
# Cloud-TM

Specific Targeted Research Project (STReP)

Contract no. 257784

# D5.8: Exploitation Plan

**Date of preparation:** 31 July 2013
**Start date of project:** 1 June 2010
**Duration:** 38 Months

## Contributors

Emanuel Bernard, Red Hat
Joao Cachopo, INESC-ID
Daniele Calisi, ALGORITHMICA
Bruno Ciciani, CINI
Fabio Cottefoglie, ALGORITHMICA
Francesca Giannone, ALGORITHMICA
Mark Little, Red Hat
Mircea Markus, Red Hat
Manik Surtani, Red Hat
Francesco Quaglia, CINI
Luis Rodrigues, INESC-ID
Paolo Romano, INESC-ID
Pedro Ruivo, Red Hat
Francesco Zaratti ALGORITHMICA
Vittorio A. Ziparo, ALGORITHMICA

# Contents

# Chapter 1

# Introduction

The Cloud-TM project joined two academic partners (CINI and INESC-ID) and 2 industrial partners (Algorithmica and Red Hat).

The academic partners have a strong background in distributed data management, transactional memory systems, parallel programming, performance modelling and autonomic computing.

Algorithmica is a SME software company, with a portfolio of highly innovative applications and services in the Cloud for both mobile and desktop devices as Software as a Service (SaaS). Red Hat is the world-wide leading company in the open source software arena, and has participated in the project with a team of industry experts in the area of cloud data management and data-oriented middleware.

Although the project's goals, research and development strategies were (in the essential) shared by all partners, the relevance of the various results and the opportunities to exploit them varied and vary among the partners.

For academic partners, the main goals were to advance the state-of-the-art in transactional data management and autonomic computing, to strengthen their position among the research communities, and to build proof-of-concept prototypes to support pedagogical and scientific activities. As detailed in the following sections, these objectives were entirely achieved. More than 60 research papers have been published, several of which in top-tier conferences and journals. Moreover, partners organized and participated in several highly reputed scientific meetings, gave invited talks in several universities and EU projects collaboration meetings, contributed to several EU-centric Working Groups, developed system prototypes, and reinforced their universities' courses with topics related to transactional data management in the cloud.

Algorithmica, on the other hand, will exploit the project's results in several ways. First, Algorithmica will continue to develop the pilots that it developed in the context of the Cloud-TM project, with the goal of commercializing them in the short future. Furthermore, the know-how gathered in the development of data-centric, highly scalable cloud applications with strong consistency requirements will represent a valuable asset for the company, both in the context of its consulting activities as well as for the development of future projects. Finally, Algorithmica is frequently involved in national and international research-oriented collaborations/projects, even in fields beyond cloud

computing (such as robotics, mechanism design, and gaming). These projects are often challenged with issues related to the real-time management of massive data streams. The high scalability, combined with the ease of programming, offered by the Cloud-TM Platform make it a perfect candidate to tackle these challenges.

Finally, Red Hat has great interest in integrating the most promising research results achieved within the Cloud-TM project in its suite of middleware projects for cloud computing. As it shall be detailed in the following, the results of the Cloud-TM project have been, or are being integrated (at the time of writing) in several strategic projects for the JBoss platform, such as Infinispan, Hibernate OGM, Hibernate Search, and OpenShift.

# Chapter 2

# Algorithmica

The main goal of Algorithmica in the Cloud-TM project has been to drive user requirements and to evaluate the Cloud-TM Platform through the development of two pilot applications. To this end, Algorithmica has acted as an end-user of the platform. During the project, Algorithmica acquired an extensive know-how in the development of Cloud-based applications. In particular, it gained specific expertise on topics related to data consistency, data distribution vs replication, response latency and high levels of concurrency. In addition, Algorithmica became an expert user of a number of Red Hat enterprise technologies, including application servers and enterprise messaging systems.

During the course of this project Algorithmica developed the MADMASS framework. MADMASS[1] embodies the lessons learned, the best practices, and the know-how that Algorithmica acquired during the project. MADMASS is fully integrated with the Cloud-TM Platform and with the JBoss middleware. Thus, MADMASS is the building block for Algorithmica's exploitation plans of the Cloud-TM Platform. Specifically, MADMASS is an open source framework that Algorithmica plans to use as a base for its consulting activities in order to simplify the development of complex applications (thus reducing costs). One interesting market, from a consulting perspective, is the one of geo-social or Location-based Mobile Social Networking (LMSN) applications, and the GeoGraph pilot provides a very powerful tool that can be used as a basel framework to develop (and benchmark) a wide range of different geo-social applications. GeoGraph combines the benefits of MADMASS with a set of ready-to-use implemented services that can be found in most geo-social services.

