



# Best Practices For Disaster Recovery and High Availability

Stan Lukyanov

Customer Solutions, GridGain Systems




# Agenda



- Disaster recovery essentials
- Disaster recovery options for Apache Ignite and GridGain
- Choosing the solution


# What is Disaster Recovery?



# What is Disaster Recovery?




App




GridGain  
Cluster

# What is Disaster Recovery?



# What is Disaster Recovery?




# How to Compare DR Solutions?



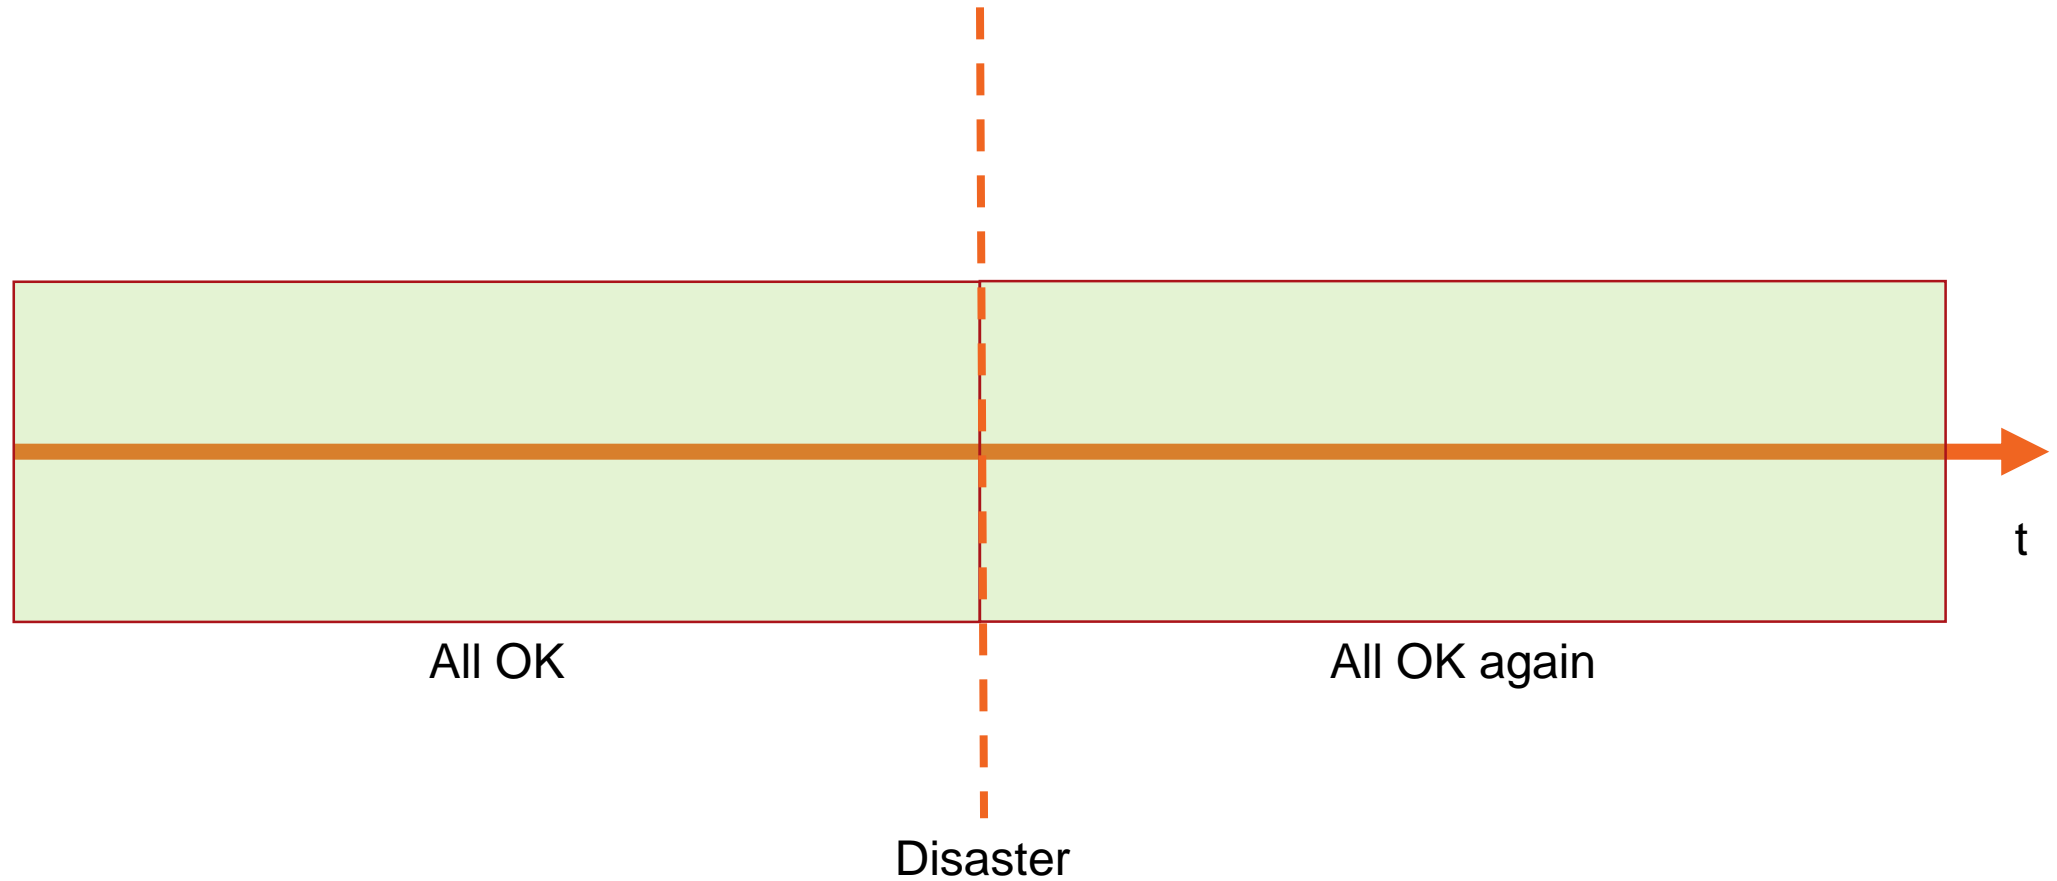
- How long the **service is down** in case of a disaster?
- Is there a **data loss** in case of a disaster?
- How much does the solution **cost**?
- Any additional **benefits** to implementing that?

