

Semantic Import Versioning

Go & Versioning, Part 3

Russ Cox

February 21, 2018

research.swtch.com/vgo-import

How do you deploy an incompatible change to an existing package? This is the fundamental challenge, the fundamental decision, in any package management system. The answer decides the complexity of the resulting system. It decides how easy or difficult package management will be to use. (It also decides how easy or difficult package management will be to implement, but the user experience is more important.)

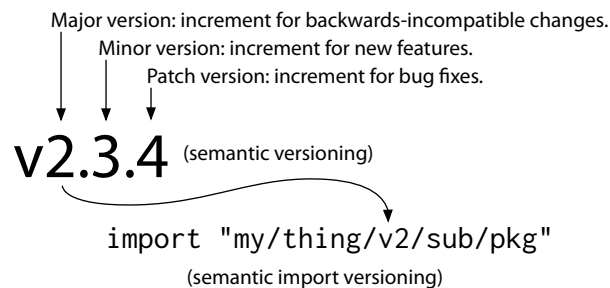
To answer this question, this post first presents the *import compatibility rule* for Go:

If an old package and a new package have the same import path, the new package must be backwards compatible with the old package.

We've argued for this principle from the start of Go, but we haven't given it a name or such a direct statement.

The import compatibility rule dramatically simplifies the experience of using incompatible versions of a package. When each different version has a different import path, there is no ambiguity about the intended semantics of a given import statement. This makes it easier for both developers and tools to understand Go programs.

Developers today expect to use semantic versions to describe packages, so we adopt them into the model. Specifically, a module `my/thing` is imported as `my/thing` for `v0`, the incompatibility period when breaking changes are expected and not protected against, and then also during `v1`, the first stable major version. But when it's time to add `v2`, instead of redefining the meaning of the now-stable `my/thing`, we give it a new name: `my/thing/v2`.



I call this convention *semantic import versioning*, the result of following the import compatibility rule while using semantic versioning.

A year ago, I believed that putting versions in import paths like this was ugly, undesirable, and probably avoidable. But over the past year, I've come to understand just how much clarity and simplicity they bring to the system. In this post I hope to give you a sense of why I changed my mind.

A Dependency Story

To make the discussion concrete, consider the following story. The story is hypothetical, of course, but it's motivated by a real problem. When `dep` was released, the team at Google that wrote the `OAuth2` package asked me how they should go about introducing some incompatible improvements they've wanted to do for a long time. The more I thought about it, the more I realized that this

was not as easy as it sounded, at least not without semantic import versioning.

Prologue

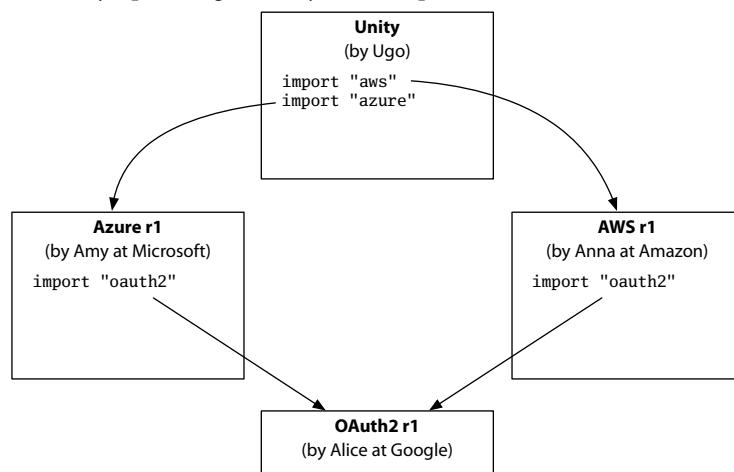
From the perspective of a package management tool, there are Authors of code and Users of code. Alice, Anna, and Amy are Authors of different code packages. Alice works at Google and wrote the OAuth2 package. Amy works at Microsoft and wrote the Azure client libraries. Anna works at Amazon and wrote the AWS client libraries. Ugo is the User of all these packages. He's working on the ultimate cloud app, Unity, which uses all of those packages and others.

