
Ocelot Documentation

Release 1.0.0

Tom Pallister

Sep 25, 2018

1	Big Picture	3
1.1	Basic Implementation	4
1.2	With IdentityServer	4
1.3	Multiple Instances	5
1.4	With Consul	5
1.5	With Service Fabric	6
2	Getting Started	7
2.1	.NET Core 2.1	7
2.2	.NET Core 1.0	8
3	Contributing	11
4	Not Supported	13
5	Configuration	15
5.1	Multiple environments	16
5.2	Merging configuration files	17
5.3	Store configuration in consul	17
5.4	Reload JSON config on change	18
5.5	Follow Redirects / Use CookieContainer	18
5.6	SSL Errors	19
6	Routing	21
6.1	Catch All	22
6.2	Upstream Host	23
6.3	Priority	23
6.4	Dynamic Routing	24
6.5	Query Strings	24
7	Request Aggregation	27
7.1	Advanced register your own Aggregators	27
7.2	Basic expecting JSON from Downstream Services	29
8	GraphQL	31
9	Service Discovery	33
9.1	Consul	33

9.2	Eureka	35
9.3	Dynamic Routing	35
10	Service Fabric	39
11	Authentication	41
11.1	JWT Tokens	42
11.2	Identity Server Bearer Tokens	42
11.3	Okta	43
11.4	Allowed Scopes	43
12	Authorisation	45
13	Websockets	47
13.1	SignalR	47
13.2	Supported	48
13.3	Not Supported	48
14	Administration	51
14.1	Providing your own IdentityServer	51
14.2	Internal IdentityServer	52
14.3	Administration API	53
15	Rate Limiting	55
16	Caching	57
16.1	Your own caching	58
17	Quality of Service	59
18	Headers Transformation	61
18.1	Add to Request	61
18.2	Add to Response	61
18.3	Find and Replace	62
18.4	Pre Downstream Request	62
18.5	Post Downstream Request	62
18.6	Placeholders	62
18.7	Handling 302 Redirects	63
18.8	X-Forwarded-For	63
18.9	Future	63
19	Claims Transformation	65
19.1	Claims to Claims Transformation	65
19.2	Claims to Headers Transformation	66
19.3	Claims to Query String Parameters Transformation	66
20	Logging	67
20.1	Warning	67
21	Tracing	69
21.1	Butterfly	69
22	Request Id / Correlation Id	71
23	Middleware Injection and Overrides	73

24 Configuration	75
25 Service Discovery	77
26 CookieStickySessions	79
27 Configuration	81
28 Service Discovery	83
29 CookieStickySessions	85
30 Configuration	87
31 Service Discovery	89
32 CookieStickySessions	91
33 Delegating Handlers	93
33.1 Usage	93
34 Raft (EXPERIMENTAL DO NOT USE IN PRODUCTION)	95
35 Overview	97
36 Building	99
37 Tests	101
38 Release process	103

Thanks for taking a look at the Ocelot documentation. Please use the left hand nav to get around. I would suggest taking a look at introduction first.

CHAPTER 1

Big Picture

Ocelot is aimed at people using .NET running a micro services / service orientated architecture that need a unified point of entry into their system.


In particular I want easy integration with IdentityServer reference and bearer tokens.

Ocelot is a bunch of middlewares in a specific order.


Ocelot manipulates the HttpRequest object into a state specified by its configuration until it reaches a request builder middleware where it creates a HttpRequestMessage object which is used to make a request to a downstream service. The middleware that makes the request is the last thing in the Ocelot pipeline. It does not call the next middleware. The response from the downstream service is stored in a per request scoped repository and retrived as the requests goes back up the Ocelot pipeline. There is a piece of middleware that maps the HttpResponseMessage onto the HttpResponse object and that is returned to the client. That is basically it with a bunch of other features.

The following are configuration that you use when deploying Ocelot.


1.1 Basic Implementation



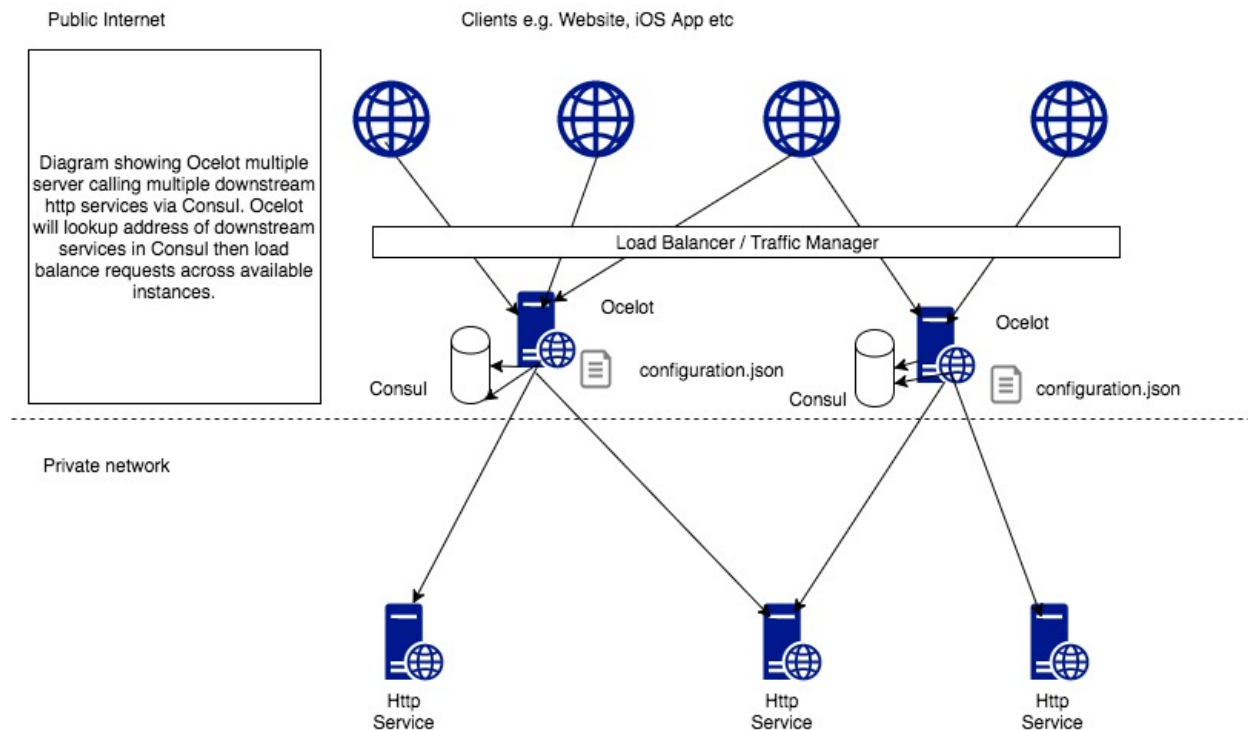
1.2 With IdentityServer




1.3 Multiple Instances



1.4 With Consul



1.5 With Service Fabric



Ocelot is designed to work with .NET Core only and is currently built to netstandard2.0 [this](#) documentation may prove helpful when working out if Ocelot would be suitable for you.

2.1 .NET Core 2.1

Install NuGet package

Install Ocelot and it's dependencies using nuget. You will need to create a netstandard2.0 project and bring the package into it. Then follow the Startup below and *Configuration* sections to get up and running.

```
Install-Package Ocelot
```

All versions can be found [here](#).

Configuration

The following is a very basic ocelot.json. It won't do anything but should get Ocelot starting.

```
{
  "ReRoutes": [],
  "GlobalConfiguration": {
    "BaseUrl": "https://api.mybusiness.com"
  }
}
```

The most important thing to note here is BaseUrl. Ocelot needs to know the URL it is running under in order to do Header find & replace and for certain administration configurations. When setting this URL it should be the external URL that clients will see Ocelot running on e.g. If you are running containers Ocelot might run on the url <http://123.12.1.1:6543> but has something like nginx in front of it responding on <https://api.mybusiness.com>. In this case the Ocelot base url should be <https://api.mybusiness.com>.

If for some reason you are using containers and do want Ocelot to respond to client on <http://123.12.1.1:6543> then you can do this but if you are deploying multiple Ocelot's you will probably want to pass this on the command line in some kind of script. Hopefully whatever scheduler you are using can pass the IP.

Program

Then in your Program.cs you will want to have the following. The main things to note are AddOcelot() (adds ocelot services), UseOcelot().Wait() (sets up all the Ocelot middleware).

```
public class Program
{
    public static void Main(string[] args)
    {
        new WebHostBuilder()
            .UseKestrel()
            .UseContentRoot(Directory.GetCurrentDirectory())
            .ConfigureAppConfiguration((hostingContext, config) =>
            {
                config
                    .SetBasePath(hostingContext.HostingEnvironment.ContentRootPath)
                    .AddJsonFile("appsettings.json", true, true)
                    .AddJsonFile($"appsettings.{hostingContext.HostingEnvironment.
↪EnvironmentName}.json", true, true)
                    .AddJsonFile("ocelot.json")
                    .AddEnvironmentVariables();
            })
            .ConfigureServices(s => {
                s.AddOcelot();
            })
            .ConfigureLogging((hostingContext, logging) =>
            {
                //add your logging
            })
            .UseIISIntegration()
            .Configure(app =>
            {
                app.UseOcelot().Wait();
            })
            .Build()
            .Run();
    }
}
```

2.2 .NET Core 1.0

Install NuGet package

Install Ocelot and its dependencies using nuget. You will need to create a netcoreapp1.0+ project and bring the package into it. Then follow the Startup below and *Configuration* sections to get up and running. Please note you will need to choose one of the Ocelot packages from the NuGet feed.

All versions can be found [here](#).

Configuration

The following is a very basic ocelot.json. It won't do anything but should get Ocelot starting.

```
{
  "ReRoutes": [],
  "GlobalConfiguration": {}
}
```

Program

Then in your Program.cs you will want to have the following.

```
public class Program
{
    public static void Main(string[] args)
    {
        IWebHostBuilder builder = new WebHostBuilder();

        builder.ConfigureServices(s => {
        });

        builder.UseKestrel()
            .UseContentRoot(Directory.GetCurrentDirectory())
            .UseStartup<Startup>();

        var host = builder.Build();

        host.Run();
    }
}
```

Startup

An example startup using a json file for configuration can be seen below.

```
public class Startup
{
    public Startup(IHostingEnvironment env)
    {
        var builder = new ConfigurationBuilder()
            .SetBasePath(env.ContentRootPath)
            .AddJsonFile("appsettings.json", optional: true, reloadOnChange: true)
            .AddJsonFile($"appsettings.{env.EnvironmentName}.json", optional: true)
            .AddJsonFile("ocelot.json")
            .AddEnvironmentVariables();

        Configuration = builder.Build();
    }

    public IConfigurationRoot Configuration { get; }

    public void ConfigureServices(IServiceCollection services)
    {
        services.AddOcelot(Configuration);
    }

    public void Configure(IApplicationBuilder app)
    {
        app.UseOcelot().Wait();
    }
}
```

This is pretty much all you need to get going.

CHAPTER 3

Contributing

Pull requests, issues and commentary welcome! No special process just create a request and get in touch either via gitter or create an issue.

Not Supported

Ocelot does not support...

