

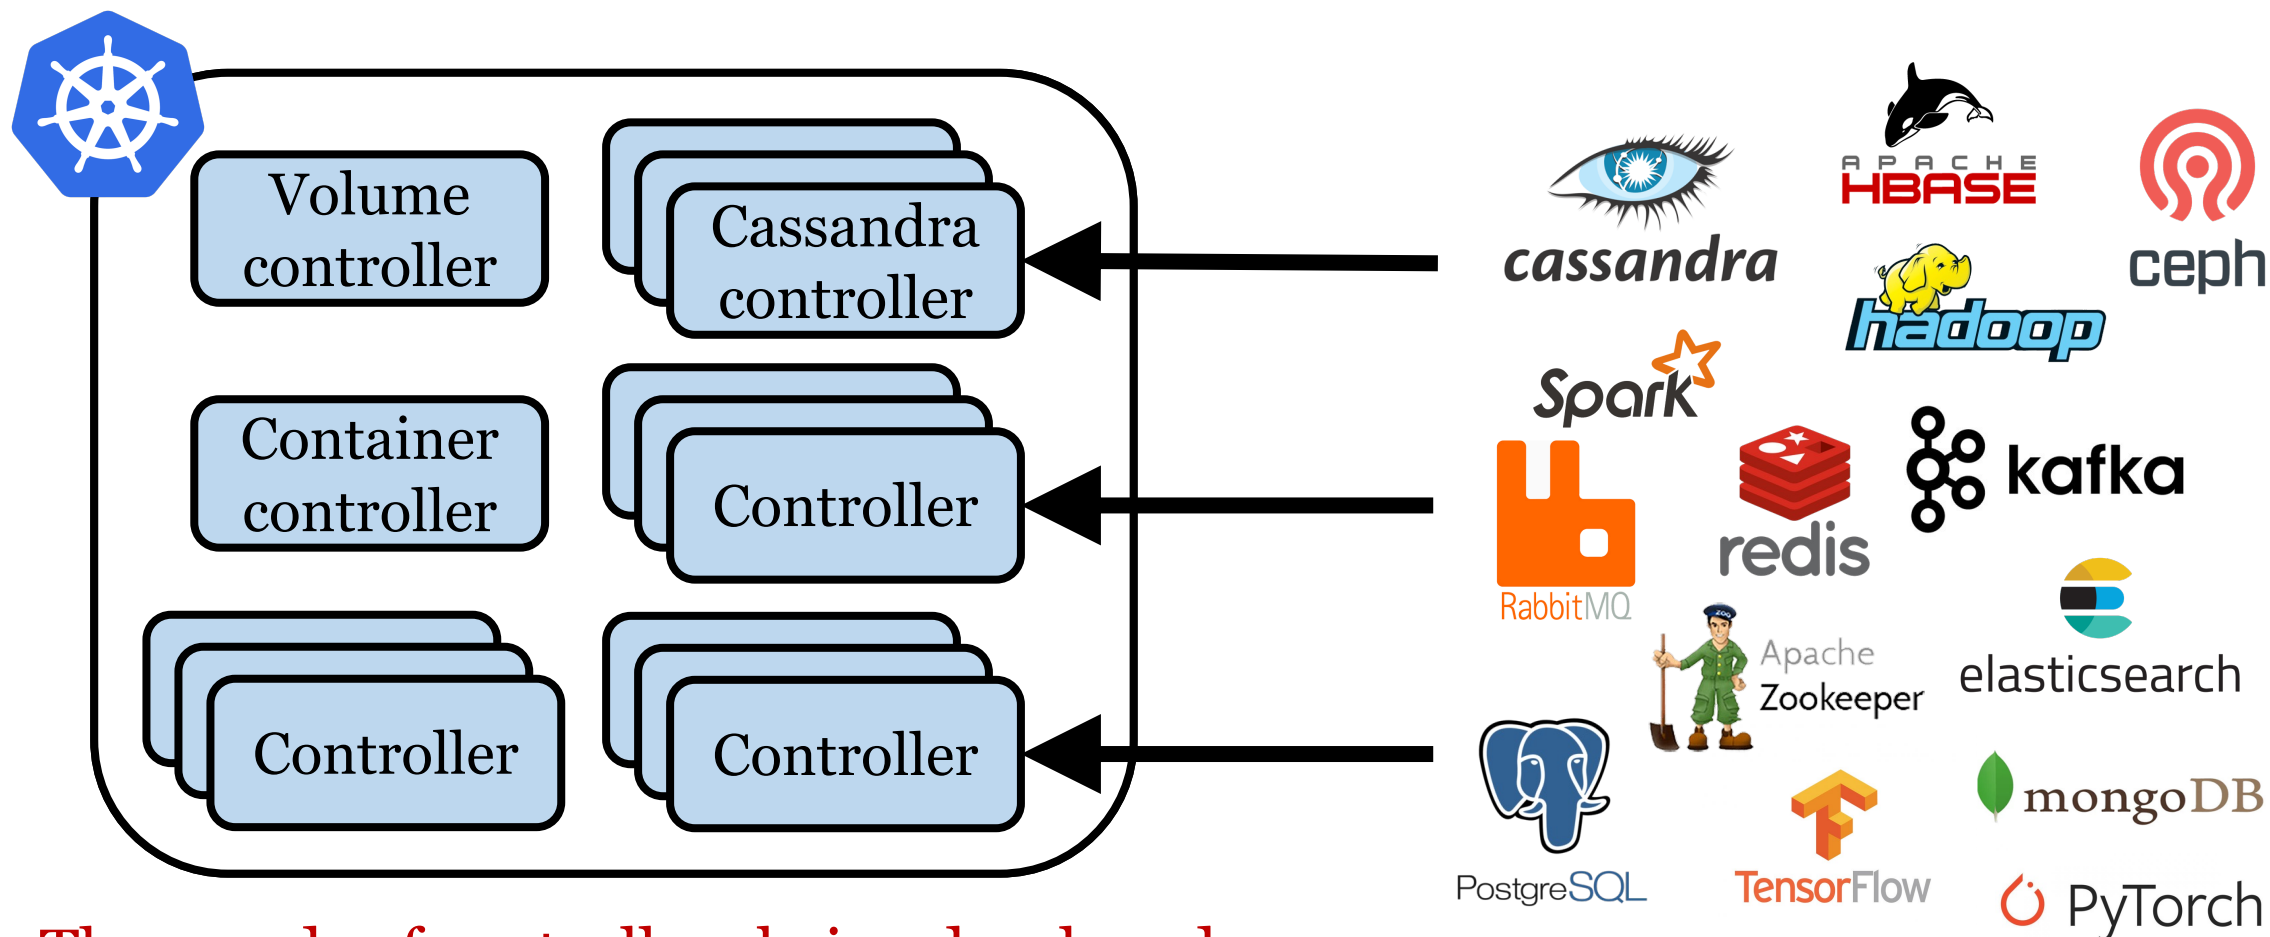


# Automatic Reliability Testing for Cluster Management Controllers

**Xudong Sun**, Wenqing Luo, Jiawei Tyler Gu, Aishwarya Ganesan, Ramnathan Alagappan, Michael Gasch, Lalith Suresh, Tianyin Xu

