

# IN-MEMORY DATA FABRIC: Data Grid

#### COPYRIGHT AND TRADEMARK INFORMATION

© 2014 GridGain Systems. All rights reserved. This document is provided "as is". Information and views expressed in this document, including URL and other web site references, may change without notice. This document does not provide you with any legal rights to any intellectual property in any GridGain product. You may copy and use this document for your internal reference purposes. GridGain® is a registered trademark of GridGain Systems, Inc. Windows, .NET and C# are either registered trademarks or trademarks of Microsoft Corporation in the United States and/ or other countries. JEE and Java are either registered trademarks or trademarks of SUN Microsystems and/or Oracle Corporation in the United States and/or other countries. All other trademarks and trade names are the property of their respective owners and used here for identification purposes only.

**GRIDGAIN.COM** 

WHITE PAPER



## Table of Contents

| Re-Imagining Ultimate Performance                        |
|--|
| Introducing the GridGain In-Memory Data Fabric           |
| Features   |
| In-Memory Data Grid at a Glance                          |
| In-Memory Data Grids and Fast Data                       |
| Real-time Fraud Detection                                |
| Biometrics and Border Security                           |
| Financial Risk Analytics                                 |
| GridGain In-Memory Data Grid vs. Other Solutions         |
| Data Grid Feature of the GridGain In-Memory Data Fabric7 |
| Local, Replicated and Partitioned Distribution Modes     |
| Local Mode   |
| Replicated Mode  |
| Partitioned Mode   |
| Off-heap Memory  |
| Distributed ACID Transactions                            |
| HyperLocking Technology                                  |
| Multi Version Concurrency Control (MVCC)                 |
| In-Memory SQL Queries                                    |
| Data Preloading  |
| Delayed and Manual Preloading                            |
| Pluggable Persistent Store                               |
| Read-Through and Write-Through                           |
| Refresh-Ahead  |
| Write-Behind Caching                                     |
| Fault Tolerance and Data Resiliency                      |
| Datacenter Replication                                   |
| Management   |
| End-to-End Stack & Total Integration                     |
| Core Technology  |
| GridGain Foundation Layer                                |
| Hyper Clustering   |
| Zero Deployment  |
| Advanced Security  |
| SPI Architecture and PnP Extensibility13                 |
| Remote Connectivity                                      |
| About GridGain   |



## Re-Imagining Ultimate Performance

#### What is In-Memory Computing?

Data volumes and ever decreasing SLAs have overwhelmed existing diskbased technologies for many operational and transactional data sets, requiring the industry to alter its perception of performance and scalability. In order to address these unprecedented data volumes and performance requirements a new solution is required.

With the cost of system memory dropping 30% every 12 months In-Memory Computing is rapidly becoming the first choice for a variety of workloads across all industries. In fact, In-Memory Computing paves the way to a lower TCO for data processing systems while providing an undisputed performance advantage.

## Introducing the GridGain In-Memory Data Fabric

In-Memory Computing is characterized by using high-performance, integrated, distributed memory systems to manage and transact on large-scale data sets in real time, orders of magnitude faster than possible with traditional disk-based technologies.

The GridGain In-Memory Data Fabric is a proven software solution, which delivers unprecedented speed and unlimited scale to accelerate your business and time to insights. It enables high-performance transactions, real-time streaming and fast analytics in a single, comprehensive data access and processing layer. The In-Memory Data Fabric is designed to easily power both existing and new applications in a distributed, massively parallel architecture on affordable, industry-standard hardware.

The GridGain In-Memory Data Fabric provides a unified API that spans all key types of applications (Java, .NET, C++) and connects them with multiple data stores containing structured, semi-structured and unstructured data (SQL, NoSQL, Hadoop). It offers a secure, highly available and manageable data environment that allows companies to process full ACID transactions and generate valuable insights from real-time, interactive and batch queries.

The In-Memory Data Fabric offers a strategic approach to in-memory computing that delivers performance, scale and comprehensive capabilities far above and beyond what traditional in-memory databases, data grids or other in-memory-based point solutions can offer by themselves.







#### **FEATURES**

The GridGain In-Memory Data Fabric accesses and processes data from distributed enterprise and cloud-based data stores orders of magnitudes faster, and shares them with today's most demanding transactional, analytical and hybrid applications with varying SLA requirements (real-time, interactive, batch jobs). To that effect, the GridGain In-Memory Data Fabric includes the following four features:







Clustering and Compute Grid



Hadoop Acceleration

The data grid feature in the GridGain In-Memory Data Fabric supports local, replicated, and partitioned data sets and allows to freely cross query between these data sets using standard SQL syntax. It supports standard SQL for querying in-memory data including support for distributed SQL joins. The GridGain In-Memory Data Fabric offers an extremely rich set of data grid capabilities, including off-heap memory support, load-balancing, fault tolerance, remote connectivity, support for full ACID transactions and advanced security.

