

API Foundations in Go

by Tit Petric



API foundations in Go

You've mastered PHP and are looking at Node.js? Skip it and try Go.

Tit Petric

This book is for sale at <http://leanpub.com/api-foundations>

This version was published on 2019-01-15



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.

© 2016 - 2019 Tit Petric

Tweet This Book!

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

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

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

[#apifoundations](#)

I'm dedicating this book to my wife, Anastasia. Even if you say that you don't suffer, I'm sure that I'm killing your inner child by explaining what Docker is. I'm sorry for that.

Contents

Introduction	1
About me	1
Why Go?	1
Who is this book for?	1
How should I study it?	2
Setting up your environment	3
Networking	3
Setting up a runner for your Go programs	3
Setting up Redis	4
Other services	5
Data structures	6
Declaring structs	6
Casting structs	7
Declaring interfaces	7
Abusing interfaces	8
Embedding and composition	9
Limiting goroutine parallelization	11
Slices	12
The slice operator	13
Allocation by append	14
Copying slices	15
Using slices in channels	16

Introduction

About me

I'm one of those people with about two decades of programming experience under my belt. I started optimizing code in the '90s, discovered PHP in the 2000s, and have built several large-scale projects to date, all while discovering other programming language families like Java, Node.js and ultimately, Go.

I have built numerous APIs for my own content management products. Several products I've been involved with as a lead developer have been sold and are used in multiple countries. I've written a professional, dedicated API framework, which doubles as a software development kit for the Slovenian national TV and Radio station website, RTV Slovenia. I've also been a speaker at several local PHP user group events and conferences.

Why Go?

Go has been on my list for a while now. About a year or two ago, my coworker created a protocol converter that emulates a Memcache server, using Redis as the backend datastore. Since our biggest project has a code base built up over some 13 years, it was preferable to keep the working code as-is, and just replace the software around it as needed. Go made this possible.

One of the reasons for choosing Go was the constant comparison of Go with Node. Node has a much more vocal community, which seems to religiously advocate it. We have carried out several tests in recent months, and Node, while not very hard to start development with, had poorer performance than pretty much everything except PHP. I'm not saying that Go is better than Node, or that anything is better than anything else, but from what we've seen, it seems advisable to skip Node.js and go straight to Go. This might change as ES6 and ES7 get more traction - but there are immediate benefits of switching to Go. If you don't want to move from Node.js, this book is not for you. If you have an open mind - find out what Go can do.

Who is this book for?

This book is for senior developers who might not have had a chance to try Go, but are familiar with concepts of API development in languages like PHP or Node. Any reader must have a good understanding of REST APIs and server architecture. While I'll try to make everything as clear as possible, realize that if you're a novice programmer, there might be a gap between what you know, and what I'm trying to explain here.

I'm not going to be explaining the Go programming language in detail. I'm going to be diving head first into using it with just a note here and there. This is why familiarity and strong knowledge of programming concepts are required.

In the book, I will cover these subjects:

1. [Setting up your environment](#)
2. [Data structures](#)
3. [Organizing your code](#)
4. [Encoding and decoding JSON](#)
5. [Serving HTTP requests](#)
6. [Parallel fetching of data](#)
7. [Using external services \(Redis\)](#)
8. [Using external services \(MySQL\)](#)
9. [Test driven API development](#)
10. [Your first API](#)
11. [Running your API in production](#)
12. [Resources](#)

Covering these concepts should give you a strong foundation for your API implementation. The book doesn't try to teach you Go; the book tries to give you a strong software foundation for APIs, using Go.

How should I study it?

Through the book, I will present several examples on how to do common things when developing APIs. The examples are published on GitHub, you can find the link in the last chapter of the book.

You should follow the examples in the book, or you can look at each chapter individually, just to cover the knowledge of that chapter. The examples are stand-alone, but generally build on work from previous chapters.

Setting up your environment

Setting up a development environment, as well as a production environment, is an important topic today. While spinning up a virtual machine and installing software by hand is perhaps the accepted way of doing things, recently I've learned to systematize my development by using Docker containers.

The biggest benefit of Docker containers is a “zero install” way of running software - as you're about to run a container, it downloads the image containing all the software dependencies you need. You can take this image and copy it to your production server, where you can run it without any change in environment.

Networking

When you have docker set up, we will need to create a network so the services that we're going to use can talk to each other. Creating a network is simple, all you need to do is run the following command:

```
$ docker network create -d bridge --subnet 172.25.0.0/24 party
```

This command will create a network named **party** on the specified subnet. All docker containers which will run on this network will have connectivity to each other. That means that when we run our Go container, it will be able to connect to another Redis container on the same network.

