"></script>
31     <script src="bower_components/angular-grid/build/ng-grid.js"></script>
32
33     <script src="app/app.js"></script>
34     <script src="app/lib/resolveEntity.js"></script>
35     <script src="app/lib/RepositoryFactory.js"></script>
36     <script src="app/keywords/KeywordsController.js"></script>
37 </body>
38
39 </html>
```

When running the backend server locally, this will be served at <http://localhost:8080/>, and will host our single-page app.

As you can see on line 32 and following, we need to provide several JavaScript files - these are the components of our AngularJS application.

Let's start with `src/frontend/app/app.js`:

```
1  'use strict';
2
3  (function() {
4
5      var app = angular.module('app', ['ngRoute', 'ngGrid', 'restangular']);
6
7      app.config(['$routeProvider',
8          function($routeProvider) {
9
10         // This makes app/keywords/KeywordsController.js handle the http://localho\
11 st:8080/#/ URL
12         $routeProvider.
13             when('/', {
14                 templateUrl: 'app/keywords/partials/editor.html',
15                 controller: 'KeywordsController'
16             });
17
18         });
19
20     })();
```

This defines the main entry point for the AngularJS applications. It sets up dependencies and in-app routing.

The key component here is the *KeywordController*. Put this into *src/frontend/app/keywords/KeywordsController.js*:

```
1  'use strict';
2
3  (function() {
4
5      var app = angular.module('app');
6
7      app.controller('KeywordsController', function($scope, RepositoryFactory, resolveEntity) {
8
9
10         /* == Frontend Initialization == */
11
12         /* All of this happens as soon as the page loads */
13
14         /* resolveEntity is a helper function which is used in partials/keywordCategoryGridCell.html
15 in order to display the name of a keyword category based on its id */
16         $scope.resolveEntity = resolveEntity;
17
18
19         /* A repository is the connection between this controller and the REST Api.
20 We use one for keyword categories... */
21         var KeywordCategoriesRepository = new RepositoryFactory({
22             endpoint: 'keywords/categories',
23             retrieveItems: function (data) {
24                 return data._items;
25             }
26         });
27
28         /* ...and one for keywords */
29         var KeywordsRepository = new RepositoryFactory({
30             endpoint: 'keywords',
31             retrieveItems: function(data) {
32                 return data._items;
33             }
34         });
35
36         /* When the frontend loads, we want the controller to immediately
37 load all keyword categories and categories from the API */
38         KeywordCategoriesRepository.readAll().then(function(keywordCategories) {
39             $scope.keywordCategories = keywordCategories;
40             KeywordsRepository.readAll().then(function(keywords) {
```

```
41     $scope.keywords = keywords;
42   });
43 });
44
45   /* The grid. This part is best coded while listening to the Tron: Legacy sou\
46 ndtrack. */
47   $scope.keywordsGridOptions = {
48     data: 'keywords', // This makes the grid use the data in $scope.keywords
49     enableCellSelection: false, // breaks edit of cells with select element if\
50 true
51     enableCellEdit: true,
52     keepLastSelected: false,
53     enableRowSelection: false,
54     multiSelect: false,
55     enableSorting: true,
56     enableColumnResize: true,
57     enableColumnReordering: true,
58     showFilter: false,
59     rowHeight: '40',
60     columnDefs: [
61       {
62         field: 'id',
63         displayName: 'ID',
64         enableCellEdit: false,
65         width: '80px'
66       },
67       {
68         field: 'value',
69         displayName: 'Value'
70       },
71       {
72         /* The keyword category field does not use the build-in cell template,\
73 but our own */
74         field: 'keywordCategoryID',
75         displayName: 'Category',
76         cellTemplate: 'app/keywords/partials/keywordCategoryGridCell.html',
77         editableCellTemplate: 'app/keywords/partials/keywordCategoryGridCellEd\
78 itor.html'
79       },
80       {
81         /* Same goes for the operations column */
82         field: '',
```

```
83     displayName: 'Operations',
84     cellTemplate: 'app/keywords/partials/operationsGridCell.html',
85     enableCellEdit: false,
86     sortable: false
87   }
88 ]
89 };
90
91
92 /* == Frontend Operations == */
93
94 /* These functions are called when the frontend is operated, e.g., if a butt\
95 on is clicked */
96
97 $scope.createKeyword = function(newKeyword) {
98   $scope.$broadcast('ngGridEventEndCellEdit');
99   if (newKeyword.value.length > 0) {
100     KeywordsRepository.createOne(newKeyword).then(function () {
101       KeywordsRepository.readAll().then(function (keywords) {
102         $scope.keywords = keywords;
103       });
104     });
105   }
106 };
107
108 $scope.updateKeyword = function(keyword) {
109   $scope.$broadcast('ngGridEventEndCellEdit');
110   KeywordsRepository.updateOne(keyword);
111 };
112
113 $scope.deleteKeyword = function(keyword) {
114   $scope.$broadcast('ngGridEventEndCellEdit');
115   KeywordsRepository.deleteOne(keyword).then(function() {
116     KeywordsRepository.readAll().then(function(keywords) {
117       $scope.keywords = keywords;
118     });
119   });
120 };
121
122 /* These are here to make the grid behave cleanly in regards to the keyword \
123 category select */
124 $scope.stopEditingKeywordCategory = function() {
```

