GridGain

Moving Ignite into Production: An Initial Checklist

Denis Magda,
Ignite PMC Chair,
GridGain VP of Product

Agenda

- High-availability and fault-tolerance
- Capacity planning and basic tuning
- Upgrades and patch releases
- Production support
- Q&A

Apache Ignite and GridGain Overview

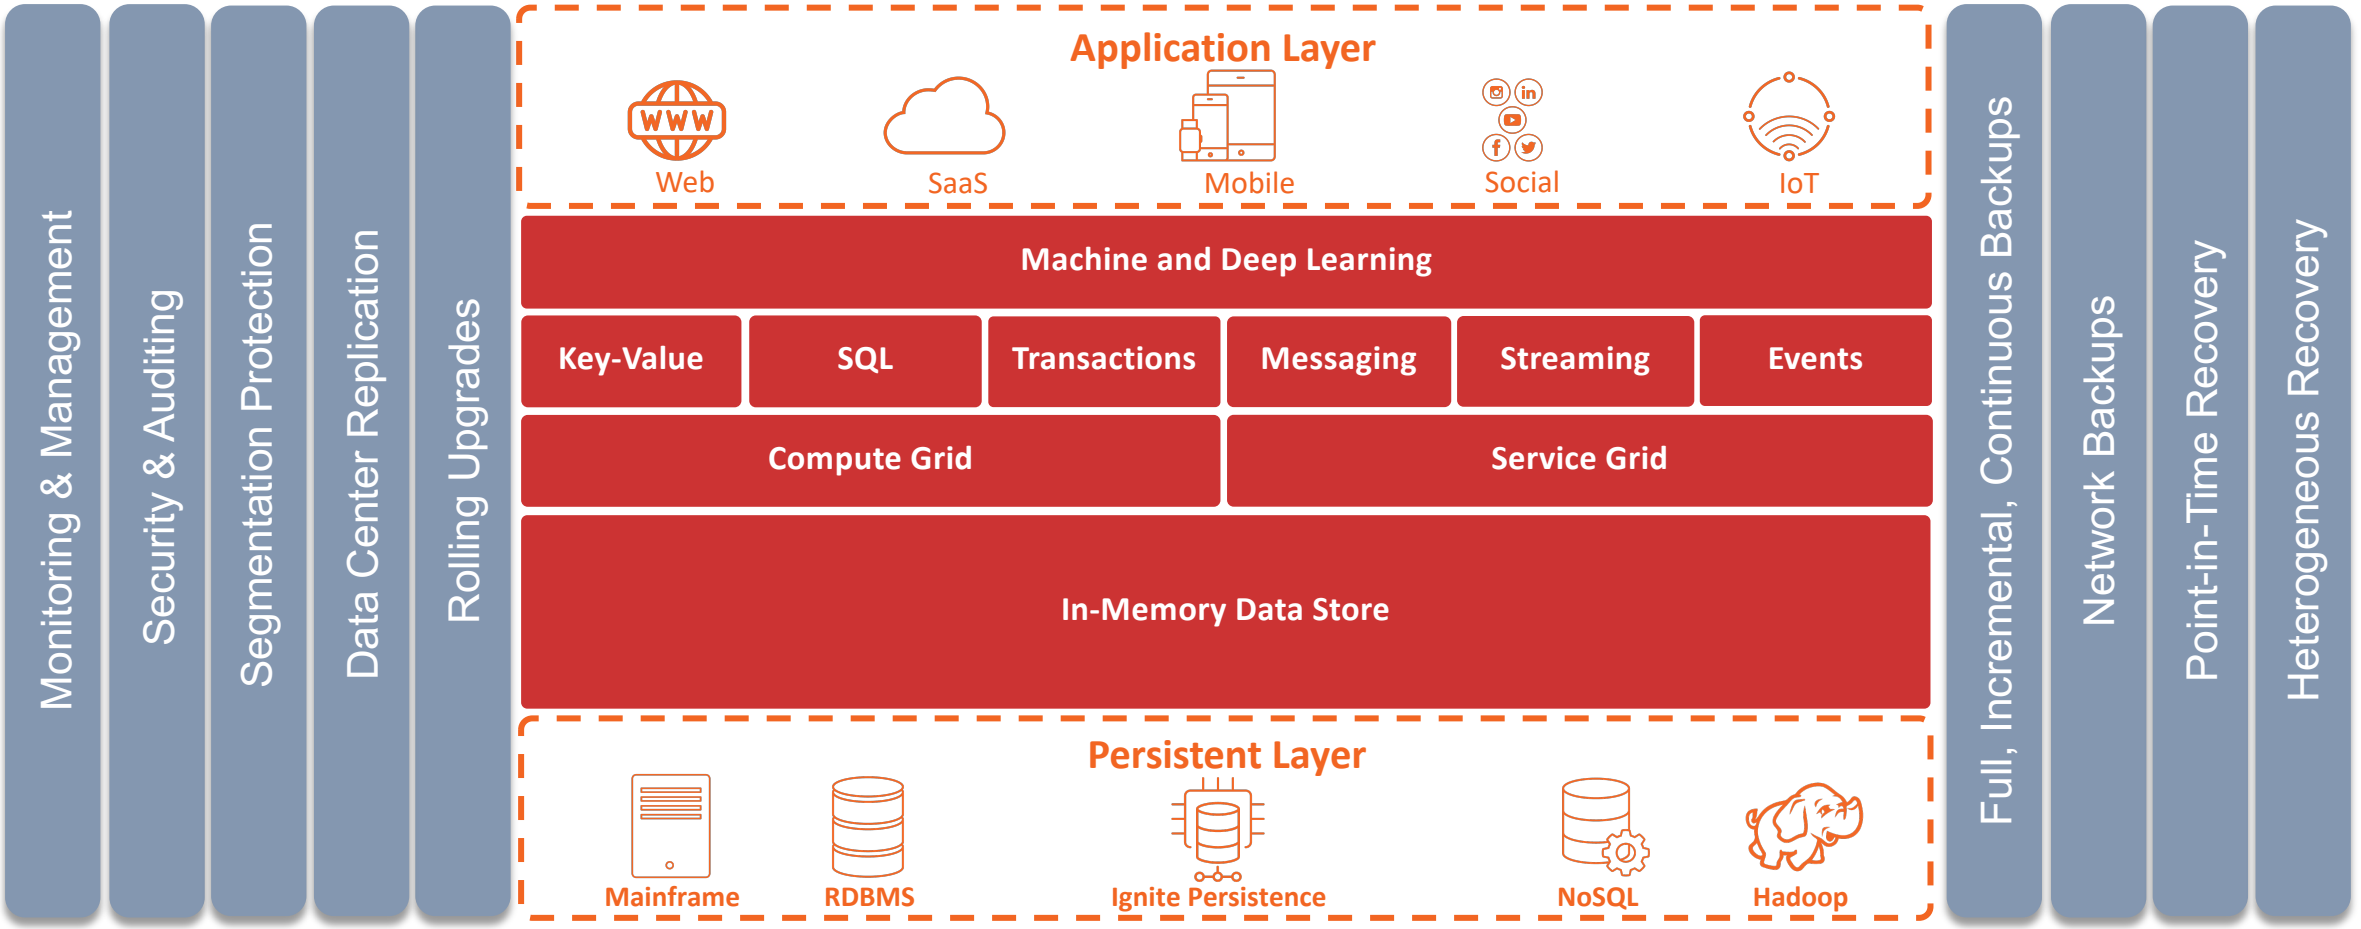

Apache Ignite Features
GridGain Enterprise Features

GridGain

High-availability and fault-tolerance

High-Availability Checklist

- Replication factor (aka. backups)
- Rack-awareness (aka. rack-safety)
- Persistence
- Cluster Snapshots and PITR
- Data Center Replication