Setting up a runner for your Go programs

There is an official Go image available on Docker. Getting a Docker image and running it is very simple, requiring just one line:

```
$ docker run --net=party -p 8080:80 --rm=true -it -v `pwd`: /go/src/app -w /go/src/ap\
p golang go "$@"
```

Save this code snippet as the file 'go', make it executable and copy it to your execution path (usually /usr/local/bin is reserved for things like this).

Let's quickly go over the arguments provided:

- **-net=party** - runs the container on the shared network

- **-p** - network forwarding from host:8080 to container:80 (HTTP server)
- **-rm=true** - when the container stops, clean up after it (saves disk space)
- **-v option** - creates a volume from the current path (pwd) in the container
- **"\$@"** - passes all arguments to go to the container application

Very simply, what this command does is run a Docker container in your current folder, execute the go binary in the container, and clean up after itself when the program finishes.

Note: as we only expose the current working path to the container, it limits access to the host machine - if you have code outside the current path, for example in “..” or “/usr/share”, this code will not be available to the container.

An example of running go would be:

```
$ go version
go version go1.6 linux/amd64
```

And, what we will do through most of this book is run the following command:

- **go run [file]** - compile and run Go program

```
$ go run hello_world.go
Hello world!
```

A minimal example of a Go program

```
1 package main
2
3 import "fmt"
4
5 func main() {
6     fmt.Printf("Hello world!\n")
7 }
```

Setting up Redis

We will also use Docker to run Redis, which is as simple as setting up Go. Redis is an in-memory data structure store. It provides functionality to store and retrieve data in various structures, beyond a simple key-value database like Memcache. We will use this service later in the book when implementing our example API endpoints.

To run an instance of Redis named ‘redis’:

```
$ docker run --restart=always -h redis --name redis --net=party -d redis
```

Just like Go, Redis provides an official build on the Docker Hub. Interacting with Redis will be covered in detail in later chapters.

Other services

In addition to Go and Redis, other services are also available on the [Docker Hub](https://hub.docker.com/)¹. Depending on your use case, you might want to install additional services via docker.

Popular projects which I use from Docker on a daily basis:

- nginx (and nginx-extras flavours)
- Percona (MySQL, MariaDB) - also covered later in the book
- letsencrypt
- redis
- samba
- php

Docker is a very powerful tool which gives you all the software you might need. It's very useful also for testing, as you can use it to set up a database, populate it with test data, and then tear down and clean up after it. It's a very convenient way to run programs that are isolated from your actual environment, and may only be active temporary.

¹<https://hub.docker.com/>

Data structures

Defining and handling data structures is a key activity in any programming language. When talking about object oriented programming, it's worth noting some differences between Go and other programming languages.

- Go doesn't have classes, it has structs
- Methods are bound to a struct, not declared within one
- The "interface" type can be an intersection of many or all types
- Packages behave like namespaces, but everything in a package is available

In short, it means that the functions you'll define for your structs will be declared outside of the struct they are working on. If you want to declare several types, the only way to assign them to a same variable, without many complications, is to declare a common interface.

Declaring structs

When we define a structure in Go, we set types for every member. Let's say we would like to define the usual "Petstore", which in turn has a list of pets. You could define the structure like this:

```
1 type Petstore struct {
2     Name string
3     Location string
4     Dogs []*Pet
5     Cats []*Pet
6 }
7
8 type Pet struct {
9     Name string
10    Breed string
11 }
```

The example is simple, in the sense that I'm not defining a "Pet", which can be a "Dog", or can be a "Cat". I'm just using one type for all.

Note: Members of defined structures begin with an upper case letter. This means that they are visible outside the package they are defined in. All members, which you will encode to JSON in the next chapter, need to be public, that is - start with an upper case letter. The same applies to functions that you expose in packages.

Casting structs

We could declare a Dog and Cat struct, and could cast one to one from the other. A requirement for this is that the Dog and Cat types *must have identical underlying types*. So, if I was to extend the above example, I would copy the Pet struct into Dog and Cat so they are identical.

If they are not identical, the value can't be cast.