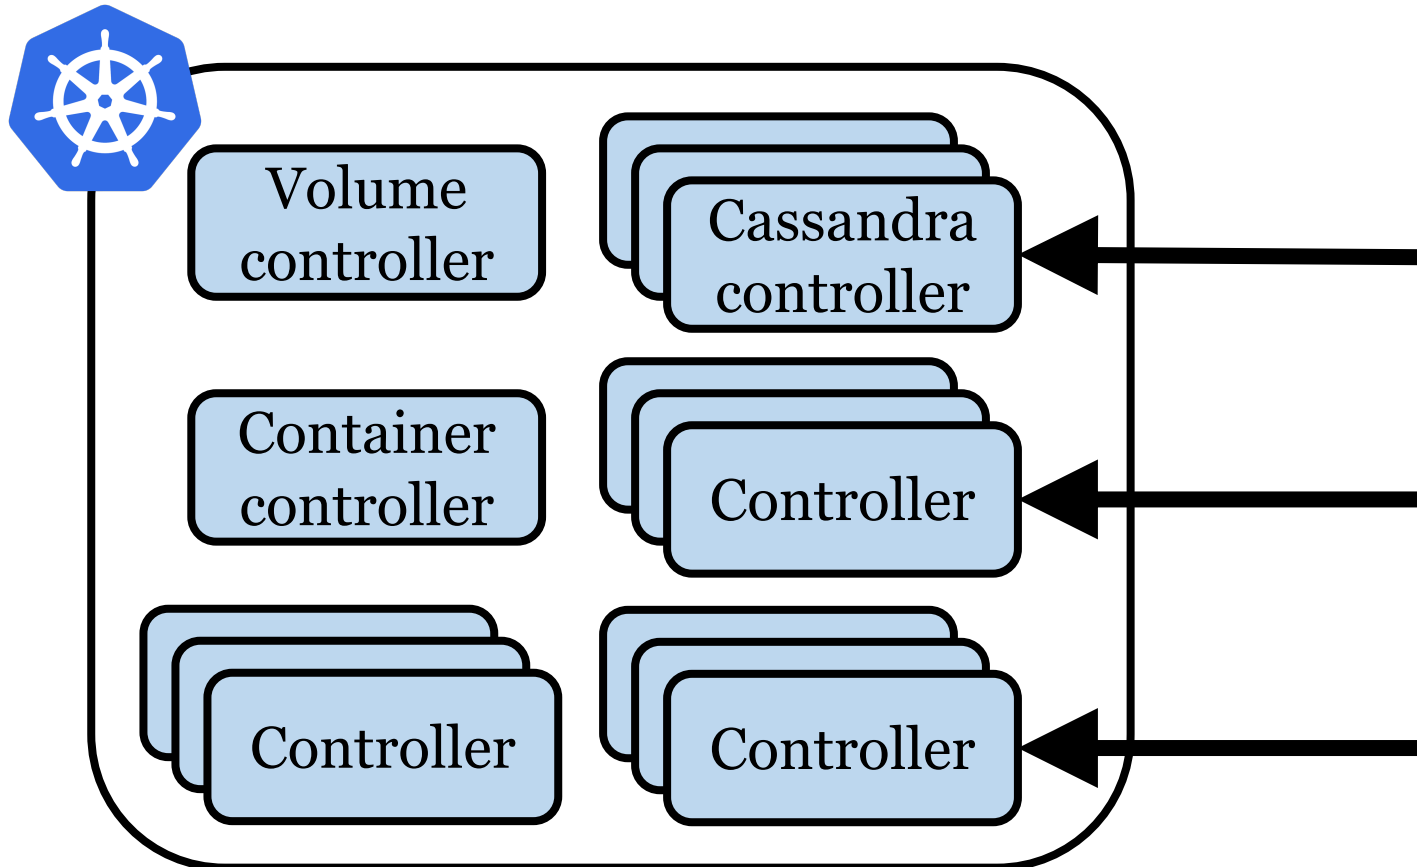


# Cluster management is realized by controllers



Thousands of controllers being developed  
by the Kubernetes community

# Cluster management is realized by controllers



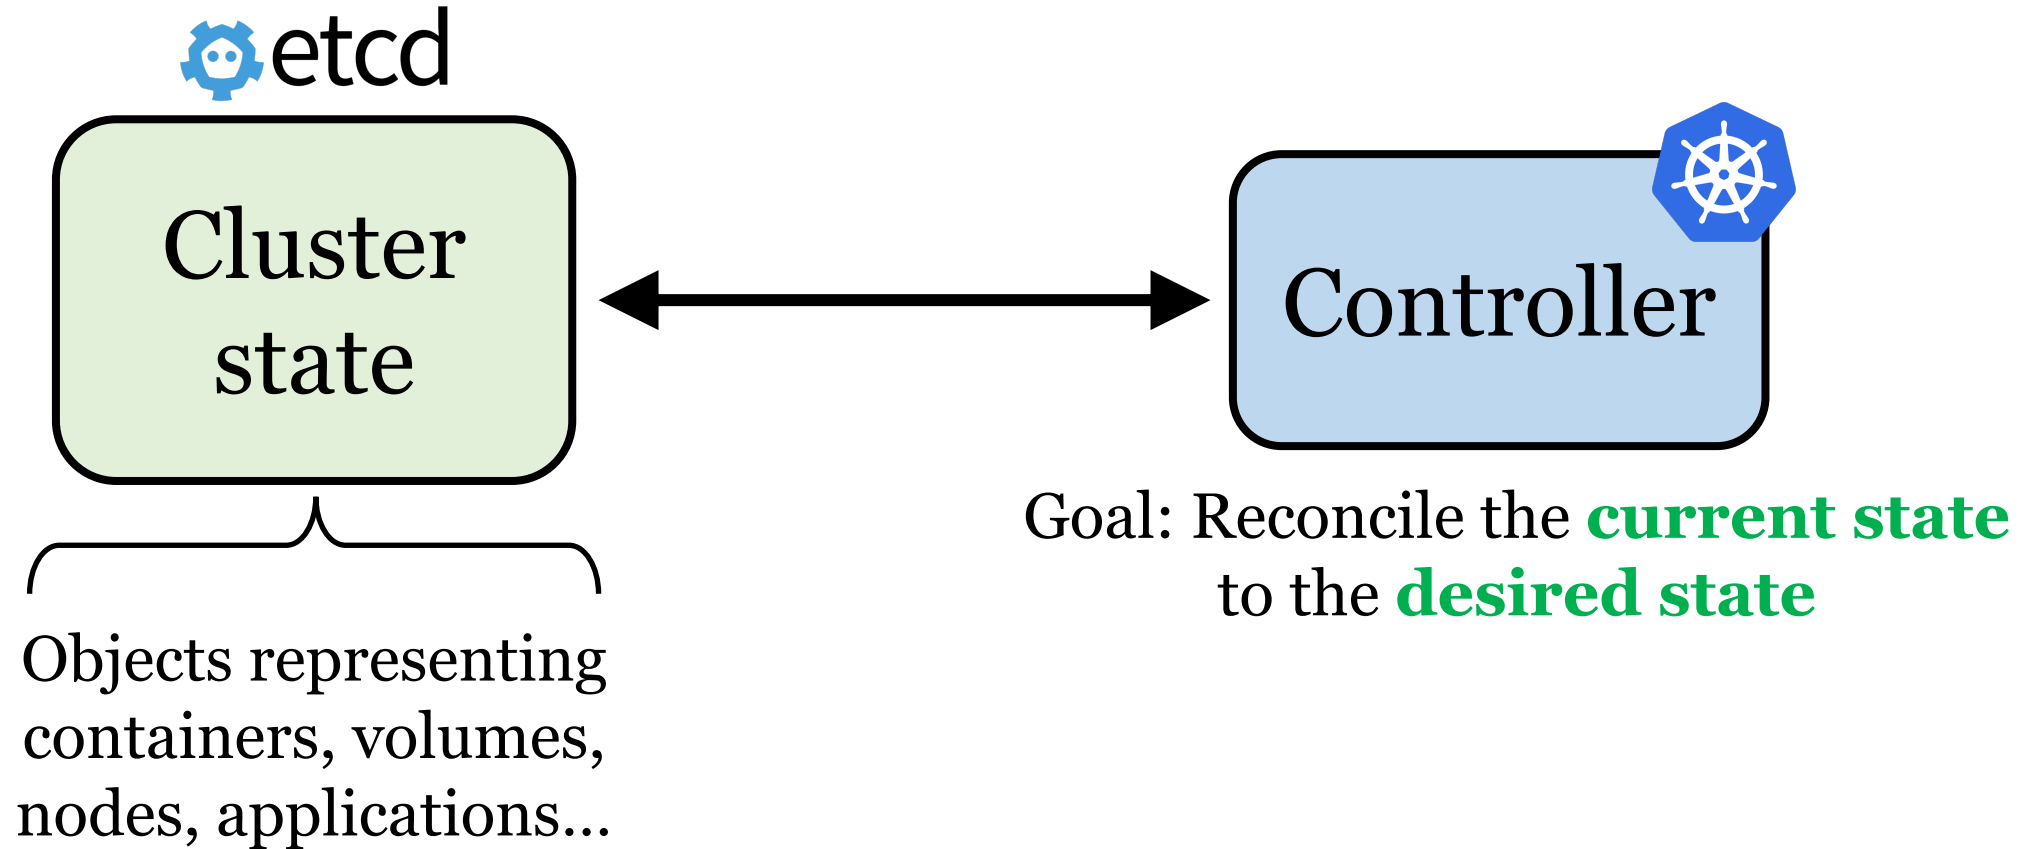
Twitter joins the rest of the world  
– moves to Kubernetes

