Grant Agreement No.: 610764
Instrument: Collaborative Project
Call Identifier: FP7-ICT-2013-10

PANACEA

Proactive Autonomic Management of Cloud Resources

PANACEA Final Report

Version: v.1

Grant Agreement number	610764
Project acronym:	PANACEA
Project title:	Proactive Autonomic Management of Cloud Resources
Funding Scheme:	FP7-ICT-2013-10
Date of latest version of Annex I against which the assessment will be made:	2013-07-08
Periodic report:	Final report
Period covered:	From: M12 To: M24
Name, title and organisation of the scientific representative of the project's coordinator:	Dr Olivier Brun, CNRS
Tel:	335 613 369 18
Fax:	335 613 369 69
Email:	brun@laas.fr
Project website:	http://panacea-cloud.eu

Abstract	This deliverable contains a publishable summary of the PANACEA project, highlighting the main S&T results of the project. It also includes a description of our plan for the dissemination and exploitation of the main project outcomes.
Keywords	Final report, project management, main S&T outcomes.

Document Revision History

Version	Date	Description of change	List of contributor(s)
V1.0	26.04.2016	Initial draft released to the consortium	Olivier Brun (CNRS)
V2.0	04.05.2016	Description of the main S&T results by individual partners	All partners
V3.0	04.05.2016	Final version	Olivier Brun (CNRS)

Disclaimer

Copyright notice

Acknowledgment (if needed)

*R: report, P: prototype, D: demonstrator, O: other

Project co-funded by the European Commission in the 7th Framework Programme (2007-2013)		
Nature of the deliverable:	R*	
Dissemination Level		
PU	Public	**X**
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to bodies determined by the PANACEA project	
CO	Confidential to PANACEA project and Commission Services	

EXECUTIVE SUMMARY

This deliverable is the final report of the PANACEA project. It provides a publishable summary of the main project outcomes, as well as a description of our plan for the dissemination and exploitation of the main project outcomes.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABBREVIATIONS

CA	Consortium Agreement
CFS	Certificate on the Financial Statement (formerly referred to as Audit Certificate)
CFU	Cost Follow-up table
DOW	Description of Work (Technical Annex or Annex I to the Grant Agreement)
EC	European Commission
EC-GA	Grant Agreement
EU	European Union
IPR	Intellectual Property Rights
MB	Management Board
Partner	Project Partner Institution, also referred to as Beneficiary and Party of the Consortium Agreement
PO	Project Officer (EC supervisor of the contract)
PM	Project Month
PM	Project Manager
RB	Research Board
RTD	Research and Technology Development
WP	Work Package
WPL	Work Package Leader

1 FINAL PUBLISHABLE SUMMARY

1.1 Executive Summary

Modern cloud services have reached a level of complexity where the human effort required to get them up and running at expected performance levels is becoming prohibitively expensive. Yet, even small degradations in the performance and availability of cloud services can have a considerable business impact in terms of damage to brand reputation, lost revenue and reduced productivity. The goal of PANACEA was to significantly reduce the cost of managing services deployed in the cloud, while at the same time improving their performance and availability, by enabling cloud services to manage themselves with minimal direct human intervention and to react autonomously and dynamically to various inevitable anomalies. To this end, PANACEA has developed several innovative solutions:

- **Proactive Service Management Using Machine Learning**: By the time a service failure or QoS violation is detected, it is too late and end-users are already experiencing service downtime or decreased QoS. A proactive approach is therefore required, and the proactive service management solution developed by PANACEA improve service availability and QoS by predicting anomalies such as software failures and response time violations, and then reconfiguring the service before the predicted anomaly occurs (e.g., by transferring part of the service to a stand-by VM).
- **Cloud Management Platform for Autonomic Services**: The cloud management solution of PANACEA provides the effector (self-configuration) and sensor (self-awareness) mechanisms necessary for autonomic management of services, which enables unattended services to request or provide information about individual VMs or the service itself and request reconfiguration operations, like migrating to another host or attaching a new device. In particular, it enables services to autonomously grow and shrink resources allocated to them as and when needed.
- **Pervasive Cloud Monitoring**: The proposed monitoring solution *self-optimizes by adaptively prioritizing which nodes to monitor in the cloud*, thereby achieving low overhead while providing timely delivery of fine-grained monitoring data, and enabling an accurate representation of the state of the cloud, which is critical for good management and allocation decisions.
- **Online QoS-driven Task Allocation**: The online QoS-driven task allocation system uses machine learning techniques and online measurements to dispatch jobs to available servers in order to improve user experience by maintaining or improving QoS, such as response time, in the face of overload conditions and changes in demand and the infrastructure.
- **Routing Overlay**: In order to improve the Internet communication paths between clouds, we have developed *a self-healing and self-optimizing overlay network* that quickly detects path failures and QoS degradations, and adapts the overlay paths in response to failures and congestion

The main expected impact of PANACEA is improved availability and QoS of cloud services, as well as reduced OPEX for cloud providers since autonomic services will require less supervision and manual reconfiguration and intervention. The PANACEA technologies will

also have a positive impact on the service providers and their end-users, both of which will benefit from higher availability and improved QoS.

1.2 PANACEA Context and Objectives

Cloud computing may be the most innovative technology development in decades, and every day more and more companies are changing their overall information technology strategies to embrace cloud computing [6][7]. However, availability and reliability are major concerns for most users [8], and many large companies are still reluctant to deploy services that require continuous operation over time in the cloud. As a matter of fact, even small degradations in the performance and availability of these services can have a considerable business impact in terms of damage to brand reputation, lost revenue and reduced productivity [1][2]. Present cloud computing platforms do not support mechanisms for recovering from many inevitable anomalies. Among these anomalies, software failures are long known to be a dominant source of unplanned system outages [5,6]. For long running applications, these software failures are usually related to the so-called software aging phenomenon [7], that is the accumulation of errors (e.g., memory bloats or leaks, unterminated threads, etc.) that causes a gradual performance degradation and, eventually, a service hang or crash.

Autonomic computing is emerging as a significant new approach to the design of computer services. The main motivation for Autonomic Computing is the high level of complexity reached by modern computer services, where the human effort required to get them up and running at expected performance levels is becoming prohibitively expensive. Inspired by the autonomic nervous system of the human body, Autonomic Computing aims at enabling computer systems to manage themselves with minimal direct human intervention, while at the same time improving their performance and availability.

The PANACEA project started on the 1st of October, 2013 and lasted 30 months. Its primary goal was to significantly increase the availability and performance of services deployed in the cloud. To this end, we have designed and developed innovative solutions for a proactive and autonomic management scheme of cloud services. PANACEA-enabled services are agile services that are able to autonomously and dynamically

- grow and shrink resources allocated to them as and when needed to meet demand and to draw those resources from the most optimal location,
- dispatch incoming jobs to the best available resources in order to maintain and improve the response time of jobs in response to changing workload conditions,
- detect and quickly recover from path outages or performance degradations of inter-cloud IP routes, always routing flows over the most optimal routes,
- heal themselves by predicting performance degradations or service crash due to the accumulation of errors (e.g., memory leaks) and proactively reconfiguring themselves.

In order to provide these new capabilities to cloud services, the following innovations have been developed in PANACEA:

- **Proactive Service Management Using Machine Learning**: By the time a service failure or QoS violation is detected, it is too late and end-users are already experiencing service downtime or decreased QoS. A proactive approach is therefore required, and the proactive service management solution developed by PANACEA improve service availability and QoS by predicting anomalies such as software failures

and response time violations, and then reconfiguring the service before the predicted anomaly occurs (e.g., by rejuvenating a VM or by transferring part of the service to a stand-by VM).

- **Cloud Management Platform for Autonomic Services**: The cloud management solution of PANACEA provides the effector (self-configuration) and sensor (self-awareness) mechanisms necessary for autonomic management of services, which enables unattended services to request or provide information about individual VMs or the service itself and request reconfiguration operations, like migrating to another host or attaching a new device. In particular, it enables services to autonomously grow and shrink resources allocated to them as and when needed.
- **Pervasive Cloud Monitoring**: The proposed monitoring solution *self-optimizes by adaptively prioritizing which nodes to monitor in the cloud*, thereby achieving low overhead while providing timely delivery of fine-grained monitoring data, and enabling an accurate representation of the state of the cloud, which is critical for good management and allocation decisions.
- **Online QoS-driven Task Allocation**: The online QoS-driven task allocation system uses machine learning techniques and online measurements to dispatch jobs to the most optimal servers in order to improve user experience by maintaining or improving QoS, such as response time, in the face of overload conditions and changes in demand and the infrastructure.
- **Routing Overlay**: In order to improve the Internet communication paths between clouds, we have developed *a self-healing and self-optimizing overlay network* that quickly detects path failures and QoS degradations, and adapts the overlay paths in response to failures and congestion

The main expected impact of these new capabilities is improved availability and QoS of cloud services. This will translate into the provisioning of better services by the cloud providers, thus improving revenue, brand name, company visibility, reputation, and competitiveness in the market. In addition, the provider will also see a reduction in its OPEX since autonomic services will require less supervision and manual reconfiguration and intervention. The PANACEA technologies will also have a positive impact on the service providers and their end-users, both of which will benefit from higher availability and improved QoS.

1.3 Description of the main S&T results

1.3.1 Machine Learning Framework for Proactive Self-* (Healing, Optimizing and Configuring)

By relying on the Intra/Inter Autonomic Cloud Management (ACM) System based on Machine Learning (ML) for proactive cloud resources management, cloud providers and application developers can deploy their application in a highly available fashion, with great ease. In particular, the application can be deployed as well on a hybrid cloud infrastructure, where cloud regions (namely, homogeneous sets of virtual machines hosted by a given cloud provider in a given location) can be several and even geographically distributed. In our ACM System with the Intra/Inter ACM, we rely on a geographically *distributed hybrid clouds*. Each location will be referred to as *a cloud region*. The proposed solution enables to enforce high availability and scalability of the deployed application. In fact, scalability can be controlled at

different levels. *Intra-ACM* controllers can decide upon the number of VMs, which are currently active. By controlling at runtime the execution dynamics of VMs into a cloud region (specifically, by controlling the Mean Time to Failure), the controller is able to make a prediction (based on *Machine Learning (ML) prediction models*) on the number of VMs that are necessary so as to keep the rejuvenation rate below a certain threshold. Considering that the rejuvenation rate is directly proportional to the current workload seen by the system, the controller can decide whether new VMs are required to fulfil the requests coming to the region. In this way, the *Intra-ACM* controller is able to proactively *scale up/down* the number of VMs composing the region i.e. - *elasticity property* in a given region. *Inter-ACM* can connect multiple cloud regions together, even at runtime. This gives the user of the ACM System an additional level of *scalability.* In fact, rather than adding single VMs to the system, entire regions, composed of any number of VMs, could be connected. This allows, for example, to inject a higher amount of computing power, which could be possibly required only during certain operating periods of the system. Overall, our solution gives great freedom in the composition of the system, and enforces autonomic and proactive scalability properties, with minor (or null) manual intervention.

The overall Intra ACM system is based on the *ML Framework* initially presented in Deliverables (3.1 , 3.2 and 3.3) and is shown in Figure 1. Basic software modules at a Node level and a Local Controller are shown in Figure 2. The Intra-ACM architecture is reported in Figure 3. In its most general form, it includes a VM acting as a controller (VMC) and k couples of VMs (slave VMs) acting as (replicated) servers. The slave VMs of a couple c_x (with $x \in \{0, \dots, k-1\}$) are named $VM1_x$ and $VM2_x$, respectively. VMC and the slave VMs communicate via message exchange**.** Once the selected *Remaining Time to Failure* (RTTF/RTTC) from the prediction model has been fed in PU, Intra ACM can perform the *on-line control loop*. VMC predicts when a given VM should be proactively rejuvenated and all actions for controlling the states of the pair VMs are executed by the Dynamic Reconfiguration Flow Diagram, as shown in Figure 4. In Figure 5, we show some results related to a time window extracted from the whole experiment. We report various measured features, namely number of active threads, free memory, used swap memory, and (total) CPU usage. The average response time is kept below a selected threshold <3 seconds (max 4.5). To complete the training, and to show the accuracy ACM System, we measured the number of *false negatives*.

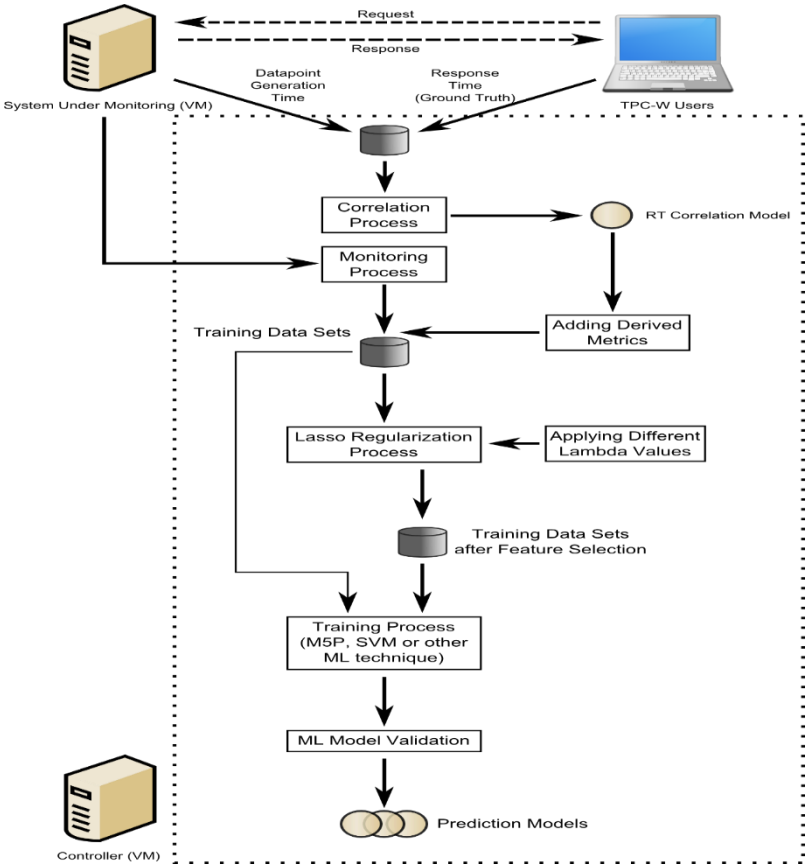

Figure 1 Machine Learning Framework.

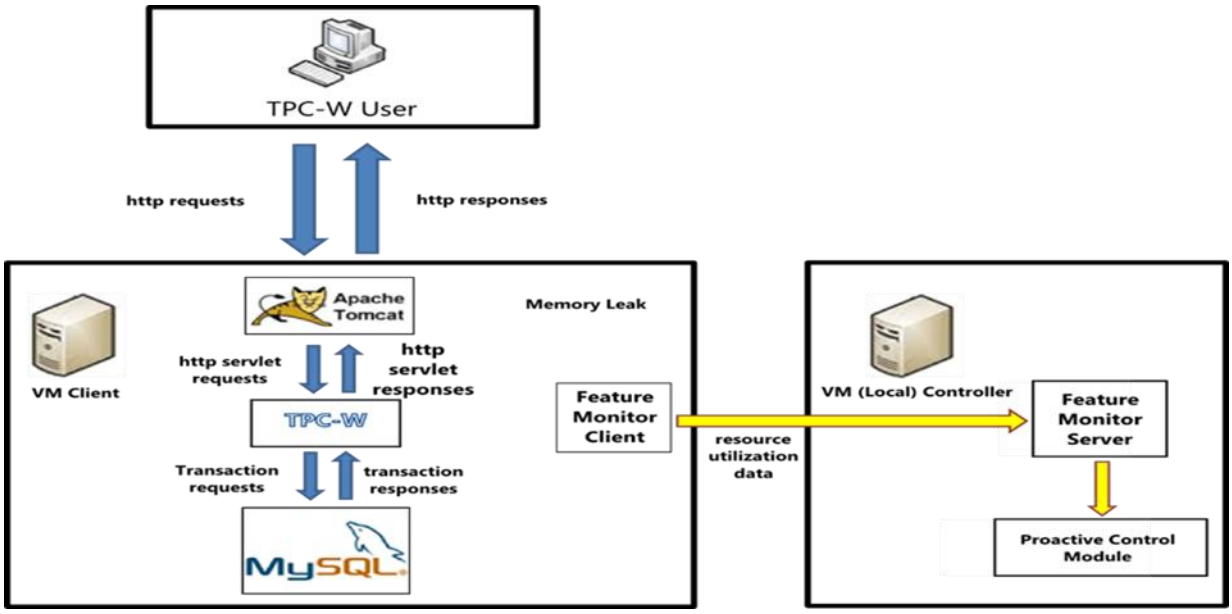

Figure 2 Software Configuration at node level and Local Controller.

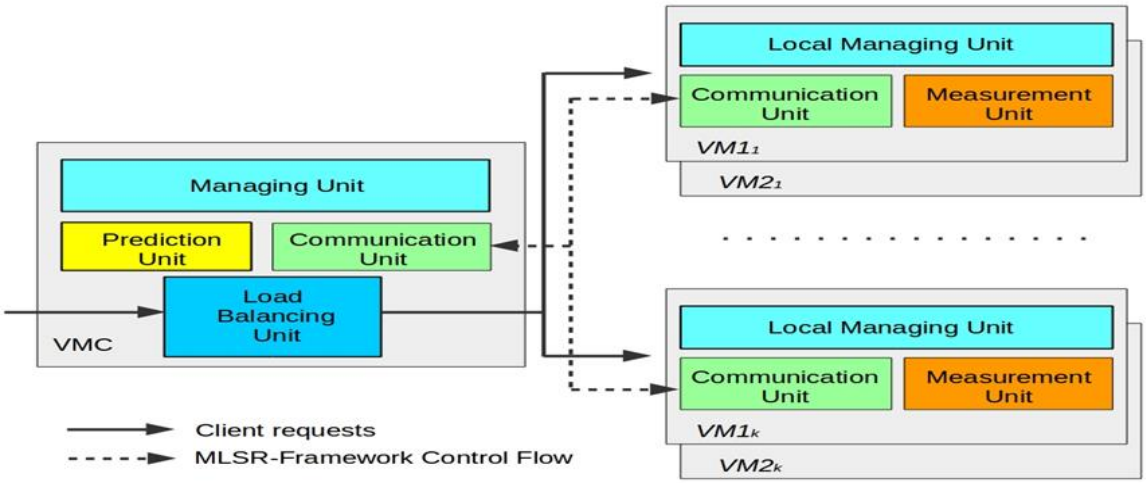

Figure 3 Intra ACM Architecture.