Error when casting

```
1 type Dog struct {  
2     name string  
3 }  
4 type Cat struct {  
5     name          string  
6     hypoallergenic bool  
7 }  
8  
9 func main() {  
10     dog := Dog{name: "Rex"}  
11     cat := Cat(dog)  
12 }
```

The above example results in:

```
./type1.go:12: cannot convert dog (type Dog) to type Cat
```

Even if “Cat” has all the properties of “Dog”, type conversion is not possible. It makes no sense to declare strong types of individual animals, if casting to a generic type is not possible. But we can assign these kind of strong types to a common interface type.

Declaring interfaces

An interface defines zero or more methods that need to be implemented on types/objects that can be assigned to that interface. If we extend the above example to declare explicit types for Cat and Dog, we will end up with something like this:

Structure for an agnostic pet list

```
1 type Dog struct {
2     name string
3     breed string
4 }
5 type Cat struct {
6     name string
7     hypoallergenic bool
8 }
9 type Petstore struct {
10     name string
11     pets []interface{}
12 }
```

We are using the declaration of `pets` as “`interface{}`”. The interface doesn’t define any common methods, so anything can be assigned to it. Usually, one would implement getter or setter methods which should be common to all the types an interface can hold. With our example, we could have declared a function `GetName` that would return the name of the `Pet`:

Structure for an agnostic pet list

```
1 type Pet interface {
2     getName() string
3 }
4
5 func (d Dog) getName() string {
6     return d.name
7 }
8 func (c Cat) getName() string {
9     return c.name
10 }
```

Abusing interfaces

Interfaces are a powerful and flexible way to work with objects of varying types. You should however tend to prefer static type declarations when possible - it makes the code faster and more readable. The Go compiler catches a lot of your common errors with static analysis. If you’re using interfaces, you’re exposing yourself to risks and errors which the compiler cannot catch.

All the variables which have been declared in the object are unreachable from the interface and will result in an error. But that doesn’t mean that you can’t access them.

If you absolutely must, you can use a feature called “reflection”. Reflection exposes the members of any interface via an API.

“Hello world!” example using Reflection

```
1 package main
2
3 import "fmt"
4 import "reflect"
5
6 type Test1 struct {
7     A string
8 }
9
10 func main() {
11     var t interface{}
12     t = Test1{"Hello world!"}
13
14     data := reflect.ValueOf(t)
15     fmt.Printf("%s\n", data.FieldByName("A"))
16 }
```

Interfaces can be a powerful part of Go, but care should be taken not to overuse them. They make many tasks simpler and can be a powerful tool when they are used correctly.

Reflection is used by many packages to provide generic function interfaces, like the `encoding/json` package does, which we will use in a later chapter.

Embedding and composition

You can embed types into a struct by listing them in the struct declaration. This makes all the values and functions of the embedded struct available for use on your struct. For example, the `sync.Mutex` struct implements `Lock` and `Unlock` functions.

```
1 type Services struct {
2     sync.Mutex
3     ServiceList map[string]Service
4 }
5
6 func (r *Services) Add(name string, service Service) {
7     r.Lock()
8     r.ServiceList[name] = service
9     r.Unlock()
10 }
```

This pattern enables composition. You can build your APIs for a larger program from smaller structures that can deal with specific problem domains. The example also shows a beneficial side-effect: you can reason that the embedded `sync.Mutex` is used to lock your structure fields below the embed.

You can use embedding to your advantage with, for example, `sync.Mutex`, `sync.RWMutex` or `sync.WaitGroup`. You can actually embed *many* structs, so your structure may perform the functions of them all.

An example from a project I'm working on uses two embedded structs:

```
1 type RunQueue struct {
2     sync.RWMutex
3     sync.WaitGroup
4     // ...
5     flagIsDone bool
6 }
7 func (r *RunQueue) Close() {
8     r.Lock()
9     defer r.Unlock()
10    r.flagIsDone = true
11 }
12 func (r *RunQueue) IsDone() bool {
13     r.RLock()
14     defer r.RUnlock()
15     return r.flagIsDone
16 }
```

Leveraging `sync.RWMutex`

The declaration of the `RunQueue` struct above leverages the `sync.RWMutex` to provide synchronous access to the object from many goroutines. A goroutine may use `Close` to finish the execution of the goroutine queue. Each worker in the queue would call `IsDone` to check if the queue is still active.

