



Using Distributed Tracing to Resolve Performance Issues in Apache Ignite

Greg Stachnick

Director of Product Management

August 26, 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



Monitoring & Management

Security & Auditing

Segmentation Protection

Data Center Replication

Rolling Upgrades

Connectors (Kafka, GoldenGate, Hadoop)

Application Layer



Web-Scale Apps



IoT



Mobile Apps



Social Media

Machine and Deep Learning

Key-Value

SQL

Transactions

Messaging

Streaming

Events

Compute Grid

Service Grid

In-Memory Data Store

Data Layer



RDBMS



GridGain



NoSQL



Hadoop

Full, Incremental, Continuous Backups

Network Backups

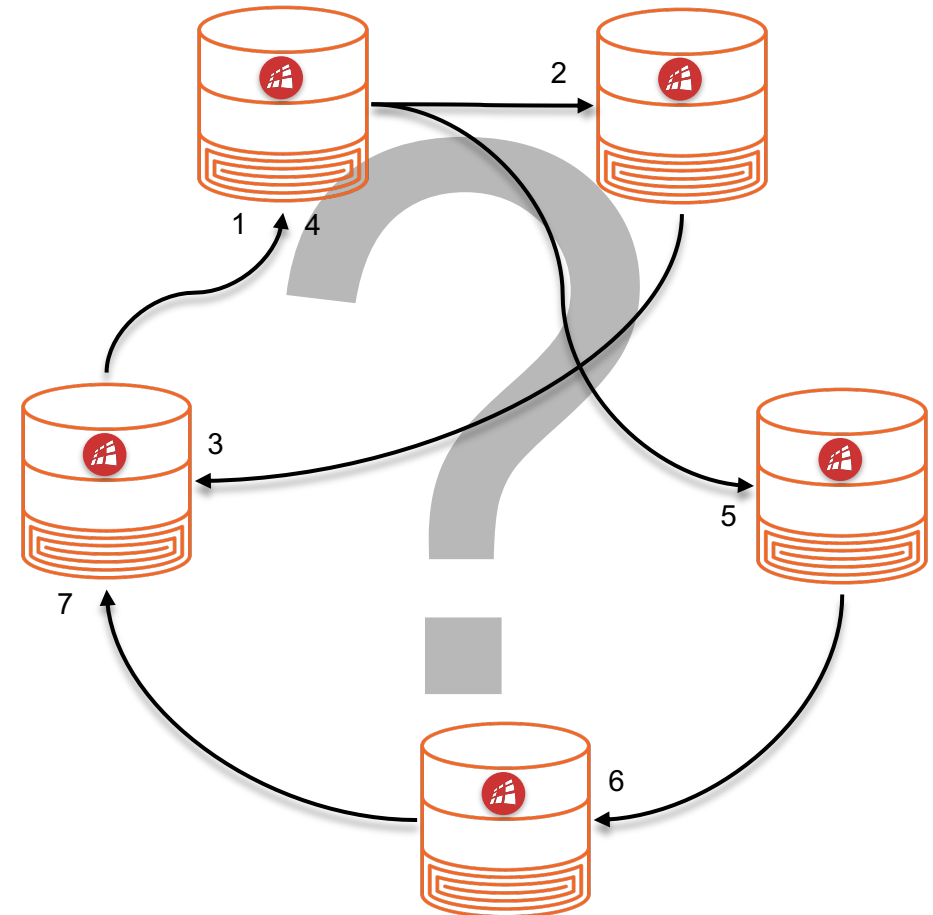
Point-in-Time Recovery

Heterogeneous Recovery

Why is Tracing Important?



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



What is OpenCensus?