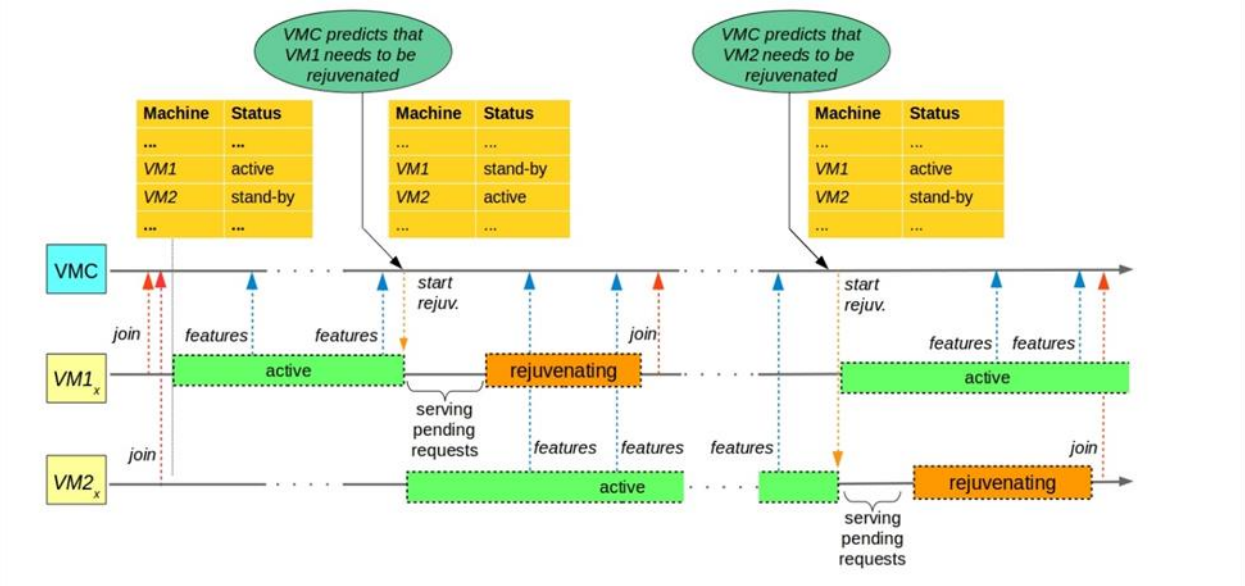

Figure 4 Dynamic Reconfiguration Flow Diagram.

In order to count the number of false negatives, we modified VMC in order to check if, based on values of features received by the active VM, the failure condition has been met. In Figure xxx6, we report the percentage of *false negatives* we measured, for *Threshold* (T) equal to 300, 420 and 540 seconds, respectively, and related to ML prediction models generated using M5P, Lasso as a predictor and Linear Regression. Results show that the most-effective ML algorithm is M5P, providing a lower percentage *of false negatives* with respect to both Lasso and Linear Regression. Nevertheless, we can see that the higher the threshold T, the lower the number of false negatives i.e. 0.16 %. This result contributes for growing applications availability and reducing the outages per year considerably.

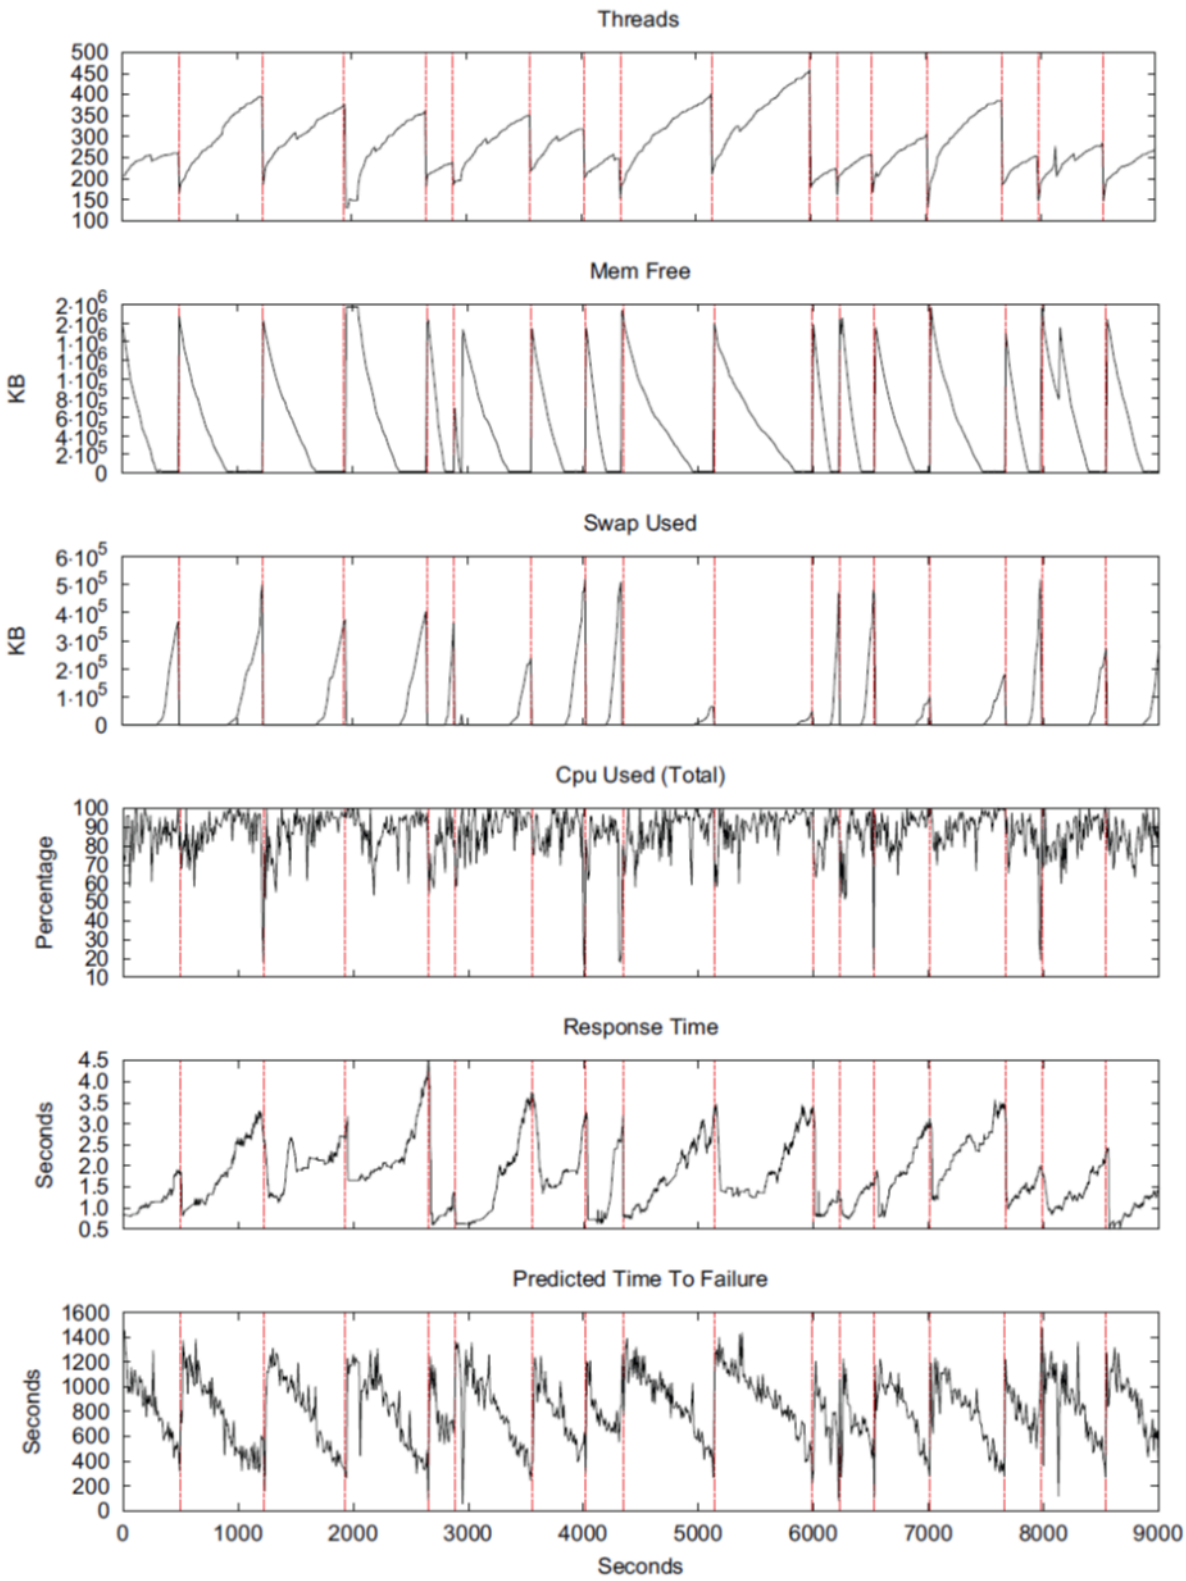

Figure 5 System features, response time and predicted RTTF for the scenario with 2 VMs and Lasso (with reduced parameters) as a predictor.

Results we discussed above show that can improve some system properties, including ones

belonging to the category of *self*(healing, optimizing and configuring)* and significantly *increase availability* of cloud applications. The impact of threshold T is presented in Figure 6.

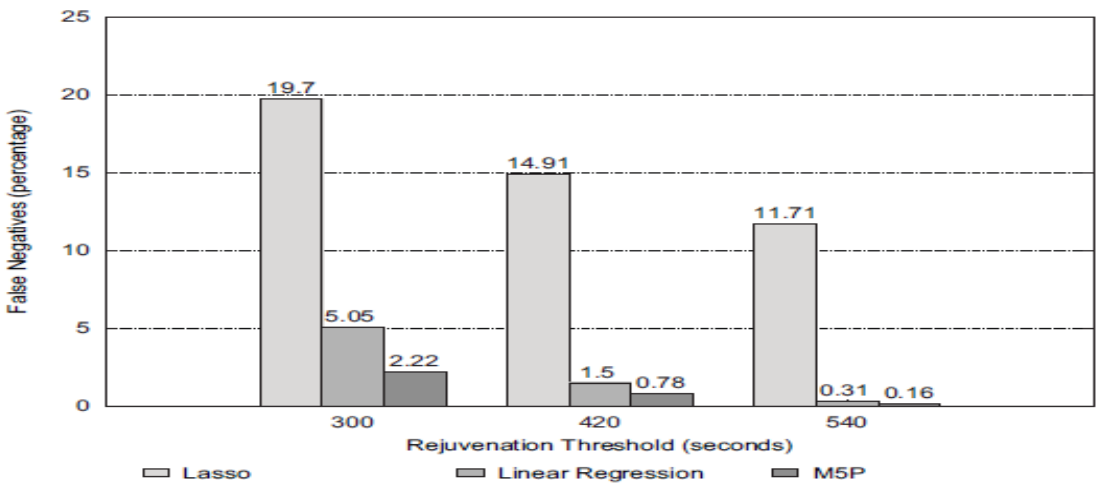

Figure 6 False negatives in the case of memory leaks and unterminated threads.

Inter Autonomic Management

As mentioned, the goal of the proposed framework is to maximize the availability of a given application hosted on a geographically-distributed (hybrid) cloud infrastructure, and at the same time to keep the *response time* as seen by remote clients of the application under a given threshold. The available cloud resources can be distributed over the services offered by different cloud providers, and/or over services offered by the same cloud provider in different geographical locations, and/or over private cloud infrastructures. Each set of VMs, namely VMs hosted in a given region by one single cloud provider (or private infrastructure) is referred to as a *cloud region*. Each cloud region hosts at least the following set of VMs: Controller, Load Balancer and Computing Nodes. As well, the Private Cloud has a leader that orchestrates the work of all region controllers. The System Architecture for scalable distributed clouds based on ML is shown in Figure 7. This scalability is guaranteed by design due to the nature of the protocol diagram, which is reported in Figure 4. In fact, new VMs can be added asynchronously, in any order, without the risk for deadlock or starvation. The results obtained in (Deliverable 3.3 – 2015-07-24.docx,) demonstrate that Inter Autonomic Cloud Manager (Inter ACM) is able to properly coordinate the behaviour of multiple regions. In particular, Inter ACM is able to promptly detect the state of entire regions, and its internal metrics quickly converge to the actual values, even in case of strong variations of the system (an entire region is removed). This allows for proactive scalability of Inter ACM on a larger scale. In fact, this configuration (exactly using Inter ACM) allows to proactively add entire regions of VMs (each region can be composed of any number of VMs) and promptly respond to the execution dynamics of the application. As well, the ML Framework is capable of building accurate ML prediction models for supporting the Intra/Inter ACM. The ML prediction model is created for Ireland Amazon Region and produced by *ML framework* (defined in Figure 1.) The ML prediction, made by M5P learning algorithm, has minimum prediction error(s), smaller training and validation times in comparison with other ML algorithms (Linear Regression, SVM, SVM2 and Lasso) and is shown in Figure 9**.**

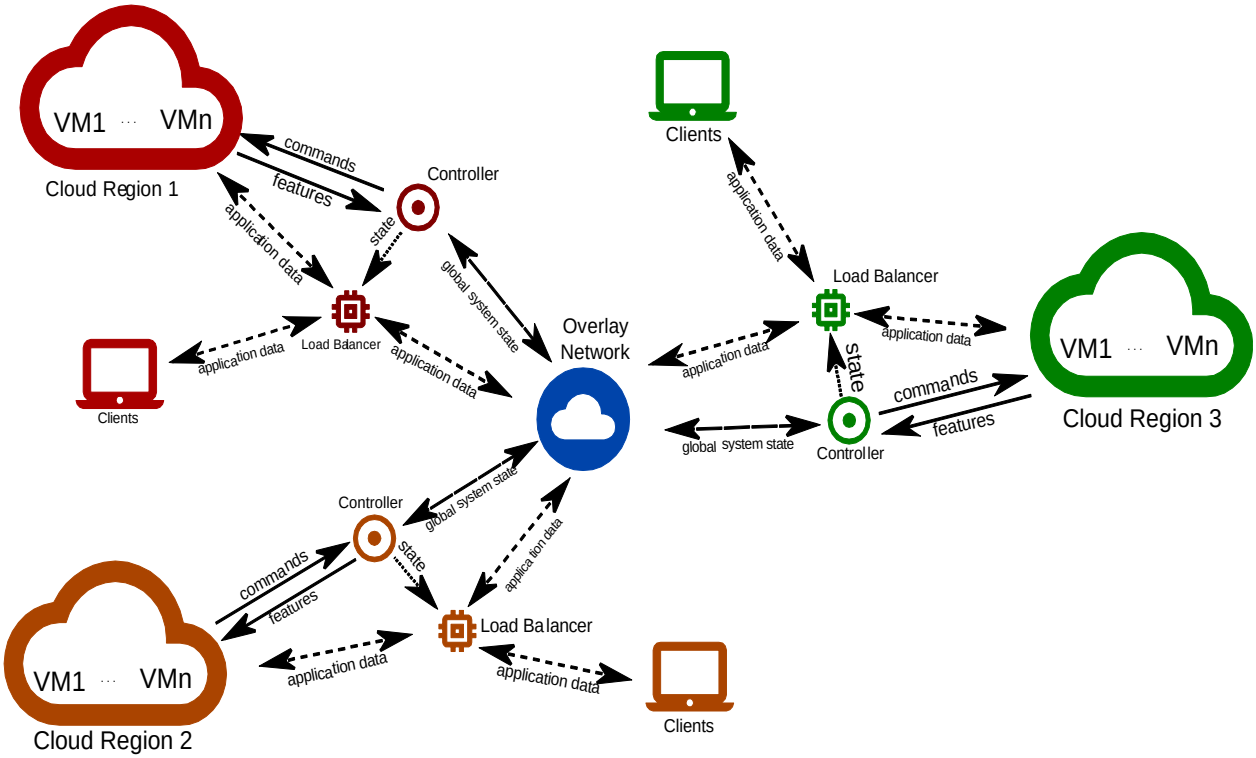

Figure 7 System Architecture for a proactive management of cloud resources based on Machine Learning.

