

Using Distributed Tracing to Resolve Performance Issues in Apache Ignite

Greg Stachnick, Director of Product Management
GridGain
10/28/2020



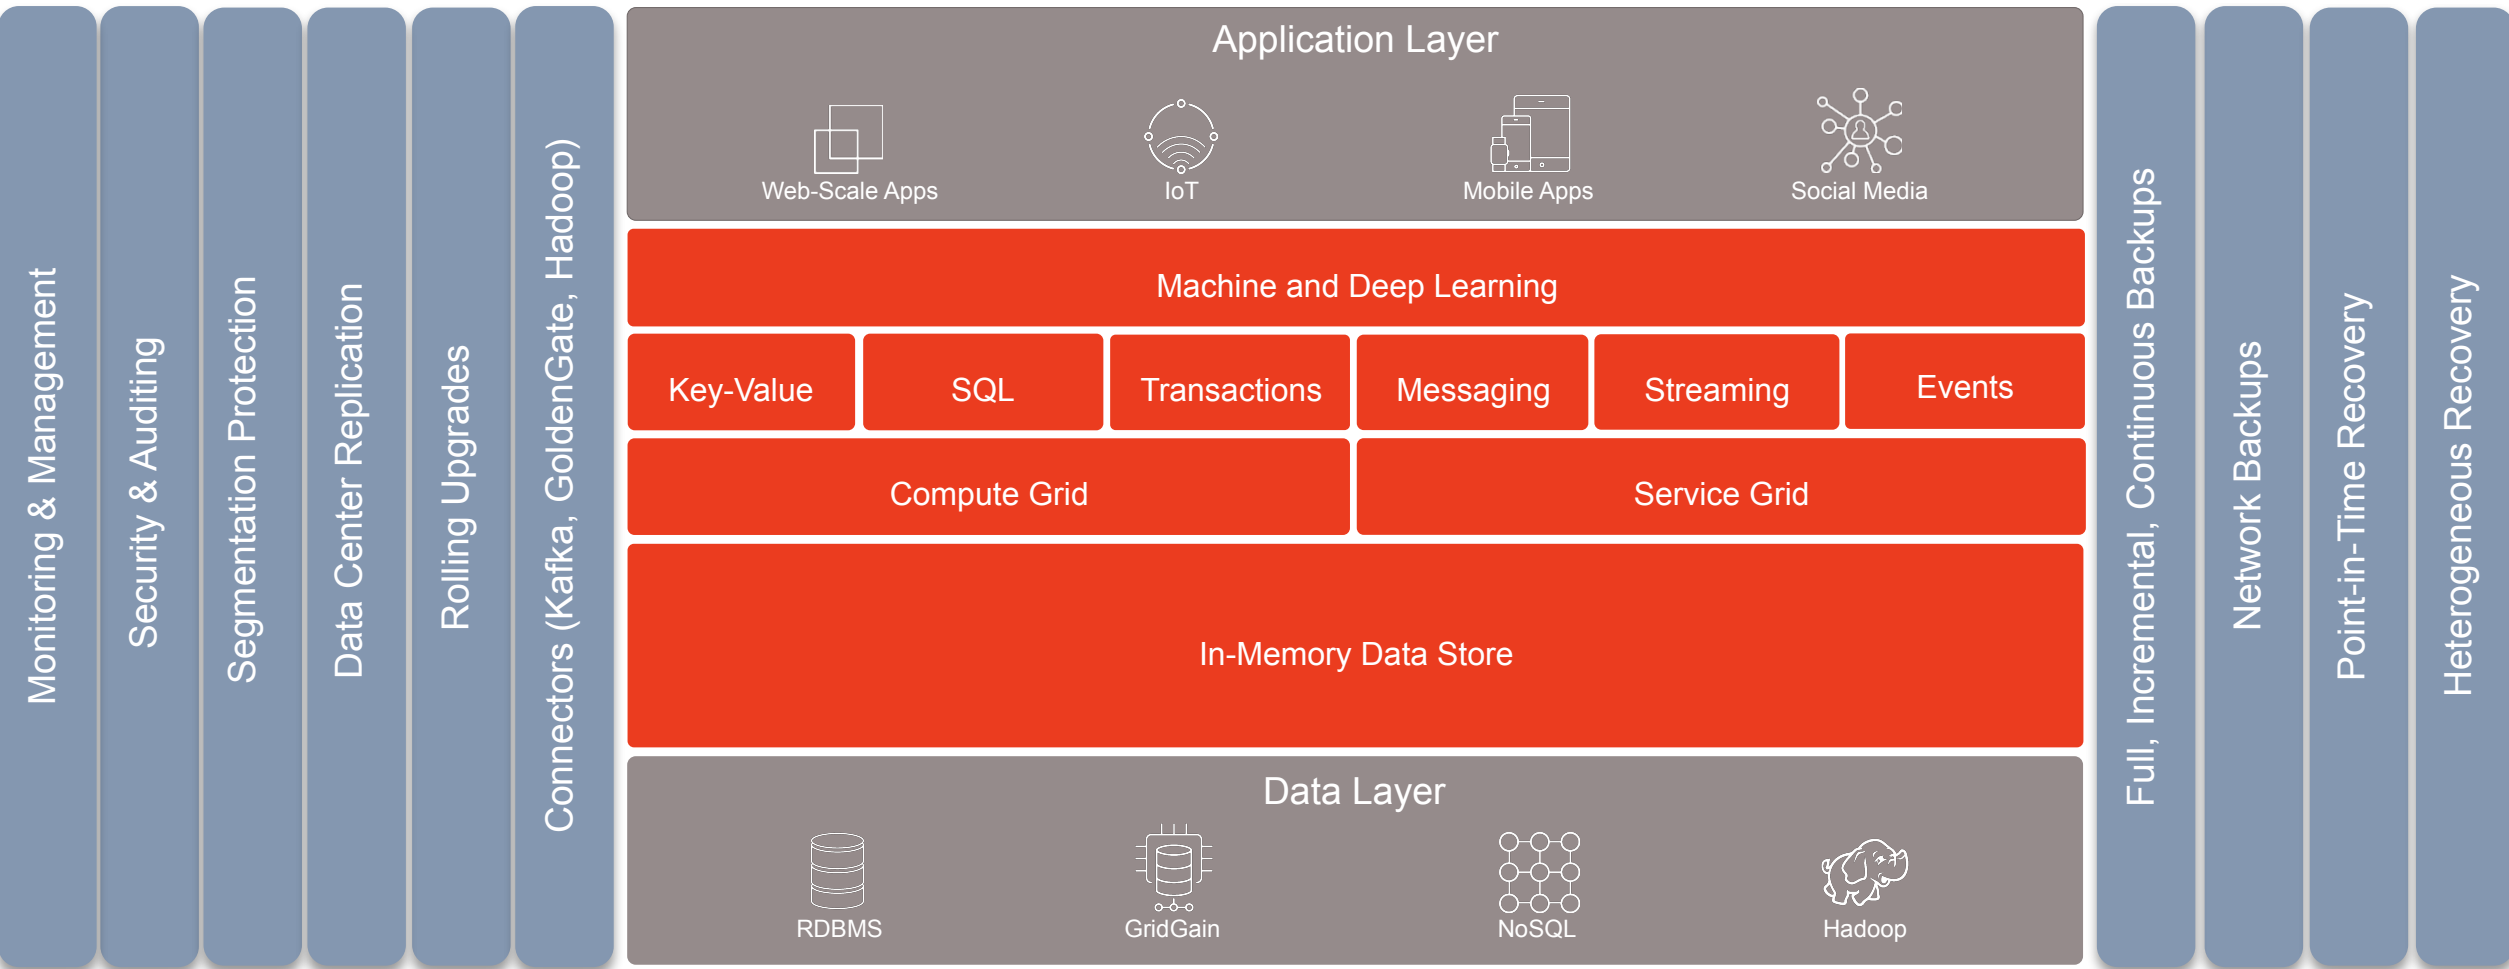
Agenda

- GridGain & Ignite Overview
- Why Tracing?
- Tracing in GridGain & Ignite
- Control Center Overview
- Demo: TX Troubleshooting
- Q&A




Photo by [Taylor Simpson](#)

Ignite & GridGain In-Memory Computing Platform



Why is Tracing Important?

- Monitor execution of running application
- Useful for analyzing and debugging distributed systems
- Accelerates root cause analysis



What is OpenCensus?