- Chunked Encoding - Ocelot will always get the body size and return Content-Length header. Sorry if this doesn't work for your use case!
- Forwarding a host header - The host header that you send to Ocelot will not be forwarded to the downstream service. Obviously this would break everything :(
- Swagger - I have looked multiple times at building swagger.json out of the Ocelot ocelot.json but it doesn't fit into the vision

I have for Ocelot. If you would like to have Swagger in Ocelot then you must roll your own swagger.json and do the following in your Startup.cs or Program.cs. The code sample below registers a piece of middleware that loads your hand rolled swagger.json and returns it on /swagger/v1/swagger.json. It then registers the SwaggerUI middleware from Swashbuckle.AspNetCore

```
app.Map("/swagger/v1/swagger.json", b =>
{
    b.Run(async x => {
        var json = File.ReadAllText("swagger.json");
        await x.Response.WriteAsync(json);
    });
});
app.UseSwaggerUI(c =>
{
    c.SwaggerEndpoint("/swagger/v1/swagger.json", "Ocelot");
});
app.UseOcelot().Wait();
```

The main reasons why I don't think Swagger makes sense is we already hand roll our definition in ocelot.json. If we want people developing against Ocelot to be able to see what routes are available then either share the ocelot.json with them (This should be as easy as granting access to a repo etc) or use the Ocelot administration API so that they can query Ocelot for the configuration.

In addition to this many people will configure Ocelot to proxy all traffic like `/products/{everything}` to there product service and you would not be describing what is actually available if you parsed this and turned it into a Swagger path. Also Ocelot has no concept of the models that the downstream services can return and linking to the above problem the same endpoint can return multiple models. Ocelot does not know what models might be used in POST, PUT etc so it all gets a bit messy and finally the Swashbuckle package doesnt reload swagger.json if it changes during runtime. Ocelot's configuration can change during runtime so the Swagger and Ocelot information would not match. Unless I rolled my own Swagger implementation.

If the user wants something to easily test against the Ocelot API then I suggest using Postman as a simple way to do this. It might even be possible to write something that maps ocelot.json to the postman json spec. However I don't intend to do this.

CHAPTER 5

Configuration

An example configuration can be found [here](#). There are two sections to the configuration. An array of ReRoutes and a GlobalConfiguration. The ReRoutes are the objects that tell Ocelot how to treat an upstream request. The Global configuration is a bit hacky and allows overrides of ReRoute specific settings. It's useful if you don't want to manage lots of ReRoute specific settings.

```
{
  "ReRoutes": [],
  "GlobalConfiguration": {}
}
```

Here is an example ReRoute configuration, You don't need to set all of these things but this is everything that is available at the moment:

```
{
  "DownstreamPathTemplate": "/",
  "UpstreamPathTemplate": "/",
  "UpstreamHttpMethod": [
    "Get"
  ],
  "AddHeadersToRequest": {},
  "AddClaimsToRequest": {},
  "RouteClaimsRequirement": {},
  "AddQueriesToRequest": {},
  "RequestIdKey": "",
  "FileCacheOptions": {
    "TtlSeconds": 0,
    "Region": ""
  },
  "ReRouteIsCaseSensitive": false,
  "ServiceName": "",
  "DownstreamScheme": "http",
  "DownstreamHostAndPorts": [
    {
      "Host": "localhost",
```

(continues on next page)

```

        "Port": 51876,
    }
  ],
  "QoSOptions": {
    "ExceptionsAllowedBeforeBreaking": 0,
    "DurationOfBreak": 0,
    "TimeoutValue": 0
  },
  "LoadBalancer": "",
  "RateLimitOptions": {
    "ClientWhitelist": [],
    "EnableRateLimiting": false,
    "Period": "",
    "PeriodTimespan": 0,
    "Limit": 0
  },
  "AuthenticationOptions": {
    "AuthenticationProviderKey": "",
    "AllowedScopes": []
  },
  "HttpHandlerOptions": {
    "AllowAutoRedirect": true,
    "UseCookieContainer": true,
    "UseTracing": true
  },
  "DangerousAcceptAnyServerCertificateValidator": false
}

```

More information on how to use these options is below..

5.1 Multiple environments

Like any other asp.net core project Ocelot supports configuration file names such as configuration.dev.json, configuration.test.json etc. In order to implement this add the following to you

```

.ConfigureAppConfiguration((hostingContext, config) =>
{
    config
        .SetBasePath(hostingContext.HostingEnvironment.ContentRootPath)
        .AddJsonFile("appsettings.json", true, true)
        .AddJsonFile($"appsettings.{hostingContext.HostingEnvironment.
↵EnvironmentName}.json", true, true)
        .AddJsonFile("ocelot.json")
        .AddJsonFile($"configuration.{hostingContext.HostingEnvironment.
↵EnvironmentName}.json")
        .AddEnvironmentVariables();
})

```

Ocelot will now use the environment specific configuration and fall back to ocelot.json if there isnt one.

You also need to set the corresponding environment variable which is ASPNETCORE_ENVIRONMENT. More info on this can be found in the [asp.net core docs](#).

5.2 Merging configuration files

This feature was requested in [Issue 296](#) and allows users to have multiple configuration files to make managing large configurations easier.

Instead of adding the configuration directly e.g. `AddJsonFile("ocelot.json")` you can call `AddOcelot()` like below.

```
.ConfigureAppConfiguration((hostingContext, config) =>
{
    config
        .SetBasePath(hostingContext.HostingEnvironment.ContentRootPath)
        .AddJsonFile("appsettings.json", true, true)
        .AddJsonFile($"appsettings.{hostingContext.HostingEnvironment.
↪EnvironmentName}.json", true, true)
        .AddOcelot(hostingContext.HostingEnvironment)
        .AddEnvironmentVariables();
})
```

In this scenario Ocelot will look for any files that match the pattern `(?)ocelot.([a-zA-Z0-9]*).json` and then merge these together. If you want to set the `GlobalConfiguration` property you must have a file called `ocelot.global.json`.

The way Ocelot merges the files is basically load them, loop over them, add any `ReRoutes`, add any `AggregateReRoutes` and if the file is called `ocelot.global.json` add the `GlobalConfiguration` aswell as any `ReRoutes` or `AggregateReRoutes`. Ocelot will then save the merged configuration to a file called `ocelot.json` and this will be used as the source of truth while ocelot is running.

At the moment there is no validation at this stage it only happens when Ocelot validates the final merged configuration. This is something to be aware of when you are investigating problems. I would advise always checking what is in `ocelot.json` if you have any problems.

You can also give Ocelot a specific path to look in for the configuration files like below.

```
.ConfigureAppConfiguration((hostingContext, config) =>
{
    config
        .SetBasePath(hostingContext.HostingEnvironment.ContentRootPath)
        .AddJsonFile("appsettings.json", true, true)
        .AddJsonFile($"appsettings.{hostingContext.HostingEnvironment.
↪EnvironmentName}.json", true, true)
        .AddOcelot("/foo/bar", hostingContext.HostingEnvironment)
        .AddEnvironmentVariables();
})
```

Ocelot needs the `HostingEnvironment` so it know's to exclude anything environment specific from the algorithm.

5.3 Store configuration in consul

The first thing you need to do is install the NuGet package that provides Consul support in Ocelot.

```
Install-Package Ocelot.Provider.Consul
```

Then you add the following when you register your services Ocelot will attempt to store and retrieve its configuration in consul KV store.

```
services
    .AddOcelot()
```

(continues on next page)

(continued from previous page)

```
.AddConsul()  
.AddConfigStoredInConsul();
```

You also need to add the following to your ocelot.json. This is how Ocelot finds your Consul agent and interacts to load and store the configuration from Consul.

```
"GlobalConfiguration": {  
  "ServiceDiscoveryProvider": {  
    "Host": "localhost",  
    "Port": 9500  
  }  
}
```

I decided to create this feature after working on the raft consensus algorithm and finding out its super hard. Why not take advantage of the fact Consul already gives you this! I guess it means if you want to use Ocelot to its fullest you take on Consul as a dependency for now.

This feature has a 3 second ttl cache before making a new request to your local consul agent.

5.4 Reload JSON config on change

Ocelot supports reloading the json configuration file on change. e.g. the following will recreate Ocelots internal configuration when the ocelot.json file is updated manually.

```
config.AddJsonFile("ocelot.json", optional: false, reloadOnChange: true);
```

5.4.1 Configuration Key

If you are using Consul for configuration (or other providers in the future) you might want to key your configurations so you can have multiple configurations :) This feature was requested in [issue 346!](#) In order to specify the key you need to set the ConfigurationKey property in the ServiceDiscoveryProvider section of the configuration json file e.g.

```
"GlobalConfiguration": {  
  "ServiceDiscoveryProvider": {  
    "Host": "localhost",  
    "Port": 9500,  
    "ConfigurationKey": "Oceolot_A"  
  }  
}
```

In this example Ocelot will use Oceolot_A as the key for your configuration when looking it up in Consul.

If you do not set the ConfigurationKey Ocelot will use the string InternalConfiguration as the key.

5.5 Follow Redirects / Use CookieContainer

Use `HttpHandlerOptions` in `ReRoute` configuration to set up `HttpHandler` behavior:

1. `AllowAutoRedirect` is a value that indicates whether the request should follow redirection responses. Set it true if the request should automatically follow redirection responses from the Downstream resource; otherwise false. The default value is false.
2. `UseCookieContainer` is a value that indicates whether the handler uses the `CookieContainer`

property to store server cookies and uses these cookies when sending requests. The default value is false. Please note that if you are using the CookieContainer Ocelot caches the HttpClient for each downstream service. This means that all requests to that DownstreamService will share the same cookies. [Issue 274](#) was created because a user noticed that the cookies were being shared. I tried to think of a nice way to handle this but I think it is impossible. If you don't cache the clients that means each request gets a new client and therefore a new cookie container. If you clear the cookies from the cached client container you get race conditions due to inflight requests. This would also mean that subsequent requests don't use the cookies from the previous response! All in all not a great situation. I would avoid setting UseCookieContainer to true unless you have a really really good reason. Just look at your response headers and forward the cookies back with your next request!

5.6 SSL Errors

If you want to ignore SSL warnings / errors set the following in your ReRoute config.

```
"DangerousAcceptAnyServerCertificateValidator": false
```

I don't recommend doing this, I suggest creating your own certificate and then getting it trusted by your local / remote machine if you can.

Ocelot's primary functionality is to take incoming http requests and forward them on to a downstream service. Ocelot currently only supports this in the form of another http request (in the future this could be any transport mechanism).

Ocelot's describes the routing of one request to another as a ReRoute. In order to get anything working in Ocelot you need to set up a ReRoute in the configuration.

```
{
  "ReRoutes": [
  ]
}
```

To configure a ReRoute you need to add one to the ReRoutes json array.

```
{
  "DownstreamPathTemplate": "/api/posts/{postId}",
  "DownstreamScheme": "https",
  "DownstreamHostAndPorts": [
    {
      "Host": "localhost",
      "Port": 80,
    }
  ],
  "UpstreamPathTemplate": "/posts/{postId}",
  "UpstreamHttpMethod": [ "Put", "Delete" ]
}
```

The DownstreamPathTemplate, DownstreamScheme and DownstreamHostAndPorts define the URL that a request will be forwarded to.

DownstreamHostAndPorts is a collection that defines the host and port of any downstream services that you wish to forward requests to. Usually this will just contain a single entry but sometimes you might want to load balance requests to your downstream services and Ocelot allows you add more than one entry and then select a load balancer.

The UpstreamPathTemplate is the URL that Ocelot will use to identify which DownstreamPathTemplate to use for a given request. The UpstreamHttpMethod is used so Ocelot can distinguish between requests with different HTTP

verbs to the same URL. You can set a specific list of HTTP Methods or set an empty list to allow any of them.