In Figure 8, we report the results of controlled experiments using three regions: Munich (the leader, on a private infrastructure), Frankfurt (on a public Amazon infrastructure), Ireland(on a public Amazon infrastructure). Again, this controlled experiment is on a hybrid cloud. By the result, we can see that when the workload starts to increase, Inter ACM promptly detects that the predicted Region MTTF decreases. Nevertheless, the distributed knowledge algorithm detects, as well, that this increased workload is affecting, at the same time, both active regions. Therefore, Inter ACM correctly decides that the client requests should not be diverted to other regions. This shows that Inter ACM, relying on the prediction models which were built using ML algorithms correctly avoids making wrong decisions, in case of sudden bursts of load from remote clients. Nevertheless, when we reach minute 90 and the third region is added, we can see that (similarly to the previous experiments) the MTTF of the three regions starts to increase, as long as the forward probability grow. As well, the forwarding probability for three regions reaches *equilibrium* at minute 127 i. e. 33% . This is again an indication that Inter ACM is able to promptly cope with the variation of the conditions of the geographically-distributed deploy of the application, both when there is a variation on the conditions concerning the external clients connected to the region, and when there is an internal variation (due, e.g., to scale up/down of the system i.e. *elasticity property*). In the meantime, the Inter ACM maintains the *average response time for all regions below 3 seconds.*

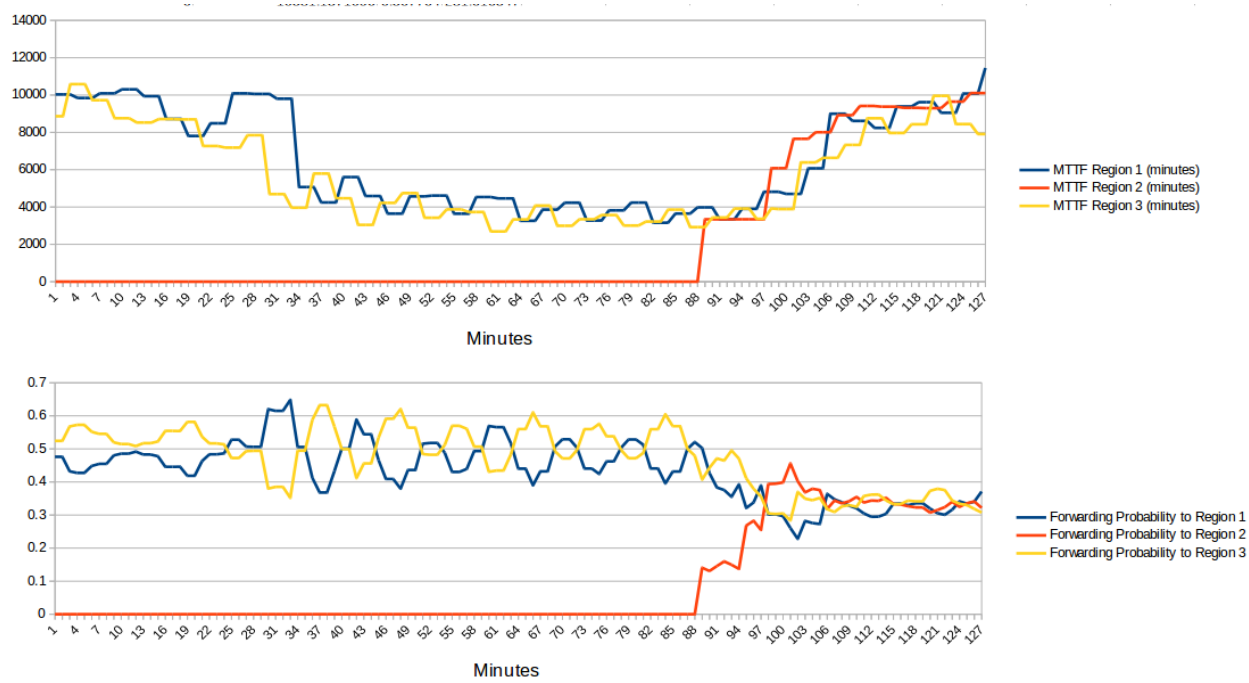

Figure 8 Request rate, MTTF, and Forward Probability, 3 regions, variable rate.

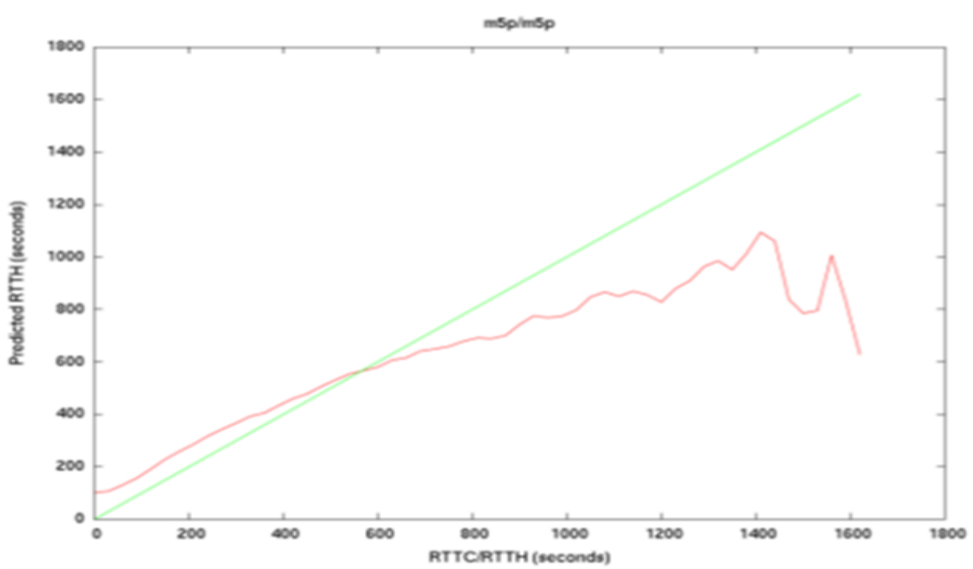

Figure 9 ML Prediction Model for Ireland Amazon Region.

Integrated Validation of Machine Learning (ML) framework with Autonomic Overlay Network

The ACM System driven by ML framework is successfully integrated with SMART overlay network. In Figure 10, we present the results related to the response time of the integrated systems without and without SMART. In Figure 11 results for the RMTTF without and with SMART. Graphs show that, in both cases, the RMMTF of the three regions are quite similar. However, with SMART, values of RMMTF are slightly higher (about 7%). The presented results demonstrate the capability of the system, after the integration with SMART, to handle the

cases where congestion or errors may occur within the Internet. As a result, the response time and the RMTTF will be maintained at the defined level.

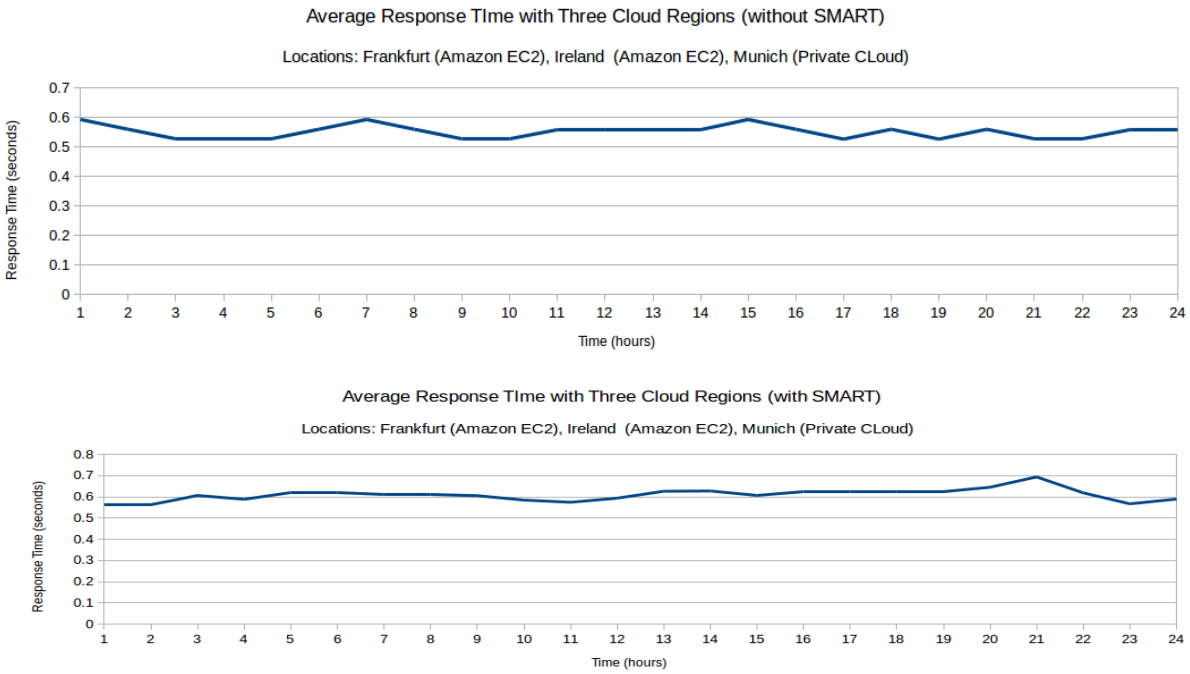

Figure 10 Response Time without and with SMART.

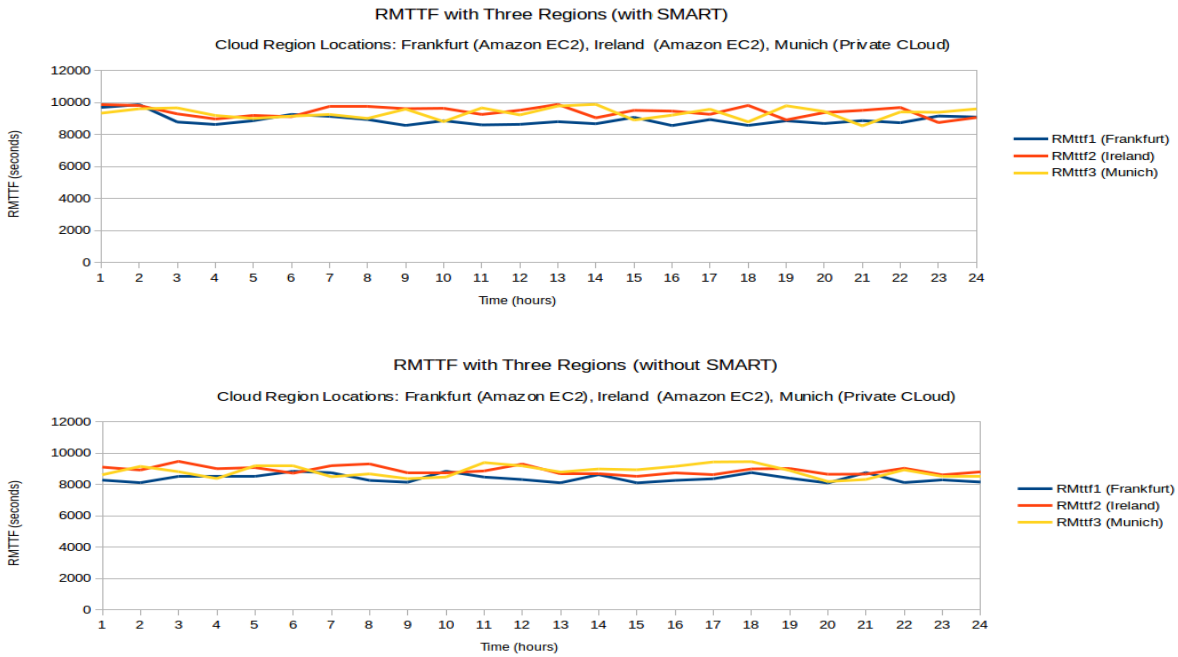

Figure 11 RMTTF without and with SMART.

Hadoop/IRIANC Autonomic Cloud Management (ACM) Integration

The Hadoop cluster is installed by IRIANC in Frankfurt region by sharing the Hadoop configuration, provided by ATOS, and using Amazon AMI (public cloud). The core of the ACM System is installed in Munich Private Cloud for generating the Machine Learning models. The architecture of the deployment of the Hadoop Cluster consist in two master nodes and three slaves. We have installed our ACM System on one of the Hadoop nodes dedicated to computation. In particular, on top of that node, we have installed our *Feature Monitor Client (FMC)*, which was connected to the *Feature Monitor Server (FMS)* running on the *namenode server* of the Hadoop setup. We have varied the incidence of memory leaks on the π application. We have then run the ACM ML Framework to build the ML-based prediction models. The first step entails running Lasso Regularization to select only the most relevant parameters to the construction of prediction models. Lasso is correctly able to reduce the number of parameters (from 35 to 4), which are considered as relevant to the construction. The smallest prediction error, training and validation time is obtained by REPTree ML technique. Lasso is more efficient in *real-time critical applications* in clouds.

Benefits from the Integration

In this section, we discuss what are the benefits on application availability in case of applications built on top of Hadoop when integrated with the ACM System and by relying on ML-based prediction models, is able to detect that an application in Hadoop 2.x is about to crash due to memory leaks (as well, ACM can make prediction for another anomalies). In the process of integration, a middleware can be designed. It can be notified by ACM about what is the about-to-fail node (VM – Worker.) Therefore, the middleware can ask to the *Hadoop Resource Manager* to decommission it. At this point ACM can provide a new VM (worker node). By using the middleware, the new node is assigned to Hadoop and it can start assigning jobs to it. Therefore, we are able to provide a new VM (worker node) to be added to an existing Hadoop cluster. The obtained results of the Integration of Hadoop with ACM System for generating ML prediction models can significantly contribute for increasing *availability of applications* running in Hadoop 2.x and maintaining the *response time* of Hadoop applications below a given threshold.

1.3.2 Cloud Management Platform for Autonomic Services

The cloud management solution of PANACEA provides the effector (self-configuration) and sensor (self-awareness) mechanisms necessary for autonomic management of services, which enables unattended services to request or provide information about individual VMs or the service itself and request reconfiguration operations, like migrating to another host or attaching a new device. In particular, it enables services to autonomously grow and shrink resources allocated to them as and when needed.

PANACEA uses OpenNebula as Cloud Management Platform and OneFlow as Service Manager. Self-awareness and self-configuration are provided by extending the OneGate component provided by OpenNebula.

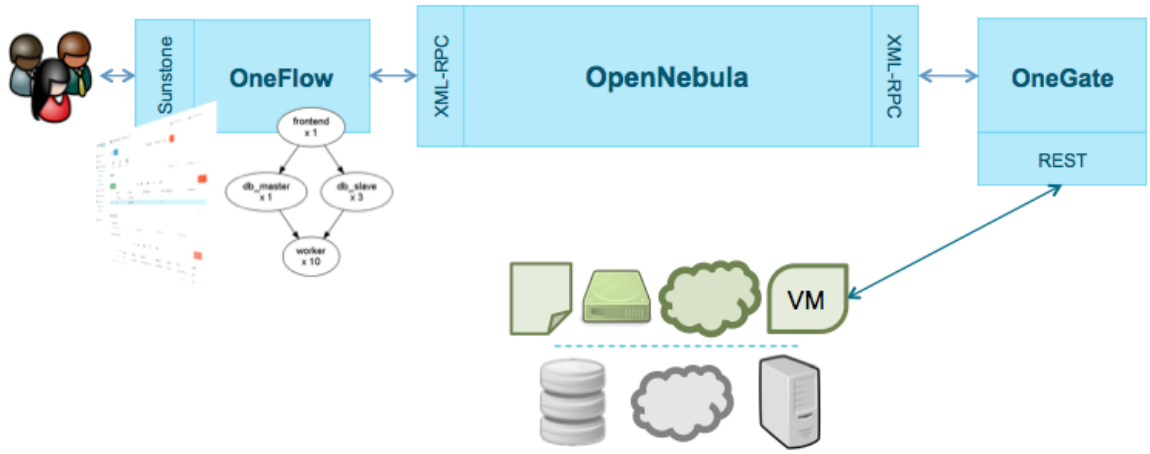

Figure 12: OpenNebula-based architecture for a Cloud Management Platform for Autonomic Services.

OpenNebula controls the physical and virtual resources of the infrastructure, monitoring hosts to know if they have available resources and automatically managing the lifecycle of VMs on them.

OneFlow allows users and administrators to define, execute and manage services composed of interconnected VMs with deployment dependencies between them. Each group of VMs is deployed and managed as a single entity, and is completely integrated with the advanced OpenNebula user and group management.

OneGate provides a REST interface for VMs to interact with OpenNebula. OpenNebula assigns an individual security token to each VM instance, and makes it accessible to applications running inside the VM through the contextualization mechanism. Applications use this token to contact OneGate in order to request operations to OpenNebula. OneGate checks the VM ID and the token sent and, if valid, the operation is forwarded to OpenNebula through its XML-RPC interface.

OneGate allows VMs to pull information from OpenNebula. For VMs that are part of a service, they can also retrieve information from other VMs or retrieve service information. This can be used to automate the deployment of services and VMs. OneGate also allows VMs to push information to OpenNebula. This can be used to provide configuration information or performance metrics.

Moreover, OneGate allows VMs to request operations on VMs, like `shutdown`, `resched`, `stop`, `suspend`, `resume`, `hold`, `release`, `resize`, `migrate`…. Also, the `scale` operation on the service allows a VM to modify the cardinality (number of VMs) of a given role, which is the basis for autonomic elasticity.

For the multi-cloud scenario, OpenNebula already had the capability to use resources from remote Cloud providers, such as Amazon EC2, IBM SoftLayer or Microsoft Azure. Since Amazon EC2 has been extensively used in PANACEA, its integration has been smoothed and a new contextualization mechanism has been added, making it possible to access the security token and URL needed to access OneGate from Amazon EC2 regions.

All features done in PANACEA have been already contributed to OpenNebula and publicly released. For example, self-awareness is included since version 4.8 and self-configuration is included since version 4.14.

OpenNebula has been integrated with SMART to optimize bandwidth or latency in hybrid cloud environments based on Amazon. It has been also used to deploy the Pervasive Monitoring System and the DAaaS platform.

1.3.3 Pervasive Cloud Monitoring

The role of the pervasive monitoring system (PMon) within the PANACEA cloud management solution is to collect measurements from physical hosts, virtual machines, services and applications, and provide the monitoring data to the entities that make management and allocation decisions based on the observed conditions of the cloud, such as the autonomous service and the cloud managers developed within PANACEA. The innovation of the pervasive monitoring system is to adaptively prioritize which nodes to monitor in the cloud, thereby achieving low overhead while providing timely delivery of fine-grained monitoring data and high accuracy of the representation of the cloud, which enables better, i.e., more responsive and closer to optimal, management and allocation decisions to be made. The innovation is primarily achieved by using machine learning techniques based on random neural networks (RNN) that continually learn which parts of the cloud are changing the most, and focusing the monitoring on those areas as needed.

