

# Addressing the limitations of Kubernetes' Ingress object

David Cheney - Heptio<sup>†</sup>



Abstract: The Kubernetes Ingress object has a number of limitations which over the years have been papered over with annotations. Contour, the Ingress controller my team at Heptio are building, recently introduced a new Ingress object which addresses the existing limitations and unlocks the ability for teams and operators to have more control over ingress deployments in multi team and multi tenant scenarios. In this short talk I'll explain the limitations of the current ingress object and how our new Ingress object addresses those shortcomings while making it possible for multiple teams to collaborate and delegate responsibility using various routing patterns and strategies that our new Ingress object makes possible.

g'day



My name is David, for the last year, and perhaps one more week, I work for Heptio.

We're a plucky little startup based seattle working on building tools, automation, and visualisations on top of Kubernetes to make Kubernetes easier to use.

Connaissez-vous  
Kubernetes?



More importantly is a talk about Kubernetes, so as a quick straw poll, who here is using or thinking about using kubernetes to deploy their applications?

# Ingress-*what* controller?



The part of kubernetes that I spend my time in is something called an ingress controller.

Ingress controllers is responsible for getting traffic from the outside world down to your pods.

But practical terms; HTTP, TLS, load balancers, reverse proxies all at Layer 7.

**An ingress controller should take  
care of the 90% use case for  
deploying HTTP middleware**



That's quite a broad remit so the way I approach the design space is I think that a good ingress controller should take care of 90% of the cases that traditionally you would have used an apache, or nginx, or squid, sidecar container or something in the request flow to your app.

# Getting to the 90% case

- Traffic consolidation
- TLS management
- Abstract configuration
- Path based routing



Here are some things that I think contribute to an ingress controller getting to that 90% level of functionality.

The first one is consolidation. If you use a service load balancer then every service you deploy has an ELB in front of it, that's a cost, and also a maintenance issue. They consume a public IP which are a scarce resource, and depending on your company your security team may not be cool with hundreds of public IPs funnelling trading into their kubernetes cluster.

The second is TLS management. It's 2018, you need to be talking TLS. Chrome 68 is out and non https sites "insecure". We have projects like cert manager and lets encrypt that take care of obtaining a certificate, and an ingress controller covers presenting that certificate on port 443, so there shouldn't be any reason to not be secure in 2018.

The third is a notion of being able to describe the properties of your web application in an abstract manner, or at least to have some portability between different ingress controllers (and clouds). You should be able to talk about the host name, tls configuration, route names and backends for those routes without having to write an apache configuration, or an nginx configuration.

Path based routing; with a service, all the traffic goes straight to the cluster IP, if you wanted to serve your static images from one service, and your dynamic data from another, you can't do that with a service.

# What is Contour?



So what is contour?

Contour is an ingress controller that I've built at heptio. We use Lyft's envoy proxy as our data plane

Contour exists to fulfil the requirements I just described.

I want to be clear that this talk isn't a product pitch — well it is — but it's not for contour.  
The fact that the stuff I'm going to talk to you about is implemented in contour is incidental