- Open source set of libraries to collection application metrics and distributed traces.
- Instrumented for many popular languages



- Java
- Go
- Node.js
- C#
- C++
- Ruby
- Python
- PHP

- Can extend to your custom code
- Supported by many monitor tools through custom Exporters

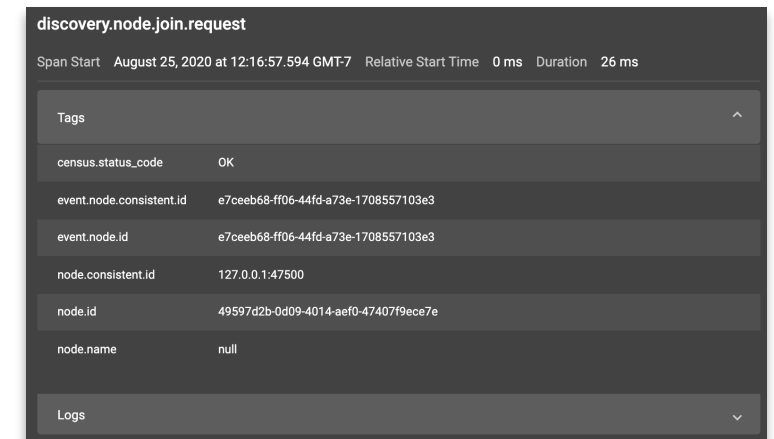
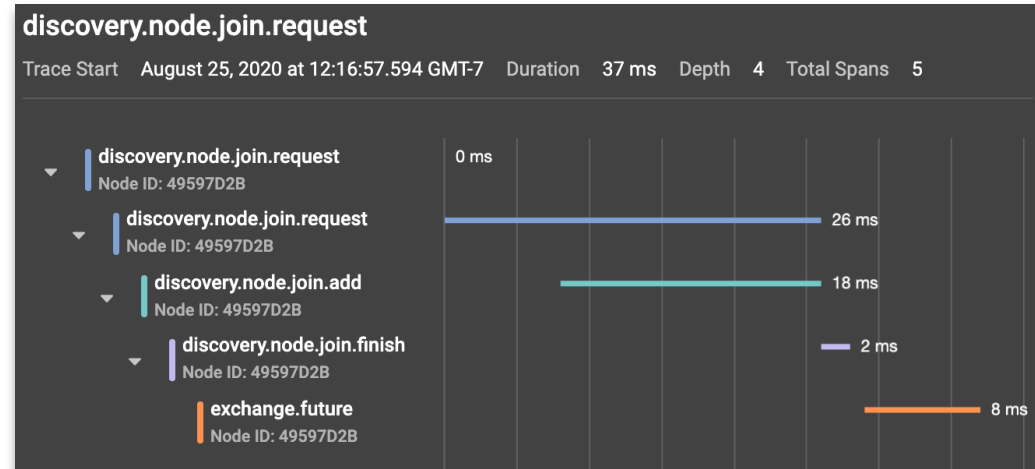
OpenCensus Tracing In GridGain & Ignite

Anatomy of a Trace

- API Name
- Node
- Start Time & Duration
- Status
- API Specific Tags
- Logs

Instrumented APIs

- Transactions
- Communication
- Discovery
- Cache Write
- Exchange
- Cache Read



Enabling Tracing in GridGain

- GridGain 8.7.x or Ignite 2.9
- ignite-opencensus in libs/
- Enable span collection from control.sh --tracing configuration
- scope – specifies API (TX, DISCOVERY, EXCHANGE, COMMUNICATION, CACHE)
- sampling-rate – Between 0 & 1
- Add tracingSpi to config

```
→ bin JVM_OPTS="-DIGNITE_ENABLE_EXPERIMENTAL_COMMAND=true" ./control.sh  
--tracing-configuration set --scope DISCOVERY --sampling-rate 1
```

```
Control utility [ver. 8.7.23#20200728-sha1:4e015639]  
2020 Copyright(C) GridGain Systems, Inc. and Contributors  
User: gstachnick  
Time: 2020-08-25T17:05:04.583  
Command [TRACING-CONFIGURATION] started  
Arguments: --tracing-configuration set --scope DISCOVERY --sampling-rate  
1  
-----  
-----  
Scope, Label, Sampling Rate, included scopes  
DISCOVERY,,1.0,[]  
Command [TRACING-CONFIGURATION] finished with code: 0  
Control utility has completed execution at: 2020-08-25T17:05:04.821  
Execution time: 238 ms
```

```
<bean class="org.apache.ignite.spi.tracing.  
opencensus.OpenCensusTracingSpi"/>
```