The main purpose of the pervasive monitoring system is to support the proactive and autonomous management of cloud resources by providing fine-grained monitoring data in a timely manner without excessive overhead. The requirements of such a pervasive monitoring system were identified as *pervasiveness, elasticity, accuracy, timely delivery*, and *scalability*. *Pervasiveness* refers to the requirement that the monitoring system needs to collect measurements from a variety of physical and virtual resources at different cloud layers. *Elasticity* is the ability of the monitoring system to accommodate macro-changes in the cloud effortlessly, including run-time discovery of resources to monitor, e.g., new virtual machines (VMs), and VMs that migrate. *Accuracy* is defined as the correctness of the representation of the cloud provided by the monitoring data, and is mainly influenced by whether important changes in the cloud are reflected in the monitoring data in a timely manner. This requirement is therefore linked to the requirement of *timely delivery* of monitoring data, where data from the monitored resource should reach the users of the data quickly in order to enable responsive management. *Scalability* is the ability of the monitoring system to accommodate a growing workload, which is due to both the high number of resources and the number of metrics that need to be monitored. The pervasive monitoring system needs to efficiently collect data from many probes without impairing the normal operations of the cloud.

We have adopted an agent-based design approach for PMon in order to meet these requirements. Agent-based monitoring solutions are attractive as they provide flexible monitoring architectures that can be adapted to changes in the cloud. In our proposed design, PMon consists of monitoring agents (MA) and monitoring managers (MM), as shown in Figure 13. The MAs are located on the physical hosts and the VMs, and therefore have access to measurements from the physical infrastructure and also from within the virtual machines, i.e., from the services and applications running within the VMs.

The monitoring agent is the core component of the pervasive monitoring system, and it is a lightweight software module that is installed on the physical hosts and the virtual machines that need to be monitored. The main functions of the MA are to periodically collect measurements from its node, to store these measurements to a local round-robin database

(RRD), and to communicate with its monitoring manager (MM) for control and data collection operations. Since we want to be able to provide fine-grained monitoring data from a number of metrics, the agent needs to be lightweight, i.e., it needs to be able to run a large number of probes for frequent collection of data without incurring a significant overhead on its node.

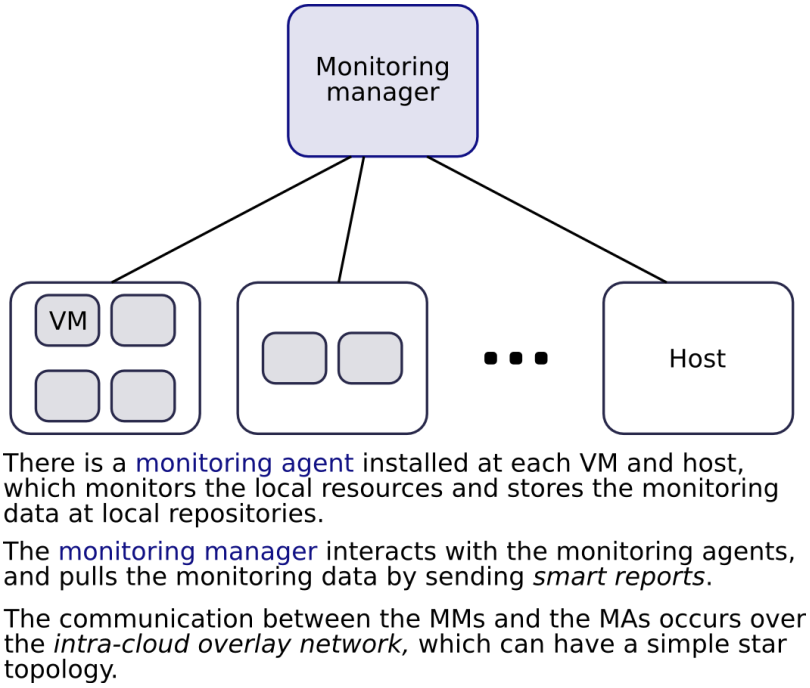

Figure 13: The agent-based architecture of the pervasive monitoring system (PMon). PMon consists of monitoring agents (MA) and monitoring managers (MM), connected through an intra-cloud monitoring overlay network. Each monitoring overlay constitutes a virtual monitoring system (VMS).

We use collectd as the basis of the MA, and leverage related existing tools in order to ease the development process while focusing on the main innovation of PMon, i.e., adaptive monitoring of the cloud by the application of machine learning methods for self-configuration and self-optimization. The design of the MA is presented in **Erreur ! Source du renvoi introuvable.**, which depicts the components that make up the MA as described below.

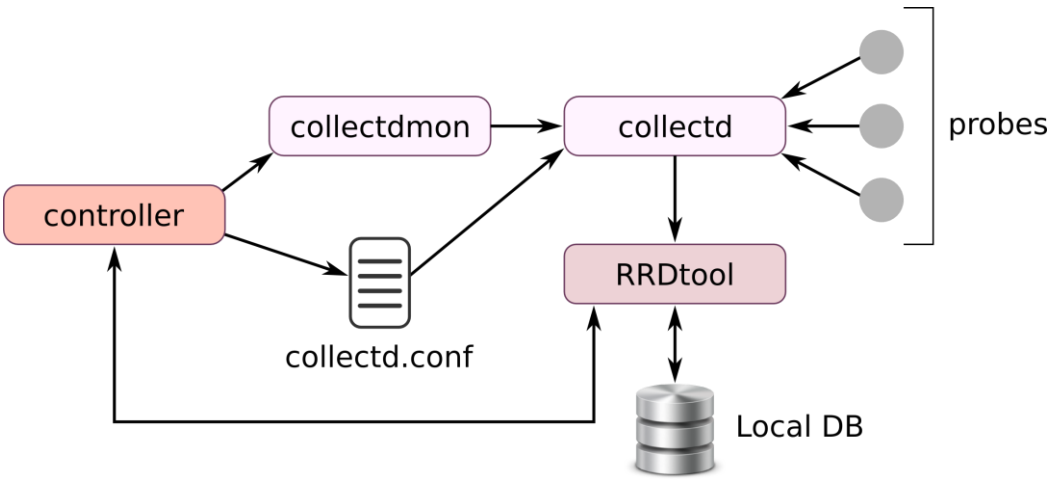

Figure 14: The design of the monitoring agent. The main functions of the MA are to collect periodic measurements from its hosting node, i.e., the physical host or the VM, to record these measurements to a local

round-robin database (RRD), and to communicate with its monitoring manager (MM) for control and data collection operations. We leverage existing software components for the design of the MA, in particular the collectd daemon and RRDtool, and supplement these components with the central controller component that implements the agent management functions and communication operations for interaction with the MM and other MAs.

The MM implements the control and data collection functions of the monitoring system, and provides an interface to PMon. The MMs collect the monitoring data from the MAs in an innovative way, and provide it to the interested entities, e.g. the autonomic cloud managers and the service managers, via a local round robin database. The MMs employ the innovative RNN-based monitoring method in order to dynamically select which parts of the cloud to monitor. Each MM controls a set of MAs, activating, reconfiguring, and deactivating measurements at MAs as necessary.

1.3.4 Online QoS-driven Task Allocation

We have developed a task allocation platform (TAP) that is internal to a Cloud can be used to dynamically make task allocation decisions that optimize QoS, that uses both model based and learning based approaches, each of which exploits measuerements collected in real time in the Cloud. Other recent work (below) has shown that a limited amount of re-routing over an overlay network using a measurement and big data approach can substantially average end-to-end delay for internet traffic and also reduce the perceived packet loss.

Thus we have combined these two ideas to focus on the QoS that tasks receive, in particular the overall response time which is determined by the network latency to access, forward data and programs, retrieve results and data, and which includes the local or remote Cloud processing delay. Thus we have developed and tested a practical system for adaptive workload distribution across multiple Clouds over wide area networks. The system includes a TAP deployed in each Cloud that optimises user perceived QoS and exploits the routing overlay SMART over Clouds for improving network delay incurred by data transfer. User requests are routed to a designated local Cloud, which may well be the geographically closest one, provided that it has enough available capacity to handle the request. When the workload at the local Cloud increases, the TAP at the local Cloud can decide to forward requests to remote Clouds. In the process, TAP will consider the effect of both the data transfer delay and Cloud processing delay, each being weighted for their relative importance. The estimate of data transfer delay used in our system also takes into account the measured packet loss which will result in extra network delay for applications that use TCP for data transmission. In order to optimise Cloud delay and network latency, our approach is meant to be easily deployable over a large population of machines, and it is able to make fast online decisions resulting in good QoS with low computational overhead. Although it requires measurement and monitoring both of network characteristics and of local and remote Cloud delays, we limit the frequency and overhead related to the monitoring effort, and also limit the computational complexity of decision making by using reinforcement learning.

Our approach benefits from limited measurement overhead as it probes the performance of subsystems which provide better QoS, while still exploring less frequently a wider range of alternative systems that can in the future prove to provide improved QoS if the current set of frequently used subsystems result in poor QoS.

Each TAP is an autonomic job allocation service, which provides self-optimization functions for cloud-based services. The job allocation service is intended for cloud-based services that

receive or generate many short-lived jobs frequently, and assigns each job to the best node, e.g. VM, that is among the resources available to the service and able to execute the job at the time. The selection of the "best" node is determined based on the job type and its quality-of-service (QoS) requirements, which will be at least partially based on its type. Therefore, the job allocation service will receive incoming jobs as input, together with their type and QoS requirements, and will dispatch each job to the "best" node.

The job allocation service continuously monitors the conditions of the available resources in the cloud, e.g. the load of the VMs, and also monitors the performance, i.e. the QoS, that jobs receive once dispatched to their respective nodes.

The allocation service uses these monitoring data in order to update its internal representation of the state of the cloud, which is captured by one or more random neural networks (RNNs) stored at the job allocation service provider, i.e. the job controller. The controller keeps an RNN for each job type/QoS, and each RNN has *n* neurons, with each neuron representing a node to which jobs can be assigned. The RNNs are updated using the collected monitoring data and a learning algorithm, and are subsequently used for job allocation decisions. We are currently using RNNs with reinforcement learning, but other learning algorithms are also applicable.

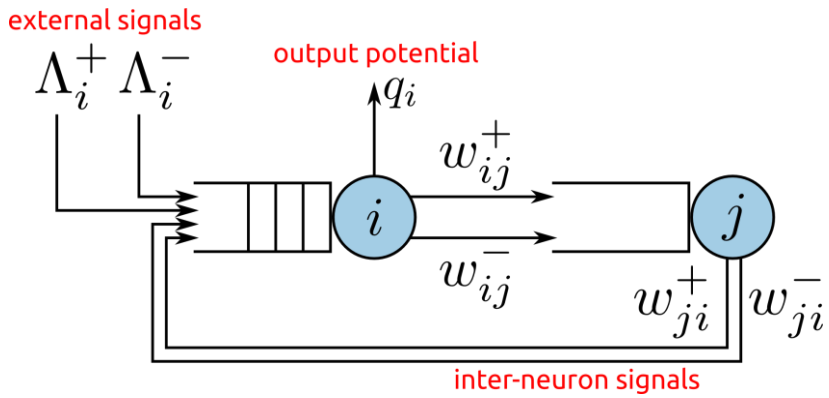

Figure 15: Random neural network

Figure 15 shows the basic structure of a random neural network, depicting the connections between the neurons, the input and output signals of the neurons, and their output potential.

Figure 22 shows the overall structure of TAP which is the basic building block of our task dispatching architecture.

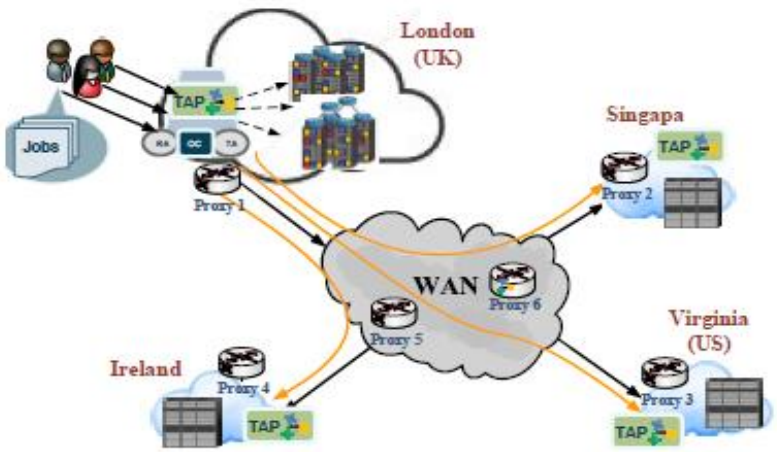

Figure 16 The Task Allocation Platform (TAP) and its insertion in the overall System Architecture

Experiments were conducted on a real large scale system operating on the Internet at a global scale and we empirically evaluate the potential of our system when there is great diversity both in the types of jobs, the class of QoS criteria and the resources they request. The experimental results that we obtain, validate the adaptiveness and effectiveness of our proposed system for dynamic environments.

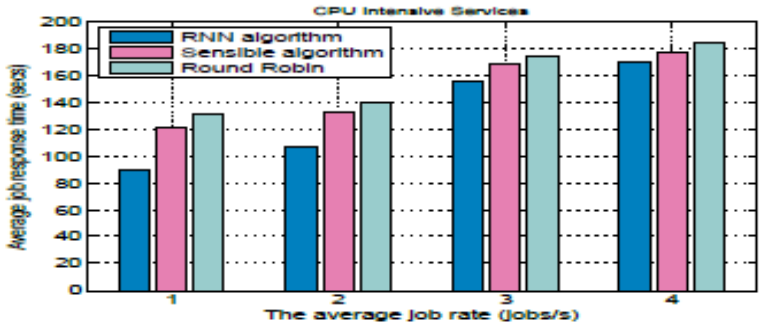

(a)

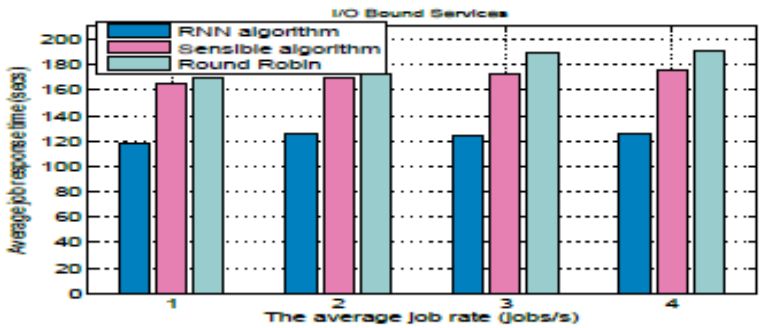

Figure 17 TAP Performance using the RNN based and other algorithms operating in a single Cloud; the observed task(job) response time is shown as a function of overall load (x-axis) on the Cloud, for different allocation algorithms that are tested. We observe the superiority of our proposed RNN (and Sensible) algorithms, over the standard Round Robin approach.

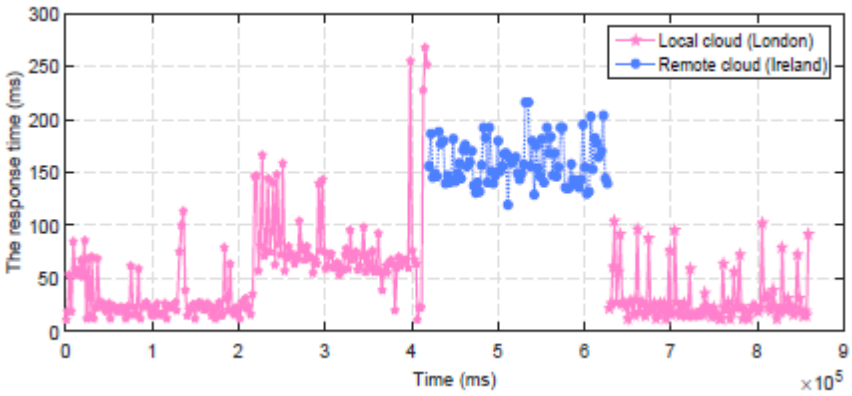

Figure 18 The system behaviour shown describes multiple TAPs operating with multiple Cloud services as in Figure 16. Measurements are conducted over a long time interval of 100,000 seconds. We observe that TAP offloads tasks to remote Clouds dynamically as the local load increases, and then it reallocates tasks locally when the local load drops down.

1.3.5 Autonomic Routing Overlay

After Internet routing was shown in a number of classic measurement papers to result in paths that are sub-optimal with respect to a number of metrics, routing overlays were proposed as a method for improving performance, without the need to re-engineer the underlying network. We have developed SMART, a self-healing, self-optimizing and highly scalable routing overlay, which has a number of advantages with respect to existing solutions. First, SMART can run with off-the-shelf applications and does not require any kernel modification. In addition, SMART can be widely deployed over a sizable population of routers, because it can quickly learn and efficiently track the optimal path with a limited monitoring effort.

The SMART architecture is shown in Figure 19. In this architecture, two types of agents are used:

- **Proxy**. An agent, called a Proxy, is executed in each cloud and acts as an intermediary for communications with other clouds. It basically realizes two functions. First, it monitors the quality of the Internet path between the local cloud and the other clouds. The monitoring information is used to determine the optimal path in the overlay network to each possible destination. Second, the source Proxy routes each incoming packet sent by a local application to its destination on the optimal path, while intermediate Proxies forward the packet to the next hop on the path chosen by the source.
- **Transmission/Reception agents**. They are located on each VM running an application controlled by the PANACEA system. These agents are started automatically at start-up of their respective VM. The role of the Transmission Agent (TA) is to intercept the packets sent by the application and to forward them to the local Proxy, whereas the role of the Reception Agent (RA) is to receive the packets sent by the Proxy and to deliver them to the local application.