In Ocelot you can add placeholders for variables to your Templates in the form of {something}. The placeholder variable needs to be present in both the DownstreamPathTemplate and UpstreamPathTemplate properties. When it is Ocelot will attempt to substitute the value in the UpstreamPathTemplate placeholder into the DownstreamPathTemplate for each request Ocelot processes.

You can also do a catch all type of ReRoute e.g.

```
{
  "DownstreamPathTemplate": "/api/{everything}",
  "DownstreamScheme": "https",
  "DownstreamHostAndPorts": [
    {
      "Host": "localhost",
      "Port": 80,
    }
  ],
  "UpstreamPathTemplate":("/{everything})",
  "UpstreamHttpMethod": [ "Get", "Post" ]
}
```

This will forward any path + query string combinations to the downstream service after the path /api.

The default ReRouting configuration is case insensitive!

In order to change this you can specify on a per ReRoute basis the following setting.

```
"ReRouteIsCaseSensitive": true
```

This means that when Ocelot tries to match the incoming upstream url with an upstream template the evaluation will be case sensitive.

6.1 Catch All

Ocelot's routing also supports a catch all style routing where the user can specify that they want to match all traffic.

If you set up your config like below, all requests will be proxied straight through. The placeholder {url} name is not significant, any name will work.

```
{
  "DownstreamPathTemplate":("/{url})",
  "DownstreamScheme": "https",
  "DownstreamHostAndPorts": [
    {
      "Host": "localhost",
      "Port": 80,
    }
  ],
  "UpstreamPathTemplate":("/{url})",
  "UpstreamHttpMethod": [ "Get" ]
}
```

The catch all has a lower priority than any other ReRoute. If you also have the ReRoute below in your config then Ocelot would match it before the catch all.

```
{
  "DownstreamPathTemplate": "/",
  "DownstreamScheme": "https",
  "DownstreamHostAndPorts": [
    {
      "Host": "10.0.10.1",
      "Port": 80,
    }
  ],
  "UpstreamPathTemplate": "/",
  "UpstreamHttpMethod": [ "Get" ]
}
```

6.2 Upstream Host

This feature allows you to have ReRoutes based on the upstream host. This works by looking at the host header the client has used and then using this as part of the information we use to identify a ReRoute.

In order to use this feature please add the following to your config.

```
{
  "DownstreamPathTemplate": "/",
  "DownstreamScheme": "https",
  "DownstreamHostAndPorts": [
    {
      "Host": "10.0.10.1",
      "Port": 80,
    }
  ],
  "UpstreamPathTemplate": "/",
  "UpstreamHttpMethod": [ "Get" ],
  "UpstreamHost": "somedomain.com"
}
```

The ReRoute above will only be matched when the host header value is somedomain.com.

If you do not set UpstreamHost on a ReRoute then any host header will match it. This means that if you have two ReRoutes that are the same, apart from the UpstreamHost, where one is null and the other set Ocelot will favour the one that has been set.

This feature was requested as part of [Issue 216](#).

6.3 Priority

You can define the order you want your ReRoutes to match the Upstream HttpRequest by including a “Priority” property in ocelot.json See [Issue 270](#) for reference

```
{
  "Priority": 0
}
```

0 is the lowest priority, Ocelot will always use 0 for `/{catchAll}` ReRoutes and this is still hardcoded. After that you are free to set any priority you wish.

e.g. you could have

```
{
  "UpstreamPathTemplate": "/goods/{catchAll}"
  "Priority": 0
}
```

and

```
{
  "UpstreamPathTemplate": "/goods/delete"
  "Priority": 1
}
```

In the example above if you make a request into Ocelot on `/goods/delete` Ocelot will match `/goods/delete` `ReRoute`. Previously it would have matched `/goods/{catchAll}` (because this is the first `ReRoute` in the list!).

6.4 Dynamic Routing

This feature was requested in [issue 340](#).

The idea is to enable dynamic routing when using a service discovery provider so you don't have to provide the `ReRoute` config. See the docs [service-discovery](#) if this sounds interesting to you.

6.5 Query Strings

Ocelot allows you to specify a querystring as part of the `DownstreamPathTemplate` like the example below.

```
{
  "ReRoutes": [
    {
      "DownstreamPathTemplate": "/api/subscriptions/{subscriptionId}/updates?
↪unitId={unitId}",
      "UpstreamPathTemplate": "/api/units/{subscriptionId}/{unitId}/updates",
      "UpstreamHttpMethod": [
        "Get"
      ],
      "DownstreamScheme": "http",
      "DownstreamHostAndPorts": [
        {
          "Host": "localhost",
          "Port": 50110
        }
      ]
    }
  ],
  "GlobalConfiguration": {
  }
}
```

In this example Ocelot will use the value from the `{unitId}` in the upstream path template and add it to the downstream request as a query string parameter called `unitId`!

Ocelot will also allow you to put query string parameters in the `UpstreamPathTemplate` so you can match certain queries to certain services.

```
{
  "ReRoutes": [
    {
      "DownstreamPathTemplate": "/api/units/{subscriptionId}/{unitId}/updates",
      "UpstreamPathTemplate": "/api/subscriptions/{subscriptionId}/updates?
↔unitId={unitId}",
      "UpstreamHttpMethod": [
        "Get"
      ],
      "DownstreamScheme": "http",
      "DownstreamHostAndPorts": [
        {
          "Host": "localhost",
          "Port": 50110
        }
      ]
    }
  ],
  "GlobalConfiguration": {
  }
}
```

In this example Ocelot will only match requests that have a matching url path and the querystring starts with unitId=something. You can have other queries after this but you must start with the matching parameter. Also Ocelot will swap the {unitId} parameter from the query string and use it in the downstream request path.

Request Aggregation

Ocelot allow's you to specify Aggregate ReRoutes that compose multiple normal ReRoutes and map their responses into one object. This is usual where you have a client that is making multiple requests to a server where it could just be one. This feature allows you to start implementing back end for a front end type architecture with Ocelot.

This feature was requested as part of [Issue 79](#) and further improvements were made as part of [Issue 298](#).

In order to set this up you must do something like the following in your ocelot.json. Here we have specified two normal ReRoutes and each one has a Key property. We then specify an Aggregate that composes the two ReRoutes using their keys in the ReRouteKeys list and says then we have the UpstreamPathTemplate which works like a normal ReRoute. Obviously you cannot have duplicate UpstreamPathTemplates between ReRoutes and Aggregates. You can use all of Ocelot's normal ReRoute options apart from RequestIdKey (explained in gotchas below).

7.1 Advanced register your own Aggregators

Ocelot started with just the basic request aggregation and since then we have added a more advanced method that let's the user take in the responses from the downstream services and then aggregate them into a response object.

The ocelot.json setup is pretty much the same as the basic aggregation approach apart from you need to add an Aggregator property like below.

```
{
  "ReRoutes": [
    {
      "DownstreamPathTemplate": "/",
      "UpstreamPathTemplate": "/laura",
      "UpstreamHttpMethod": [
        "Get"
      ],
      "DownstreamScheme": "http",
      "DownstreamHostAndPorts": [
        {
          "Host": "localhost",
```

(continues on next page)

(continued from previous page)

```

        "Port": 51881
      }
    ],
    "Key": "Laura"
  },
  {
    "DownstreamPathTemplate": "/",
    "UpstreamPathTemplate": "/tom",
    "UpstreamHttpMethod": [
      "Get"
    ],
    "DownstreamScheme": "http",
    "DownstreamHostAndPorts": [
      {
        "Host": "localhost",
        "Port": 51882
      }
    ],
    "Key": "Tom"
  }
],
"Aggregates": [
  {
    "ReRouteKeys": [
      "Tom",
      "Laura"
    ],
    "UpstreamPathTemplate": "/",
    "Aggregator": "FakeDefinedAggregator"
  }
]
}

```

Here we have added an aggregator called FakeDefinedAggregator. Ocelot is going to look for this aggregator when it tries to aggregate this ReRoute.

In order to make the aggregator available we must add the FakeDefinedAggregator to the OcelotBuilder like below.

```

services
    .AddOcelot ()
    .AddSingletonDefinedAggregator<FakeDefinedAggregator> ();

```

Now when Ocelot tries to aggregate the ReRoute above it will find the FakeDefinedAggregator in the container and use it to aggregate the ReRoute. Because the FakeDefinedAggregator is registered in the container you can add any dependencies it needs into the container like below.

```

services.AddSingleton<FooDependency> ();

services
    .AddOcelot ()
    .AddSingletonDefinedAggregator<FooAggregator> ();

```

In this example FooAggregator takes a dependency on FooDependency and it will be resolved by the container.

In addition to this Ocelot lets you add transient aggregators like below.

```
services
    .AddOcelot ()
    .AddTransientDefinedAggregator<FakeDefinedAggregator> ();
```

In order to make an Aggregator you must implement this interface.

```
public interface IDefinedAggregator
{
    Task<DownstreamResponse> Aggregate(List<DownstreamResponse> responses);
}
```

With this feature you can pretty much do whatever you want because DownstreamResponse contains Content, Headers and Status Code. We can add extra things if needed just raise an issue on GitHub. Please note if the HttpClient throws an exception when making a request to a ReRoute in the aggregate then you will not get a DownstreamResponse for it but you would for any that succeed. If it does throw an exception this will be logged.

7.2 Basic expecting JSON from Downstream Services

```
{
  "ReRoutes": [
    {
      "DownstreamPathTemplate": "/",
      "UpstreamPathTemplate": "/laura",
      "UpstreamHttpMethod": [
        "Get"
      ],
      "DownstreamScheme": "http",
      "DownstreamHostAndPorts": [
        {
          "Host": "localhost",
          "Port": 51881
        }
      ],
      "Key": "Laura"
    },
    {
      "DownstreamPathTemplate": "/",
      "UpstreamPathTemplate": "/tom",
      "UpstreamHttpMethod": [
        "Get"
      ],
      "DownstreamScheme": "http",
      "DownstreamHostAndPorts": [
        {
          "Host": "localhost",
          "Port": 51882
        }
      ],
      "Key": "Tom"
    }
  ],
  "Aggregates": [
    {
      "ReRouteKeys": [
        "Tom",
```

(continues on next page)

(continued from previous page)

```
        "Laura"
      ],
      "UpstreamPathTemplate": "/"
    }
  ]
}
```

You can also set `UpstreamHost` and `ReRouteIsCaseSensitive` in the `Aggregate` configuration. These behave the same as any other `ReRoutes`.

If the `ReRoute /tom` returned a body of `{"Age": 19}` and `/laura` returned `{"Age": 25}` the the response after aggregation would be as follows.

```
{"Tom": {"Age": 19}, "Laura": {"Age": 25}}
```

At the moment the aggregation is very simple. Ocelot just gets the response from your downstream service and sticks it into a json dictionary as above. With the `ReRoute` key being the key of the dictionary and the value the response body from your downstream service. You can see that the object is just JSON without any pretty spaces etc.

All headers will be lost from the downstream services response.

Ocelot will always return content type `application/json` with an aggregate request.

If you downstream services return a 404 the aggregate will just return nothing for that downstream service. It will not change the aggregate response into a 404 even if all the downstreams return a 404.

7.2.1 Gotcha's / Further info

You cannot use `ReRoutes` with specific `RequestIdKeys` as this would be crazy complicated to track.

Aggregation only supports the `GET` HTTP Verb.

CHAPTER 8

GraphQL

OK you got me Ocelot doesn't directly support GraphQL but so many people have asked about it I wanted to show how easy it is to integrate the [graphql-dotnet](#) library.