# What are the problems with Ingress?



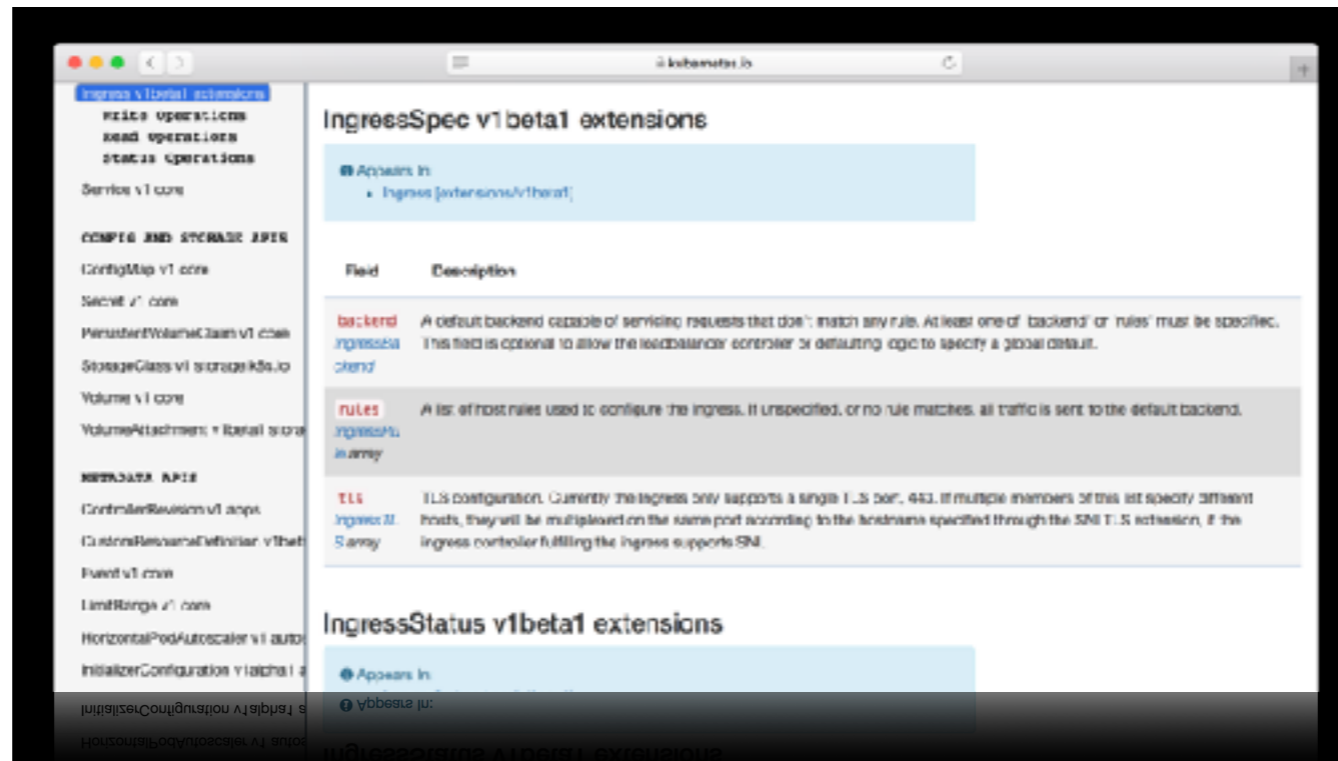
Let's talk about the problems the Ingress object.



JSON is not a spec



I think my biggest complaint with ingress is the “specification” is a bunch of text in comments on the api data structure.



This is literally the specification implementors have to work from.

This is perfectly fine for someone who needs a schema to send a message to the k8s API server, but as an implementor, or someone building tools on top of this API, we need a level of formality that just isn't there.

What I would like to see is something akin to the level of detail of an RFC, because there are, as we'll see the 'spec' is rife with ambiguity,

# Gosh darned default backend



Take for example the default backend.

“A default backend capable of servicing requests that don't match any rule.”



Each ingress document has a notion of default backend. Requests that don't match any rule are routed to the default backend

This would make sense if there is only ever one single ingress document, but that is almost never the case, also as we'll see routes can fail to match for reasons other than the path

# Default backend ambiguity

- Default backend conflates the notion of a vhost, the Host: header traffic arrives on, from the backend to serve it
- The host key in spec.rules is *optional* — does this mean the rule matches *any* host?  $\overline{\setminus}(\overline{\_}(\overline{\_} \overline{\_}))\overline{\_}$
- Default backend can be present in multiple Ingress objects — which takes precedence?



Default backend conflates ...

A default *vhost* is a notion of a http handler for traffic that fails all other routing rules — that needn't be a backend, it could be a simple 302 redirect

Host key is optional ...

Default backend present in multiple—all?—ingress objects, which takes precedence? Should they be merged together? That's a bit tricky because the default backend is a service, not something you can do a route match on.