July 30, 2019 1 Minute

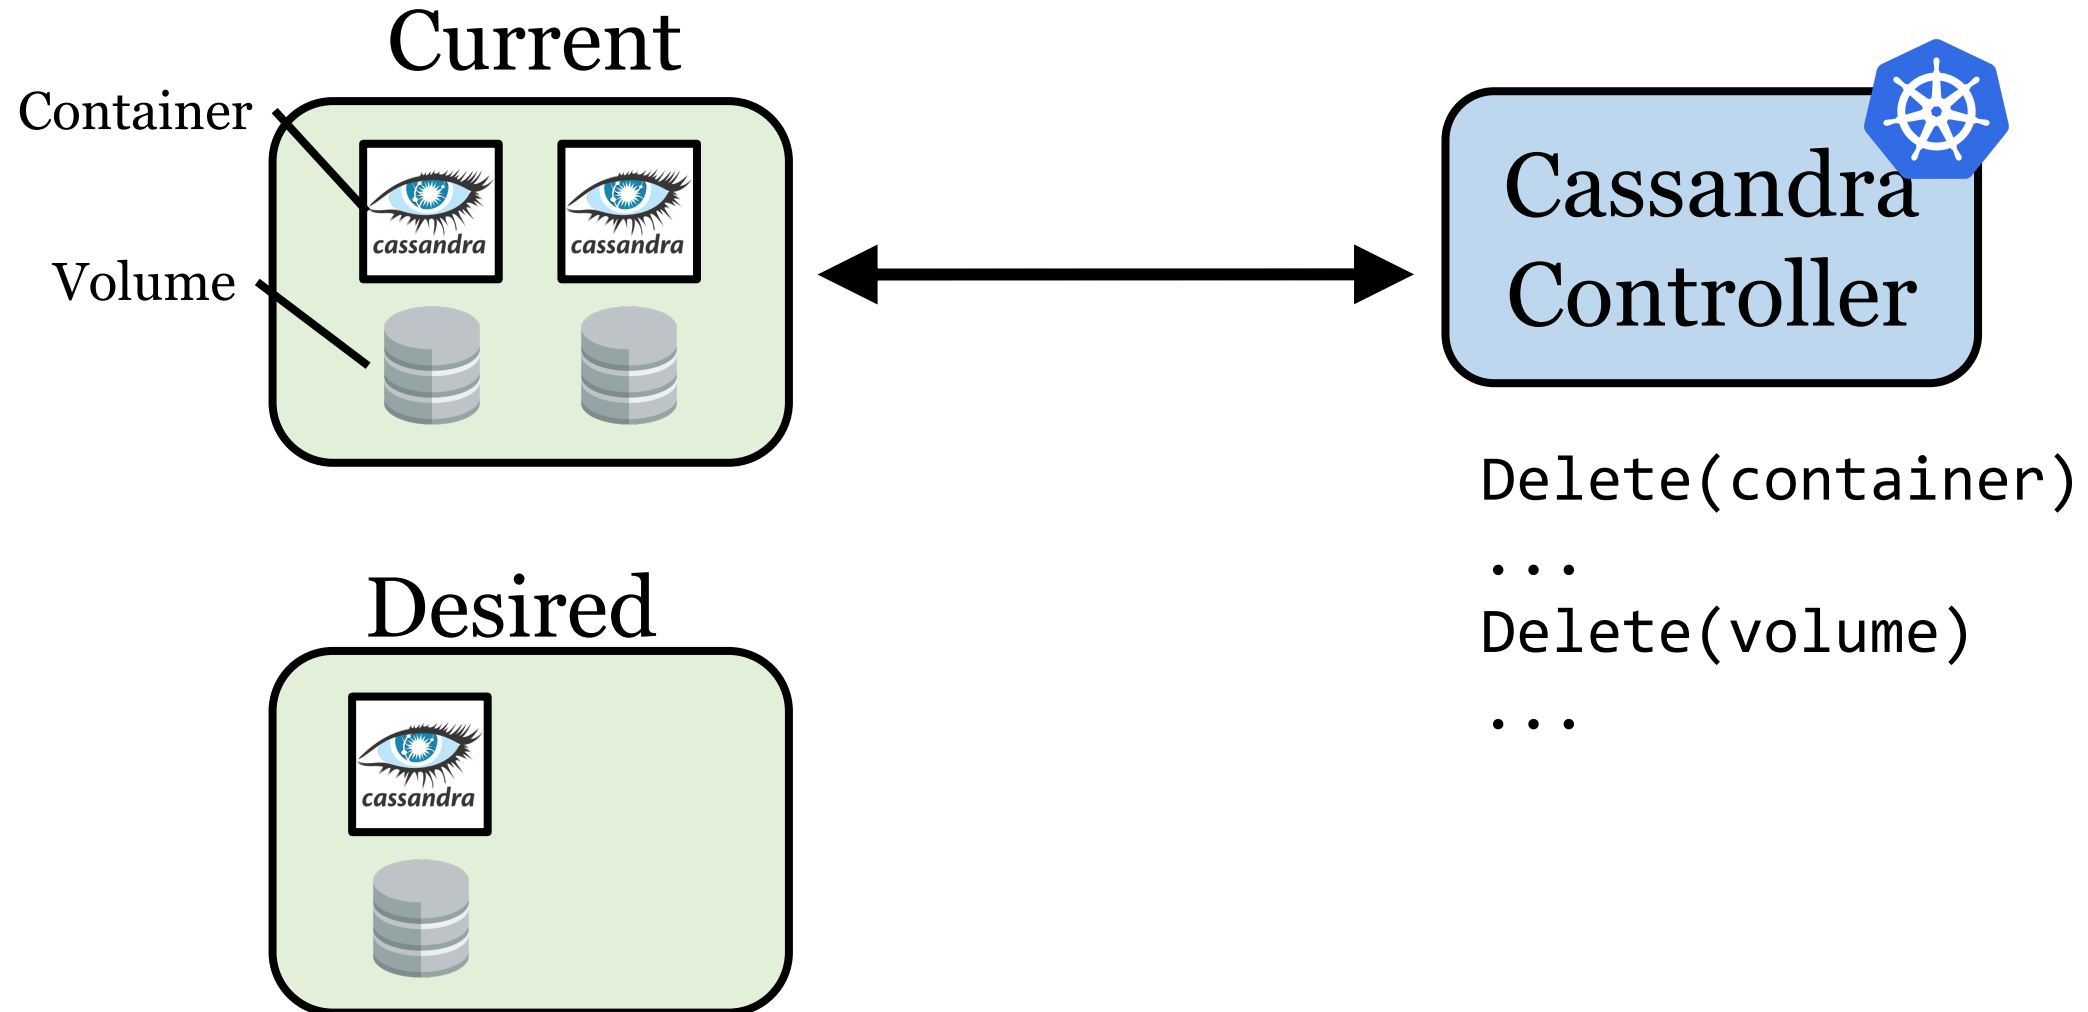


Thousands of controllers being developed  
by the Kubernetes community

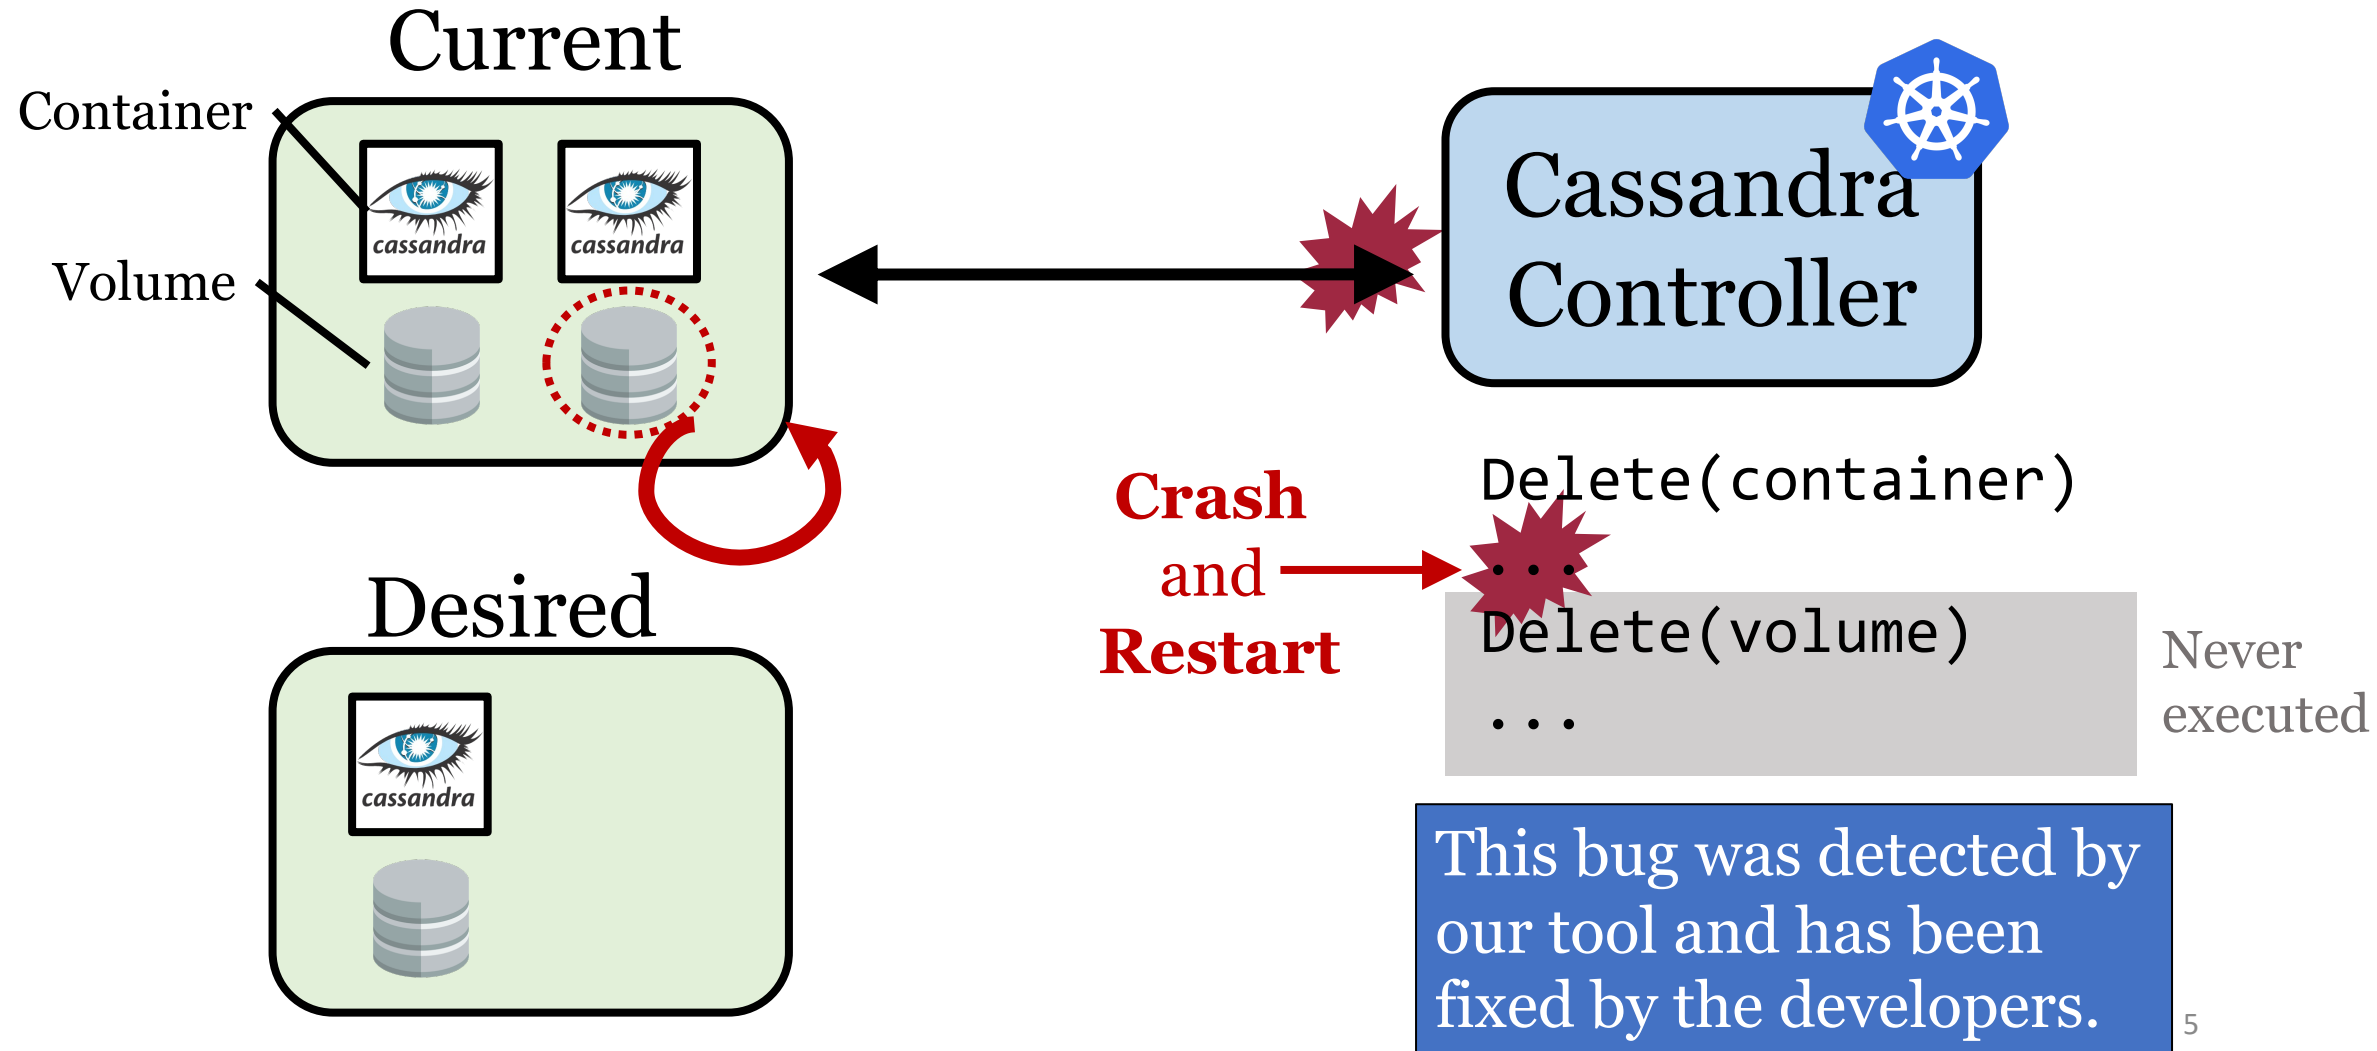
# Controllers implement state reconciliation