In-Memory clustering and compute grids are characterized by using high-performance, integrated, distributed memory systems to compute and transact on large-scale data sets in real-time, orders of magnitude faster than possible with traditional disk-based or flash technologies.

The real-time streaming feature of the GridGain In-Memory Data Fabric uses programmatic coding with rich data indexing support to provide CEP querying capabilities over streaming data. The GridGain In-Memory Data Fabric also provides comprehensive support for customizable event workflow.

Hadoop acceleration included in the GridGain In-Memory Data Fabric features the GridGain in-memory file system (GGFS). It has been designed to work in dual mode as either a standalone primary file system in the Hadoop cluster, or in tandem with HDFS, serving as an intelligent caching layer with HDFS configured as the primary file system.

## In-Memory Data Grid at a Glance

#### What is an In-Memory Data Grid?

As traditional approaches to application architecture based on spinning disk technologies struggle to keep up with the ever expanding data volumes and velocities inherent in today's enterprise applications, a faster, scalable alternative is required, and organizations are increasingly considering In-Memory Data Grid features as the cornerstone of their next generation development efforts.

In-Memory Data Grids (IMDG) are characterized by the fact that they store all of their data in-memory as opposed to traditional Database Management Systems that utilize disk as their primary storage mechanism. By utilizing system memory rather than spinning disk, IMDGs are typically orders of magnitude faster than traditional DBMS systems.

Keeping data in memory is not the only reason why IMDGs perform significantly faster than disk-based databases. The main reason for performance difference are the actual differences in the architecture. IMDGs are specifically designed with memory-first and disk-second approach where memory is utilized as a primary storage and disk as a secondary storage for backup and persistence. Since memory is a much more limited resource than disk, IMDGs are



built from ground up with a notion of horizontal scale and ability to add nodes on demand in real-time. IMDGs are designed to linearly scale to hundreds of nodes with strong semantics for data locality and affinity data routing to reduce redundant data movement.

Disk based databases, on the other hand, are designed with disk-first and memory-second architecture and are primarily optimized for disk-based, block device access. Even though disk-based databases do employ caching, generally caching is just a thin layer holding a small portion of overall data in memory, while absolute majority of the data is still stored on disk. When the data access pattern goes beyond of what little can fit in cache, disk-based system suffer from heavy paging overhead of data being brought in and out of memory. This is true for traditional SQL RDBMS as well as for the vast majority of new NoSQL and NewSQL systems that are predominately disk-based.



Since disk is a virtually unlimited resource when compared to memory, most disk-based systems can often hold the whole or most of the data set and are rarely designed to horizontally scale simply because there is no need for it from data storage standpoint. Such architecture makes disk based systems much less suited for parallel processing and often results in possible database overloading and thrashing.

IMDGs, however, due to linear scalability and in-memory data partitioning, can localize processing logic to the nodes where the data actually resides. This makes IMDGs much more suited for parallel processing as multiple CPUs across the cluster can be utilized to process computations that work on different data sets in parallel.

Ability to collocate computations with the data makes IMDGs much more suited for parallel processing than traditional disk-based systems.

## In-Memory Data Grids and Fast Data

#### What problems does an In-Memory Data Grid solve?

Gartner defines the "three V's" of Big Data as: Velocity, Volume and Variety. Being able to easily handle these three considerations are core to a Fast Data strategy. For example, Fast Data often encompasses the ingestion and analysis of streaming and/or transactional data (velocity). While there is certainly an ever-growing amount of data (volume) and the sources and types of data in the enterprise are expanding (variety), the key notion to understand about Fast Data is the type of the latencies at which data can be processed to produce actionable insights or machine-driven decision-making. Succinctly, the key notion that makes data "Fast" rather than just "Big" is the actionable analysis of information in near real-time.



It is helpful to consider some of representative Fast Data use cases where GridGain's IMDG functionality is used:

#### **REAL-TIME FRAUD DETECTION**

The ever-increasing speed of commerce and the growing sophistication and organization of the people with a criminal intent create an extremely difficult problem for financial institutions. Transaction processing gateways no longer simply have to deal with a high-volume of transactions, they now must be able to determine, in milliseconds, the legitimacy of a given transaction. This task is further complicated by the fact that transactions cannot be considered in isolation, but must be considered in the context of multiple transactions as well as transaction history. As the sophistication of the people committing fraud has increased, traditional post facto analysis has become effectively useless. By the time an alert is generated, the fraudsters have discontinued use of the credit card and have moved on to the next target. With a Fast Data solution, transactions can be analyzed in real-time for suspicious patterns as compared against historical purchase history to detect fraud in real time and deliver a decision on the legitimacy of a transaction in real time.