Leveraging `sync.WaitGroup`

The `RunQueue` struct leverages a `sync.WaitGroup` to provide queue clean up and statistics, such as elapsed time. While I can't provide all the code, the basic usage is as follows:

```

1 func (r *RunQueue) Runner() {
2     fmt.Printf("Starting %d runners\n", runtime.NumCPU())
3     for idx := 1; idx <= runtime.NumCPU(); idx++ {
4         go r.runOnce(idx)
5     }
6 }
7 func NewRunQueue(jobs []Command) RunQueue {
8     q := RunQueue{}
9     for idx, job := range jobs {
10         if job.SelfId == 0 {
11             q.Dispatch(&jobs[idx])
12         }
13     }
14     q.Add(len(q.jobs)) // sync.WaitGroup
15     return q
16 }
17 runnerQueue := NewRunQueue(commands)
18 go runnerQueue.Finisher()
19 go runnerQueue.Runner()
20 runnerQueue.Wait() // sync.WaitGroup

```

The main idea of the program I'm building is that it starts `runtime.NumCPU()` runners, which handle execution of a fixed number of commands. The `WaitGroup` comes into play very simply:

```

1 NewRunQueue calls *wg.Add(number of jobs)
2 Individual jobs are processed with RunQueue.runOnce, and they call *wg.Done()
3 RunnerQueue.Wait() (*wg.Wait()) will wait until all jobs have been processed

```

Limiting goroutine parallelization

At one point I struggled to create a queue manager, which would parallelize workloads to a fixed limit of parallel goroutines. My idea was to register a slot manager, which would provide a pool of running tasks. If no pool slot is available, it'd sleep for a few seconds before trying to get a slot again. It was frustrating.

Just look at the loop from the `Runner` function above:


```
1 for idx := 1; idx <= runtime.NumCPU(); idx++ {  
2     go r.runOnce(idx)  
3 }
```

This is an elegant way to limit parallelization to N routines. There is no need to bother yourself with some kind of routine allocation pool structs. The `runOnce` function should only do a few things:

1. Listen for new jobs in an infinite loop, reading jobs from a channel
2. Perform the job without new goroutines

The reason to read the jobs from a channel is that the read from a channel is a blocking operation. The function will just wait there until a new job appears on the channel it's reading from.

```
1 func (r *RunQueue) runOnce(idx int) {  
2     for {  
3         queueJob, ok := <-r.runQueue  
4         if !ok {  
5             return  
6         }  
7         // run tasks  
8     [...]
```

The job needs to be executed without a goroutine, or with nested `*WaitGroup.Wait()` call. The reason for this should be obvious - as soon as you start a new goroutine, it gets executed in parallel and the `runOnce` function reads the next job from the queue. This means that the limit on the number of tasks running in parallel would not be enforced.

Slices

A Slice is a Go-specific data type, which consists of:

1. a pointer to an array of values,
2. the capacity of the array,
3. the length of the array

If you're working with slices, which you mostly are, you'll inevitably find yourself in a situation where you'll have to merge two slices into one. The naïve way of doing this is something like the following:

```
1 alice := []string{"foo", "bar"}
2 bob := []string{"verdana", "tahoma", "arial"}
3 for _, val := range bob {
4     alice = append(alice, val)
5 }
```

A slightly more idiomatic way of doing this is:

```
1 alice = append(alice, bob...)
```

The ellipsis or variadic argument expands `bob` in this case to all its members, effectively achieving the same result, without that loop. You can read more about it in the official documentation: [Appending to and copying slices²](https://golang.org/ref/spec#Appending_and_copying_slices)

You could use it to wrap logic around `log.Printf` that, for example, trims and appends a newline at the end of the format string.

```
1 import "log"
2 import "strings"
3
4 func Log(format string, v ...interface{}) {
5     format = strings.TrimSuffix(format, "\n") + "\n"
6     log.Printf(format, v...)
7 }
```

Since slices are basically a special type of pointer, this means there are a number of scenarios where you may inadvertently modify a slice without intending to. Thus, it's important to know how to manipulate slices.

The slice operator

The first thing one should be aware of is that the slice operator, *does not copy the data to a newly-created slice*. Not being aware of this fact can lead to unexpected results:

```
1 a := []string{"r", "u", "n"}
2 b := a[1:2]
3 b[0] = "a"
4 fmt.Println(a)
```

²https://golang.org/ref/spec#Appending_and_copying_slices

Does it print `[r u n]` or `[r a n]`? Since the slice `b` is not a copy of the slice, but just a slice with a modified pointer to the array, the above will print `[r a n]`.