Please see the sample project [OcelotGraphQL](#). Using a combination of the [graphql-dotnet](#) project and Ocelot's DelegatingHandler features this is pretty easy to do. However I do not intend to integrate more closely with GraphQL at the moment. Check out the samples readme and that should give you enough instruction on how to do this!

Good luck and have fun :>

Service Discovery

Ocelot allows you to specify a service discovery provider and will use this to find the host and port for the downstream service Ocelot is forwarding a request to. At the moment this is only supported in the GlobalConfiguration section which means the same service discovery provider will be used for all ReRoutes you specify a ServiceName for at ReRoute level.

9.1 Consul

The first thing you need to do is install the NuGet package that provides Consul support in Ocelot.

```
Install-Package Ocelot.Provider.Consul
```

Then add the following to your ConfigureServices method.

```
s.AddOcelot ()  
    .AddConsul ();
```

The following is required in the GlobalConfiguration. The Provider is required and if you do not specify a host and port the Consul default will be used.

```
"ServiceDiscoveryProvider": {  
  "Host": "localhost",  
  "Port": 8500,  
  "Type": "Consul"  
}
```

In the future we can add a feature that allows ReRoute specific configuration.

In order to tell Ocelot a ReRoute is to use the service discovery provider for its host and port you must add the ServiceName and load balancer you wish to use when making requests downstream. At the moment Ocelot has a RoundRobin and LeastConnection algorithm you can use. If no load balancer is specified Ocelot will not load balance requests.

```
{
  "DownstreamPathTemplate": "/api/posts/{postId}",
  "DownstreamScheme": "https",
  "UpstreamPathTemplate": "/posts/{postId}",
  "UpstreamHttpMethod": [ "Put" ],
  "ServiceName": "product",
  "LoadBalancerOptions": {
    "Type": "LeastConnection"
  },
}
```

When this is set up Ocelot will lookup the downstream host and port from the service discover provider and load balance requests across any available services.

A lot of people have asked me to implement a feature where Ocelot polls consul for latest service information rather than per request. If you want to poll consul for the latest services rather than per request (default behaviour) then you need to set the following configuration.

```
"ServiceDiscoveryProvider": {
  "Host": "localhost",
  "Port": 8500,
  "Type": "PollConsul",
  "PollingInterval": 100
}
```

The polling interval is in milliseconds and tells Ocelot how often to call Consul for changes in service configuration.

Please note there are tradeoffs here. If you poll Consul it is possible Ocelot will not know if a service is down depending on your polling interval and you might get more errors than if you get the latest services per request. This really depends on how volatile your services are. I doubt it will matter for most people and polling may give a tiny performance improvement over calling consul per request (as sidecar agent). If you are calling a remote consul agent then polling will be a good performance improvement.

You services need to be added to Consul something like below (c# style but hopefully this make sense)... The only important thing to note is not to add http or https to the Address field. I have been contacted before about not accepting scheme in Address and accepting scheme in address. After reading [this](#) I don't think the scheme should be in there.

Or

```
"Service": {
  "ID": "some-id",
  "Service": "some-service-name",
  "Address": "localhost",
  "Port": 8080
}
```

9.1.1 ACL Token

If you are using ACL with Consul Ocelot supports adding the X-Consul-Token header. In order so this to work you must add the additional property below.

```
"ServiceDiscoveryProvider": {
  "Host": "localhost",
  "Port": 8500,
  "Token": "footoken",
  "Type": "Consul"
}
```


Ocelot will add this token to the consul client that it uses to make requests and that is then used for every request.

9.2 Eureka

This feature was requested as part of [Issue 262](#) . to add support for Netflix's Eureka service discovery provider. The main reason for this is it is a key part of [Steeltoe](#) which is something to do with [Pivotal!](#) Anyway enough of the background.

The first thing you need to do is install the NuGet package that provides Eureka support in Ocelot.

```
Install-Package Ocelot.Provider.Eureka
```

Then add the following to your `ConfigureServices` method.

```
s.AddOcelot()  
    .AddEureka();
```

Then in order to get this working add the following to `ocelot.json`..

```
"ServiceDiscoveryProvider": {  
  "Type": "Eureka"  
}
```

And following the guide [Here](#) you may also need to add some stuff to `appsettings.json`. For example the json below tells the steeltoe / pivotal services where to look for the service discovery server and if the service should register with it.

```
"eureka": {  
  "client": {  
    "serviceUrl": "http://localhost:8761/eureka/",  
    "shouldRegisterWithEureka": false,  
    "shouldFetchRegistry": true  
  }  
}
```

I am told that if `shouldRegisterWithEureka` is false then `shouldFetchRegistry` will default to true so you don't need it explicitly but left it in there.

Ocelot will now register all the necessary services when it starts up and if you have the json above will register itself with Eureka. One of the services polls Eureka every 30 seconds (default) and gets the latest service state and persists this in memory. When Ocelot asks for a given service it is retrieved from memory so performance is not a big problem. Please note that this code is provided by the `Pivotal.Discovery.Client` NuGet package so big thanks to them for all the hard work.

9.3 Dynamic Routing

This feature was requested in [issue 340](#). The idea is to enable dynamic routing when using a service discovery provider (see that section of the docs for more info). In this mode Ocelot will use the first segment of the upstream path to lookup the downstream service with the service discovery provider.

An example of this would be calling ocelot with a url like `https://api.mywebsite.com/product/products`. Ocelot will take the first segment of the path which is `product` and use it as a key to look up the service in consul. If consul returns a service Ocelot will request it on whatever host and port comes back from consul plus the remaining path segments in this case `products` thus making the downstream call `http://hostfromconsul:portfromconsul/products`. Ocelot will append any query string to the downstream url as normal.

In order to enable dynamic routing you need to have 0 ReRoutes in your config. At the moment you cannot mix dynamic and configuration ReRoutes. In addition to this you need to specify the Service Discovery provider details as outlined above and the downstream http/https scheme as DownstreamScheme.

In addition to that you can set RateLimitOptions, QoSOptions, LoadBalancerOptions and HttpHandlerOptions, DownstreamScheme (You might want to call Ocelot on https but talk to private services over http) that will be applied to all of the dynamic ReRoutes.

The config might look something like

```
{
  "ReRoutes": [],
  "Aggregates": [],
  "GlobalConfiguration": {
    "RequestIdKey": null,
    "ServiceDiscoveryProvider": {
      "Host": "localhost",
      "Port": 8500,
      "Type": "Consul",
      "Token": null,
      "ConfigurationKey": null
    },
    "RateLimitOptions": {
      "ClientIdHeader": "ClientId",
      "QuotaExceededMessage": null,
      "RateLimitCounterPrefix": "ocelot",
      "DisableRateLimitHeaders": false,
      "HttpStatusCode": 429
    },
    "QoSOptions": {
      "ExceptionsAllowedBeforeBreaking": 0,
      "DurationOfBreak": 0,
      "TimeoutValue": 0
    },
    "BaseUrl": null,
    "LoadBalancerOptions": {
      "Type": "LeastConnection",
      "Key": null,
      "Expiry": 0
    },
    "DownstreamScheme": "http",
    "HttpHandlerOptions": {
      "AllowAutoRedirect": false,
      "UseCookieContainer": false,
      "UseTracing": false
    }
  }
}
```

Ocelot also allows you to set DynamicReRoutes which lets you set rate limiting rules per downstream service. This is useful if you have for example a product and search service and you want to rate limit one more than the other. An example of this would be as follows.

```
{
  "DynamicReRoutes": [
    {
      "ServiceName": "product",
      "RateLimitRule": {
```

(continues on next page)

(continued from previous page)

```
        "ClientWhitelist": [],
        "EnableRateLimiting": true,
        "Period": "1s",
        "PeriodTimespan": 1000.0,
        "Limit": 3
    }
}
],
"GlobalConfiguration": {
    "RequestIdKey": null,
    "ServiceDiscoveryProvider": {
        "Host": "localhost",
        "Port": 8523,
        "Type": "Consul"
    },
    "RateLimitOptions": {
        "ClientIdHeader": "ClientId",
        "QuotaExceededMessage": "",
        "RateLimitCounterPrefix": "",
        "DisableRateLimitHeaders": false,
        "HttpStatusCode": 428
    }
    "DownstreamScheme": "http",
}
}
```

This configuration means that if you have a request come into Ocelot on `/product/*` then dynamic routing will kick in and ocelot will use the rate limiting set against the product service in the `DynamicReRoutes` section.

Please take a look through all of the docs to understand these options.

CHAPTER 10

Service Fabric

If you have services deployed in Service Fabric you will normally use the naming service to access them.

The following example shows how to set up a ReRoute that will work in Service Fabric. The most important thing is the ServiceName which is made up of the Service Fabric application name then the specific service name. We also need to set up the ServiceDiscoveryProvider in GlobalConfiguration. The example here shows a typical configuration. It assumes service fabric is running on localhost and that the naming service is on port 19081.

The example below is taken from the samples folder so please check it if this doesnt make sense!

```
{
  "ReRoutes": [
    {
      "DownstreamPathTemplate": "/api/values",
      "UpstreamPathTemplate": "/EquipmentInterfaces",
      "UpstreamHttpMethod": [
        "Get"
      ],
      "DownstreamScheme": "http",
      "ServiceName": "OcelotServiceApplication/OcelotApplicationService",
    }
  ],
  "GlobalConfiguration": {
    "RequestIdKey": "OcRequestId",
    "ServiceDiscoveryProvider": {
      "Host": "localhost",
      "Port": 19081,
      "Type": "ServiceFabric"
    }
  }
}
```

If you are using stateless / guest exe services ocelot will be able to proxy through the naming service without anything else. However if you are using statefull / actor services you must send the PartitionKind and PartitionKey query string values with the client request e.g.

GET <http://ocelot.com/EquipmentInterfaces?PartitionKind=xxx&PartitionKey=xxx>

There is no way for Ocelot to work these out for you.

Authentication

In order to authenticate ReRoutes and subsequently use any of Ocelot's claims based features such as authorisation or modifying the request with values from the token. Users must register authentication services in their Startup.cs as usual but they provide a scheme (authentication provider key) with each registration e.g.

```
public void ConfigureServices(IServiceCollection services)
{
    var authenticationProviderKey = "TestKey";

    services.AddAuthentication()
        .AddJwtBearer(authenticationProviderKey, x =>
            {
            });
}
```

In this example TestKey is the scheme that this provider has been registered with. We then map this to a ReRoute in the configuration e.g.

```
"ReRoutes": [{
    "DownstreamHostAndPorts": [
        {
            "Host": "localhost",
            "Port": 51876,
        }
    ],
    "DownstreamPathTemplate": "/",
    "UpstreamPathTemplate": "/",
    "UpstreamHttpMethod": ["Post"],
    "ReRouteIsCaseSensitive": false,
    "DownstreamScheme": "http",
    "AuthenticationOptions": {
        "AuthenticationProviderKey": "TestKey",
        "AllowedScopes": []
    }
}]
```

When Ocelot runs it will look at this `ReRoutes.AuthenticationOptions.AuthenticationProviderKey` and check that there is an Authentication provider registered with the given key. If there isn't then Ocelot will not start up, if there is then the ReRoute will use that provider when it executes.

If a ReRoute is authenticated Ocelot will invoke whatever scheme is associated with it while executing the authentication middleware. If the request fails authentication Ocelot returns a http status code 401.