```
125     $scope.$broadcast('ngGridEventEndCellEdit');
126   };
127
128   $scope.$on('ngGridEventRows', function(newRows) {
129     $scope.$broadcast('ngGridEventEndCellEdit');
130   });
131
132 });
133
134 })();
```

This controller has two dependencies, *RepositoryFactory* and *resolveEntity*. Let's define those.

In file *src/frontend/app/lib/RepositoryFactory.js*:

```
1  'use strict';
2
3  (function() {
4
5    var app = angular.module('app');
6
7    app.factory(
8      'RepositoryFactory',
9      ['Restangular', '$q', RepositoryFactory]
10   );
11
12   function RepositoryFactory(Restangular, $q) {
13
14     Restangular.setBaseUrl('/api/');
15
16     var Repository = function(options) {
17       this.endpoint = options.endpoint;
18       this.retrieveItems = options.retrieveItems;
19     };
20
21     Repository.prototype.readAll = function() {
22       var self = this;
23       var deferred = $q.defer();
24       Restangular.all(self.endpoint + '/').doGET().then(function(data) {
25         var items = self.retrieveItems(data);
26         deferred.resolve(items);
27       });
28       return deferred.promise;
```

```
29     };
30
31     Repository.prototype.createOne = function(newItem) {
32         var self = this;
33         var deferred = $q.defer();
34         Restangular.one(self.endpoint + '/', '').post('', newItem).then(function(r\
35 esponse) {
36             deferred.resolve(response);
37         });
38         return deferred.promise;
39     };
40
41     Repository.prototype.updateOne = function(item) {
42         var self = this;
43         var deferred = $q.defer();
44         Restangular.one(self.endpoint, item.id).post('', item).then(function(respo\
45 nse) {
46             deferred.resolve(response);
47         });
48         return deferred.promise;
49     };
50
51     Repository.prototype.deleteOne = function(item) {
52         var self = this;
53         var deferred = $q.defer();
54         Restangular.one(self.endpoint, item.id).remove().then(function(response) {
55             deferred.resolve(response);
56         });
57         return deferred.promise;
58     };
59
60     return Repository;
61 }
62 }
63
64 }());
```

This is basically a convenience wrapper for all the different API operations we are going to need for keywords and keyword categories: Reading, creating, updating, deleting, and merging.

Note how each operation is asynchronous and therefore returns promises. You can see how these operations are used in the controller on lines 37, 39, 95, 96, 105, 110, 111, 123 and 126.



Of course, the webservice routing and logic for these operations does not yet exist in our backend server (with the exception of *GET /api/keywords*. We will implement those soon.

We also need to create *src/frontend/app/lib/resolveEntity.js*, which is a helper function that is needed to get from category IDs to category names in the view of our app:

```
1  'use strict';
2
3  (function() {
4
5      var app = angular.module('app');
6
7      app.factory(
8          'resolveEntity',
9          [resolveEntity]
10     );
11
12     function resolveEntity() {
13
14         return function(queryParams) {
15             for (var i=0; i < queryParams.from.length; i++) {
16                 if (queryParams.from[i][queryParams.where] === queryParams.is) {
17                     return queryParams.from[i][queryParams.what];
18                 }
19             }
20             return false;
21         };
22
23     }
24
25 })();
```

Adding AngularJS view templates

With this, the JavaScript code for our AngularJS app is in place, but we aren't done yet. We also need four HTML views for the app: The main view with the editor, plus three custom grid cell templates, one for the operations column that holds the *Save* and *Delete* controls, one for the *Category* column cells (because there we need to translate from category IDs to category names), and one for the *Category* cells in edit mode, because when we edit the category of a keyword, we don't want to edit its name as a string, but rather choose from a list of existing categories through a dropdown control.

We start with *src/frontend/app/keywords/partials/editor.html*:

```

1 <div class="span12">
2   <form>
3     <input
4       class="input-medium"
5       ng-model="newKeyword.value"
6       ng-init="newKeyword.value = ''"
7       type="text"
8       placeholder="Name of new keyword"
9     >
10    <select
11      ng-model="newKeyword.categoryID"
12      ng-init="newKeyword.categoryID = 1"
13      ng-options="keywordCategory.id as keywordCategory.name for keywo\
14 rdCategory in keywordCategories" >
15    </select>
16    <button ng-class="newKeyword.value.length > 0 && 'btn btn-success' || 'b\
17 tn'" ng-click="newKeyword.value.length > 0 && createKeyword(newKeyword)">Create<\
18 /button>
19  </form>
20 </div>
21
22 <div class="gridStyle" ng-grid="keywordsGridOptions"></div>

```

This provides us with the form for creating new keywords, plus it provides the div element where *ngGrid* will render the table. Note how on line 15 the *createKeyword* method of our *KeywordsController* is called upon clicking the *Create* button.

Now we create *src/frontend/app/keywords/partials/operationsGridCell.html*:

```

1 <button class="btn btn-primary" ng-click="updateKeyword(row.entity)">Save</butto\
2 n>
3 <button class="btn btn-warning" ng-click="deleteKeyword(row.entity)">Delete</but\
4 ton>

```

These buttons simply call the according controller methods, passing the keyword entity of the according row as a parameter.

We can now create the HTML template for the grid cells displaying the category name of a keyword, in file *src/frontend/app/keywords/partials/keywordCategoryGridCell.html*:

```
1 <div class="ngCellText" ng-class="col.colIndex()">
2   <span ng-cell-text>
3     {{
4       resolveEntity(
5         {
6           'what' : 'name',
7           'from' : keywordCategories,
8           'where' : 'id',
9           'is'   : row.entity.categoryID
10        }
11      );
12    }}
13 </span>
14 </div>
```

This makes use of the helper function we created earlier.

Last but not least, we need a custom grid cell template for editing the category of a keyword, in file *src/frontend/app/keywords/partials/keywordCategoryGridCellEditor.html*:

```
1 <select
2   class="form-control"
3   ng-model="row.entity.categoryID"
4   ng-options="keywordCategory.id as keywordCategory.name for keywordCategory in \
5 keywordCategories"
6   ng-blur="stopEditingTagHeading()">
7 </select>
```

With this, all components of our frontend app are in place, and we could give it a try in our browser. But not yet, because the backend still needs a lot of changes.

Milestone 5: More work on the backend

Adding a route for retrieving categories

We currently only provide keywords data through our backend API. The app requests keywords data as well as categories data, and the latter would fail for now. Let's implement the required route and logic now.

We first need to extend the spec in *spec/e2e/apiSpec.js*:

```
1  'use strict';
2
3  var request = require('request');
4  var dbSession = require('../../src/backend/dbSession.js');
5  var Server = require('../../src/backend/server.js').Server;
6  var resetDatabase = require('../resetDatabase.js');
7  var async = require('async');
8
9  describe('The API', function () {
10
11     var server;
12
13     beforeEach(function (done) {
14         server = Server('8081');
15         server.listen(function (err) {
16             resetDatabase(dbSession, function() {
17                 done(err);
18             });
19         });
20     });
21
22     afterEach(function (done) {
23         server.close(function () {
24             resetDatabase(dbSession, function() {
25                 done();
26             });
27         });
28     });
29 }
```

```
28 });
29
30
31 it('should respond to a GET request at /api/keywords/', function (done) {
32   var expected = {
33     "_items": [
34       {'id': 1, 'value': 'Aubergine', 'categoryID': 1},
35       {'id': 2, 'value': 'Onion', 'categoryID': 1},
36       {'id': 3, 'value': 'Knife', 'categoryID': 2}
37     ]
38   };
39
40   async.series(
41     [
42
43     function(callback) {
44       resetDatabase(dbSession, callback);
45     },
46
47     function(callback) {
48       dbSession.insert(
49         'keyword',
50         {'value': 'Aubergine', 'categoryID': 1},
51         function(err) { callback(err) });
52     },
53
54     function(callback) {
55       dbSession.insert(
56         'keyword',
57         {'value': 'Onion', 'categoryID': 1},
58         function(err) { callback(err) });
59     },
60
61     function(callback) {
62       dbSession.insert(
63         'keyword',
64         {'value': 'Knife', 'categoryID': 2},
65         function(err) { callback(err) });
66     }
67   ],
68   done);
69
```