#### **BIOMETRICS AND BORDER SECURITY**

Border security is a growing concern in today's global climate. The ability for front-line border security personnel to have at their disposal the most up-to-date information and risk models is critical. With a Fast Data solution, things such as passport scans and biometric information for every incoming passenger presenting themselves at a port of entry can be quickly processed against vast data sets, as well as run through complex risk models to identify potential security risks before entry is granted.

#### FINANCIAL RISK ANALYTICS

The speed with which market data can be ingested and more importantly analyzed against open positions in a given portfolio, is a key component to a given firm's competitiveness and ultimately profitability. As market data ages, it becomes less useful for driving trading decisions leaving the slower firm at a disadvantage. Consequently, data must be ingested and risk calculations performed as quickly as possible. With a Fast Data solution, the streaming data can be consumed as fast as it is available from the information provider and risk calculations can be performed in parallel in near real-time. Furthermore, more complex modeling can be done with no drop in latencies simply by adding additional nodes.

The commonalities across these use cases are apparent in that they all share a low-latency service-level agreement (SLA) and require complex analytics to be performed across a high volume of data. In the following section, the functionality of GridGain's IMDG feature – included in the GridGain In-Memory Data Fabric – is compared to other solutions.

## GridGain In-Memory Data Grid vs. Other Solutions

What makes GridGain's In-memory data grid feature a unique solution?

In-Memory data grids have been around for a while, but recently the availability of comparatively inexpensive memory has triggered development of a variety of new in-memory data management platforms.

The GridGain In-Memory Data Fabric supports local, replicated, and partitioned data sets and allows to freely cross query between these data sets using standard SQL syntax.



One of the distinguishing features for most IMDGs is the way data replication is supported. There are several common approaches ranging from no replication at all to full replication where every node sees absolutely identical data set. Perhaps the most interesting approach is partitioned approach where each node gets a chunk of data which it is responsible for. Systems that support this approach are usually most scalable, as with the addition of more nodes in the cluster more data can be stored in memory. However, what makes an IMDG a powerful capability is not efficient support for some replication mode, but the ability to use different replication modes together on different data sets within the same application, as well as the ability to easily cross-query between these data sets in real-time.

Another important characteristic of an IMDG system is data accessibility and search features of the provided query language. Most IMDGs will provide access to the data by a primary key. Much fewer systems will provide a custom expressive query language for richer data querying capabilities. Very few systems will allow to use standard SQL for data querying. Systems which allow for execution of distributed

The GridGain In-Memory Data Fabric supports standard SQL for querying in-memory data including support for distributed SQL joins.

joins on the data are usually the most rare. Ability to execute standard SQL joins in a distributed environment is usually one of the hardest features to support, but when done right, provides the most elegant and efficient way for accessing data.

## Data Grid Feature of the GridGain In-Memory Data Fabric

#### What are the key capabilities of GridGain's in-memory data grid feature?

The goal of an in-memory data grid is to provide extremely high availability of data by keeping it in memory and in highly distributed (i.e. parallelized) fashion. The GridGain in-memory data grid subsystem is fully integrated into the core of the GridGain In-Memory Data Fabric, and is built on top of the existing functionality such as pluggable auto-discovery, communication, marshaling, on-demand class loading, and support for functional programming.

#### LOCAL, REPLICATED AND PARTITIONED DISTRIBUTION MODES

The GridGain In-Memory Data Fabric provides 3 different modes of cache operation: Local, Replicated, and Partitioned:

#### LOCAL MODE

Local mode is the most light-weight mode of cache operation, as no data is distributed to other cache nodes. It is ideal for scenarios where data is either read-only, or can be periodically refreshed at some expiration frequency. It also works very well with read-through behavior where data is loaded from persistent storage on misses. Other than distribution, local caches still have all the features of a distributed cache, such as automatic data eviction, expiration, disk swapping, data querying, and transactions.

#### **REPLICATED MODE**

In REPLICATED mode all data is replicated to every node in the grid. This cache mode provides the utmost availability of data as it is available on every node. However, in this mode every data update must be propagated to all other nodes which can have an impact on performance and scalability.

As the same data is stored on all grid nodes, the size of a replicated cache is limited by the amount of memory available on the node with the smallest amount of RAM. This means that no matter how many grid nodes you have in your data grid, the maximum amount of data you can cache does not change.



This mode is ideal for scenarios where cache reads are a lot more frequent than cache writes, and data availability is the most important criteria for your use case.

