#### Unless otherwise specified...



/\*

Copyright 2017 The Kubernetes Authors.

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

http://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.





#### KubeCon | CloudNativeCon

North America 2018

## Why Are We Copying and Pasting So Much?

Philip Wittrock (@pwittrock) & Solly Ross (@directxman12), Google



- Respond to changes to Resources made by users
- Respond to changes to Cluster State

#### **Controllers Are Simple**



### ✓ watch resources,

- ✓ do some business logic,
- ✓ reconcile differences,
- ... and a bit of plumbing

How complicated could it be?

#### Sample Controller at a Glance



~/go/src/k8s.io/sample-controller \$ wc -l < controller.go
429
~/go/src/k8s.io/sample-controller \$ tree -I vendor
<snip>
28 directories, 52 files

#### return controller

#### Sample

Run will set up the event handlers for types we are interested in, as well as syncing informer caches and starting workers. It will block until stopC is closed, at which point it will shutdown the workqueue and wait for workers to finish processing their current work items. we (c \*Controller) Run(threadiness int, stopCh <-chan struct{}) error { defer runtime.HandleCrash() defer c.workqueue.ShutDown()

// Start the informer factories to begin populating the informer caches
glog.Info("Starting Foo controller")

// Wait for the caches to be synced before starting workers
glog.Info("Waiting for informer caches to sync")
if ok := cache.WaitForCacheSync(stopCh, c.deploymentsSynced, c.foosSynced); !ok {
 return fmt.Errorf("failed to wait for caches to sync")

glog.Info("Started workers")
<-stopCh
glog.Info("Shutting down workers")</pre>

#### return nil

runWorker is a long-running function that will continually call the processNextWorkItem function in order to read and process a message on the workqueue. nc (c \*Controller) runWorker() {

or c.processNextWorkItem() {

processNextWorkItem will read a single work item off the workqueue and attempt to process it, by calling the synchandler. mc (c \*Controller) processNextWorkItem() bool { obj, shutdown := c.workqueue.Get()

if shutdown { return false

// We wrap this block in a func so we can defer c.workqueue.Done.
err := func(obj interface{}) error {

// We call Done here so the workqueue knows we have finished // processing this item. We also must remember to call Forget if w // do not want this work item being re-queued. For example, we do // not call Forget if a transient error occurs, instead the item i // out being an the communication determined coming after a back off.

// period.

defer c.workqueue.Done(obj)

var key string

var ok bool

(/ We expect strings to come off the workqueue. These are of the (/ form namespace/name. We do this as the delayed nature of the // workqueue means the items in the informer cache may actually be // more up to date that when the item was initially put onto the // workqueue.

if key, ok = obj.(string); !ok {

// As the item in the workqueue is actually invalid, we call // Forget here else we'd go into a loop of attempting to

/ process a work item that is invalid.



### Sample Controller





#### This is not a flame chart





Note: Boilerplate code to copy-paste highlighted in Red

#### At least the comments are thorough...

. . .

. . .



// We call Done here so the workqueue knows we have finished // processing this item. We also must remember to call Forget if we // do not want this work item being re-queued. For example, we do // not call Forget if a transient error occurs, instead the item is // put back on the workqueue and attempted again after a back-off // period.

// We expect strings to come off the workqueue. These are of the // form namespace/name. We do this as the delayed nature of the // workqueue means the items in the informer cache may actually be // more up to date that when the item was initially put onto the // workqueue.

// As the item in the workqueue is actually invalid, we call
// Forget here else we'd go into a loop of attempting to
// process a work item that is invalid.

#### **Controllers are Simple?**

- 1. Copy Clients to Reconciler Struct
- 2. Create a Queue for Reconcile Requests (namespace/name)
- 3. Add EventHandlers for all Object Types that you are interested in
  - Don't Enqueue for Deletion Events on the Reconciled Type а.
    - Unless you have a finalizer, then you should
  - Do Enqueue for Deletion Events on Owned Types b.
    - But find the owner of the object to Reconcile instead İ.
      - Write Handlers to walk owners References to find parents 1
        - Be sure to handle Tombstones correctly а.

٠V

all BIELBUBE

Actually maybe we don't need to do that anymore? Ĭ.

9

·0/

- Write Logic to Turn Objects Into Keys C.
  - Type Cast Request to String i.
    - Error Handling if this fails
- Start Worker Threads 4.
  - But not too many а.
- Start Informers 5.
- Wait For Cache To Sync 50 31 31 а.



#### Which is the real Sample Controller?







Sample Controler + Core Kubernetes Controller + Third Party Operator





65% of the sample controller code is...

