

App Connect and Knative

John Hosie

Architect: IBM App Connect



Contents

Introduction to Knative

Auto Scaling

Scale to Zero

App Connect on Knative

Developer Friendly

Developer Distractions

ingress networking scaling breaking API changes

glue code footprint data formats

Knative

helps to take the toil out of deployment

App Connect

helps to avoid specialist knowledge in the integration layer

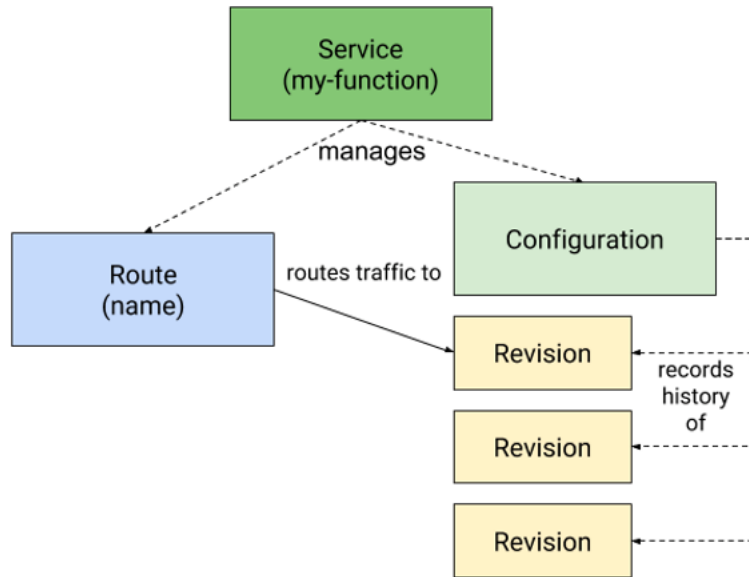
Knative Basics

Demo

What Just Happened?

Developer Declares a Service

Knative Operator creates
Route
Configuration
Revision



Explanation of terms

Service

The primary resource that the developer interacts with.
Declaration of the code and config needed to run the app.

Route

Externally addressable endpoint.
Provides configuration to the ingress gateway for routing inbound traffic.

Configuration

Specification of runtime Kubernetes resources (pods) required to run the app.

Revision


Snapshot of a configuration.
Allows rollback, canary testing etc...

Demo

What Just Happened?

Concurrency Levels increase

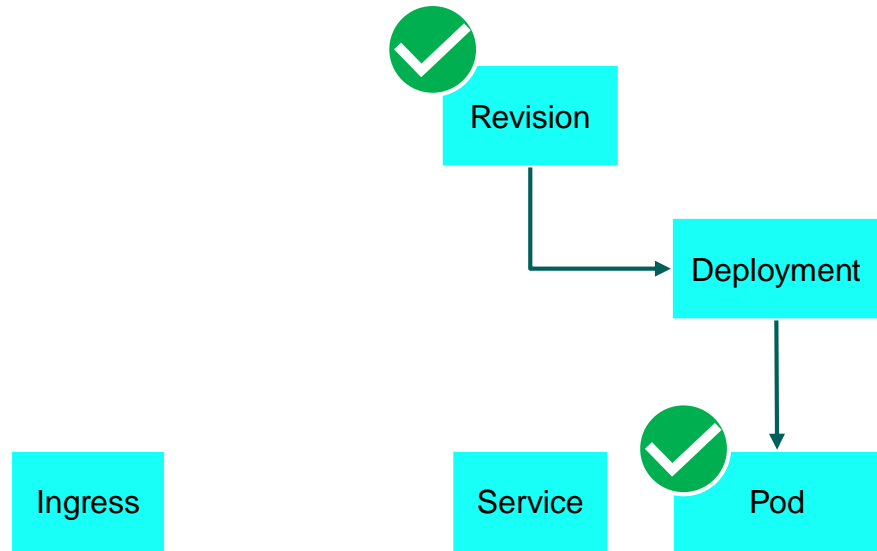
Knative activator scales the pod



What Just Happened?

Concurrency Levels increase

Knative activator scales the pod



Revision is marked as ready once it has created a deployment with a single pod and that pod is ready

What Just Happened?

Concurrency Levels increase

Knative activator scales the pod

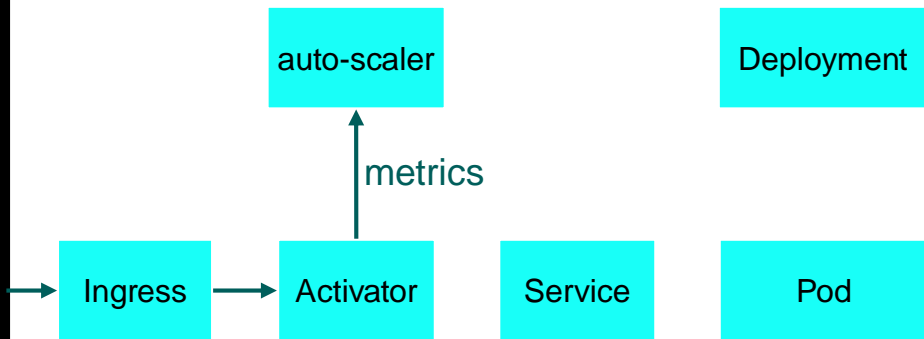


Ingress Gateway forwards traffic, for all routes, to the Activator in the knative-serving namespace

What Just Happened?