And this ambiguity suggests that a default backend can be namespaced because ingress objects can be namespaces, but of course that doesn't work.

# Ingress objects can span namespaces



Speaking of namespaces, the ingress spec permits the definition of a virtual host to span more than one ingress object.

I can see the argument for this; ingress objects can only use services in the same namespace, but what if you want to have `/finance` managed in the finance namespace and `/ads` in the ads namespace? You do this by putting part of the vhost definition in the finance namespace and part of the vhost definition in the ads namespace. They both refer to the same virtual host, so the ingress controller stitches them all together for you.

However, this means that if someone has RBAC permission to add an ingress object in their namespace, they can inject a route onto the ingress you defined in your namespace.

**Cert-manager *relies* on this  
feature to support Let's  
Encrypt's HTTP-01  
challenge!**



And just in case you were thinking — hmm this sounds like a massive security hole, I'd like to disable this please — you cannot because in a perfect example of Hyrum's Law, projects like Kube lego and cert-manager rely on the ability to inject a route onto your vhost from another namespace so they can route the HTTP-01 challenge to a service running in their namespace

**Ingress makes shared tenancy  
difficult if your tenants aren't  
incentivised to play nicely with  
each other**



What this boils down too is Ingress is very difficult to use in a shared kubernetes cluster.

There are no safeguards to prevent anyone with RBAC permission to create or edit ingress objects from accidentally, or maliciously, injecting conflicting or invalid configuration onto the vhost for another tenant.



# One route. One Service.



Lets talk about some other problems with the ingress spec that affect people trying to use the modern patterns for web applications.

Kubernetes services are mapped onto http routes via an ingress document, however the ingress spec only permits `_one_` service per route.

IngressBackend describes all endpoints for a given service and port.

- i** Appears In:
- [HTTPIngressPath](#) extensions/v1beta1
  - [IngressSpec](#) extensions/v1beta1

Field	Description
<code>serviceName</code> <i>string</i>	Specifies the name of the referenced service.
<code>servicePort</code>	Specifies the port of the referenced service.

Now a kubernetes service can match multiple pods if they share the same label, they'll all get mixed into the same endpoint document, but at best you're going to get a weighted distribution across the deployments that make up the service.

If you want to send 1% of your traffic to the new version of your application, you'll need to have 99 pods running the old version of your app to make the ratios work out.

# Annotation potpourri



Another big problem with the ingress spec is the schema is so limited the only place you can stuff extra parameters or attributes about your web application is in annotations

## Cambrian explosion of Ingress annotations

- Allow port 80 and/or 301 upgrade to HTTPS
- Request timeout (applies to all the entries in the Ingress document)
- Retry parameters (also applies to all entries in the Ingress document)
- TLS minimum protocol version
- Websocket enabled routes



This has led to a Cambrian explosion of ingress document annotations.

here are just some that contour implements, this is barely scratching the surface of what has been shoehorned into an untyped map of annotations by various ingress controllers.

If this Lasso fair approach wasn't bad enough, the configuration of a vhost may be spread across several ingress documents, so how these annotations are applied is confusing

301 upgrade settings, request timeouts and retry parameters likely apply per ingress document not per vhost, therefore if you want those settings to apply to some routes and not to others for a vhost, you have to split them across two ingress documents

TLS minimum protocol version has to be specified in an annotation because the TLS stanza of the ingress document has no place for it; same with cipher specs.

if you split a host across several ingress documents, do things like TLS min protocol apply to all the ingresses that match that host, or just the one where the annotation is present? There is no right answer here. If you say TLS min proto applies only to the ingress spec in a single document, you're committing to altering the TLS handshaking operation based on the request line of the HTTP request which you don't have at that point.

If you say that annotations like TLS min proto apply across any ingress with that host, because ingresses can span namespaces, someone in another namespace can alter the TLS parameters for your virtual host even though they don't have permission to write into your namespace.