GridGain Control Center Overview

- Comprehensive Cluster Management and Developer Tool for GridGain and Apache Ignite
- Supported Products
- Built for GridGain 8.7+ and Apache Ignite 2.8+
- Apache Ignite requires Control Center Agent install



Modern Interface with Drag & Drop Flexibility

The screenshot displays the GridGain Control Center interface. The top navigation bar includes the GridGain logo, a menu icon, and the current cluster name 'trusting_feistel'. The main dashboard area is titled 'Default Dashboard' and shows a 'Last 24 hour' time range. A sidebar on the left contains navigation icons for Dashboard, Alerting, Tracing, SQL, Baseline, Snapshots, DR, and Clusters. The dashboard is composed of several widgets:

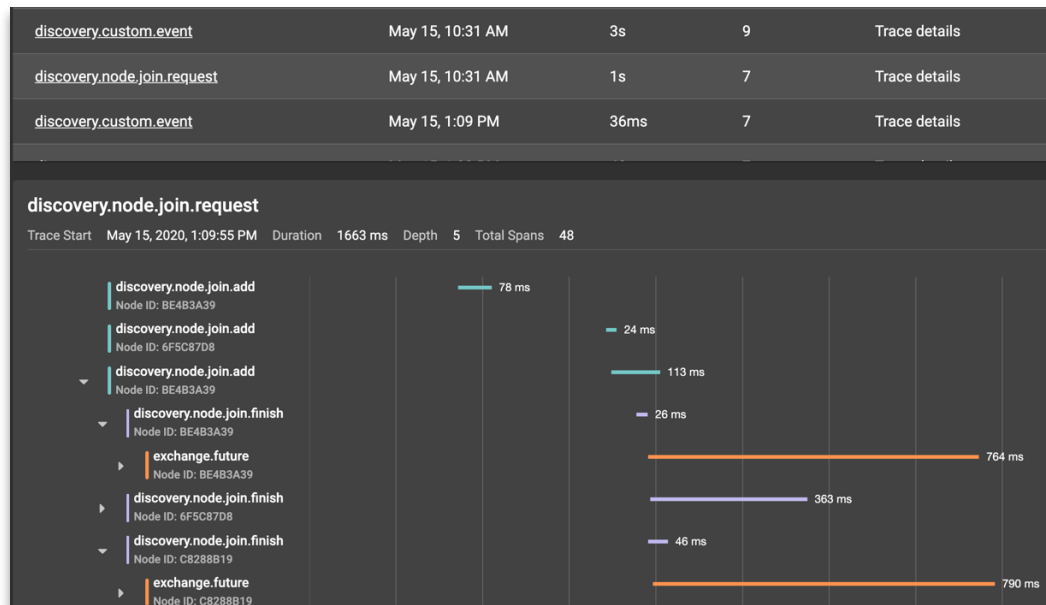
- Nodes:** A table listing three nodes with their consistency IDs, online status, types, versions, and IP addresses.
- Thread Time:** A line chart showing the current thread CPU time for five nodes over a period from 10:41 AM to 11:11 AM.
- Heap Usage:** A visual representation of heap usage for three nodes, with a color scale from 0% (orange) to 100% (yellow).
- Memory:** A table showing the heap used memory for three nodes.
- Heap Size:** A line chart showing the heap size for five nodes over a period from 10:41 AM to 11:21 AM.

A 'Rebalance' dialog box is overlaid on the right side of the dashboard, showing a progress indicator at 100% and a color scale for the rebalancing process.

Order	Consisten...	Online	Type	Version	IP
1	b9a42525...	✓	Coor...	8.7.14	192.168.8...
2	36ae0c82...	✓	Server	8.7.14	192.168.8...
3	b992d4a4...	✓	Server	8.7.14	192.168.8...

Consisten...	Heap Used
E88F3FDA	62.96MB
B9A42525	69.66MB
B992D4A4	49.69MB

Active Tracing & Root Cause Analysis



Easily Track API Call Execution

- Identify all response times for API calls
- Single view of stack traces and logs from all nodes
- Follow execution across the cluster
- Combine with application traces
- Open Census compliant

Setup:

- Copy *ignite-opencensus* from */libs/optional* to */libs/*
- Add Tracing to Cluster Configuration

```
<bean class="org.apache.ignite.spi.tracing.  
opencensus.OpenCensusTracingSpi"/>
```