As the Authors, Alice, Anna, and Amy need to be able to write and release new versions of their packages. Each version of a package specifies a required version for each of its dependencies.

As the User, Ugo needs to be able to build Unity with these other packages; he needs control over exactly which versions are used in a particular build; and he needs to be able to update to new versions when he chooses.

There's more that our friends might expect from a package management tool, especially around discovery, testing, portability, and helpful diagnostics, of course, but those are not relevant to the story.

As our story opens, Ugo's Unity build dependencies look like:



Chapter 1

Everyone is writing software independently.


At Google, Alice has been busy designing a new, simpler, easier-to-use API for the OAuth2 package. It can still do everything that the old package can do, but with half the API surface. She releases it as OAuth2 r2. (The 'r' here stands for revision. For now, the revision numbers don't indicate anything other than sequencing: in particular, they're not semantic versions.)

At Microsoft, Amy is on a well-deserved long vacation, and her team decides not to make any changes related to OAuth2 r2 until she returns. The Azure package will keep using OAuth2 r1 for now.

At Amazon, Anna finds that using OAuth2 r2 will let her delete a lot of ugly code from the implementation of AWS r1, so she changes AWS to use OAuth2 r2. She fixes a few bugs along the way and issues the result as AWS r2.

Ugo gets a bug report about behavior on Azure and tracks it down to a bug in the Azure client libraries. Amy already released a fix for that bug in Azure r2 before leaving for vacation. Ugo adds a test case to Unity, confirms that it fails, and asks the package management tool to update to Azure r2.

After the update, Ugo's build looks like:



He confirms that the new test passes and that all his old tests still pass. He locks in the Azure update and ships an updated Unity.


Chapter 2

To much fanfare, Amazon launches their new cloud offering, Amazon Zeta Functions. In preparation for the launch, Anna added Zeta support to the AWS package, which she now releases as AWS r3.


When Ugo hears about Amazon Zeta, he writes some test programs and is so excited about how well they work that he skips lunch to update Unity. Today's update does not go as well as the last one. Ugo wants to build Unity with Zeta support using Azure r2 and AWS r3, the latest version of each. But Azure r2 needs OAuth2 r1 (not r2), while AWS r3 needs OAuth2 r2 (not r1). Classic diamond dependency, right? Ugo doesn't care what it is. He just wants to build Unity.

Worse, it doesn't appear to be anyone's fault. Alice wrote a better OAuth2 package. Amy fixed some Azure bugs and went on vacation. Anna decided AWS should use the new OAuth2 (an internal implementation detail) and later added Zeta support. Ugo wants Unity to use the latest Azure and AWS packages. It's very hard to say any of them did something wrong. If these people aren't wrong, then maybe the package manager is. We've been assuming that there can be only one version of OAuth2 in Ugo's Unity build. Maybe that's the problem: maybe the package manager should allow different versions to be included in a single build. This example would seem to indicate that it must.

Ugo is still stuck, so he searches StackOverflow and finds out about the package manager's `-multiverse` flag, which allows multiple versions, so that his program builds as:

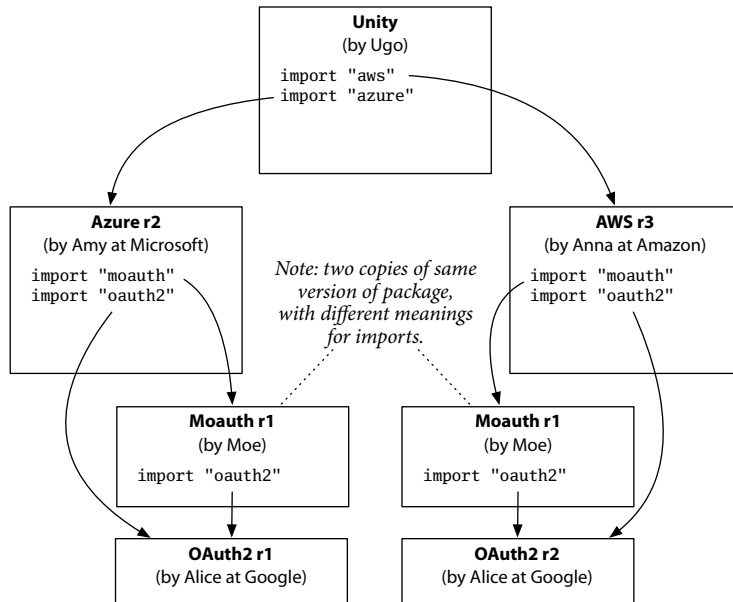


Ugo tries this. It doesn't work. Digging further into the problem, Ugo discovers that both Azure and AWS are using a popular OAuth2 middleware library called Moauth that simplifies part of the OAuth2 processing. Moauth is not a complete API replacement: users still import OAuth2 directly, but they use Moauth to simplify some of the API calls. The details that Moauth helps with didn't change from OAuth2 r1 to r2, so Moauth r1 (the only version that exists) is compatible with either. Both Azure r2 and AWS r3 use Moauth r1. That works fine in programs using only Azure or only AWS, but Ugo's Unity build actually looks like:



Unity needs both copies of OAuth2, but then which one does Moauth import?

In order to make the build work, it would seem that we need two identical copies of Moauth: one that imports OAuth2 r1, for use by Azure, and a second that imports OAuth2 r2, for use by AWS. A quick StackOverflow check shows that the package manager has a flag for that: `-fc:clone`. Using this flag, Ugo's program builds as:



This actually works and passes its tests, although Ugo now wonders if there are more problems lurking. He heads home for a late dinner.


Chapter 3

Back at Microsoft, Amy has returned from vacation. She decides that Azure can keep using OAuth2 r1 for a while longer, but she realizes that it would help users to let them pass Moauth tokens directly into the Azure API. She adds this to the Azure package in a backwards-compatible way and releases Azure r3. Over at Amazon, Anna likes the Azure package's new Moauth-based API and adds a similar API to the AWS package, releasing AWS r4.


Ugo sees these changes and decides to update to the latest version of both Azure and AWS in order to use the Moauth-based APIs. This time he blocks off an afternoon. First he tentatively updates the Azure and AWS packages without modifying Unity at all. His program builds!

Excited, Ugo changes Unity to use the Moauth-based Azure API, and that builds too. When he changes Unity to also use the Moauth-based AWS API, though, the build fails. Perplexed, he reverts his Azure changes, leaving only the AWS changes, and the build succeeds. He puts the Azure changes back, and the build fails again. Ugo returns to StackOverflow.

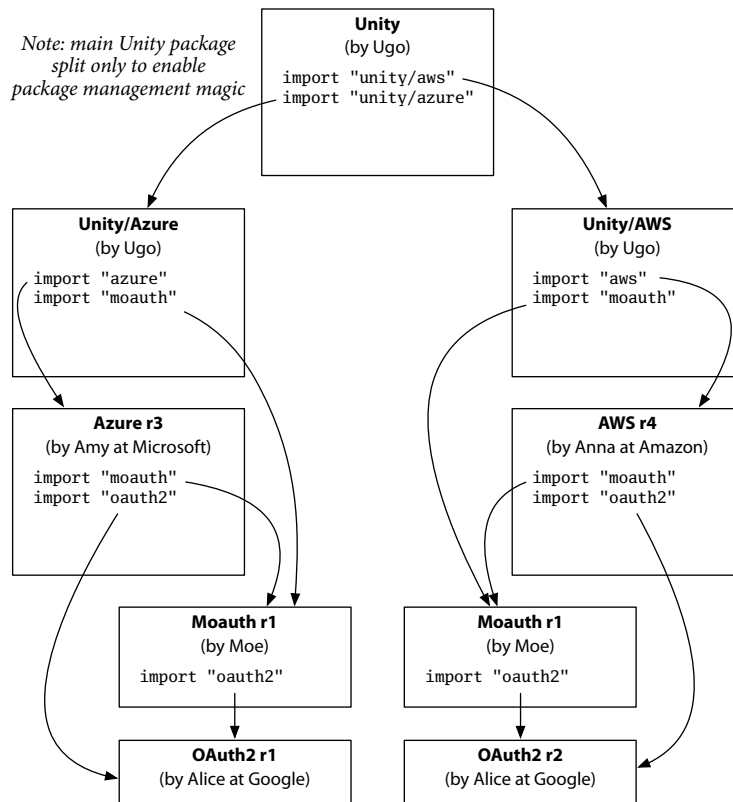
Ugo learns that when using just one Moauth-based API (in this case, Azure) with `-fmultiverse -fclone`, Unity implicitly builds as:



but when he is using both Moauth-based APIs, the single `import "moauth"` in Unity is ambiguous. Since Unity is the main package, it cannot be cloned (in contrast to Moauth itself):



A comment on StackOverflow suggests moving the Moauth import into two different packages and having Unity import them instead. Ugo tries this and, incredibly, it works:



Ugo makes it home on time. He's not terribly happy with his package manager, but he's now a big fan of StackOverflow.

A Retelling with Semantic Versioning

Let's wave a magic wand and retell the story with semantic versions, assuming that the package manager uses them instead of the original story's 'r' numbers.

Here's what changes:

- OAuth2 r1 becomes OAuth2 1.0.0.
- Moauth r1 becomes Moauth 1.0.0.
- Azure r1 becomes Azure 1.0.0.
- AWS r1 becomes AWS 1.0.0.
- OAuth2 r2 becomes OAuth2 2.0.0 (partly incompatible API).
- Azure r2 becomes Azure 1.0.1 (bug fix).
- AWS r2 becomes AWS 1.0.1 (bug fix, internal use of OAuth2 2.0.0).
- AWS r3 becomes AWS 1.1.0 (feature update: add Zeta).
- Azure r3 becomes Azure 1.1.0 (feature update: add Moauth APIs).
- AWS r4 becomes AWS 1.2.0 (feature update: add Moauth APIs).

Nothing else about the story changes. Ugo still runs into the same build problems, and he still has to turn to StackOverflow to learn about build flags and refactoring techniques just to keep Unity building. According to semver, though, Ugo should have had no trouble at all with any of his updates: not one of

the packages that Unity imports changed its major version during the story. Only OAuth2 did, deep in Unity's dependency tree. Unity itself does not import OAuth2. What went wrong?

The problem here is that the semver spec is really not much more than a way to choose and compare version strings. It says nothing else. In particular, it says nothing about how to handle incompatible changes after incrementing the major version number.

The most valuable part of semver is the encouragement to make backwards-compatible changes when possible. The FAQ correctly notes:

“Incompatible changes should not be introduced lightly to software that has a lot of dependent code. The cost that must be incurred to upgrade can be significant. Having to bump major versions to release incompatible changes means you'll think through the impact of your changes and evaluate the cost/benefit ratio involved.”

I certainly agree that “incompatible changes should not be introduced lightly.” Where I think semver falls short is the idea that “having to bump major versions” is a step that will make you “think through the impact of your changes and evaluate the cost/benefit ratio involved.” Quite the opposite: it's far too easy to read semver as implying that as long as you increment the major version when you make an incompatible change, everything else will work out. The example shows that this is not the case.

From Alice's point of view, the OAuth2 API needed backwards-incompatible changes, and when she made them, semver seemed to promise it would be fine to release an incompatible OAuth2 package, provided she gave it version 2.0.0. But that semver-approved change triggered the cascade of problems that befell Ugo and Unity.