In addition to consulting, Algorithmica develops its own products, such as the AI-Colony game. AI-Colony is a great example of a potentially successful MADMASS based product. Online games are indeed an ever-growing market and, specially HTML5 games, are steadily conquering bigger market shares. Although AI-Colony is still a research prototype, our preliminary experiments with focus groups suggest that there is a potential for the development of a successful business.

[1] Source code available at: `https://github.com/algorithmica/madmass`

## 2.1 MADMASS

The MAssively Distributed Multi Agent System Simulator (MADMASS) is an open source framework for developing rich-client web applications that require scalability, strong consistency, and that can generate very heterogeneous and dynamic workloads. Target applications include, but are not limited to, Multi-Player Online Games and Location-based Mobile Social Networks (or geo-social networks). MADMASS has been developed in the context of the Cloud-TM project and is at the core of the two Cloud-TM pilots: GeoGraph and AI-Colony.

MADMASS is at the core of the Cloud-TM pilots, and as such, it was designed from the ground up for the cloud. It integrates seamlessly with the Cloud-TM Platform and the JBoss enterprise middleware (e.g. TorqueBox [1]). Further, MADMASS has its roots in Artificial Intelligence and Multi-Agent Systems research. These disciplines provide the foundations of the MADMASS programming paradigm, making it a versatile and intuitive framework for developing complex applications that feature scalability and concurrency. The MADMASS project stands on the shoulders of a number of existing open source projects. Namely, it is a Ruby On Rails Engine [2]. Ruby On Rails is a very popular web framework that has been used, for example, to develop Twitter and Github. Finally, MADMASS also supports WebSockets [3, 4] and rich browser GUIs through HTML5 [5].

Algorithmica's exploitation plans for MADMASS are twofold:

1. On the one hand, Algorithmica plans to use MADMASS as a base framework for the systems developed for its customers. MADMASS itself is stable and MADMASS applications (i.e., GeoGraph) have been proved to scale during the evaluation of the Cloud-TM Platform. Moreover, MADMASS builds upon Red Hat's middleware allowing for developing enterprise level applications and integrating with legacy systems through Hibernate OGM. Despite this, the framework has been used to build two research prototypes and will need further refinements and testing to make sure it can be used for applications in production. To this end, Algorithmica will introduce MADMASS functionalities gradually into its systems, having in mind both the reduction of development costs and the iterative refinement of the MADMASS framework through its validation on an increasing number of real applications.

2. On the other hand, being MADMASS an open source tool, Algorithmica also plans to benefit from MADMASS specific consulting activities. Clearly, this will be possible only if the MADMASS framework is adopted by a critical mass of developers. To achieve this, Algorithmica will pursue a strategy based on showcases, which will hopefully attract developers to the framework. The showcases will feature banners advertising the framework (e.g., "MADMASS inside"), screencasts, and demos at open source conferences. Showcases will include both systems developed for our customers and our own products, e.g. AI-Colony (see Section 2.3).

## 2.2 GeoGraph

Apps for smartphones are considered to be the core of the next generation of applications for the Internet [6]. Within the category of mobile apps, LMSN ones, also known as geo-social networks, are gaining increasing market shares [7]. The basic idea behind LMSN is to provide a second generation of Social Networks that can take advantage of the position of the users in order to provide innovative services. The LMSN market is already very rich and features a wide spectrum of products developed both by well known corporations and by independent players (see [8] for an exhaustive list of products). For example, both Google and Microsoft developed their own LMSNs, called Latitude [9] and Vine [10], respectively. Nokia, instead, recently bought a similar product from an independent player, called Plazes [11]. Likewise, there are many independent companies which are developing their own innovative products [12, 13, 14, 15, 16, 17].

Consistently with the worldwide expansion of the LMSN market, Algorithmica is facing an increasing amount of requests for development of LMSN applications. To this end, GeoGraph [18] is a very appropriate tool that, by relying on MADMASS and on a wide number of implemented LMSN services, can support Algorithmica in the development of such applications for its customers. One important feature, both from a technological and from a marketing perspective, is that GeoGraph provides a scalable workload generator that can be used for testing scalability and estimating infrastructure (IaaS/PaaS) costs.

GeoGraph is a research prototype that has been used to evaluate the Cloud-TM Platform. As for MADMASS, Algorithmica plans to gradually introduce GeoGraph into our customers' LMSN products, while incrementally fine-tuning and improving its functionalities for real world applications. For example, Algorithmica is developing a sailing social network for an italian startup named SeeSailing,[2] which features, among other services, real-time GPS tracking of boats during trainings and official regattas. The application will be launched in September and is currently using a relational database to persist position data. Algorithmica is considering with its customers the possibility of introducing GeoGraph in future releases, to improve the scalability and reduce the latency of the overall system. This would be one of the very first showcases of MADMASS.

