

# Best Practices for Monitoring Distributed In-Memory Systems

Denis Mekhanikov

June 4, 2019



# What communication with GridGain support often looks like



Customer: The cluster is hanging.

**GG:** Please send logs.

Customer: We don't have logs.

**GG:** Did you take thread dumps?

Customer: Nope.

**GG:** The problem is probably in GC.

What is the memory consumption level?

Customer: ...



# Why should we monitor?




- Check if everything is fine
  - Prevent upcoming issues
  - Discover and react to the issues that already happened
- Dashboarding**
- Find a reason for an issue and prevent it from happening again
- Logging**

# What to monitor?




- Every node in isolation
- Connection between nodes
- System as a whole



# Every node is...



- Hardware (hypervisor)
- Operating System
- Virtual machine
- Application



# Hardware / Hypervisor / OS


- ! CPU
- ! Memory
- ! Disk
  - System logs
  - Cloud Provider's logs



# Network



- Ping monitoring
- Network hardware monitoring
-  TCP dumps





- ! GC logs
  - JMX
- 🔥 Java Flight Recorder
- 🔥 Thread Dumps
- 🔥 Heap Dumps
- ! `java -XX:+HeapDumpOnOutOfMemoryError ...`





# Application

- ! Logs
  - JMX
  - Throughput / Latency
  - Test queries



# Tools



# Tools



## Metrics



# Tools



## Logs

splunk<sup>®</sup>>



# Tools



## JVM



Java Flight Recorder

GCeasy




MAT

VisualVM

# Tools



## Network



# Tools



## Benchmarking



# Tools



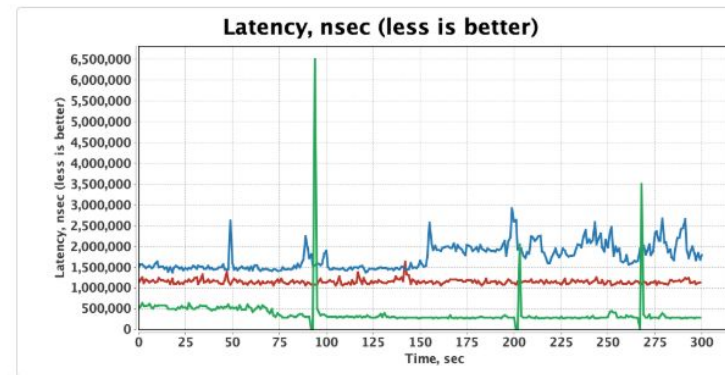
## Benchmark Compound Results

Color	Benchmark	Configurations
■	IgniteGetBenchmark	RELEASE-atomic-get-1-backup
■	IgnitePutGetBenchmark	RELEASE-atomic-put-get-1-backup
■	IgnitePutBenchmark	RELEASE-atomic-put-1-backup

### ThroughputLatencyProbe



	Avg	Min	Max	SD
■	188,831.77	0.00	237,678.00	52,727.19
■	37,976.28	22,007.00	46,892.00	5,925.58
■	56,020.37	44,341.00	60,427.00	2,457.01