Semantic versions are an important way for authors to convey expectations to users, but that's all they are. By itself, it can't be expected to solve these larger build problems. Instead, let's look at an approach that does solve the build problems. Afterward, we can consider how to fit semver into that approach.


A Retelling with Import Versioning

Once again, let's retell the story, this time using the import compatibility rule:

In Go, if an old package and a new package have the same import path, the new package must be backwards compatible with the old package.

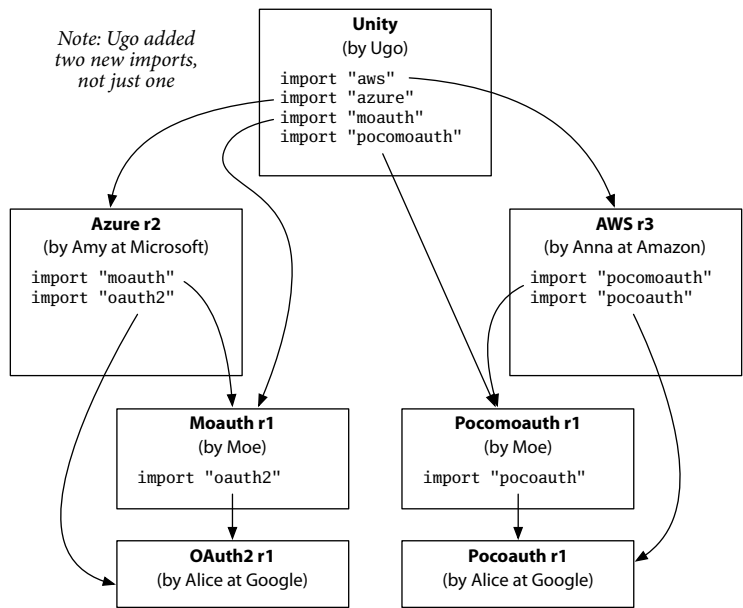
Now the plot changes are more significant. The story starts out the same way, but in Chapter 1, when Alice decides to create a new, partly incompatible OAuth2 API, she cannot use "oauth2" as its import path. Instead, she names the new version Pocomo and gives it the import path "pocomo". Presented with two different OAuth2 packages, Moe (the author of Moauth) must write a second package, Moauth for Pocomo, which he names Pocomo and gives the import path "pocomo". When Anna updates the AWS package to the new OAuth2 API, she also changes the import paths in that code from "oauth2" to "pocomo" and from "moauth" to "pocomo". Then the story proceeds as before, with the release of AWS r2 and AWS r3.

In Chapter 2, when Ugo eagerly adopts Amazon Zeta, everything just works. The imports in all the packages code exactly match what needs to be built. He doesn't have to look up special flags on StackOverflow, and he's only five minutes late to lunch.



In Chapter 3, Amy adds Moauth-based APIs to Azure while Anna adds equivalent Pocomoauth-based APIs to AWS.

When Ugo decides to update both Azure and AWS, again there's no problem. His updated program builds without any special refactoring:



At the end of this version of the story, Ugo doesn't even think about his package manager. It just works; he barely notices that it's there.

In contrast to the semantic versioning translation of the story, the use of import versioning here changed two critical details. First, when Alice introduced her backwards-incompatible OAuth2 API, she had to release it as a new package (Pocoauth). Second, because Moe's wrapper package Moauth exposed the OAuth2 package's type definitions in its own API, Alice's release of a new package forced Moe's release of a new package (Pocomoauth). Ugo's final Uni-

ty build went well because Alice's and Moe's package splits created exactly the structure needed to keep clients like Unity building and running. Instead of Ugo and users like him needing incomplete package manager complexity like `-fmultiverse -fclone` aided by extraneous refactorings, the import compatibility rule pushes a small amount of additional work onto package authors, and all users benefit.

There is certainly a cost to needing to introduce a new name for each backwards-incompatible API change, but as the semver FAQ says, that cost should encourage authors to more clearly consider the impact of such changes and whether they are truly necessary. And in the case of Import Versioning, the cost pays for significant benefits to users.