- Unmodified Copy-Pasted Functions
- Initialization and Instantiation of Deps
- Plumbing and Wiring





#### controller-runtime (proper noun)

a Go standard library for controllers

backage main

Just Kidding

#### Seriously though, use this instead



#### controller-runtime (proper noun)

a Go standard library for controllers

backage main

import (
 "sigs.k8s.io/controller-runtime/pkg/builder"

#### **Create Controller**



pe Controller struct {
 workqueue workqueue.RateLimitingInterface

nc Start() {
 stopCh := signals.SetupSignalHam

controller.Run(2, stopCh

inc (c \*Controller) Run(threadiness int, stopCh <-chan struct{}) error {
 defer runtime.HandleCrash()
 defer (sovrkueues.ShutDown()</pre>

klog.Info("Starting Foo controlle

klog\_Info("Waiting for informer caches to sync") if ok := cache.MaitForCacheSync(stopCh, c.deploymentsSynced, c.foosSynced); lok { return fmt.Errorf("failed to wait for caches to sync")

klog.Info("Starting workers")
for i := 0; i < threadiness; i++ {
 go wait.Until(c.runWorker, time.Second, stop()</pre>

klog.Info("Started workers")
<-stopCh
klog.Info("Shutting down workers".</pre>

return nil

inc (c \*Controller) runMorker() {
 for c.processNextWorkIten() {

nc (c \*Controller) processNextWorkItem() bool { obj, shutdown := c.workgueue.Get()

if shutdown {
 return false

err := func(obj interface{}) error { defer c.workqueue.Done(obj)

var key stri

f key, ok = obj.(string);

runtime.HandleError(fmt.Errorf("expected string in workqueue but got %#v",
runtime.HandleError(fmt.Errorf("expected string in workqueue but got %#v",

} if err := c.syncHandler(key); err != nil {

return fmt.Errorf("error syncing '%s': %s, requeuing", key, er

c.workqueue.Forge(100])
klog.Infof("Successfully synced "%s", key)
return nil

if err != nil {
 runtime.HandleErrc
 return true

return true

#### type Controller struct {}

func Start() {
 ctrl, \_ := builder.SimpleController().Build(&Controller{})
 ctrl.Start(signals.SetupSignalHandler())

#### Watch Object



<pre>sype Controller struct {     sampleclientset clie     foosLister lie     foosSynced ca     workqueue workqueue.</pre>	intset.Interface Lsters.FooLister che.InformerSynced RateLimitingInterface	
stopCh := signals.Se	tupSignalHandler()	
	nd.BuildConfigFromFlags(masterURL, kubeconfig)	type (
	and the first state of the stat	rybe c
klog.Fatalf("Eri	or building Rubecontig: %5", err.Error())	c1
		1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
exampleClient, err :	= clientset.NewForConfig(cfg)	3
<pre>if err != nit {     klog Fatalf("Fri </pre>	or building example clientset: se" orr Error())	
	or burning example councient of , enterior(//	
exampleInformerFacto fooInformer := examp	ory := informers.NewSharedInformerFactory(exampleClient, time.Second*30) leInformerFactory.Samplecontroller().Vlalphal().Foos()	func (
controller := &Contr	-oller{	
sampleclientset:	sampleclientset,	6.
	<pre>fooInformer.Lister(),</pre>	100
	fooInformer.Informer().HasSynced,	1.22
	<pre>workqueue.NewNamedRateLimitingQueue(workqueue.DefaultControllerRateLimiter(), "Foos"),</pre>	3
fooInformer.Informer	-().AddEventHandler(cache.ResourceEventHandlerFuncs{	
AddFunc: control	ler.enqueueFoo,	
UpdateFunc: fund	:(old, new interface{}) {	func S
controller.e	enqueueFoo (new)	Constraint of the
		ct
	ory.Start(stopCh)	
	con(b)	
	weveEnalabi interface()) [	
var key string	Instantion) Three Large A. C	ct
	alante - the state - of West	
if key, err = cache.	MetaNamespaceKeyFunc(obj); err != nil {	}
runtime.HandleEr	ror(err)	
c.workqueue.AddRatel	.imited(key)	

type Controller struct {
Client client.Client

```
func (ct *Controller) InjectClient(c client.Client) error {
    ct.Client = c
    return nil
```

```
inc Start() {
    ctrl, _ := builder.SimpleController().
        ForType(&v1alpha1.Foo{}).
        Build(&Controller{})
        ctrl.Start(signals.SetupSignalHandler())
```