11.1 JWT Tokens

If you want to authenticate using JWT tokens maybe from a provider like Auth0 you can register your authentication middleware as normal e.g.

```
public void ConfigureServices(IServiceCollection services)
{
    var authenticationProviderKey = "TestKey";

    services.AddAuthentication()
        .AddJwtBearer(authenticationProviderKey, x =>
        {
            x.Authority = "test";
            x.Audience = "test";
        });

    services.AddOcelot();
}
```

Then map the authentication provider key to a ReRoute in your configuration e.g.

```
"ReRoutes": [{
    "DownstreamHostAndPorts": [
        {
            "Host": "localhost",
            "Port": 51876,
        }
    ],
    "DownstreamPathTemplate": "/",
    "UpstreamPathTemplate": "/",
    "UpstreamHttpMethod": ["Post"],
    "ReRouteIsCaseSensitive": false,
    "DownstreamScheme": "http",
    "AuthenticationOptions": {
        "AuthenticationProviderKey": "TestKey",
        "AllowedScopes": []
    }
}]
```

11.2 Identity Server Bearer Tokens

In order to use IdentityServer bearer tokens register your IdentityServer services as usual in `ConfigureServices` with a scheme (key). If you don't understand how to do this please consult the IdentityServer documentation.

```
public void ConfigureServices(IServiceCollection services)
{
```

(continues on next page)

(continued from previous page)

```
var authenticationProviderKey = "TestKey";
Action<IdentityServerAuthenticationOptions> options = o =>
{
    o.Authority = "https://whereyouridentityserverlives.com";
    o.ApiName = "api";
    o.SupportedTokens = SupportedTokens.Both;
    o.ApiSecret = "secret";
};

services.AddAuthentication()
    .AddIdentityServerAuthentication(authenticationProviderKey, options);

services.AddOcelot();
}
```

Then map the authentication provider key to a ReRoute in your configuration e.g.

```
"ReRoutes": [{
  "DownstreamHostAndPorts": [
    {
      "Host": "localhost",
      "Port": 51876,
    }
  ],
  "DownstreamPathTemplate": "/",
  "UpstreamPathTemplate": "/",
  "UpstreamHttpMethod": ["Post"],
  "ReRouteIsCaseSensitive": false,
  "DownstreamScheme": "http",
  "AuthenticationOptions": {
    "AuthenticationProviderKey": "TestKey",
    "AllowedScopes": []
  }
}]
```

11.3 Okta

I have not had time to write this up but we have [Issue 446](#) that contains some code and examples that might help with Okta integration.

11.4 Allowed Scopes

If you add scopes to AllowedScopes Ocelot will get all the user claims (from the token) of the type scope and make sure that the user has all of the scopes in the list.

This is a way to restrict access to a ReRoute on a per scope basis.

Ocelot supports claims based authorisation which is run post authentication. This means if you have a route you want to authorise you can add the following to you ReRoute configuration.

```
"RouteClaimsRequirement": {  
  "UserType": "registered"  
}
```

In this example when the authorisation middleware is called Ocelot will check to see if the user has the claim type UserType and if the value of that claim is registered. If it isn't then the user will not be authorised and the response will be 403 forbidden.

Ocelot supports proxying websockets with some extra bits. This functionality was requested in [Issue 212](#).

In order to get websocket proxying working with Ocelot you need to do the following.

In your Configure method you need to tell your application to use WebSockets.

```
Configure (app =>
{
    app.UseWebSockets ();
    app.UseOcelot ().Wait ();
})
```

Then in your ocelot.json add the following to proxy a ReRoute using websockets.

```
{
  "DownstreamPathTemplate": "/ws",
  "UpstreamPathTemplate": "/",
  "DownstreamScheme": "ws",
  "DownstreamHostAndPorts": [
    {
      "Host": "localhost",
      "Port": 5001
    }
  ],
}
```

With this configuration set Ocelot will match any websocket traffic that comes in on / and proxy it to localhost:5001/ws. To make this clearer Ocelot will receive messages from the upstream client, proxy these to the downstream service, receive messages from the downstream service and proxy these to the upstream client.

13.1 SignalR

Ocelot supports proxying SignalR. This functionality was requested in [Issue 344](#).

In order to get websocket proxying working with Ocelot you need to do the following.

Install Microsoft.AspNetCore.SignalR.Client 1.0.2 you can try other packages but this one is tested.

Do not run it in IISExpress or install the websockets feature in the IIS features

In your Configure method you need to tell your application to use SignalR.

```
Configure(app =>
{
    app.UseWebSockets();
    app.UseOcelot().Wait();
})
```

Then in your ocelot.json add the following to proxy a ReRoute using SignalR. Note normal Ocelot routing rules apply the main thing is the scheme which is set to "ws".

```
{
  "ReRoutes": [
    {
      "DownstreamPathTemplate": "{catchAll}",
      "DownstreamScheme": "ws",
      "DownstreamHostAndPorts": [
        {
          "Host": "localhost",
          "Port": 50000
        }
      ],
      "UpstreamPathTemplate": "/gateway/{catchAll}",
      "UpstreamHttpMethod": [ "GET", "POST", "PUT", "DELETE", "OPTIONS" ]
    }
  ]
}
```

With this configuration set Ocelot will match any SignalR traffic that comes in on / and proxy it to localhost:5001/ws. To make this clearer Ocelot will receive messages from the upstream client, proxy these to the downstream service, receive messages from the downstream service and proxy these to the upstream client.

13.2 Supported

1. Load Balancer
2. Routing
3. Service Discovery

This means that you can set up your downstream services running websockets and either have multiple DownstreamHostAndPorts in your ReRoute config or hook your ReRoute into a service discovery provider and then load balance requests... Which I think is pretty cool :)

13.3 Not Supported

Unfortunately a lot of Ocelot's features are non websocket specific such as header and http client stuff. I've listed what won't work below.

1. Tracing

2. RequestId
3. Request Aggregation
4. Rate Limiting
5. Quality of Service
6. Middleware Injection
7. Header Transformation
8. Delegating Handlers
9. Claims Transformation
10. Caching
11. Authentication - If anyone requests it we might be able to do something with basic authentication.
12. Authorisation

I'm not 100% sure what will happen with this feature when it get's into the wild so please make sure you test thoroughly!

Ocelot supports changing configuration during runtime via an authenticated HTTP API. This can be authenticated in two ways either using Ocelot's internal IdentityServer (for authenticating requests to the administration API only) or hooking the administration API authentication into your own IdentityServer.

The first thing you need to do if you want to use the administration API is bring in the relevant NuGet package..

```
Install-Package Ocelot.Administration
```

This will bring down everything needed by the admin API.

14.1 Providing your own IdentityServer

All you need to do to hook into your own IdentityServer is add the following to your ConfigureServices method.

```
public virtual void ConfigureServices(IServiceCollection services)
{
    Action<IdentityServerAuthenticationOptions> options = o => {
        // o.Authority = ;
        // o.ApiName = ;
        // etc....
    };

    services
        .AddOcelot()
        .AddAdministration("/administration", options);
}
```

You now need to get a token from your IdentityServer and use in subsequent requests to Ocelot's administration API.

This feature was implemented for [issue 228](#). It is useful because the IdentityServer authentication middleware needs the URL of the IdentityServer. If you are using the internal IdentityServer it might not always be possible to have the Ocelot URL.

14.2 Internal IdentityServer

The API is authenticated using bearer tokens that you request from Ocelot itself. This is provided by the amazing [Identity Server](#) project that I have been using for a few years now. Check them out.

In order to enable the administration section you need to do a few things. First of all add this to your initial Startup.cs.

The path can be anything you want and it is obviously recommended don't use a url you would like to route through with Ocelot as this will not work. The administration uses the MapWhen functionality of asp.net core and all requests to {root}/administration will be sent there not to the Ocelot middleware.

The secret is the client secret that Ocelot's internal IdentityServer will use to authenticate requests to the administration API. This can be whatever you want it to be!

```
public virtual void ConfigureServices(IServiceCollection services)
{
    services
        .AddOcelot()
        .AddAdministration("/administration", "secret");
}
```

In order for the administration API to work, Ocelot / IdentityServer must be able to call itself for validation. This means that you need to add the base url of Ocelot to global configuration if it is not default (<http://localhost:5000>). Please note if you are using something like docker to host Ocelot it might not be able to call back to localhost etc and you need to know what you are doing with docker networking in this scenario. Anyway this can be done as follows..

If you want to run on a different host and port locally..

```
"GlobalConfiguration": {
  "BaseUrl": "http://localhost:55580"
}
```

or if Ocelot is exposed via dns

```
"GlobalConfiguration": {
  "BaseUrl": "http://mydns.com"
}
```

Now if you went with the configuration options above and want to access the API you can use the postman scripts called ocelot.postman_collection.json in the solution to change the Ocelot configuration. Obviously these will need to be changed if you are running Ocelot on a different url to <http://localhost:5000>.

The scripts show you how to request a bearer token from ocelot and then use it to GET the existing configuration and POST a configuration.

If you are running multiple Ocelot instances in a cluster then you need to use a certificate to sign the bearer tokens used to access the administration API.

In order to do this you need to add two more environmental variables for each Ocelot in the cluster.

OCELOT_CERTIFICATE The path to a certificate that can be used to sign the tokens. The certificate needs to be of the type X509 and obviously Ocelot needs to be able to access it.

OCELOT_CERTIFICATE_PASSWORD The password for the certificate.

Normally Ocelot just uses temporary signing credentials but if you set these environmental variables then it will use the certificate. If all the other Ocelot instances in the cluster have the same certificate then you are good!

14.3 Administration API

POST {adminPath}/connect/token

This gets a token for use with the admin area using the client credentials we talk about setting above. Under the hood this calls into an IdentityServer hosted within Ocelot.

The body of the request is form-data as follows

`client_id` set as `admin`

`client_secret` set as whatever you used when setting up the administration services.

`scope` set as `admin`

`grant_type` set as `client_credentials`

GET {adminPath}/configuration

This gets the current Ocelot configuration. It is exactly the same JSON we use to set Ocelot up with in the first place.

POST {adminPath}/configuration

This overwrites the existing configuration (should probably be a put!). I recommend getting your config from the GET endpoint, making any changes and posting it back... simples.

The body of the request is JSON and it is the same format as the `FileConfiguration.cs` that we use to set up Ocelot on a file system.

Please note that if you want to use this API then the process running Ocelot must have permission to write to the disk where your `ocelot.json` or `ocelot.{environment}.json` is located. This is because Ocelot will overwrite them on save.

DELETE {adminPath}/outputcache/{region}

This clears a region of the cache. If you are using a backplane it will clear all instances of the cache! Giving your the ability to run a cluster of Ocelots and cache over all of them in memory and clear them all at the same time / just use a distributed cache.

The region is whatever you set against the `Region` field in the `FileCacheOptions` section of the Ocelot configuration.

Rate Limiting

Thanks to [@catcherwong](#) article for inspiring me to finally write this documentation.

Ocelot supports rate limiting of upstream requests so that your downstream services do not become overloaded. This feature was added by [@geffzhang](#) on GitHub! Thanks very much.

OK so to get rate limiting working for a ReRoute you need to add the following json to it.

```
"RateLimitOptions": {
  "ClientWhitelist": [],
  "EnableRateLimiting": true,
  "Period": "1s",
  "PeriodTimespan": 1,
  "Limit": 1
}
```

ClientWhitelist - This is an array that contains the whitelist of the client. It means that the client in this array will not be affected by the rate limiting. **EnableRateLimiting** - This value specifies enable endpoint rate limiting. **Period** - This value specifies the period that the limit applies to, such as 1s, 5m, 1h,1d and so on. If you make more requests in the period than the limit allows then you need to wait for **PeriodTimespan** to elapse before you make another request. **PeriodTimespan** - This value specifies that we can retry after a certain number of seconds. **Limit** - This value specifies the maximum number of requests that a client can make in a defined period.