An advantage of Import Versioning here is that package names and import paths are well-understood concepts for Go developers. If you tell an author that making a backwards-incompatible change requires creating a different package with a different import path, then—without any special knowledge of versioning—the author can reason through the implications on client packages: clients are going to need to change their own imports one at a time; Moauth is not going to work with the new package; and so on.

Able to predict the effects on users more clearly, authors might well make different, better decisions about their changes. Alice might look for way to introduce the new, cleaner API into the original OAuth2 package alongside the existing APIs, to avoid a package split. Moe might look more carefully at whether he can use interfaces to make Moauth support both OAuth2 and Pocomoauth, avoiding a new Pocomoauth package. Amy might decide it's worth updating to Pocomoauth and Pocomoauth instead of exposing the fact that the Azure APIs use outdated OAuth2 and Moauth packages. Anna might have tried to make the AWS APIs allow either Moauth or Pocomoauth, to make it easier for Azure users to switch.

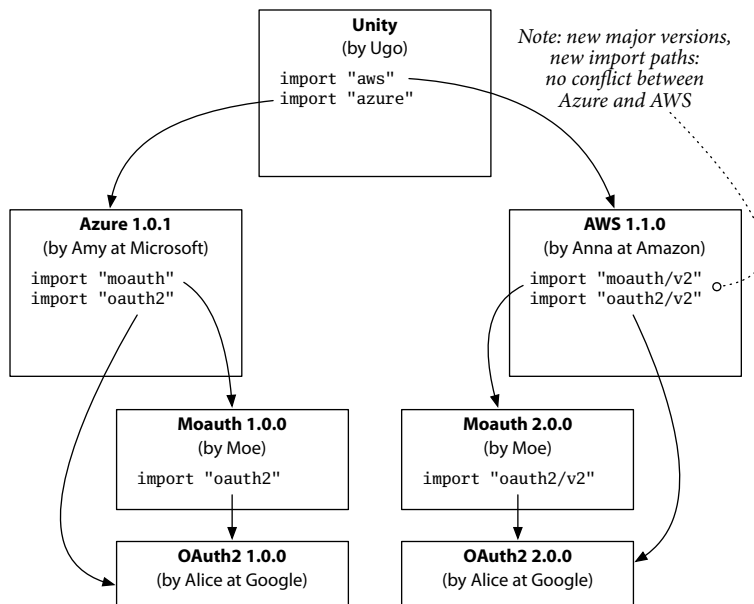
In contrast, the implications of a semver “major version bump” are far less clear and do not exert the same kind of design pressure on authors. To be clear, this approach creates a bit more work for authors, but that work is justified by delivering significant benefits to users. In general, this balance makes sense, because packages aim to have many more users than authors, and hopefully all packages at least have as many users as they do authors.

Semantic Import Versioning


The previous section showed how import versioning leads to simple, predictable builds during updates. But choosing a new name at every backwards-incompatible change is difficult and unhelpful to users. Given the choice between OAuth2 and Pocomoauth, which should Amy use? Without further investigation, there's no way to know. In contrast, semantic versioning makes this easy: OAuth2 2.0.0 is clearly the intended replacement for OAuth2 1.0.0.

We can use semantic versioning *and* follow the import compatibility rule by including the major version in the import path. Instead of needing to invent a cute but unrelated new name like Pocomoauth, Alice can call her new API OAuth2 2.0.0, with the new import path `"oauth2/v2"`. The same for Moe: Moauth 2.0.0 (imported as `"moauth/v2"`) can be the helper package for OAuth2 2.0.0, just as Moauth 1.0.0 was the helper package for OAuth2 1.0.0.

When Ugo adds Zeta support in Chapter 2, his build looks like:



Because "moauth" and "moauth/v2" are simply different packages, it is perfectly clear to Ugo what he needs to do to use "moauth" with Azure and "moauth/v2" with AWS: import both.