# Ingress isn't broken, it's just limited



I want to take a moment to say that while I personally have a bunch of gripes with Ingress coming from my position as an implementor ingress isn't broken.

I don't want you to come away from this talk thinking "welp, dave says I can't use that at all"

Ingress isn't broken, it's just limited, and if you're not hitting those limits then far be it from me to tell you need to change what you're doing.

# What are Heptio doing about these problems?



However if you have experienced some of these problems, then let me tell you about what we're doing in contour to try to improve the situation.

At the start of the year Heptio signed a joint development deal with Yahoo Japan to build them a large load balancing solution for them using a kubernetes cluster almost like an appliance.

# Ingress IngressRoute



Realising that yahoo Japan were encountering many of the issues with multi tenancy that I described above we set out to define a new kind of ingress document, which we call ingressroute

Every Ingressroute document has  
one hostname



The first thing that we changed is each ingress route document refers to one hostname and one hostname only.



```
apiVersion: contour.heptio.com/v1beta1
kind: IngressRoute
metadata:
  name: blog
  namespace: marketing
spec:
  virtualhost:
    fqdn: blog.heptio.com
    tls:
      secretName: blog-secret
  routes:
  - match: /
    services:
    - name: blog-svc
      port: 80
```

**The virtualhost key indicates  
this is a root ingressroute**



This means all the properties of a virtual host, its name, its tls parameters, the secret that holds the tls certificate are in one namespace alone.

We call this the ingress route document a root, for reasons I'll explain in a little bit

## Load balancing strategies can be specified per backend service



Because we're no longer limited to the schema of the kubernetes ingress object we now have a place to hang configuration attributes that used to be smuggled into annotations.

```
apiVersion: contour.heptio.com/v1beta1
kind: IngressRoute
metadata:
  name: blog
  namespace: marketing
spec:
  virtualhost:
    fqdn: blog.heptio.com
    tls:
      secretName: blog-secret
  routes:
  - match: /blog
    services:
    - name: blog-svc
      port: 80
      strategy: WeightedLeastRequest
```

**For this route, use  
WeightedLeastRequest  
across the endpoints  
matching blog-svc**



For example, per route, per service, you can control the load balancing strategy that will be used across the endpoints that make up this service.

# Websocket support



A key reasons for choosing Envoy as our data plane was Envoy's support of long running websocket sessions across configuration changes.

```
apiVersion: contour.heptio.com/v1beta1
kind: IngressRoute
metadata:
  name: chat
  namespace: default
spec:
  virtualhost:
    fqdn: chat.example.com
  routes:
    - match: /
      services:
        - name: chat-app
          port: 80
    - match: /websocket
      enableWebsockets: true
      services:
        - name: chat-app
          port: 80
```

**Only permit  
Upgrade: websocket  
on /websocket**



Enabling websocket support per route is as simple as adding the `enableWebsockets: true` key to your route.

In general where something can be enabled for all use cases rather than having a parameter or flag that people have to know to turn on, my policy is to turn it on across the board; http compression is a good example of this.

However I wasn't comfortable permitting `Upgrade: websocket` by default for all routes, so we made a deliberate decision to make it opt in only.

# Multiple service backends



The kubernetes ingress document limits routes to a single backend service. Using ingressroute we have the ability to say instead of a single service, allow a list of services.

```
apiVersion: contour.heptio.com/v1beta1
kind: IngressRoute
metadata:
  name: blog
  namespace: marketing
spec:
  virtualhost:
    fqdn: blog.heptio.com
    tls:
      secretName: blog-secret
  routes:
  - match: /
    services:
    - name: service1
      port: 8080
    - name: service2
      port: 8080
```

**Traffic will be load balanced  
across service1 and service2**



# Weighted services



The main reason you'd want to have more than one backend service per route is to enable patterns like canary deploys or blue/green deployments



```
apiVersion: contour.heptio.com/v1beta1
kind: IngressRoute
metadata:
  name: gmail
  namespace: google
spec:
  virtualhost:
    fqdn: gmail.google.com
  routes:
  - match: /
    services:
    - name: gmail-v1.3.1
      port: 80
      weight: 90
    - name: gmail-v2.0.0
      port: 80
      weight: 10
```

Shift traffic from  
v1.3.1 to v2.0.0 by altering  
the service weights



In this example, 90% of the requests to gmail.google.com are routed to the version 1.3.1 and 10% are routed to version 2.0.0. As you gain confidence in the the deployment you can edit the document to increase the weighting towards version 2.