#### PARTITIONED MODE

Partitioned mode is the most scalable distributed cache mode. In this mode the overall data set is divided equally into partitions and all partitions are split equally between participating nodes, essentially creating one huge distributed memory for caching data. This approach allows you to store as much data as can be fit in the total memory available across all nodes. Essentially, the more nodes you have, the more data you can cache.

Unlike Replicated mode, where updates are expensive because every node in the grid needs to be updated, with Partitioned mode, updates become cheap because only one primary node (and optionally 1 or more backup nodes) need to be updated for every key. However, reads become somewhat more expensive because only certain nodes have the data cached. In order to avoid extra data movement, it is important to always access the data exactly on the node that has that data cached. This approach is called affinity colocation and is strongly recommended when working with partitioned caches.

#### **OFF-HEAP MEMORY**

Off-Heap overflow provides a mechanism in grid by which grid cache or any other component can store data outside of JVM heap (i.e. in off-heap memory). By allocating data off-heap, Java Virtual Machine Garbage Collection (JVM GC) does not know about it and hence does not slow down. In fact you can start your Java application with a relatively small heap, e.g. below 512M, and then let GridGain utilize 100s of Gigabytes of memory as off-heap data cache. One of the distinguishing characteristics of GridGain off-heap memory is that the on-heap memory foot print is constant and does not grow as the size of off-heap data grows.

The GridGain In-Memory Data Fabric provides a tiered approach to memory by allowing data to migrate between On-Heap, Off-Heap and Swap storages. By doing so, the GridGain product avoids the well known issues with JVM Garbage Collection (GC) when attempting to utilize large on-heap memory configurations.

Traditionally, GC pauses were mitigated by starting multiple JVMs on a single server. However, this is less than ideal for applications that need to co-locate large amounts of data in a single JVM for low latency processing requirements.

#### DISTRIBUTED ACID TRANSACTIONS

The GridGain In-Memory Data Fabric supports Distributed ACID Transactions that use a two-phase commit (2PC) protocol to ensure data consistency within the cluster. The 2PC protocol in GridGain's product has been significantly optimized to ensure minimal network chattiness and lock-free concurrency.

Distributed data grid transactions in the GridGain In-Memory Data Fabric span data on both local and remote nodes. While automatic enlisting into JEE/JTA transactions is supported, the GridGain In-Memory Data Fabric also allows users to create more light-weight transactions which are often more convenient to use.

GridGain transactions support all ACID properties that you would expect from any transaction, including support for Optimistic and Pessimistic concurrency levels and Read-Committed, Repeatable-Read, and Serializable isolation levels. If a persistent data store is configured, then the transactions will also automatically span the data store.

#### HYPERLOCKING® TECHNOLOGY

HyperLocking is one of ways to achieve even better performance with GridGain 2-phase-commit (2PC) transactions.



With the standard 2PC approach, a separate lock will be acquired, either optimistically or pessimistically, for every key within transaction. If keys are backed up on other nodes, then these locks will have to be propagated to backup nodes, sequentially, which may impose additional latencies and network overhead.

When HyperLocking is used, only one lock per transaction is acquired even though transaction may be modifying 1000s of cache entries. In pessimistic mode the lock is acquired at the beginning of the transaction, and in optimistic mode it is acquired at commit phase. The requirement is that all elements within transaction are grouped by the same affinity key or partition ID.

HyperLocking can provide up to 100x performance boost over a standard 2PC approach and is one of the main reasons why the GridGain In-Memory Data Fabric outperforms most data grid or database systems by an order of magnitude.

#### MULTI VERSION CONCURRENCY CONTROL (MVCC)

The GridGain In-Memory Data Fabric utilizes an advanced and optimized implementation of MVCC (Multi Version Concurrency Control). MVCC provides practically lock-free concurrency management by maintaining multiple versions of data instead of using locks with a wide scope. Thus, MVCC in the GridGain In-Memory Data Fabric provides a backbone for high performance and overall system throughput for systems under load.

#### **IN-MEMORY SQL QUERIES**

The GridGain In-Memory Data Fabric provides the ability to query data using standard SQL. SQL can be issued via API-based mechanisms, or via a read-only JDBC interface.

There are almost no restrictions as to which SQL syntax can be used. All inner, outer, or full joins are supported, as well as the rich set of SQL grammar and functions. The ability to join different classes of objects stored in the DBMS or across different caches makes GridGain queries a very powerful tool. All indices are usually kept in memory, either on-heap or off-heap, resulting in very low latencies for query execution.

In addition to SQL queries, the GridGain In-Memory Data Fabric also supports Text queries by utilizing Lucene text engine for efficient indexing of text data.