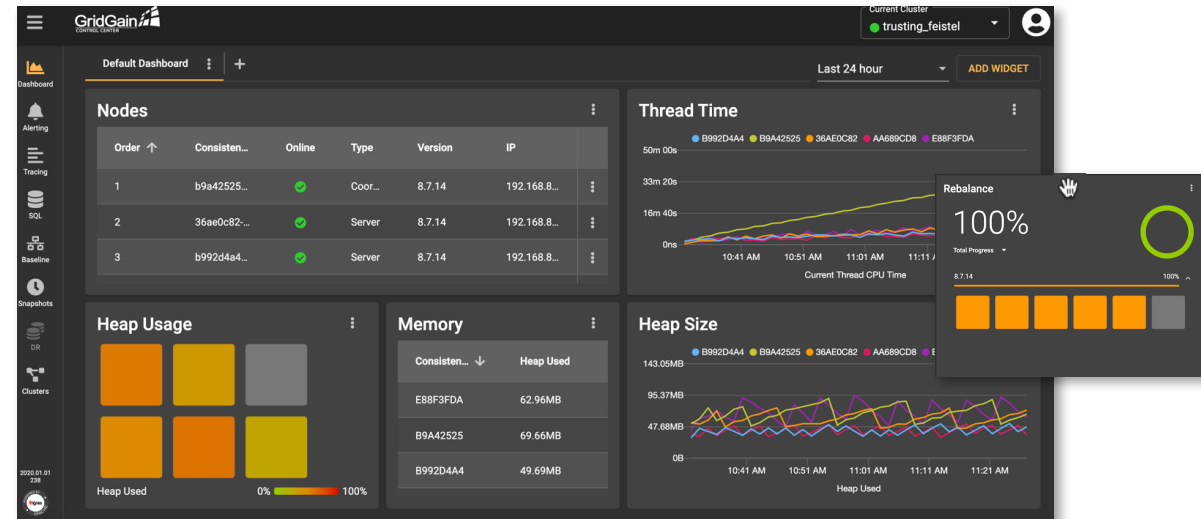
Customizable, Intuitive Dashboard

Easy to Use with Visual Cues

- Visualize cluster status and manage behavior

Highly Flexible

- Drag and drop panes to view metrics of interest
- Tracks over 200 cluster metrics
 - Open Census compliant



User-Defined Production Alerts

CONFIGURATIONS		EVENTS		NOTIFICATION CHANNELS	
Tag	Type	Condition	Count	Status	
Node Count	Cluster metric	Server Nodes Count < 10 for 0s	0	OK	
Memory	Node metric	Heap Used > 95.37MB for 0s	0	OK	
CPU	Node metric	CPU Load > 2% for 1s	0	OK	

CONFIGURATIONS		EVENTS		NOTIFICATION CHANNELS	
Created On ↓	State	Message		Alert	
5/15/20, 10:41 AM	Alerting	Node "Heap Used" metric was > 10000000		Memory	
5/15/20, 10:41 AM	Alerting	Node "Heap Used" metric was > 10000000		Memory	
5/15/20, 10:41 AM	Alerting	Node "Heap Used" metric was > 10000000		Memory	
5/15/20, 10:41 AM	Alerting	Node "Heap Used" metric was > 10000000		Memory	
5/15/20, 10:41 AM	Alerting	Node "Heap Used" metric was > 10000000		Memory	
5/15/20, 10:41 AM	Alerting	Cluster "Server Nodes Count" metric was < 10		Node Count	

Quickly Identify and Resolve Issues with Configurable Alerts

- Create custom active alerts on any metric
- Monitor cluster, node, and cache events
- Configure flexible notifications
 - Email and SMS