## 2.3 AI-Colony

Online gaming is nowadays one of the major businesses of the web, especially if considered in terms of revenue and penetration [19]. Online casual and social gaming takes $39\%$ of all $215$ million hours spent on gaming each day in the US and $29\%$ of money. Share of time and money is comparable in Asia, Europe, and emerging markets. $87\%$ of all $145M$ US gamers aged between $10$ and $65$ play on casual websites or social networks. Despite its size, this market is extremely wide and dynamic. For example, the web-gaming company Zynga after only four years had over 232 million active users and was valued between $15 billion and $20 billion.

[2] `http://seesailing.com`

Algorithmica believes that the next step of online gaming is web real-time multiplayer games. To this end, it is interested in developing a new class of games that take inspiration from classical board game and create real-time variants of them. Indeed, one such example is AI-Colony [20]. AI-Colony has been a valuable prototype for evaluating the ease of programming of the Cloud-TM Platform and is an interesting starting point for developing a commercial product. AI-Colony is an example of a MADMASS showcase. To maximize the impact of the showcase to developers, the entire project has been released as open source for non commercial use.[3]

During the development of AI-Colony, in parallel with the Cloud-TM project, we have released the game under the commercial name The Harvestar. The goal of this release was to gather feedback on the game from focus groups and alpha testers. To this end, Algorithmica invested resources outside the Cloud-TM budget to provide the game with an appealing user interface and published it on the Amazon IaaS.[4] The game has been used by a small group of users ($\sim 500$) that can be found at `http://www.facebook.com/theharvestar`.

Although the beta testers have expressed great appreciation for the game, analyzing their comments and their behavior made clear that there is the need to make significant improvements for the product to be competitive in the online gaming market. The following roadmap summarizes the foreseen steps for commercial exploitation of AI-Colony:

1. The software prototype needs several enhancements to be appealing to broad audiences. First of all, there is a need for interactive tutorials. Being a variant of a board game, AI-Colony has a rich and complex set of rules that take time to be understood. Text and video tutorials are not an option, as gamers are usually willing to take action immediately or they will leave the game.

2. The gameplay must be enhanced to include also "solo playing" in addition to multi-player. Moreover, for the sake of viral uptake, there is the need for introducing social features and social distribution channels such as the Facebook App Center.

3. There is the need to explore possible revenue models, in order to make the product sustainable. From our preliminary market analysis it is clear that the free to play model is the dominant strategy in this segment. Within this direction, we plan to explore several possibilities, including free with advertisement and free with premium accounts, although free playing with paid virtual goods, according to state of the art, seems to be the most effective strategy.

4. There is the need to devise a business plan and gather the required funding for executing it. The latter goal, will be pursued though a canonical "startup" path that involves incremental rounds of venture capitalist funding. Once the business plan is in place, we will participate to dedicated events and competitions.

[3] The code is available at `https://github.com/algorithmica/theharvestar`

[4] `http://www.theharvestar.com`

# Chapter 3

# CINI

Upon starting up the Coud-TM project, members of the CINI research team already had experience in the fields of distributed systems, transactional systems, and high performance computing systems. However, the perspective from which the above systems were studied was a classical one, where the focus was mainly posed on conventional transactional systems, such as databases, and conventional computing architectures not relying on virtualization technologies. The research effort that members of the CINI team put in the Cloud-TM project allowed them to greatly improve their expertise in relation to fairly innovative paradigms for transactional data management, such as in-memory transactional systems, as well as virtualization technologies/approaches.

For what concerns in-memory transactional data management, the research done in the context of the Cloud-TM project allowed the team members to definitely improve/expand their knowledge and capabilities on the side of system modeling and optimization. Particularly, the shift towards in-memory transactions led to focus the modeling process on aspects and runtime dynamics that where historically considered less relevant, given the predominance of other aspects (such as per-transaction disk based interactions, which are/can be avoided in the context of in-memory transactional systems). Such a new experience will be (and partially has already been, see [21, 22, 23, 24]) exploited for research on modeling and optimization of general contexts entailing concurrency/parallelism and accesses to shared data placed within the live memory layout of the active threads. On the other hand, the deep study on data replication strategies and protocols, which has been carried out during the development of the Cloud-TM project, allowed the team members to gain knowledge and expertise in the context of distributed data management systems characterized by a very reduced computation-to-communication ratio (given the typical reduced granularity of in-memory transactions). Such an experience will form a background for the research that the team members plan to carry out in the future along the path of performance optimization of generic parallel/distributed applications, possibly entailing fine grain tasks, run on top of large scale platforms.

On the other hand, the effort put into the Cloud-TM project allowed the members of CINI to gain knowledge on the effects (such as performance effects) associated with virtualization technologies, which will constitute the base for addressing complex

performance optimization problems in general parallel/distributed applications run on top of Cloud based systems. In relation to this aspect, some of the CINI team members are currently involved in a National (PRIN) project proposal coping with performance optimization of scientific computing applications run in the cloud, the preparation of which greatly benefited by the knowledge gained thanks to the work carried out in the context of the Cloud-TM.