# Disaster Timeline

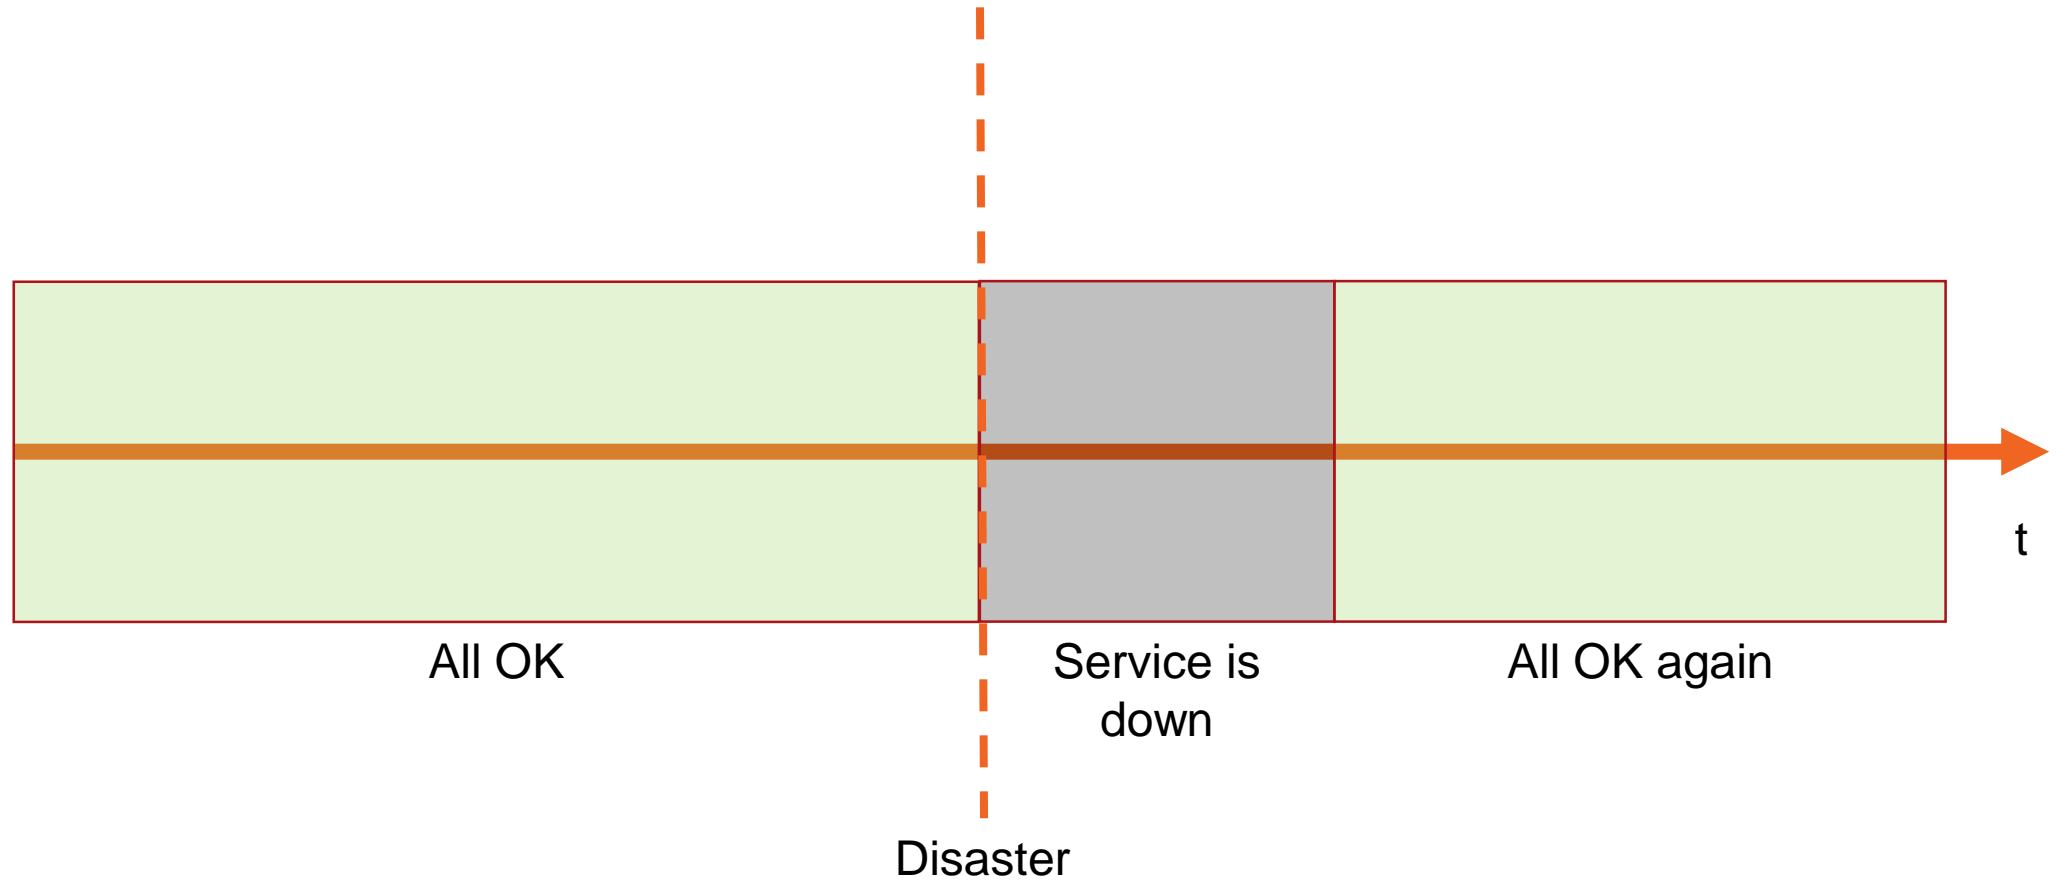





# Disaster Timeline



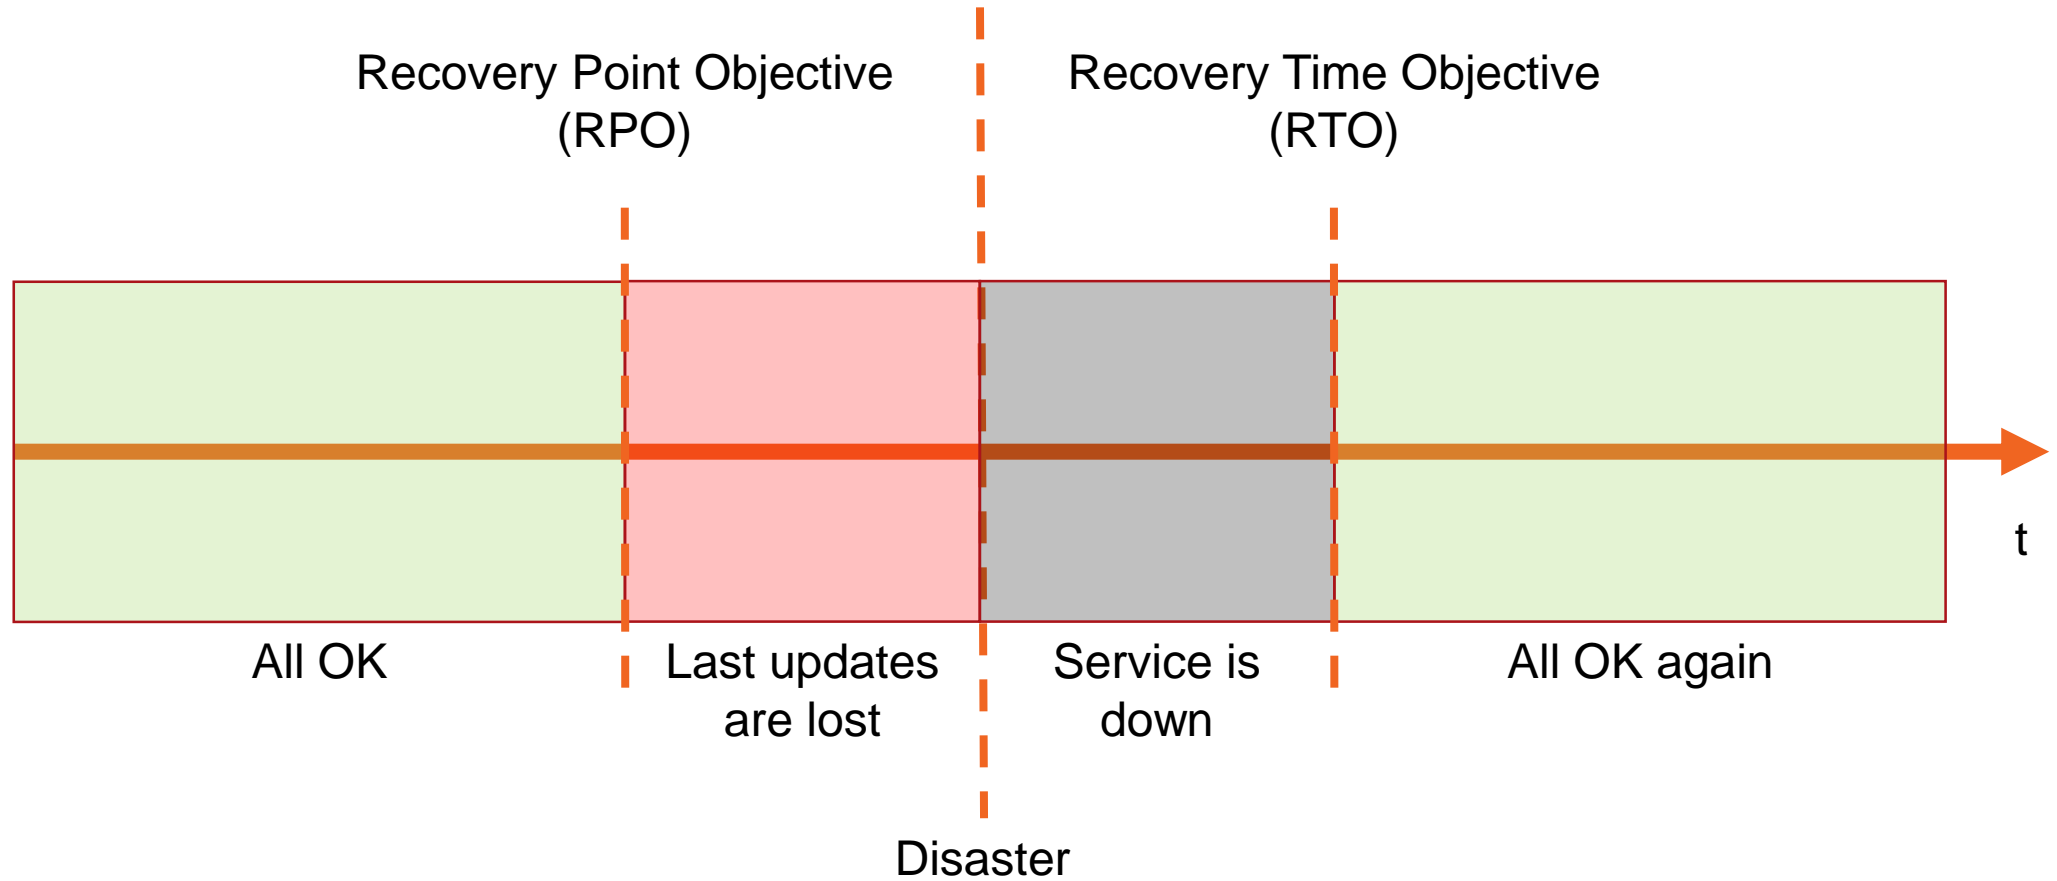
# Disaster Timeline



# Disaster Timeline



# Disaster Timeline



# Recovery Point Objective (RPO)



- Maximum time for which data is allowed to be lost
- Defined by the replication lag
- $RPO = 0$  – updates are replicated immediately, i.e. replication is synchronous

