

ARCHITECTURE FOR A FULLY DECENTRALIZED PEER-TO-PEER
COLLABORATIVE COMPUTING PLATFORM

by

Dany Wilson

A thesis submitted in conformity with the requirements
for the degree of Master of Computer Science
Graduate Department of Computer Science
University of Ottawa

© Dany Wilson, Ottawa, Canada, 2015

Abstract

We present an architecture for a fully decentralized peer-to-peer collaborative computing platform, offering services similar to Cloud Service Provider’s Platform-as-a-Service (PaaS) model, using volunteered resources rather than dedicated resources. This thesis is motivated by three research questions: (1) Is it possible to build a peer-to-peer collaborative system using a fully decentralized infrastructure relying only on volunteered resources?, (2) How can light virtualization be used to mitigate the complexity inherent to the volunteered resources?, and (3) What are the minimal requirements for a computing platform similar to the PaaS cloud computing platform?

We propose an architecture composed of three layers: the *Network layer*, the *Virtual layer*, and the *Application layer*.

We also propose to use *light virtualization* technologies, or containers, to provide a uniform abstraction of the contributing resources and to isolate the host environment from the contributed environment. Then, we propose a minimal API specification for this computing platform, which is also applicable to PaaS computing platforms.

The findings of this thesis corroborate the hypothesis that peer-to-peer collaborative systems can be used as a basis for developing volunteer cloud computing infrastructures. We outline the implications of using light virtualization as an integral virtualization primitive in public distributed computing platform. Finally, this thesis lays out a starting point for most volunteer cloud computing infrastructure development effort, because it circumscribes the essential requirements and presents solutions to mitigate the complexities inherent to this paradigm.

Acknowledgements

I would like to thank my supervisor, Professor Stéphane Somé, for his continuous support throughout the different stages of this research effort, and for his patience while reading the various drafts. Thank you.

I would also like to thank the Professor Gilbert Arbez, for the sound advice on pursuing graduate studies, and ultimately for making this opportunity a reality.

I owe special thanks to the DLPS team at IBM Canada, especially Gordon McDonald, Jeff Jobb, Glynn Kneebone, and Eric Teutsch; for providing me with excellent learning experiences and numerous research opportunities, as part of the MITACS Accelerate Program.

Je voudrais remercier mes parents, Guy Wilson et Lucie Dupelle, ainsi que mon frère Jeff, pour avoir cru en moi et m'avoir soutenu malgré les multiples délais. Je vous dédie cette monographie, car sans vous elle n'aurait pas été possible.

Me gustaría dar las gracias a mi novia por todo el apoyo y la comprensión. Estoy eternamente agradecido, Nina.

Contents

Abstract	ii
Acknowledgements	iii
1 Introduction	1
1.1 Motivations	1
1.2 Problem Definition	3
1.3 Contributions	4
1.4 Main Assumptions	4
1.5 Outline	5
2 Background	7
2.1 Collaborative Systems	7
2.1.1 SETI@Home	8
2.1.2 BOINC: Berkeley Open Infrastructure for Network Computing . .	9
2.2 Cloud Computing	10
2.2.1 Deployment Models	11
2.2.2 Service Models	11
2.2.3 Volunteer Cloud Computing	13
2.3 Peer-to-Peer Computing	13
2.3.1 Overview	13
2.3.2 Overlay Networks	14
2.3.3 Structured Overlay Network	15
2.3.4 Unstructured	16
2.3.5 Slicing	17
2.3.6 Comparison	19
2.4 Virtualization Technologies	20
2.4.1 Overview	20

2.4.2	Full Virtualization vs. Light Virtualization	21
2.4.3	Containers	23
2.4.4	LxC	23
2.4.5	Docker	24
3	Related Work	26
3.1	Evaluation Framework	26
3.2	Cloud@Home	28
3.2.1	Architecture	29
3.2.2	Evaluation	30
3.3	Peer-to-Peer Cloud System	33
3.3.1	Architecture	34
3.3.2	Evaluation	36
3.4	Discussion	39
4	Architecture	42
4.1	Overview	42
4.2	Network Layer	43
4.2.1	The Ring	45
4.2.2	The Fellowships	47
4.3	Virtual Layer	49
4.4	Application Layer	51
4.4.1	Overview	52
4.4.2	Databases and Storages	53
4.4.3	Communication and Networking	55
4.4.4	Load Balancing and Scalability	56
4.4.5	Security	58
4.4.6	Application Deployment and Management	60
4.5	Component Interaction	62
4.6	Discussion	64
4.6.1	Collaborative Peer-to-Peer System Framework Implementation	64
4.6.2	Research Requirements	67
5	Implementation	70
5.1	Technology Used	70
5.2	Constructs	73
5.2.1	Task	73

5.2.2	ApplicationNode	75
5.2.3	The Ring	76
5.2.4	The Fellowships	76
5.2.5	Conclusions	77
5.3	Workflows and Protocols	77
5.3.1	Initializing a Node	78
5.3.2	Fellowship's Protocol	78
5.3.3	Web Protocol	79
5.3.4	Error-Handling	80
5.4	Proof of Concept: Calculator	81
5.4.1	Overview	82
5.4.2	Application Deploying Node	82
5.4.3	Contributing Node	83
5.5	Conclusion	84
6	Case Study: Multi-Document Text Summarization using Genetic Algorithms	86
6.1	Problem at Hand	87
6.1.1	Automatic Summarization	87
6.1.2	Genetic Algorithms	88
6.1.3	Extensions Proposed for MDTS	88
6.2	Translation of the Problem for this Architecture	90
6.3	Implementation Details	94
6.3.1	Application Deploying Node	94
6.3.2	Data Node	96
6.3.3	Worker Node	98
6.4	Experimentation	100
6.5	Conclusions	102
7	Discussion	104
7.1	Positioning of the System	104
7.2	Problems Encountered	106
8	Conclusion	110
8.1	Requirements Fulfilled	110
8.2	Contributions	112
8.3	Future Work	113

Appendices	115
A API Specification	115
B Code Samples	118
B.1 Configuration Files	118
B.1.1 Architecture Configuration File	118
B.1.2 Docker Configuration File	119
B.2 Calculator	120
B.2.1 Web Server	120
B.2.2 Web Protocol	122
B.3 Multi-Document Text Summarization	124
B.3.1 Web Server	124
B.3.2 Web Protocol	126

List of Figures

2.1	Ontological Representation of Cloud [YBDS08]	11
2.2	Cloud Services Separation of Responsibilities [Lew11]	12
2.3	Example of Overlay Network Virtual Connections.	15
2.4	VM vs. Containers [DIRB14]	22
3.1	Framework for Peer-to-Peer Resource Collaboration Problem [BJ13]	27
3.2	Cloud@Home Architecture	29
3.3	Cloud@Home Infrastructure Middleware [AAC ⁺ 11]	31
3.4	Peer-to-Peer Cloud System Architecture [BJ13]	34
4.1	Overview of the Proposed Architecture.	43
4.2	Network Layer Overview	44
4.3	Abstract Representation The Ring.	45
4.4	Fellowship Abstraction.	47
4.5	API Specification Overview	52
4.6	Proportional-Integral-Derivative (PID) Controller [GSL14]	57
4.7	Advertise and Discover Phases Sequence Diagram.	66
4.8	Selection and Match Phases Sequence Diagram.	66
5.1	Comparison between Programming Models. [MF13]	71
5.2	Graphical Representation of a Deferred. [MF13]	72
5.3	The Fellowship Protocol	79
5.4	Web Request Translation into Task(s).	80
5.5	Calculator’s Web Interface	82
5.6	Application Deployer Node as part of the Calculator Application.	83
5.7	Contributing Node as part of the Calculator Application.	84
5.8	Calculator Application Architecture Overview.	85
6.1	Event Flow: Incoming Request.	92

6.2	Event Flow: Processing a Request.	93
6.3	Event Flow: Returning a Response.	93
6.4	MDTS: General Workflow.	94
6.5	Application Deployer Node as part of the MDTS Application.	95
6.6	MDTS Application's Web Interface	96
6.7	Data Node as part of the MDTS Application.	97
6.8	Files table schema.	97
6.9	Results table schema.	98
6.10	Worker Node as part of the MDTS Application.	99
6.11	MDTS Application Heap Memory Profile.	101

List of Tables

3.1	Summary of the Solutions and the Requirements	41
6.1	First 8 Snapshots: Heap Memory Usage Peaking at Snapshot 8.	100
6.2	Memory Statistics Before Container Creation and Execution.	102
6.3	Memory Statistics After Container Creation with Active Container. . . .	102

Acronyms

API Application Programmable Interface. 15, 24, 45–47, 50, 54, 62–64, 66, 68–70, 79, 80, 120–122

CPU Central Processing Unit. 35, 58, 60

CSP Cloud Service Providers. 14, 15, 23, 24, 39, 41, 43, 63, 64, 66, 67, 72

DHT Distributed Hash Table. 28, 57, 76, 121, 126

DNS Domain Naming System. 117, 118

EDA Event-Driven Architecture. 74, 100, 101

GA Genetic Algorithm. 96, 98, 99

HTL Hops-To-Live. 29

HTTPS HyperText Transfer Protocol Secure. 70

IaaS Infrastructure-as-a-Service. 24, 25, 39, 44, 45, 49–51, 62, 65, 122

LxC Linux Containers. 32

MAC Mandatory Access Control. 70, 79

MADQ Multi-Attribute Dominated Query. 29, 116

NIST National Institute of Standards and Technology, U.S.A.. 22

OS Operating Systems. 14, 24, 32–35, 60–62, 78

P2PCS Peer-to-Peer Cloud System. 44, 50, 51, 53, 55, 62

- PaaS** Platform-as-a-Service. 14, 24, 25, 39, 44, 50, 51, 62, 63, 65, 68, 71, 80, 115, 120, 122, 123
- PID** Proportional-Integral-Derivative. 68, 121, 123
- PKI** Public-Key Infrastructure. 60
- QoS** Quality of Service. 39, 41, 44, 50
- RDBMS** Relational DataBase Management System. 64, 65
- REST** REpresentational State Transfer. 66
- SaaS** Software-as-a-Service. 23–25, 39
- SADQ** Single-Attribute Dominated Query. 29
- SLA** Service-Level Agreement. 39, 41, 42, 44, 68, 78
- SLO** Service-Level Objective. 68, 69
- SOA** Service Oriented Architecture. 22, 101
- SSL** Secure Sockets Layer. 69, 70
- SSO** Single-Sign-On. 70, 71
- TTL** Time-To-Live. 29
- UC** Utility Computing. 14
- URL** Uniform Resource Locator. 66
- VM** Virtual Machine. 33–35, 40, 43–46, 48, 50, 51, 61, 62, 110

Chapter 1

Introduction

This thesis presents an architecture for a fully decentralized computing infrastructure built using volunteered resources. This computing infrastructure is composed of a collection of geographically distributed commodity devices (personal computers) connected over the Internet. The applications deployed to this computing infrastructure are hosted by volunteered resources, resulting in a computing platform similar to the Platform-as-a-Service (PaaS) offered by Cloud Service Providers (CSP).

In this chapter we present the motivations that drove this research effort and a definition of the problem we are addressing in this thesis. Then, we outline our contributions, and present our main assumptions. Finally, we present the content of the different chapters included in this thesis.

1.1 Motivations

In this section we present the motivations that drove our research effort, including the current state of CSP and the current state of independent web communities.

Cloud Service Providers

CSP offer one of the most prominent computing platform online to develop and deploy applications. Cloud computing is characterized by its service provisioning model, offering to the consumer a *seemingly* "infinite pool" of resources that can be contracted and released instantly to mitigate the fluctuations in the workload. Using a business model centered around consumption-based pricing, known as Utility Computing (UC), it provides to the consumer a wide selection of service models to choose from depending on their needs. Service models vary in proportion to the management responsibilities outsourced to the CSP, from completely outsourced software to hardware infrastructure capable of hosting a whole Operating Systems (OS), effectively reducing the cost of the

on-premise IT infrastructure. But some security concerns still persist, impeding this paradigm shift despite its popularity. Privacy is a prime concern, because for a consumer to use the cloud computing infrastructure they must migrate their software assets and data to this CSP. This migration implies establishing a trust relationship between the consumer and the CSP, which includes their employees having (in)direct access to the physical infrastructure. Malicious administrators (insiders) is the other major security concern. An *Orwellian* example of this security concern is the CSP potential to exercise (in)direct censorship, either out of strategic positioning or as a tyrannical authority. This is extreme, but not completely farfetched.

For example: two applications providing similar services are hosted on the same CSP, if one of them can persuade (by any means) the CSP to favor its application, then the provider can (in)directly exercise censorship against the other application. For instance, it could prioritize the traffic of that application, or deprioritize the other application access to the computing resources (or Application Programmable Interface (API)).

Consequently, the benevolence of the individuals employed by this provider is *blindly* assumed by the consumers, because nothing forces these providers to be *completely* transparent in their operations anymore than any other private service providers. This is due to the private nature of these CSP, because they are not providing public services they are not legally bound to reveal any of their mechanisms and/or protocols enforcing the security and privacy within their organization.

Web Communities and Communities in General

Local and Web Communities are supported by a small core of dedicated individuals, and usually host applications including discussion forums, image boards, and websites. Normally, these communities rely on donations to sustain their activities, including the fees incurred to host their applications. Given that all the members of these communities possess a personal computer with Internet access, to access the application. Can their operation costs be reduced, taking into account the increase in domestic bandwidth of Internet Service Providers (ISP) and the increase of performance of personal computers?

Another interesting example of communities is the education system, where different school boards form micro-organizations grouping multiple schools together. These micro-organizations are equipped with computing resources, primarily composed of personal computers and usually accessed during office (school) hours. In some cases they remain idling through the night, or are simply turned off. Could recycling these resources effectively reduce the cost of operations, either by generating income and covering their

initial purchase cost; or by extending their usefulness?

A commonality is shared within these two examples of communities, a potential desire to recycle their computing resources to mitigate their cost of operation; or to extend the usefulness of their resources. Recycling computing resources exhibits many incentives, including financial restitution of the investment to acquire the resources. But more importantly, in this era of consumerism, prolonging their usefulness is an important incentive, especially given the eventual scarcity of the resources used for mass production. In a similar line of thought, recycling computing resources can help intellectual advancement by offering them to different scientific enterprises to run experiments or computationally intensive tasks. There are benefits to provide to smaller communities the technology allowing them to recycle their computing resources and make them profitable, for themselves or for the greater good. We also see potential benefits in solutions catering to all types of requirements, be it in terms of privacy and data ownership or simply to enable a private organization to be part of a fair competitive market and avoid possible censorship from a tyrannical oligarchy.