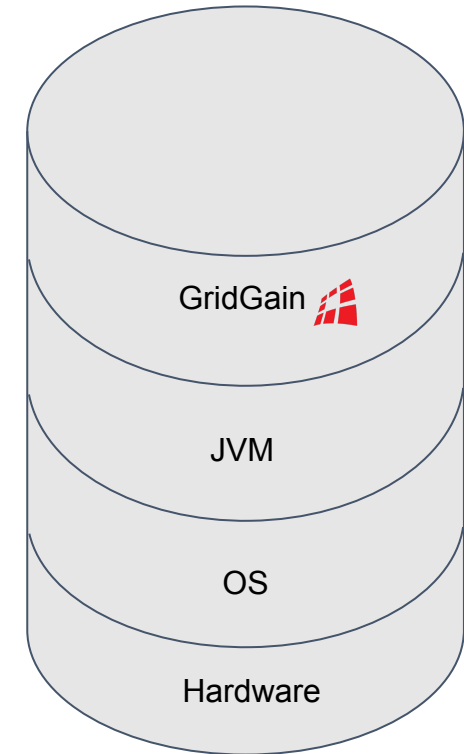
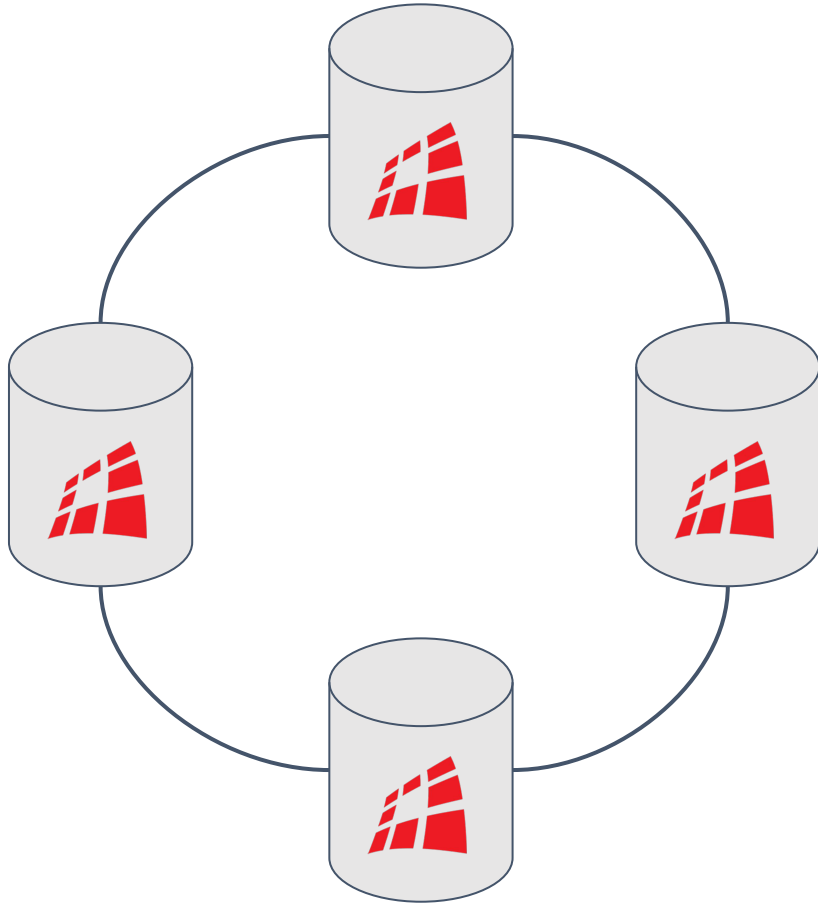
	Avg	Min	Max	SD
■	383,936.50	0.00	6,494,206.73	425,184.85
■	1,734,200.80	1,365,531.29	2,907,411.96	304,816.73
■	1,146,314.69	1,058,484.76	1,623,905.21	56,636.81



# GridGain



# GridGain



# GridGain: Cache Metrics



## CacheMetricsMXBean

- CacheGets
- AverageGetTime
- AverageTxCommitTime
- ...

## CacheGroupMetricsMXBean

- LocalNodeMovingPartitionsCount
- ClusterMovingPartitionsCount
- ClusterOwningPartitionsCount
- ...

# GridGain: Cache Metrics



## How to enable cache metrics

```
CacheConfiguration<K, V> cacheCfg = new CacheConfiguration<>("cache");
```

```
// Enable metrics.
```

```
cacheCfg.setStatisticsEnabled(true);
```

```
ignite.createCache(cacheCfg);
```

# GridGain: Discovery and Communication




## TcpDiscoverySpiMBean

- MessageWorkerQueueSize
- AvgMessageProcessingTime
- Coordinator
- NodesFailed
- ...