Backup Copies

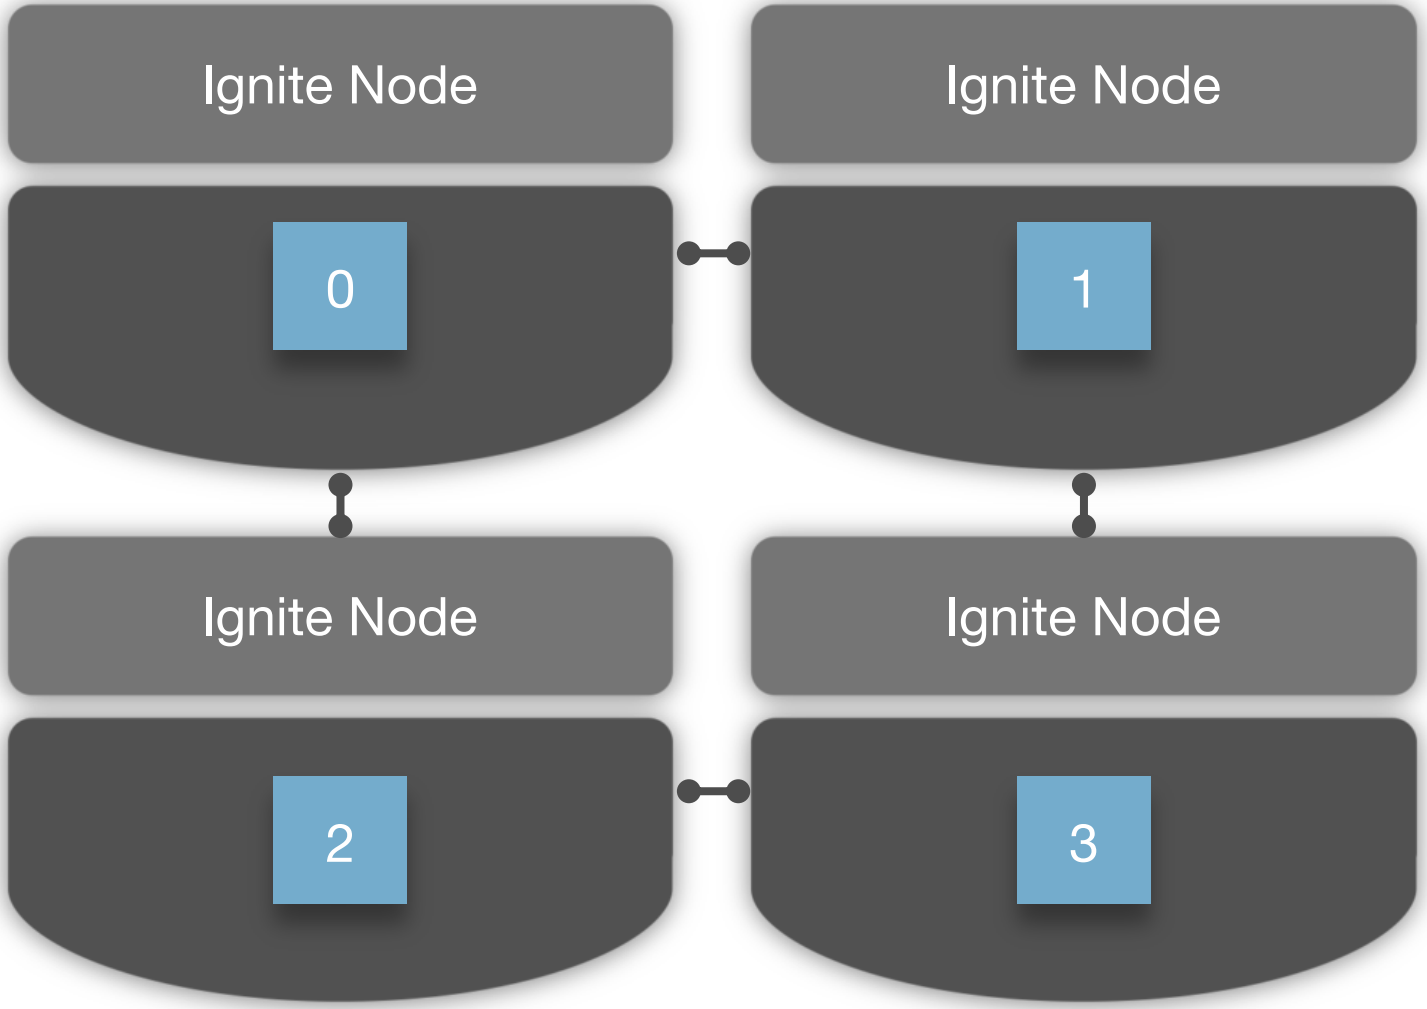

Backup Copies

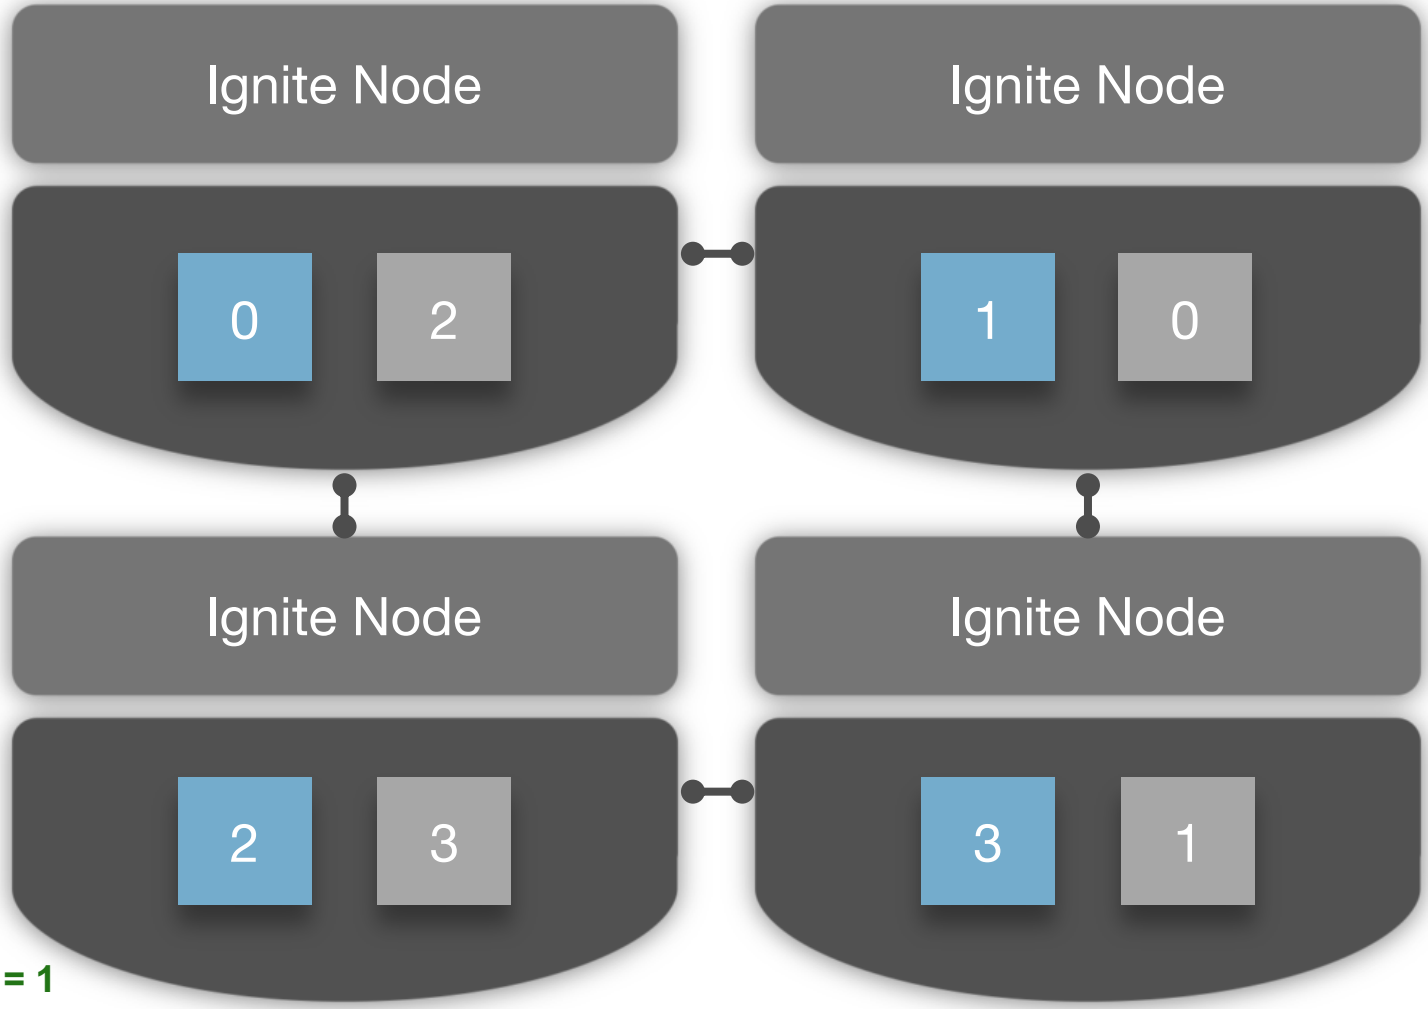

CacheConfiguration.backups = 1

Rack-Awareness: Problem

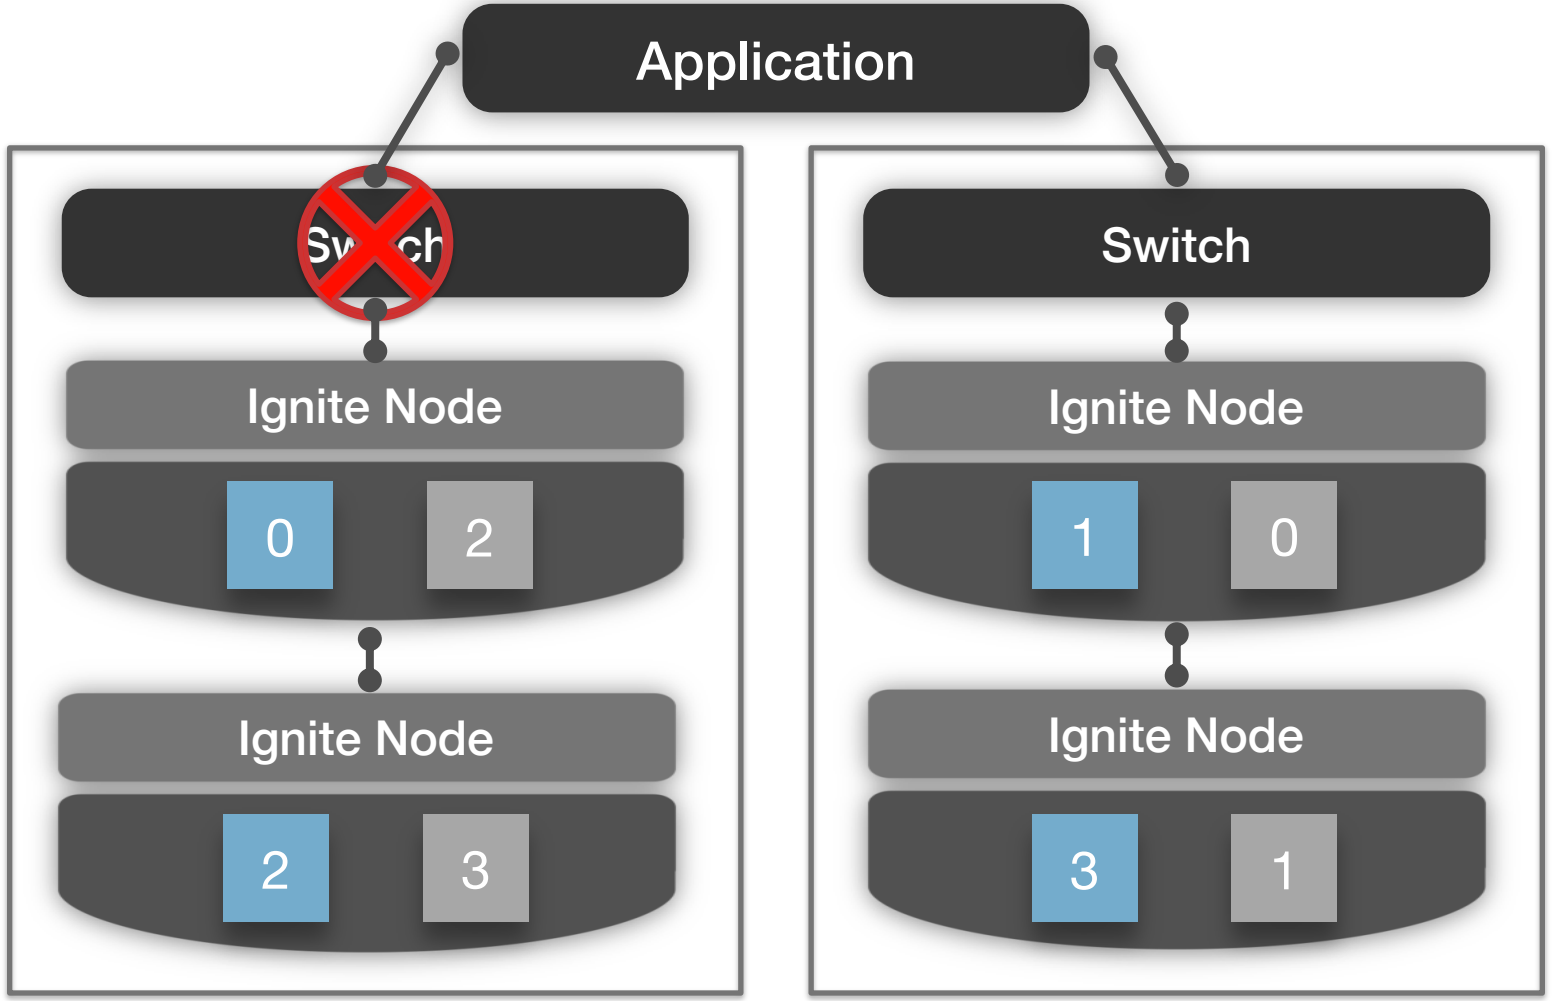

Rack-Awareness: Solution

- Assign attribute to each node:

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

- Ensure affinity function considers the racks:

```
public class RackFilter implements IgniteBiPredicate<ClusterNode, List<ClusterNode>> {
    @Override public boolean apply(ClusterNode candidate, List<ClusterNode> assigned) {
        String candidateRack = candidate.attribute("rack");
        String primaryRack = assigned.get(0).attribute("rack");

        return !Objects.equals(candidateRack, primaryRack);
    }
}
```