# Recovery Time Objective (RTO)



- Maximum allowed service interruption time
- $RTO = 0$  – second cluster is in standby AND client switches instantly
- In real life it is either:
  - $RTO = \text{failure detection time (seconds)}$  – if the second cluster is in standby
  - $RTO = \text{cluster startup time (minutes to hours)}$  – if the second cluster is dormant


# Comparison Parameters



Solution	RPO	RTO	Cost and Benefits
???	0 to hours	0* to hours	Minimal cost: <ul style="list-style-type: none"><li>- Second DC</li><li>- Cluster switchover logic</li><li>- Data replication logic</li></ul>


\* “RTO = 0” actually means “RTO = failure detection time”

# Disaster Recovery Options






# Disaster Recovery Options



# Option 1: Backup-Based DR



# Option 1: Backup-Based DR



# Option 1: Backup-Based DR



- RPO = backup period (hours)
- RTO = 0 if second DC is in standby, minutes to hours otherwise
- Requires a backup solution and a disk storage
- Requires custom cluster switchover
- Backups are generally useful
  - E.g. can be restored in a development environment

# Option 1: Backup-Based DR




## GridGain Solution: Snapshots

- Backups of a live cluster – no service disruption
- Incremental backups for more frequent backups – better RPO
- Automatic backup management (scheduling)
- Point-in-Time Recovery
- Only works with Native Persistence
- Available in GridGain Ultimate Edition

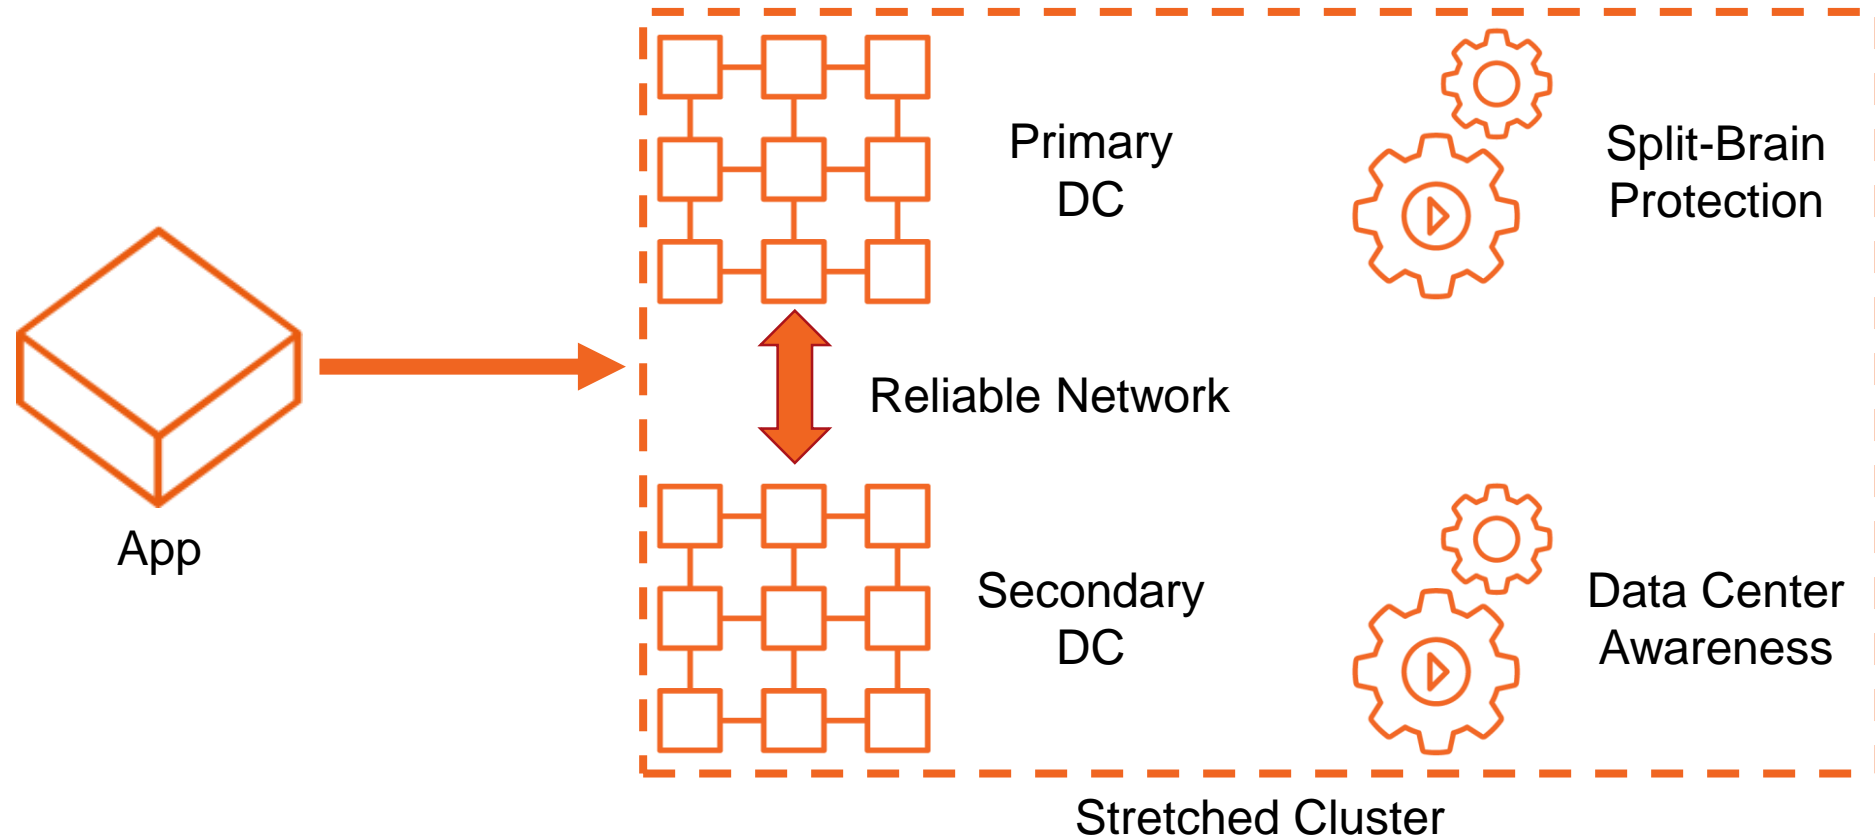
## Option 2: Stretched Cluster



# Option 2: Stretched Cluster



# Option 2: Stretched Cluster





## Option 2: Stretched Cluster




- RPO = 0 – if writes in the cluster are synchronous (they often are)
- RTO = 0 – client switches as if on a node failure
- Doesn't require cluster switchover nor replication solution
- Requires fast and reliable network between DCs
- Requires data center awareness
- Requires split-brain protection