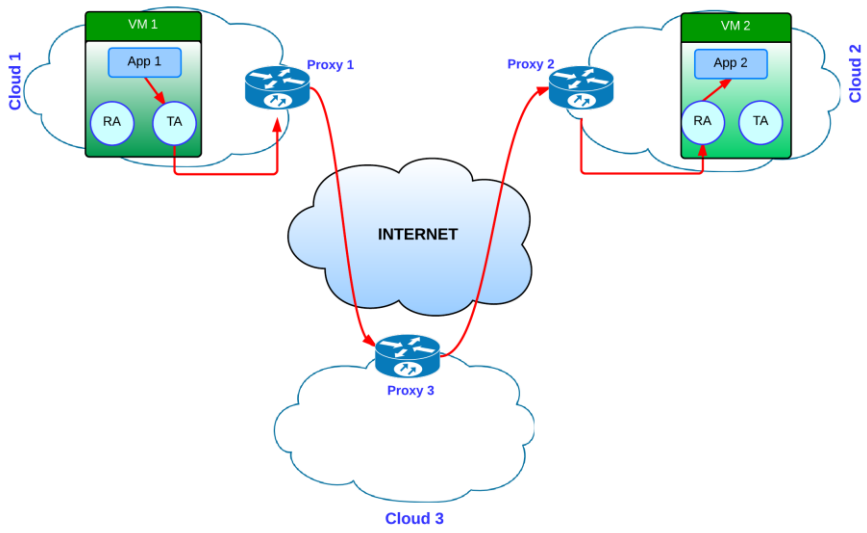

Figure 19: Architecture of the PANACEA overlay routing system.

In this architecture, the routing/forwarding of a packet proceeds as shown in Figure 20.

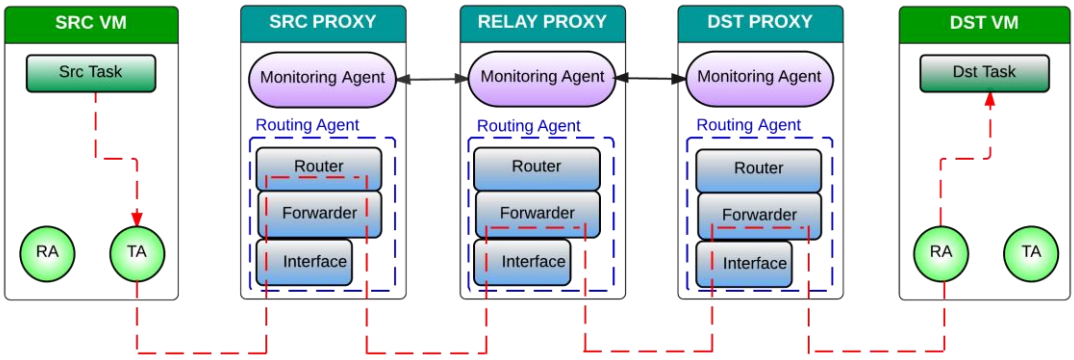

Figure 20: Routing process.

When a packet is sent by a source task to a destination task located in a different cloud, it is first intercepted[1] and forwarded to the TA. The TA uses IP-in-IP encapsulation to forward an altered packet to the Proxy. The payload of the altered packet is the original packet, plus an additional SMART header. Upon reception of the SMART packet, the routing agent of the Proxy looks-up its routing table in order to determine the path to the destination. The sequence of intermediate proxies is written in the SMART header, and then the Panacea packet is forwarded to the first one of these proxies. Each intermediate proxy then forwards the packet to the next hop on the path, until the final proxy is reached. When this occurs, the packet is forwarded to the RA of the destination VM. The RA desencapsulates the SMART packet and forwards the original IP packet to the destination task using a raw socket.

[1] **Packet interception is realized using a filtering mechanism known as NetFilter NFQUEUE. NFQUEUE is an iptables target which delegate the decision on packet to a user space software**

The optimal path to a destination cloud is found using active monitoring. The monitoring agent of a Proxy regularly measures the latency or throughput to the destination cloud by sending probe packets along paths to that destination. Existing routing overlays use all-pairs probing, which has the advantage that it is guaranteed that an optimal path is discovered; but the downside is that this approach does not scale very well due to its costly $O(n^2)$ monitoring effort in an overlay of n nodes. Since we wish to build a routing overlay that can be widely deployed over a sizable population of routers, it makes sense to minimize the monitoring effort required to discover the shortest path.

We have first studied this problem as a graph-theoretic problem, known as the Shortest Path Discovery (SPD) problem [13]. In this problem, we are given a complete graph, a source node and a destination node. The lengths of the edges are unknown, but they can be queried. The SPD problem aims at finding the shortest path between the source and the destination querying the minimum number of edges. In addition to designing an online algorithm that guarantees an approximation ratio of 2, we have also shown a number of results for this problem, including a negative result which states that for any algorithm, there exists a bad instance for which the number of queries will be $O(n^2)$ [14].

In view of this result, we have designed online algorithms, which instead of guaranteeing that an optimal path is found at each time measurement epoch, guarantees that the average performance is asymptotically the same as the optimal routing. These algorithms exploit past observations so as to quickly learn and efficiently track the optimal path. In fact, the problem of learning the optimal route between two given nodes can be cast as a multi-armed bandit problem, for which many efficient algorithms have been proposed. We have investigated the use of the EXP3 algorithm proposed by Auer et al. [27], as well as an approach inspired from Cognitive Packet Network (CPN), which uses Reinforcement Learning to adjust the parameters of a Random Neural Network (RNN), acting as an adaptive critic [28].

As an example of the results obtained, Figure 21 shows the results obtained in a live experiment over the Internet, in which two traffic flows were transmitted in parallel from Sao Paulo to Tokyo, one of them being routed over the IP route whereas the other was routed through the overlay. As can be seen, in this experiment SMART allowed to increase the throughput by 51%.

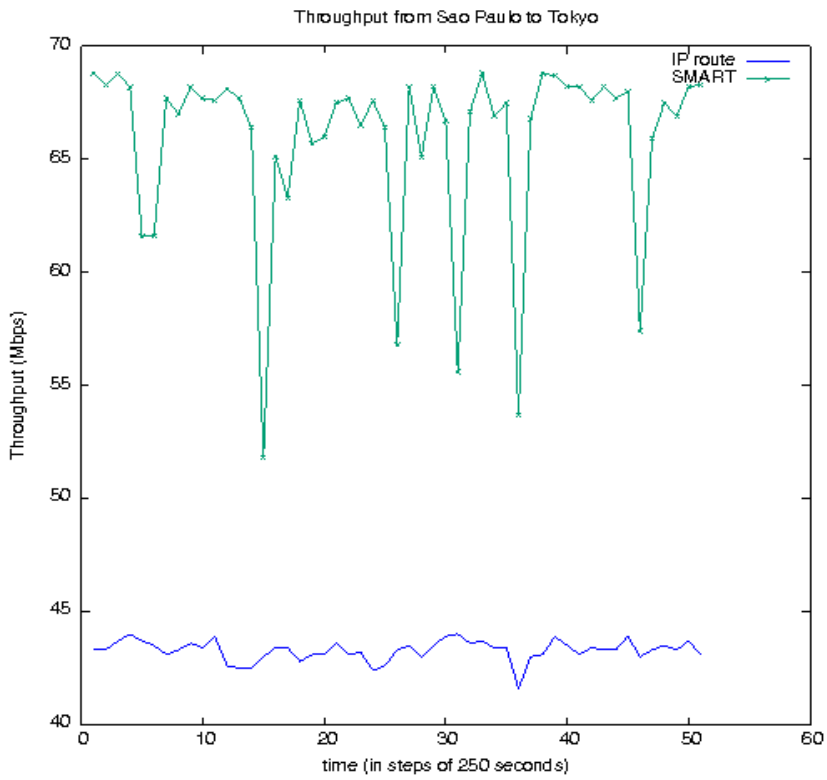

Figure 21: Results obtained in the live experiment with SMART.

1.3.6 Simulation/Emulation Environment

Overlay networks are considered as a flexible solution for traffic engineering over IP networks. In fact, IP networks are suffering from significant congestion and slow recovery from failures. This can have a major impact on distributed applications (over clouds).

QoS Design's objective in the PANACEA project is focused on the validation of the SMART module, which is developed by LAAS-CNRS and is part of the PANACEA solution. SMART's main role is to monitor the communication links of the overlay network and to discover and track the optimal paths according to application-specific metrics.

We are mainly interested in the ability of SMART to bypass and recover from many inevitable anomalies, and to autonomously optimize performances of the overlay network under changing conditions, thus providing a self-healing and self-optimizing communication system on top of the Internet network.

The validation process was carried out using two approaches (simulation and emulation). Both approaches aim at providing large-scale experimentation and more flexibility in controlling the system behaviour, which is not possible in a real-life environment (Internet).

The simulation approach is based on the interaction between NEST's environment and SMART. After simulating the network behaviour and evaluating its key performances, the simulation kernel exchanges with SMART the needed information, i.e. the overlay link metrics, such as end-to-end delay or available bandwidths. SMART analyzes the metrics of monitored

overlay links, based on its learning algorithm, and detects whether some paths can be optimized.

As an example let us assume that, after being informed of the evolution of end-to-end delay of the monitored overlay links, SMART detects degradation in performance for a certain path between two proxies. Consequently, by analyzing the overlay network, it is able to find an alternative solution to overpass this degradation. This modification would eventually ameliorate the end-to-end delay, as this new delay was coming from a certain network event, such as a path congestion or a router failure.

The emulation environment differs from the simulation one in that the network infrastructure is emulated and real VMs hosting real applications are used. This environment will have SMART configured and operational, meaning that we will deploy and activate on each VM a SMART proxy. The VMs hosting the real applications will be linked to the emulated environment as if they were connected to routers of a real network. The emulation of the network is carried out using CORE[2] emulator.

The emulation process is based on the interaction between the emulated network, the application VMs and the simulation kernel (emulation driven by simulation). The NEST simulation kernel models the emulated network, generates background traffic based on user-specified parameters, and evaluates the performance metrics. This simulation kernel then sets the interface metrics in CORE so as to correspond to the simulation results. The overlay paths between client VMs and application VMs are consequently influenced by these link/interface metric changes.

SMART, whenever necessary, modifies the overlay paths in a manner to enhance the overlay links performances and informs the simulation kernel of such modifications.

Validation for both the simulation and emulation processes were carried out as follows.

- Traffic is generated and routed in the network.
- Adverse events are introduced (failures, traffic perturbations) during certain time-frames to degrade performances.
- The performance metrics are computed and logged for the overlay links
- Gain in performance is evaluated based on the comparison of the two situations where SMART is activated or not
    - We consider the following situations in our use cases
        - the end-to-end delay metric which influences the response time perceived by the users,
        - the available bandwidth on the overlay link which influences the volume of requests that could be served.

The first experimentation illustrates clients injecting http requests toward a web server. The idea is to evaluate the impact of the network on the response-time of each request. After activating SMART, we check whether it can alleviate this impact or not.

We ran several simulations that span over a 24-hour period. We extracted a window of 3 hours (cf. Figure 22) containing the most significant results. During this window a huge perturbation that lasts for one hour occurs.

[2] http://www.nrl.navy.mil/itd/ncs/products/core

We compared the response times of http requests sent during this period with and without SMART. The results are compiled in the following chart.

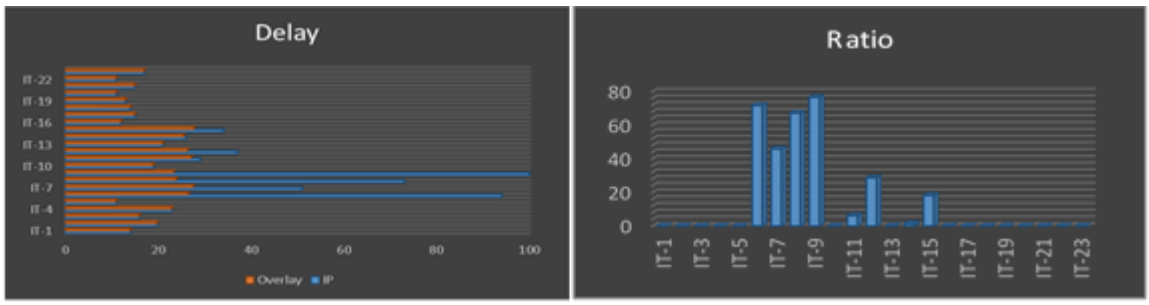

Figure 22: Comparison of response time over the IP and SMART Path.

What is significant in this case is the improvement ratio. While the IP path can deliver good performance in general, sometimes SMART can achieve ~70% improvement during perturbations and maintain good performances for the end user.

We also conducted a test campaign focusing on the bandwidth metric (cf. Figure 23). During this experimentation, we used the same scenario and the same environment. Rather than injecting http requests, we injected new file transfer requests that download 10MB files.

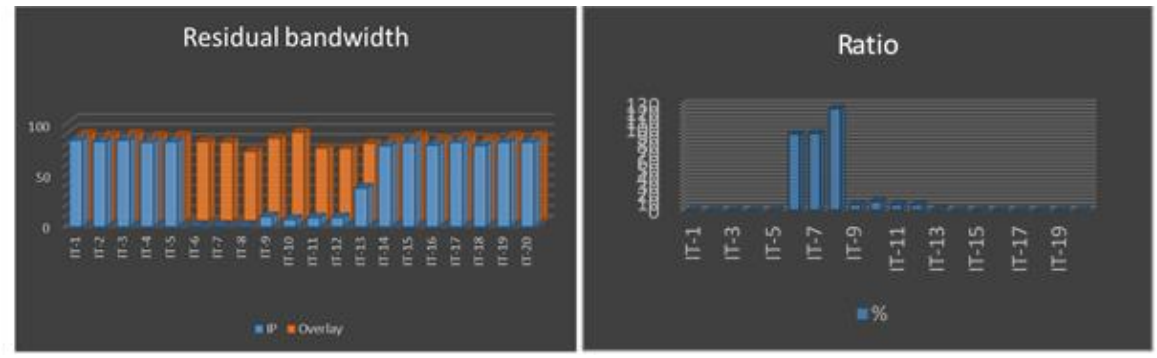

Figure 23: Comparison of residual bandwidth over the IP and SMART Path.

It appears that while the IP path suffers congestion, SMART was able to find better paths using intermediate proxies.

From a user perspective, on the iteration "IT8", an improvement of 124% is achieved using SMART. We measured the time needed for a user to download a 10MB file.

The results that we obtained are:

- On the IP path : 1m52s (737.4 Kb/s)
- Over SMART Path : 33s (2.42Mb/s)

This means a 70.5% improvement on the user side.

Finally we can summarize the main outcomes of QoS Design's work within PANACEA as:

- An enriched **simulation environment** for designing and optimizing overlay networks.
    - A simulation kernel, that allows long run simulations taking into account the underlying network (routers, links, data-centers, protocols, etc.).
    - A what-if analysis for several adverse events (link failures, router failures,

server failures, congestion, etc.).

- A **hypervisor** environment for SMART system that allows the monitoring of the health status of deployed services, routing decisions over the time and network metrics.

A **Test platform** providing an emulation driven by simulation approach, this platform, allowed the validation of self-healing and self-optimizing features of overlay networks. This platform allows the deployment of a distributed application in a quasi-realistic environment, with the benefit of cutting-off all costs related to a real deployment in a real cloud infrastructure.

1.3.7 Autonomic Elasticity

Elasticity has been and still is one of the main attractive features of Cloud solutions. The ability to dimension the size of a platform to the incoming load from the clients has a double benefit: from one side, the user experience is better, since the infrastructure grows with the number of requests by the users, from the other side, the infrastructure adapts to the number of user requests peaks that typically happens to web-services type apps. For the companies running those services, this mean that they only pay for what they need at a giving time, not failing in the risk of over provisioning resources –paying more of what it is necessary– or, under provisioning the number of resources, affecting the user experience in a negative way. Also, since this elasticity process happens in an automatic way, there are no operational costs, the infrastructure grows of shrinks as needed.

The Autonomic Elasticity in PANACEA is the result of combining two previously introduced technologies: the Pervasive Cloud Monitoring described in Section 1.3.3 and the Cloud Management Platform for Autonomic Services described in Section 1.3.2. The monitoring helps to understand the status of the application and if it is necessary to perform a change to it to adapt to the incoming load and the Cloud Management platform will help to perform those actions in an easy and quick way.

Of course, the way the number of resources, in this case VMs, are increased, depends on the Cloud application. To demonstrate the usefulness of this technology we employed the UC2, Data Analytics as a Service (DAaaS). DAaaS is Cloud platform designed at Atos to help our customers to easily build data processing workflows from different big data sources. Since this application is offered as a service to Atos costumers, the load incoming to the application changes overtime, depending on the type and number of costumers Atos gets for the service. Having the ability to automatically scale the service to the incoming load will help Atos save both CAPEX and OPEX costs.