```
70     function(err, results) {
71         if (err) throw(err);
72         request.get(
73             {
74                 'url': 'http://localhost:8081/api/keywords/',
75                 'json': true
76             },
77             function (err, res, body) {
78                 expect(res.statusCode).toBe(200);
79                 expect(body).toEqual(expected);
80                 done();
81             }
82         );
83     }
84
85 );
86
87 });
88
89
90 it('should respond to a GET request at /api/keywords/categories/', function (done) {
91     var expected = {
92         "_items": [
93             { 'id': 1, 'name': 'Vegetable' },
94             { 'id': 2, 'name': 'Utility' }
95         ]
96     };
97
98     async.series(
99         [
100             [
101
102                 function(callback) {
103                     resetDatabase(dbSession, callback);
104                 },
105
106                 function(callback) {
107                     dbSession.insert(
108                         'category',
109                         { 'name': 'Vegetable' },
110                         function(err) { callback(err) });
111                 },
```

```
112
113     function(callback) {
114         dbSession.insert(
115             'category',
116             {'name': 'Utility'},
117             function(err) { callback(err) });
118     }
119
120 ],
121
122     function(err, results) {
123         if (err) throw(err);
124         request.get(
125             {
126                 'url': 'http://localhost:8081/api/keywords/categories/',
127                 'json': true
128             },
129             function (err, res, body) {
130                 expect(res.statusCode).toBe(200);
131                 expect(body).toEqual(expected);
132                 done();
133             }
134         );
135     }
136
137 );
138
139 });
140
141 });
```

Additional code starts at line 86. In order to fulfill the new expectation, we need to extend *src/backend/server.js* as follows:

```
1  'use strict';
2
3  var Percolator = require('percolator').Percolator;
4  var dbSession = require('../../src/backend/dbSession.js');
5
6  var Server = function(port) {
7    var server = Percolator({'port': port, 'autoLink': false, 'staticDir': __dirna\
8 me + '/../frontend'});
9
10   server.route('/api/keywords',
11     {
12       GET: function (req, res) {
13         dbSession.fetchAll('SELECT id, value, categoryID FROM keyword ORDER BY id',
14 function(err, rows) {
15           if (err) {
16             console.log(err);
17             res.status.internalServerError(err);
18           } else {
19             res.collection(rows).send();
20           }
21         });
22       }
23     }
24   );
25
26   server.route('/api/keywords/categories',
27     {
28       GET: function (req, res) {
29         dbSession.fetchAll('SELECT id, name FROM category ORDER BY id', function\
30 (err, rows) {
31           if (err) {
32             console.log(err);
33             res.status.internalServerError(err);
34           } else {
35             res.collection(rows).send();
36           }
37         });
38       }
39     }
40   );
41
42   return server;
```

```
43 };
44
45 module.exports = {'Server': Server};
```

One caveat: We also need to extend the *resetDatabase* helper in order to ensure that sqlite primary key sequencing is resetted for each test case:

```
1  'use strict';
2
3  var async = require('async');
4
5  var resetDatabase = function (dbSession, callback) {
6
7    async.series(
8      [
9
10     function (callback) {
11       dbSession.remove('keyword', '1', function (err) {
12         callback(err)
13       });
14     },
15
16     function (callback) {
17       dbSession.remove('category', '1', function (err) {
18         callback(err)
19       });
20     },
21
22     function (callback) {
23       dbSession.remove('sqlite_sequence', '1', function (err) {
24         callback(err)
25       });
26     }
27
28   ],
29
30   function (err, results) {
31     callback(err);
32   }
33 );
34
35 };
```

```
36
37 module.exports = resetDatabase;
```

We can now run

```
$ ./node_modules/.bin/jasmine-node --verbose --captureExceptions ./spec/
```

from the top folder of the project in order to verify that the new route behaves as expected.

Making the backend recognize different environments

But we are not yet done and need to refrain from opening our app just yet.

What's the problem?

Until now, we only used the database when running the specs. Now we have another client for the webservice (and the underlying database), and this is the frontend application. We probably don't want our specs to wipe out the data we add through the application. We can circumvent this problem by using the different environments we created for *db-migrate*. Each environment has its own database setup, and this way we can keep things neatly separated - our specs work with one database setup, and the application with another one.

We need to add a helper script that detects the current environment when executing our code. Place this into *src/backend/env.js*:

```
1  'use strict';
2
3  (function() {
4    var env;
5
6    if (process.env.KW_ENV) {
7      env = process.env.KW_ENV;
8    } else {
9      env = 'test';
10   }
11
12   if (!( env === 'test'
13         || env === 'dev'
14         || env === 'production')) {
15     throw new Error('"' + env + '" is not an allowed environment');
16   }
17
18   module.exports = env;
19 })();
```

This will detect the environment based on a shell environment variable named *KW_ENV*. If no environment parameter can be detected, then the fallback is to assume the *test* environment.

Next, we need to make the database connection logic aware of the environment, and make it use the matching database connection configuration. We therefore need to change *src/backend/dbSession.js* as follows:

```
1  'use strict';
2
3  var env = require('./env');
4  var dbOptions = require('.././database.json')[env];
5  var DBWrapper = require('node-dbi').DBWrapper;
6
7  var dbWrapper;
8  if (dbOptions.driver === 'sqlite3') {
9    var dbWrapper = new DBWrapper('sqlite3', {'path': dbOptions.filename});
10 } else if (dbOptions.driver === 'mysql') {
11   dbWrapper = new DBWrapper('mysql', {
12     'host': dbOptions.host,
13     'user': dbOptions.user,
14     'password': dbOptions.password,
15     'database': dbOptions.database
16   });
17 } else {
18   throw(new Error('No suitable database config found.'));
19 }
20
21 dbWrapper.connect();
22 module.exports = dbWrapper;
```

We also need to make our *resetDatabase.js* helper in folder *spec* aware of the current environment, and make it handle each database type correctly:

```
1  'use strict';
2
3  var async = require('async');
4  var env = require('../src/backend/env');
5  var dbOptions = require('.././database.json')[env];
6
7  var resetDatabase = function (dbSession, callback) {
8
9    if (dbOptions.driver === 'sqlite3') {
10
```



```
11     async.series(  
12         [  
13  
14             function (callback) {  
15                 dbSession.remove('keyword', '1', function (err) {  
16                     callback(err)  
17                 });  
18             },  
19  
20             function (callback) {  
21                 dbSession.remove('category', '1', function (err) {  
22                     callback(err)  
23                 });  
24             },  
25  
26             function (callback) {  
27                 dbSession.remove('sqlite_sequence', '1', function (err) {  
28                     callback(err, null);  
29                 });  
30             }  
31  
32         ],  
33  
34         function (err, results) {  
35             callback(err);  
36         }  
37     );  
38  
39 }  
40  
41 if (dbOptions.driver === 'mysql') {  
42  
43     async.series(  
44         [  
45  
46             function (callback) {  
47                 dbSession.remove('TRUNCATE keyword', [], function (err) {  
48                     callback(err)  
49                 });  
50             },  
51  
52             function (callback) {
```

```
53         dbSession.remove('TRUNCATE category', [], function (err) {
54             callback(err)
55         });
56     }
57
58     ],
59
60     function (err, results) {
61         callback(err);
62     }
63 );
64
65 }
66
67 };
68
69 module.exports = resetDatabase;
```

We are now able to execute the specs as well as the server process under different environments:

```
# Apply migrations in dev environment
$ ./node_modules/.bin/db-migrate up --env dev

# Apply migrations in test environment
$ ./node_modules/.bin/db-migrate up --env test

# Run specs in dev environment
$ KW_ENV=dev ./node_modules/.bin/jasmine-node --verbose --captureExceptions ./spec/

# Run specs in test environment
$ KW_ENV=test ./node_modules/.bin/jasmine-node --verbose --captureExceptions ./spec/

# Start server in dev environment
$ KW_ENV=dev node src/backend/index.js

# Start server in test environment
$ KW_ENV=test node src/backend/index.js
```

This shows how *db-migrate* supports environment detection based on the `-env` parameter, while our specs and our server make use of the `KW_ENV` variable.

We now have the backend and frontend code in place, we are able to run the backend server with a database config different from our specs runs - what's missing is some data to play around with in the dev environment. To fix this, simply run the following on the command line:

```
rm -f /var/tmp/keyword-wrangler.dev.sqlite
./node_modules/.bin/db-migrate up --env dev
echo "DELETE FROM category;" | sqlite3 /var/tmp/keyword-wrangler.dev.sqlite
echo "DELETE FROM keyword;" | sqlite3 /var/tmp/keyword-wrangler.dev.sqlite
echo "DELETE FROM sqlite_sequence;" | sqlite3 /var/tmp/keyword-wrangler.dev.sqlite
echo "INSERT INTO category (name) VALUES ('Vegetable');" | sqlite3 /var/tmp/keyword-wrangler.dev.sqlite
echo "INSERT INTO category (name) VALUES ('Utility');" | sqlite3 /var/tmp/keyword-wrangler.dev.sqlite
echo "INSERT INTO keyword (value, categoryID) VALUES ('Aubergine', 1);" | sqlite3 /var/tmp/keyword-wrangler.dev.sqlite
echo "INSERT INTO keyword (value, categoryID) VALUES ('Onion', 1);" | sqlite3 /var/tmp/keyword-wrangler.dev.sqlite
echo "INSERT INTO keyword (value, categoryID) VALUES ('Knife', 2);" | sqlite3 /var/tmp/keyword-wrangler.dev.sqlite
```

From now on, always start the server as follows, from the root folder of the project:

```
KW_ENV=dev node src/backend/index.js
```

Milestone 6: Completing the backend and finalizing the application

After starting the server in dev mode, have a look at <http://localhost:8080/>. You should now see the frontend with the sample data.

Playing around with this frontend quickly reveals that we are still lacking the API backend operations needed to create, update and delete keyword entries. Let's fix that.

Creating the route for adding new keywords

Creating a new keyword results in a POST request to `/api/keywords`. Let's extend our spec accordingly:



For the sake of brevity, I'm only showing the additional *it* block instead of the whole spec file.