Concurrency Levels increase

Knative activator scales the pod

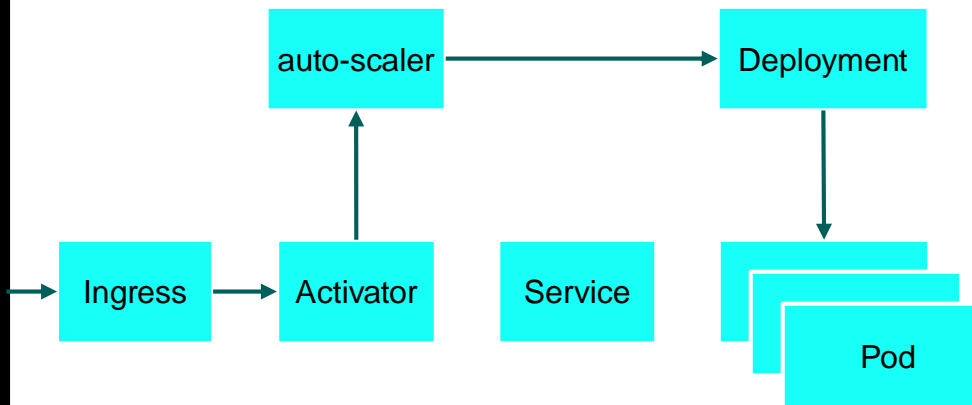


Activator buffers the request if there are no revisions ready and also sends metrics to the auto-scaler

What Just Happened?

Concurrency Levels increase

Knative activator scales the pod

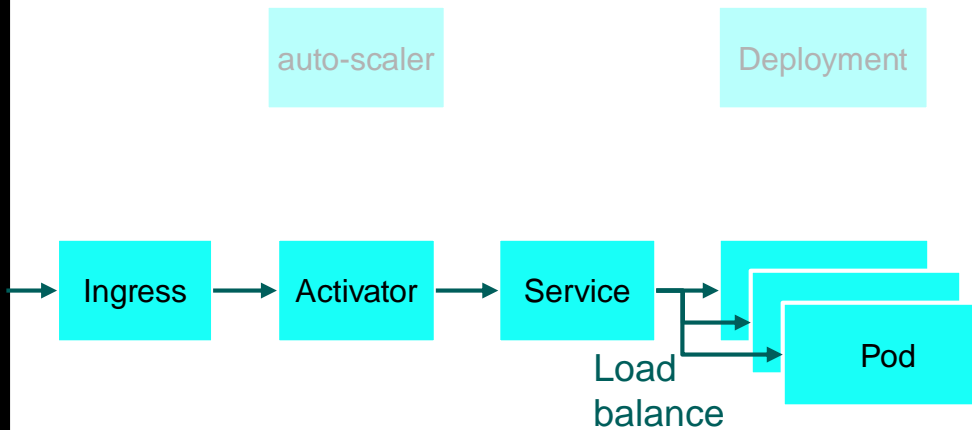


Based on the metrics, auto-scaler will configure the number of replicas on the Deployment. Deployment then creates Pods

What Just Happened?

Concurrency Levels increase

Knative activator scales the pod



Activator forwards the traffic to the Service that then load balances to the pods

Demo

What Just Happened?

Service is idle for 2 minutes

Pod scales to zero



App Connect on Knative

Demo

Making App Connect Fit

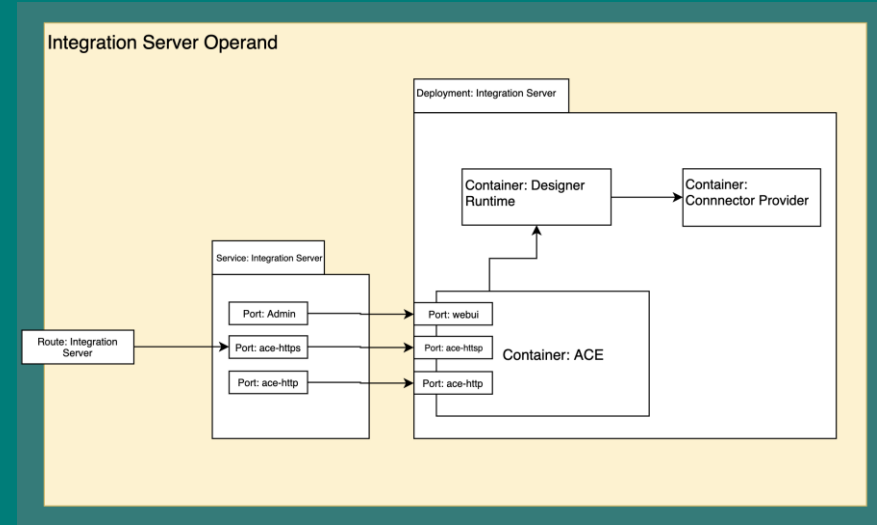
Knative only supports single container pods

Knative only supports a single port

Readiness and Liveness

Scale to nearly zero

HTTP(S) inbound only



Try it yourself

<https://github.com/ot4i/knapp-connect>

Instructions, sample code, discussion

Watch this space

Thank you

John Hosie
Architect: IBM App Connect

jhosie@uk.ibm.com
Twitter: @hosie31504

© Copyright IBM Corporation 2019. All rights reserved. The information contained in these materials is provided for informational purposes only, and is provided AS IS without warranty of any kind, express or implied. Any statement of direction represents IBM's current intent, is subject to change or withdrawal, and represent only goals and objectives. IBM, the IBM logo, and ibm.com are trademarks of IBM Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available at [Copyright and trademark information](#).