# Controllers implement state reconciliation



# Controller reliability is critical, but challenging!



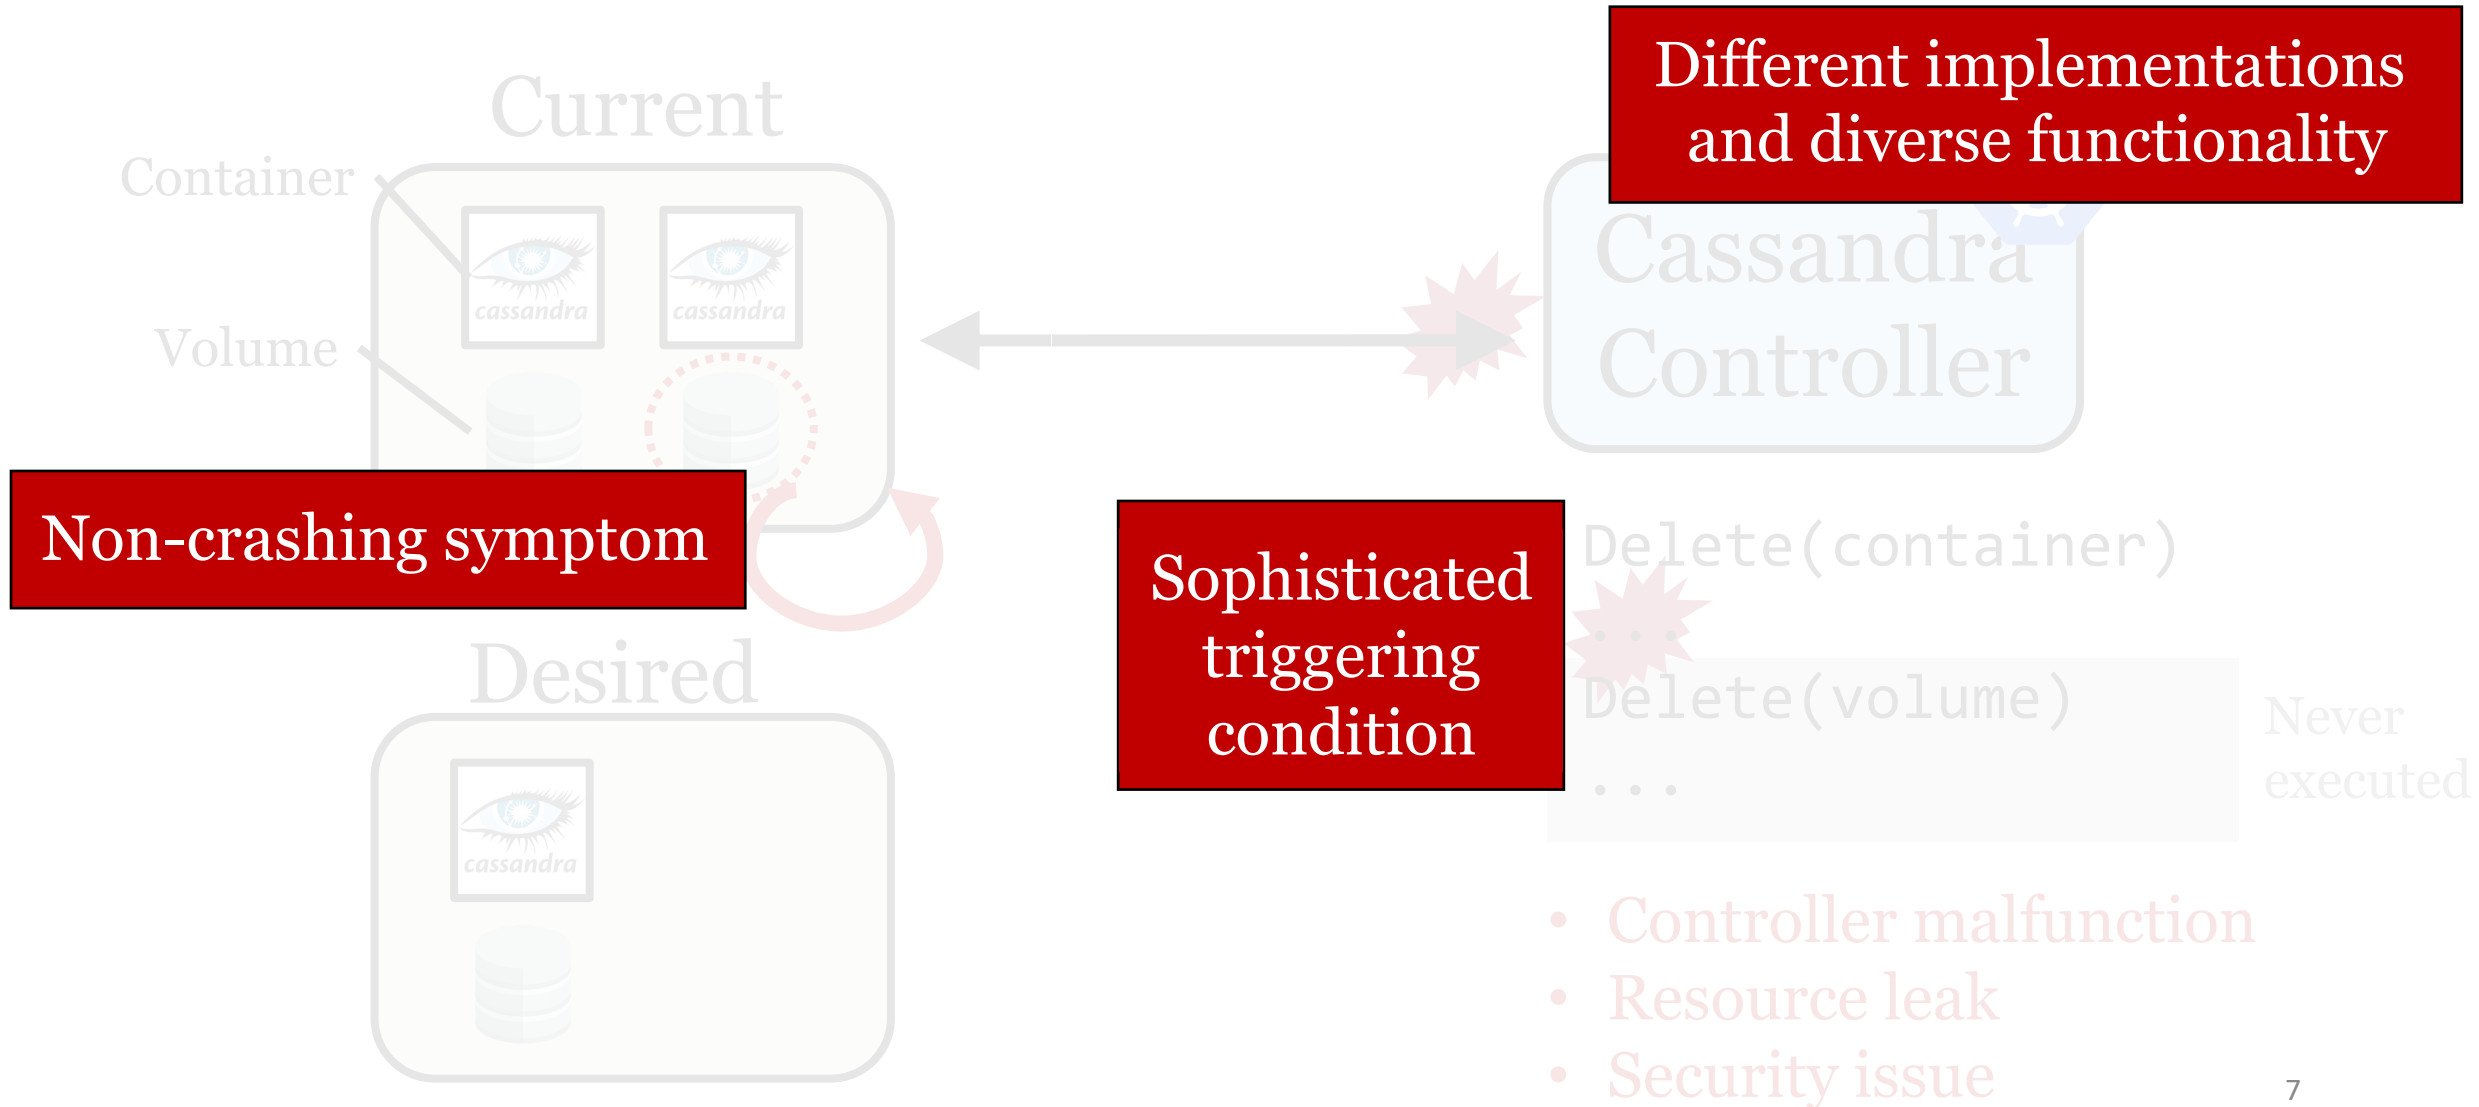


# Contributions

- Sieve: automatic reliability testing for Kubernetes controllers
  - **Key idea:** Perturbing the controller's view of cluster state
  - **Usability:** Testing unmodified controllers
  - **Reproducibility:** Reproducing detected bugs reliably
  - **Open sourced** at <https://github.com/sieve-project/sieve>
- Detected **46** serious bugs in 10 popular Kubernetes controllers
  - **Severe consequences:** System outage, data loss, security issues, etc.
  - **35** confirmed and **22** fixed