```
1  it('should create a new keyword when receiving a POST request at /api/keywords/'\
2  , function (done) {
3    var expected = {
4      "_items": [
5        {'id': 1, 'value': 'Aubergine', 'categoryID': 1},
6        {'id': 2, 'value': 'Onion', 'categoryID': 1}
7      ]
8    };
9
10   var body = {
11     'value': 'Onion',
12     'categoryID': 1
13   };
14
15   async.series(
16     [
17
18     function(callback) {
19       dbSession.insert(
```

```
20     'category',
21     {'name': 'Vegetable'},
22     function(err) { callback(err) });
23 },
24
25 function(callback) {
26     dbSession.insert(
27         'keyword',
28         {'value': 'Aubergine', 'categoryID': 1},
29         function(err) { callback(err) });
30     }
31
32 ],
33
34 function(err, results) {
35     if (err) throw(err);
36     request.post(
37         {
38             'url': 'http://localhost:8081/api/keywords/',
39             'body': body,
40             'json': true
41         },
42         function (err, res, body) {
43             if (err) throw(err);
44             expect(res.statusCode).toBe(200);
45             request.get(
46                 {
47                     'url': 'http://localhost:8081/api/keywords/',
48                     'json': true
49                 },
50                 function (err, res, body) {
51                     expect(res.statusCode).toBe(200);
52                     expect(body).toEqual(expected);
53                     done();
54                 }
55             );
56         });
57     }
58 );
59 }
60
61 );
```

```
62
63 });
```

As you can see, this spec makes a complete round-trip. It adds the new keyword, and then makes a GET request to verify that the added keyword is contained in the list of all keywords.

Here is our trusty *src/backend/server.js* file updated with the code that is necessary to support the new POST operation:

```
1  'use strict';
2
3  var Percolator = require('percolator').Percolator;
4  var dbSession = require('../src/backend/dbSession.js');
5
6  var Server = function(port) {
7    var server = Percolator({'port': port, 'autoLink': false, 'staticDir': __dirna\
8 me + '/../frontend'});
9
10   server.route('/api/keywords',
11     {
12       GET: function (req, res) {
13         dbSession.fetchAll('SELECT id, value, categoryID FROM keyword ORDER BY i\
14 d', function(err, rows) {
15           if (err) {
16             console.log(err);
17             res.status.internalServerError(err);
18           } else {
19             res.collection(rows).send();
20           }
21         });
22       },
23
24       POST: function (req, res) {
25         req.onJson(function(err, newKeyword) {
26           if (err) {
27             console.log(err);
28             res.status.internalServerError(err);
29           } else {
30             dbSession.query('INSERT INTO keyword (value, categoryID) VALUES (?, \
31 ?);', [newKeyword.value, newKeyword.categoryID], function (err, result) {
32               if (err) {
33                 console.log(err);
34                 res.status.internalServerError(err);
```

```
35         } else {
36             res.object({'status': 'ok', 'id': result.insertId}).send();
37         }
38     });
39     }
40     });
41 }
42 }
43 );
44
45 server.route('/api/keywords/categories',
46 {
47     GET: function (req, res) {
48         dbSession.fetchAll('SELECT id, name FROM category ORDER BY id', function\
49 (err, rows) {
50             if (err) {
51                 console.log(err);
52                 res.status.internalServerError(err);
53             } else {
54                 res.collection(rows).send();
55             }
56         });
57     }
58 }
59 );
60
61 return server;
62 };
63
64 module.exports = {'Server': Server};
```

As you can see on line 23, we make use of Percolators *onJSON* method, which provides us with the POST request body data in a very convenient way.

Creating the route for updating keywords

We also need to be able to save changes to an existing keyword entry. This is another POST request, this time at `/api/keywords/:i/*`. Here is the spec for this:

```
1  it('should update a keyword when receiving a POST request at /api/keywords/:id\  
2  /', function (done) {  
3    var expected = {  
4      "_items": [  
5        {'id': 1, 'value': 'Onion', 'categoryID': 2}  
6      ]  
7    };  
8  
9    var body = {  
10     'id': 1,  
11     'value': 'Onion',  
12     'categoryID': 2  
13   };  
14  
15   async.series(  
16     [  
17  
18     function(callback) {  
19       dbSession.insert(  
20         'category',  
21         {'name': 'Vegetable'},  
22         function(err) { callback(err) });  
23     },  
24  
25     function(callback) {  
26       dbSession.insert(  
27         'category',  
28         {'name': 'Utility'},  
29         function(err) { callback(err) });  
30     },  
31  
32     function(callback) {  
33       dbSession.insert(  
34         'keyword',  
35         {'value': 'Aubergine', 'categoryID': 1},  
36         function(err) { callback(err) });  
37     }  
38  
39   ],  
40  
41   function(err, results) {  
42     if (err) throw(err);
```



```
43     request.post(  
44         {  
45             'url': 'http://localhost:8081/api/keywords/1',  
46             'body': body,  
47             'json': true  
48         },  
49         function (err, res, body) {  
50             if (err) throw(err);  
51             expect(res.statusCode).toBe(200);  
52             request.get(  
53                 {  
54                     'url': 'http://localhost:8081/api/keywords/',  
55                     'json': true  
56                 },  
57                 function (err, res, body) {  
58                     expect(res.statusCode).toBe(200);  
59                     expect(body).toEqual(expected);  
60                     done();  
61                 }  
62             );  
63         }  
64     );  
65 }  
66 }  
67  
68 );  
69  
70 });
```

And here is the updated server code:

```
1  'use strict';  
2  
3  var Percolator = require('percolator').Percolator;  
4  var dbSession = require('../src/backend/dbSession.js');  
5  
6  var Server = function(port) {  
7      var server = Percolator({'port': port, 'autoLink': false, 'staticDir': __dirna\  
8  me + '/../frontend'});  
9  
10     server.route('/api/keywords',  
11         {
```

```
12     GET: function (req, res) {
13         dbSession.fetchAll('SELECT id, value, categoryID FROM keyword ORDER BY id',
14 d', function(err, rows) {
15             if (err) {
16                 console.log(err);
17                 res.status.internalServerError(err);
18             } else {
19                 res.collection(rows).send();
20             }
21         });
22     },
23
24     POST: function (req, res) {
25         req.onJson(function(err, newKeyword) {
26             if (err) {
27                 console.log(err);
28                 res.status.internalServerError(err);
29             } else {
30                 dbSession.query('INSERT INTO keyword (value, categoryID) VALUES (?, \
31 ?);', [newKeyword.value, newKeyword.categoryID], function (err, result) {
32                     if (err) {
33                         console.log(err);
34                         res.status.internalServerError(err);
35                     } else {
36                         res.object({'status': 'ok', 'id': result.insertId}).send();
37                     }
38                 });
39             }
40         });
41     }
42 }
43 );
44
45 server.route('/api/keywords/categories',
46 {
47     GET: function (req, res) {
48         dbSession.fetchAll('SELECT id, name FROM category ORDER BY id', function\
49 (err, rows) {
50             if (err) {
51                 console.log(err);
52                 res.status.internalServerError(err);
53             } else {
```

```
54         res.collection(rows).send();
55     }
56     });
57 }
58 }
59 );
60
61 server.route('/api/keywords/:id',
62 {
63     POST: function(req, res) {
64         var keywordId = req.uri.child();
65         req.onJson(function(err, keyword) {
66             if (err) {
67                 console.log(err);
68                 res.status.internalServerError(err);
69             } else {
70                 dbSession.query('UPDATE keyword SET value = ?, categoryID = ? WHERE \
71 keyword.id = ?;', [keyword.value, keyword.categoryID, keywordId], function (err, \
72 result) {
73                 if (err) {
74                     console.log(err);
75                     res.status.internalServerError(err);
76                 } else {
77                     res.object({'status': 'ok'}).send();
78                 }
79             });
80         }
81     });
82 }
83 }
84 );
85
86 return server;
87 };
88
89 module.exports = {'Server': Server};
```



Note how we handle the `/api/keywords/categories` route first, and then the `/api/keywords/:id` route. This is important, because if the `/api/keywords/:id` route would be defined first, it would handle `/api/keywords/categories` requests, interpreting `categories` as the `:id`.

Creating the route for deleting keywords

Now it's time for the final step. We need to add route handling for DELETE requests, and our application is complete.

Once again, we start with the spec:

```
1  it('should remove a keyword when receiving a DELETE request at /api/keywords/:\
2  id/', function (done) {
3    var expected = {
4      "_items": [
5        {'id': 1, 'value': 'Aubergine', 'categoryID': 1}
6      ]
7    };
8
9    async.series(
10   [
11
12     function(callback) {
13       dbSession.insert(
14         'category',
15         {'name': 'Vegetable'},
16         function(err) { callback(err) });
17     },
18
19     function(callback) {
20       dbSession.insert(
21         'keyword',
22         {'value': 'Aubergine', 'categoryID': 1},
23         function(err) { callback(err) });
24     },
25
26     function(callback) {
27       dbSession.insert(
28         'keyword',
29         {'value': 'Onion', 'categoryID': 1},
30         function(err) { callback(err) });
31     }
32   ],
33
34   function(err, results) {
35     if (err) throw(err);
36   }
```