Finally, concurrent programming and virtualization are becoming core topics also in the context of BS/MS courses at European Universities. In this context, the experience done with the Cloud-TM project can definitely help in the preparation of both basic and advanced courses on computing systems. More in detail, results achieved within Cloud-TM can be exploited by CINI members for both refining design methodologies to be presented in the courses and for setting up specific case studies.

# Chapter 4

# INESC-ID

During the project, INESC-ID greatly reinforced its team expertise on fundamental and practical aspects of distributed data management in cloud environments, transactional memory (both in multicore and distributed settings), self-tuning of large scale systems, machine learning applied to performance forecasting of distributed platforms, control theory and benchmarking. The knowledge and experience amassed in the large codebase developed during the period of the project constitutes an immensely competitive business position that can be exploited by all the involved members. These include, in addition to the various senior researchers that coordinated the research activities, also the following 7 PhD students, which have either recently graduated, or whose graduation is planned in a time window of two years since the time of writing:

- Nuno Carvalho, graduated in December 2011 at the Technical University of Lisbon, PhD thesis on "Dependable Software Transactional Memories," currently employed at Fujitsu Laboratories of Europe, Middlesex, UK, on the design of distributed platforms for Big Data systems;
- Stoyan Garbatov, INESC-ID/Technical University of Lisbon, PhD Thesis on "Exploring Data-Access Patterns for Optimizing Large-Scale Applications,", graduation planned for end 2013;
- Sergio Fernandes, INESC-ID/Technical University of Lisbon, PhD Thesis on "Efficient and Strongly Consistent Software Transactional Memory for Clustered Enterprise Applications," graduation planned for early 2014;
- Sebastiano Peluso, joint PhD by Sapienza Rome University and Technical University of Lisbon, PhD Thesis on "Strong consistent replication in large-scale transactional data grids," graduation planned for June 2014;
- Maria Couceiro, INESC-ID/Technical University of Lisbon, PhD Thesis on "Autonomic Replicated Software Transactional Memory," graduation planned for end 2014;

- Diego Didona, INESC-ID/Technical University of Lisbon, PhD Thesis on "Performance Modeling and Self-Tuning of Distributed Transactional Platforms," graduation planned for end 2014;

- Nuno Diegues, INESC-ID/Technical University of Lisbon, PhD Thesis on "Protocols and Abstractions for Efficient Large-Scale Transactional Systems," graduation planned for end 2015.

These PhD students have inherited the know-how acquired by the research team throughout the project, and are investigating research areas (one of them in an industrial research laboratory) that represent the natural continuation of the results achieved in the context of Cloud-TM.

In addition to these PhD students, there have been also 5 MsC students who have joined the INESC-ID team and performed their thesis on subjects closely related to the Cloud-TM project:

- Pedro Ruivo, graduated in Oct. 2011 at the Technical University of Lisbon with a thesis on the integration of Total Order-based Replication Schemes in Infinispan, currently employed in the Infinispan's team of Red Hat.

- João Fernandes, graduated in Oct. 2011 at the Technical University of Lisbon with a thesis on speculative replication techniques in replicated transactional memory systems, currently employed as R& D engineer at Coriant, Portugal.

- Hugo Pimentel, MsC thesis student at the Technical University of Lisbon, graduation planned for Oct. 2013, with a thesis on the development of L1 Caching Scheme in Infinispan.

- Fabio Perfetti, MsC thesis student at the Sapienza Rome University and INESC-ID, graduation planned for Oct. 2013, with a thesis on the development of the Workload Analyzer and Adaptation Manager of the Cloud-TM Platform.

- Ennio Torre, MsC thesis student at the Sapienza Rome University and INESC-ID, graduation planned for Jan. 2014, with a thesis on elastic scaling methods for Red Hat's OpenShift PaaS platform.

Particularly noteworthy is the case of Pedro Ruivo, who has been recently hired by Red Hat's Infinispan team (see Section 5.5). Pedro represents a valuable link towards Red Hat, and we have already planned to keep on collaborating to finalize the integration of several recent research results (like Autoplacer [25] and GMU [26]) in the official version of Infinispan. Collaborations with Red Hat will continue also thanks to the MsC thesis of Ennio Torre, which is aimed at integrating the Cloud-TM's Autonomic Manager with the Red Hat's PaaS, namely OpenShift.