1.2 Problem Definition

In this section we present the definition of the problem addressed in this thesis. We define this problem by formulating a series of requirements that a solution should fulfill.

From the motivations, presented in the previous section, we can define a list of essential requirements for a candidate solution:

Requirement 1 Given a collection of heterogeneous commodity personal computers, we need to be able to use this system to deploy an application or multiple applications.

Requirement 2 Using this system should not force any third-party to provide services (except domain hosting, everything else should be accomplished by the collection of contributing resources), and thus be self-contained as much as possible.

Requirement 3 The system should be resilient to resources failing, and/or leaving. Thus, it should be fault-tolerant and should not introduce a single point of failure, for security purposes to resist DoS type of attacks and possible censorship.

Requirement 4 The system should not require from its user to acquire any special (or dedicated) equipment, but rather recycle the current resources available.

Sub-Requirement 4.1 The memory footprint should be small, as to be able to construct this system from lower-end devices.

Requirement 5 The system should be able to provide scalability to the application deployed, or the ability to dynamically adjust the amount of resources available and account for the fluctuation in demand. As a consequence, it must provide dynamic membership capabilities to all applications.

1.3 Contributions

Within the context of this thesis we propose a solution for this problem in the form of an architecture and with this proposition we make the following contributions:

Contribution 1 Propose a fully-decentralized collaborative system that provides a web computing platform.

Contribution 2 Propose a candidate API that addresses the minimal requirements of this computing platform, but more generally one that is suitable for (PaaS) Platform-as-a-Service.

Contribution 3 The system we propose make use of light virtualization to abstract the specificities of the contributing resources and to isolate the hosting environment from the contributing environment.

Contribution 4 The system that we propose is minimally intrusive and leaves a small memory footprint, enabling it to be used with lower-end computers and even micro-computers.

1.4 Main Assumptions

In this section we present the main assumptions made in the context of this research effort. We present the implications of these assumptions in terms of the **privacy** of the system and its users; of the **data integrity** and the obvious **threats**.

Our most important assumption is the **semi-trusted** environment in which the system is designed to operate. In the context of this research we define a *semi-trusted* environment as follows:

A semi-trusted environment requires both parties (consumer and contributor) to know and trust each other *a-priori*, but they are required to communicate over an untrusted network infrastructure, i.e., the Internet.

This assumption relaxes the requirements to maintain the system and the user's privacy, but operating using public untrusted communication medium still poses possible threats. The threats do not originate from the users themselves, rather it originates from the possible *man-in-the-middle* attack, or intercepting and tampering of the data. We will discuss how to ensure privacy using security mechanisms and privacy enhancing technologies in Section 4.4.5.

This assumption has consequences with respect to the data integrity, which relates to preserving and verifying the consistency of the data, as well as its accuracy. In the context of this system, the users are responsible to maintain and verify data integrity at the application-level, which simplifies this infrastructure from a design perspective but also provides a flexible solution that can be tailored to any required level of data integrity. For example, an application may require to verify the computation for possible (in)voluntary data corruption, i.e., bit rot or computation cheating, whereas another may only require to verify the content of a message. In the former case, it is possible to resort to sophisticated solutions tailored to a specific problem domain [?]; the latter may only require a simple error-correction code, given that the message is certified and signed (using a Public-Key Infrastructure (PKI)).

Finally, this system operates in an environment for which the trust-relationship between the consumer and contributor is previously established, but still it requires diligence from the application developer to ensure that the data transported using the Internet, as a communication medium, was not tampered with. We designed this architecture to provide a flexible solution that can be easily adapted to suit various degrees of security and privacy requirements.

1.5 Outline

In this section we present the outline of this thesis and the content of each chapters.

Chapter 2 Background

This chapter presents the background material concerning Collaborative Systems, the basics of Cloud Computing, the basics of Peer-to-Peer technologies, and current Virtualization technologies.

Chapter 3 Related Work

This chapter presents two similar systems, namely Cloud@Home and P2P Cloud System. It present an evaluation of these systems with respect to our requirements and the framework proposed by [BJ13] for peer-to-peer collaborative systems.

Chapter 4 Architecture

This chapter presents the architecture we designed and the design decisions for each layer of this architecture, namely: the Network Layer, the Virtual Layer and the Application Layer.

Chapter 5 Implementation

This chapter presents the current implementation of our architecture and the technologies used. We then discuss the implementation-choices and the relevant underlying mechanisms that provides the functionalities requested in the requirements. We conclude with the presentation of a proof of concept: a calculator.

Chapter 6 Use Case: Multi-Document Text Summarization using GA

This chapter presents a use-case, namely Multi-Document Text Summarization using Genetic Algorithm, and how to translate it into a distributed web application. We analyze the problem at hand, and walk-through the thought process of re-factoring an existing application into a distributed version compatible with our architecture.

Chapter 7 Discussion

This chapter presents the positioning of this system in the distributed computing platform landscape and the main problems encountered throughout this thesis.

Chapter 8 Conclusions and Future Work

This final chapter assesses the extent to which we fulfill the research mandate, in the scope of our requirements and contributions. It also present the future work the areas to focus on from a short-term and from a long-term perspective.

Chapter 2

Background

In this chapter we present an overview of the material required to understand the context of this thesis. We first present Collaborative Systems, then Cloud Computing, followed by an overview of the current Peer-to-Peer technologies and we finish by glancing at the latest technologies in terms of Virtualization.

2.1 Collaborative Systems

In this section we present collaborative systems and outline the differences between two deployment models: public and private. Then we focus on public collaborative systems and present an example, *SETI@Home*. We conclude this section with the presentation of a middleware used to develop collaborative systems, namely *BOINC*.

Collaborative Systems can have different meanings in different contexts. For example, in a Business context, they often refer to Groupware, or software enabling teams to cooperatively work on a project. Where the word *collaborative* is used to describe the interactions between the users of the system [CS99]. Whereas in the context of distributed systems, collaborative systems refers to a system belonging to a sub-class of distributed computing known as *volunteer computing* (or public-resource computing), for which a collection of volunteered resources accomplishes a common task [AF06]. Here, the word *collaborative* is used to describe the compositional structure of the system, being composed of disparate components collaborating together to perform a computational task.

In general, *volunteer computing* and *public-resource computing* are interchangeable when referring to this sub-class of distributed computing. We prefer the latter over the former, since it exhibits the public nature of this type of computing explicitly. It is possible to draw further distinctions between collaborative systems, especially in the context of distributed systems, relating to the deployment model. A collaborative system that is

deployed on a grid or a cluster, has a *private* deployment model, hence it can be referred to as a *private collaborative system*. Whereas a collaborative system deployed using public resources can be referred to as a *public collaborative system*. Both types share some similarities in structure, but public collaborative systems have different requirements related to the use of volunteered resources in their infrastructure, rather than dedicated resources. We can formulate the differences as follows:

1. Resources are unreliable: they may leave or fail at any moment and the computations performed may be incorrect.
2. Quality of the resources are usually on the lower end of the spectrum (commodity hardware), compared to high performance computing clusters.
3. Control over the resources is not centralized, and participation is incentive-driven.

These differences illustrate challenges that arise from leveraging volunteered resources in a public environment and are not as present in a private (controlled) environment. This is particularly relevant with respect to the types of problems applicable to public collaborative system platforms, as we will illustrate in Section 2.1.1.

An extensive body of work already exists on the subject, for more information on public collaborative systems in general see [And03] [AF06], and for more information on private collaborative systems, but more specifically on grid computing see [BFH03].

2.1.1 SETI@Home

SETI@Home [ACK⁺02] is a notable example of a public collaborative system, built using the BOINC system framework [And04]. It is used to analyze radio transmissions for extra-terrestrial intelligence across startlingly different frequency domains. Its architecture consists of a central server which distribute work units to participants in order to perform the computations, then the participants return the result to the server and request another work unit. Due to the unreliability of the publicly volunteered computing resources and the unreliability of the network supporting the communication, it is necessary for the design to factor these in, as stated in Section 2.1.

1. Accuracy of the computations is ensured by dispatching the work units to several different participants and by using a consensus mechanism in which a majority of the results returned by the participants must agree.
2. The nature of the problem domain mitigates the unreliability of the network, because it has a high computation-to-data ratio. Work units are composed of small

data-pieces (350kb) providing considerable workload and yielding smaller result data-pieces (1kb).

Consequently, this architecture minimizes the possibility of bottlenecks at the central server while preserving a centralized structure.

This is one of the most successful project of this nature, in terms of aggregate computing power achieving upwards of 660 teraFLOPs [Wik15]. The success lies in how well their problem was tailored for distributed computations, or how easily parallelizable it was. Generally speaking, problems for which the dataset can be segregated into independent subsets of the problem are best suited for this paradigm.

2.1.2 BOINC: Berkeley Open Infrastructure for Network Computing

BOINC [And04] is a middleware system that enables researchers, with limited computer knowledge, to easily create and deploy public-resource computing projects such as SETI@Home. A typical instance of a BOINC project, consists of 4 major components:

- *Master URL* of the project presents a landing page to register, contribute, and track the progress of the project.
- *Data Server* is responsible for uploading and provisioning the data to and from the participants.
- *Scheduling Server* is responsible for handling the incoming RPCs from the participants.
- *Client Application* enables the participants to connect to the server and request work units, but it also offers a graphical interface to represent any progress made.

The ease of use lies in the abstraction of the underlying housekeeping mechanisms required to coordinate the agglomeration of contributing nodes. But also by abstracting the underlying logic required to distribute work units and retrieve the results, and presenting to the researchers a fill-in-the-blanks type of user experience.

The simplicity of setting up and understanding the infrastructure provided by BOINC is the key to its success. Essentially every scientific enterprise resorting to public-resource computing have similar requirements. By creating a middleware that fulfills these requirements out of the box, it enables the researchers to devote more time to significant research problems rather than to computer-related technicalities.

More collaborative projects are created with ease every year using this infrastructure, and this renders the use of public-resource computing more apprehensible from a non-technical standpoint. Ultimately, BOINC removes the technological and technical barriers that possibly impedes scientific data *crunching* when resorting to public computing resources.

2.2 Cloud Computing

In this section, we present the cloud computing paradigm, by identifying the different deployment models and service models it offers. Then, we present *volunteer cloud computing*, which is the intersection between public-resource computing and cloud computing.

Cloud computing is the natural evolution of Service Oriented Architecture (SOA), the Web 2.0 and the virtualization technologies, resulting in a paradigm using the *network as a platform* to provide a variety of services directly to the users. This effectively reduces the gap between content producers and content consumers, by augmenting the availability of the services and reducing the cost of contracting such services. Cost reduction is the result of offering services using a consumption-based pricing, or as a utility, completely removing the cost of acquiring and maintaining the IT infrastructure required to provide these services. A more formal definition of the characteristics of cloud computing is provided by the National Institute of Standards and Technology, U.S.A. (NIST) [MG11]:

1. **On-Demand Self-Service** : provides the ability to a user to configure his services automatically through the infrastructure portal by himself.
2. **Broad Network Access** : enables clients to connect to the cloud over the Internet.
3. **Resource Pooling** : provides the ability to clients to abstract the underlying specificities of the resources, and to simply access an infinite pool of homogeneous resources (virtual or physical).
4. **Rapid Elasticity** : provides the ability to augment the amount of resources to respond to the fluctuations in the workload.
5. **Measured Service** : services can be measured or monitored in a manner that is transparent to both the provider and the consumer.

2.2.1 Deployment Models

A cloud computing infrastructure can manifest itself using one of 4 different deployment models [MG11]. A *deployment model* identifies the intended consumer and intended producer of the services offered by the CSPs. *Public* clouds are the most popular deployment model, intended to be used by the general public they are offered by privately owned companies, or CSP (like Google, or Amazon). Another popular deployment model, are *private* clouds, and they are intended to be used exclusively by one entity or organization, which maintain or not the cloud infrastructure, in which case they delegate to a 3rd-party CSP. The two previous deployment models represent the extrema on the continuum of possible deployment models. In the middle of these two lies a variant which provides exclusive access to organizations sharing common interests, this deployment model refers to *community* clouds. The fourth, and last deployment model, are *hybrid* clouds, where the cloud infrastructure is composed of several sub-clouds, be it private or public. A common use-case for this deployment model is an organization that requires the security provided by private clouds and requires the scalability and availability of public clouds. Thus, the organization can use a private cloud to store sensitive information and use a public cloud to leverage a business intelligence application to manipulate their information, using a Software-as-a-Service (SaaS).

2.2.2 Service Models

An ontological representation of cloud computing has been proposed [YBDS08], it represents a categorization of the services generally offered by CSP.

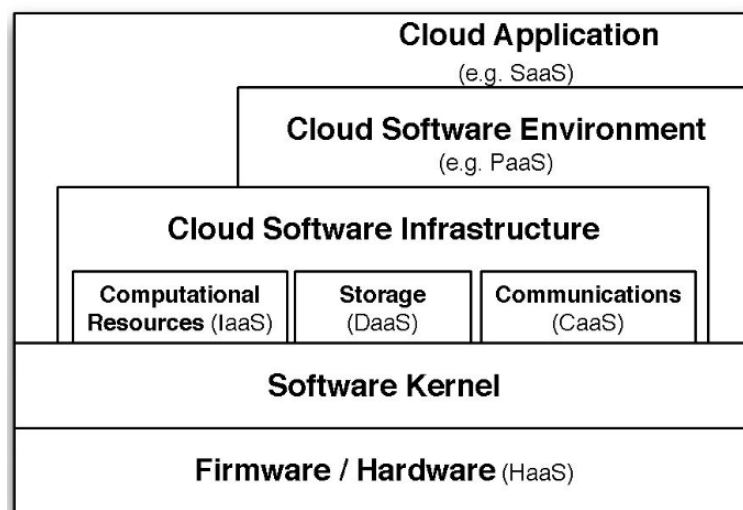


Figure 2.1: Ontological Representation of Cloud [YBDS08]

In Figure 2.1, several service models are presented for completeness, but within the scope of our research only three service models are relevant, namely: Infrastructure-as-a-Service (IaaS), PaaS and SaaS. The differences between each service model is better understood when it is illustrated by the separation of responsibilities between the CSP and the consumer.