Rack-Awareness: Solved

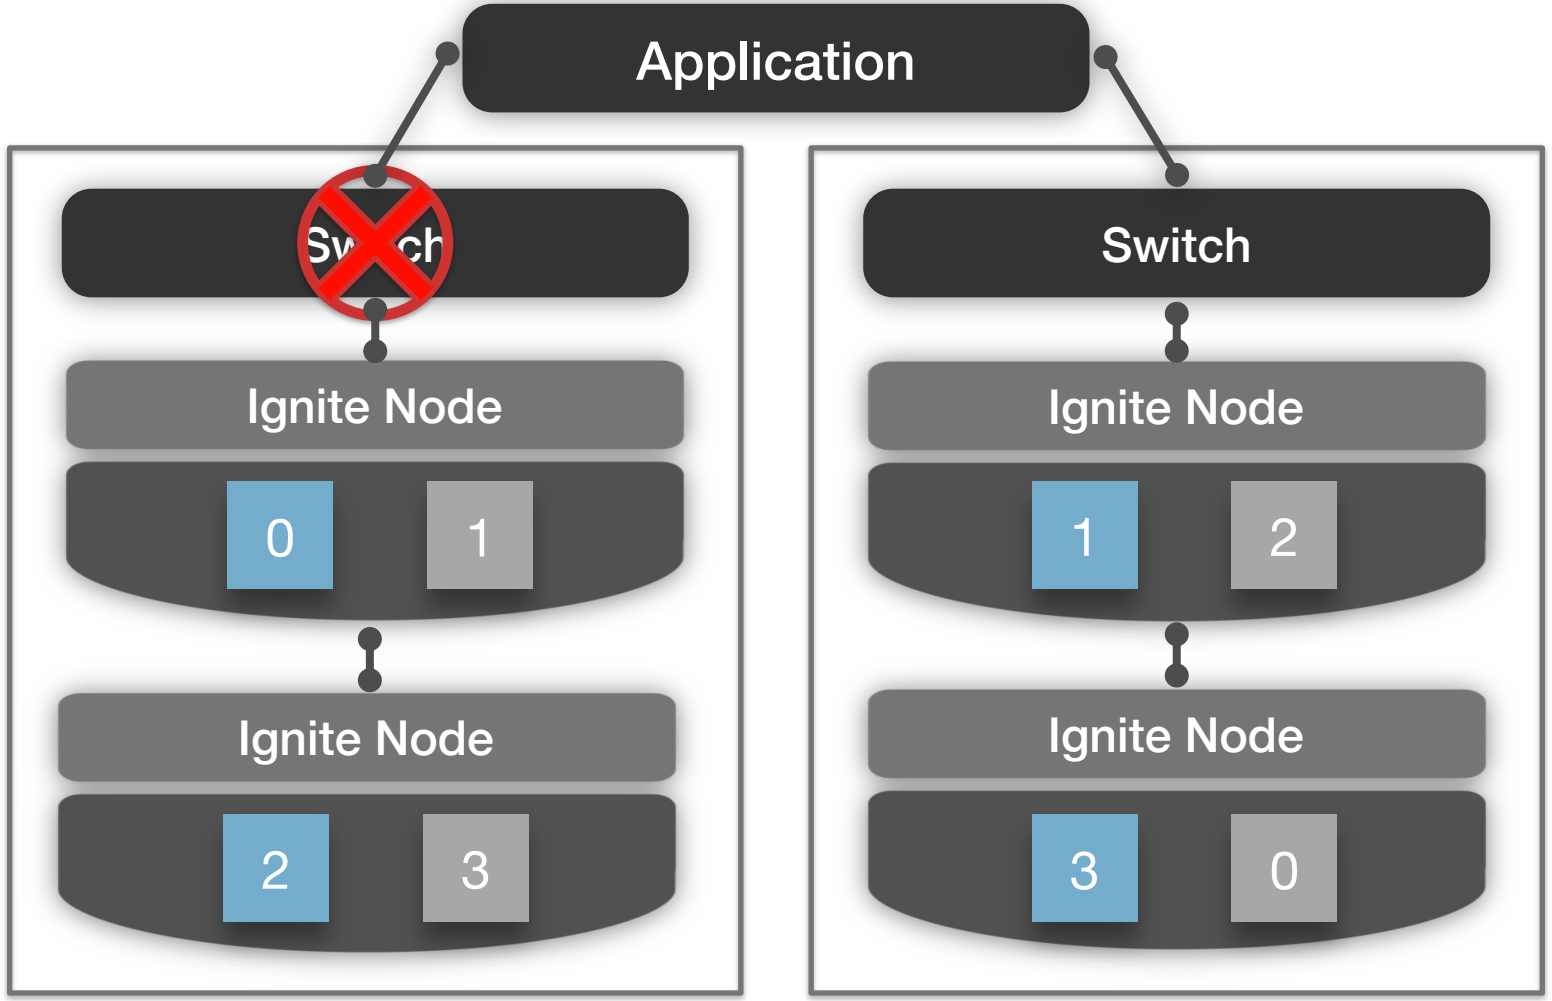

Deployment Modes: Choice of Persistence

Enhance Legacy Architecture - IMDG

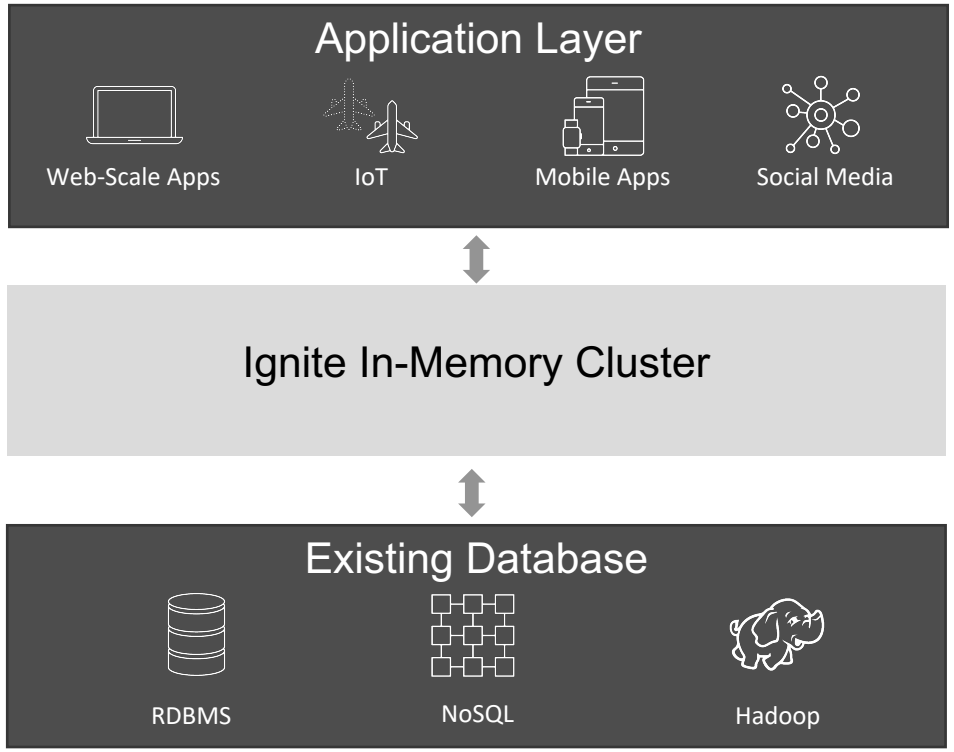

Simplified Modern Architecture - IMDB

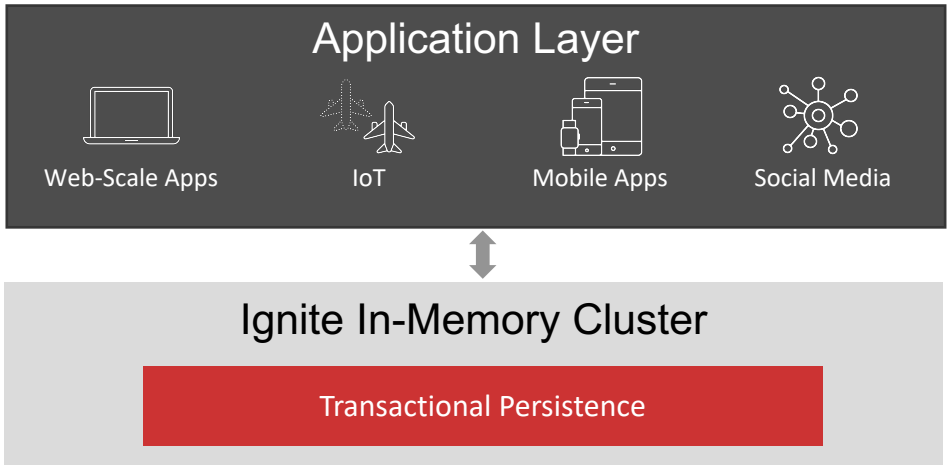

Centralized Backup and Recovery Management