```

37     request.del(
38       {
39         'url': 'http://localhost:8081/api/keywords/2/',
40         'json': true
41       },
42       function (err, res, body) {
43         if (err) throw(err);
44         request.get(
45           {
46             'url': 'http://localhost:8081/api/keywords/',
47             'json': true
48           },
49           function (err, res, body) {
50             expect(res.statusCode).toBe(200);
51             expect(body).toEqual(expected);
52             done();
53           }
54         );
55       }
56     );
57   }
58 }
59
60 );
61
62 });

```

Then, we extend the server code:

```

1  'use strict';
2
3  var Percolator = require('percolator').Percolator;
4  var dbSession = require('../src/backend/dbSession.js');
5
6  var Server = function(port) {
7    var server = Percolator({'port': port, 'autoLink': false, 'staticDir': __dirna\
8    me + '/../frontend'});
9
10   server.route('/api/keywords',
11     {
12       GET: function (req, res) {
13         dbSession.fetchAll('SELECT id, value, categoryID FROM keyword ORDER BY i\

```

```
14 d', function(err, rows) {
15     if (err) {
16         console.log(err);
17         res.status.internalServerError(err);
18     } else {
19         res.collection(rows).send();
20     }
21 });
22 },
23
24 POST: function (req, res) {
25     req.onJson(function(err, newKeyword) {
26         if (err) {
27             console.log(err);
28             res.status.internalServerError(err);
29         } else {
30             dbSession.query('INSERT INTO keyword (value, categoryID) VALUES (?, \
31 ?);', [newKeyword.value, newKeyword.categoryID], function (err, result) {
32                 if (err) {
33                     console.log(err);
34                     res.status.internalServerError(err);
35                 } else {
36                     res.object({'status': 'ok', 'id': result.insertId}).send();
37                 }
38             });
39         }
40     });
41 }
42 }
43 );
44
45 server.route('/api/keywords/categories',
46 {
47     GET: function (req, res) {
48         dbSession.fetchAll('SELECT id, name FROM category ORDER BY id', function\
49 (err, rows) {
50             if (err) {
51                 console.log(err);
52                 res.status.internalServerError(err);
53             } else {
54                 res.collection(rows).send();
55             }
56         }
57     }
58 }
```

```
56     });
57   }
58 }
59 );
60
61 server.route('/api/keywords/:id',
62   {
63     POST: function(req, res) {
64       var keywordId = req.uri.child();
65       req.onJson(function(err, keyword) {
66         if (err) {
67           console.log(err);
68           res.status.internalServerError(err);
69         } else {
70           dbSession.query('UPDATE keyword SET value = ?, categoryID = ? WHERE \
71 keyword.id = ?;', [keyword.value, keyword.categoryID, keywordId], function (err, \
72 result) {
73           if (err) {
74             console.log(err);
75             res.status.internalServerError(err);
76           } else {
77             res.object({'status': 'ok'}).send();
78           }
79         });
80       }
81     });
82   },
83
84   DELETE: function(req, res) {
85     var keywordId = req.uri.child();
86     dbSession.query('DELETE FROM keyword WHERE keyword.id = ?;', [keywordId] \
87 , function(err, result) {
88       if (err) {
89         console.log(err);
90         res.status.internalServerError(err);
91       } else {
92         res.object({'status': 'ok'}).send();
93       }
94     });
95   }
96 }
97 );
```

```
98
99   return server;
100 };
101
102 module.exports = {'Server': Server};
```

We can now restart the backend server (remember, *KW_ENV=dev!*), and with this, we should be able to fully enjoy all the operations our frontend provides.

Summary

In this chapter, we created a complete web-based CRUD application. It consists of a Node.js backend and an AngularJS frontend, connected through a RESTful webservice API. We learned how to create a continuous delivery environment where end-to-end specs provide a tests suite that verifies the correctness of the webservice API and a migrations setup ensures that database schema and code are in sync. The application can be run in the context of different configuration environments, allowing to keep test, dev and production data separated.

Epilogue

You have reached the end of *The Node Craftsman Book*. I sincerely hope you enjoyed the journey and feel confident to use Node.js in your own projects small and large.

If you have any feedback or questions, please feel free to use the community forum at <https://leanpub.com/nodecraftsman/feedback>, or simply drop me a line at manuel@kiessling.net.