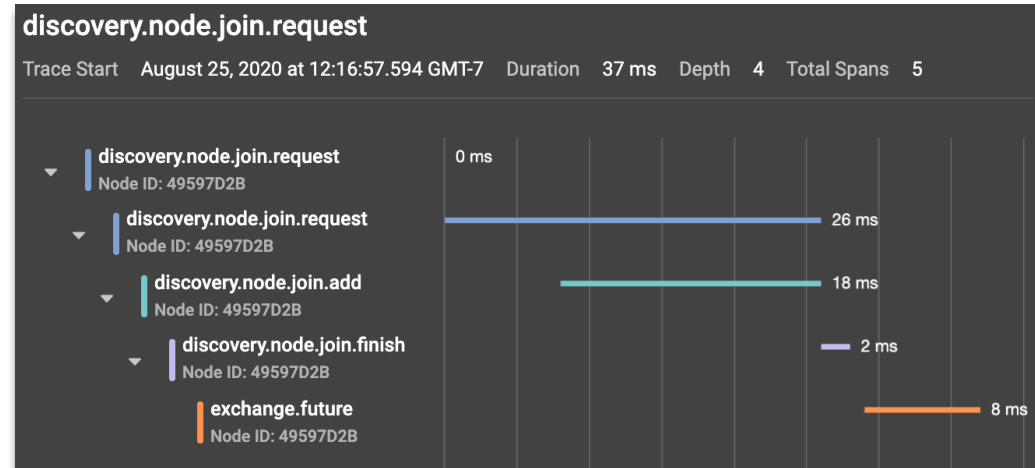


- Open source set of libraries to collection application metrics and distributed traces.
- Instrumented for many popular languages
 - Java
 - Go
 - Node.js
 - C#
 - C++
 - Ruby
 - Python
 - PHP
- Can extend to your custom code
- Supported by many monitor tools through custom Exporters

OpenCensus Tracing In GridGain & Ignite



- Anatomy of a Trace
 - API Name
 - Node
 - Start Time & Duration
 - Status
 - API Specific Tags
 - Logs
- Instrumented APIs
 - Transactions
 - Discovery
 - Exchange
 - Communication



discovery.node.join.request

Span Start August 25, 2020 at 12:16:57.594 GMT-7 Relative Start Time 0 ms Duration 26 ms

Tags	
census.status_code	OK
event.node.consistent.id	e7ceeb68-ff06-44fd-a73e-1708557103e3
event.node.id	e7ceeb68-ff06-44fd-a73e-1708557103e3
node.consistent.id	127.0.0.1:47500
node.id	49597d2b-0d09-4014-ae0-47407f9ece7e
node.name	null