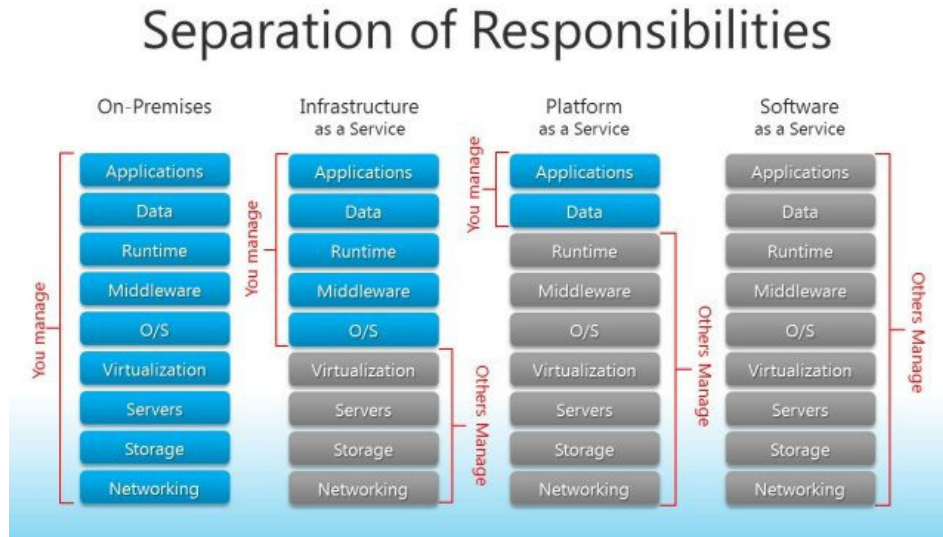


Figure 2.2: Cloud Services Separation of Responsibilities [Lew11]

Figure 2.2 presents this separation of responsibilities for each of these service models.

The IaaS service model imposes the majority of the responsibilities above the virtualization layer to the consumer, whereas the CSP is responsible for providing the physical and virtualization layers. Compared to the other service models, it grants greater flexibility to the consumer, using a virtual machine image to encapsulate the *complete* executing environment, from the OS to the applications and everything in between. An example of such a service model is *Amazon EC2* [Inc15a].

PaaS is not as flexible, but it is easier to configure and use. It provides a set of abstractions (in the form of an API where only the Application layer and the Data layer are presented) for the consumer to use to develop applications. The *Google App Engine* is a popular example of PaaS [Inc14b].

Finally, SaaS consists of providing an application as a service, where everything else is the responsibility of the CSP. Using this service model the consumer accesses a software through the CSP, and is not required to own a copy, but rather leverages its functionalities as services. An example of a SaaS provider is *Salesforce*, which offers a multitude of business applications as SaaS [sal15].

2.2.3 Volunteer Cloud Computing

Volunteer cloud computing is an emerging paradigm. Conceptually, it is the intersection of the public-resource computing paradigm and the cloud computing paradigm, it can also be defined as the 5th deployment model of cloud computing. This new deployment model consist of constructing a cloud computing infrastructure by relying solely on (volunteered) public-resources. It is a fairly recent idea, as it is not mentioned in earlier publications such as [RCL09], but it is in later publications such as [ZCB10].

This paradigm suffers from different challenges compared to the traditional cloud computing paradigm [ASK15], which are similar to the challenges of public collaborative systems. Volunteer cloud computing provides a cloud computing infrastructure built using volunteered computing resources, whereas public collaborative systems provides an infrastructure to perform distributed computationally intensive tasks. The dichotomy of their purposes, is what really differentiates volunteer cloud computing from public collaborative systems.

Many different research efforts that attempts to provide this type of infrastructure exists with varying degrees of success, and adopting different approaches while addressing these challenges. Note that some research efforts, such as [CDPS10c], target the full spectrum of service models: IaaS, PaaS, and SaaS. Whereas, other focuses on only one service model, IaaS [BMT12] [CW09]. Finally, others adopt a completely different approach, such as building a transparent infrastructure using peer-to-peer interception techniques [MGLPPJ13].

2.3 Peer-to-Peer Computing

In this section we present an overview of peer-to-peer computing, we introduce the concept of overlay networks, and then detail the differences between two types of topologies. Finally, we present a comparative analysis of the topologies in the form of a summary table.

2.3.1 Overview

Peer-to-peer computing has many different characteristics that makes it an interesting network primitive for a public environment, such as the Internet, when constructing a distributed system.

A system built using a peer-to-peer structure is referred to as a *peer-to-peer system*. Essentially, peer-to-peer systems are defined as a distributed system where every node in

the system is at the same time, a consumer of the services offered by the system and a producer, a server and a client (servent). The cost of operations is then amortized and distributed among the participants, because every participant in the system assumes the cost of operating their own computers, where the aggregate cost of operating the system is the combination of all the costs incurred for operating the participants computers. Peer-to-peer systems are generally decentralized, removing any single point of failure and resulting in systems that are very resilient to the failures of the nodes composing its infrastructure. Peer-to-peer systems are also scalable, inasmuch as more nodes are available to join the system, theoretically *ad infinitum*. This underlines that there are no explicit scalability limitations for this type of system, given that the business logic it implements also does not explicitly prevents it (by design). The equality amongst nodes is usually achieved by creating a decentralized network, where each node assume equal responsibilities in terms of routing and discovery of other nodes within the network. Generally, there are no central servers and all the nodes in the system are equal, but practically some system relies on a bootstrapping server to ease joining new nodes to the system, which is the case for mobile networks [Olk06].

In order to join a system, a node is required to know at least one other node in it. Initially, if there are no other nodes in the system, the first node to join will become the only node to contact to join this system. Subsequently, the next node that wishes to join this system contacts this node, but the third node has the choice to contact either the first or the second node, and so on. Once a node has joined the system, it usually has only a partial view of the entire underlying network and in order to contact any node not contained within this view, it must interact with neighboring nodes for indications on how to proceed. This co-operative location mechanism differs in implementation, but conceptually a node requires the assistance of the other nodes, in some ways, to navigate the underlying network in its entirety. Consequently, peer-to-peer systems usually construct a virtual *overlay network* on top of the physical underlying network, to mitigate the complexities inherent to this decentralized architecture [MKL⁺02] [Bar01].

2.3.2 Overlay Networks

Overlay networks are a fundamental primitive for peer-to-peer systems and can be defined as virtual networks superposed on a physical network (the Internet). The virtual network represents virtual connections between the different nodes in the physical network. It abstracts the underlying physical connections, and exposes the logical (virtual) connections between the nodes of the system.




Figure 2.3: Example of Overlay Network Virtual Connections.

Figure 2.3 exemplifies the concept of virtual connections using an overlay network. For example: If **Node A** is connected to **Node B**, which is connected to **Node C**, which is finally connected to **Node D** in the underlying physical network topology. Then, it is possible to express the indirect connection between **Node A** and **Node D** as a direct virtual connection in the virtual overlay network. Not only, overlay networks abstract away the details of the underlying network, but also allows nodes to communicate between them using this virtual topology.

Overlay Network topologies are categorized relative to their structure [LCP⁺05]. In the next sub-sections we present the distinction between a *structured overlay network* and a *unstructured overlay network* respectively.

2.3.3 Structured Overlay Network

What characterizes the structured overlay networks is the fact that they are constructed by organizing the peers into a structured graph.

An abstraction known as *keyspace* is used to organize the participating peers into a structure. Each peer is assigned a portion of the keyspace and is responsible for locating and indexing keys within it. Partitioning of the keyspace is done according to the *keyspace partitioning scheme*, which dictates the structure of the resulting network. Some keyspace partitioning schemes can produce a ring topology, see [SMK⁺01]; whereas others can produce various different topologies, including a tree-based topology, see [JOV05]

As a consequence of its static structure, this type of overlay network is not suited for complex multi-attribute queries. Because the organization of the keys is structured according to a single metric, it only supports single-attribute dominated queries. These queries are deterministic, because the attribute for which the queries are tested against corresponds to the key and will return upon exact or partial match. Ultimately these networks are built around the idea of being able to efficiently locate any key within the network.

Distributed Hash Table (DHT) is a class of structured overlay networks. It consists of creating an overlay network by dividing the keys of an associative array (hash table) among the peers, then each peer is responsible for a portion of the associative array. A key in this context represents either a node or an entry in the DHT. To retrieve a specific key, a node queries the other nodes to find which one is responsible for this key, or finds the node which is responsible for the portion of the associative array closest to the key. If neither has the key, the query returns because there are no value associated with this key, hence the deterministic nature of DHT. Many implementation of DHT exists [Sar10] [LCP⁺05] [BJ13].

Structured overlay networks are constructed to fulfill a specific functionality, such as storing data in a distributed fashion, and they are very good at it. But they are not as versatile as their unstructured counterparts in terms of applications, resistance to failures, and querying abilities [BJ13] [LCP⁺05].

2.3.4 Unstructured

Unstructured Overlay Networks differ with respect to their construction, resulting in a unstructured topology, namely a flat or hierarchical random graph. More importantly, there are no relations between the topology and the partitioning of the keyspace, since technically the keyspace is irrelevant in this context.

Rather than utilizing a sophisticated partitioning mechanism for the graph, each peer connect to another peer in the network, and by performing cyclic exchanges of information among them (generally in pair-wise fashion) they update their views of the network. Peers only have a local view of the overlay network, but they perform cyclic exchanges to achieve local convergence between their views of the network, resulting in global convergence of the network topology. As a result of global convergence (or network-wide convergence), several self-* properties emerges such as: self-configuring, self stabilizing and self-healing properties [JVG⁺07] [Bir07]. *Self-configuration* refers to the ability to autonomously configure the deployment of a system, or to respond to changes in the topology of the composing components [KC03] [BG09]. The *self-stabilization* property refers to the ability for a system to start in any arbitrary configuration and eventually converge to a desired configuration [BG09] [Dol00]. Whereas, the *self-healing* property provides autonomous detection, diagnosis and repair of localized problems [KC03], but also autonomous component-level failure recovery [BG09]. Maintaining the overall health of the system, can be referred to as the *survivability*, which is the prime objective of self-healing components [PD11] [GSRU07]. Given the emergence of these properties, this type

of overlay network is very robust in highly dynamic environments, such as the Internet [Bir07].

Generally, the routing is done by performing random-walks or flooding the network, and consequently querying is not deterministic. Since the structure of the graph, or the lack thereof, is not dependent on any particular arrangement of the peers, the network supports complex queries because of the arbitrary nature of the routing mechanism, which can reinterpret the network accordingly. This is commonly referred to as *slicing* and we present it in the following subsection.

Because of the routing mechanism, location of specific data managed by a peer in the network is rather difficult if it has not been replicated on several nodes. Replication schemes usually target the most popular data in the network, and replicates it across several nodes to ensure availability. Thus, rarely queried data is unlikely to be found by queries if it is managed by a single or very small collection of peers. Because queries generally have a Time-To-Live (TTL), which can be represented in terms of hops or Hops-To-Live (HTL), after which the query will simply be dropped and no result will be returned.

Examples of unstructured overlay networks are usually based on epidemic protocols, or gossip-based protocols [RV11], for which a peer joins the network and periodically exchange its local view of the network with another, randomly selected, peer.

For more information on overlay networks, and comparisons on the different types see [LCP⁺05] and [BJ13].

2.3.5 Slicing

Slicing is a primitive in distributed systems, which dictates the querying abilities of the underlying network structure. The ability and degree to which the system can perform slicing, is relative to the complexity of the queries it supports. It can be formulated as the following:

Given a graph, representing a network of peers, can we partition it according to a set of node local-attribute(s).

Several techniques exists to solve this problem as depicted in [JK06] and [PMRS14]. Techniques varies in terms of the type and number of attributes considered, and consists of ordering the nodes according to these attributes providing different perspectives of the same collection of nodes.

Based on this definition of the slicing problem, we can now formulate the differentiation between Single-Attribute Dominated Query (SADQ) and Multi-Attribute Domi-

nated Query (MADQ). The distinction arises from the ability to query resources, according to the specifications of their attributes, whether it support querying a single attribute per query or multiple. Or we can also distinguish the two, based on the ability to slice the current network into slices (partitions) according to the(se) attribute(s).

This problem helps deciding which topology of overlay network to use, because it illustrate the querying ability provided by the different topologies.

2.3.6 Comparison

We can summarize the advantages and disadvantages of the different types of overlay networks according to different characteristics, inspired by [LCP⁺05] and [BJ13]:

Characteristics	Structured	Unstructured
Construction	Generate abstract keyspace, which is divided among the peers. Topology results in a structured graph.	Peers query a participating node and retrieve information about the network. Repeats periodically, results in a flat or hierarchical random graph.
Routing	Key-based routing, each node contacts the closest node to the key (according to its local view), and then the closest node to this node, and so on until the node responsible for the key is reached.	Perform a random-walk through the network or using flooding techniques. May not return a result (query could time-out).
Lookup	Deterministic and has a general time complexity of $O(\log n)$, where n is the number of keys in the keyspace.	Non-deterministic and is generally a best-effort attempt to locate data, popular data can be easily located due to replication. No defined boundary in terms of time complexity, uses a parameter that represent the TTL, (in seconds or hops).

Continued on next page

Continued from previous page

Characteristics	Structured	Unstructured
Join/Leave	Join: Node contacts a live node and a portion of the keyspace is attributed to it, according to its identifier (or key). Leave: Node leaves, eventually the keys will be remapped to the neighboring nodes in the network and all the neighbors table will be updated.	Join: Node contact an arbitrary live node and exchange information, and repeats periodically forever achieving a dynamic view of the network. Leave: Node leaves, and then since it won't participate in the periodic information exchange, it will be discarded from the dynamic view of the network.
Reliability/Fault-Tolerance	Resist churn and normal level of node failures.	Highly resistant to churn and very high level of node failures.
Slicing	(SADQ) Single-Attribute Dominated Queries	(MADQ) Multi-Attributes Dominated Queries
Concluded		

2.4 Virtualization Technologies

In this section we present a general overview of *virtualization*. We differentiate between *full virtualization* and *light virtualization*. Then, we present what are *containers*, and take a look a two different instances of it such as Linux Containers (LxC) and Docker containers.

2.4.1 Overview

Virtualization provides the ability to allocate the physical resources needed to accomplish a task prescribed by the software. It is *generally* achieved through the decoupling of software from hardware [Tav12].

The simplest example of virtualization that is used by a majority of OS are processes [Chi08]. Processes are isolated into virtual environments that exposes the resources as if they are the sole consumer, by abstracting the other processes away and effectively decoupling the software from the underlying hardware providing the computational resources.