# Option 2: Stretched Cluster




## Split-Brain Explained



# Option 2: Stretched Cluster




## Split-Brain Explained



# Option 2: Stretched Cluster



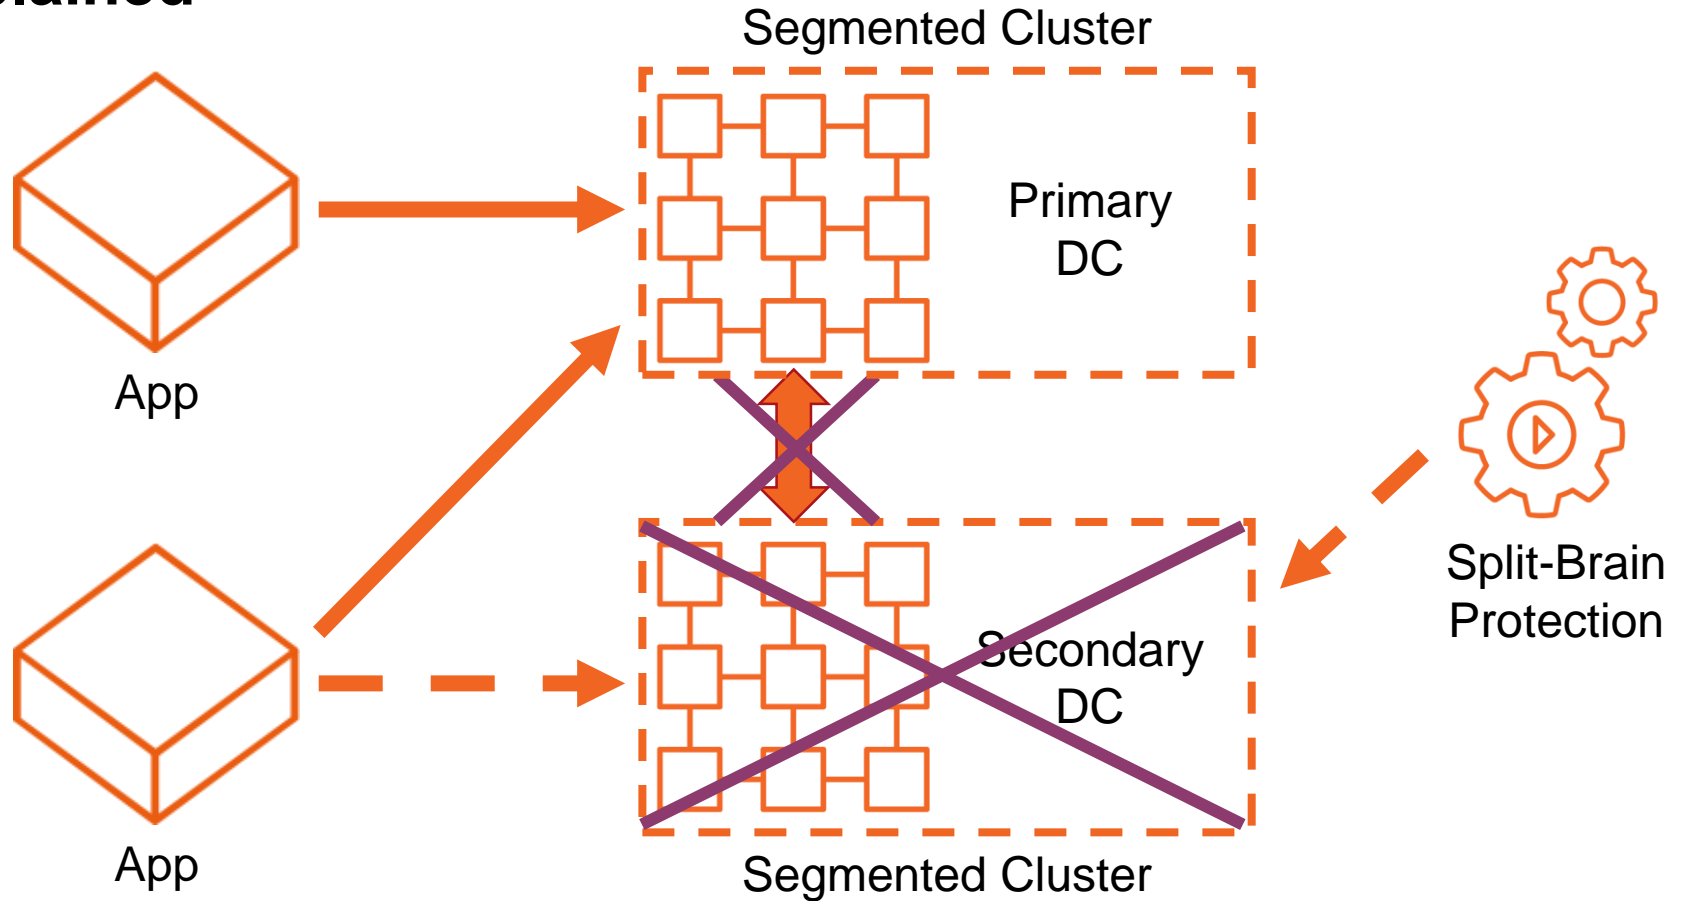
## Split-Brain Explained



# Option 2: Stretched Cluster



## Split-Brain Explained



# Option 2: Stretched Cluster



## GridGain Solution For Split-Brain Protection

### Solution 1: `TopologyValidator` + `SegmentationResolver`

- `TopologyValidator` prevents updates in the segmented part
- `SegmentationResolver` stops the segmented part
- Available in GridGain Enterprise Edition