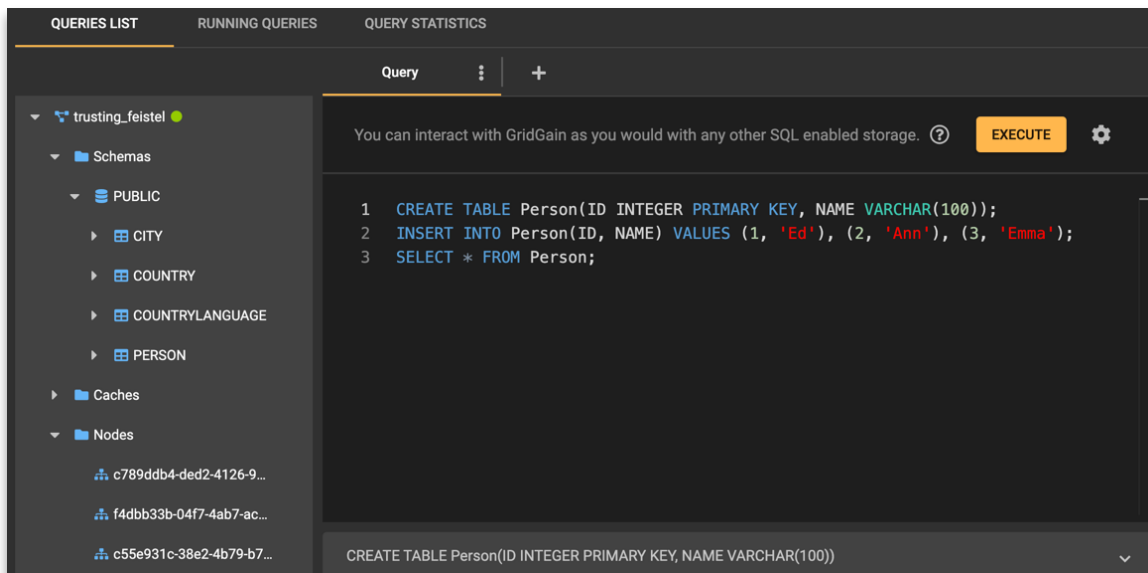
Query Development Tools

Accelerate Query Development

- Fine tune for optimal performance

Develop and Run SQL Scripts

- Introspect SQL execution
- Analyze performance of long running queries and query history
- Display all clusters, caches, and schemas in real-time



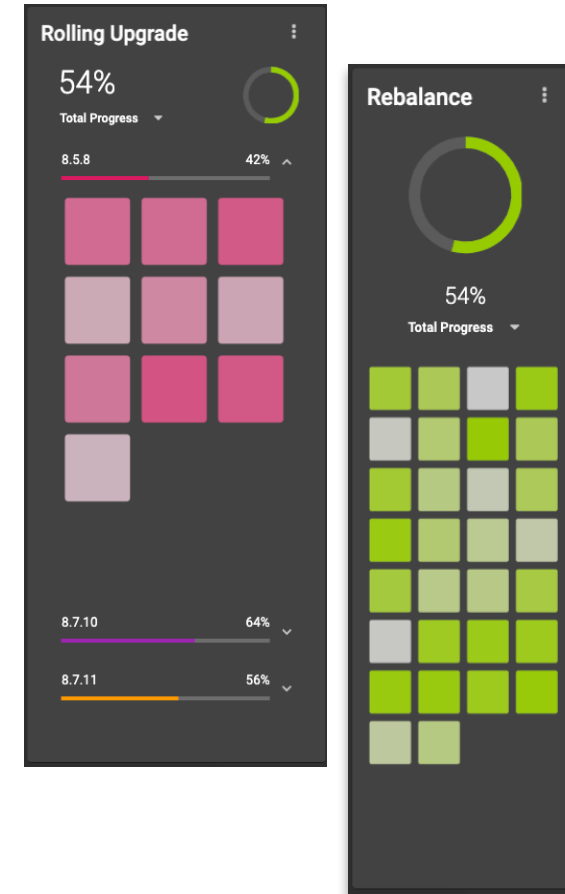
Monitor Rolling Upgrades & Cluster Rebalancing