## TcpCommunicationSpiMBean

- OutboundMessagesQueueSize
- SentMessagesCount
- ReceivedMessagesCount
- ...

# GridGain: Data Storage



# GridGain: Data Storage Metrics



## Data volume

### DataStorageMetricsMXBean

- WalTotalSize
- TotalAllocatedSize
- OffheapUsedSize
- ...

### DataRegionMetricsMXBean

- TotalAllocatedPages
- AllocationRate
- PagesFillFactor
- ...


# GridGain: Data Storage Metrics



## Checkpoints

### DataStorageMetricsMXBean

- DirtyPages
- CheckpointTotalTime
- LastCheckpointDuration
- UsedCheckpointBufferSize
- LastCheckpointPagesWriteDuration
- LastCheckpointMarkDuration
- LastCheckpointTotalPagesNumber
- ...





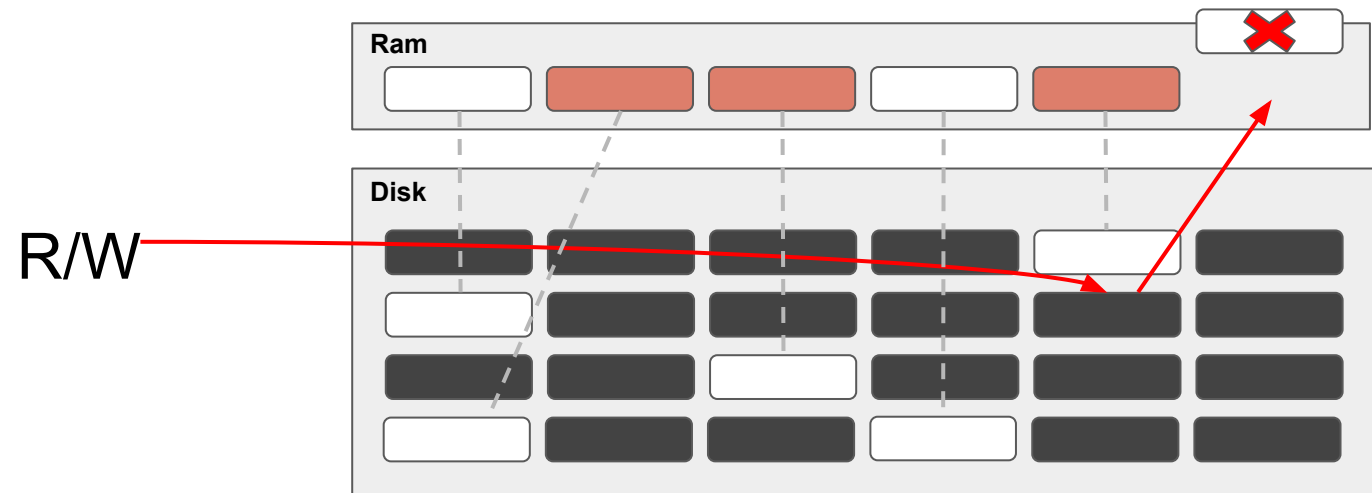
# GridGain: Data Storage Metrics



## Page replacement

### DataRegionMetricsMXBean

- PagesReplaceRate
- PagesReplaceAge
- PagesReplaced



# GridGain: Data Storage Metrics



## How to enable data storage metrics

```
DataStorageConfiguration storageCfg = new DataStorageConfiguration();  
DataRegionConfiguration regionCfg = new DataRegionConfiguration();  
regionCfg.setName("myDataRegion");
```

```
// Enable metrics.
```

```
storageCfg.setMetricsEnabled(true); // Metrics for data storage.  
regionCfg.setMetricsEnabled(true); // Metrics for a particular data region.
```

```
storageCfg.setDataRegionConfigurations(regionCfg);
```