DAaaS with Autonomic Elasticity architecture

DAaaS service can employ several big data analysis tools, for this specific demonstration we are going to focus in one of the most used around the world, Apache Hadoop[3]. The following figure presents the architecture of DAaaS adapted to PANACEA:

[3] **http://hadoop.apache.org/**

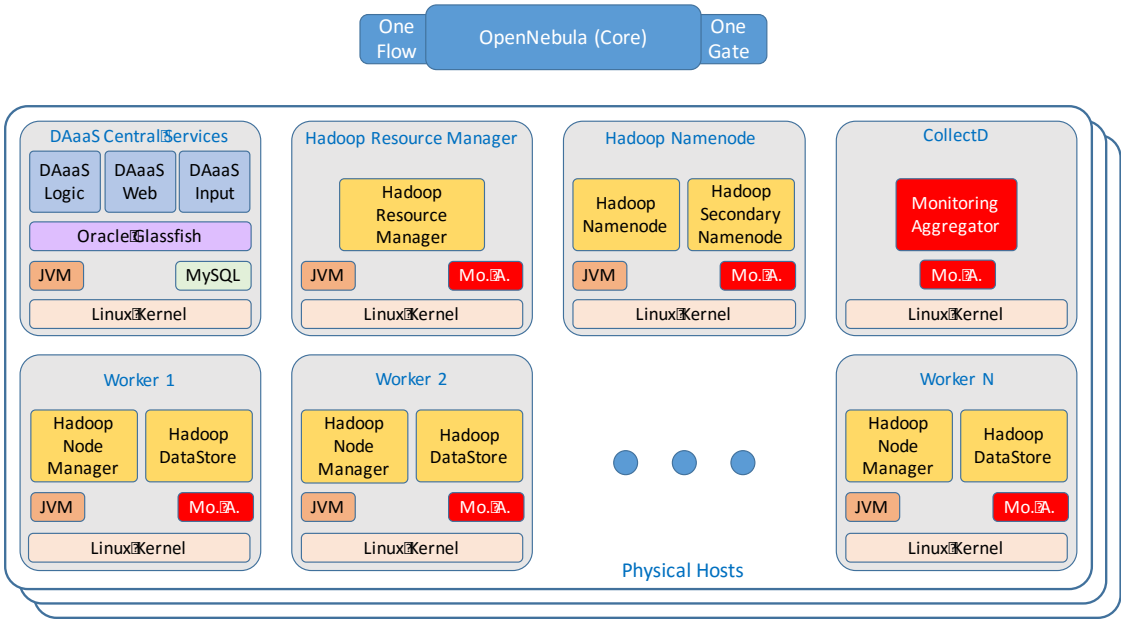

Figure 24: The initial deployment schema for the DAaaS UC.

DAaaS is composed of 5 types of different VMs. DAaaS Central Services that compose the input point for the users. Hadoop Resource Manager that manages the Hadoop Cluster resources. Hadoop Namenode that manages the Hadoop distributed filesystem. Monitoring that centralize the monitoring information to be used by the logic that scales up or down the system. Finally, the worker nodes, the ones that keep a copy of the data and perform the big-data calculations for the DAaaS solutions. It is this last type of nodes the ones that can scale up or down. Also, on the top of the figure we can observe the OpenNebula OneGate and OneFlow services that will help to control the elasticity process.

Hadoop Elasticity

As we commented previously, we wanted to increase the number of working nodes depending in the incoming load. With our previous experience with Hadoop we speculated this could be achieved by watching the pending tasks queue in the Hadoop Resource Manager (we are only scaling in the computation jobs, ignoring the data side of things to reduce the complexity of the problem). So, based in this variable, that we got via the Pervasive Monitoring, and using the OneGate and OneFlow technologies of OpenNebula we scale-up nodes (increase the size) or scale down worker nodes (shrink their size). Each of those actions have several implications:

- **Scale up –** This is quite simple and straightforward process. Create a new working node via OneGate/OneFlow and register it to the Hadoop Resource Manager and NameNode.
- **Scale down –** This is a more complicated problem. First we need to notify Hadoop that the node is going to be decommissioned. Doing this, Hadoop stops sending jobs to it and starts checking if it needs to replicate the data that it stores in other place. Once those actions are done, the node can be safely deleted from the Cloud.

Experimental results

For the cluster detailed previously, using 3 workers nodes as starting point, we start to send

jobs composed of 32 maps actions and 10000000 reduce actions. For the amount of resources that we had, this ensures that it will take several hours for the jobs to finish. All those jobs were sent to the Resource Manager into the shape of 100 small map-reduce applications, with the idea to stress incrementally the system. The following graph shows how long takes to finish this cluster without allowing it to scale up or down (6 hours and 59 minutes):

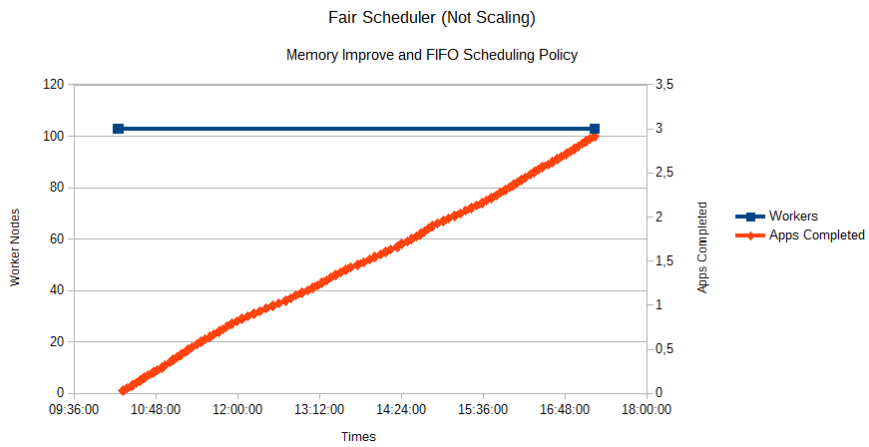

Figure 25: No allowing scaling result.

If we allow the system to scale-up and down, we see that the elasticity logic decides to increase the number of nodes in while in the middle of the processing, increasing for a bit the number of jobs finished. After the queue gets smaller, it reduces again the number of nodes for the remaining set of jobs. Doing this the process took more than 10 minutes less than before:

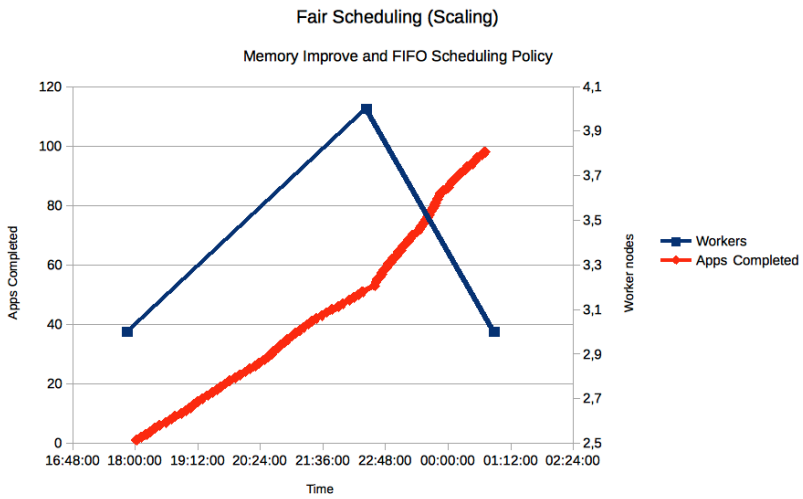

Figure 26: Allowing scaling results.

After fine tuning a bit more the configuration of the elasticity engine, it was possible to have savings of nearly two hours with respect the previous case:

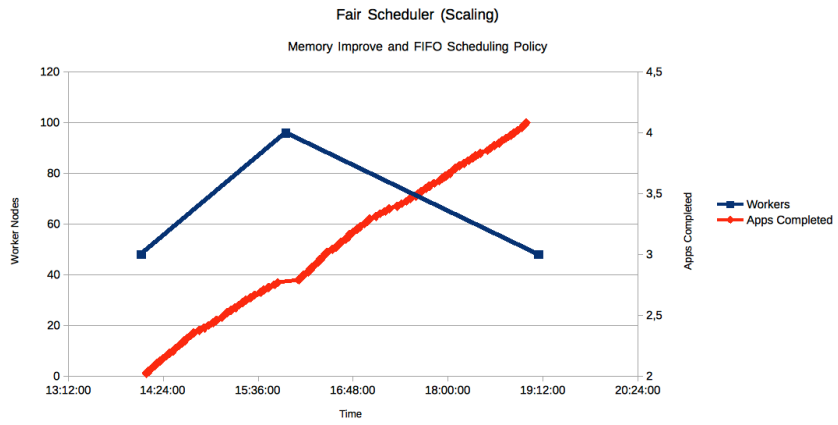

Figure 27: Allowing scaling results, with optimal memory configuration.

Conclusions:

The work in understanding elasticity in Hadoop has just started internally at Atos. The objective now it is to learn how to correlate those metrics with the type of incoming applications coming to the cluster.

Also the study of the parameters indicates that other possible solution it is to create clusters adhoc. Thanks to the self-configuration techniques in PANACEA it is quite quick to deploy a complete Hadoop cluster in matter of minutes. Depending on the type of map-reduces jobs that comes to our system, it will be easy to deploy and adhoc Hadoop cluster, optimize in configuration for that type of job, and when it finish, delete it completely. **Saving both OPEX and CAPEX costs.**

In both cases, it is important to note the necessity to separate Hadoop computation resources and HDFS resources, with the possibility to use object-storage solutions instead of HDFS ones. This decoupling will allow to have independent lifecycles for data part, that it would probably have a longer lifetime, to the computation part where, each computation resource, has an smaller lifetime, even in the scale of few hours.

1.4 Potential impact, main dissemination activities and exploitation of results

Potential Impact

For cloud providers, the expected benefits of the technologies developed in PANACEA are multiple, being the main reference the reduction of operating costs through a better management of failing resources. In addition, it can be expected that by significantly increasing the availability of cloud services, PANACEA will reduce one of the main barriers to adoption of cloud computing.

For cloud users, the proactive autonomic management scheme developed by PANACEA will lead to higher availability of cloud services, by predicting and reacting to failures before they occur, and to better performances, by reconfiguring in advance when performance degradations are predicted. In addition, it can be expected that the lower operating costs for cloud providers will translate into lower usage costs for users.

Main dissemination activities

The dissemination strategy defined at month 3 identified the stakeholders to be targeted, the list of activities to be performed and also the dissemination tools and channels to be used. In order to be aligned with the project maturity level, these activities were divided into three different phases:

- **First phase - Brand and awareness (October 2014 – March 2014).** The first phase of our dissemination strategy was mainly focused on raising awareness around the PANACEA brand. In order to achieve that, the consortium set up different tools like the project´s website, the different dissemination assets (posters, flyers, etc.), the creation of the social media (Twitter [29], LinkedIn [30] and Facebook [31]) accounts, etc.
- **Second phase - Communication (March 2014 – March 2015).** The second stage was more devoted to communicate the first achievements of the project. In order to do so, several scientific articles and papers were generated. Additionally, the consortium attended multiple conferences and workshops where PANACEA was publicized and first contacts with other projects and related initiatives were established.
- **Third phase - Implementation (March 2015 – March 2016).** The last phase of our dissemination strategy has focused on strengthening the dissemination and communication efforts of the previous stage, aiming at maximizing the impact of the project. In order to achieve that, we have attended to more conferences, published more scientific articles, increased the social media activity, generated specific demos, press releases and also created a commercial video for PANACEA.

Figure 28 summarizes the dissemination strategy for PANACEA.

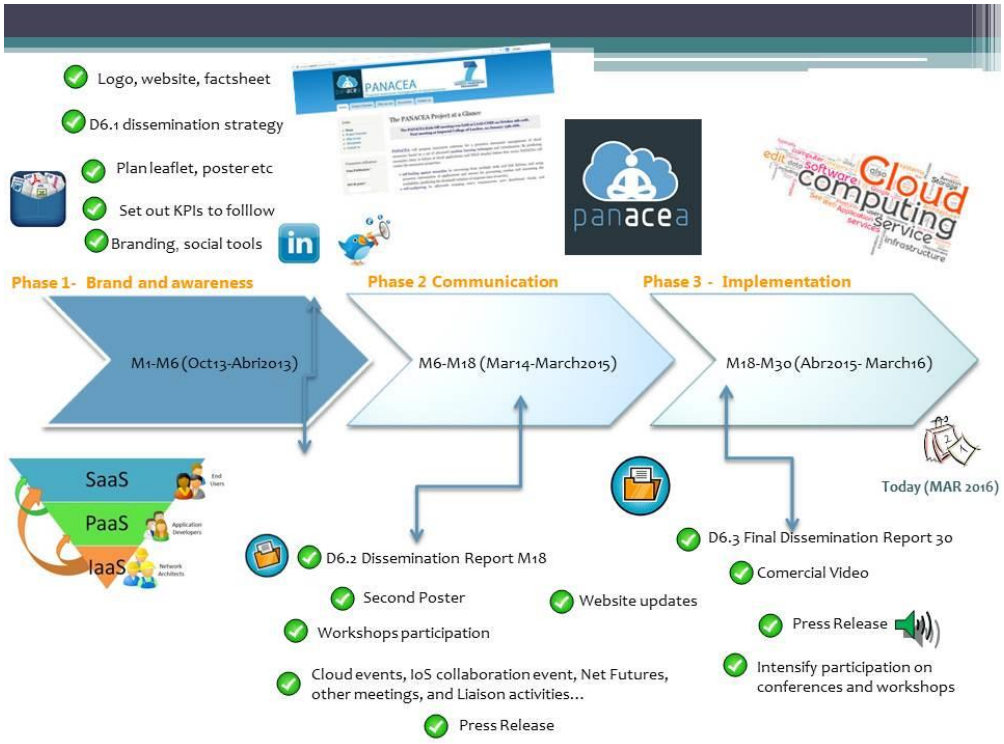

Figure 28: Dissemination status and recap

In order to measure the effectiveness of the dissemination strategy defined for PANACEA, we use the initial set of KPIs as a lower threshold to evaluate the work done. The results are presented in **Table 1**.

KPI	PLANNED	DONE	EVALUATION
Scientific Conference proceedings	14	17	EXCELLENT
Peer reviewed papers	3	16	EXCELLENT
Whitepaper	1	1	**GOOD**
Press releases	3	3	**GOOD**
Website visits	1000 visits per year	Average: More than 3.000 visitors per year	EXCELLENT
Presence in social media	**2 Groups:** LinkedIn and Twitter	**5 Groups:** Twitter LinkedIn Facebook *AppHub and OCD	**GOOD**
Significant presence at events	2	5	EXCELLENT

Table 1: Project´s KPIs

As can be seen on the table, the overall results are quite positive. On one hand we have some KPIs that had an excellent result (attendance to conferences, paper publications, website visitors, etc.) as they far exceeded the initial expectations. On the other hand, the number of

press releases and whitepapers generated meet the objectives set at the beginning of the project, so we evaluate them as good. Although the availability of a PoC to showcase to our potential stakeholders would have been extremely useful for the dissemination (and also for the exploitation) of the project, some interesting achievements have been obtained. Clear examples are the collaboration with other research projects (who could be interested on adopting PANACEA results) or the consortium decision of creating the **PANACEA Alliance**, which will continue evolving and supporting the solution for at least a year beyond the funding period.

Exploitation of results

The dissemination strategy defined for the PANACEA project has already yielded their first results. These results include the impact generated by the project at different levels.

One of the main impacts of PANACEA can be seen through the improvements provided to **Open Nebula.** The new features provided by the project, like the self-awareness and self-configuration mechanisms to support unattended services with improved availability and performance, have been already integrated on **OpenNebula 4.8,** codenamed **Lemon Slice**, and publicly released in August 2014. The contribution to Open Nebula is crucial for PANACEA, as it provides access to an Open Source Community with more than 900 registered users from both academia and industry, which could contribute to different aspects of the project.

Another important achievement of PANACEA, from the exploitation/impact standpoint, is the reuse of some of the components developed within the project in other initiatives. Example of this are **Prove** (a **STIC AmSud** research project that involves organizations from France, Uruguay and Chile) and **STEM** (a CELTIC project involving QoS Design), which will reuse and extend **SMART** (the Self-Managing Routing Overlay System developed by PANACEA) in their projects.

Additionally to this, the **knowledge generated** and the innovative capabilities provided by PANACEA have been used to improve the **Atos DAaaS platform** through use case 2. The maturity level of the platform is such that has now been included on an EIT ICT initiative called **MCloudDaaS** (Multi-cloud for Data Analytics as a Service) led by Atos.

On the academic side, besides the knowledge transfer generated by PANACEA (that will derive in new courses or PhD programs, and which could grant access to new grants and R&D projects), there is a very important achievement directly related to PANACEA. This milestone consists of a **12-months contract recently started by Imperial with Huawei**, which aims at designing specific data networks that support internal communications for Cloud systems and Cloud Data Centres. This work exploits the know-how developed in PANACEA, especially in the area of dynamic task allocation and dispatching using smart adaptive techniques.

In the case of **IBM**, the knowledge obtained from **PANACEA** is being used to improve the quality of their products and solutions. Indeed, the PANACEA outcomes are already contributing to IBM´s internal developments (mainly **FRAPPE**) which will be made available as a service on **BlueMix** (IBM' s cloud platform). Different patents have been submitted by IBM on this regard, as detailed on the exploitation deliverable D5.3.

Another important outcome of PANACEA is the **enriched simulation environment** for designing and optimizing overlay networks and developed by QoS Design during the project. Additionally to this, they also developed a **hypervisor environment for SMART**, which will allow the monitoring of the health status of the deployed services, routing decisions over the time and the network metrics.

Finally, a **test platform** has been also produced, which will allow simulating specific conditions to validate the self-healing and the self-optimizing features of overlay networks.

These are just some examples of the impact generated by PANACEA on both the academic and industrial environments. If we take into account the initial objective of the project, which was mainly focus on research, these achievements are even more important.

These achievements have been possible thanks to the effort and commitment of the consortium partners, together with the collaborations established with other projects and related initiatives (like SeaClouds at ESOCC, the OCEAN workshop, AppHub, Ocean Cloud Directory or CloudWave).

1.5 PANACEA Contact Information

The project website URL is: panacea-cloud.eu

Any of the contacts listed below in Table 2 will be happy to give more information about the project, answer any technical or press enquiries and discuss licensing or R&D collaboration opportunities.

For project administration enquiries, please contact the project coordinator, Olivier Brun (brun@laas.fr), at LAAS-CNRS. For project research enquiries, please contact the Research Board coordinator, Dimiter Avresky (autonomic@irianc.com), IRIANC. For contractual, legal and financial matters, please contact Laurent Perez (lperez@laas.fr).