For compatibility with existing Go usage and as a small encouragement not to make backwards-incompatible API changes, I am assuming here that major version 1 is omitted from import paths: `import "moauth"`, not `"moauth/v1"`. Similarly, major version 0, which explicitly disavows compatibility, is also omitted from import paths. The idea here is that by using a v0 dependency, users are explicitly acknowledging the possibility of breakage and taking on the responsibility to deal with it when they choose to update. (Of course, it's then important that updates don't happen automatically. We'll see in the next post how minimal version selection helps with that.)

Functional Names & Immutable Meanings

Twenty years ago, Rob Pike and I were modifying the internals of a Plan 9 C library, and Rob taught me the rule of thumb that when you change a function's behavior, you also change its name. The old name had one meaning. By using a different meaning for the new name and eliminating the old one, we ensured the compiler would complain loudly about every piece of code that needed to be examined and updated, instead of silently compiling incorrect code. And if people had their own programs using the function, they'd get a compile-time failure instead of a long debugging session. In today's world of distributed version control, that last problem is magnified, making the name change even more important. A merge of concurrently-written code expecting the old semantics should not silently get the new semantics instead.

Of course, deleting an old function works only when all the uses can be found, or when users understand that they are responsible for keeping up with changes, as was the case in a research system like Plan 9. For exported APIs, it's usually much better to leave the old name and old behavior intact and only add a new name with new behavior. Rich Hickey made the point in his "Spec-ulation" talk in 2016 that this approach of only adding new names and behaviors, never removing old names or redefining their meanings, is exactly what functional programming encourages with respect to individual variables or data structures. The functional approach brings benefits in clarity and predictability in small-scale programming, and the benefits are even larger when applied, as in the import compatibility rule, to whole APIs: dependency hell is really just mutability hell writ large. That's just one small observation in the talk; the whole thing is worth watching.

In the early days of "go get", when people asked about making backwards-incompatible changes, our response—based on intuition derived from years of experience with these kinds of software changes—was to give the import versioning rule, but without a clear explanation why this approach was better than not putting the major version in the import paths. Go 1.2 added a FAQ entry about package versioning that gave this basic advice (unchanged as of Go 1.10):

Packages intended for public use should try to maintain backwards compatibility as they evolve. The Go 1 compatibility guidelines are a good reference here: don't remove exported names, encourage tagged composite literals, and so on. If different functionality is required, add a new name instead of changing an old one. If a complete break is required, create a new package with a new import path.

One motivation for this blog post is to show, using a clear, believable example, why following the rule is so important.

Avoiding Singleton Problems

One common objection to the semantic import versioning approach is that package authors today expect that there is only ever one copy of their package in a given build. Allowing multiple packages at different major versions may cause problems due to unintended duplications of singletons. An example would be registering an HTTP handler. If `my/thing` registers an HTTP handler for `/debug/my/thing`, then having two copies of the package will result in duplicate registrations, which causes a panic at registration time. Another problem would be if there were two HTTP stacks in the program. Clearly only one HTTP stack can listen on port 80; we wouldn't want half the program registering handlers that will not be used. Go developers are already running into problems like this due to vendoring inside vendored packages.

Moving to vgo and semantic import versioning clarifies and simplifies the current situation though. Instead of the uncontrolled duplication caused by vendoring inside vendoring, authors will have a guarantee that there is only one instance of each major version of their packages. By including the major version into the import path, it should be clearer to authors that `my/thing` and `my/thing/v2` are different and need to be able to coexist. Perhaps that means exporting debug information for v2 on `/debug/my/thing/v2`. Or perhaps it means coordinating. Maybe v2 can take charge of registering the handler but also provide a hook for v1 to supply information to display on the page. This would mean `my/thing` importing `my/thing/v2` or vice versa; with different import paths, that's easy to do and easy to understand. In contrast, if both v1 and v2 are `my/thing` it's hard to comprehend what it means for one to import its own import path and get the other.