Logs

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)
 - 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
```

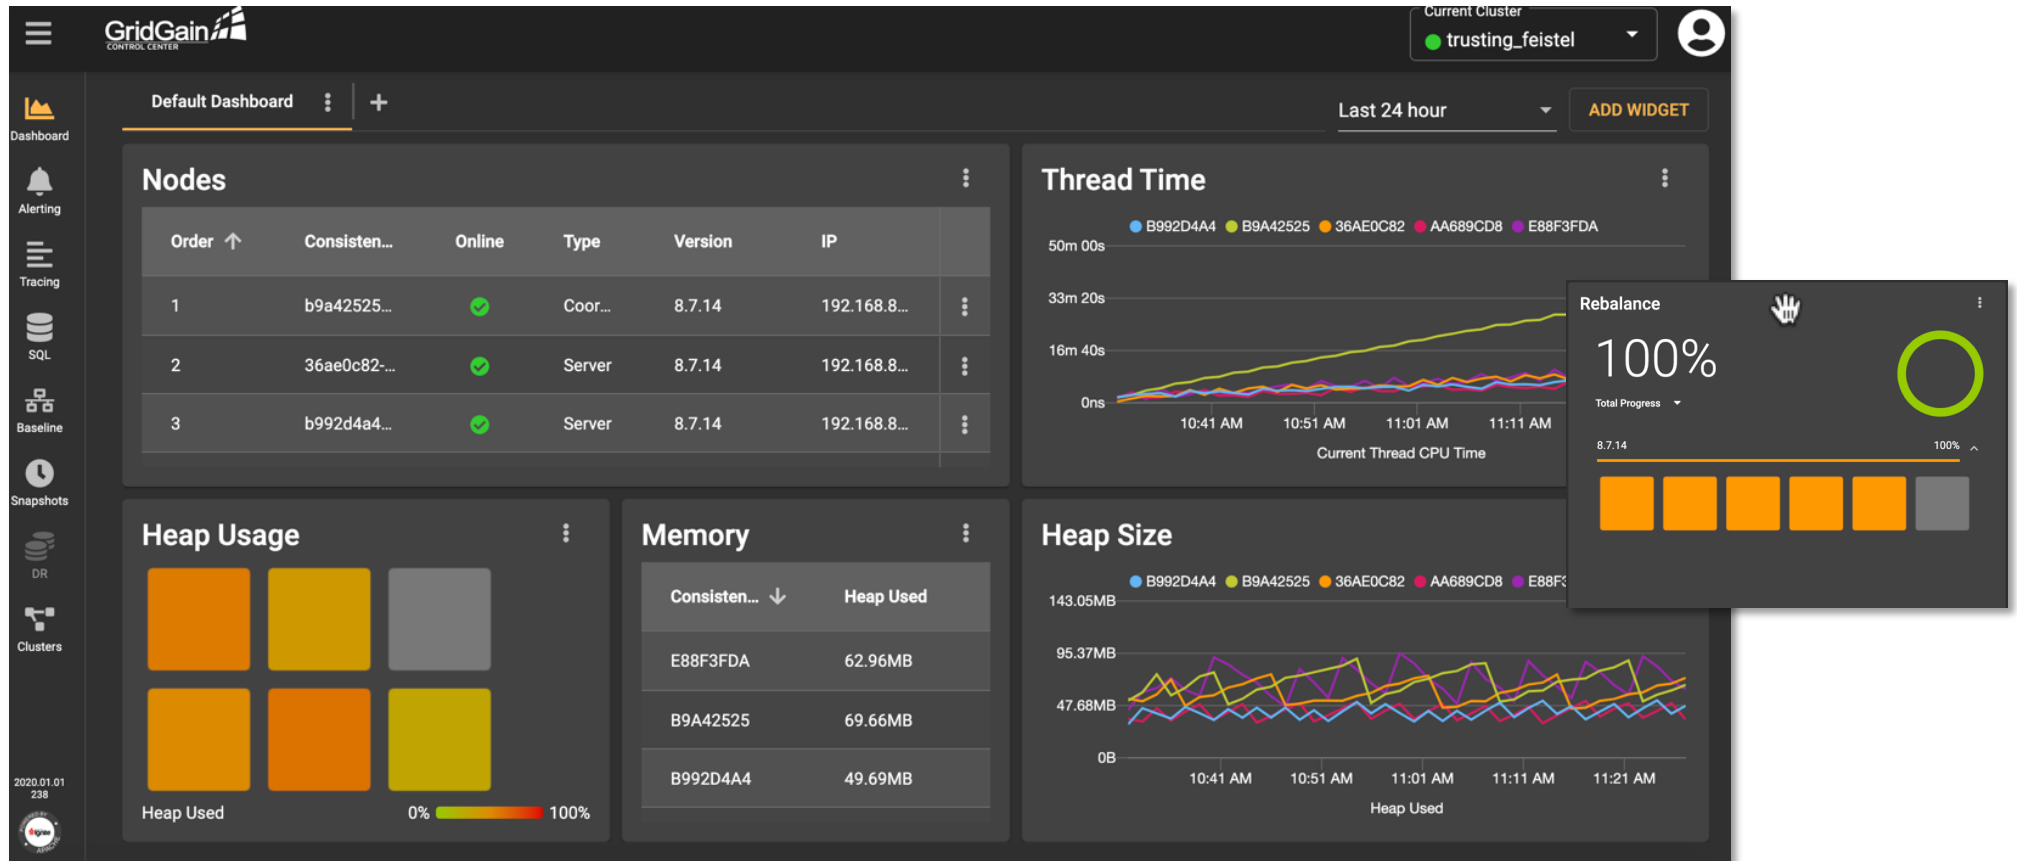
GridGain Control Center Overview



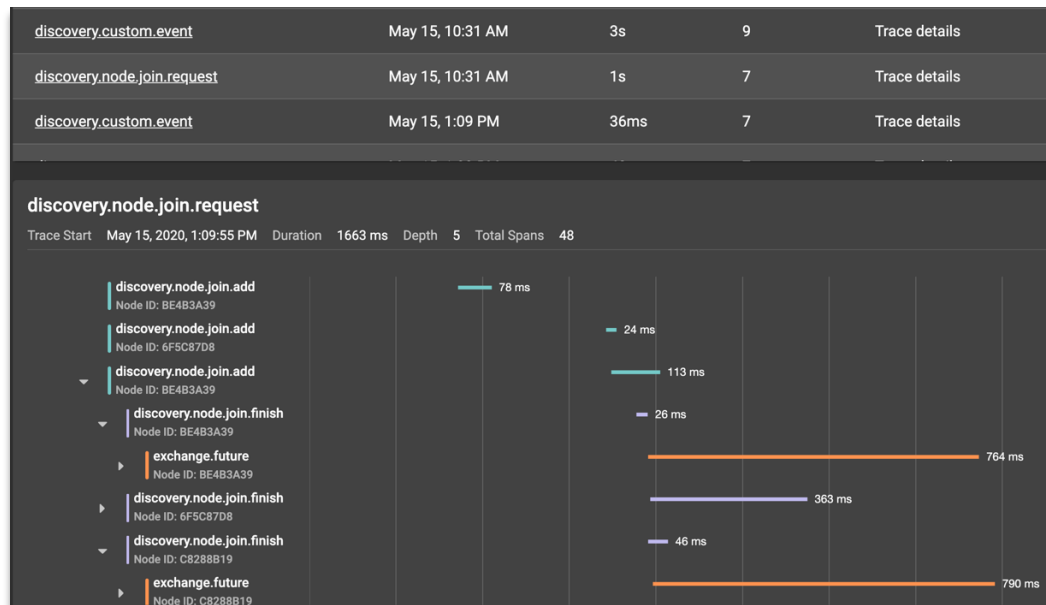
- Comprehensive Cluster Management and Developer Tool for GridGain and Apache Ignite
- Supported Products