### Solution 2: Zookeeper Discovery

- Zookeeper is responsible for keeping the cluster together
- Available in Apache Ignite and GridGain Community Edition

# Option 2: Stretched Cluster




## Data Center Awareness Explained



# Option 2: Stretched Cluster



## Data Center Awareness Explained






# Option 2: Stretched Cluster



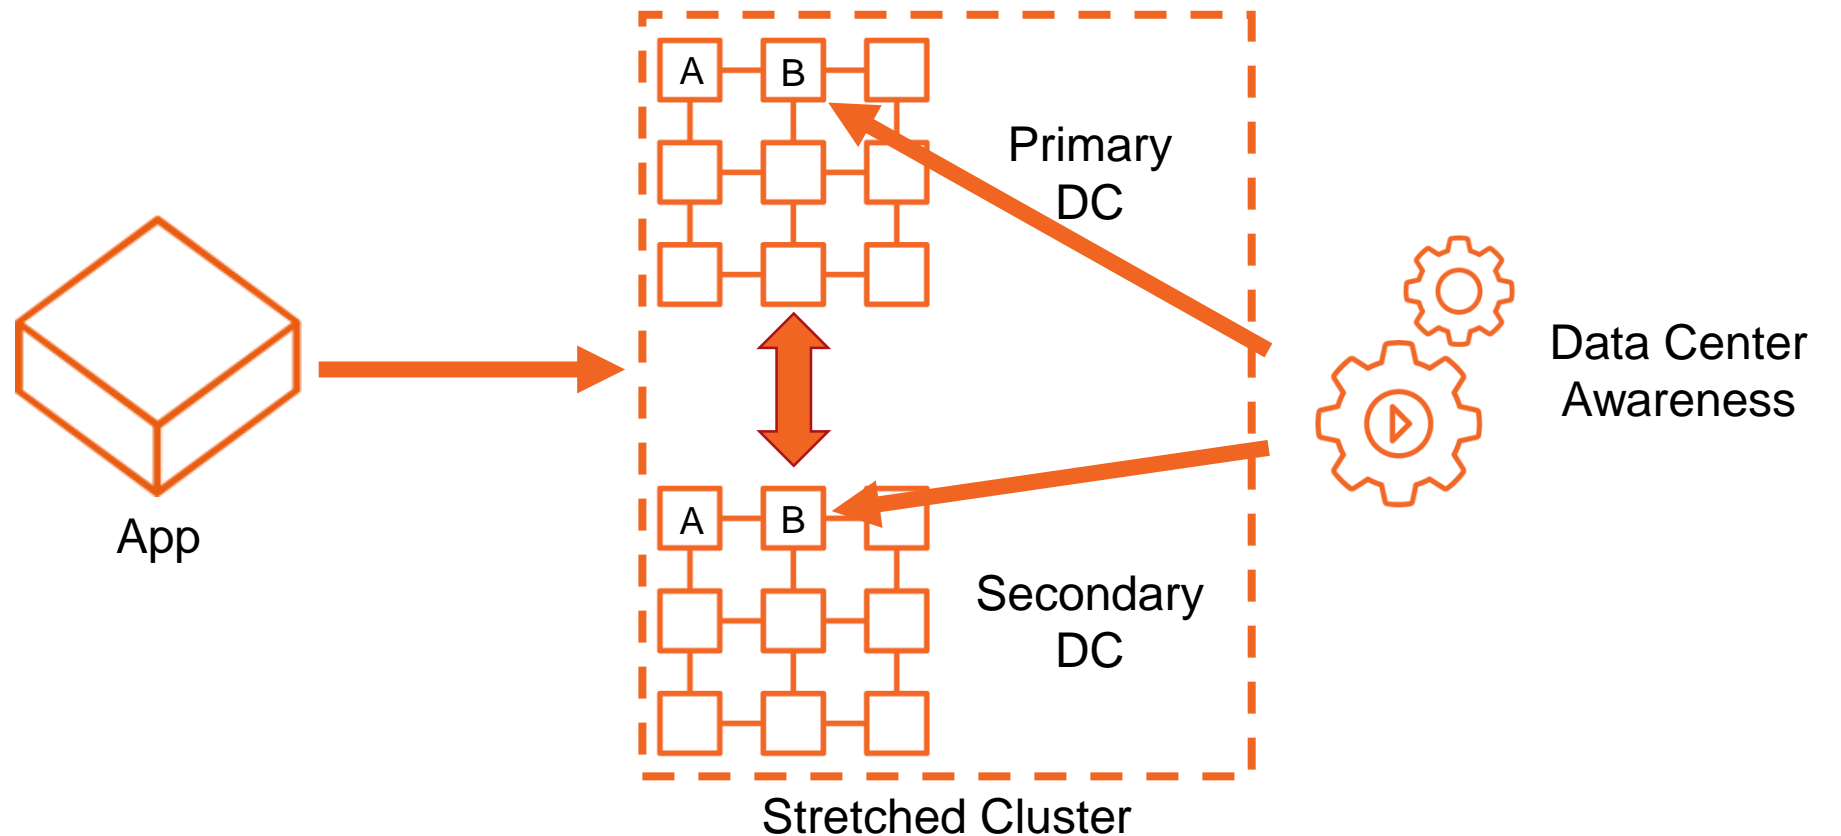
## Data Center Awareness Explained



# Option 2: Stretched Cluster



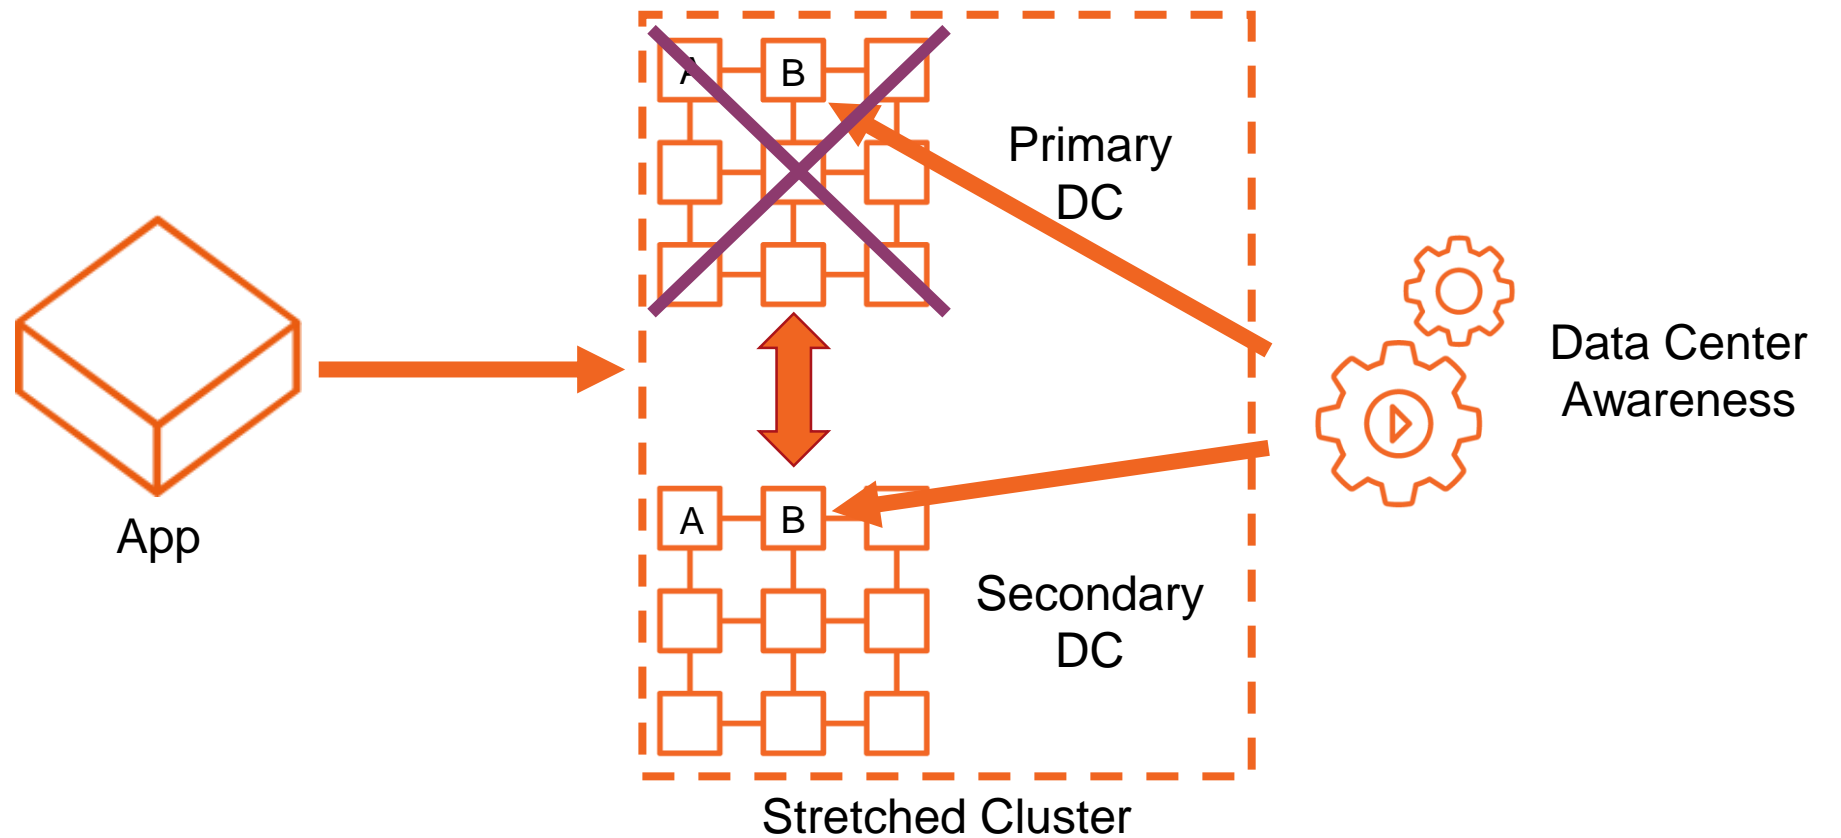
## Data Center Awareness Explained



# Option 2: Stretched Cluster



## Data Center Awareness Explained



# Option 2: Stretched Cluster



## GridGain Solution For Data Center Awareness

- `RendezvousAffinityFunction.affinityBackupFilter` controls distribution of backups
- Available in Apache Ignite and GridGain Community Edition

# Option 2: Stretched Cluster



## GridGain Solution For Data Center Awareness

### Step 1: Implement backup filter

```
class DcFilter implements IgniteBiPredicate<ClusterNode, List<ClusterNode>> {  
    @Override  
    public boolean apply(ClusterNode candidate, List<ClusterNode> assigned) {  
        String candidateDc = candidate.attribute("dc");  
        String primaryDc = assigned.get(0).attribute("dc");  