Leader of	Name	Organisation	Contact
WP1 - Project Management	Olivier Brun	CNRS	brun@laas.fr
WP2 –Monitoring System	Eliezer Dekel	IBM	dekel@il.ibm.com
WP3 – Machine Learning	Dimiter Avresky	IRIANC	autonomic@irianc.com
WP 4 - Evaluation	Erol Gelenbe	IMPERIAL	e.gelenbe@imperial.ac.uk
WP5 – Exploitation	Ivan Febles	ATOS	ivan.febles@atos.net
WP6 – Dissemination	Ivan Febles	ATOS	ivan.febles@atos.net

Table 2: WPs and responsibilities

2 PLAN FOR USE AND DISSEMINATION OF FOREGROUND

2.1 Dissemination Measures

Elaborated material for dissemination purposes

We briefly summarize below the dissemination assets generated during the project lifecycle:

- **PANACEA Factsheet**. One of the first documents generated on PANACEA was the project´s factsheet. This document is available on the project´s website and provides an overview about PANACEA.
- **Project' s logo**. Another important element generated at the beginning of the project is the PANACEA logo. If the factsheet was one of the first attempts towards raising awareness about the project, the logo is a core element for creating the **PANACEA brand**.
- **Flyer**. This brief PANACEA "snapshot" was an easy way to publicise the project and to reach a wider audience, while supporting the project partners on the different events and conferences they attended (FIA, Closer, CloudScape, the two Concertation meetings, etc.). The flyer described the project and its main functionalities, together with the potential impact of the different stakeholders of our solution.
- **Project' s poster**. We have generated two different posters. The objective of both posters was to provide information about PANACEA in a fast and eye-catching fashion, aiming at grabbing the attention of the potential stakeholders of the PANACEA solution.
- **Press releases**. The PANACEA consortium has elaborated three different press releases as a part of the dissemination and communication strategy.
- **Commercial video**. A very important dissemination (and also exploitation) asset that has been generated in the last phase of the project´s development is the PANACEA commercial video. The video, that uses graphics and animations, presents PANACEA by first identifying the problem addressed and then, introducing the solution, its main features and benefits, etc.
- **PANACEA White Paper**. Another important asset generated on the last phase of the dissemination strategy is the PANACEA whitepaper. This document is intended as a reference guide for the project, and its objective is to describe PANACEA to the interested audience, explaining the problem it addresses, introducing how the solution works, its main functionalities and added values, the expected impact, et

These assets have been generated for multiple purposes. On one hand they have contributed to raise awareness about the project and to communicate the main achievements of PANACEA, while on the other hand, they have been extremely useful to support the consortium members on the events and conferences they attended. At the same time, assets like the press releases and the project whitepaper have been a powerful tool to better explain the benefits of PANACEA and to reach a wider audience. An example of the impact generated by this is appearance of PANACEA on the online news bulleting **IT Digital Media Group**, a direct consequence of the first press release published by ATOS.

Figure 29: PANACEA on IT Television

PANACEA Website

Throughout the entire project´s lifecycle, the project´s website has been the main tool for dissemination and communication used in PANACEA.

This core element has always been at the central hub for the project´s dissemination and communication activities. Available at http://panacea-cloud.eu, this website contains the main achievements (produced deliverables, main innovations, etc.) together with the latest news (events and conferences attended, published papers and scientific articles, etc.) of the project.

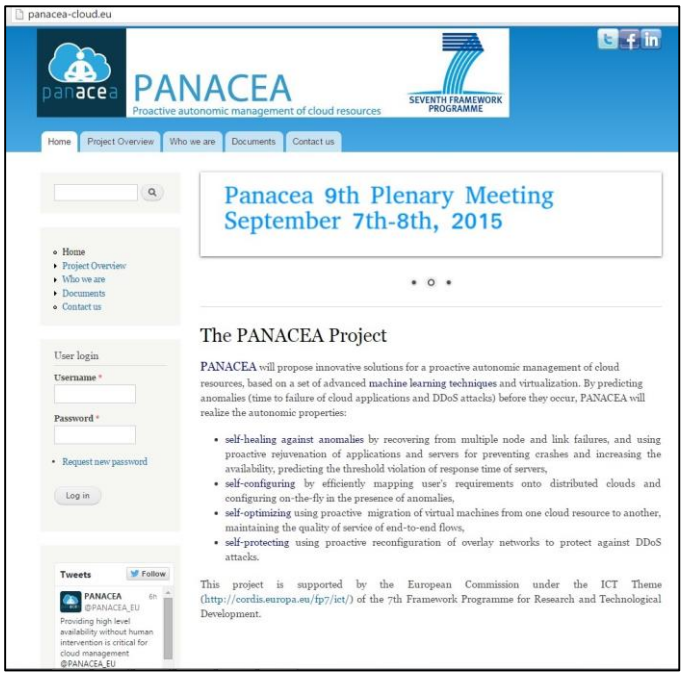

Figure 30: Project´s website

In the last months, the project´s website has been constantly updated. For instance, Twitter feed has been integrated on the main page, a new section regarding project´s innovations have been added, new deliverables and assets have been uploaded, a glossary of terms have been included, etc.

As regards to the website statistics, the results have been quite impressive.
If during the first year of the project we already overpassed the estimated KPIs (a total of 1.290 visitors for an initial estimation of **1.**000 visitors per year), these figures have grown exponentially during the second and third year: PANACEA website has far exceeded the 7.600 visitors (whereas our initial expectations were 2.500 visitors) from 122 different countries in 30 months.

Participation to events and conferences

An important element for the visibility of the project and also for the transfer of knowledge is the attendance to events and conferences. This attendance has been extremely important not only for raising awareness about PANACEA and to engage potential stakeholders, but also to promote potential collaborations with other research projects. For that reason (and for the inherent profile of the involved partners), the consortium has been quite prolific on this topic, attending to more venues than required (18 in total) among which we can find international conferences, congresses, exhibition fairs and workshops. The following table shows a recap on the different events attended by the consortium along the three years of the project lifecycle.

Title	Date/Period	Place	Audience type	Size of audience	Countries addressed	Attending Partner
Cloud Technology Days	March 25th 2015	Prague	Scientific Community	30	Europe	**UCM**
Net Futures Conference	March 25th – 26th, 2015	Brussels	Scientific Community & Industry	700	Europe	**UCM IRIANC CNRS**
OpenNebula Conference 2014	December 2nd – 4th , 2014	Berlin	Scientific Community	150	Europe	**UCM**
Software &Services, Cloud Computing Concertation Meeting	September 11th - 12th, 2014	Brussels	Scientific Community	150	Europe	**IRIANC UCM**
SeaClouds Workshop in ESOCC2014	September 2nd, 2014	Manchester	Scientific Community & Industry	100	Europe	**IMPERIAL ATOS**
EGI Community Forum 2014	May 19th – 23rd, 2014	Helsinki	Scientific Community	300	Europe	**UCM**
Closer	April 3rd – 5th, 2014	Barcelona	Scientific Community	200	Europe	**IRIANC UCM ATOS**
EU-México R&D Workshop	March 20th, 2014	Athens	Scientific Community & Industry	50	Europe & Mexico	**UCM**
FIA	March 18th – 20th, 2014	Athens	Scientific Community & Industry	400	Europe	**IRIANC UCM**
4th European Conference on Service-Oriented on Cloud Computing (ESOCC 2015)	September 15th – 17th, 2015	Taormina, Italy	Scientific Community & Industry	50	Europe	**IMPERIAL**
OpenNebula Conf. 2015	October 20th – 22th, 2015	Barcelona	Scientific Community & Industry	150	Europe	**UCM**
Software & Services, Cloud Computing Concertation Meeting	March 12th – 13th, 2014	Brussels	Scientific Community	300	Europe	**IRIANC**

CloudScape VI	February 24th–25th,2014	Brussels	Scientific Community & Industry	50	Europe	**UCM**
OCEAN Workshop	February 24th, 2014	Brussels	Scientific Community & Industry	10	Europe	**UCM**
IEEE 4th Symposium on Network Cloud Computing and Applications (NCCA 2015)	June 11th – 12th, 2015	Munich	Scientific Community & Industry	50	Europe	**IRIANC**
30th International Symposium on Computer and Information Sciences (ISCIS)	September 21st – 25th, 2015	London	Scientific Community & Industry	50	Europe	**IMPERIAL**
32nd International Symposium on Computer Performance, Modelling, Measurements and Evaluation	October 7th – 9th, 2015	Turin	Scientific Community & Industry	50	Europe	**CNRS**
IFIP Networking 2016	May 17th-19th, 2016	Vienna	Scientific Community	80	World	**CNRS**

Table 3 Events and conferences

Scientific publications

As regards to the publication of papers and scientific articles (another important pillar of the dissemination strategy), the results have been also quite impressive. While the initially set KPIs required only the generation of 3 peer-reviewed papers, the consortium has quintupled that number, reaching 16 publications among journal publications and conference proceedings. The following table summarizes all the information related to these publications:

Title	Main author	Title of periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages
Energy-QoS Trade-Offs in Mobile Service Selection	Erol Gelenbe Ricardo Lent	Future Internet	May 9th, 2013	MDPI	Online	2013	128-139
A resource-sharing game with relative priorities	J. Doncel, U. Ayesta, O. Brun B. Prabhu	Performance Evaluation	September 2014	Elsevier: Performance Evaluation Volume 79	Netherlands	2014	287 –305
Optimal Path Discovery Problem with Homogeneous Knowledge	C. Thraves-Caro, J. Doncel O. Brun	LAAS report	N° 15181	Theoretical Computer Science	-	-	-

A Software Defined Self-Aware Network: The Cognitive Packet Network	Erol Gelenbe	Proceedings of the IEEE 3rd Symposium on Network Cloud Computing and Applications	5-7 February 2014	IEEE Computer Society	Rome, Italy	2014	9 -14
A Resource-Sharing Game with Relative Priorities	*J. Doncel, U. Ayesta, O. Brun, B.J. Prabhu*	*IFIP Performance Conference*	*October 2014*	*IFIP Performance 2014, The 32nd Int. Symposium on Computer Performance, Modelling, Measurements and Evaluation*	*Turin, Italy*	*2014*	Volume 79
Principles of Pervasive Cloud Monitoring	*G. Gorbil D. Garcia E Cuesta*	*Proceedings of the 29th Int. Symposium on Computer and Information Sciences*	*27-28 October 2014*	*Springer*	*Krakow, Poland*	*2014*	*117 - 124*
Experiments with Smart Workload Allocation to Cloud Servers	*L. Wang E. Gelenbe*	*4th Symposium on Network Cloud Computing and Applications*	*11th - 12th June, 2015*	*IEEE Computer Society*	*Munich, Germany*	*2015*	-
Machine Learning for Achieving Self- Properties and Seamless Execution of Applications in the Cloud*	*P. Di Sanzo, A. Pellegrini and D. R. Avresky*	*4th Symposium on Network Cloud Computing and Applications (NCCA 2015)*	*11th -12th June, 2015*	*IEEE Computer Society*	*Munich, Germany*	*2015*	-
Shortest Path Discovery Problem Revisited (Query Ratio, Upper and Lower Bounds	*Olivier Brun Josu Doncel Christopher Thraves Caro*	*Centre pour la communication Scientifique directe*	*22nd October 2014*	*HAL*	*Online*	*2014*	*Web publication*[4]
Adaptive Dispatching of Tasks in the Cloud	*Lan Wang Erol Gelenbe*	*IEEE Transactions on Cloud Computing*	-	*IEEE Computer Society*	-	-	-
A Machine Learning based Framework for Building Application Failure Prediction Models	*A. Pellegrini P. Di Sanzo D. R. Avresky*	*Proceedings of the 20th IEEE Workshop on Dependable Parallel, Distributed and Network-Centric Systems*	*May 25th – 29th , 2015*	*IEEE Computer Society*	*IDPS*	*2015*	-

[4] https://hal.archives-ouverte.fr/hal-01076628v1

Proactive Scalability And Management of Resources in Hybrid Clouds via Machine Learning	*Dimiter R. Avresyky P. Di Sanzo Alessandro Pellegrinly Bruno Ciclani Luca Fortex*	*Proceedings of the IEEE 14th International Symposium on Network Computing and Applications*	*September 28th – 30th, 2015*	*IEEE NCA 2015*	*Cambridge, MA. USA*	*2015*	-
TAP: A Task Allocation Platform for the EU FP7 PANACEA Project	*Erol Gelenbe Lan Wang*	*4th European Conference on Service-Oriented and Cloud Computing (ESOCC 15)*	*September 15th – 17th, 2015*	*Springer*	*Taormina, Italy*	*2015*	-
Big Data Driven SMART Intercontinental Overlay Networks	*O. Brun L. Wang E. Gelenbe*	*IEEE Jour. Selected Areas in Communications – 2016 Special Issue on emerging Technologies in Communications – Big Data 2015*	-	*IEEE journal on selected areas in communicatio n*	-	*2015*	-
Scalable Self-Healing and Self-Optimizing Routing Overlays	*O. Brun H. Hassan J. Vallet*	*IFIP Networking*	*May 17th - 19th, 2016*	-	*Vienna, Austria*	*2016*	-

Table 4: Publications

As most of the project´s results have been obtained on the last year, it is expected that the number of peer-reviewed papers will increase during the next year. Besides that, due to PANACEA Alliance, the project partners will continue supporting and improving PANACEA for at least one more year, where multiple events will be targeted and more publications will be generated.

Collaboration

To create impact and foster the adoption of the PANACEA technologies, the collaboration activities performed along the project are as important (or even more) as the attendance to the aforementioned events and the publication of scientific articles. Therefore, our academic partners led by UCM, have participated in different workshops and cooperative events, as summarized in **Table 5**.

Event name	Date & Place	Projects involved	Topics
ESOCC 2015	September 15th – 17th, 2015 Taormina (Italy)	DICE, MODAClouds, CloudWave, AppHub, CAMEL, Broker@Cloud, BEACON, EUBrazil Cloud Connect, CLIPS, FrontierCities	Cloud Computing Cloud Management App Development Big Data DevOps
PANACEA – STEM collaboration	Online	STEM, PANACEA	Presentation of SMART for potential adoption by Prove

PANACEA – Prove collaboration	Online	Prove, PANACEA	Presentation of SMART for potential adoption by Prove
Online workshop AppHub - PANACEA	January 20th, 2015 Skype	AppHub, PANACEA	Presentation of AppHub Marketplace for potential adoption by PANACEA
ESOCC 2014	September 2nd – 4th, 2014 London	SeaClouds, PaaSage, MODAClouds, Artist, ECO2Cloud, ASCETiC, CoherentPaaS	Cloud computing Interoperability Cloud standards
EGI-CF 2014	May 19th – 23rd, 2014 Helsinki	EGI-InSPIRE DIRAC EGCF (former IGE)	Collaboration with EGI and infrastructure projects
EU-México R&D Workshop	March 20th, 2014 Athens	CELAR, CoherentPaaS CumuloNimbo OPTIMIS PaaSage SeaClouds Spribo	Concertation meeting organized by the European Commission to explore collaborations with related Mexican projects
FIA 2014	March 18th – 20th, 2014 Athens	More than 150 ICT-FP7 projects	Meeting with other Future internet projects
CloudScape VI	February 24th–25th,2014 Brussels	CloudSpaces OCEAN CoherentPaaS	Discussions on standardization
OCEAN Workshop	February 24th, 2014 Brussels	SPRERS mOSAIC	PANACEA in the Open Cloud Directory

Table 5: Collaboration

2.2 Exploitable foreground and plans for exploitation

The PANACEA plans for exploitation distinguishes between two different approaches, the internal and the external.

The **internal exploitation focuses on the adoption of the PANACEA outcomes by the consortium partners**. This approach has mainly focused on the exploitation of individual assets rather than the solution as a whole. Thus, the different innovations provided by PANACEA and the outcomes obtained from them, are being exploited by its owners.

As we have previously explained on different exploitation deliverables, 2 of the innovations provided by PANACEA are released as open source (the ones provided by CNRS and UCM), while the others are considered proprietary. The two open source innovations are an important **vehicle to drive project´s adoption** (or at least some of its outcomes), as it provides a direct communication channel with the open source communities (the ones related to Open Nebula and the CNRS) which could contribute to the promotion of the project and also to improve the quality of the final solution. The proprietary-licensed components generated in PANACEA and their associated knowledge, belong to their direct

owners, who will establish the exploitation mechanisms and business models that they consider appropriated for their exploitation.

These individual plans are detailed on the last deliverable of the exploitation series, D5.3 delivered on M30. Among these plans, we can distinguish between the academic and industry oriented. For the **academic partners** like CNRS, Imperial, IRIANC and UCM, the path to follow usually focuses on creating impact on the research community. This is achieved by the research activities, the discovery of innovative algorithms and techniques, attendance to conferences or the publications of scientific articles in reputed magazines, etc.
These achievements provide a prestige for the involved organizations, which can be later exploited for accessing to new research grants or obtaining new R&D contracts with external organizations. A clear example of this is Imperial, which has recently started a 12-month contract with Huawei for designing specific data networks that support internal communications for Cloud systems and Cloud Data Centres.

The aforementioned prestige can be also exploited by attracting new applicants to the new courses and PhD programs generated with the knowledge obtained along the project development. This new courses are a revenue source for the academic institutions, becoming an important factor for their sustainability and indirectly, for the sustainability of the project.

As regards to the **industrial partners**, the exploitation path focuses on using the PANACEA outcomes to improve their existing portfolio or to extend it. A clear example of this are ATOS, who has used the innovative capabilities provided by PANACEA and the knowledge generated from it, to improve the **ATOS DAaaS platform** through the use case 2. This improvement has led to the inclusion of this platform in other EIT ICT initiatives.

Another important example of industrial exploitation for PANACEA is IBM, which has improved its tool FRAPPE to later make it available as a service on **BlueMix** (IBM' s cloud platform). Prove of the interest of IBM on this are the different patents that IBM has already submitted for filing and which are described on D5.3.