# Challenges of testing controllers





# Perturb the controller's view of cluster state

Reference run




Cluster state: Objects in  etcd

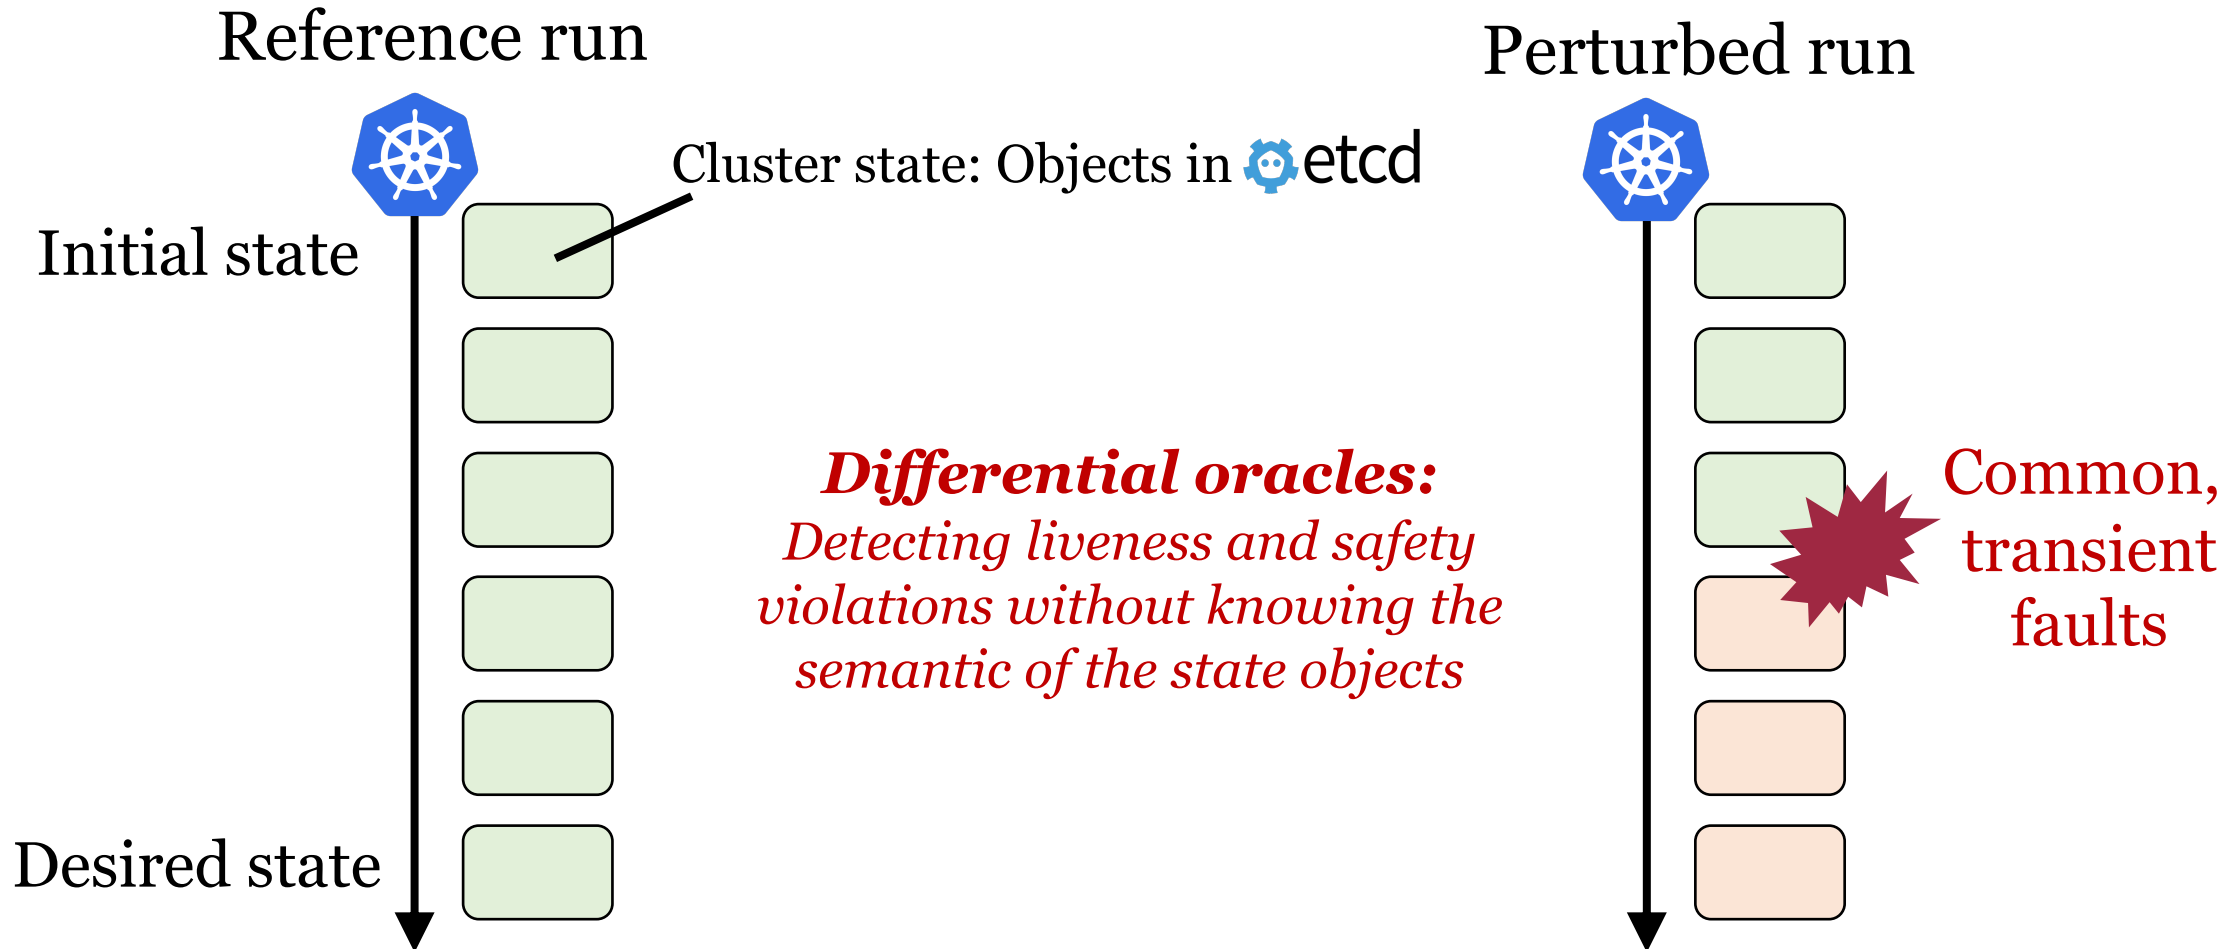
Every object creation/update/deletion advances the state

A controller makes reconciliation decisions based on its view of the current cluster state.