        return !Objects.equals(candidateDc, primaryDc);  
    }  
}
```

# Option 2: Stretched Cluster



## GridGain Solution For Data Center Awareness

### Step 2: Configure the filter in each cache

```
<property name="affinity">
  <bean
class="org.apache.ignite.cache.affinity.rendezvous.RendezvousAffinityFunction">
    <property name="affinityBackupFilter">
      <bean class="com.mycompany.DcFilter"/>
    </property>
  </bean>
</property>
```

# Option 2: Stretched Cluster




## GridGain Solution For Data Center Awareness

### Step 3: Assign a value to each node


```
<property name="userAttributes">  
  <map>  
    <entry key="dc" value="dc-1"/>  
  </map>  
</property>
```

# Option 3: Streaming





# Option 3: Streaming



# Option 3: Streaming



- RPO = seconds – replication lag
  - Could be 0 but synchronous replication is usually a bad idea
- RTO = 0
- Requires a streaming platform and/or embedded change-data-capture functionality

# Option 3: Streaming



## GridGain Solution: Data Center Replication

- No additional software
- Active-Passive or Active-Active
- Allows for complex topologies, up to 32 data centers
- Available in GridGain Enterprise Edition
- Data Center Replication docs:  
<https://docs.gridgain.com/docs/data-center-replication>