Monitor Rolling Upgrades

- View cluster upgrade progress
- GridGain Enterprise and Ultimate Editions only

Monitor Data Rebalancing

- Visualize data resilience processes
- Identify and isolate problematic nodes



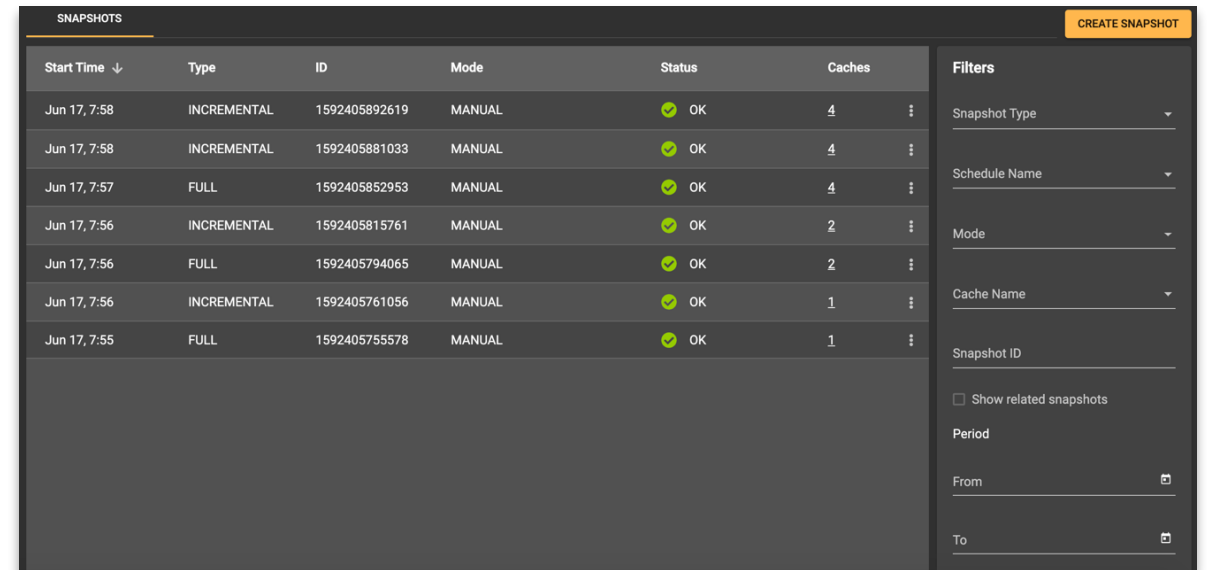
Disaster Recovery & Backup Management

Easy Configuration of Backups

- Fully managed backups and comprehensive data recovery tools
- GridGain Ultimate Edition only

Comprehensive Backup Monitoring and Management

- Create full and incremental backups
- Validate backup integrity
- Recover state from a specific point in time
- Automate backup creation lifecycle



The screenshot displays a 'SNAPSHOTS' management interface. At the top right, there is a 'CREATE SNAPSHOT' button. The main area contains a table with the following columns: Start Time (with a downward arrow), Type, ID, Mode, Status, and Caches. The table lists seven snapshots, all with a status of 'OK' and a green checkmark icon. To the right of the table is a 'Filters' sidebar with several dropdown menus: Snapshot Type, Schedule Name, Mode, Cache Name, and Snapshot ID. Below these filters, there is a checkbox for 'Show related snapshots' and a 'Period' section with 'From' and 'To' date pickers.

Start Time ↓	Type	ID	Mode	Status	Caches
Jun 17, 7:58	INCREMENTAL	1592405892619	MANUAL	✓ OK	4
Jun 17, 7:58	INCREMENTAL	1592405881033	MANUAL	✓ OK	4
Jun 17, 7:57	FULL	1592405852953	MANUAL	✓ OK	4
Jun 17, 7:56	INCREMENTAL	1592405815761	MANUAL	✓ OK	2
Jun 17, 7:56	FULL	1592405794065	MANUAL	✓ OK	2
Jun 17, 7:56	INCREMENTAL	1592405761056	MANUAL	✓ OK	1
Jun 17, 7:55	FULL	1592405755578	MANUAL	✓ OK	1

Getting Started with Ignite & Gridgain

Apache Ignite

- Download Agent
 - Maven
 - Gridgain.com
- Extract To /libs
- Start Cluster
- Add Cluster ID to Control Center



GridGain

- Copy Agent from /libs/
optional to /libs
- Start Cluster
- Add Cluster ID to
Control Center



Demo

- Local GridGain Community Edition Cluster
- Client performs simple transaction
- Local Control Center



Photo by [Josh Redd](#)

Control Center Options



Online Version

Limited functionality

Available for free

<https://control.gridgain.com>



Developer Version

Full functionality up to 2 nodes and 36 cores

Available for free download

Installable on-premises



Commercial Version

Full functionality for any size cluster

Installable on-premises

Price based on nodes in your cluster(s)

Q&A



Photo by [Emily Morier](#)

GridGain Resources

GridGain Control Center

- Visit <https://control.gridgain.com>

Webinars

- Visit <https://www.gridgain.com/resources/webinars>

Developer Portal

- Visit <https://www.gridgain.com/developer>

Documentation

- Visit <https://www.gridgain.com/docs/control-center/latest/overview>

Free 30-Day Ultimate, Enterprise or Professional Edition Trial

- Visit <https://www.gridgain.com/resources/download>

Thank you

Greg Stachnick, Director of Product Management

greg.stachnick@gridgain.com

@gstachni