INESC-ID carried several initiatives in order to disseminate and apply Cloud-TM results and experience. Three tutorials [27, 28, 29] on Distributed Transactional Memories were given during the execution of the project, and one joint workshop was organized with the Euro-TM COST Action with a number of talks about Cloud-TM. At the Technical University of Lisbon a course dedicated to cloud platforms for big data systems is going to be created in the doctoral program. The project's coordinator gave

a number (e.g., [30, 31, 32]) of talks on Cloud-TM, and has already been invited to give two additional talks over the next two months ([33], [34]). Further, Paolo Romano is also chair of the Euro-TM COST Action, a networking project supported by the EU commission that brings together the leading experts in the area of Transactional Memory. Via Euro-TM's activities (workshops, working group meetings, demonstrations to the broad public), INESC-ID researchers will keep on actively disseminating the results of the Cloud-TM project to the scientific community and to the broad mass of programmers. For instance, we plan to demonstrate the Cloud-TM Platform at FOSDEM 2014, in Bruxelles, February 1-2 2014, an international meeting for open source developers that gathers every year 5000+ participants.

Finally, the work carried out in the context of the Cloud-TM project has paved the way for current and future research lines that INESC-ID is already working on or plans to target: energy efficiency in distributed data platforms, exploitation of recently introduced hardware support for transactional memory, design of highly scalable consistency/concurrency control protocols, distributed indexing on transactional key-value stores (to support efficient queries beyond the primary key), self-tuning mechanisms for cloud data grids for applications with real-time constraints. These topics have been object of a recently submitted national proposal (Green-TM, still under evaluation) and will be at the basis of two proposals for next Horizon 2020 ICT call (currently under preparation).

# Chapter 5

# Red Hat

Before discussing the specific exploitation plans of Red Hat regarding the results of the Cloud-TM project, it is worth recalling what is the business model adopted by Red Hat, that is, the so called, Professional Open Source business model.

In this approach, all of Red Hat's products are based on upstream open source communities (e.g., Fedora, JBoss.org, Apache, Eclipse). Red Hat takes these upstream projects and productises them, to create its products. Red Hat then sells subscriptions to these products: a subscription is not the same as a licence. Red Hat subscriptions have no client access licenses. No support incident limits. No unbudgeted upgrade costs. No hidden charges. Instead, your subscription gives you Red Hat's enterprise open source solutions and everything you need to use it effectively. Red Hat products are provided on a per-instance or per-installation subscription basis, which gives you access to all subscription benefits during your subscription term. These benefits include access to Red Hat's:

- Software
- Updates
- Upgrades
- Technical support
- Security fixes
- Open source assurance program
- Proven, enterprise-ready solutions

Red Hat's main goal in the Cloud-TM project is to provide a vehicle for exploiting Cloud-TM research via several open source projects and the products that Red Hat supports enterprise customers on, which are based on said open source projects.

In the following we detail how the Cloud-TM's project results will be (and, in some cases, have already been) integrated in 4 strategic Red Hat's projects—leading to 3 Red Hat products—in the area of cloud/distributed data management: Inifinispan

(see Section 5.1), Hibernate OGM (Section 5.2), Hibernate Search (Section 5.3), and OpenShift (Section 5.4).

It is worth to highlight that, as a part of integrating Cloud-TM research into these projects, Red Hat has hired Pedro Ruivo, a researcher from INESC-ID who was working on Cloud-TM. This has not only helped Infinispan as a project but also increases our capacity to exploit Cloud-TM research at a faster rate (Section 5.5).

## 5.1 Infinispan

Infinispan is an open source data grid platform that can be used as a distributed, in-memory cache to front a database or NoSQL datastore, or can also be used as a persistent, durable, distributed key/value NoSQL-like database in its own right. It is transactional and usable both as an embedded library in a Java application, or as a standalone cluster of data storage nodes, accessed remotely over a socket. In the latter case, it is usable not only by Java applications, but by any platform for which client libraries exist, including Ruby, Python, C++, and .NET.

Within the lifetime of Cloud-TM, Red Hat has moved the Infinispan project from a community effort (no commercial support) into a product. This product, the JBoss Data Grid[1] has proven to be extremely successful and some of this is because of the changes that have come directly from the Cloud-TM work. Although Red Hat does not publicly publish its list of customers, the following are public: The Chicago Board of Options Exchange (see `http://goo.gl/kSJSE8/`) and Cisco.

### 5.1.1 Exploitation of Cloud-TM results

Infinispan has benefited from Cloud-TM in a variety of ways, which we detail in the following.

#### Performance and quality

The constant performance benchmarks performed by Cloud-TM has uncovered performance issues in Infinispan that have since been addressed, leading to a higher quality codebase. Repeated stress testing in a variety of access patterns by Algorithmica have contributed significantly to this.