The concept can be extended to the virtualization of whole OSs. The interactions from the OS with the physical hardware are done through a software abstraction layer, the resulting machine is called a Virtual Machine (VM) [SPEW11].

Within the context of distributed systems, one type of virtualization pre-dominantly exist, the *full virtualization* or desktop virtualization. Semantically it is different from *para-virtualization* and both types shouldn't be conflated. The latter requires modification to the guest operating system to comply with the interface defined to access the physical hardware. Whereas the former requires no such modifications, and the calls from the operating system to the hardware can be interpreted directly [Tav12].

An emergent type of virtualization technology, known as *light virtualization* or operating-system level virtualization, is gaining popularity. It is akin to the type virtualization used processes. It is important to draw the differences between these two types of virtualizations, in order to understand which is best suited to fulfill our requirements. We neglect para-virtualization in the context of this thesis since it is not as relevant for our research, as we do not intend to impose any modification on the OS. For more details on the basics of virtualization, see [Tav12] [BDF⁺03].

2.4.2 Full Virtualization vs. Light Virtualization

Several distinctions exist between full virtualization and light virtualization and some are advantageous, others are disadvantageous in different contexts.

Full virtualization can be defined as a virtualization technique for which the entire hardware is virtualized, providing an abstract computing base, as a software layer, on which it is possible to execute a complete OS without any modifications [BDF⁺03].

The use of VM, rather than physical machines, offers several key features that are desirable in the context of distributed systems. One of these features is the ability to clone, which enables the replication of the complete execution environment with ease. Another very desirable feature is the ability to test changes before applying them, and the ability to migrate the VMs across different hardware, which can also be done without much interruption of service. The penultimate feature of using VMs, as opposed to physical machines, is undoubtedly the ability to execute multiple VMs in parallel on a single physical machine. It grants the ability to have multiple different execution environments without having to dedicate a physical machine to each. When applied to servers, this feature is referred to as *server consolidation* [Tav12].

A VM is controlled by an *hypervisor*, which is responsible for orchestrating the VMs access to the underlying physical resources of the host, and it is also known as Virtual

Machine Monitor (VMM). It is possible to distinguish between two types of hypervisors. The distinction depends on whether or not there is an OS running the hypervisor (Type-2) or if it is directly interfacing with the hardware (Type-1) [PG74].

VirtualBox [Wat08], is a prime example of Type-2 hypervisor. Figure 2.4 shows the difference between full virtualization (on the left) and light virtualization (on the right). It illustrates how each type of virtualization interacts with the hardware through the OS.

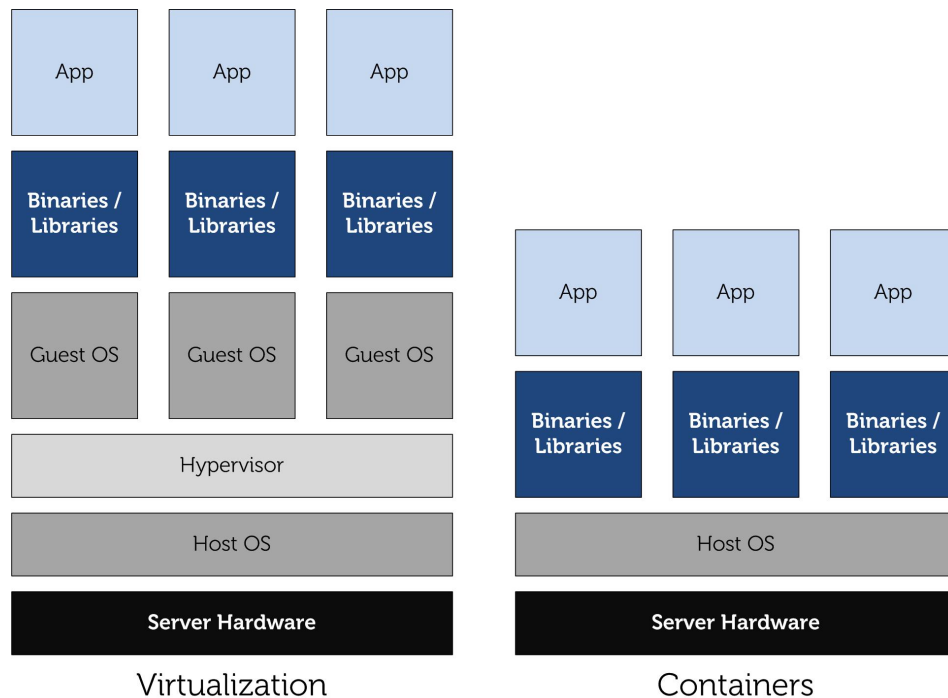


Figure 2.4: VM vs. Containers [DIRB14]

Light virtualization or operating system-level virtualization, does not attempt to execute a complete OS on top of the current OS, but rather it shares the kernel and the libraries to create fully self-contained environments through isolation mechanisms. As a result, light virtualization technologies are generally referred to as *containers*.

These containers, benefits from similar advantages as their fully virtualized counterparts, such as the ability to be cloned, and the ability to run multiple instances on a single hardware base. But the feature that sets them apart, is the ability to host much more instances on a single machine. This feature is the result of using a non-redundant virtualization scheme compared to a full virtualization redundant scheme (OS on top of OS). This scheme grants a full order of magnitude of additional deployment to lightweight virtualization, on identical hardware, from 10-100 VMs to 100-1000 containers [DIRB14].

Operating system-level virtualization originated from a desire to compartmentalize different services within the OS, and one of the earliest example is FreeBSD Jails for the

FreeBSD OS [?]. Another notable example is the Solaris Containers for the Solaris OS [?]; or for Linux, the Linux-VServer [?]. All of these examples offer virtualization mechanisms providing various capabilities, but recent and more capable light virtualization technologies now exists, including Docker and LxC, both of which are presented in the following subsections.

Contrary to full-virtualization, containers do not perform any emulation, but rather through namespace isolation re-uses the libraries and kernel of the host operating system. This makes it an interesting candidate when using low-end hardware, to maximize the hosting capability. But from a security perspective, this entails that full-virtualization has a software layer protecting the host OS from the VM, namely the hypervisor. Whereas, light virtualization has no such layer to protect the host from the container, which has access to the kernel interfaces.

2.4.3 Containers

In this subsection we present the technologies that make containers possible in Linux. Introduced as a kernel patch [Men11], *c-groups* consists of an aggregation and partition mechanism for groups of tasks, that contains (within those aggregations/partitions) all their children. The central technology behind containers in Linux is *c-groups*, or *control-groups*. A *c-group* allows different specialized behaviors to be explicitly defined, by generating a hierarchy of groups. The strength of this concept lies in the ability for it to be used with other Linux subsystems and to provide additional properties for groups of processes. One example that illustrates the definition of specialized behaviors, consist of pairing *c-groups* with the *cpuset* subsystem to restrict a group of processes to a specific Central Processing Unit (CPU)-core. Through the *c-groups* technology, and by extending the definition of the behaviors of processes in Linux, emerged the container technology. We present in the following subsections two different implementations of Linux containers, LxC and Docker.

2.4.4 LxC

LxC or Linux Containers, is an open-source implementation of the containerization technology for Linux. It focuses on providing the tools to facilitate the development of system containers in a distribution agnostic fashion [Los15].

LxC leverages *c-groups* partitioning and aggregating capabilities to provide resource and namespace isolation without any extra virtualization mechanisms, but rather by using the Linux kernel native subsystems. Namespace isolation provides the ability to

isolate running applications completely from the operating system execution environment, albeit not an exclusive feature of LxC. A general example of namespace isolation is the *PID namespace*, which provides the means to create sets of tasks such that each set is completely independent from one another [EK07]. In other words two tasks belonging to different set of tasks, can have identical ID without incurring any ambiguity as to which task belongs to which set. Resource isolation provides the means (via cgroups) to allocate system resources to different groups of tasks.

Thus, LxC provides capabilities analogous to an operating systems (filesystem, network, processes) and dedicated access to the physical resources of its underlying host, in a isolated environment in the form of a container [DIRB14].

2.4.5 Docker

Docker was released on March 2013 and is ever-growing in popularity. It already has its own convention: DockerCon [Hyk14], which is endorsed by major technological companies such as IBM, Microsoft, Google, and RedHat (to name a few).

Docker is open-source and provide means to automate deployment of applications within containers. It provides a high-level interface to containers, by abstracting the intricacies and specificities of the containers into an intuitive and easy to use high-level API, providing management, configuration and monitoring capabilities.

Docker is not a replacement for LxC, rather it is an addition of high-level features, on top of the containerization primitives (LxC) of Linux, to interact with containers. It implements a custom version of these primitives, largely based on the LxC primitives. As does LxC, Docker leverages namespace isolation and c-groups to provide isolation to the container, but also uses *union file-systems* to provide a lightweight templating system, known as images. *Docker Images* are used to specify the operating system environment to be instantiated within the container [Hyk15].

The union file system is a Unix filesystem service, with which it is possible to create virtual filesystems from separate filesystems through branching and layering. One of the advantages of this service is the ability to update a filesystem by simply applying the difference from the previous versions. In Docker's context, this means the ability to update a container by simply uploading the difference (or new layer) and applying it, rather than uploading the entire container again [Hyk15].

Docker uses a client-server model to communicate between the container (client) and the originator/deployer (server), which can reside on a single host or not. Easing the workload for the client, because it is possible to offload computationally intensive tasks,

such as container construction, to more powerful machines while conserving lightweight clients [Hyk15].

Chapter 3

Related Work

In the previous chapter we have presented the background information required to understand the following chapters and this information is useful to understand the context in which we approach the following related works.

In this chapter we present the work that relates to our research, notably the two most relevant projects: Cloud@Home and Peer-to-Peer Cloud System. But before, we introduce a framework designed to evaluate collaborative systems [BJ13], and use it to evaluate the two projects. Finally, we discuss how these solutions cater to our requirements and what is the scope of these projects.

3.1 Evaluation Framework

The framework aims at formalizing the requirements for peer-to-peer collaborative systems and it expresses the functional life-cycle of its participants. More specifically, this framework circumscribes the resource collaboration problem, where each phase cover a functional requirement [BJ13], as illustrated in Figure 3.1.

1. **Advertise:** Each node advertises its resources and their capabilities using formal specifications defined over a set of attributes.
2. **Discover:** A mechanism to discover and keep track of the useful specification advertised by the other participating nodes; it accelerates the querying mechanisms and it preserves the inter-resource relationship information.
3. **Select:** A mechanism for a node to select a group of resources that satisfies the formal specification of their requirements, defined in a query containing the attributes and their respective domains.




Figure 3.1: Framework for Peer-to-Peer Resource Collaboration Problem [BJ13]

4. **Match:** A mechanism to formally specify, in the selection query, the inter-resource relationship requirements and to ensure that the selected nodes satisfy them.
5. **Bind:** Provides a binding mechanism between the resources and the applications, preventing two applications from selecting the same resources. It is useful to cope with the dynamic nature of the peer-to-peer environment, because a node may not (still) be available at the time of use, but it was at the time of selection.
6. **Use:** Utilize the best subset of available resources, from the resources acquired, to execute the application (and all its tasks) while respecting the utilization agreement between the application and the resources.
7. **Release:** Provides a mechanism that allows to release (some) resources in relation to the application demand, and/or the contractual binding (if time-sensitive). This mechanism can also provide the means to partially release a resource and to enable that resource to collaborate with multiple applications simultaneously.

Thus, if one were to implement each one of these functional requirements into a

separate module, then the resulting system would be modular and would address the resource collaboration problem. For more details on this framework and for an evaluation conducted on different peer-to-peer collaborative systems using this framework, see [BJ13] [BJ12].

For our purposes, this framework acts as a checklist to identify how each of the functional requirements inherent to peer-to-peer collaborative systems are implemented in the following two projects. Ultimately, we will use it in further chapters to assess the fulfillment of these functional requirements in *our* own architecture.

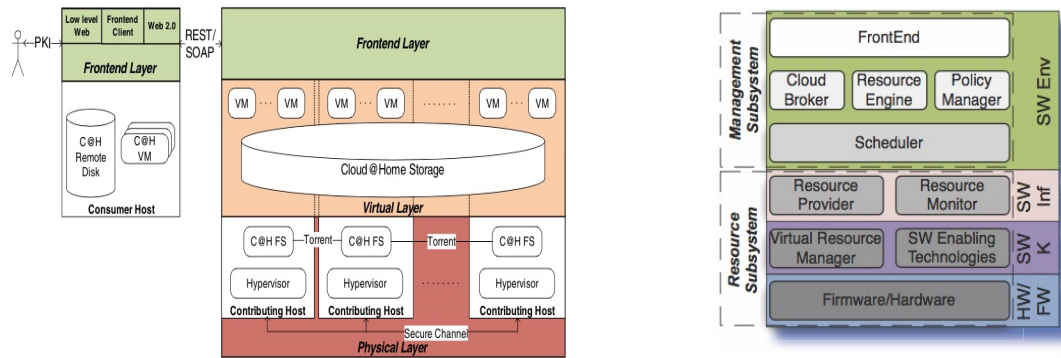
3.2 Cloud@Home

Cloud@Home is a new paradigm which spawned at the intersection of two existing paradigms: the *public-resource computing paradigm* and the *cloud computing paradigm*. It is one of the earliest attempt to create a *volunteer cloud computing infrastructure*. Over several publications, the authors have analyzed extensively the challenges relating to volunteer cloud computing, [CDPS09b] [ABC⁺11] [AAC⁺11] [CDPS10a] [CDPS10b] [DCP10] [CDPS10d] [CDP10] [DFP11] [DP12] [DPR⁺11].

Their vision is to offer a full-fledged cloud computing infrastructure constructed using volunteered resources. Full-fledged, in this context, entails the provisioning of the following service models: IaaS, PaaS and SaaS. The Cloud@Home system aims at offering performances and Quality of Service (QoS) similar to those offered by major CSP.

To fulfill the performance requirements, the authors focused on *interoperability* between CSP, providing a user with the ability to contract resources from commercial CSPs, given that the volunteered resources available are insufficient or inadequate. Whereas to fulfill the QoS requirements, they require that both parties negotiate the expected or desired QoS contained in a Service-Level Agreement (SLA), created by the resource consumer. In other words, the consumers provide their intended terms of use, relative to the resources, and the providers can decide whether or not they accept the terms of use. In case of disagreement with the SLA and the QoS it contains, the negotiation process is aborted. For more details pertaining to the QoS and SLA in the context of this system, see [DPR⁺11].