And of course, weighting, load balancing strategy, websockets, etc can be combined per service, per route, depending on your applications needs.

It's important to note that modifying the weights triggers an immediate shift of traffic pattern in Envoy (via Contour).

# Delegation



Delegation is our answer to helping multi tenant clusters stay manageable.

All the ingressroute documents we've seen so far are what we call "root documents", because they are at the root of a delegation tree.

To explain why we think delegation is powerful let me lay out a scenario for you.

# https://google.com/finance

- You want to delegate control of https://google.com/finance to the Google Finance developers working the google-finance namespace.
- The Google Finance team should not be able to alter the configuration for the rest of https://google.com/
- None of the teams working on https://google.com/ should have access to the TLS secret for https://google.com/



```
apiVersion: contour.heptio.com/v1beta1
kind: IngressRoute
metadata:
  name: google-com
  namespace: google
spec:
  virtualhost:
    fqdn: google.com
    tls:
      secret: google-com-secret
  routes:
  - match: /
    delegate:
      name: search
      namespace: google-search
  - match: /finance
    delegate:
      name: finance
      namespace: google-finance
```

The configuration for /  
finance is found in the  
finance/finance ingressroute




In this example we have a standard ingressroute root; its for google.com, and references google-com-secret in the google namespace.

However all the routes, / and /finance refer to ingressroute documents in other namespaces. You can think of this as an “include” macro, contour will find the configuration fragment for the search service in the search namespace, and the finance service in the finance namespace.

Let's have a look at finance

```
apiVersion: contour.heptio.com/v1beta1
kind: IngressRoute
metadata:
  name: finance
  namespace: google-finance
spec:
  routes:
  - match: /finance
    services:
      name: finance-v1.0.1
      port: 8080
```

Only routes that are a sub match of /finance



Here is the finance ingressroute document.

It does not have a virtual host stanza, which means it is not a root. It is a delegate and can only be referenced by other ingressroute documents that delegate to it explicitly. And contour is only going to reference routes that start with the prefix that was delegated too

DNS example

# Restricted root namespaces



The final piece of the multi tenant puzzle is the ability to restrict the namespaces that contour will look for ingressroute roots.

This is an opt in feature, by default anyone with RBAC permission to create an ingressroute can do so, but if you want to make creating a new root a administrative event, you can configure contour to only look for roots in a set of namespaces which you have arranged that only administrators can write too.

**Are you going to upstream  
Ingressroute?**



I'm sure the first question out of your mouths is going to be, are you planning on upstreaming this work?

# Well ... it's complicated



And the answer is it's complicated.

One analogy to explain my thinking about this is climate change.

There is universal agreement, at least from people who pay attention, that climate change is a serious problem.

However, the solutions are scattered, the political will and consensus is just not there yet.

I don't think that there is much disagreement that the current Kubernetes Ingress object is a dead end, but I'm not seeing any strong movement upstream to replace Ingress.

I'm sure they're well aware of the problem, but there are bigger problems in kubernetes, and this isn't the pot boiling over at the moment