#### DATA PRELOADING

Whenever a node joins or leaves the topology, the remaining nodes will make an attempt to preload or repartition all values from the other nodes. This way, the data in the grid remains consistent and equally balanced. In the case where data is replicated to all nodes the new node will try to load the full set of data from the existing nodes. When data is partitioned across the grid, the current node will only try to load the entries for which the current node is either primary or back up.

Data preloading can be either synchronous or asynchronous. In synchronous mode distributed caches will not start until all necessary data is loaded from other available grid nodes. Essentially, existing data grid nodes will keep operating as usual, but the new nodes will not be allowed to start until they have all the necessary data.

If asynchronous mode is turned on, which is the default behavior, distributed caches will start immediately and will load all necessary data from other available grid nodes in the background. This mode is the most efficient preloading mode and should be utilized unless application logic has strong requirements for synchronous mode.



#### DELAYED AND MANUAL PRELOADING

There are some cases when it's not efficient to start preloading right after a new node starts or an existing node leaves. For example, if you start multiple nodes one by one, you may want to wait until all nodes are started to avoid multiple preloading processes to take effect one after another for each individual node start. Or, if you plan to restart existing nodes, it's better to start pre-loading only after nodes are restarted, and not right after they are stopped.

To support such cases, the GridGain In-Memory Data Fabric provides a way to delay pre-loading. Purposely delaying preloading process often allows to significantly reduce network overhead and improve system performance.

Note that delaying preloading does not have any effect on data consistency. The data in the grid always remains consistent, regardless of when preloading starts or completes.

#### PLUGGABLE PERSISTENT STORE

In-memory data grids are often used in conjunction with an external persistent data store, such as a disk-based database or a file system. Whenever a persistent store is configured, loading data and updating data into the data store is automatically handled by the system. Moreover, the persistent store plugs into data grid transactions, so updating the underlying persistent store is part of the same transaction as the data grid update; if one fails the other fails as well.

#### **READ-THROUGH AND WRITE-THROUGH**

Properly integrating with persistent stores is important whenever read-through or write-through behavior is desired. Read-through means that data will be automatically read from the persistent store whenever it is not available in cache, and write-through means that data will be automatically persisted whenever it is updated in cache.

All read-through and write-through operations will participate in overall cache transactions and will be committed or rolled back as a whole.

#### **REFRESH-AHEAD**

Refresh-ahead automatically reloads data from the persistent store whenever it has expired in memory. Data can expire in memory whenever its time-to-live value is passed. To prevent reloading data every time data is expired, refresh-ahead ensures that entries are always automatically re-cached whenever they are nearing expiration.

#### WRITE-BEHIND CACHING

In a simple write-through mode each cache put and remove operation will involve a corresponding request to the storage and therefore the overall duration of the cache update might be relatively long. Additionally, an intensive cache update rate can cause an extremely high load on the underlying persistent store.

For such cases, the GridGain In-Memory Data Fabric offers an option to perform an asynchronous storage update also known as write-behind. The key concept of this approach is to accumulate updates and then asynchronously flush them to the persistent store as a bulk operation.

In addition to obvious performance benefits, because cache writes simply become faster, this approach scales a lot better, assuming that your application can tolerate delayed persistence updates. When the number of nodes in the topology grows and every node performs frequent updates, it is very easy to overload the underlying persistent store. By using a write-behind approach one can maintain a high throughput of writes in the In-Memory Data Fabric without



bottlenecking at the persistence layer. Moreover, the In-Memory Data Fabric can continue operating even if your underlying persistence store crashes or goes down. In this case the persistence queue will keep storing all the updates until the data grid comes back up.

#### FAULT TOLERANCE AND DATA RESILIENCY

A common misconception about in-memory data grids in general is that you lose all data in the event of a node failure. This is simply not the case if you have architected your grid topology correctly. In a partitioned scenario, you are free to configure as many back-ups as you wish. In the event of a primary node failure, the grid will automatically promote the backup node to primary, and data access will continue without interruption.

Furthermore, you can write your data through to underlying persistent storage in a variety of ways. By writing through to persistent storage, you can then ensure that data is not lost in the event of node failures.

#### DATACENTER REPLICATION

When working with multiple data centers, it is important to make sure that if one data center goes down, another data center is fully capable of picking up its load and data. When data center replication is turned on, the GridGain In-Memory Data Fabric will automatically make sure that each data center is consistently backing up its data to other data centers (there can be one or more).

The GridGain In-Memory Data Fabric supports both active-active and active-passive modes for replication. In active-active mode both data centers are fully operational online and act as a backup copy of each other. In active-passive node, only one data center is active and the other data center serves only as a backup for the active data center.