Contrary to the CSP, because Cloud@Home uses volunteered resources, the authors were obligated to ponder on how to entice people to contribute their resources to this infrastructure. They devised an incentive-based business model that offers financial restitution to the users for their contributions. This model proposes to recycle current idling resources into a utility that can be sold, with respect to the quality and capabilities of



(a) Overview [CDPS09a]

(b) Functional Representation [DFP11]

Figure 3.2: Cloud@Home Architecture

these resources.

In the following subsections we will present and discuss the architecture proposed for this solution, and discuss its conformance to our requirements and the evaluation framework.

3.2.1 Architecture

The Cloud@Home system architecture, shown in Figure 3.2a, is composed of three layers: the *Frontend Layer*, the *Virtual Layer* and the *Physical Layer*. The functional representation of the infrastructure, shown in Figure 3.2b is divided into the *Management Subsystem* and the *Resource Subsystem*.

Starting by the architectural overview, the first layer or *Frontend Layer*, is responsible for providing a high-level service-oriented point of view of the underlying infrastructure. The relation between the Consumer Host¹ and this layer is based on a client-server relationship, providing a centralized point of access to the infrastructure.

The *Virtual Layer* provides a homogeneous perspective of a set of heterogeneous resources, by means of virtualization. Because it uses virtualization, it is possible to completely abstract the underlying hardware specificities of the different contributed resources and to exploit them in a platform-agnostic way. This layer is further divided into two services: the *Execution Service* and the *Storage Service*. The former enables the consumers to host and execute VMs, according to their needs, providing a similar service as the Infrastructure-as-a-Service model provides. Contributors are then assigned an arbitrary virtual machine to host, or multiple VMs to host, relative to their intended

¹Consumer Host, in this context, represents a consumer of the cloud service provider and not a consumer of the application deployed in the cloud.

contributions. Whereas the latter provides a distributed storage facility by implementing a distributed file system such as [GGL03]. Consumers can then mount locally a remote disk that corresponds to a portion of the distributed file system. As for contributors, they host an arbitrary portion of this distributed file system, relative to the amount of disk space they wish to contribute, providing a unified view of the entire distributed file system to the consumers.

The *Physical Layer* is responsible for connecting the contributing resources together and providing means for communication. Negotiations for the resources between the contributors and the consumers are performed at this layer. The negotiation process consists of a request for resources, emitted by the consumer, followed by an attempt to match the parameters of this request with the parameters of the volunteered resources available. The authors specified several mechanisms to ensure and negotiate the Quality of Service and Service-Level Agreements between consumers and contributors (or CSP) [ABC⁺11][AAC⁺11].

The *Management Subsystem* translates the user's requests into multiple sub-requests. Whether it is possible for the requests to be satisfied, using volunteer resources or their commercial resources, depends on the QoS specified by the user and if it is compatible with the resources available. These request(s) are passed on to the Resource Subsystem which binds and matches the resources to the user from which the(se) request(s) originated. The Management Subsystem is centralized in order to manage the infrastructure, and to manage the QoS and the SLA of all the resources, but also to provide dynamic provisioning. The authors claim that it is the only way to aggregate the global information of the infrastructure and to provide a reliable perspective of it [CDPS10c].

Finally, the authors proposed a middleware that implements the ideas from the architectural overview and the functional representation, covering the different goals posed by the Cloud@Home paradigm. Figure 3.3 showcases the deployment middleware of the Cloud@Home infrastructure.

In this overview, we have presented a simplification of the functionalities of the two subsystems for more details consult [CDPS10c]. For more details on the different layers composing its architecture consult [AAC⁺11], and finally, for more information on the middleware consult [CDPS10a].

3.2.2 Evaluation

In this subsection we evaluate the system according to the evaluation framework that we introduced in Section 3.1. Then, we evaluate the system against the functional require-

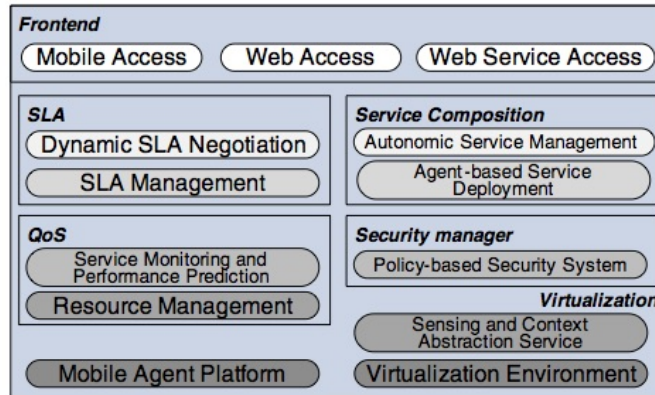


Figure 3.3: Cloud@Home Infrastructure Middleware [AAC⁺11]

ments that we have identified in Section 1.2.

Evaluation Framework

We use the evaluation framework from Section 3.1 to identify any functional deficiencies in this system related to the essential requirements of peer-to-peer collaborative systems and the resource collaboration problem. The implementation of these functional requirements does not require mutual exclusion, and a subset of these requirements can be combined into a single module.

As a matter of fact, such is the case for the first three phases. **Advertise**, **Discover**, and **Select** are accomplished by the Management Subsystem. They are grouped together to provide the ability to a user to enroll and advertise its resources, or to a consumer to find resources that are available according to their needs.

It is also the case for the remaining phases, **Match**, **Bind**, **Use**, and **Release**, all being accomplished by the Resource Subsystem. It provides the ability, given a request, to match the resources which satisfies the request, then to allocate the resources to the consumer in order to use them. Upon using the resources, it is possible to release or to deallocate the resources when the consumer no longer requires them. As a consequence, it provides *dynamic membership* capabilities to the resources, or it allows preemption of the resources according to the SLA or the current workload.

Essentially all the functional requirements presented in the evaluation framework are fulfilled in one way or another, directly or indirectly. As a consequence, this project is appropriately devised to address the essential functional requirements of a peer-to-peer collaborative system and the resource collaboration problem.

Functional Requirements

We now evaluate how this system addresses our functional requirements, as presented in Section 1.2, enlightening us on how it could address the research questions presented in this thesis.

Starting with the first requirement, which entails the use of commodity hardware to form the infrastructure on which it is possible to deploy multiple applications. We can assess that this requirement is completely satisfied, by leveraging public-resources as part of the resources available to deploy an application, and by supporting the deployment of multiple applications simultaneously.

As for the second requirement, which states that no third-party should be introduced to provide any services, in order to maintain the self-containment property of the system. We can only assess partial satisfaction of this requirement, because it is possible using this system to rely exclusively on contributed resources, and not contract resources from any commercial CSP. But in order to use this infrastructure, both consumers and contributors are required to interface with the Management Subsystem, and to an extent this is a third-party in relation to the consumer and contributor.

The third requirement, states that the system should be resistant to resource failures and should not introduce any single point of failure, for security purposes. The Cloud@Home system focuses on creating a lightweight middleware built around the concept of service migrations. Upon failure of any resources, the system can swiftly migrate the corresponding workload to any other available resources, resulting in an infrastructure optimized for lightweight services and failure-resistant. Unfortunately, this system is built around the concept that a central server is required to manage its infrastructure, as stated previously in Section 3.2.1. Although, some fault-tolerant mechanisms are put in place to mitigate this single point of failure, including redundant servers, and the authors claim it is sufficient. Further tests and benchmarks are required to corroborate their claims, in a realistic environment, and to satisfy this requirement.

The fourth requirement, states that no special equipment should be required to access and participate in this system, but rather encourages resource recycling. Consequently, it implies, as a sub-requirement, that the memory footprint should be small enough to enable legacy, or lower-end, devices to participate in this system. Because this system relies on full virtualization technology, it forces the contributing peers to host VMs in order to contribute a computational resource. Thus, limited resources are expected to be consumers rather than contributors, except in very specific cases where the application lends itself to it. Such cases are when the limited resources are used for sensorial inputs, such geo-location, accelerometers, barometers, and other various sensors that could be

used as a stream of information available to any applications in the cloud. These stream of informations are analogous to nodes in a sensor network. It is clear to which extent the memory footprint is small in the context of the proposed IaaS model, but not so much in the context of the other service models. Cloud@Home offers PaaS capabilities by using their *SLA Engine*, ensuring the availability and capacity of the resources, and *CHASE (Cloud@Home Autonomic Service Engine)*, which determines the optimal configuration parameters for the service application using the current configuration of the cloud. Given that this processing occurs at the centralized server, it shouldn't incur a large memory footprint on the contributing resources. But the contributing resources are still required to host VMs. This requirement is not completely satisfied.

The fifth and final requirement, states that the system should provide dynamic membership capabilities to all applications, which means that for every application it should provide means to scale according to the fluctuations in the workload. The Management Subsystem enforces the resources SLA and guarantees the QoS, and consequently it will preempt the necessary resources to ensure that they are only used to the extent prescribed in their SLA. This system supports dynamic membership and provides scalability to the applications deployed, resulting in the satisfaction of this requirement.

Out of the five requirements: one remains unsatisfied, three are partially satisfied and one is completely satisfied. We can conclude that this system does not address our research question sufficiently for it to be considered a viable solution, and thus we must find a more adequate solution or create one.

3.3 Peer-to-Peer Cloud System

The Peer-to-Peer Cloud System (P2PCS), was developed as part of a doctoral thesis [BMT12], and proposes a slightly different approach to volunteer cloud computing than the previous system. One of the defining characteristic of this system is the fully decentralized structure connecting the different peers, constructed using epidemic and gossip-based protocols.

The authors propose an infrastructure with no centralized server(s) to manage the resources, conversely to the Cloud@Home system. Rather each node interact directly with the other nodes when performing any system operations. Consequently, the consumer nodes interacts directly with the (potential) contributing nodes in order to select the fittest candidates to host the application. Fitness of a candidate is evaluated using the response of a specially defined query, which contains the desired criteria for candidate resources (e.g., 10 nodes with at least 8 gigabytes of RAM...). The collection of

nodes that fulfill these criteria are selected and forms a *Slice*. This slicing mechanism is an *attempt* at solving the slicing problem presented in Section 2.3.5. The consumer would communicate with this slice, or pool of nodes, via an API that is similar to what current IaaS provider (Amazon EC2 or S3) uses. Using this API, the consumer is able to instruct the contributing nodes when to start or stop the VMs according to the application requirements. The authors present this system as a peer-to-peer cloud computing architecture focused on providing only one service model, the IaaS.

In the following subsections we present the architecture of the system, and proceed to evaluate the system against the evaluation framework and our functional requirements, as we did for the previous system.

3.3.1 Architecture

The architecture is divided into 3 functional components, although the authors did not explicitly use nomenclature to distinguish between correlated components, we do.

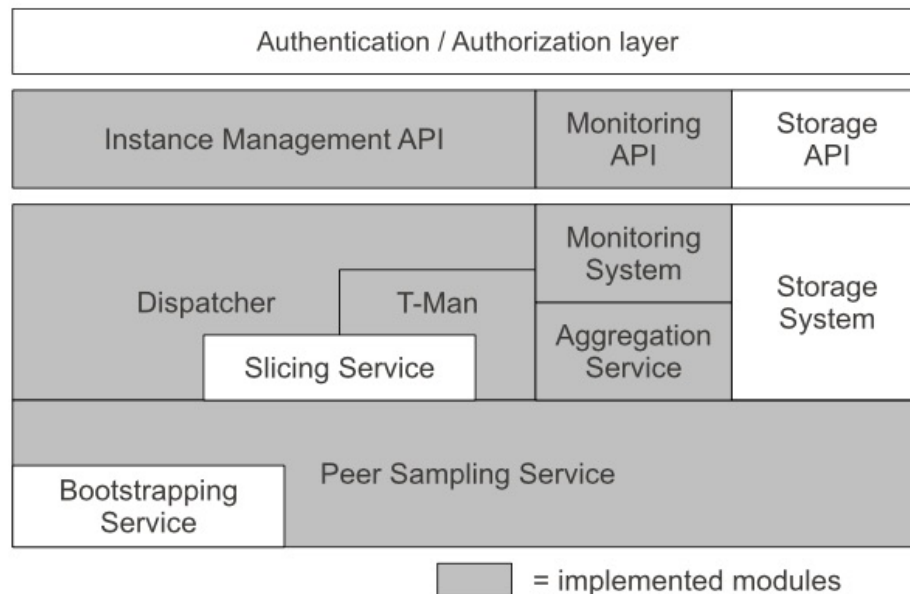


Figure 3.4: Peer-to-Peer Cloud System Architecture [BJ13]

Figure 3.4 presents the different components of this system. The authors distinguished between what has been implemented (in gray) and what has been left as future work (in white).

Using a top-down approach, we present the first component, representing the point of access to the system. The *Authentication/Authorization layer* is where a user (contributor/consumer) is required to identify themselves in order to access the system.

The second component represents the principal access point to the cloud for the consumers, it is named the *API layer*. It is composed of three sub-components providing different capabilities through various interfaces. The *Instance Management API* sub-component provides an interface that circumscribes the possible interactions between the consumer and the contributing nodes, with respect to the VM instances. Whereas, the *Monitoring API* sub-component provides a visual representation of the slice topology, and the *Storage API* sub-component provides the means to configure the distributed storage system for the current slice.

The third component represents the *Networking layer*. This component is composed of three sub-components, which in turns are composed of three or less components.

The first component of the Networking layer is composed of two services, which are the *Bootstrapping Service* and the *Peer Sampling Service*. As its name implies, the former provides the bootstrapping mechanism for a node to join this system. Whereas the latter, provides a list of peers in the network to exchange messages with. The creation and maintenance of the list is accomplished by creating and maintaining an overlay network over all the peers in the system. This overlay network is created using a simple gossip-based protocol as presented in [JVG⁺07].

The second component of the Networking layer is composed of three sub-components, which are the *Dispatcher*, the *T-Man* and the *Slicing Service*. The *Dispatcher* sub-component is responsible for the translation of the commands issued through the various interfaces by the consumer into commands that are compliant with the gossip-based protocol(s) used in this system. The following sub-component, is *T-Man*, a gossip-based protocol providing the ability to create and manage structured overlay networks, using various topologies [JMB09]. Whereas the last sub-component, the *Slicing Service*, provides the slicing capabilities for this system, as presented in Section 2.3.5, by ordering the nodes according to node-local attributes, then dividing the nodes according to various thresholds. Previously, we have presented the slicing capabilities of this system as being an attempt to solve the *slicing problem* because of how it is implemented. The (current) implementation enables the creation of slices using only one metric, the number of nodes a consumer wishes to contract, which is completely independent of any node-local attribute.