As explained, this is because the slice operator just provides a new slice with an updated reference to the same array as used by the original slice. From the official reference:

Slicing does not copy the slice's data. It creates a new slice value that points to the original array. This makes slice operations as efficient as manipulating array indices. Therefore, modifying the elements (not the slice itself) of a re-slice modifies the elements of the original slice.

Source: [Go Blog - Slices usage and internals](https://blog.golang.org/go-slices-usage-and-internals)³.

Allocation by append

Appending to a slice is simple enough. As mentioned, the slice has a length which you can get with a call to `len()`, and a capacity which you can get with a call to `cap()`.

```
1 a := []string{"r", "u", "n"}
2 fmt.Println(a, len(a), cap(a))
3 a = append(a, []string{"e"}...)
4 fmt.Println(a, len(a), cap(a))
5 a = a[1:len(a)]
6 fmt.Println(a, len(a), cap(a))
```

The expected output would be:

```
1 [r u n] 3 3
2 [r u n e] 4 4
3 [u n e] 3 3
```

Of course, that's not how `append` works. `Append` will double the existing capacity. This means you'll end up with output like this:

```
1 [r u n] 3 3
2 [r u n e] 4 6
3 [u n e] 3 5
```

If you wanted the previous result, you would have to create your own function, which would allocate only the required amount of items into a new slice, and then copy over the data from the source slice. This function might look something like this:

³<https://blog.golang.org/go-slices-usage-and-internals>

```
1 func suffix(source []string, vars ...string) []string {
2     length := len(source) + len(vars)
3     ret := make([]string, length, length)
4     copy(ret, source)
5     index := len(source)
6     for k, v := range vars {
7         ret[index+k] = v
8     }
9     return ret
10 }
11
12 func main() {
13     a := []string{"r", "u", "n"}
14     fmt.Println(a, len(a), cap(a))
15     a = suffix(a, []string{"e"}...)
16     fmt.Println(a, len(a), cap(a))
17     a = a[1:len(a)]
18     fmt.Println(a, len(a), cap(a))
19 }
```

Copying slices

As explained above, “slicing does not copy the slice’s data”. This may sometimes have unintended consequences. This applies not only to slicing, but also to passing slices into functions.

```
1 func flip(source []string) {
2     source[1] = "a"
3 }
4
5 func main() {
6     a := []string{"r", "u", "n"}
7     flip(a)
8     fmt.Println(a)
9 }
```

The above example will print `[r a n]`. This is unfortunate, because people intuitively expect slices to behave much like structs do. Passing a struct will create a copy. A slice will still point at the same memory that holds the data. Even if you pass a slice within a struct, you should be aware of this:

```
1 type X struct {
2     c string
3     source []string
4 }
5
6 func flip(x X) {
7     x.c = "b"
8     x.source[1] = "a"
9 }
10
11 func main() {
12     a := X{"a", []string{"r", "u", "n"}}
13     flip(a)
14     fmt.Println(a)
15 }
```

The above will print out a {a [r a n]}. Slices always behave as if they are passed by reference. In a way they are, as one of the parts of the slice is a pointer to the array of data it holds.

Using slices in channels

If you're using buffered channels, and are trying to be smart with slices, reusing a slice to keep allocations down, then you might find that this buffered channel is basically filled with the same slice over and over. As the slice contains the pointer to the array holding the data, the following example will have unexpected results:

```
1 func main() {
2     c := make(chan []string, 5)
3
4     go func() {
5         item := []string{"hello"}
6         for i := 0; i < 5; i++ {
7             item[0] = fmt.Sprintf("hello %d", i)
8             c <- item
9             //time.Sleep(100 * time.Millisecond)
10        }
11    }()
12
13    for i := 0; i < 5; i++ {
14        item := <-c
15        fmt.Println(item)
```

```
16     }  
17 }
```

The output is:

```
1  [hello 4]  
2  [hello 4]  
3  [hello 4]  
4  [hello 4]  
5  [hello 4]
```

If you uncomment the `time.Sleep` command in the goroutine, you will most likely get the correct result:

```
1  [hello 0]  
2  [hello 1]  
3  [hello 2]  
4  [hello 3]  
5  [hello 4]
```

At any time when the consumer is reading from the channel, the retrieved slice will be identical, because only the item in the underlying array is changing. The only solution for this, I believe, is either finding a way to work with an unbuffered channel (a consumer must *always* be waiting to read from the channel), or to explicitly copy the slice which is being put into the channel, like in [this playground example](https://play.golang.org/p/ORPyzf0zwY)⁴.

⁴<https://play.golang.org/p/ORPyzf0zwY>