Datacenter replication can be either transactional or eventually-consistent. In transactional mode, a transaction will be considered complete only when all the data has been replicated to another datacenter. If the replication step were to fail, then the whole transaction will be rolled back on both data centers. In eventually-consistent mode, the transaction will usually complete before the replication is completed. In this mode the data is usually concurrently buffered on one data center and then gets flushed to another data center either when the buffer fills up, or when a certain time period elapses. Eventually-consistent mode is generally a lot faster, but it also introduces a lag between updates on one data center and replication to another.

If one of the data centers goes offline, then another will immediately take responsibility for its data. Whenever the crashed data center goes back online, it will receive all the updates it has missed from another data center.

#### MANAGEMENT

The in-memory data grid feature of the GridGain In-Memory Data Fabric comes with a comprehensive and unified GUI-based management and monitoring tool called GridGain Visor. It provides deep operations, management and monitoring capabilities.

A starting point of the Visor management console is the Dashboard tab which provides an overview on grid topology, many relevant graphs and metrics, as well as event panel displaying all relevant grid events. To manage and monitor the data nodes you can select the Data Grid Tab which will display detailed information about the data grid feature. You can also manage and monitor individual caches within the data grid.



| <u>F</u> ile <u>V</u> iew |                                  |                      |  |                          |                            |                      |                              |                          |                      |           |                 |              |              |       |
|---------------------------|----------------------------------|----------------------|--|--------------------------|----------------------------|----------------------|------------------------------|--------------------------|----------------------|-----------|-----------------|--------------|--------------|-------|
| 📓 Log: F153B              | DF6 - gridgain-f15               | 3bdf6.log\.?\d       | d* × 🔄 🕤 File M  | Manager: 2 × 4           | Node: AF553C7              | 6 × 🔛 Pr             | ofiler: ggfs ×               | 💠 Node: F                | 153BDF6 ×            | Cache     | e: meta ×       |              |              |       |
| 🔊 Dashboard               | d 🗊 Databas                      | se 🏼 🍓 Cor           | npute 🛛 🔒 SC   | QL Viewer 🛛 🥰            | Telemetry                  | GGFS 🥠               | Streaming                    | 🟐 File Ma                | nager: 1 ×           | 🚺 Hos     | st: 192.168.1.2 | 2 × 🔳        | Cache: data  | ×     |
| Start<br>Vers             | ed: Jul 06, 2013 2<br>ion: 4.0.0 | 0:57:07 Conf<br>Grid | fig Path: modules<br>Name: <default< td=""><td>/visor/tests/config<br/>&gt;</td><td>j/visor-gui-test-lau</td><td>incher.xml T</td><td>otal CPUs: 8<br/>otal RAM: 16</td><td>Total H<br/>5.0GB Total N</td><td>losts: 1<br/>lodes: 2</td><td></td><td></td><td></td><td>📮 Ref</td><td>fresh</td></default<> | /visor/tests/config<br>> | j/visor-gui-test-lau       | incher.xml T         | otal CPUs: 8<br>otal RAM: 16 | Total H<br>5.0GB Total N | losts: 1<br>lodes: 2 |           |                 |              | 📮 Ref        | fresh |
|                           | Show 10                          | min 🖌                | CPU %  | GC 🛪 🛛 🗾 Hea             | р %                        |                      | Show 10                      | ▼ min ⊭                  | 🔲 Max Heap           | мв 🖌      | Alloc Heap мв   | 🗾 🗌 Used     | Неар мв      |       |
| 75% -                     |                                  |                      |  |                          |                            | ZCR                  |                              |                          |                      |           |                 |              |              |       |
| 50%                       |                                  |                      |  |                          |                            | 5GB -                |                              |                          |                      |           |                 |              |              |       |
| 25%-                      | 1                                |                      |  | AA                       | - Are                      | 2GB -                |                              |                          |                      |           |                 |              |              |       |
| 20                        | :58 20:59                        | 21:00 2:             | 1:01 21:02   | 21:03 21:                | 04 21:05                   | <u>~</u>             | 20:58                        | 20:59                    | 21:00 2              | 1:01 21   | :02 21:03       | 21:04        | 21:05        |       |
| Total Metrics (           | last update: 21:06               | :05)                 |  |                          |                            |                      |                              |                          |                      |           |                 |              |              |       |
| н                         | losts Nodes CF                   | Us                   | Curr. Avg  | g. GC                    | Pr                         | Px Ds                | Da                           | Part.                    | Repl. Loc.           |           | St 📔            | Ti           | Fa Na        |       |
| H/N/C:                    | 1 2                              | 8 CPU:               | 1.74% 5  | .43% 0%                  | GGFS: 0                    | 0                    | 0 1 0                        | Caches: 1                | 1 0                  | Sessions: | 6 736           | 0            | 0            | 0     |