You can also set the following in the **GlobalConfiguration** part of ocelot.json

```
"RateLimitOptions": {
  "DisableRateLimitHeaders": false,
  "QuotaExceededMessage": "Customize Tips!",
  "HttpStatusCode": 999,
  "ClientIdHeader" : "Test "
}
```

DisableRateLimitHeaders - This value specifies whether X-Rate-Limit and Rety-After headers are disabled. **QuotaExceededMessage** - This value specifies the exceeded message. **HttpStatusCode** - This value specifies the returned HTTP Status code when rate limiting occurs. **ClientIdHeader** - Allows you to specify the header that should be used to identify clients. By default it is "ClientId"

Ocelot supports some very rudimentary caching at the moment provided by the [CacheManager](#) project. This is an amazing project that is solving a lot of caching problems. I would recommend using this package to cache with Ocelot. The following example shows how to add CacheManager to Ocelot so that you can do output caching.

First of all add the following NuGet package.

```
Install-Package Ocelot.Cache.CacheManager
```

This will give you access to the Ocelot cache manager extension methods.

The second thing you need to do something like the following to your ConfigureServices..

```
s.AddOcelot()
    .AddCacheManager(x =>
    {
        x.WithDictionaryHandle();
    })
```

Finally in order to use caching on a route in your ReRoute configuration add this setting.

```
"FileCacheOptions": { "TtlSeconds": 15, "Region": "somename" }
```

In this example ttl seconds is set to 15 which means the cache will expire after 15 seconds.

If you look at the example [here](#) you can see how the cache manager is setup and then passed into the Ocelot AddCacheManager configuration method. You can use any settings supported by the CacheManager package and just pass them in.

Anyway Ocelot currently supports caching on the URL of the downstream service and setting a TTL in seconds to expire the cache. You can also clear the cache for a region by calling Ocelot's administration API.

16.1 Your own caching

If you want to add your own caching method implement the following interfaces and register them in DI e.g. `services.AddSingleton<IOcelotCache<CachedResponse>, MyCache>()`

`IOcelotCache<CachedResponse>` this is for output caching.

`IOcelotCache<FileConfiguration>` this is for caching the file configuration if you are calling something remote to get your config such as Consul.

Please dig into the Ocelot source code to find more. I would really appreciate it if anyone wants to implement Redis, memcache etc..

CHAPTER 17

Quality of Service

Ocelot supports one QoS capability at the current time. You can set on a per ReRoute basis if you want to use a circuit breaker when making requests to a downstream service. This uses the an awesome .NET library called Polly check them out [here](#).

The first thing you need to do if you want to use the administration API is bring in the relevant NuGet package..

```
Install-Package Ocelot.Provider.Polly
```

Then in your ConfigureServices method

```
public virtual void ConfigureServices(IServiceCollection services)
{
    services
        .AddOcelot()
        .AddPolly();
}
```

Then add the following section to a ReRoute configuration.

```
"QoSOptions": {
  "ExceptionsAllowedBeforeBreaking":3,
  "DurationOfBreak":5,
  "TimeoutValue":5000
}
```

You must set a number greater than 0 against ExceptionsAllowedBeforeBreaking for this rule to be implemented. Duration of break is how long the circuit breaker will stay open for after it is tripped. TimeoutValue means if a request takes more than 5 seconds it will automatically be timed out.

You can set the TimeoutValue in isolation of the ExceptionsAllowedBeforeBreaking and DurationOfBreak options.

```
"QoSOptions": {
  "TimeoutValue":5000
}
```

There is no point setting the other two in isolation as they affect each other :)

If you do not add a QoS section QoS will not be used however Ocelot will default to a 90 second timeout on all downstream requests. If someone needs this to be configurable open an issue.

Headers Transformation

Ocelot allows the user to transform headers pre and post downstream request. At the moment Ocelot only supports find and replace. This feature was requested [GitHub #190](#) and I decided that it was going to be useful in various ways.

18.1 Add to Request

This feature was requested in [GitHub #313](#).

If you want to add a header to your upstream request please add the following to a ReRoute in your ocelot.json:

```
"UpstreamHeaderTransform": {  
  "Uncle": "Bob"  
}
```

In the example above a header with the key Uncle and value Bob would be sent to the upstream service.

Placeholders are supported too (see below).

18.2 Add to Response

This feature was requested in [GitHub #280](#).

If you want to add a header to your downstream response please add the following to a ReRoute in ocelot.json..

```
"DownstreamHeaderTransform": {  
  "Uncle": "Bob"  
},
```

In the example above a header with the key Uncle and value Bob would be returned by Ocelot when requesting the specific ReRoute.

If you want to return the Butterfly APM trace id then do something like the following..

```
"DownstreamHeaderTransform": {  
  "AnyKey": "{TraceId}"  
},
```

18.3 Find and Replace

In order to transform a header first we specify the header key and then the type of transform we want e.g.

```
"Test": "http://www.bbc.co.uk/, http://ocelot.com/"
```

The key is “Test” and the value is “<http://www.bbc.co.uk/>, <http://ocelot.com/>”. The value is saying replace <http://www.bbc.co.uk/> with <http://ocelot.com/>. The syntax is {find}, {replace}. Hopefully pretty simple. There are examples below that explain more.

18.4 Pre Downstream Request

Add the following to a ReRoute in ocelot.json in order to replace <http://www.bbc.co.uk/> with <http://ocelot.com/>. This header will be changed before the request downstream and will be sent to the downstream server.

```
"UpstreamHeaderTransform": {  
  "Test": "http://www.bbc.co.uk/, http://ocelot.com/"  
},
```

18.5 Post Downstream Request

Add the following to a ReRoute in ocelot.json in order to replace <http://www.bbc.co.uk/> with <http://ocelot.com/>. This transformation will take place after Ocelot has received the response from the downstream service.

```
"DownstreamHeaderTransform": {  
  "Test": "http://www.bbc.co.uk/, http://ocelot.com/"  
},
```

18.6 Placeholders

Ocelot allows placeholders that can be used in header transformation.

{RemoteIpAddress} - This will find the clients IP address using `_HttpContextAccessor.HttpContext.Connection.RemoteIpAddress.ToString()` so you will get back some IP. {BaseUrl} - This will use Ocelot’s base url e.g. <http://localhost:5000> as its value. {DownstreamBaseUrl} - This will use the downstream services base url e.g. <http://localhost:5000> as its value. This only works for DownstreamHeaderTransform at the moment. {TraceId} - This will use the Butterfly APM Trace Id. This only works for DownstreamHeaderTransform at the moment.

18.7 Handling 302 Redirects

Ocelot will by default automatically follow redirects however if you want to return the location header to the client you might want to change the location to be Ocelot not the downstream service. Ocelot allows this with the following configuration.

```
"DownstreamHeaderTransform": {
  "Location": "http://www.bbc.co.uk/, http://ocelot.com/"
},
"HttpHandlerOptions": {
  "AllowAutoRedirect": false,
},
```

or you could use the `BaseUrl` placeholder.

```
"DownstreamHeaderTransform": {
  "Location": "http://localhost:6773, {BaseUrl}"
},
"HttpHandlerOptions": {
  "AllowAutoRedirect": false,
},
```

finally if you are using a load balancer with Ocelot you will get multiple downstream base urls so the above would not work. In this case you can do the following.

```
"DownstreamHeaderTransform": {
  "Location": "{DownstreamBaseUrl}, {BaseUrl}"
},
"HttpHandlerOptions": {
  "AllowAutoRedirect": false,
},
```

18.8 X-Forwarded-For

An example of using `{RemoteIpAddress}` placeholder...

```
"UpstreamHeaderTransform": {
  "X-Forwarded-For": "{RemoteIpAddress}"
}
```

18.9 Future

Ideally this feature would be able to support the fact that a header can have multiple values. At the moment it just assumes one. It would also be nice if it could multi find and replace e.g.

```
"DownstreamHeaderTransform": {
  "Location": "[{one,one},{two,two}"
},
"HttpHandlerOptions": {
  "AllowAutoRedirect": false,
},
```

If anyone wants to have a go at this please help yourself!!

Claims Transformation

Ocelot allows the user to access claims and transform them into headers, query string parameters and other claims. This is only available once a user has been authenticated.

After the user is authenticated we run the claims to claims transformation middleware. This allows the user to transform claims before the authorisation middleware is called. After the user is authorised first we call the claims to headers middleware and Finally the claims to query string parameters middleware.

The syntax for performing the transforms is the same for each process. In the ReRoute configuration a json dictionary is added with a specific name either AddClaimsToRequest, AddHeadersToRequest, AddQueriesToRequest.

Note: I'm not a hotshot programmer so have no idea if this syntax is good...

Within this dictionary the entries specify how Ocelot should transform things! The key to the dictionary is going to become the key of either a claim, header or query parameter.

The value of the entry is parsed to logic that will perform the transform. First of all a dictionary accessor is specified e.g. Claims[CustomerId]. This means we want to access the claims and get the CustomerId claim type. Next is a greater than (>) symbol which is just used to split the string. The next entry is either value or value with and indexer. If value is specified Ocelot will just take the value and add it to the transform. If the value has an indexer Ocelot will look for a delimiter which is provided after another greater than symbol. Ocelot will then split the value on the delimiter and add whatever was at the index requested to the transform.

19.1 Claims to Claims Transformation

Below is an example configuration that will transform claims to claims

```
"AddClaimsToRequest": {  
  "UserType": "Claims[sub] > value[0] > |",  
  "UserId": "Claims[sub] > value[1] > |"  
}
```

This shows a transform where Ocelot looks at the user's sub claim and transforms it into UserType and UserId claims. Assuming the sub looks like this "usertypevalueuseridvalue".

19.2 Claims to Headers Transformation

Below is an example configuration that will transform claims to headers

```
"AddHeadersToRequest": {  
  "CustomerId": "Claims[sub] > value[1] > |"  
}
```

This shows a transform where Ocelot looks at the user's sub claim and transforms it into a CustomerId header. Assuming the sub looks like this "usertypevalueuseridvalue".

19.3 Claims to Query String Parameters Transformation

Below is an example configuration that will transform claims to query string parameters

```
"AddQueriesToRequest": {  
  "LocationId": "Claims[LocationId] > value",  
}
```

This shows a transform where Ocelot looks at the user's LocationId claim and adds it as a query string parameter to be forwarded onto the downstream service.

Ocelot uses the standard logging interfaces `ILoggerFactory` / `ILogger<T>` at the moment. This is encapsulated in `IOcelotLogger` / `IOcelotLoggerFactory` with an implementation for the standard asp.net core logging stuff at the moment. This is because Ocelot add's some extra info to the logs such as request id if it is configured.

There is a global error handler that should catch any exceptions thrown and log them as errors.

Finally if logging is set to trace level Ocelot will log starting, finishing and any middlewares that throw an exception which can be quite useful.

The reason for not just using bog standard framework logging is that I could not work out how to override the request id that get's logged when setting `IncludeScopes` to true for logging settings. Nicely onto the next feature.

20.1 Warning

If you are logging to Console you will get terrible performance. I have had so many issues about performance issues with Ocelot and it is always logging level Debug, logging to Console :) Make sure you are logging to something proper in production :)

This page details how to perform distributed tracing with Ocelot. At the moment we only support Butterfly but other tracers might just work without anything Ocelot specific.