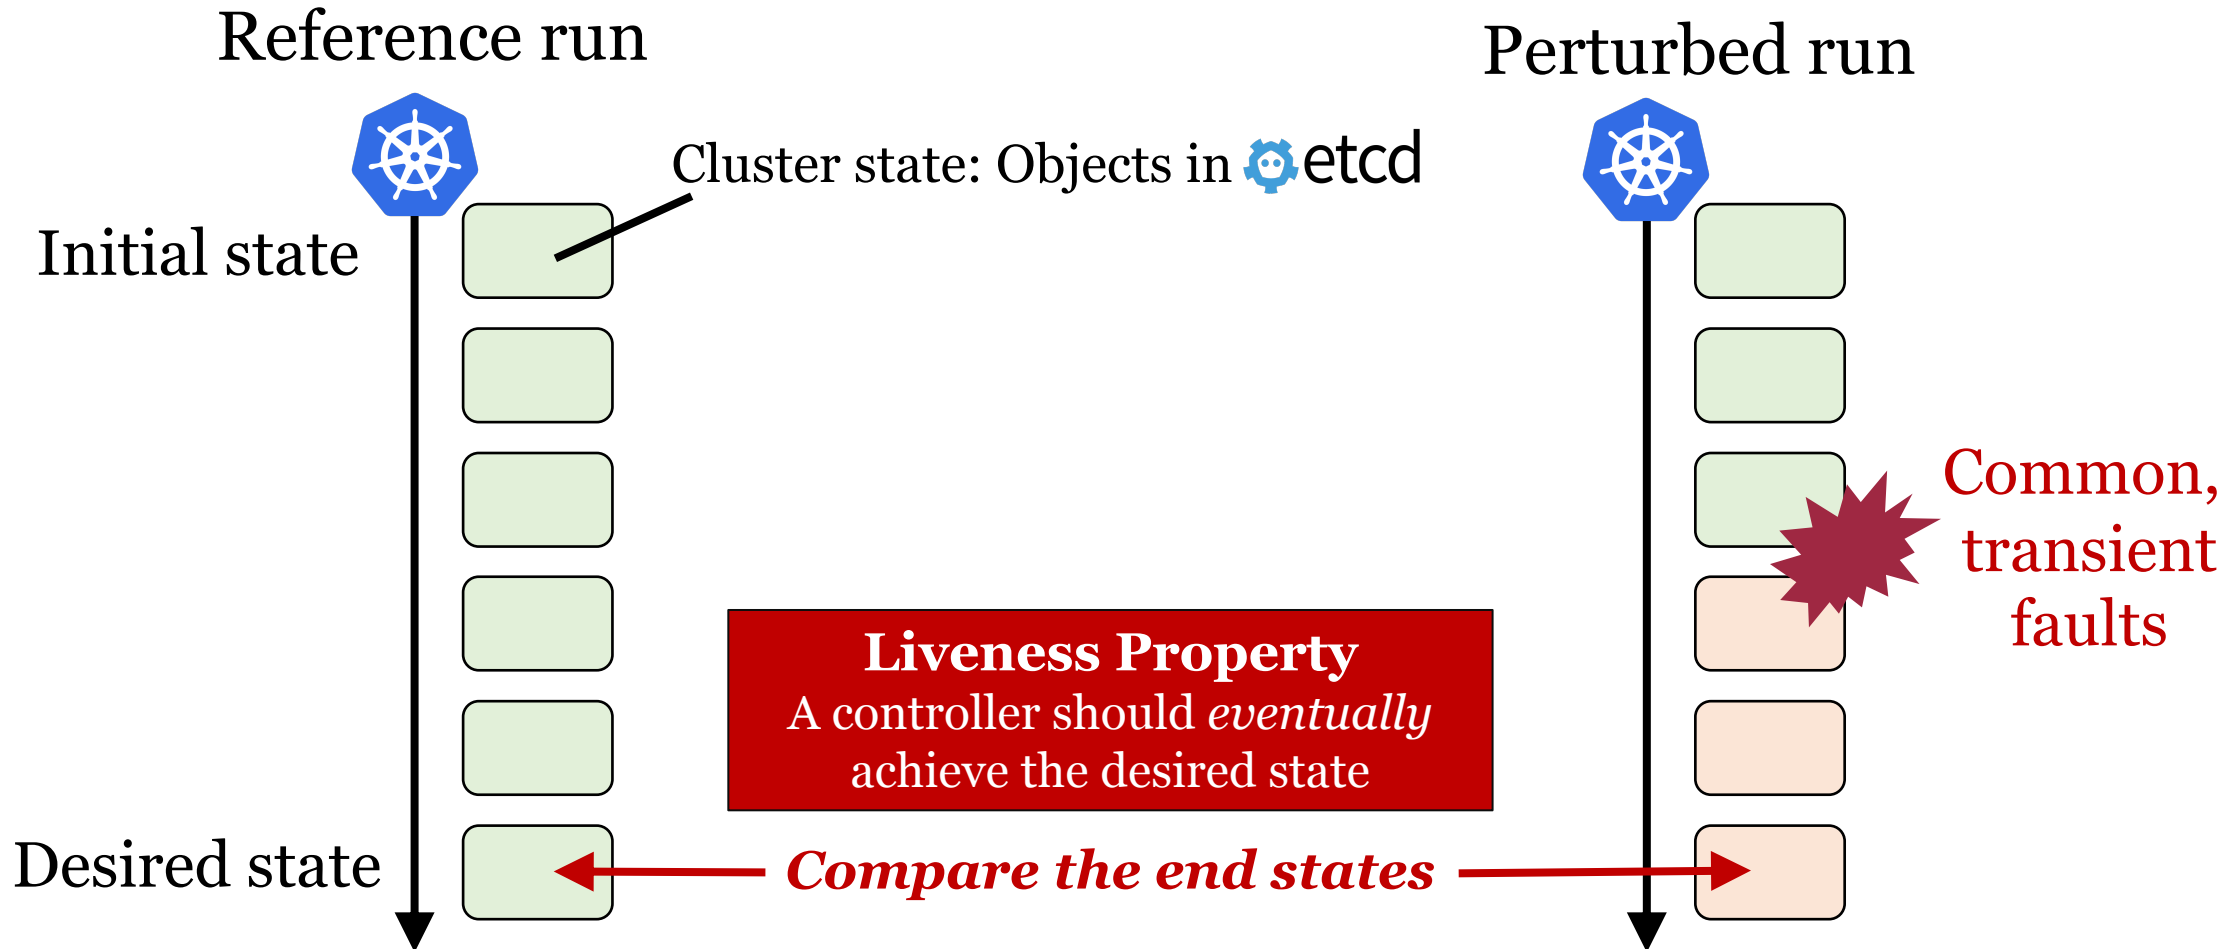
Perturbed run



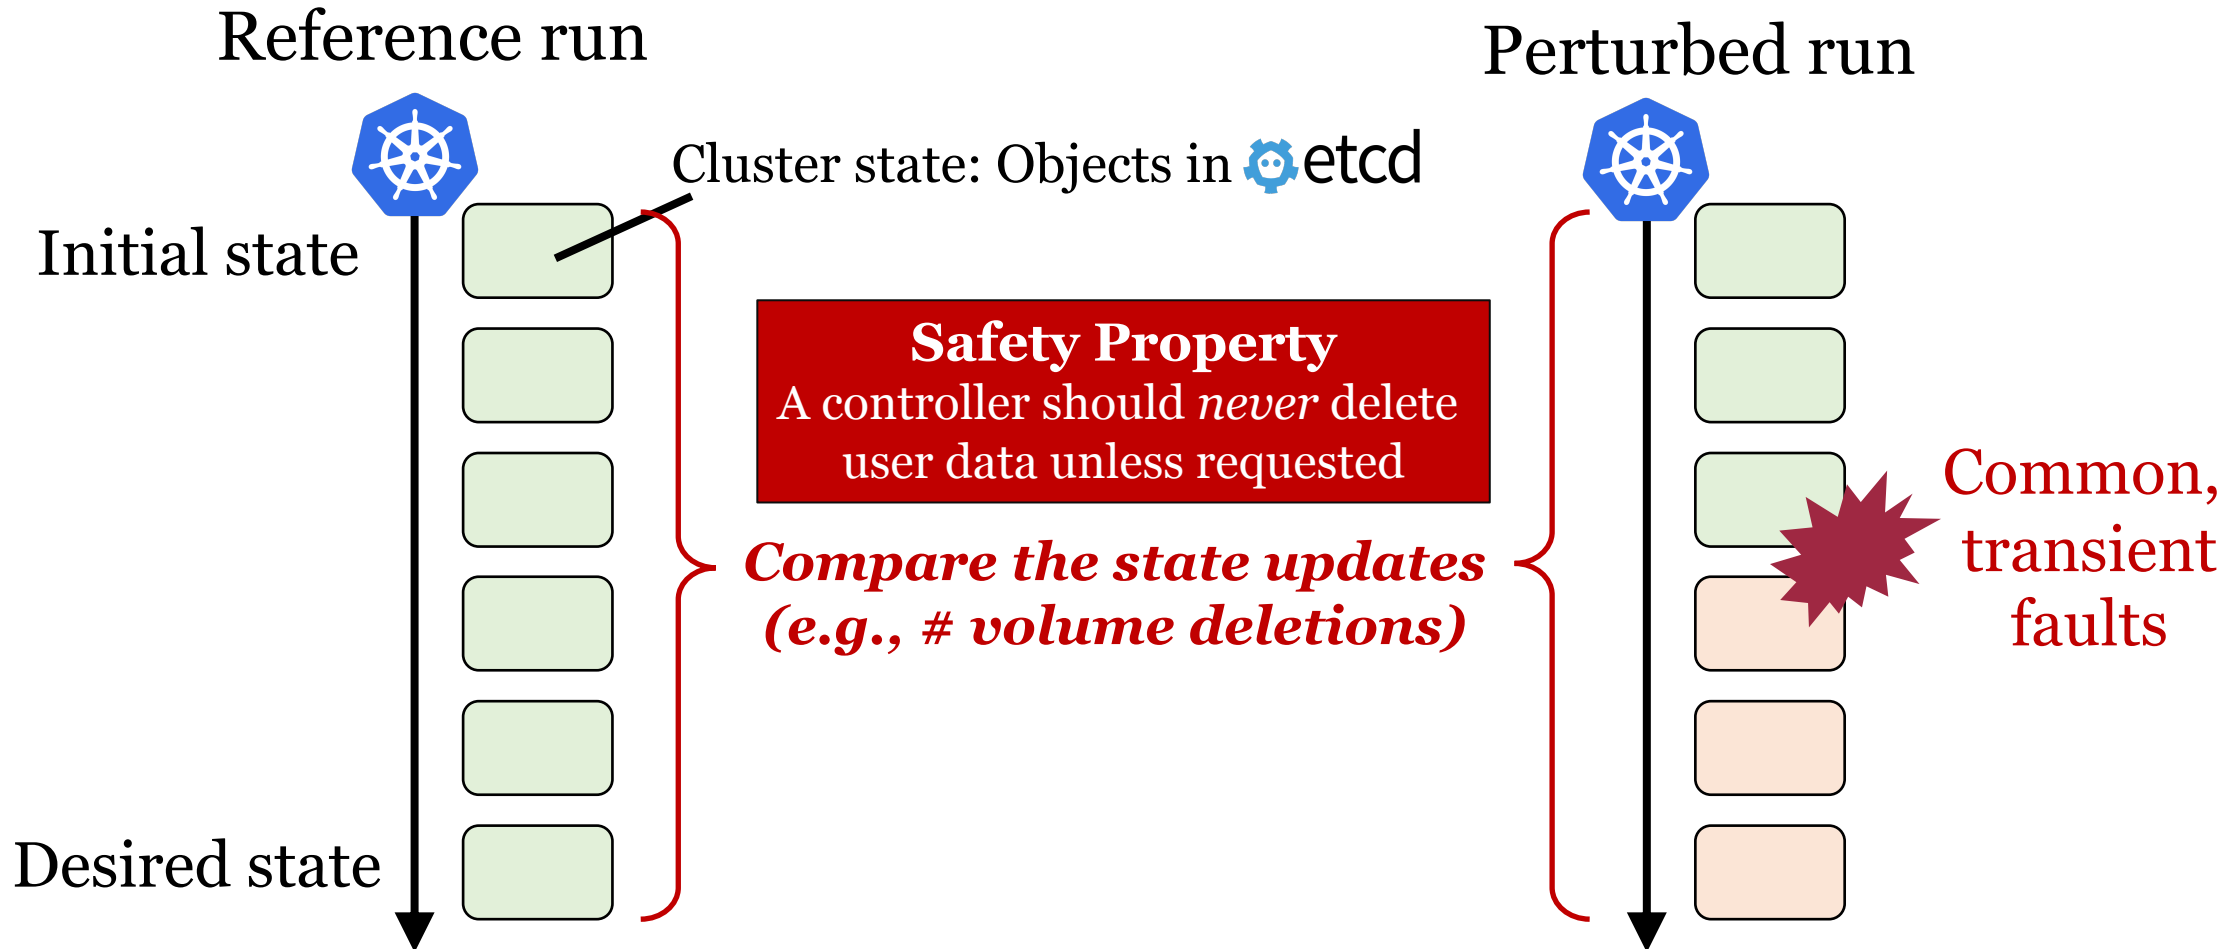
# Flag buggy behavior with *differential oracles*



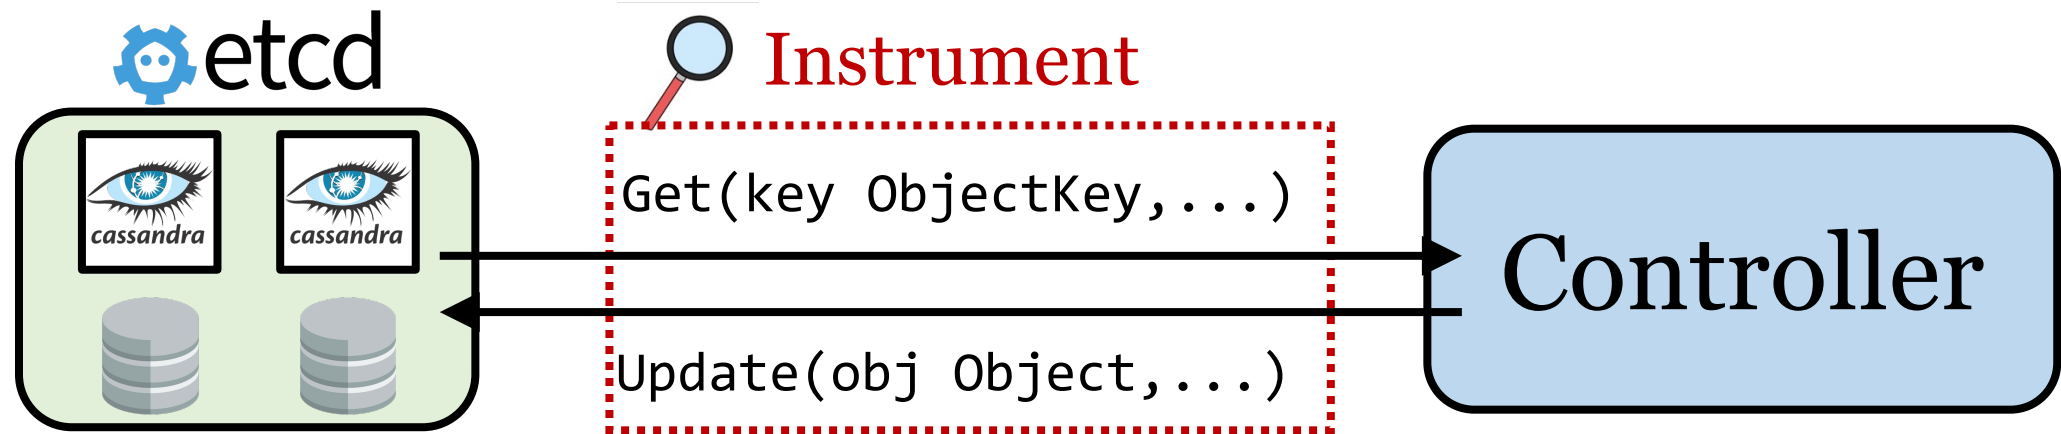
# Flag buggy behavior with *differential oracles*



# Flag buggy behavior with *differential oracles*



# Interposition around state-centric interface



- State-centric interface is used to read/write state objects
- Automatic interposition around cluster state transitions
- Allow Sieve to test *unmodified* controllers