The last component of the Networking layer is also composed of three sub-components, the *Monitoring System*, the *Aggregation Service* and the *Storage System*. Through the corresponding API, the *Monitoring System* provides to the consumer access to global system information collected and computed by the *Aggregation Service*. The *Aggregation Service* provides system-wide parameters to the peers of the system. These parameters

are computed by aggregating the various local parameters from each peer, through local message exchange among peers. To accomplish this in a decentralized and dynamic environment, this service uses a *push-pull* gossip-based protocol, as presented in [JMB05]. The last sub-component is the *Storage System*. It provides the distributed storage capabilities of this system and it is accessed through its *Storage API*.

From this overview of the architecture of this system, we can observe that gossip-based protocols are prevalent, which is *the* defining characteristic of this system. These protocols are praised for their ability to thrive in highly dynamic and volatile environments, and a considerable amount of literature has been published on the subject. For a discussion on the limitations inherent to these protocols, and its strengths see [Bir07]. For a more introductory approach to gossip-based and epidemic protocols in the context of distributed systems see [Jel].

3.3.2 Evaluation

In this subsection we evaluate this system against the evaluation framework, as presented in Section 3.1, then against our requirements, as presented in Section 1.2.

Evaluation Framework

Again, we rely on the evaluation framework to identify any functional deficiencies in this system, with respect to the essential requirements of peer-to-peer collaborative systems and the resource collaboration problem. As stated previously, the requirements of this framework can be implemented alongside other functional requirements as part of a single module.

The first two phases, **Advertise** and **Discover**, are coupled together within the *Peer Sampling Service*. But they are not exhaustively implemented. We underline the fact that the system provides no mechanism to advertise the resources a node has to offer, to the network explicitly, either using formal specifications, as proposed in [BJ13], or using any other strategies. Rather, this service presents a node to the other nodes using a gossip-based protocol. We deem these functional requirements to be partially fulfilled, because it is desirable to be able to advertise and discover resources based on their capabilities.

The **Select** phase is accomplished by the *Slicing Service*. The authors designed this system to use a gossip-based protocol to automatically partition the network into slices according to a metric [JK06], but was not implemented and slices are created based on the desired number of nodes. Thus, we deem this functional requirement to be partially

fulfilled, because it is not possible to select the resources according to their capabilities.

The **Match** phase is suspected to be realized by the *Slicing Service*, by furthering the specification of the selection query. But the limitations of the gossip-based protocol used are not clear, because the slicing mechanism proposed by the authors is based on a utility function that describes the usefulness of a node with respect to the other nodes, using node-local attributes; and to introduce additional dimensions to perform multi-dimensional slicing using a utility function is not the best approach since it introduces all sorts of statistical distortions [PMRS14]. Thus, we also deem this functional requirement to be partially fulfilled, because support for multi-dimensional slicing is not explicitly provided.

Bind, the fifth phase, is achieved using the *T-Man* gossip-based protocol. This protocol is used to create a ring overlay of all the peers contained in a specific slice, 1 ring per slice, resulting in mutually exclusive slices. There are no indications on how to cope with concurrent attempts to bind overlapping sets of resources, given that these resources satisfies multiple selection queries. This is due to the fact that the slicing mechanism uses a single metric, namely the number of requested nodes. Then instances of the concurrent binding problem would only occur in very a specific case. Manifesting itself in the form of two or more concurrent requests, for which the total requested number of nodes combined exceeds the total number of available nodes, but the number of nodes per request is inferior to the total available nodes. The problem then becomes, which application is entitled to have their request fulfilled? Because no a-priori ranking mechanism (or reputation mechanism) to decide which consumer should be favored is provided, the behavior of the system in response to this characterization of the problem is undefined. Again, we deem this functional requirement to be only partially fulfilled, because of the lack of concurrent requests binding support.

The sixth phase, referred to as the **Use** phase, is satisfied by the *Dispatcher*. It translates the higher-lever API requests into the appropriate low-level gossip protocol commands to be sent to the nodes of the slice, resulting in the utilization of the resources. This functional requirement is completely fulfilled, because this system provide the ability to use the selected resources.

The last phase, **Release**, can be accomplished by the *Instance Management API*, granting the ability to the consumer to control the VM instances it was assigned. The automation of this phase would be the result of the co-operation between the *Monitoring System* and the *Aggregation Service*, collecting global information about the system, such as the current load in a slice, and acting on these results by preempting the necessary amount of resources. This functional requirement is also completely fulfilled, because

using the *Instance Management API* in concert with the *Monitoring System* and *Aggregation Service* provides the ability to release resources when they are no longer required.

According to the evaluation framework, this architecture exhibits *some* of the characteristics that are essential in a peer-to-peer collaborative system. Only *two* out of the *seven* functional requirements are completely fulfilled, and *five* are only partially fulfilled either conceptually or because of the lack implementation. Nonetheless, this system presents a solid foundation that could be extended to meet all the requirements, because there are no apparent design decision that would explicitly prevent it.

Functional Requirements

We evaluate this system against our functional requirements, as defined in Section 1.2.

The first requirement, states that it should be possible for the infrastructure to be constructed using only commodity hardware and that it should be able to host multiple applications at once. By relying solely on user contributed resources and by providing the ability, to the consumers, to create mutually exclusive slices for their applications, this requirement is deemed to be completely fulfilled.

As for the subsequent requirement, which states that every services should be provided by the contributing resources, not 3rd-parties, we can conclude that it is fulfilled. We arrive at this conclusion by observing that the proposed system relies solely on contributed resources and does not introduce any 3rd-parties.

The following requirement states that system should not introduce any single point of failure and the system should be (somewhat) resistant to ubiquitous failures. As a consequence of being fully decentralized this architecture does not introduce any single point of failure. Using gossip-based protocols to create and maintain the overlays, this system is able to cope with a very dynamic networking environment. Failures or node leaving/joining the network, are well supported by these protocols, thus we deem this functional requirement to be fulfilled [BJ13].

The following requirement states that no special or dedicated hardware should be necessary to participate in this system, and consequently that memory footprint should be as low as possible as to enable a wider range of (weaker or lower-end) resources to participate. Nothing requires from a participant to acquire any special or dedicated equipment to participate in the system. For the IaaS model it is required to use full virtualization technologies, increasing the memory footprint, and it is not clear as how it would fare on lower-end resources. Thus we cannot deem this functional requirement to completely fulfilled, because it uses full-virtualization technologies and the authors have not shown the memory footprint to be small.

The last requirement states that the system should be able to provide scalability to the applications deployed, and consequently provide dynamic membership capabilities. As stated previously, the use of gossip-based protocols facilitate the support for dynamic membership capabilities because of their resistance to failures and arrival/departure of nodes. Consequently, this functional requirement is deemed to be fulfilled.

Ultimately, this system fulfills *four* out of the *five* functional requirements explicitly, making it almost a viable solution, but not quite. We put the emphasis on the fact that the system adopts sensible approach to solve our problem, but the limitations present dissipates any hopes of redeeming it. To back these claims, we want to outline the fact that a lot of the problems of providing a distributed computing platform are simply relayed to the end-user when adopting a IaaS model. Because it reverts the responsibilities to mitigate the problems, back to the user, in the form of (proper) configuration of the VM instances required for any applications. Thus, we think that this system does not, and could not without substantial modifications, provide a PaaS computing platform or a system capable of addressing our research requirements.

3.4 Discussion

In this section, we discuss the various details and implications of both projects, and provide a reflection on what they accomplish from a high-level overview.

The Cloud@Home system offers a decent solution to the volunteer cloud computing paradigm but also with respect to the functional requirements of peer-to-peer collaborative systems. This solution trades off full decentralization and the capacity to leverage lower-end resources for augmentation of performance and QoS. It introduces possible third-parties, in order to provide these performance and QoS guarantees, and thus it breaks our perspective of being fully self-contained with respect to the consumers and contributors. It aims at providing a business model to transform current computing resources into a utility that can be monetized.

In contrast, P2PCS provides an architecture that can transform a group of computing resources into a application deployment platform, corresponding to the IaaS model. This system provides the ability to leverage resources using full-virtualization, but it does not provide a business model to justify the viability of the infrastructure, or to provide any incentive to entice possible contributors. Like many open-source and community-driven project it anticipates a self-policing behavior from the community. The API proposed to interface with the resources is analogous to traditional IaaS APIs. The strength of this system lies in the adoption of gossip-based protocols to accomplish the underlying

network creation and maintenance. Consequently, by using these protocols it mitigates a lot of the complexities and difficulties of operating a distributed system over unreliable network infrastructure, such as the Internet.

Both of these projects offer interesting architectures supporting a multitude of desirable features for volunteer cloud computing and public-resource computing, but from our perspective some key features are still missing. Notably the ability to leverage a collection of commodity devices to provide a multi-application computing platform akin to a PaaS model, by limiting the memory footprint. The PaaS model has different concerns compared to the IaaS model, with respect to security and isolation. The latter uses VMs, to isolate the contributing environment from the hosting environment. And it is not clear how this could be accomplished in either projects without resorting to VMs. Consequently, providing a PaaS model, using the contributed resources rather than leveraging a commercial service provider, remains unresolved for P2PCS and problematic for Cloud@Home when using lower-end resources.

Ultimately, due to the deficiencies outlined, we express the need of creating an architecture based on the observations we made while evaluating these projects. In Table 3.1 we present a summary of how well each project fared against the requirements and the evaluation framework. For illustration purposes, we have included the rating for our architecture, we will present and justify in the next chapter.

	Cloud@Home	P2PCS	Our Architecture
Evaluation Framework			
Advertise	+	+	++
Discover	+	+	+
Select	+	+	++
Match	+	+	+
Bind	+	+	+
Use	+	++	++
Release	+	++	+
Functional Requirements			
Requirement 1	++	++	++
Requirement 2	+	++	++
Requirement 3	-	++	++
Requirement 4	++	++	++
Sub-Requirement 4.1	+	+	+
Requirement 5	++	++	++
Legend:			
	++ = completely fulfilled		
	+ = partially fulfilled		
	- = not fulfilled		

Table 3.1: Summary of the Solutions and the Requirements

Chapter 4

Architecture

In this chapter we present the architecture we have proposed, and illustrate the different decisions that influenced our design. We designed this architecture based on the observations made in the related works, and our design takes into account our requirements, as presented in Section 1.2, and the essential functional requirements for peer-to-peer collaborative systems, as presented in Section 3.1.

The first section provides a high-level perspective of the architecture, whereas the subsequent sections present and express the design rationale behind each layer. We then present an exploration of the various components of this architecture and their interactions. We conclude with a discussion of the requirements that were addressed by designing this architecture.

4.1 Overview

We used a multi-tiered approach for the architecture, similar to Cloud@Home and P2PCS, because it provides high modularity and loose coupling of the concerns addressed by each tier. In this architecture we differentiate between 3 types of nodes, namely a *Worker node*, a *Data node* and an *Application Deployer node*. A *Worker node* consists of a computational resource, executing tasks encompassing *all* the functionalities of the applications. Whereas a *Data node* consists of a storage resource, performing all the database and storage related tasks. The *Application Deployer node*, is the node that contracts contributing nodes to host an application in this system.

Starting at the lowest-level of Figure 4.1, we have the **Network Layer**. It is responsible for all the network responsibilities including the creation and maintenance of the overlays, the communication amongst the participants, and to provide the high-level layers with an abstraction of the underlying physical network structure.

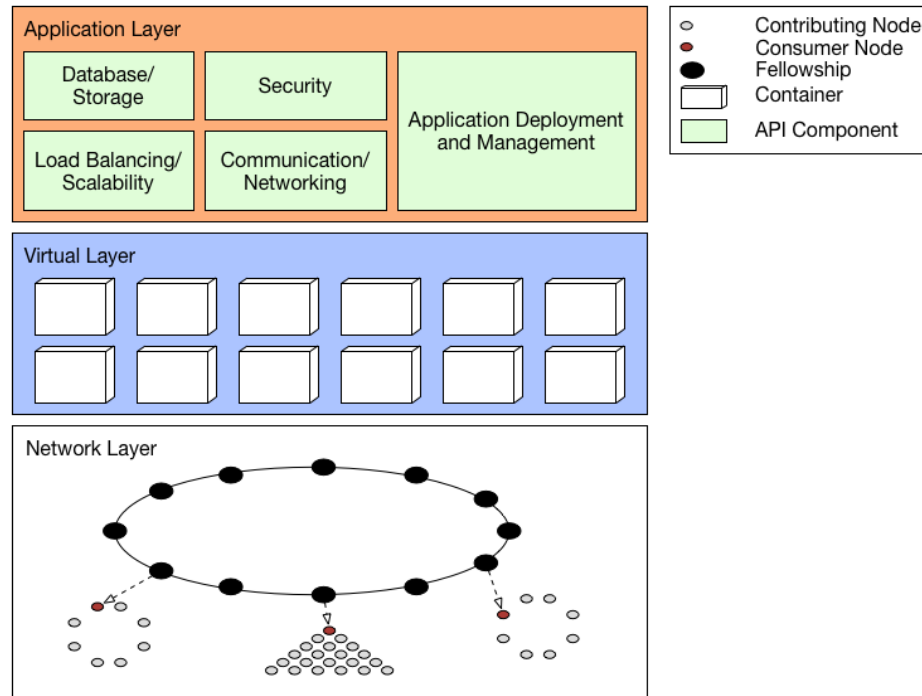


Figure 4.1: Overview of the Proposed Architecture.

On top of the Network Layer, we built the **Virtual Layer**. It is responsible for abstracting the physical characteristics of each node and providing a homogeneous interface to a collection of heterogeneous resources using virtualization technologies.

Finally, the top-most layer is the **Application Layer**. It exposes an API to the consuming node used to build applications. This API is a minimal specification of the services required to transform raw resources into a complete computing platform, similar to PaaS computing platforms.

4.2 Network Layer

In this section we present the *Network Layer* and the rationale from which it originates. As stated in Section 4.1, this layer is responsible for providing all the networking capabilities for the system. It accomplishes this using two abstractions, *the Ring* and *the Fellowship*. We present the two abstractions and explain their relevance for this architecture.