|                           | Age Avg. Upt                     | ime                  | Used Fre   | e Comm.                  | Total                      | Used                 | Free                         | Size                     | e Keys               |           | Execs           | Avg. Time    | Unique       |       |
| Topology:                 | 9m 4s 9m 4                       | s Heap: 4            | \$12.7MB 95  | .47% 1006.0MB            | Space: 14.2GB              | 158.4KB              | 100.00%                      | Metrics: 6.0             | 21.5K                | Tasks:    | 742             | 3s 45ms      |              | 3     |
| All Nodes: 2              | Selected: 1 Hosts                | 1 Filter:            |  |                          |                            | -                    |                              | <i>4</i> 9               | 9 🚊                  | 2         |                 |              |              | ٨     |
| Node, ID8                 | Node IF                          | •                    | Start Time   | Up Time                  | CPUs C                     | PU Load, %           | GC Load, 9                   | % Free Hea               | ap, %                |           | OS Infor        | mation       |              |       |
| F153BDF6                  | 127.0.0.                         | 1 Ju                 | ul 06, 20:56:59  | 9m 4s                    | 8                          | 1.74%                |                              | 0% 9                     | 5.47%                |           | 💣 Mac OS X x8   | 86_64 10.8.4 |              |       |
| AF553C76                  | 127.0.0.                         | 1 Ju                 | ul 06, 20:56:59  | 9m 4s                    | 8                          | 1.02%                |                              | 0% 9                     | 5.47%                |           | 💣 Mac OS X x8   | 86_64 10.8.4 |              |       |
|                           |                                  |                      |  |                          |                            |                      |                              |                          |                      |           |                 |              |              |       |
|                           |                                  |                      |  |                          |                            |                      |                              |                          |                      |           |                 |              |              |       |
|                           |                                  |                      |  |                          |                            |                      |                              |                          |                      |           |                 |              |              |       |
| Event Log: 5              | Since: 20:57:07 N                | lessage Filter:      |  |                          |                            |                      |                              |                          | Q                    | Details   | 🙀 Clear All     |              |              | ٨     |
| Timestam                  | np 🛋 Nod                         | e, ID8 Typ           | e  |                          |                            |                      | Ν                            | Message                  |                      |           |                 |              |              |       |
| Jul 06, 20:               | 57:07 Vi                         | sor 🚺 INF            | New version 'y   | .y.y' is available. C    | lick <u>www.gridgain.c</u> | <u>om</u> to open in | your browser.                |                          |                      |           |                 |              |              |       |
| Jul 06, 21:               | 02:09 Vi                         | sor WAF              | R Telemetry: cao   | he partitions devia      | tion is more than 1        | 0%                   |                              |                          |                      |           |                 |              |              |       |
| Jul 06, 21:               | 02:12 Vi                         | sor WAF              | R Telemetry: cao   | he misses ratio is       | more than 30%              |                      |                              |                          |                      |           |                 |              |              |       |
| Jul 06, 21:               | 02:15 Vi                         | sor WAF              | R Telemetry: Stre  | eaming maximum v         | waiting queue size i       | s more than 1        | ,000                         |                          |                      |           |                 |              |              |       |
| Jul 06, 21:               | 04:18 Vi                         | sor ERF              | R Throttling: Viso   | or error: unhandled      | exception in EDT           |                      |                              |                          |                      |           |                 |              |              |       |
| I                         |                                  |                      |  |                          |                            |                      |                              |                          |                      |           |                 |              |              |       |
| ar Connected: 🕻           | ON Last Update                   | 21:06:05             | H/N/C: 1/2/8   | CPU Load: 5.             | 43% Free Heap              | : 95.47%             | Jp Time: <b>9m</b> 4         | 4s                       | Update: 3            | s 🖉 4     |                 |              | 511M of 1006 | м     |

## End-to-End Stack & Total Integration

#### What other GridGain technologies are available?

The GridGain In-Memory Data Fabric provides a comprehensive in-memory computing solution, from high performance computing, streaming, and data grids to Hadoop acceleration. It is a complete and fully integrated solution for low-latency, high performance computing for each and every category of payloads and data processing requirements. Total integration is further extended with a single unified management and monitoring console.