 - Built for GridGain 8.7+ and Apache Ignite 2.8+
 - Apache Ignite requires Control Center Agent install



Modern Interface with Drag & Drop Flexibility



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

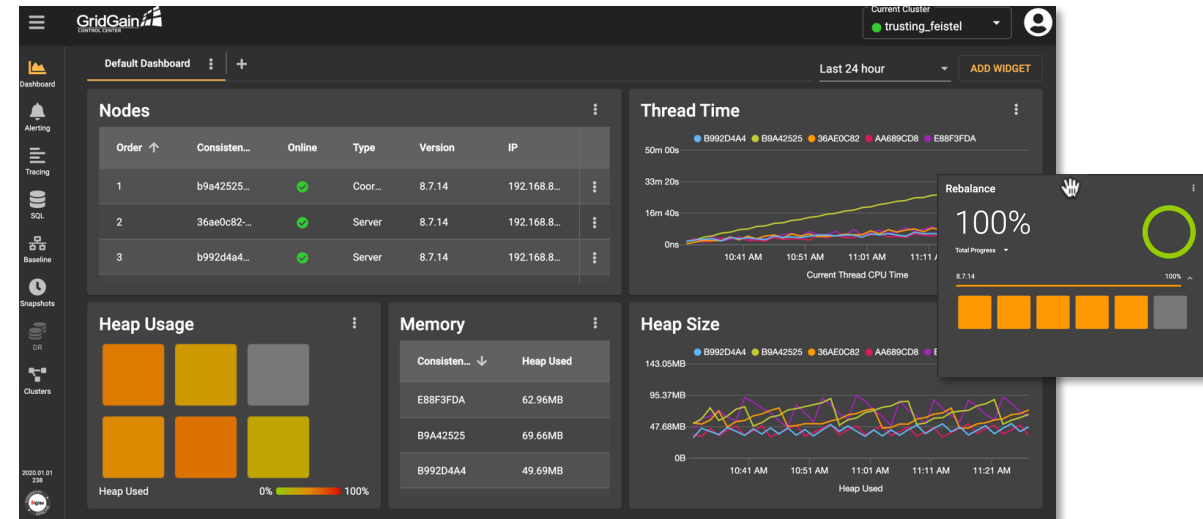
- Setup:
 - Copy *ignite-opencensus* from */libs/optional* to */libs/*
 - Add Tracing to Cluster Configuration