When devising a distributed system, one is confronted with fundamental problems such as: *How can we connect a collection of computers, using the Internet, and maintain connectivity among them?* From what we have presented in Section 2.3.2, we can address

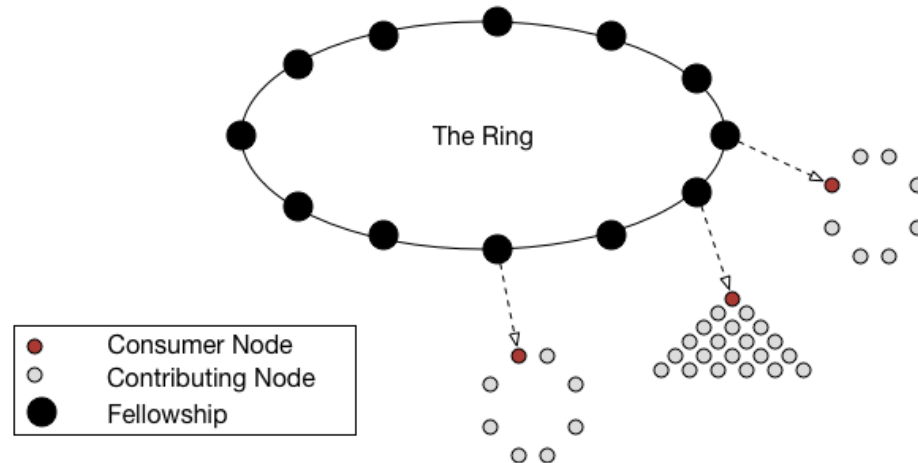


Figure 4.2: Network Layer Overview

this problem by creating an overlay network to connect a collection of computers and maintain connectivity among them. But then the question becomes, which topology of overlay network is best suited for our requirements, or is it not relevant in this context?

A case-study presented in [BJ13], outlines the difference between the various topologies an overlay network can have in the context of a peer-to-peer collaborative system. The authors concluded that *both* topologies of overlay networks, structured and unstructured, exhibit different characteristics which can be desirable for peer-to-peer collaborative systems, but none of the topologies provides a comprehensive solution addressing the requirements identified by the authors. Consequently, we know that a single overlay network will not address all our requirements of peer-to-peer collaborative systems. Before proceeding further lets recall quickly what the two related projects have done to address these requirements. If we look at the networking layer for both projects, Cloud@Home and P2PCS respectively, we notice that the former favors a centralized management entity. This entity is used as a central connection end-point, each node contacts this entity to be part of the network, and this entity maintains the connectivity in the network. The latter focuses primarily on epidemic and gossip-based protocols to generate the overlay networks and to maintain them. Adopting a slightly different strategy, each application generate and maintain their own overlay network (slices) resulting in a federated management structure, where each slice is responsible for its nodes.

In order to respect **Requirement 3**, as presented in Section 1.2, we must not opt for a centralized solution, because it introduces a single point of failure in the system. We chose a structured overlay network topology for our first abstraction, the *Ring*, providing the characteristics and features that are desirable for its intended purpose. For our second

abstraction, the *Fellowships*, we open the possibility to use any overlay network topology.

In the following subsections we present these abstractions, and we present the different design decisions made and the rationale behind them.

4.2.1 The Ring

We have created the *Ring* abstraction to separate the concerns of the publicly available portion of the networking infrastructure from the privately available portion or the application environment. In this subsection we present the design rationale, followed by the conceptual manifestation of this abstraction and its implications for the architecture.

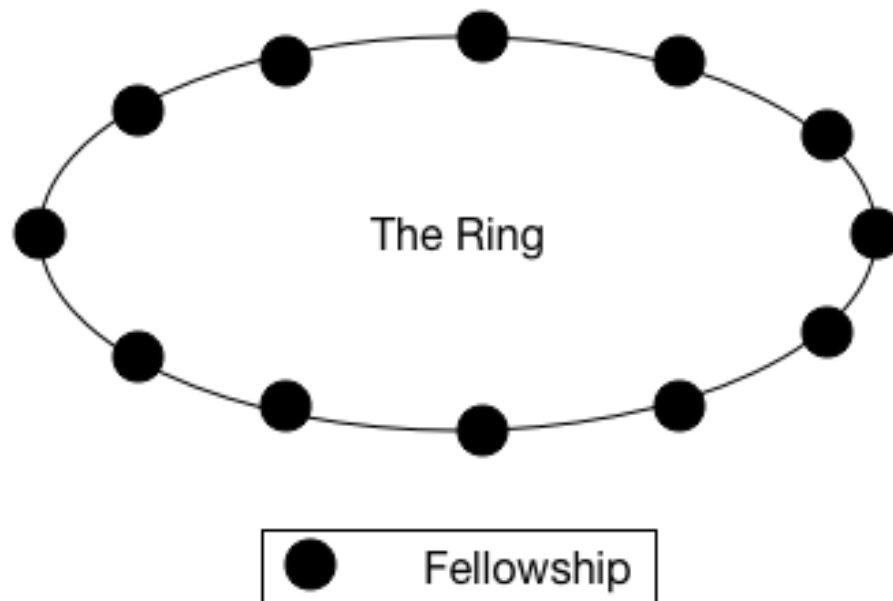


Figure 4.3: Abstract Representation The Ring.

The necessity of devising an abstraction to encapsulate the *public* environment in which our architecture operates, is an attempt to provide a construct that exhibits some of the desirable characteristics of a peer-to-peer collaborative system. We have identified four distinct features that we deemed desirable for this construct:

1. Connecting the participants together in a public environment, such as the Internet, resulting in a *public meeting point*.
2. The management responsibility should not be centralized, as to prevent a global system failure occurring as the result of a(ny) subset of participants failing or leaving.

3. Provide querying mechanism to *publicly* locate applications that are deployed using this architecture, and ensure its reliability.
4. Any information that might be required by a(ny) participant to join an application should be contained in this construct, and the consistency of the information should be ensured to prevent malicious participants from corrupting the information and paralyzing the system.

In order to provide the first feature, it is clear that an overlay network would suffice, independent of its topology. We can use the same rationale for the second feature, because none of the topologies imposes the centralization of the management responsibilities and they can operate in fully decentralized environments. In order to provide the third feature, it is necessary to opt for a *structured* topology, rather than an unstructured topology. Because, as we have presented in the Section 2.3.2, unstructured overlay networks do not provide a deterministic querying mechanism. Furthermore, the last feature entails that this construct is fully self-contained and thus, it must possess some storage capabilities to store the information.

Because of its storage capabilities, its deterministic querying mechanism and its fully decentralized architecture we opted for a structured overlay network. As a matter of fact, we are using *Kademlia* DHT, which provides the ability to locate any node, deterministically, in a time complexity of $O(\log n)$ [MM02].

Using this DHT to store information is intuitive, but using this technology into the context of this architecture imposes further constraints. We need to ensure that malicious nodes cannot compromise the system by polluting or corrupting the information stored in the DHT. This corresponds to a *storage attack*. Essentially, all DHT are inherently vulnerable to this type of attacks, because the design assumes a cooperative environment, where each node is benevolent and most implementations do not address this directly, rather they leave it to the application developer. For an extensive survey of the security techniques applicable to DHTs, see [UPS11]. Thus we require the following changes to the DHT, to ensure consistency and correctness of the information stored in the DHT to an *acceptable* degree. Restraining the writing access of the participants, to only those that are *application deployers*¹, reduces the threat model posed by possible storage attacks.

Then, the problem becomes: how can we ensure the benevolence of the *application deployers*? It is not possible to ensure absolute benevolence, but it is possible to limit the potential consequences of a storage attack by some malicious writer. We limit the application deployers ability to write values, to only two keys. One of those keys is named,

¹Further information on the types of participants is provided in Section 4.1.

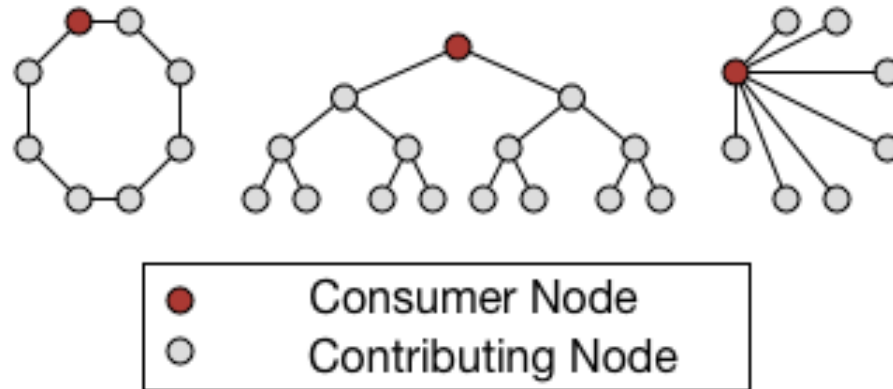


Figure 4.4: Fellowship Abstraction.

the *template key*, and it corresponds to a public repository containing the attributes a node can use to advertise its resources, be it dynamic attributes such as CPU usage, or static attributes such as total amount of physical memory. Whereas the other key, refers to the application deployed by the participant, and can only be written to by this participant, but it can be read by all. It contains the list of nodes composing the *Fellowship*, providing the ability to a node to reconnect to a previous application following a failure by contacting any nodes in the list.

Ultimately, the implementation of the DHT remains unchanged, and we enforce these restrictions in the interface defined for this layer².

4.2.2 The Fellowships

This abstraction is responsible for the *private* portion of the networking infrastructure. We use the word *private* to distinguish between publicly available networks, such as the *Ring*, and privately available networks, such as the network of contributing nodes for a single application. In this subsection we present the design rationale for this abstraction, its conceptual manifestation, and its implications for this architecture.

We design this abstraction to have a clear separation between the nodes that *are* contributing and the nodes that *wishes* to contribute, by segregating them into different networks. This confers the private characteristics to this abstraction, and accentuates the modularity of this architecture. Consequently, this abstraction needs to account for the following essential features:

1. The ability to discriminate between contracted nodes and malicious nodes (pre-

²More information about the interfaces to the different layers in Chapter 5.

- tending to be a contracted contributing node).
2. Allowing any nodes to join and leave in a graceful fashion, and consequently remove any single point of failure.
 3. Providing means for secure communications among the nodes contributing to the same application.

The *Fellowship* consists of a private collection of nodes interacting together exclusively. They are not required to interface with any public environment directly, in contrast to the *Ring* when new nodes are arriving.

Our primary focus, with respect to this abstraction, is the isolation of these private application networks, by controlling which participant can join and which can't, as expressed by the first feature. We can ensure this using a *whitelisting mechanism*. The application deployer will create a list, by selecting candidate contributors, and will **only** contract these contributors to host the application. This is especially useful in *semi-private* or *semi-trusted* environments. A *semi-trusted* environment, is an environment in which the different parties (i.e., contributors and consumers) are known to each other a-priori, but are connected together through public networks, such as the Internet. If it is impossible to establish a semi-trusted environment then other means of authentications are required. In a fully untrusted environment, such as the open Internet, it is very difficult to ensure the identity of a specific participant without resorting to any 3rd-party acting as an authority. Although, we could mitigate this by implementing a reputation system that would gradually increase the trust component for each participant as they contribute [JC10], we leave it as future work.

For the second feature, we are inclined to implement a similar strategy as Peer-to-Peer Cloud System did with the slices. Similarly, we suggest to create one overlay network per application, and we do not impose any topology, but rather leave it as an application dependent design decision. Since we are not restraining the topology of the overlay network, we extend greatly the extensibility and flexibility of this architecture, because depending on the application's networking requirements a topology might present more advantages over another topology. For example, if your application interacts with a plethora of sensing devices distributed geographically, an *unstructured overlay network* would provide the ability to aggregate the information while flooding or performing random walks through the network, thereby providing the ability to perform complex queries. Whereas, a *structured overlay network* would require multiple simple queries and post-processing of the results in order to aggregate them, achieving similar results less efficiently.

The last feature, can be provided by implementing a PKI, where the application deployer is the certificate authority providing public keys to every node, given that their identity has been verified by the registration authority, which can (and should) also be the application deployer. Then, by using this technology we can ensure that the communications are encrypted, certified, and secured.

Finally, this abstraction provides the ability to control which participant is allowed to participate in the collaborative application, but also the ability to secure the communication channels using public-key cryptography. Depending on the networking requirements of each application, it is possible to tailor the *Fellowships* structure to fulfill these requirements.

4.3 Virtual Layer

In this section we present the *Virtual Layer*, defining its purpose and the implications for this architecture. In other words, we present what this layer should provide for our architecture and how should it provide it in order to respect and fulfill our requirements. We then showcase how our approach differs from the previous approaches, used by the two projects evaluated in the Chapter 3.

There are three essential features that this layer should provide. The first and foremost feature this layer should provide, is a security mechanism to isolate a dedicated execution environment in a contributing host system. It should *sandbox* the contributed resources from the contributors operating system, in such a way that no matter what is executed within this *sandbox*, it can never access the OS hosting it or better yet, it cannot know whether it is a virtualized resource or a physically dedicated resource.

Another important feature this layer should provide, is the ability to abstract the resources from their underlying physical architectures. Using virtualization, it is possible to abstract away the specificities of the underlying physical resource, and to present all the resources, in an agnostic fashion, to the possible consumers. In other words, it is possible to present two physically different resources, such as a computer using an *Intel* CPU and a computer using an *AMD* CPU or even an *ARM* CPU, as identical computational resources, distinguishing them solely based on their capacities rather than their proprietary physical architectures.

The last feature this layer should provide, is the ability to control and restrict the amount of resources contributed to the system. That is, a contributing node should be able to restrict its contribution to a desired threshold. Consequently, when allowing only a small percentage of the available resources to be contributed, the virtualization

technology should aim at minimizing its overhead memory footprint as to maximize the usage of the contributed resources.

Then, the design question becomes, which virtualization technology offers these features without compromising our initial research requirements, presented in Section 1.2. *Light virtualization* technologies, are a potential solution to this design question, but let's examine to which extent they are different from *full virtualization* technologies.

As prescribed in the first feature, light virtualization provides the isolation required to securely execute applications without interfering or compromising the host execution environment. Due to their sandboxing properties, executions done inside a container are opaque to the host OS, as would be the case with full virtualization technologies. Thus both technologies, light and full virtualization, fulfills this requirement completely. Same goes for the second feature, which is also provided by both types of virtualization technologies. Because, both types provide homogeneous abstractions for the physically heterogeneous computing resources, and both types present the physical resources in an agnostic fashion. Still, both technologies are equally desirable in the context of the first two features.