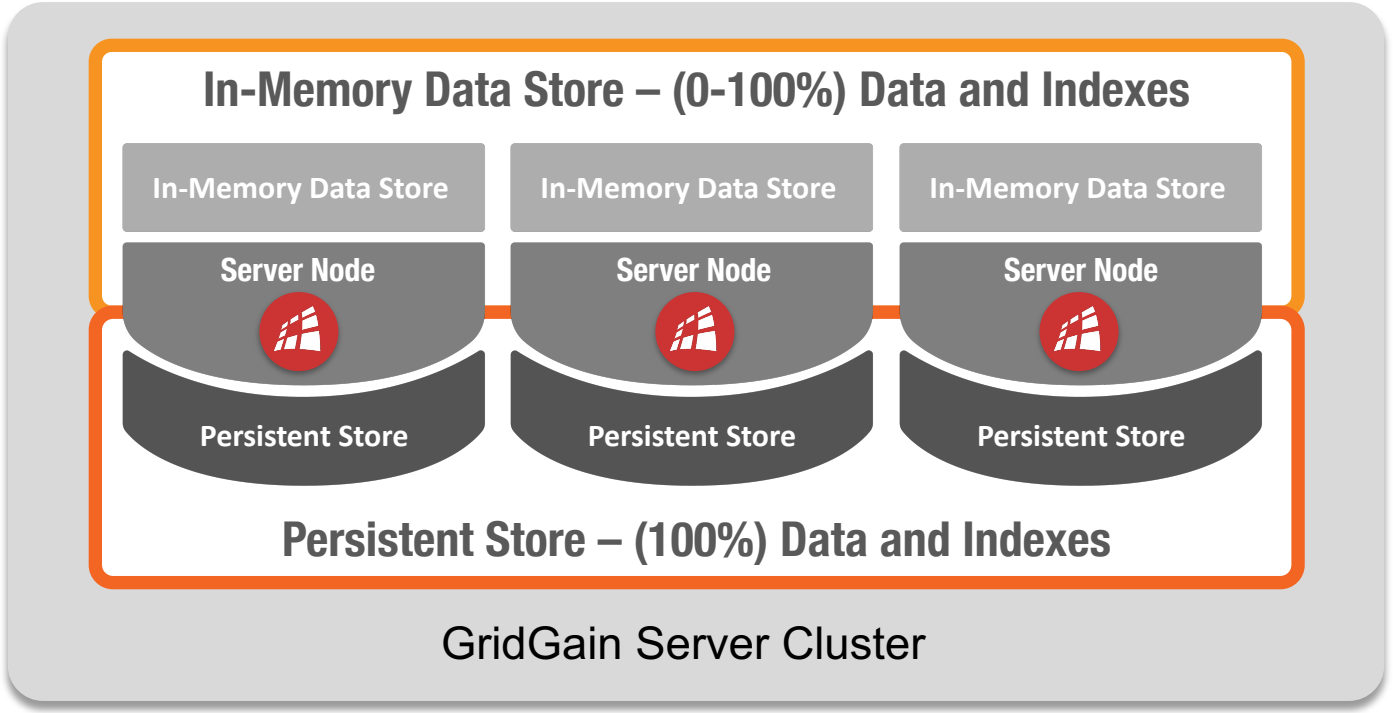

- Full and incremental snapshots
- Continuous archiving (WAL)
- Network backups

- Point-in-time Recovery
- Heterogeneous Recovery

GridGain

Multi-Datacenter Replication

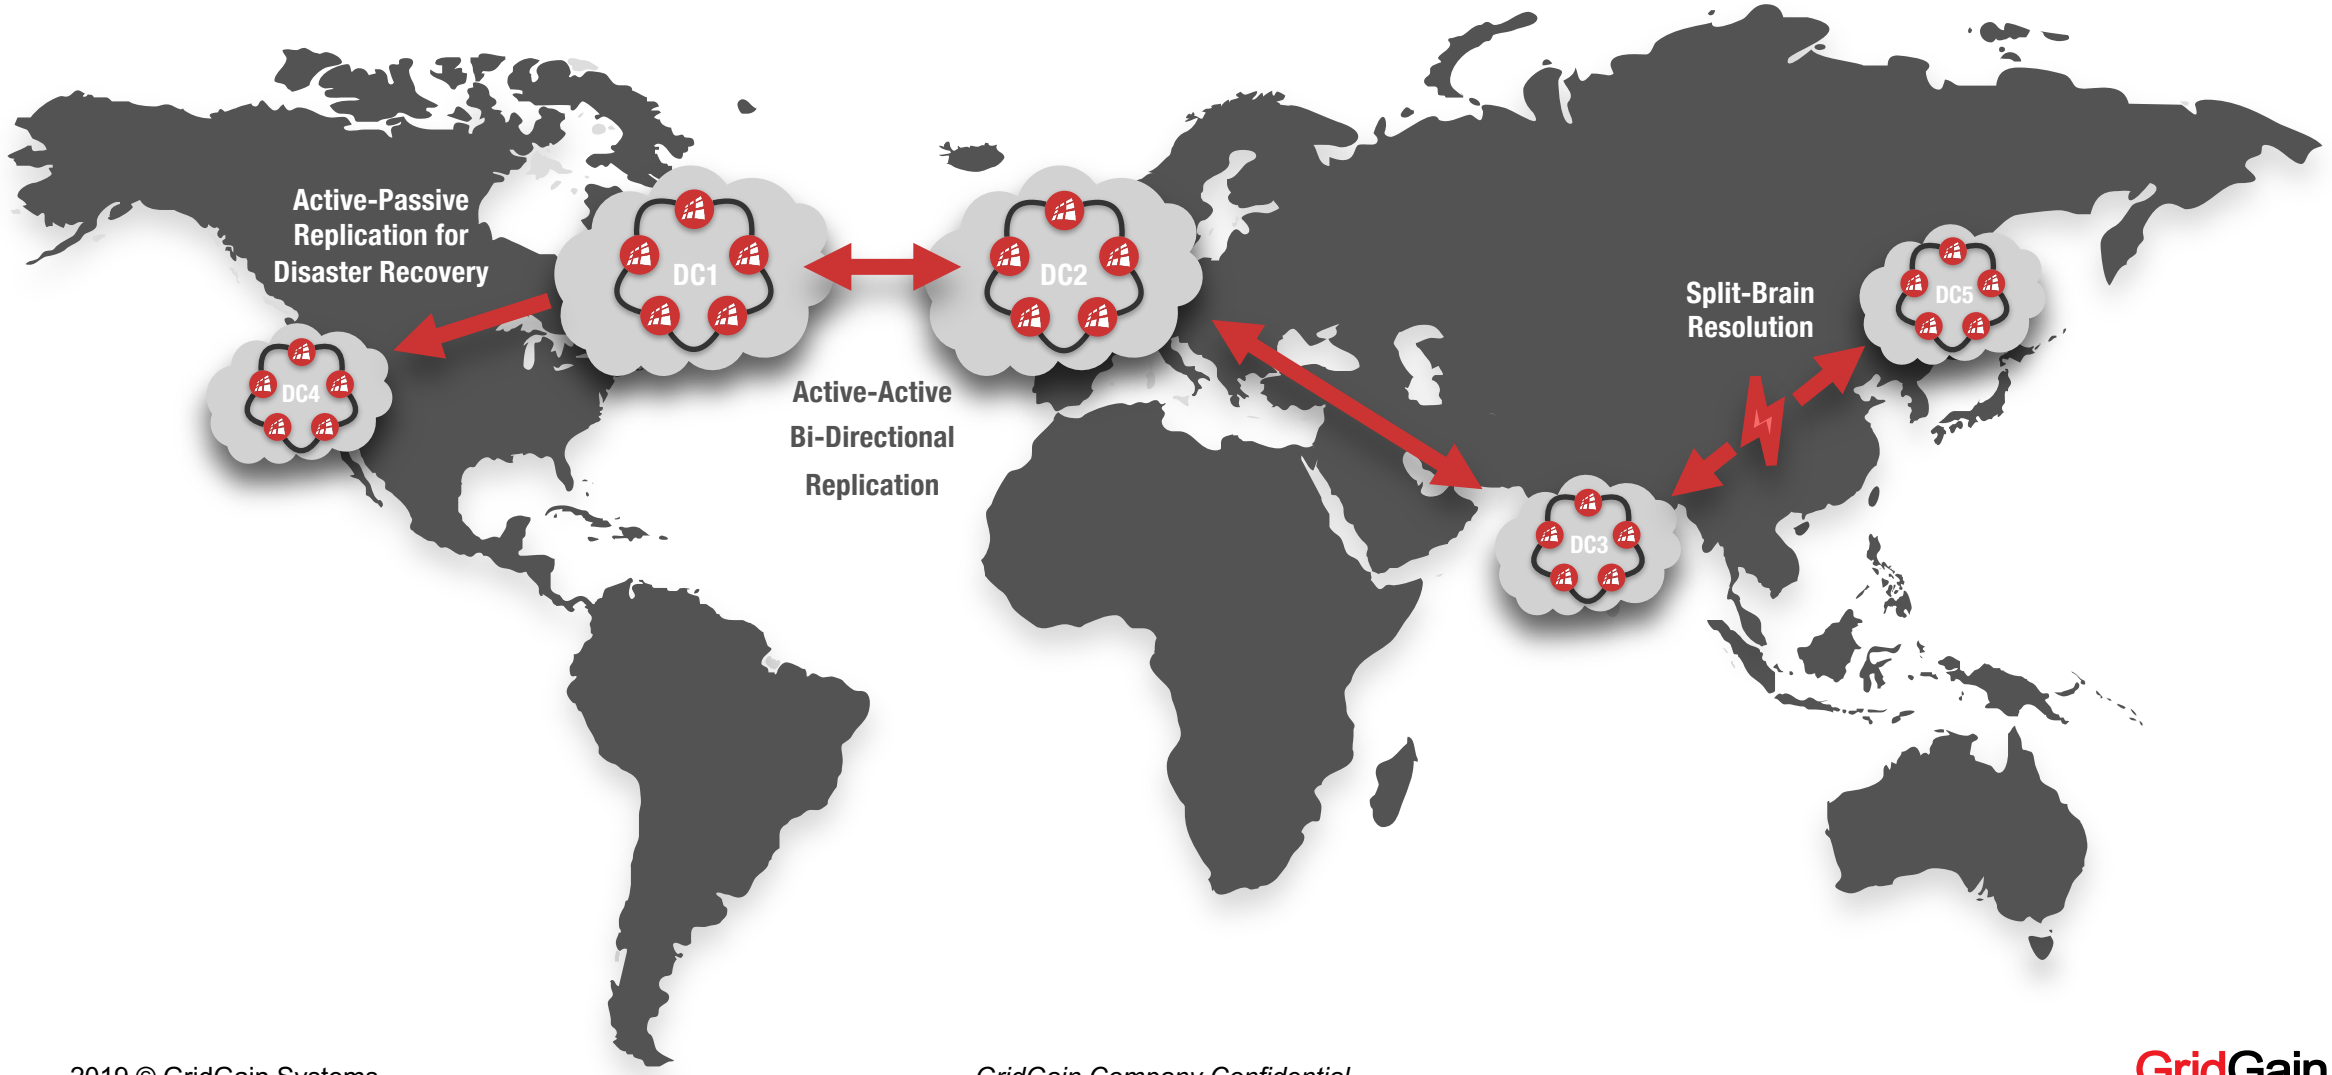

GridGain

Capacity planning and basic tuning

Capacity Planning

- Use Capacity Calculator
  - https://apacheignite.readme.io/docs/capacity-planning
- Size off-heap space:
  - For Data and indexes
  - Load data and adjust estimates
- Size Java heap space:
  - Specific for your app
  - For temporarily objects
  - Do load test
- Consider backups and persistence
  - HA and fault-tolerance