# GridGain: IO metrics

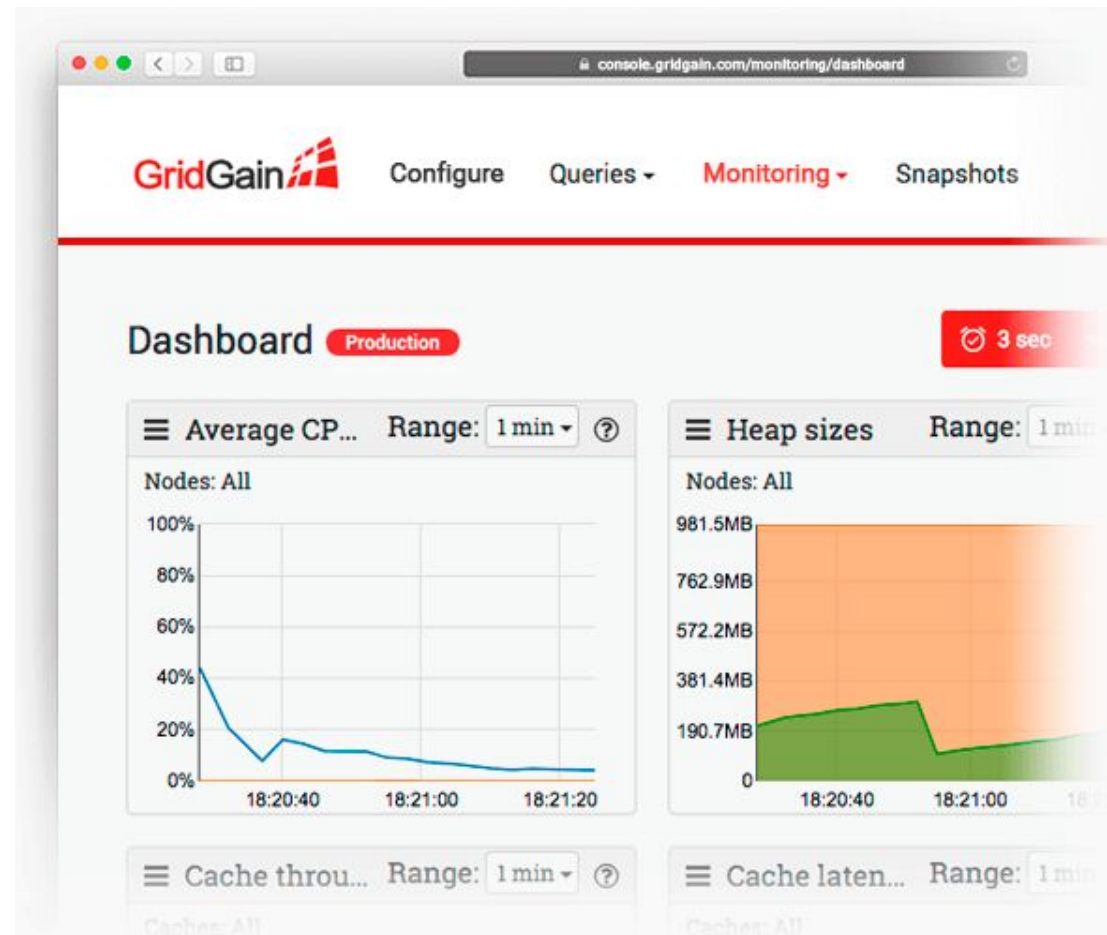


## Coming in 2.8

### IoStatisticsMetricsMXBean

- CacheGroupLogicalReads
- CacheGroupPhysicalReads
- IndexLogicalReads
- IndexPhysicalReads
- ...

# GridGain: WebConsole






# Demo

# Checklist for monitoring




- CPU / Memory / Disk / Network
  - GC logs
  - Application logs
- + Problematic places specific to your setup

# Q&A



 <https://console.gridgain.com/>

 <https://github.com/dmekhanikov/ignite-elk/>

 [dmekhanikov@gmail.com](mailto:dmekhanikov@gmail.com)