# Option 3: Streaming



## GridGain Solution: Data Center Replication

### Step 1: Configure caches on the Sender side

```
<bean class="org.apache.ignite.configuration.IgniteConfiguration">
  <property name="cacheConfiguration">
    <bean class="org.apache.ignite.configuration.CacheConfiguration">
      <property name="name" value="crossDrCache"/>
      <property name="pluginConfigurations">
        <bean class="org.gridgain.grid.configuration.GridGainCacheConfiguration">
          <property name="drSenderConfiguration">
            <bean class="org.gridgain.grid.cache.dr.CacheDrSenderConfiguration"/>
          </property>
        </bean>
      </property>
    </bean>
  </property>
</bean>
```

# Option 3: Streaming



## GridGain Solution: Data Center Replication

### Step 2: Configure a Sender Hub

```
<!-- In IgniteConfiguration.pluginConfigurations -->
<bean class="org.gridgain.grid.configuration.GridGainConfiguration">
  <property name="dataCenterId" value="1"/>
  <property name="drSenderConfiguration">
    <bean class="org.gridgain.grid.configuration.DrSenderConfiguration">
      <property name="connectionConfiguration">
        <bean class="org.gridgain.grid.dr.DrSenderConnectionConfiguration">
          <property name="dataCenterId" value="2"/>
          <property name="receiverAddresses" value="172.16.2.100:50001"/>
        </bean>
      </property>
    </bean>
  </property>
</bean>
```

# Option 3: Streaming



## GridGain Solution: Data Center Replication

### Step 3: Configure caches on the Receiver side

```
<bean class="org.apache.ignite.configuration.IgniteConfiguration">
  <property name="cacheConfiguration">
    <bean class="org.apache.ignite.configuration.CacheConfiguration">
      <property name="name" value="crossDrCache"/>
      <property name="pluginConfigurations">
        <bean class="org.gridgain.grid.configuration.GridGainCacheConfiguration">
          <property name="drReceiverEnabled" value="true"/>
        </bean>
      </property>
    </bean>
  </property>
</bean>
```

# Option 3: Streaming



## GridGain Solution: Data Center Replication

### Step 4: Configure a Receiver Hub

```
<!-- In IgniteConfiguration.pluginConfigurations -->
<bean class="org.gridgain.grid.configuration.GridGainConfiguration">
  <property name="dataCenterId" value="2"/>
  <property name="drReceiverConfiguration">
    <bean class="org.gridgain.grid.configuration.DrReceiverConfiguration">
      <property name="localInboundHost" value="172.16.2.100"/>
      <property name="localInboundPort" value="50001"/>
    </bean>
  </property>
</bean>
```

# Option 3: Streaming




## GridGain Solution: Kafka Connector


- Certified by Confluent
  - Requires a Kafka instance deployed separately
  - Maximum flexibility
  - Available in GridGain Enterprise Edition
- 
- Detailed guide for GridGain DR using Kafka:  
<https://docs.gridgain.com/docs/certified-kafka-connector-examples-dr>



# Option 4: System-Level Replication



# Option 4: System-Level Replication



# Option 4: System-Level Replication




- RPO = 0 to minutes – depends on the vendor
- RTO = minutes – usually requires cluster restart
- Requires virtualized/cloud environment – significant operations effort and costs

# Option 5: Application-Level Replication



# Option 5: Application-Level Replication




# Option 5: Application-Level Replication



- RPO = 0 – synchronous writes to both DCs
- RTO = 0 – always connected to both DCs
- Requires significant development effort – full DIY

# Choosing The Solution



# Solutions Comparison



Solution	RPO	RTO	Cost and Benefits
Backup-Based	Hours	0 to minutes	<ul style="list-style-type: none"><li>- Cost: Backup solution and disk storage</li><li>- Cost: Custom switchover</li><li>- Benefit: Backups are generally useful</li></ul>
Stretched Cluster	0	0	<ul style="list-style-type: none"><li>- Cost: Huge reliance on network</li><li>- Cost: Split-brain protection and data center awareness</li></ul>
GridGain DCR	Seconds	0	<ul style="list-style-type: none"><li>- Cost: Custom switchover</li><li>- Benefit: No additional components for replication</li></ul>
GridGain Kafka Connector	Seconds	0	<ul style="list-style-type: none"><li>- Cost: Kafka</li><li>- Cost: Custom switchover</li><li>- Benefit: Flexible, allows heterogenous consumers</li></ul>
System-Level	0 to minutes	Minutes	<ul style="list-style-type: none"><li>- Cost: VM/Cloud solution</li><li>- Cost: Huge reliance on network</li></ul>
Application-Level	0	0	<ul style="list-style-type: none"><li>- Cost: Full DIY</li></ul>




# Choosing The Solution For GridGain DR



- Default – go with GridGain Data Center Replication
  - Works out of the box, no additional components
- If RPO requirements allow – go with GridGain Snapshots
  - Simple and powerful
- If heterogenous receivers are required – go with GridGain Kafka Connector
  - Flexible and robust
- If already running on VM or in the Cloud – check their guarantees
  - May fit out of the box – but if not, don't worth it just for DR

# Q&A



# Apache Ignite Resources



- Apache Ignite documentation
  - <https://apacheignite.readme.io/docs>
  - Apache Ignite community resources
  - [user@ignite.apache.org](mailto:user@ignite.apache.org) – the mailing list
  - <https://ignite.apache.org/community/resources.html> – other resources and instructions
  - <http://apache-ignite-users.70518.x6.nabble.com> – forum and archive
  - <https://stackoverflow.com/questions/tagged/ignite> – StackOverflow questions

# GridGain Resources



## Moving Apache® Ignite™ into Production webinars

- [Initial Checklist](#)
- [Best Practices for Native Persistence and Data Recovery](#)
- [Best Practices for Monitoring Distributed In-Memory Computing](#)
- [Best Practices for Deploying Apache Ignite in the Cloud](#)

**GridGain forums:** <https://forums.gridgain.com>

**GridGain documentation:** <https://docs.gridgain.com/docs>

# Contact me



- [stan@gridgain.com](mailto:stan@gridgain.com)
- [stanlukyanov@gmail.com](mailto:stanlukyanov@gmail.com)