```
<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	

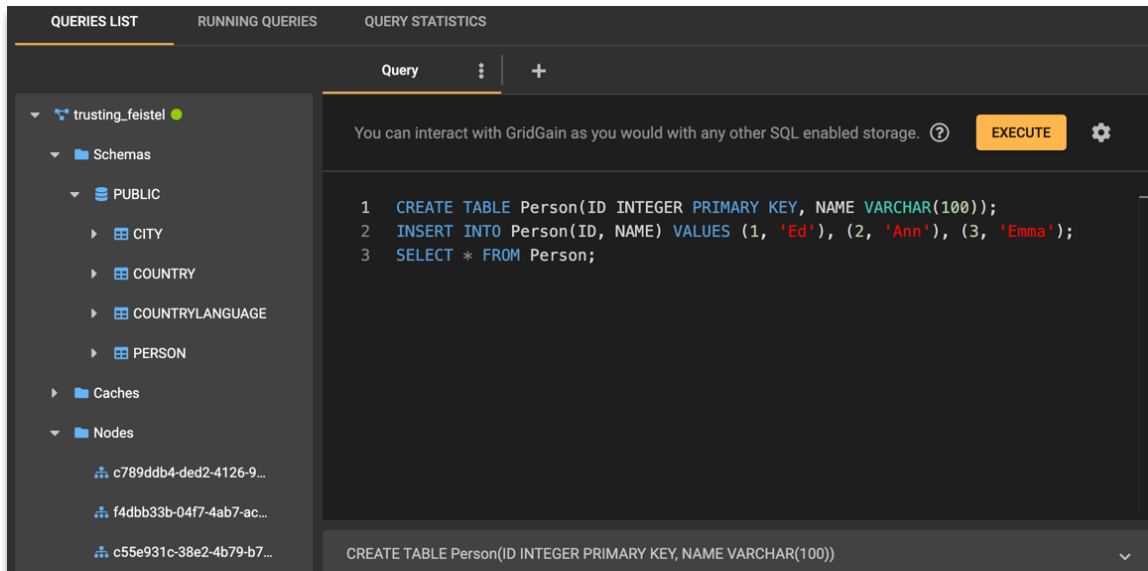
CONFIGURATIONS		EVENTS		NOTIFICATION CHANNELS			
Created On ↓	State	Message	Alert				
5/15/20, 10:41 AM	Alerting	Node "Heap Used" metric was > 10000000	Memory				
5/15/20, 10:41 AM	Alerting	Node "Heap Used" metric was > 10000000	Memory				
5/15/20, 10:41 AM	Alerting	Node "Heap Used" metric was > 10000000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



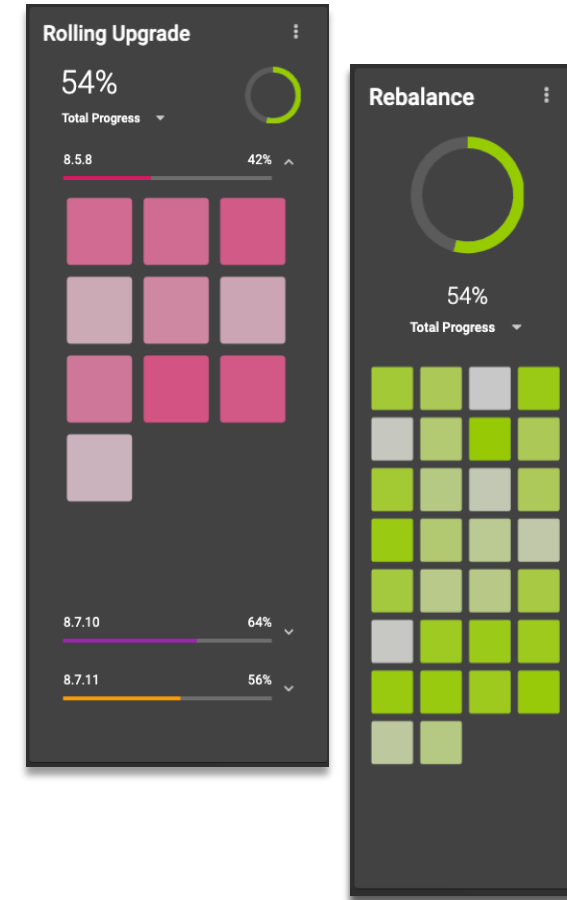
- Accelerate Query Development
 - Fine tune for optimal performance
- Develop and Run SQL Scripts
 - Introspect SQL execution
 - Analyze performance of long running queries and query history
 - Display all clusters, caches, and schemas in real-time



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

SNAPSHOTS						CREATE SNAPSHOT
Start Time ↓	Type	ID	Mode	Status	Caches	Filters
Jun 17, 7:58	INCREMENTAL	1592405892619	MANUAL	✓ OK	4	Snapshot Type
Jun 17, 7:58	INCREMENTAL	1592405881033	MANUAL	✓ OK	4	Schedule Name
Jun 17, 7:57	FULL	1592405852953	MANUAL	✓ OK	4	Mode
Jun 17, 7:56	INCREMENTAL	1592405815761	MANUAL	✓ OK	2	Cache Name
Jun 17, 7:56	FULL	1592405794065	MANUAL	✓ OK	2	Snapshot ID
Jun 17, 7:56	INCREMENTAL	1592405761056	MANUAL	✓ OK	1	<input type="checkbox"/> Show related snapshots
Jun 17, 7:55	FULL	1592405755578	MANUAL	✓ OK	1	Period
						From
						To

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 Morter](#)

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