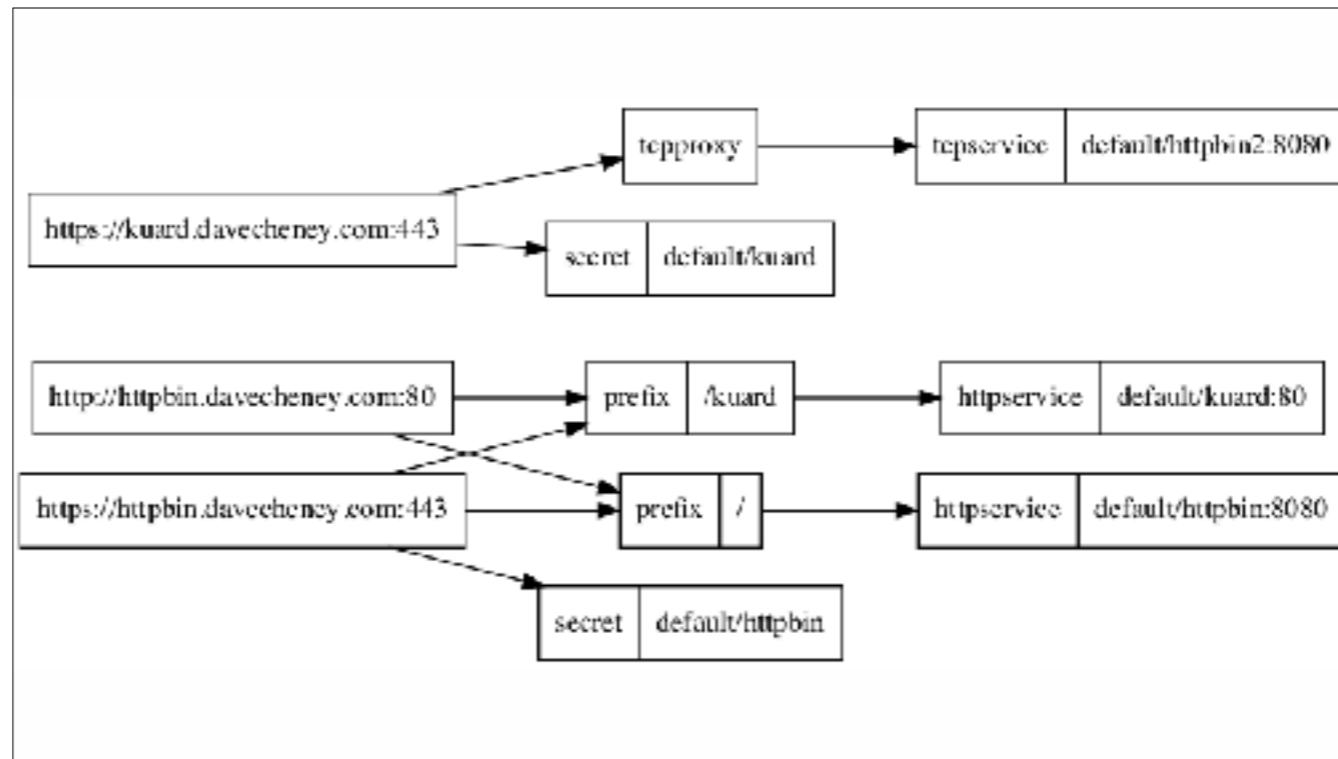
Maybe in the future when people on sig-networking are seriously discussing a replacement for Ingress it'll be the time to propose something based on our work, but my feeling is that now isn't that time

But ...



**Would you prefer to integrate  
IngressRoute into your Ingress  
controllers?**





Contour has a thing called the DAG which abstracts away both the k8s API server, and the downstream Envoy configuration.

A DAG is built from objects k8s objects, and Contour walks over the DAG to produce the various configuration tables Envoy needs.

So one option is to expose Contour's DAG as a library, so that other ingress controllers can use it. If they write visitors for the dag that produce configuration that Traefic or nginx then they would get ingress route support for free because at the level the DAG works at, the two ingress types are abstracted away

Thank you!

👉 [github.com/heptio/contour](https://github.com/heptio/contour)

👉 [@davecheney](https://twitter.com/davecheney)

👉 [dfc@heptio.com](mailto:dfc@heptio.com)



*Image: Egon Elbre*

Thank you for your time.

If you want to talk to me about anything I've said here, ingress, kubernetes, or go, come find me, I've got stickers for a company which won't exist in a week, so they might be worth something as antiques.