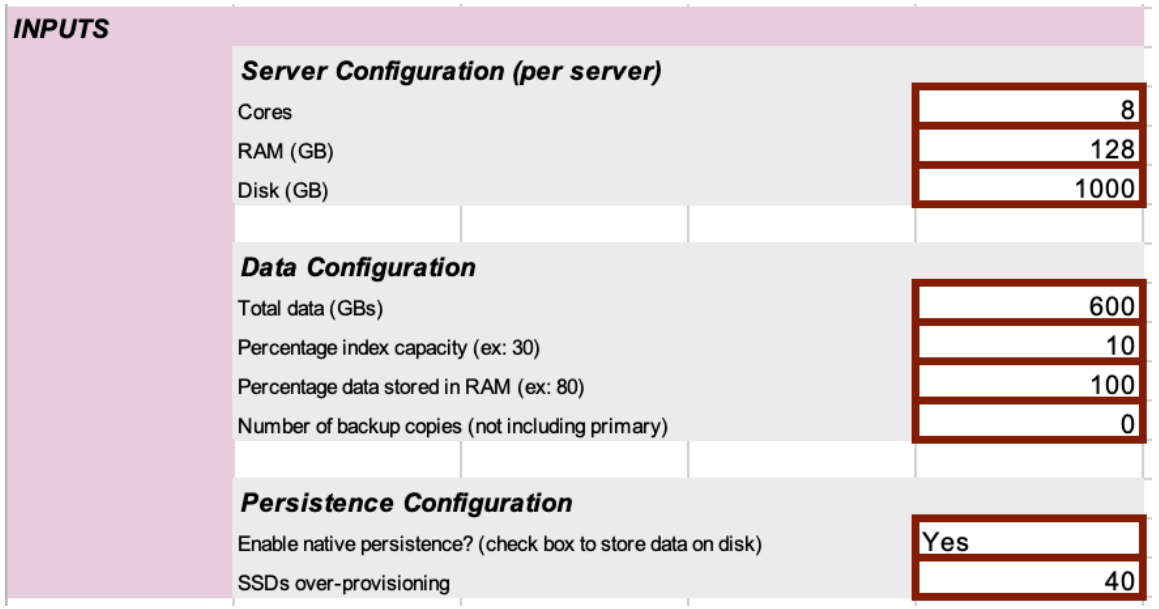

INPUTS		
	Server Configuration (per server)	
	Cores	8
	RAM (GB)	128
	Disk (GB)	1000
	Data Configuration	
	Total data (GBs)	600
	Percentage index capacity (ex: 30)	10
	Percentage data stored in RAM (ex: 80)	100
	Number of backup copies (not including primary)	0
	Persistence Configuration	
	Enable native persistence? (check box to store data on disk)	Yes
	SSDs over-provisioning	40

Getting Ready for Production Workloads

- Run solution-specific benchmarks
  - Disregard public benchmarks
- Run load-tests
- Tune memory and persistence
  - https://apacheignite.readme.io/docs/preparing-for-production
- Come up with final prod configuration

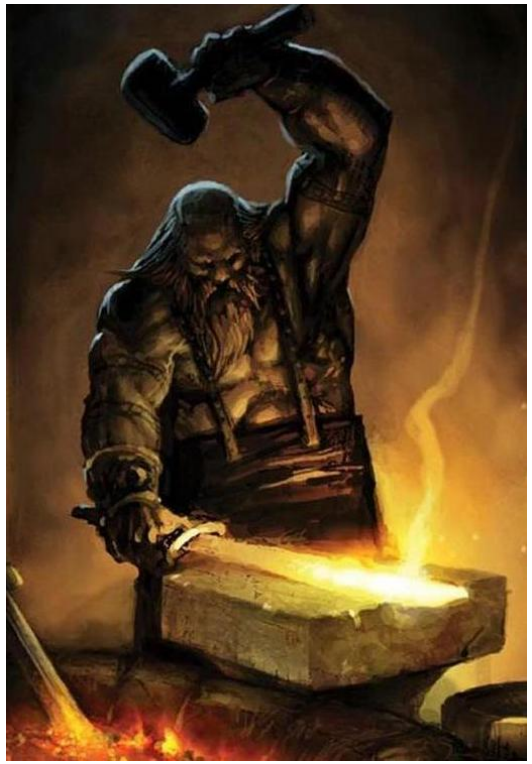

Upgrades and patching

Upgrades and Patching: Open Source Way

- Prepare for full cluster restarts
  - Versions upgrades
  - For Ignite and GridGain Community versions
- Patching and bug fixes:
  - Ignite: best effort
  - GG Community: monthly

Upgrades and Patching: Enterprise Approach

- Take a cluster snapshot (backup)
- Use Rolling Upgrades
  - No downtimes during upgrades!
- For GridGain enterprise customers
  - GridGain Enterprise Edition
  - GridGain Ultimate Edition

Product Support

Apache Ignite Support – Faster Time to Reliable Ignite

- Ignite Community User List
  - Best Effort

- Ignite Support by GridGain experts
  - Unlimited web/e-mail support
  - Identify bugs, workarounds
  - Troubleshoot performance, reliability issues

https://www.gridgain.com/products/services/support/support-apache-ignite

GridGain

Enterprise-grade Support for Business-critical Apps

- 24x7 global support
- Fully tested hot fixes, updates, new releases
- Integrated multichannel issue submission, support and resolution with SLAs
- Access to patches and features before ASF releases
- Influence product direction

Solution	Standard Community	Standard Enterprise	Premium Enterprise	Premium Ultimate
Maintenance Releases	✓	✓	✓	✓
Enterprise Features		✓	✓	✓
Backup and Recovery Management				✓
Support Hours	24x7 Phone or Online	24x7 Phone or Online	24x7 Phone or Online	24x7 Phone or Online
Hot Bug Fixes	✓	✓	✓	✓
Initial Response Time	2 hours	2 hours	1 Hour	1 Hour
Named Technical Contacts	3	3	4	4

Coming next in “Moving Ignite into Production”

- Best Practices for Native Persistence and Data Recovery
- Best Practices for Monitoring Distributed In-Memory Computing
- June – July, 2019

Q&A