#### **CORE TECHNOLOGY**

The GridGain In-Memory Data Fabric is designed to provide uncompromised performance by providing developers with a comprehensive set of APIs. Developed for the most demanding use cases, including sub-millisecond SLAs, fabric allows to programmatically fine-tune large and super-large topologies with hundreds to thousands of nodes. Other features included in the GridGain In-Memory Data Fabric, besides the In-Memory Data Grid, are:

| In-Memory HPC                    | Highly scalable distributed framework for parallel High Performance Computing (HPC).   |
|----------------------------------|--|
| In-Memory Streaming              | Massively distributed CEP and Stream Processing system with workflow and windowing support.  |
| In-Memory Hadoop<br>Acceleration | Quick and easy acceleration of existing Hadoop-based systems and products using a dual-mode, high performance in-memory file system. |



## GridGain Foundation Layer

#### What are the common components across all GridGain products?

The GridGain foundation layer is a set of components shared across all functional areas in the GridGain In-Memory Data Fabric. It provides a common set of functionality available to the end user such clustering, high performance distributed messaging, zero-deployment, security, etc.

#### **HYPER CLUSTERING®**

GridGain provides one of the most sophisticated clustering technologies on Java Virtual Machines (JVM) based on its Hyper Clustering<sup>®</sup> technology. The ability to connect and manage a heterogeneous set of computing devices is at the core of GridGain's distributed processing capabilities.

Clustering capabilities are fully exposed to the end user. The developers have full control with the following advanced features:

- > Pluggable cluster topology management and various consistency strategies
- > Pluggable automatic discovery on LAN, WAN, and AWS
- > Pluggable "split-brain" cluster segmentation resolution
- > Pluggable unicast, broadcast, and Actor-based cluster-wide message exchange
- > Pluggable event storage
- > Cluster-aware versioning
- > Support for complex leader election algorithms
- > On-demand and direct deployment
- > Support for virtual clusters and node groupings

#### ZERO DEPLOYMENT

The zero deployment feature means that you don't have to deploy any components individually on the grid – all code together with resources gets deployed automatically. This feature is especially useful during development as it removes lengthy Ant or Maven rebuild routines or copying of ZIP/JAR files. The philosophy is very simple: write your code, hit a run button in the IDE or text editor of your choice and the code will be automatically be deployed on all running grid nodes. Note that you can change existing code as well, in which case old code will be undeployed and new code will be deployed while maintaining proper versioning.

#### ADVANCED SECURITY

The GridGain security components of the foundation layer provide two levels by which security is enforced: cluster topology and client connectivity. When cluster-level security is turned on, unauthenticated nodes are not allowed to join the cluster. When client security is turned on, remote clients will not be able to connect to the grid unless they have been authenticated.

#### SPI ARCHITECTURE AND PNP EXTENSIBILITY

The Service Provider Interface (SPI) architecture is at the core of the GridGain In-Memory Data Fabric; it allows GridGain products to abstract various system level implementations from their common reusable interfaces. Essentially, instead of hard coding every decision about internal implementation of the product, GridGain products instead expose a set of interfaces that define their internal view on their various subsystem. Users then can choose to either use the built-in implementations or roll out their own when they need different functionality.

GridGain products provide SPIs for 14 different subsystems, all of which can be freely customized:

> Cluster discovery

GridGain

- > Cluster communication
- > Deployment
- > Failover
- Load balancing

- > Authentication
- > Task checkpoints
- > Task topology resolution
- > Resource collision resolution
- > Event storage

- > Metrics collection
- > Secure session
- > Swap space
- Indexing

Having the ability to change the implementation of each of these subsystems provides tremendous flexibility to how GridGain products can be used in a real-world environment. GridGain software blends naturally in almost any environment and integrates easily with practically any host eco-system.

#### **REMOTE CONNECTIVITY**

The GridGain In-Memory Data Fabric comes with a number of Remote Client APIs that allow users to remotely connect to the GridGain cluster. Remote Clients come for multiple programming languages including Java, C++, REST and .NET C#. Among many features, the Remote Clients provide a rich set of functionality that can be used without a client runtime being part of the GridGain cluster: run computational tasks, access clustering features, perform affinity-aware routing of tasks, or access the In-Memory Data Grid.

#### ABOUT GRIDGAIN<sup>™</sup>

GridGain, the leading provider of the open source In-Memory Data Fabric, offers the most comprehensive in-memory computing solution to equip the real-time enterprise with a new level of computing power. Enabling high-performance transactions, real-time streaming and ultra-fast analytics in a single, highly scalable data access and processing layer, GridGain enables customers to predict and innovate ahead of market changes. Fortune 500 companies, top government agencies and innovative mobile and web companies use GridGain to achieve unprecedented computing performance and business insights. GridGain is headquartered in Foster City, California. To download the GridGain In-Memory Data Fabric, please visit http://www.gridgain.com/download/.

## GridGain In-Memory Data Fabric

ULTIMATE SPEED AND SCALE

**COMPREHENSIVE AND PROVEN** 

**OPEN AND AFFORDABLE** 

#### Learn more at www.gridgain.com/products

© 2014 GridGain Systems, Inc. All Rights Reserved

GRIDGAIN.COM