#### Total order multicast/broadcast

Infinispan has benefited from the total order protocols developed as a part of Cloud-TM, and integrated since Infinispan 5.3.0. The total order based protocol is a lock-free commit protocol that relies on the concept of totally ordered delivery of messages which, informally, implies that each node that delivers a set of messages, delivers them in the same order. This protocol comes with two main advantages: (1) transactions can be committed in one phase, as they are delivered in the same order by the nodes that receive them, and (2) it avoids distributed deadlocks. The weaknesses of this approach

[1] `http://www.redhat.com/promo/infspn2jdg/`

are the fact that its implementation relies on a single thread per node which receives the transaction and its modification from other nodes in the cluster, and the slightly higher number of messages exchanged by JGroups. Thus, this protocol delivers best performance in scenarios of high contention, in which it can benefit from the single-phase commit and the thread that delivers/applies transactions is not the bottleneck. Currently, the Total Order based protocol is available only in transactional caches for replicated and distributed modes and it is available since Infinispan 5.3.0 Alpha1.

The results of some benchmarks comparing the total order based implementation with the locking based implementation in two different scenarios are reported in Figure 5.1, for two different scenarios:

- Contention Scenario: 1000 keys in a shared pool, 8 threads per node and 5 writes (on average) per transaction;
- No Contention Scenario: 1000 keys in a thread private pool, 8 threads per node and 5 writes (on average) per transaction.

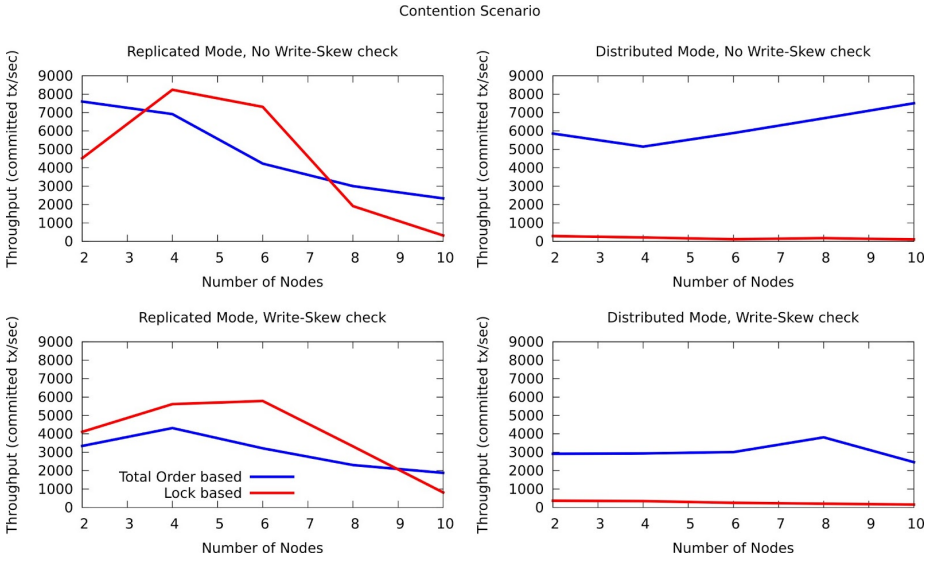


Figure 5.1: Benchmark results comparing the performance of Total Order based protocols vs protocols relying on classic Two-Phase Commit schemes.

### Extended Statistics

Additional statistics on Infinispan have been added by the Cloud-TM project, which enables GMU and adaptive data placement to work. The additional statistics captured include average transaction execution time, message sizes, number of operations per transaction, as well as top key streaming. The latter exposes which keys are the most accessed, locked, etc.

These extended statistics are being integrated into Infinispan in time for Infinispan 6.0.0.

#### GMU

GMU introduces SERIALIZABLE isolation level for data stored in Infinispan, to allow access to data as a snapshot in time.

This implementation is based on Cloud-TM research and is being integrated into Infinispan in time for Infinispan 6.0.0.

#### Adaptive data placement

Infinispan selects nodes as data owners and backups by applying a hash function (currently MurmurHash3) on the key, and then using a consistent hash wheel to determine which nodes serve as a data owner and backups. This technique works well in terms of a low-cost deterministic algorithm, however it comes with a cost of not being able to force locality of data to a process that actually requires that data.

Cloud-TM has been working on an adaptive data placement technique, named AUTOPLACER, that will allow data location to be based on additional intelligence, on where data is most frequently used. This allows for greatly reducing network roundtrips for frequently accessed data.

This work has not been integrated into Infinispan yet, and is due for integration over the course of this calendar year.

## 5.2 Hibernate OGM

Hibernate OGM is an open source project experimenting on the notion of offering Java Persistence API support for non relational databases such as Infinispan.

Hibernate OGM is not part of a current commercial offering of Red Hat's middleware portfolio but will likely be in the near future. Further development and productization is required to include Hibernate OGM in the offering.

Yet, concepts, ideas, and experience have been reused by the JBoss Data Grid platform for its upcoming implementation, in particular to offer higher level abstractions on top of the key/value store.