21.1 Butterfly

Ocelot provides tracing functionality from the excellent [Butterfly](#) project. The code for the Ocelot integration can be found [here](#).

In order to use the tracing please read the Butterfly documentation.

In ocelot you need to do the following if you wish to trace a ReRoute.

```
Install-Package Ocelot.Tracing.Butterfly
```

In your ConfigureServices method

```
services
    .AddOcelot ()
    // this comes from Ocelot.Tracing.Butterfly package
    .AddButterfly(option =>
    {
        //this is the url that the butterfly collector server is running on...
        option.CollectorUrl = "http://localhost:9618";
        option.Service = "Ocelot";
    });
```

Then in your ocelot.json add the following to the ReRoute you want to trace..

```
"HandlerOptions": {
    "UseTracing": true
},
```

Ocelot will now send tracing information to Butterfly when this ReRoute is called.

Request Id / Correlation Id

Ocelot supports a client sending a request id in the form of a header. If set Ocelot will use the requestid for logging as soon as it becomes available in the middleware pipeline. Ocelot will also forward the request id with the specified header to the downstream service.

You can still get the asp.net core request id in the logs if you set `IncludeScopes` true in your logging config.

In order to use the request id feature you have two options.

Global

In your `ocelot.json` set the following in the `GlobalConfiguration` section. This will be used for all requests into Ocelot.

```
"GlobalConfiguration": {  
  "RequestIdKey": "OcRequestId"  
}
```

I recommend using the `GlobalConfiguration` unless you really need it to be `ReRoute` specific.

ReRoute

If you want to override this for a specific `ReRoute` add the following to `ocelot.json` for the specific `ReRoute`.

```
"RequestIdKey": "OcRequestId"
```

Once Ocelot has identified the incoming requests matching `ReRoute` object it will set the request id based on the `ReRoute` configuration.

This can lead to a small gotcha. If you set a `GlobalConfiguration` it is possible to get one request id until the `ReRoute` is identified and then another after that because the request id key can change. This is by design and is the best solution I can think of at the moment. In this case the `OcelotLogger` will show the request id and previous request id in the logs.

Below is an example of the logging when set at `Debug` level for a normal request..

```
debug: Ocelot.Errors.Middleware.ExceptionHandlerMiddleware[0]  
      requestId: asdf, previousRequestId: no previous request id, message: ocelot_  
↳pipeline started,
```

(continues on next page)

(continued from previous page)

```
dbug: Ocelot.DownstreamRouteFinder.Middleware.DownstreamRouteFinderMiddleware[0]
      requestId: asdf, previousRequestId: no previous request id, message: upstream_
↳url path is {upstreamUrlPath},
dbug: Ocelot.DownstreamRouteFinder.Middleware.DownstreamRouteFinderMiddleware[0]
      requestId: asdf, previousRequestId: no previous request id, message: downstream_
↳template is {downstreamRoute.Data.ReRoute.DownstreamPath},
dbug: Ocelot.RateLimit.Middleware.ClientRateLimitMiddleware[0]
      requestId: asdf, previousRequestId: no previous request id, message:_
↳EndpointRateLimiting is not enabled for Ocelot.Values.PathTemplate,
dbug: Ocelot.Authorisation.Middleware.AuthorisationMiddleware[0]
      requestId: 1234, previousRequestId: asdf, message: /posts/{postId} route does_
↳not require user to be authorised,
dbug: Ocelot.DownstreamUrlCreator.Middleware.DownstreamUrlCreatorMiddleware[0]
      requestId: 1234, previousRequestId: asdf, message: downstream url is
↳{downstreamUrl.Data.Value},
dbug: Ocelot.Request.Middleware.HttpRequestBuilderMiddleware[0]
      requestId: 1234, previousRequestId: asdf, message: setting upstream request,
dbug: Ocelot.Requester.Middleware.HttpRequesterMiddleware[0]
      requestId: 1234, previousRequestId: asdf, message: setting http response_
↳message,
dbug: Ocelot.Responder.Middleware.ResponderMiddleware[0]
      requestId: 1234, previousRequestId: asdf, message: no pipeline errors, setting_
↳and returning completed response,
dbug: Ocelot.Errors.Middleware.ExceptionHandlerMiddleware[0]
      requestId: 1234, previousRequestId: asdf, message: ocelot pipeline finished,
```

Middleware Injection and Overrides

Warning use with caution. If you are seeing any exceptions or strange behavior in your middleware pipeline and you are using any of the following. Remove them and try again!

When setting up Ocelot in your Startup.cs you can provide some additional middleware and override middleware. This is done as follows.

```
var configuration = new OcelotPipelineConfiguration
{
    PreErrorResponderMiddleware = async (ctx, next) =>
    {
        await next.Invoke();
    }
};

app.UseOcelot(configuration);
```

In the example above the provided function will run before the first piece of Ocelot middleware. This allows a user to supply any behaviors they want before and after the Ocelot pipeline has run. This means you can break everything so use at your own pleasure!

The user can set functions against the following.

- PreErrorResponderMiddleware - Already explained above.
- PreAuthenticationMiddleware - This allows the user to run pre authentication logic and then call Ocelot's authentication middleware.
- AuthenticationMiddleware - This overrides Ocelot's authentication middleware.
- PreAuthorisationMiddleware - This allows the user to run pre authorisation logic and then call Ocelot's authorisation middleware.
- AuthorisationMiddleware - This overrides Ocelot's authorisation middleware.
- PreQueryStringBuilderMiddleware - This allows the user to manipulate the query string on the http request before it is passed to Ocelot's request creator.

Obviously you can just add middleware as normal before the call to `app.UseOcelot()`. It cannot be added after as Ocelot does not call the next middleware.

Ocelot can load balance across available downstream services for each `ReRoute`. This means you can scale your downstream services and Ocelot can use them effectively.

The type of load balancer available are:

`LeastConnection` - tracks which services are dealing with requests and sends new requests to service with least existing requests. The algorithm state is not distributed across a cluster of Ocelot's.

`RoundRobin` - loops through available services and sends requests. The algorithm state is not distributed across a cluster of Ocelot's.

`NoLoadBalancer` - takes the first available service from config or service discovery.

`CookieStickySessions` - uses a cookie to stick all requests to a specific server. More info below.

You must choose in your configuration which load balancer to use.

Configuration

The following shows how to set up multiple downstream services for a ReRoute using ocelot.json and then select the LeadConnection load balancer. This is the simplest way to get load balancing set up.

```
{
  "DownstreamPathTemplate": "/api/posts/{postId}",
  "DownstreamScheme": "https",
  "DownstreamHostAndPorts": [
    {
      "Host": "10.0.1.10",
      "Port": 5000,
    },
    {
      "Host": "10.0.1.11",
      "Port": 5000,
    }
  ],
  "UpstreamPathTemplate": "/posts/{postId}",
  "LoadBalancerOptions": {
    "Type": "LeastConnection"
  },
  "UpstreamHttpMethod": [ "Put", "Delete" ]
}
```

Service Discovery

The following shows how to set up a ReRoute using service discovery then select the LeadConnection load balancer.

```
{
  "DownstreamPathTemplate": "/api/posts/{postId}",
  "DownstreamScheme": "https",
  "UpstreamPathTemplate": "/posts/{postId}",
  "UpstreamHttpMethod": [ "Put" ],
  "ServiceName": "product",
  "LoadBalancerOptions": {
    "Type": "LeastConnection"
  },
}
```

When this is set up Ocelot will lookup the downstream host and port from the service discover provider and load balance requests across any available services. If you add and remove services from the service discovery provider (consul) then Ocelot should respect this and stop calling services that have been removed and start calling services that have been added.

CookieStickySessions

I've implemented a really basic sticky session type of load balancer. The scenario it is meant to support is you have a bunch of downstream servers that don't share session state so if you get more than one request for one of these servers then it should go to the same box each time or the session state might be incorrect for the given user. This feature was requested in [Issue #322](#) though what the user wants is more complicated than just sticky sessions :) anyway I thought this would be a nice feature to have!

In order to set up CookieStickySessions load balancer you need to do something like the following.

```
{
  "DownstreamPathTemplate": "/api/posts/{postId}",
  "DownstreamScheme": "https",
  "DownstreamHostAndPorts": [
    {
      "Host": "10.0.1.10",
      "Port": 5000,
    },
    {
      "Host": "10.0.1.11",
      "Port": 5000,
    }
  ],
  "UpstreamPathTemplate": "/posts/{postId}",
  "LoadBalancerOptions": {
    "Type": "CookieStickySessions",
    "Key": "ASP.NET_SessionId",
    "Expiry": 1800000
  },
  "UpstreamHttpMethod": [ "Put", "Delete" ]
}
```

The LoadBalancerOptions are Type this needs to be CookieStickySessions, Key this is the key of the cookie you wish to use for the sticky sessions, Expiry this is how long in milliseconds you want to the session to be stuck for. Remember this refreshes on every request which is meant to mimick how sessions work usually.

If you have multiple ReRoutes with the same LoadBalancerOptions then all of those ReRoutes will use the same load

balancer for there subsequent requests. This means the sessions will be stuck across ReRoutes.

Please note that if you give more than one DownstreamHostAndPort or you are using a Service Discovery provider such as Consul and this returns more than one service then CookieStickySessions uses round robin to select the next server. This is hard coded at the moment but could be changed. llllll merged common ancestors Load Balancer
=====

Ocelot can load balance across available downstream services for each ReRoute. This means you can scale your downstream services and Ocelot can use them effectively.

The type of load balancer available are:

LeastConnection - tracks which services are dealing with requests and sends new requests to service with least existing requests. The algorithym state is not distributed across a cluster of Ocelot's.

RoundRobin - loops through available services and sends requests. The algorithym state is not distributed across a cluster of Ocelot's.

NoLoadBalancer - takes the first available service from config or service discovery.

CookieStickySessions - uses a cookie to stick all requests to a specific server. More info below.

You must choose in your configuration which load balancer to use.

Configuration

The following shows how to set up multiple downstream services for a ReRoute using ocelot.json and then select the LeadConnection load balancer. This is the simplest way to get load balancing set up.

```
{
  "DownstreamPathTemplate": "/api/posts/{postId}",
  "DownstreamScheme": "https",
  "DownstreamHostAndPorts": [
    {
      "Host": "10.0.1.10",
      "Port": 5000,
    },
    {
      "Host": "10.0.1.11",
      "Port": 5000,
    }
  ],
  "UpstreamPathTemplate": "/posts/{postId}",
  "LoadBalancerOptions": {
    "Type": "LeastConnection"
  },
  "UpstreamHttpMethod": [ "Put", "Delete" ]
}
```

Service Discovery

The following shows how to set up a ReRoute using service discovery then select the LeadConnection load balancer.

```
{
  "DownstreamPathTemplate": "/api/posts/{postId}",
  "DownstreamScheme": "https",
  "UpstreamPathTemplate": "/posts/{postId}",
  "UpstreamHttpMethod": [ "Put" ],
  "ServiceName": "product",
  "LoadBalancerOptions": {
    "Type": "LeastConnection"
  },
}
```

When this is set up Ocelot will lookup the downstream host and port from the service discover provider and load balance requests across any available services. If you add and remove services from the service discovery provider (consul) then Ocelot should respect this and stop calling services that have been removed and start calling services that have been added.

CookieStickySessions

I've implemented a really basic sticky session type of load balancer. The scenario it is meant to support is you have a bunch of downstream servers that don't share session state so if you get more than one request for one of these servers then it should go to the same box each time or the session state might be incorrect for the given user. This feature was requested in [Issue #322](#) though what the user wants is more complicated than just sticky sessions :) anyway I thought this would be a nice feature to have!