Although cannot be consider a large company like ATOS and IBM, QoS Design has also achieved important milestones, which will be incorporated into the company´s portfolio.

The **enriched simulation environment** developed within PANACEA, together with the **hypervisor environment for SMART** or the **test platform**, are some of the PANACEA outcomes in which QoS Design will focus on its individual exploitation plan. These proprietary solutions will be licensed according the company internals' interests, becoming an important asset for their revenue stream.

As regards to the **external exploitation**, it mainly seeks the adoption of the project outcomes by third parties. Example of this is the reuse of specific components and knowledges generated within the project by other initiatives (like **Prove**, **STEM** or **MCloudDaaS**) as previously described on section 1.4.

Although it was discharged initially, the latest achievements of PANACEA at individual level, together with the great potential foreseen for the project, have led the consortium to create the **PANACEA Alliance**. The objective of this alliance is not to create a joint venture to exploit the project outcomes directly, but to promote the adoption of PANACEA (whether this is as a whole or through its individual components) as well as to continue evolving the solution. However, the exploitation opportunities originated by the PANACEA Alliance will be managed and distributed among the involved parties as described on the MoU included on D5.3.

3 REPORT ON SOCIETAL IMPLICATIONS

Replies to the following questions will assist the Commission to obtain statistics and indicators on societal and socio-economic issues addressed by projects. The questions are arranged in a number of key themes. As well as producing certain statistics, the replies will also help identify those projects that have shown a real engagement with wider societal issues, and thereby identify interesting approaches to these issues and best practices. The replies for individual projects will not be made public.

A General Information *(completed automatically when **Grant Agreement number** is entered.*

Grant Agreement Number:	610764
Title of Project:	PANACEA
Name and Title of Coordinator:	Dr Olivier BRUN

B Ethics

1. Did your project undergo an Ethics Review (and/or Screening)? • If Yes: have you described the progress of compliance with the relevant Ethics Review/Screening Requirements in the frame of the periodic/final project reports? Special Reminder: the progress of compliance with the Ethics Review/Screening Requirements should be described in the Period/Final Project Reports under the Section 3.2.2 *'Work Progress and Achievements'*	***No***
2. Please indicate whether your project involved any of the following issues (tick box) :	***NO***
RESEARCH ON HUMANS	
• Did the project involve children?	No
• Did the project involve patients?	No
• Did the project involve persons not able to give consent?	No
• Did the project involve adult healthy volunteers?	No
• Did the project involve Human genetic material?	No
• Did the project involve Human biological samples?	No
• Did the project involve Human data collection?	No
RESEARCH ON HUMAN EMBRYO/FOETUS	
• Did the project involve Human Embryos?	No
• Did the project involve Human Foetal Tissue / Cells?	No
• Did the project involve Human Embryonic Stem Cells (hESCs)?	No
• Did the project on human Embryonic Stem Cells involve cells in culture?	No
• Did the project on human Embryonic Stem Cells involve the derivation of cells from Embryos?	No
PRIVACY	
• Did the project involve processing of genetic information or personal data (eg. health, sexual lifestyle, ethnicity, political opinion, religious or philosophical conviction)?	No
• Did the project involve tracking the location or observation of people?	No
RESEARCH ON ANIMALS	
• Did the project involve research on animals?	No

• Were those animals transgenic small laboratory animals?	No
• Were those animals transgenic farm animals?	No
• Were those animals cloned farm animals?	No
• Were those animals non-human primates?	No
RESEARCH INVOLVING DEVELOPING COUNTRIES	
• Did the project involve the use of local resources (genetic, animal, plant etc)?	**No**
• Was the project of benefit to local community (capacity building, access to healthcare, education etc)?	**No**
DUAL USE	
• Research having direct military use	No
• Research having the potential for terrorist abuse	No

C Workforce Statistics

3. Workforce statistics for the project: Please indicate in the table below the number of people who worked on the project (on a headcount basis).

Type of Position	Number of Women	Number of Men
Scientific Coordinator	0	1
Work package leaders	1	4
Experienced researchers (i.e. PhD holders)	2	14
PhD Students	1	0
Other		

4. How many additional researchers (in companies and universities) were recruited specifically for this project?	**0**
Of which, indicate the number of men:	0

D Gender Aspects

5. Did you carry out specific Gender Equality Actions under the project?

○ Yes
⊗ No

6. Which of the following actions did you carry out and how effective were they?

	Not at all effective				Very effective
❑ Design and implement an equal opportunity policy	○	○	○	○	○
❑ Set targets to achieve a gender balance in the workforce	○	○	○	○	○
❑ Organise conferences and workshops on gender	○	○	○	○	○
❑ Actions to improve work-life balance	○	○	○	○	○

○ Other:

7. Was there a gender dimension associated with the research content – i.e. wherever people were the focus of the research as, for example, consumers, users, patients or in trials, was the issue of gender considered and addressed?

○ Yes- please specify

⊗ No

E Synergies with Science Education

8. Did your project involve working with students and/or school pupils (e.g. open days, participation in science festivals and events, prizes/competitions or joint projects)?

○ Yes- please specify

⊗ No

9. Did the project generate any science education material (e.g. kits, websites, explanatory booklets, DVDs)?

○ Yes- please specify

⊗ No

F Interdisciplinarity

10. Which disciplines (see list below) are involved in your project?

⊗ Main discipline[5]: 2.2

○ Associated discipline[5]:

○ Associated discipline[5]:

G Engaging with Civil society and policy makers

11a Did your project engage with societal actors beyond the research community? *(if 'No', go to Question 14)*

○ Yes
⊗ No

11b If yes, did you engage with citizens (citizens' panels / juries) or organised civil society (NGOs, patients' groups etc.)?

○ No

[5] **Insert number from list below (Frascati Manual).**

- ○ Yes- in determining what research should be performed
- ○ Yes - in implementing the research
- ○ Yes, in communicating /disseminating / using the results of the project

11c In doing so, did your project involve actors whose role is mainly to organise the dialogue with citizens and organised civil society (e.g. professional mediator; communication company, science museums)?	○ ○	Yes No

12. Did you engage with government / public bodies or policy makers (including international organisations)

- ○ No
- ○ Yes- in framing the research agenda
- ○ Yes - in implementing the research agenda
- ○ Yes, in communicating /disseminating / using the results of the project

13a Will the project generate outputs (expertise or scientific advice) which could be used by policy makers?

- ○ Yes – as a **primary** objective (please indicate areas below- multiple answers possible)
- ○ Yes – as a **secondary** objective (please indicate areas below - multiple answer possible)
- ○ No

13b If Yes, in which fields?

Agriculture		Energy		Human rights	
Audiovisual and Media		Enlargement		Information Society	
Budget		Enterprise		Institutional affairs	
Competition		Environment		Internal Market	
Consumers		External Relations		Justice, freedom and security	
Culture		External Trade		Public Health	
Customs		Fisheries and Maritime Affairs		Regional Policy	
Development Economic and Monetary Affairs		Food Safety		Research and Innovation	
Education, Training, Youth		Foreign and Security Policy		Space	
Employment and Social Affairs		Fraud		Taxation	
		Humanitarian aid		Transport	

13c If Yes, at which level? ○ Local / regional levels ○ National level ○ European level ○ International level	
H Use and dissemination	
14. How many Articles were published/accepted for publication in peer-reviewed journals?	**16**
To how many of these is open access[6] provided?	**0**
How many of these are published in open access journals?	
How many of these are published in open repositories?	
To how many of these is open access not provided?	**16**
Please check all applicable reasons for not providing open access:	
❑ publisher's licensing agreement would not permit publishing in a repository ❑ no suitable repository available ❑ no suitable open access journal available ■ no funds available to publish in an open access journal ❑ lack of time and resources ❑ lack of information on open access ❑ other[7]: ……	

15. How many new patent applications ('priority filings') have been made? *("Technologically unique": multiple applications for the same invention in different jurisdictions should be counted as just one application of grant).*		**5 (3 filled 2 being drafted)**
16. Indicate how many of the following Intellectual Property Rights were applied for (give number in each box).	Trademark	**0**
	Registered design	**0**
	Other	**0**
17. How many spin-off companies were created / are planned as a direct result of the project?		**0**
Indicate the approximate number of additional jobs in these companies:		

18. Please indicate whether your project has a potential impact on employment, in comparison with the situation before your project:			
❑	Increase in employment, or	❑	In small & medium-sized enterprises
❑	Safeguard employment, or	❑	In large companies
❑	Decrease in employment,	❑	None of the above / not relevant to the project
☒	Difficult to estimate / not possible to quantify		

[6] **Open Access is defined as free of charge access for anyone via Internet.**

[7] **For instance: classification for security project.**

19. For your project partnership please estimate the employment effect resulting directly from your participation in Full Time Equivalent (*FTE = one person working fulltime for a year*) jobs:	*Indicate figure:*
Difficult to estimate / not possible to quantify	■

I Media and Communication to the general public

20. As part of the project, were any of the beneficiaries professionals in communication or media relations?

○ Yes ⦿ No

21. As part of the project, have any beneficiaries received professional media / communication training / advice to improve communication with the general public?

○ Yes ⦿ No

22 Which of the following have been used to communicate information about your project to the general public, or have resulted from your project?

■	Press Release	■	Coverage in specialist press
❑	Media briefing	❑	Coverage in general (non-specialist) press
■	TV coverage / report	❑	Coverage in national press
❑	Radio coverage / report	❑	Coverage in international press
■	Brochures /posters / flyers	■	Website for the general public / internet
❑	DVD /Film /Multimedia	■	Event targeting general public (festival, conference, exhibition, science café)

23 In which languages are the information products for the general public produced?

❑	Language of the coordinator	■	English
■	Other language(s)		

***Question F-10*:** Classification of Scientific Disciplines according to the Frascati Manual 2002 (Proposed Standard Practice for Surveys on Research and Experimental Development, OECD 2002):

FIELDS OF SCIENCE AND TECHNOLOGY

1. NATURAL SCIENCES

1.1 Mathematics and computer sciences [mathematics and other allied fields: computer sciences and other allied subjects (software development only; hardware development should be classified in the engineering fields)]

1.2 Physical sciences (astronomy and space sciences, physics and other allied subjects)

1.3 Chemical sciences (chemistry, other allied subjects)

1.4 Earth and related environmental sciences (geology, geophysics, mineralogy, physical geography and other geosciences, meteorology and other atmospheric sciences including climatic research, oceanography, vulcanology, palaeoecology, other allied sciences)

1.5 Biological sciences (biology, botany, bacteriology, microbiology, zoology, entomology, genetics, biochemistry, biophysics, other allied sciences, excluding clinical and veterinary sciences)

2 ENGINEERING AND TECHNOLOGY

2.1 Civil engineering (architecture engineering, building science and engineering, construction engineering, municipal and structural engineering and other allied subjects)
2.2 Electrical engineering, electronics [electrical engineering, electronics, communication engineering and systems, computer engineering (hardware only) and other allied subjects]
2.3. Other engineering sciences (such as chemical, aeronautical and space, mechanical, metallurgical and materials engineering, and their specialised subdivisions; forest products; applied sciences such as geodesy, industrial chemistry, etc.; the science and technology of food production; specialised technologies of interdisciplinary fields, e.g. systems analysis, metallurgy, mining, textile technology and other applied subjects)

3. MEDICAL SCIENCES
3.1 Basic medicine (anatomy, cytology, physiology, genetics, pharmacy, pharmacology, toxicology, immunology and immunohaematology, clinical chemistry, clinical microbiology, pathology)
3.2 Clinical medicine (anaesthesiology, paediatrics, obstetrics and gynaecology, internal medicine, surgery, dentistry, neurology, psychiatry, radiology, therapeutics, otorhinolaryngology, ophthalmology)
3.3 Health sciences (public health services, social medicine, hygiene, nursing, epidemiology)

4. AGRICULTURAL SCIENCES
4.1 Agriculture, forestry, fisheries and allied sciences (agronomy, animal husbandry, fisheries, forestry, horticulture, other allied subjects)
4.2 Veterinary medicine

5. SOCIAL SCIENCES
5.1 Psychology
5.2 Economics
5.3 Educational sciences (education and training and other allied subjects)
5.4 Other social sciences [anthropology (social and cultural) and ethnology, demography, geography (human, economic and social), town and country planning, management, law, linguistics, political sciences, sociology, organisation and methods, miscellaneous social sciences and interdisciplinary , methodological and historical S1T activities relating to subjects in this group. Physical anthropology, physical geography and psychophysiology should normally be classified with the natural sciences].

6. HUMANITIES
6.1 History (history, prehistory and history, together with auxiliary historical disciplines such as archaeology, numismatics, palaeography, genealogy, etc.)
6.2 Languages and literature (ancient and modern)
6.3 Other humanities [philosophy (including the history of science and technology) arts, history of art, art criticism, painting, sculpture, musicology, dramatic art excluding artistic "research" of any kind, religion, theology, other fields and subjects pertaining to the humanities, methodological, historical and other S1T activities relating to the subjects in this group]

REFERENCES

[1] Zack Whittaker. *Amazon Web Services Suffers Partial Outage* http://www.zdnet.com/article/amazon-web-services-suffers-partial-outage/

[2] *10 Devastating Outages and Failures of Major Brands in 2011* http://www.evolven.com/blog/2011-devastating-outages-major-brands.html

[3] D. Avresky, O. Brun, E. Dekel, D. Garcia Perez, G. Gorbil, E. Huedo and A. Rachdi, "D4.1: Description of Feasible Use Cases", public deliverable of FP7 PANACEA project, February 2015.

[4] I. Febles, "D5.1: Market Analysis, Business Models and Value Chain", public deliverable of FP7 PANACEA project, February 2015

[5] Guide to Financial Issues relating to FP7 Indirect Actions, http://www.eurosfaire.prd.fr/7pc/doc/1364575718_financialguide_18_03_2013_en.pdf

[6] Gartner, "Gartner Says Cloud Computing Will Become the Bulk of New IT Spend by 2016 », http://www.gartner.com/newsroom/id/2613015

[7] Market Research Media, "Global Cloud Computing Market Forecast 2015-2020 ", http://www.marketresearchmedia.com/?p=839

[8] Cloud Survey Results, PC Connection, 2013.

[9] J. Gray and D. P. Siewiorek, "High-Availability Computer Systems," IEEE Computer 24, 39 – 48, September 1991.

[10] M. Sullivan and R. Chillarege, "Software Defects and Their Impact on System Availability—A Study of Field Failures in Operating Systems," Proceedings of the 21st IEEE International Symposium on Fault-Tolerant Computing, , pp. 2–9, 1991.

[11] Y. Huang, C. Kintala, N. Kolettis, and N. D. Fulton, "Software Rejuvenation: Analysis, Module and Applications," Proceedings of the 25th Symposium on Fault Tolerant Computer Systems, pp. 381–390, Pasadena, CA, June 1995.

[12] Saucez, Damien ; Iannone, Luigi ; Bonaventure, Olivier ; Farinacci, Dino, Designing a Deployable Future Internet: the Locator/Identifier Separation Protocol (LISP) case IEEE Internet Computing, December 2012.

[13] Host Identity Protocol (HIP) Architecture, RFC4423, https://tools.ietf.org/html/rfc4423

[14] Jesse Gross (September 2013). "Programmable Networking with Open vSwitch", LinuxCon, November 24, 2013.

[15] C. Labovitz, G. R. Malan, and F. Jahanian. 1998. Internet routing instability. IEEE/ACM Trans. Netw. 6, 5 (October 1998), 515-528.

[16] M. Dahlin, B. Chandra, L. Gao and A. Nayate, « End-to-end WAN Service Availability », In Proc. 3rd USITS, 2001.

[17] Csaba Szepesv_ari. Shortest path discovery problems: A framework, algorithms and experimental results. In Proc. of AAAI, 550-555, 2004.

[18] Olivier Brun, Josu Doncel, Christopher Thraves Caro. Shortest Path Discovery Problem, Re- visited (Query Ratio, Upper and Lower Bounds). [Research Report] LAAS-CNRS. 2014. <hal- 01076628>. https://hal.archives-ouvertes.fr/hal-01076628.

[19] E. Gelenbe. "Random neural networks with negative and positive signals and product form solution," *Neural Computation*, vol. 1, no. 4, pp. 502-510, 1989.

[20] E. Gelenbe and J. M. Fourneau. "Random neural networks with multiple classes of signals," *Neural Computation*, vol. 11, no. 4, pp. 953-963, May 1999.

[21] E. Gelenbe and S. Timotheou. "Random neural networks with synchronised interactions," *Neural Computation*, vol. 20, no. 9, pp. 2308-2324, Sep. 2008.

[22] E. Gelenbe. "Learning in the recurrent random neural network," *Neural Computation*, vol. 5, no. 1, pp. 154-164, Jan. 1993.

[23] E. Gelenbe and K. Hussain. "Learning in the multiple class random neural network," *IEEE Transactions on Neural Networks*, vol. 13, no. 6, pp. 1257-1267, Nov. 2002.

[24] D. Cheriton and R. E. Tarjan. "Finding minimum spanning trees," *SIAM Journal on Computing*, vol. 5, pp. 724-741, Dec. 1976.

[25] collectd - the system statistics collection daemon. Online: https://collectd.org/

[26] T. Oetiker. RRDtool. Online: http://oss.oetiker.ch/rrdtool/

[27] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. Schapire, "The nonstochastic multi-armed bandit problem," SIAM Journal on Computing, vol. 32, no. 1, pp. 48–77, 2002.

[28] O. Brun, L. Wang, and E. Gelenbe, "Big data for autonomic intercontinental overlays," to appear in IEEE Jour. Selected Areas in Communications (special Issue on Emerging Technologies in Communications - Big data), 2016.

[29] PANACEA Twitter feed: https://twitter.com/PANACEA_EU

[30] PANACEA LinkedIn group: https://www.linkedin.com/groups?gid=7422264

[31] PANACEA Facebook page: https://www.facebook.com/ThePanaceaProject