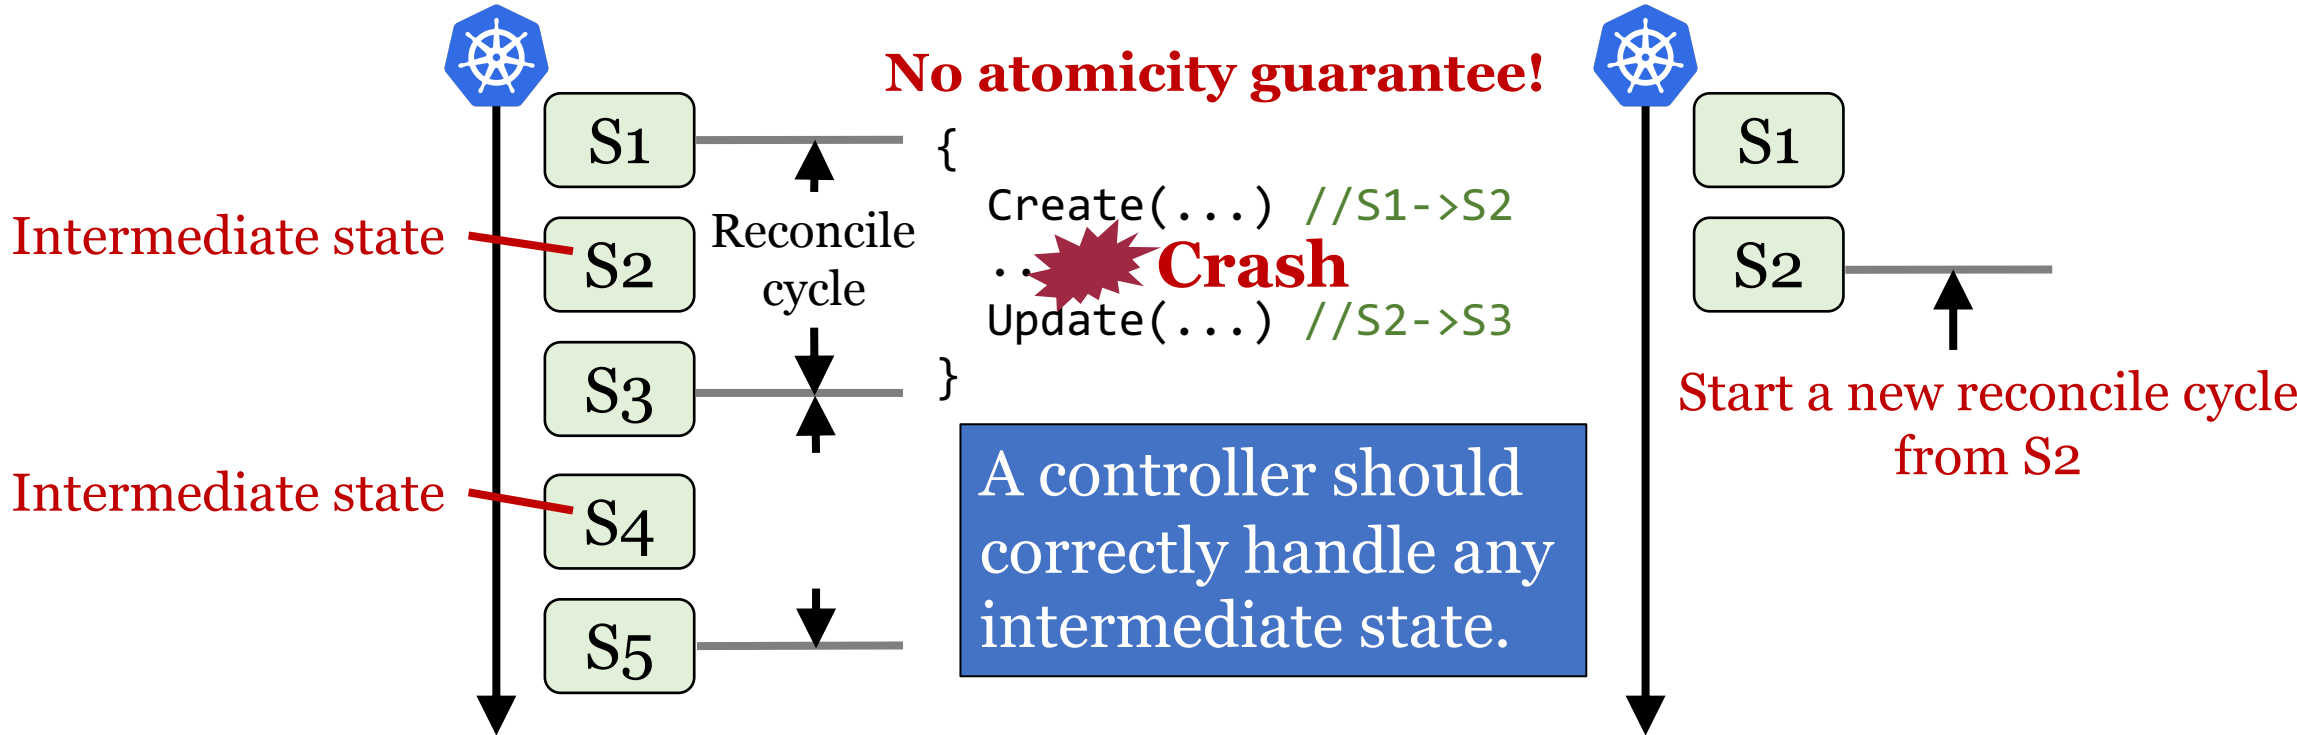
# Detect diverse controller bugs

- Employ **three perturbation patterns**
- **Exhaustively** test all bug-triggering perturbations
  - Systematically find all the targeted bugs
  - Inject faults with different timings
- Prune out ineffective perturbations to be **efficient**
  - Not every perturbation leads to bugs

# The intermediate-state pattern

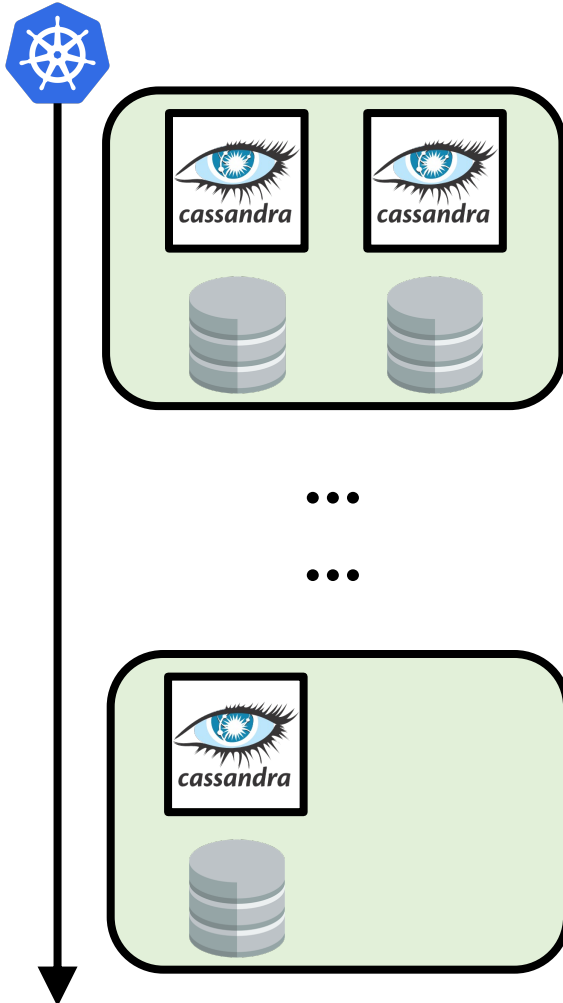
Reference run

Perturbed run



# An intermediate-state bug detected by Sieve

Reference run

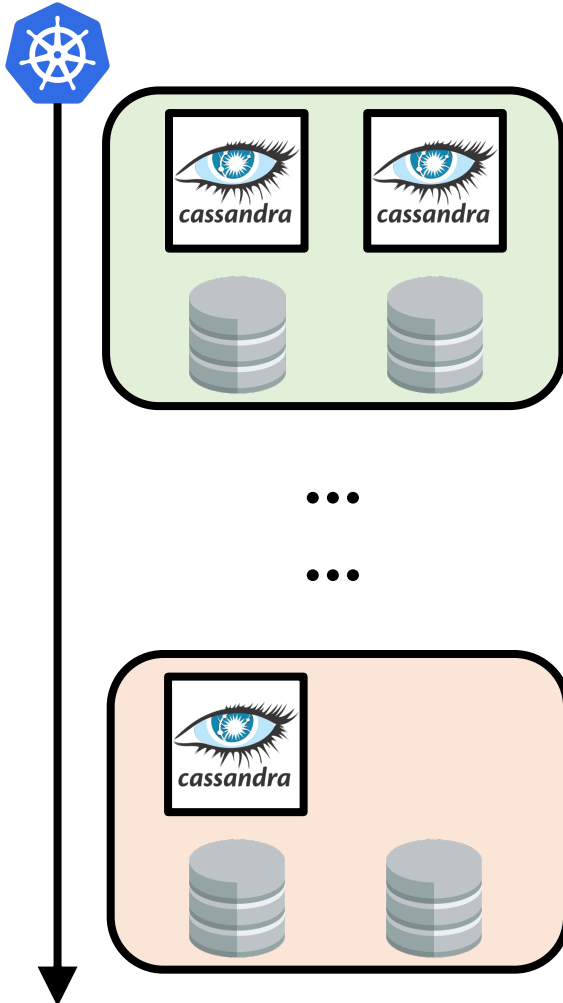


```
switch phase ← Uses phase to drive the reconciliation
case "Ongoing":
    if ContainerNotFound(container) {
        return Error("Container not found")
    }
    ...
    Delete(container)
    ...
    UpdatePhase("Finalizing") // phase <- "Finalizing"
    ...
case "Finalizing":
    ...
    Delete(volume)
    ...
    UpdatePhase("Done") // phase <- "Done"
}
// https://github.com/Orange-OpenSource/casskop
```



# An intermediate-state bug detected by Sieve

Perturbed run



```
switch phase {  
  case "Ongoing":  
    if ContainerNotFound(container) {  
      return Error("Container not found")  
    }  
    ...  
    Delete(container)  
    ...  
    UpdatePhase("Finalizing") // phase <- "Finalizing"  
    ...  
  case "Finalizing":  
    ...  
    Delete(volume)  
    ...  
    UpdatePhase("Done") // phase <- "Done"  
  }  
  // https://github.com/Orange-OpenSource/casskop
```


**Always returns here...**

**Crash**

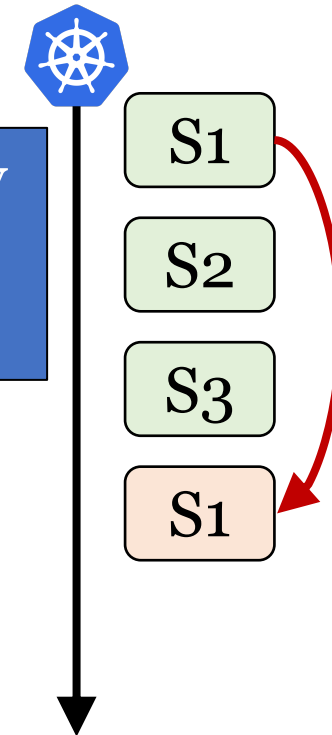
**Never executes this...**

# The stale-state pattern

Reference run



Perturbed run




A controller should correctly handle staleness caused by asynchrony and caching.