### 5.2.1 Exploitation of Cloud-TM results

Hibernate OGM has been heavily influenced and thus benefited from Cloud-TM.

Performance issues have been identified while testing the performance and scalability of the Cloud-TM Platform. Subsequent improvements in Hibernate OGM, Hibernate ORM, and Infinispan code has followed to address these problems. In particular, work has been done around object allocation and the data structure holding the objects attached to a given session.

Likewise, work on the JP-QL query parser and query conversion from the object world to the non-relational datastore has been influenced by the feedback gathered

via Cloud-TM. An initial JP-QL to Lucene query only conversion has been deemed inappropriate and a readjustment to a backend-agnostic approach has been taken.

Future exploitation plans include:

- The integration of techniques recently introduced in the Fenix Framework and aimed to enhance data locality by controlling co-location of objects in a distributed data grid (such as Infinispan).
- The reduction of conflicts by exploiting a-priori knowledge of commutativity of operations on shared objects.
- To continue the effort in JP-QL to Lucene query converter to increase the query complexity supported in OGM: Cloud-TM experiments gave us the most important use cases that we addressed but also revealed important use cases that we have not been able to address within the timeframe of the Cloud-TM project. In particular, aggregation operations.

## 5.3 Hibernate Search

Hibernate Search offers full text and geo-localized queries to data objects. Ease of use and transparent clustering are key aspects of the system.

Hibernate Search is offered as a full-text search engine inside Red Hat's JBoss Web Framework Kit platform. Features such as geo-localized queries, which have emanated from Cloud-TM, have made it through the productized version of Hibernate Search.

The project is also at the core of JBoss Data Grid's search engine, which is due in the product for this calendar year. Various Cloud-TM efforts have made this feature possible:

- The decoupling between Hibernate Search and Hibernate ORM: Given that Cloud-TM had various object sources (Hibernate OGM, Infinispan, Fenix Framework), we have decoupled the Hibernate Search engine and the indexing capabilities from Hibernate ORM. This has been key to reuse Hibernate Search from within JBoss Data Grid.
- Hibernate Search makes use of Apache Lucene for its indexing capability. We worked on the ability to store these indexes on the data grid itself, offering instant propagation and elasticity to the solution.
- Various work around master / slaves workload distribution, transparent master reelection, and multi-masters approaches have been designed for Cloud-TM distributed capabilities and are being reused as part of JBoss Data Grid.

### 5.3.1 Exploitation of Cloud-TM results

Cloud-TM led most of the recent prominent features of Hibernate Search. In particular, offering flexible backends, geo-localized queries, and simplification of configurations in cluster, particularly in Infinispan. These features are now used both by community members as well as by Red Hat customers (or will soon be).

## 5.4 OpenShift

For the past 24 months Red Hat has been working on an open source IaaS (DeltaCloud) and PaaS (using JBoss technologies). In May 2011 we released OpenShift, which is a series of PaaS implementations. Express is free for developers and offers limited configuration and capabilities. Flex is more configurable and offers a wider range of capabilities. Both of these PaaS implementations will have Infinispan running on them soon and we are clearly indicating that this is also part of the Cloud-TM efforts. OpenShift is being adopted by many of Red Hat and JBoss developers and customers, which is a good way to get the Cloud-TM message out.

### 5.4.1 Exploitation of Cloud-TM results

CINI and INESC-ID have been working on integrating QoS-based elastic scaling mechanisms using innovative technology developed in the context of the Cloud-TM project. The work could not be brought to completion due to resource constraints and to the fact that in that period the OpenShift platform was undergoing a major redesign (which ended up hindering development). However, this work allowed to confirm the viability of integrating the Cloud-TM's middleware (either in its entirety, or its individual modules) in OpenShift, via the development of custom cartridges whose automatic provisioning and tuning could be integrated with the existing OpenShift's brokering mechanism. A master thesis in computer engineering is planned to start on September 2013, co-supervised by Prof. Francesco Quaglia, CINI/Sapienza, and by Prof. Paolo Romano, INESC-ID/IST, and will aim at pursuing this additional dissemination and exploitation opportunity for the project's results.

## 5.5 Employment of personnel from INESC-ID

Over the past 3 years, Red Hat has worked closely with all of the Cloud-TM partners, but perhaps more closely with INESC-ID. This has been a bi-directional flow of information and development. The Infinispan project has grown over that period of time, with contributors from all of our partners and especially INESC-ID. This has been greatly advantageous for all, and, thus, to facilitate knowledge transfer to and from Red Hat, we decided that it was beneficial to hire Pedro Ruivo from INESC-ID as a full time Red Hat employee and member of the Infinispan team. Pedro has continued to work on Cloud-TM and has helped to make further changes to Infinispan and other JBoss technologies.

# Bibliography


[1] JBoss, "Torquebox." torquebox.org, 2012.