Automatic API Updates

One of the key reasons to allow both v1 and v2 of a package to coexist in a large program is to make it possible to upgrade the clients of that package one at a time and still have a buildable result. This is specific instance of the more general problem of gradual code repair. (See my 2016 article, “Codebase Refactoring (with help from Go),” for more on that problem.)

In addition to keeping programs building, semantic import versioning has a significant benefit to gradual code repair, which I touched on in the previous section: one major version of a package can import and be written in terms of another. It is trivial for the v2 API to be written as a wrapper of the v1 implementation, or vice versa. This lets them share the code and, with appropriate design choices and perhaps use of type aliases, might even allow clients using v1 and v2 to interoperate. It may also help resolve a key technical problem in defining automatic API updates.

Before Go 1, we relied heavily on `go fix`, which users ran after updating to a new Go release and finding their programs no longer compiled. Updating code that doesn't compile makes it impossible to use most of our program analysis tools, which require that their inputs are valid programs. Also, we've wondered how to allow authors of packages outside the Go standard library to supply “fixes” specific to their own API updates. The ability to name and work with multiple incompatible versions of a package in a single program suggests a possible solution: if a v1 API function can be implemented as a wrapper around the v2 API, the wrapper implementation can double as the fix specification. For example, suppose v1 of an API has functions `EnableFoo` and `DisableFoo` and v2 replaces the pair with a single `SetFoo(enabled bool)`. After v2 is released, v1 can be implemented as a wrapper around v2:

```
package p // v1

import v2 "p/v2"

func EnableFoo() {
    //go:fix
    v2.SetFoo(true)
}

func DisableFoo() {
    //go:fix
    v2.SetFoo(false)
}
```

The special `//go:fix` comments would indicate to `go fix` that the wrapper body that follows should be inlined into the call site. Then running `go fix` would rewrite calls to `v1 EnableFoo` to `v2 SetFoo(true)`. The rewrite is easily specified and type-checked, since it is plain Go code. Even better, the rewrite is clearly safe: `v1 EnableFoo` is *already* calling `v2 SetFoo(true)`, so rewriting the call site plainly does not change the meaning of the program.

It is plausible that `go fix` might use symbolic execution to fix even the reverse API change, from a `v1` with `SetFoo` to a `v2` with `EnableFoo` and `DisableFoo`. The `v1 SetFoo` implementation could read:

```
package q // v1

import v2 "q/v2"

func SetFoo(enabled bool) {
    if enabled {
        //go:fix
        v2.EnableFoo()
    } else {
        //go:fix
        v2.DisableFoo()
    }
}
```

and then `go fix` would update `SetFoo(true)` to `EnableFoo()` and `SetFoo(false)` to `DisableFoo()`. This kind of fix would even apply to API updates within a single major version. For example, `v1` could be deprecating (but keeping) `SetFoo` and introducing `EnableFoo` and `DisableFoo`. The same kind of fix would help callers move away from the deprecated API.

To be clear, this is not implemented today, but it seems promising, and this kind of tooling is made possible by giving different things different names. These examples demonstrate the power of having durable, immutable names attached to specific code behavior. We need only follow the rule that when you make a change to something, you also change its name.

Committing to Compatibility

Semantic import versioning is more work for authors of packages. They can't just decide to issue `v2`, walk away from `v1`, and leave users like Ugo to deal with the fallout. But authors who do that are hurting their users. It seems to me a good thing if the system makes it harder to hurt users and instead naturally steers authors toward behaviors that don't hurt users.

More generally, Sam Boyer talked at GopherCon 2017 about how package managers moderate our social interactions, the collaboration of people building software. We get to decide. Do we want to work in a community built around a system that optimizes for compatibility, smooth transitions, and working well together? Or do we want to work in a community built around a system that optimizes for creating and describing incompatibility, that makes it acceptable for authors to break users' programs? Import versioning, and in particular handling semantic versioning by lifting the semantic major version into the import path, is how we can make sure we work in the first kind of community.

Let's commit to compatibility.