In order to set up CookieStickySessions load balancer you need to do something like the following.

```
{
  "DownstreamPathTemplate": "/api/posts/{postId}",
  "DownstreamScheme": "https",
  "DownstreamHostAndPorts": [
    {
      "Host": "10.0.1.10",
      "Port": 5000,
    },
    {
      "Host": "10.0.1.11",
      "Port": 5000,
    }
  ],
  "UpstreamPathTemplate": "/posts/{postId}",
  "LoadBalancerOptions": {
    "Type": "CookieStickySessions",
    "Key": "ASP.NET_SessionId",
    "Expiry": 1800000
  },
  "UpstreamHttpMethod": [ "Put", "Delete" ]
}
```

The LoadBalancerOptions are Type this needs to be CookieStickySessions, Key this is the key of the cookie you wish to use for the sticky sessions, Expiry this is how long in milliseconds you want to the session to be stuck for. Remember this refreshes on every request which is meant to mimick how sessions work usually.

If you have multiple ReRoutes with the same LoadBalancerOptions then all of those ReRoutes will use the same load

balancer for there subsequent requests. This means the sessions will be stuck across ReRoutes.

Please note that if you give more than one DownstreamHostAndPort or you are using a Service Discovery provider such as Consul and this returns more than one service then CookieStickySessions uses round robin to select the next server. This is hard coded at the moment but could be changed. ===== Load Balancer =====

Ocelot can load balance across available downstream services for each ReRoute. This means you can scale your downstream services and Ocelot can use them effectively.

The type of load balancer available are:

LeastConnection - tracks which services are dealing with requests and sends new requests to service with least existing requests. The alorythm state is not distributed across a cluster of Ocelot's.

RoundRobin - loops through available services and sends requests. The alorythm state is not distributed across a cluster of Ocelot's.

NoLoadBalancer - takes the first available service from config or service discovery.

CookieStickySessions - uses a cookie to stick all requests to a specific server. More info below.

You must choose in your configuration which load balancer to use.

The following shows how to set up multiple downstream services for a ReRoute using ocelot.json and then select the LeastConnection load balancer. This is the simplest way to get load balancing set up.

```
{
  "DownstreamPathTemplate": "/api/posts/{postId}",
  "DownstreamScheme": "https",
  "DownstreamHostAndPorts": [
    {
      "Host": "10.0.1.10",
      "Port": 5000,
    },
    {
      "Host": "10.0.1.11",
      "Port": 5000,
    }
  ],
  "UpstreamPathTemplate": "/posts/{postId}",
  "LoadBalancerOptions": {
    "Type": "LeastConnection"
  },
  "UpstreamHttpMethod": [ "Put", "Delete" ]
}
```

Service Discovery

The following shows how to set up a ReRoute using service discovery then select the LeastConnection load balancer.

```
{
  "DownstreamPathTemplate": "/api/posts/{postId}",
  "DownstreamScheme": "https",
  "UpstreamPathTemplate": "/posts/{postId}",
  "UpstreamHttpMethod": [ "Put" ],
  "ServiceName": "product",
  "LoadBalancerOptions": {
    "Type": "LeastConnection"
  },
}
```

When this is set up Ocelot will lookup the downstream host and port from the service discover provider and load balance requests across any available services. If you add and remove services from the service discovery provider (consul) then Ocelot should respect this and stop calling services that have been removed and start calling services that have been added.

CookieStickySessions

I've implemented a really basic sticky session type of load balancer. The scenario it is meant to support is you have a bunch of downstream servers that don't share session state so if you get more than one request for one of these servers then it should go to the same box each time or the session state might be incorrect for the given user. This feature was requested in [Issue #322](#) though what the user wants is more complicated than just sticky sessions :) anyway I thought this would be a nice feature to have!

In order to set up CookieStickySessions load balancer you need to do something like the following.

```
{
  "DownstreamPathTemplate": "/api/posts/{postId}",
  "DownstreamScheme": "https",
  "DownstreamHostAndPorts": [
    {
      "Host": "10.0.1.10",
      "Port": 5000,
    },
    {
      "Host": "10.0.1.11",
      "Port": 5000,
    }
  ],
  "UpstreamPathTemplate": "/posts/{postId}",
  "LoadBalancerOptions": {
    "Type": "CookieStickySessions",
    "Key": "ASP.NET_SessionId",
    "Expiry": 1800000
  },
  "UpstreamHttpMethod": [ "Put", "Delete" ]
}
```

The LoadBalancerOptions are Type this needs to be CookieStickySessions, Key this is the key of the cookie you wish to use for the sticky sessions, Expiry this is how long in milliseconds you want to the session to be stuck for. Remember this refreshes on every request which is meant to mimick how sessions work usually.

If you have multiple ReRoutes with the same LoadBalancerOptions then all of those ReRoutes will use the same load

balancer for there subsequent requests. This means the sessions will be stuck across ReRoutes.

Please note that if you give more than one DownstreamHostAndPort or you are using a Service Discovery provider such as Consul and this returns more than one service then CookieStickySessions uses round robin to select the next server. This is hard coded at the moment but could be changed.

Delegating Handlers

Ocelot allows the user to add delegating handlers to the HttpClient transport. This feature was requested [GitHub #208](#) and I decided that it was going to be useful in various ways. Since then we extended it in [GitHub #264](#).

33.1 Usage

In order to add delegating handlers to the HttpClient transport you need to do two main things.

First in order to create a class that can be used a delegating handler it must look as follows. We are going to register these handlers in the asp.net core container so you can inject any other services you have registered into the constructor of your handler.

```
public class FakeHandler : DelegatingHandler
{
    protected override async Task<HttpResponseMessage> SendAsync(HttpRequestMessage request, CancellationTok
    en cancellationToken)
    {
        //do stuff and optionally call the base handler..
        return await base.SendAsync(request, cancellationToken);
    }
}
```

Next you must add the handlers to Ocelot's container like below...

```
services.AddOcelot ()
    .AddDelegatingHandler<FakeHandler> ()
    .AddDelegatingHandler<FakeHandlerTwo> ()
```

Both of these Add methods have a default parameter called global which is set to false. If it is false then the intent of the DelegatingHandler is to be applied to specific ReRoutes via ocelot.json (more on that later). If it is set to true then it becomes a global handler and will be applied to all ReRoutes.

e.g.

As below...

```
services.AddOcelot ()
    .AddDelegatingHandler<FakeHandler> (true)
```

Finally if you want ReRoute specific DelegatingHandlers or to order your specific and / or global (more on this later) DelegatingHandlers then you must add the following json to the specific ReRoute in ocelot.json. The names in the array must match the class names of your DelegatingHandlers for Ocelot to match them together.

```
"DelegatingHandlers": [
  "FakeHandlerTwo",
  "FakeHandler"
]
```

You can have as many DelegatingHandlers as you want and they are run in the following order:

1. Any globals that are left in the order they were added to services and are not in the DelegatingHandlers array from ocelot.json.
2. Any non global DelegatingHandlers plus any globals that were in the DelegatingHandlers array from ocelot.json ordered as they are in the DelegatingHandlers array.
3. Tracing DelegatingHandler if enabled (see tracing docs).
4. QoS DelegatingHandler if enabled (see QoS docs).
5. The HttpClient sends the HttpRequestMessage.

Hopefully other people will find this feature useful!

Raft (EXPERIMENTAL DO NOT USE IN PRODUCTION)

Ocelot has recently integrated [Rafty](#) which is an implementation of Raft that I have also been working on over the last year. This project is very experimental so please do not use this feature of Ocelot in production until I think it's OK.

Raft is a distributed consensus algorithm that allows a cluster of servers (Ocelots) to maintain local state without having a centralised database for storing state (e.g. SQL Server).

To get Raft support you must first install the Ocelot Rafty package.

```
Install-Package Ocelot.Provider.Rafty
```

Then you must make the following changes to your Startup.cs / Program.cs.

```
public virtual void ConfigureServices(IServiceCollection services)
{
    services
        .AddOcelot()
        .AddAdministration("/administration", "secret")
        .AddRafty();
}
```

In addition to this you must add a file called peers.json to your main project and it will look as follows

```
{
  "Peers": [{
    "HostAndPort": "http://localhost:5000"
  },
  {
    "HostAndPort": "http://localhost:5002"
  },
  {
    "HostAndPort": "http://localhost:5003"
  },
  {
    "HostAndPort": "http://localhost:5004"
  },
  ],
}
```

(continues on next page)

(continued from previous page)

```
    {  
      "HostAndPort": "http://localhost:5001"  
    }  
  ]  
}
```

Each instance of Ocelot must have its address in the array so that they can communicate using Raft.

Once you have made these configuration changes you must deploy and start each instance of Ocelot using the addresses in the peers.json file. The servers should then start communicating with each other! You can test if everything is working by posting a configuration update and checking it has replicated to all servers by getting their configuration.

CHAPTER 35

Overview

This document summarises the build and release process for the project. The build scripts are written using [Cake](#), and are defined in *.build.cake*. The scripts have been designed to be run by either developers locally or by a build server (currently [AppVeyor](#)), with minimal logic defined in the build server itself.

- You'll generally want to run the *.build.ps1* script. This will compile, run unit and acceptance tests and build the output packages locally. Output will go to the *.artifacts* directory.
- You can view the current commit's SemVer build information by running *.version.ps1*.
- The other *./*.ps1* scripts perform subsets of the build process, if you don't want to run the full build.
- The release process works best with GitFlow branching; this allows us to publish every development commit to an unstable feed with a unique SemVer version, and then choose when to release to a stable feed.
- Alternatively you can build the project in VS2017 with the latest .NET Core SDK.

The tests should all just run and work apart from the integration tests which need the following environmental variables setting. This is a manual step at the moment.

```
OCELOT_USERNAME=admin
```

```
OCELOT_HASH=kE/mxd1hO9h9S12VhGhwJUd9xZEv4NP6qXoN39nIqM4=
```

```
OCELOT_SALT=zzWITpnDximUNKYLiUam/w==
```

On windows you can use..

```
SETX OCELOT_USERNAME admin
```

On mac..

```
export OCELOT_USERNAME=admin
```

I need to work out a nicer way of doing this in the future.

Release process

This section defines the release process for the maintainers of the project. * Merge pull requests to the *release* branch.

- Every commit pushed to the Origin repo will kick off the `ocelot-build` project in AppVeyor. This performs the same tasks as the command line build, and in addition pushes the packages to the unstable nuget feed.
- When you're ready for a release, create a release branch. You'll probably want to update the committed `.ReleaseNotes.md` based on the contents of the equivalent file in the `.artifacts` directory.
- When the *release* branch has built successfully in Appveyor, select the build and then Deploy to the *GitHub Release* environment. This will create a new release in GitHub.
- In Github, navigate to the `release`. Modify the release name and tag as desired.
- When you're ready, publish the release. This will tag the commit with the specified release number.
- The `ocelot-release` project will detect the newly created tag and kick off the release process. This will download the artifacts from GitHub, and publish the packages to the stable nuget feed.
- When you have a final stable release build, merge the *release* branch into *master* and *develop*. Deploy the master branch to github and following the full release process as described above. Don't forget to uncheck the "This is a pre-release" checkbox in GitHub before publishing.
- Note - because the release builds are initiated by tagging a commit, if for some reason a release build fails in AppVeyor you'll need to delete the tag from the repo and republish the release in GitHub.