#### Watch Owned Objects



	CLEANSACLINE FALC
	er appslisters.DeploymentLister
	listers.FooLister
	cache.InformerSynced
	ueue.RateLimitingInterface
	ls.SetupSignalHandler()
	entend.BuildConfigTromTlags(masterURL, kubeconfig)
	("Error building Asbeconfig: %s", err.Error[])
	the function of the second
	("Error deltaing elbernetes clientset: 55", errichtor())
	err := clientset.NewForConfig1cfg)
	('Error building example clientset: %s", err.Error())
	en la sina de la seconda en de seconda en el case
exampleInformat	Tortery := informers.NewSharedInformerFactory(exampleClient, time.Second*30) Factory := informers.NewSharedInformerFactory(exampleClient, time.Second*30)
fooInformer :-	exampleInformerFactory.Samplecostroller().Vialpha1().Foos()
	(mtral)ard
	tset: sumplectiontset,
deptoyments	Synced: deploymentInformer.Informer().HasSynced,
	fooInformer.Lister(), fooInformer.Informer().HasSynced,
	<pre>workqueue.NewNamedRateLimitingDumue(workqueue.DefaultControllerRateLimiter(), "Foce"),</pre>
	ormer().AddDventHandler(coche.ResourceDventHandlerFuncs(
	ntroller.engueueFoo, func(old, new interface{}) (
	ter.enqueseFos(new)
deploymentInfor AddFunc: co	mer.Informer().AddIventHandler(cache.ResourceIventHandlerPancs( etroller.handleObject.
UpdateFanc:	func(old, new interface()) (
	:= ald.(*sppsv1.Osployment)
	am
	ler.handleCbject(new)
	controller.handle0bject,
	tory.Start(stopCh)
	factory.Start(stopTh)
	2, stopCh)
	) hasfledbject(obj interface()) (
	P1:05ject
if object, ek - tombatore.	obj.(metavl.Object); iek ( ak := obj.(cache.DeletedFinalStateUnknown)
	.HandleError(fmt.Errorf('error decoding object, invalid type'))
object, ok	e tombotone.Obj.(metavl.Object)
	.NandleError(fmt.Errorf('error decoding object tombstome, invalid type"))
	infof('Recovered deleted object '\s' from tombstone', object.GetHame())
	("Processing object: hs", object.GotName())
if conserRef in	estavl.GetControllerOf(sbject); semerRef := nil ( abject is not semed by a Fos, we should not de anything more
// with it.	Rind in Treat J
retern	
	- Annal Anton Providelants Californian (1), Californian (1) (mail
if err is a	al (
	().Isfof("ignoring orphaned object "As" of fos "As"", object.GatSelfLink(), ownerRef.Name)
	e (foo)
return	

#### type Controller struct { Client client.Client

func (ct \*Controller) InjectClient(c client.Client) error {
 ct.Client = c
 return nil

func Start() {
 ctrl, \_ := builder.SimpleController().
 ForType(&v1alpha1.Foo{}).
 Owns(&appsv1.Deploment{}).
 Build(&Controller{})
 ctrl.Start(signals.SetupSignalHandler())

### The magic of abstractions...





Helpers for Common Cases (*EnqueueRequestForObject*)

Flexible Lower-Level Interfaces (*EventHandler, Source*)

API Machinery (∗€ふኇ፟፟፟፟፟፟የ፟፟፟፟፟፟፟ የ፝፝፝፝፝፞፝ \ የ፝፟፝፝ የ፝፟፝ \ የ፝፞፝፝ \ የ፟፟፝ \ የ

#### **Sophisticated Watch APIs**



**The.** For Type **builder function is really just a call to** .Watch **with** EnqueueRequestForObject

#### **Sophisticated Watch APIs**



// .Owns(&appsv1.ReplicaSet{})
c.Watch(&source.Kind{Type: &corev1.Pod{}},
 &handler.EnqueueRequestForOwner{
 IsController: true, OwnerType: &appsv1.ReplicaSet{}})

**The.** Owns **builder function is really just a call to** .Watch **with** EnqueueRequestForOwner

#### **Sophisticated Watch APIs**



Users can implement their own patterns by providing a function to map events to objects

#### **Extending Watch APIs**



Watch(src source.Source, eventhandler handler.EventHandler, predicates ...predicate.Predicate) error

type EventHandler interface {

// Create is called in response to an create event - e.g. Pod Creation. Create(event.CreateEvent, workqueue.RateLimitingInterface)

// Update is called in response to an update event - e.g. Pod Updated. Update(event.UpdateEvent, workqueue.RateLimitingInterface)

// Delete is called in response to a delete event - e.g. Pod Deleted. Delete(event.DeleteEvent, workqueue.RateLimitingInterface)

// Generic is called in response to an event of an unknown type or a synthetic event triggered as a cron or // external trigger request - e.g. reconcile Autoscaling, or a Webhook. Generic(event.GenericEvent, workqueue.RateLimitingInterface)

type Source interface {

// Start is internal and should be called only by the Controller to register an EventHandler with the Informer
// to enqueue reconcile.Requests.

Start(handler.EventHandler, workqueue.RateLimitingInterface, ...predicate.Predicate) error

### For Developers...





### For Ops...

# KubeCon CloudNativeCon North America 2018

#### Standard...

- Behavior
- Logging
- Metrics
- Error handling
- Error degradation
- Maintenance



### For Platform Developers...



#### Build high-level abstractions with common building blocks:

- Kubebuilder
- Operator SDK
- Maestro Declarative Operator (via kubebuilder)
- AddOnOperator (via kubebuilder)

#### Remember this...





#### Remember this...





#### It fits!





c updateDeployment(deploymentName string, foo *samplev1alpha1.Foo, existing runtime.Object) error {
depl := existing.(*appsv1.Deployment)
depl.Name 🗕 deploymentName
depl.Namespace = foo.Namespace

#### In conclusion...



#### **Controllers Are Simple**

package main

import (
 "sigs.k8s.io/controller-runtime/pkg/builder"





# https://book.kubebuilder.io