1. Inject delay to make a backup cache stale
2. Reconnect the controller to the stale cache

S1 replayed in the controller's view


# The unobserved-state pattern

Reference run



A controller should function correctly without observing every state.

Perturbed run



Inject delay to make the controller miss a state

S3 missed in the controller's view

# *Exhaustive* perturbation for each pattern

- **Key principle:** Inject faults at each execution point
- Run many different tests, each performing a different perturbation
  - Intermediate-state: Crash after *every* state update
  - Stale-state: Replay *every* stale state
  - Unobserved-state: Make the controller miss *every* state

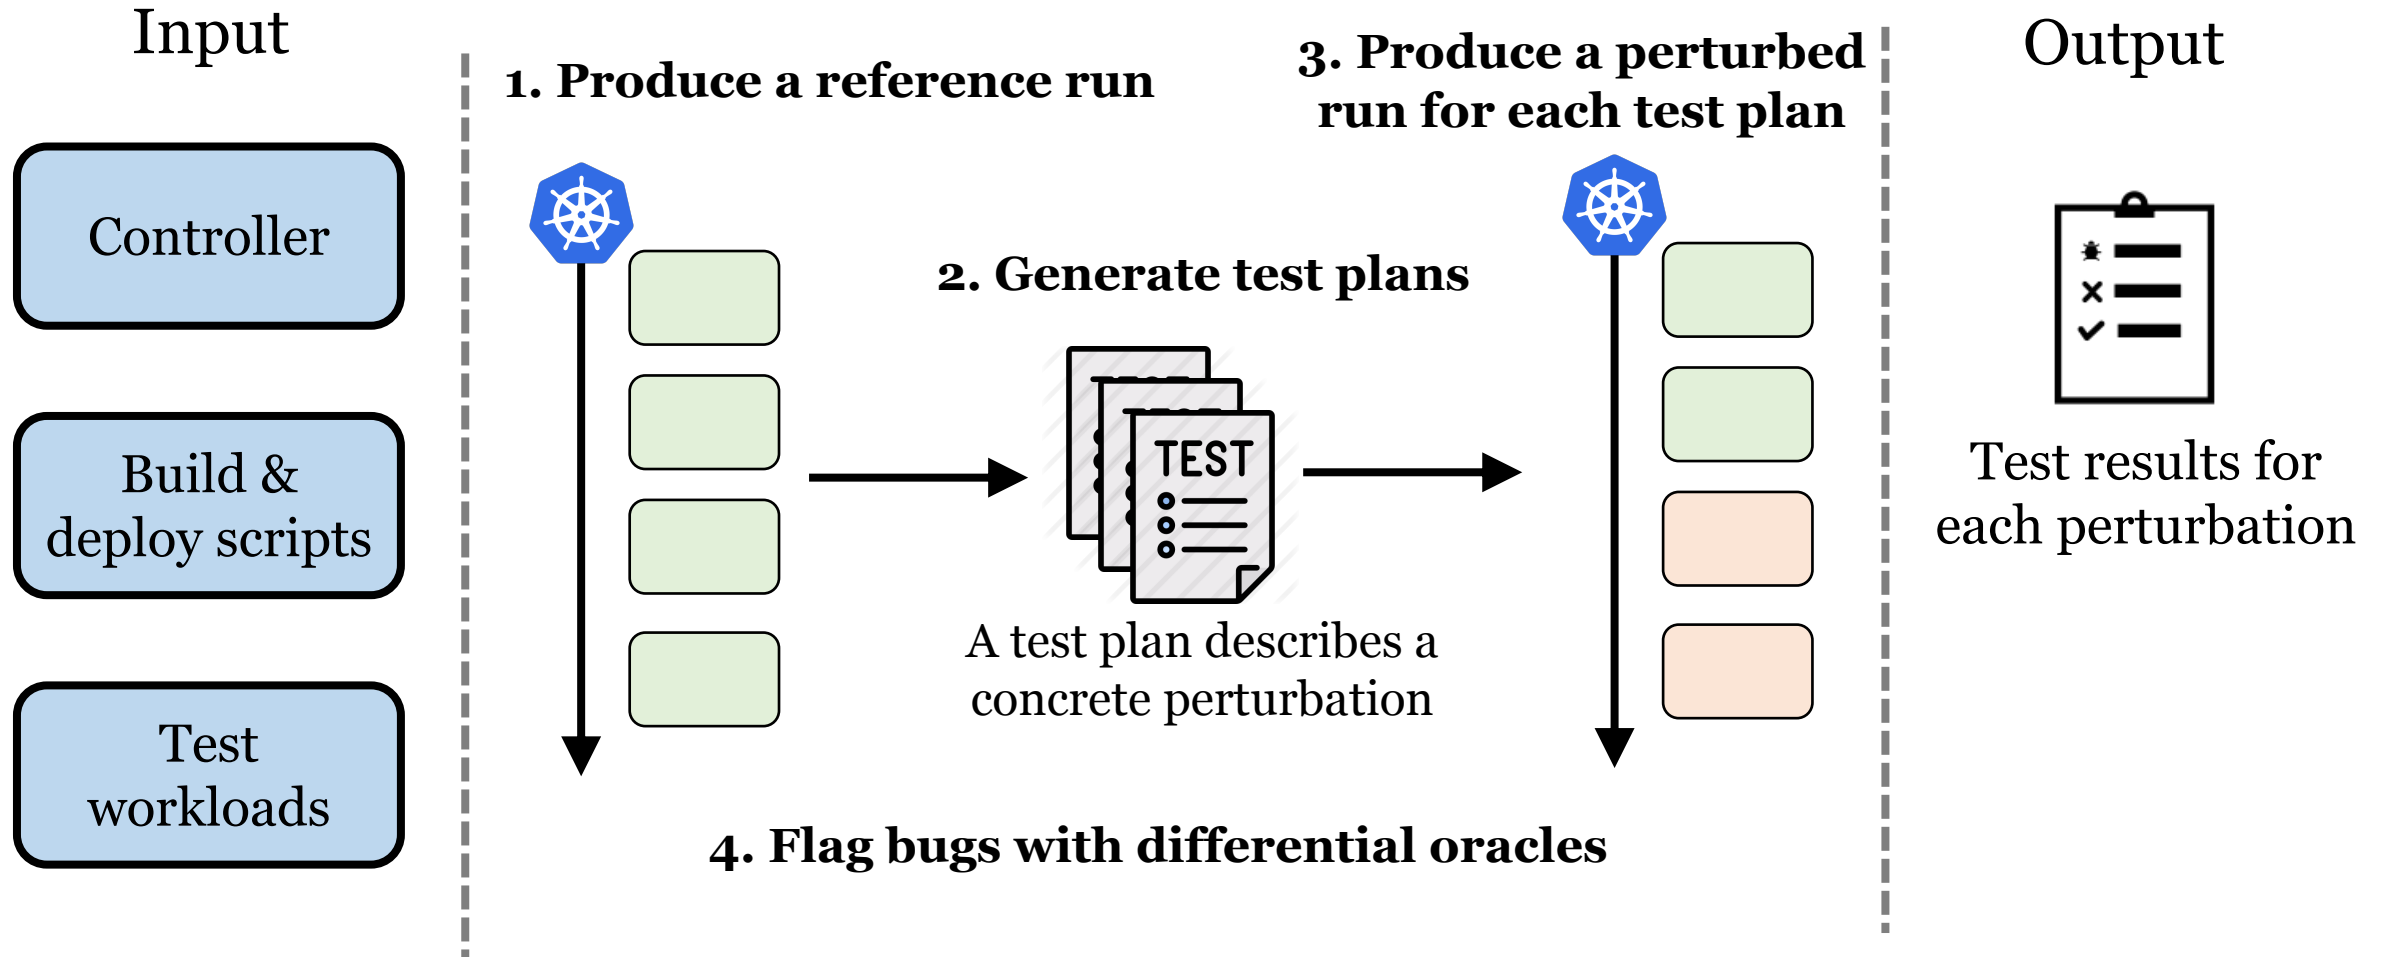
# Prune ineffective perturbations for efficiency

**Key principle:** Prune out perturbations that cannot affect a controller's behavior

- **Intermediate-state:** Prune out crashes that do not result in new intermediate states
- **Stale- and unobserved-state:** Avoid perturbing the state if observing the state does not causally lead to any controller effect
  - Reason about **causality** from state to effect

```
// Reconcile cycle
...
Create(...)
...
Update(...)
...
Delete(...) ← Delete a non-existing object
...
Update(...)
...
```

# Sieve end-to-end workflow



# Evaluation

- Applied Sieve to **10** popular Kubernetes controllers
- Can Sieve **effectively** find new bugs in real-world controllers?
  - Sieve found **46** bugs in **10** controllers
- Does Sieve do so **efficiently**?
  - Sieve pruned out **46% - 99%** of perturbations
  - Sieve tested each controller with a **nightly** run
- Are Sieve's testing results **trustworthy**?
  - Sieve had a low false positive rate of **3.5%**

# Finding *new* bugs



**35 confirmed; 22 fixed**

Controller		Intermediate state bugs	Stale state bugs	Unobserved state bugs	Indirect bugs	Total
cass-operator		2	1	0	0	3
cassandra-operator		0	2	1	2	5
casskop		1	2	1	0	4
elastic-operator		0	2	0	0	2
mongodb-operator		2	3	1	3	9
nifikop		2	0	0	1	3
rabbitmq-operator		1	2	1	0	4
xtradb-operator		3	3	1	0	7
yugabyte-operator		0	2	1	2	5
zookeeper-operator		0	2	1	1	4
<b>Total</b>		<b>11</b>	<b>19</b>	<b>7</b>	<b>9</b>	<b>46</b>





# Conclusion

- Controller reliability is critical but challenging!
- Sieve: automatic reliability testing for Kubernetes controllers
  - **Key idea:** Perturbing the controller's view of the cluster state
  - **Usability:** Testing unmodified controllers
  - **Reproducibility:** Reproducing detected bugs reliably
- **Open sourced** at <https://github.com/sieve-project/sieve>
  - Test your controller with Sieve!