[2] D. H. Hansson *et al.*, "Ruby on rails," *Website. Projektseite: http://www. rubyonrails. org*, 2009.

[3] W. W. W. C. (W3C), "W3C the websocket api." http://www.w3.org/TR/websockets/. [Online; accessed 4-May-2012].

[4] I. E. T. F. (IETF), "IETF the websocket protocol." http://tools.ietf.org/html/rfc6455. [Online; accessed 4-May-2012].

[5] M. Pilgrim, *HTML5: Up and running*. O'Reilly, 2010.

[6] WIRED Magazine, "The web is dead. long live the internet." http://www.wired.com/magazine/2010/08/ff_webrip/all/1, 2010.

[7] Abi Research, "Location-based mobile social networking." http://www.abiresearch.com/press/1204-Locationbased+Mobile+Social+Networking+Will+Generate+Global+Revenues+of+$3.3+Billion+by+2013, 2008.

[8] Claudio Schapsis, "Location based social networks links." bdnooz.com/lbsn-location-based-socialnetworking-links.

[9] Google Inc., "Latitude." http://www.google.com/intl/en_us/latitude/intro.html.

[10] Microsoft Corporation, "Vine." http://www.vine.net/default.aspx.

[11] Plazes AG, "Plazes." http://plazes.com.

[12] Pelago Inc, "Whrrl." http://whrrl.com.

[13] BrightKite Inc, "Brightkite." http://brightkite.com.

[14] Loopt Inc, "Loopt." www.loopt.com.

[15] GeoSolutions, B.V., "Gypsii." http://www.gypsii.com/.

[16] Ipoki Technologies S.L, "Ipoki." www.ipoki.com.

[17] Bliin B.V., "Bilin." bliin.com.

[18] Vittorio A. Ziparo, Fabio Cottefoglie, Daniele Calisi,Marco Zaratti, Francesca Giannone, Paolo Romano, "D4.3 companion document for the final prototype of pilot application 1," 2013.

[19] N. BV, "Casual & social games trend report," 2012.

[20] Vittorio A. Ziparo, Fabio Cottefoglie, Marco Zaratti, Francesca Giannone, Paolo Romano, "D4.2 companion document for the prototype of pilot application 2," 2013.

[21] A. Porfirio, A. Pellegrini, P. di Sanzo, and F. Quaglia, "Transparent support for partial rollback in software transactional memories," in *Euro-Par*, pp. 583–594, 2013.

[22] D. Rughetti, P. di Sanzo, B. Ciciani, and F. Quaglia, "Machine learning-based self-adjusting concurrency in software transactional memory systems," in *MASCOTS*, pp. 278–285, 2012.

[23] A. Pellegrini, R. Vitali, S. Peluso, and F. Quaglia, "Transparent and efficient shared-state management for optimistic simulations on multi-core machines," in *MASCOTS*, pp. 134–141, 2012.

[24] P. D. Sanzo, F. D. Re, D. Rughetti, B. Ciciani, and F. Quaglia, "Regulating concurrency in software transactional memory: An effective model-based approach," in *SASO*, 2013.

[25] D. Didona, P. Romano, S. Peluso, and F. Quaglia, "Transactional auto scaler: Elastic scaling of in-memory transactional data grids," in *Proc. ICAC 2012*, 2012.

[26] S. Peluso, P. Ruivo, P. Romano, F. Quaglia, and L. Rodrigues, "When scalability meets consistency: Genuine multiversion update-serializable partial data replication," in *ICDCS*, pp. 455–465, 2012.

[27] P. Romano and L. Rodrigues, "Tutorial on distributed transactional memories," 2010.

[28] M. Couceiro, P. Romano, and L. Rodrigues, "Tutorial on distributed transactional memories," 2011.

[29] L. Rodrigues, "Tutorial on distributed transactional memories," in *Doctoral School on Hot Topics In Distributed Computing (HTDC)*, 2011.

[30] P. Romano, "Sla management in elastic transactional data grids," in *Internet of Services Collaboration Meeting, Bruxelles, Belgium, 16 October 2012*, 2013.

[31] P. Romano, "Self-optimizing transactional data grids for elastic cloud environments (keynote speech)," in *CloudViews 2011, 4 November 2011, Porto, Portugal*, 2013.

[32] P. Romano, "Autonomic mechanisms for transactional replication in elastic cloud environments (keynote speech)," in *2nd Workshops on Software Services (WOSS), 4 June 2011, Timisoara, Romania*, 2013.

[33] P. Romano, "Cloud-TM Project (Keynote speech)," in *2nd Workshop on Storage, INRIA-Rennes, November 20*, 2013.

[34] P. Romano, "Cloud-TM Project (Keynote speech)," in *Coherence and scalability in the cloud, Braga, Portugal, September 30 2013*, 2013.