Now we need to look at whether or not it provides the ability to restrict the virtualization to a subset of the available computational resources. Docker, one of the containerization technology presented in Section 2.4.5, offers this functionality using their image system. Using the c-group technology, as presented in Section 2.4.3, users hosting containers are able to restrict the processes spawned within a container to only a specific subset of the available resources. As for the full virtualization technologies, it is also possible to restrict the amount of resources available to a VM, through various configuration parameters. But a *difference* ultimately persist between the two types of virtualization technologies. By using lightweight virtualization technologies, we can leverage the libraries and the kernel of the host OS, reducing the possible redundancy of the libraries and binaries to a minimum, as presented in Section 2.4. Whereas, using full virtualization technologies, it is likely that some of the libraries and binaries are installed twice, once in the VM itself and another copy could be installed on the host OS. This outlines the major difference between full and light virtualization technologies. Because, by design the overhead incurred by full virtualization, with the notion of a hypervisor residing on top of the OS having to translate the requests and commands from the VMs into intelligible commands for the OS, is far greater than by light virtualization. Because it removes the intermediate hypervisor, and instead uses namespace isolation and resource isolation to access the resources directly.

We conclude that light virtualization technologies are more apt for limited-resources,

because they provide isolation as well as minimize their memory footprints, thereby fulfilling the last requirement.

Using light virtualization technologies is justified due to our specialized requirements, and in this interlude we examine how the previous projects (Cloud@Home and P2PCS) addressed their virtualization requirements. If we recall what both projects proposed, we can observe that **both** advocated the use of full virtualization technologies. It provides the desired isolation properties, as well as possessing the abstractive capabilities necessary for *their* architecture. In the context of IaaS, it is mandatory to resort to VMs in order to provide the ability to host an entire OS on the resources, as advertised by this service model.

We beg to differ with respect to full virtualization as being adequate for *every* service model. For a service model akin to PaaS, VMs provide too much flexibility and extensibility of configuration which complicates the application creation and deployment process. It reduces usability by forcing the person writing the application to consider details about the configuration of the execution environment, across multiple layers rather than only across the application and data layers as presented in Figure 2.2. Conflating multiple concerns ultimately hinders the productivity and the usability of the computing platform. We believe that constructing an application using a computing platform, such as a PaaS, where the platform is meant to abstract away the characteristics of the underlying physical resources, enables the developer to write an application using the interfaces provided independently of their underlying implementations. Whereas full virtualization provides exactly the opposite experience, by exposing all the underlying resources, it forces the user to decide which implementation to use and how the resources should interact with each other at a lower-level.

Consequently, we choose to use operating system-level virtualization or lightweight virtualization, because it satisfies our virtualization requirements, but also vastly improves the usability of the system. To the best of our knowledge and in all humility we believe to be the first proposing this type of virtualization as an integral part of a distributed computing platform.

4.4 Application Layer

In this section we present the *Application Layer*, and we present how we devised this minimal API specification, which reflects the essential features of distributed computing platforms, such as PaaS.

Initially, we present the rationale behind this layer, and then proceed to present each of the components composing this API: the *Databases and Storages* component, the *Communication and Networking* component, the *Load Balancing and Scalability* component, the *Security* component and the *Application Deployment and Management* component.

4.4.1 Overview

The purpose of this layer is to provide to the application developer the necessary building blocks to develop a scalable distributed web application. We propose a minimal API specification that provides these building blocks. It is minimal in the sense that it is sufficient to develop most applications, but more features can be added to ease the development process of more complex applications.

By definition APIs are extensible, since they provide interfaces to the functionalities contained, while abstracting away the details of the implementations. Using an API to develop applications provides greater modularity with respect to the underlying system, because these applications will always be compatible with this system even if several update occurs, as long as it respects the interface defined initially.

We have defined the essential components of this API by investigating the major CSP that provides a PaaS service model, such as *Google, Amazon, and Microsoft*. For more details on the different services offered by the major service providers, which served as a basis for this proposed specification, and reproducibility's sake please refer to [Inc14b] [Inc15a] [Inc14c]. Then, we have identified the overlapping components, and eliminated the *redundant* components. Figure 4.5 is a pictorial representation of the resulting API specification.

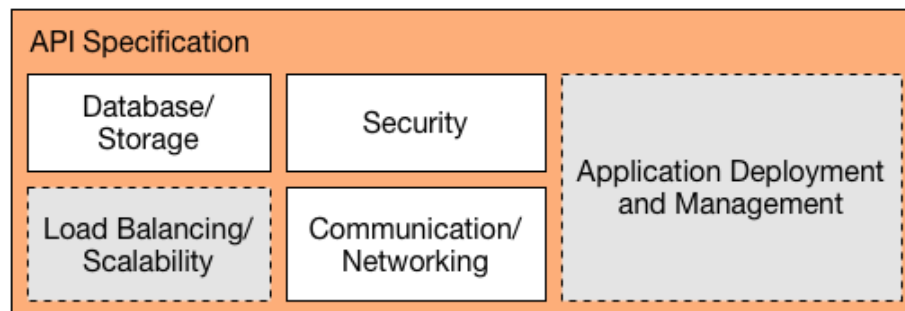


Figure 4.5: API Specification Overview

The components identified by a *solid box* are those exposed to the application developer, and those identified by a *grayed dashed box* are integral parts of the distributed computing platform. The division of concerns is not absolute, and we chose to provide

access to the components represented by a *grey dashed box* to the application developer through configuration parameters rather than building blocks.

In the following subsections we discuss in further details each of the five components identified, providing an overview of the functionalities they contribute to this architecture and the reason of their inclusion. For more details about this API specification, see Appendix A.

4.4.2 Databases and Storages

The first component of this API, is centered around the necessity to *persist* data or information in the context of a distributed application. We have defined a taxonomy of the primary storage services offered by the major CSPs, and then we discuss the concerns relative to the provision of these services and its consequences on this architecture.

- **Relational Databases:** provides traditional Relational DataBase Management System (RDBMS) facilities.
- **Non-Relational Databases:** commonly referred to as *No-SQL*, this type of databases offers schemaless database facilities, such as key-value stores.
- **Storage:** provides all storage needs with larger space requirements per entry (up to 1TB) and for heterogeneous objects that usually are represented as a binary string (for which the format and content are not relevant and all objects are represented as the same type of object). A common way of using such storage components is to pair it with a RDBMS, in which we store the meta-data for all the objects stored in the data store. Then this meta-data is indexed and associated with a key that represent the location of the actual data in the data store.
- **Caching:** provides caching capabilities for fast access to small chunks of data, storing them in memory for future access.

Relational databases, non-relational databases and storage systems, in general, share similar concerns about *availability*, *reliability*, and *consistency* in a distributed environment, and thus we can analyze them in parallel.

In all cases, due to the distributed nature of the underlying compositional resources, there is a need to properly *replicate* the persistent data of the application, in order to ensure *consistency* and *availability* of the data.

Were we offering a IaaS computing platform, we would extend concern about the distributed nature of the system not only towards databases, but also towards file systems,

because each component host their own file system. Multiple solutions exists, one of which is the distributed file system from Google, the *Google File System*³ [GGL03]. But we are focusing on a computing platform that shares more in common with PaaS than any other service models, and consequently within this context a file system is irrelevant because it is too low-level. Our concerns with respect to the distributed nature of the system persists, and reliability, availability and consistency must be ensured nonetheless.

We then resort to *distributed databases*, for which there exists different implementations⁴ that provides a plethora of features crossing the boundaries between relational databases, non-relational databases and even storage type solutions. Distributed databases provide the ability to adjust the availability of the data in response to the fluctuations of the number of incoming requests, by adding more instances to the database cluster. Distributed databases are devised around two important concepts: *replication* and *fragmentation*. The former represents the ability to replicate the data across several instances, to ensure the availability and the consistency of the data. Whereas the latter represents the ability to decompose the relations between schema entities into several sub-relations and to distribute them across several instances. After which, it is possible to reconstruct the original relations using these sub-relations (or fragments), thus balancing the workload across the instances. Much more information can be found on distributed databases, and to provide an extensive overview of all the characteristics and features of the many variations is out of the scope of this thesis, instead refer to [Lin76] [DP80] [ÖV11].

Rather than committing this architecture to a single distributed database implementation, or even to an implementation for each of these three services (*RDBMS*, *No-SQL*, and *storage solutions*), we opted for an extensible design. We provide an universal interface⁵ to these database and storage systems, and thus if the solution provided doesn't suit their need it is possible to extend the interface to account for other functionalities.

The last service comprised in this component, is *caching* and it is useful in many distributed web applications. Caching consists of a service that enables the application developer to store data in a cache for *faster* future access. It is primarily used in database-driven web applications, in order to reduce latency by keeping popular data in-memory, and thus reducing the number calls to external or physical data sources. *Memcached* is an example of a distributed caching mechanism, allowing to logically combine any unused memory in servers to form a bigger unified cache [Fit11]. It is open-source, and we can

³As a matter of fact, this is exactly what Cloud@Home proposed to implement to offer distributed file system capabilities within their infrastructure.

⁴We explore these a bit further when discussing the actual implementation in Chapter 5.

⁵We discuss this interface in greater details in Chapter 5.

provide it as a service without much problems.

In summary, we provide database and storage capabilities through the implementation of an interface, enabling the application developer to quickly integrate any open-source and readily available database or storage system implementation. We agree that caching is important for data-intensive applications, but we leave the implementation of this feature as future work due to the scope of this thesis.

4.4.3 Communication and Networking

This component is responsible for online accessibility, presentation of the information using markup languages, and the communication interface between the user and the application.

But, first, we need to underline that the problem of collaborative web hosting remains an open problem in the context of the current Internet's infrastructure. The problem revolves around the ability to name the resources in a dynamic distributed environment, but also how to provide searching and indexing capabilities in that same environment, as well as ensuring content availability. An extensive reflection already exists about this specific problem, and we diligently refer the reader to [AB14] for more information.

Therefore, we must assume that a domain is already hosted in order to access the application using a Uniform Resource Locator (URL), or it is accessed using the IP address directly. By supporting a web framework, we can provide a service to programmatically present the information of the web application.

We provide a communication interface between the application and the end-user using REpresentational State Transfer (REST) based APIs, as a result of the unanimity amongst the CSP investigated. REST, is a collection of design patterns and guidelines to create scalable web services, and it uses HTTP as the underlying communication protocol. For a complete account of the guidelines and design patterns present in REST, see [RR08].

Ultimately, we provide the services for the end-user to communicate with the web application using HTTP requests, and provide access to the underlying RESTful APIs to the end-user to interact with the web application. A web server accessible using the current Internet infrastructure receives these requests, providing means to programmatically present the information of the application to the user.

4.4.4 Load Balancing and Scalability

It is common to conflate *load balancing* and *scalability*, because of their semantic similarities relating to their purpose. In other words, both are concerned with optimizing the performance of the system, but operates at different levels. Load balancing consists of distributing the workload as evenly as possible across the different resources to optimize the resource consumption. It can be achieved by using *task queues* as a service, as shown by the CSP investigated. Whereas scalability differs in how it attempts to optimize the performance of the system. Load balancing optimizes the system performance by devising the optimal scheduling plan for the current workload. Conversely, scalability optimizes the system performance by preempting resources to respond to the fluctuations in the workload. The workload could diminish to a point that most of the resources are idling, and enforcing scalability would preempt some of the resources in order to maximize the utilization of the remaining resources, and vice-versa.

In this subsection we present *task queues*⁶, as a service or feature of a distributed computing platform. Then we present one, out of many, applicable autonomous techniques to achieve scalability in a distributed application.

Task queues are used to manage the work that happens outside of the normal request-response cycle of web applications. Tasks that are queued are handled *asynchronously*, in order to prevent interruptions or delays to the request-response cycle. Tasks can also originate from other sources, such as time-consuming maintenance operations and long-standing background processes. Here is a common workflow for using task queues within the request-response cycle:

1. Define the maximum request-response interval your application tolerates.
2. Evaluate if a task will surpass the interval.
3. Given that it surpasses the interval, schedule it in the task queue.
4. Upon completion of the task, store the result using a cache mechanism.
5. Respond, when possible, by reading the value from the cache.

Combining such a workflow with the concept of independent computational entities, such as *worker nodes* in a distributed system, monitoring the queues for any new tasks

⁶*caveat lector*: Since queues are a fundamental component of this architecture we will only address *task queues* as a service or feature to build applications, and then we will discuss the details of this architecture in further details in Chapter 5. Because it digresses too much from the service offered in the context of the application layer.

to perform, is generally sufficient to distribute the workload across the different nodes and to achieve *decent* load balancing.

These task queues are central to this architecture and this is why we do not provide them as a service in the application layer, because it would be redundant. When explicitly necessary one could extend this architecture to incorporate any task queue implementation of their choice. Although, we acknowledge that we can get into more details with respect to load balancing since sophisticated algorithms have been proposed to handle different specific cases [APHB14] [KC12], we deem it to be *application-dependent* and thus not as relevant for this minimal specification for an API.

The other aspect of this component is scalability, we rationalized scalability as being the ability to respond to the fluctuations in the workload, either manually or autonomously. The PaaS providers we have investigated, offers the ability to the user to specify *static* or *dynamic* policies. These policies usually are expressed in terms of SLA, and ensure that the different Service-Level Objective (SLO) are respected by requesting or returning resources accordingly. SLA are contracts that specifies the terms of a service between the consumer and the producer, whereas the SLO are metrics used to measure the service provisioning performance of a provider, preventing any misunderstanding between both parties. There are many important characteristics when defining SLOs and forming proper SLA, we focus on a slightly different approach to provide scalability and thus we refer the reader to [KL03] [SMJ00], for more information on service level management.

We provide scalability by implementing a decentralized autonomous controller for each type of node⁷, based on [GSL14]. In order to achieve scalability, the application developer will specify it's policies with respect to the utilization of the two types of nodes, in the form of a percentage. Using this policy we will attach an autonomous controller for each type of node, and it will be used to monitor and make the appropriate adjustments in response to the changes in the performance.




Figure 4.6: Proportional-Integral-Derivative (PID) Controller [GSL14]

The controller used is Proportional-Integral-Derivative (PID), shown in Figure 4.6, and it operates as closed-loop control system. It computes the difference between the

⁷see Section 4.1.

desired performance U_i^{SLO} , described in the SLO for that type of node, and the performance observed U_i , which is represented by $e(t)$ in this figure. It will interpret the result, $e(t)$, according to this function and return the changes required to remain within the U_i^{SLO} , denoted by $u(t)$:

$$u(t) = K_P e(t) + K_I \int_0^t e(t) dt + K_D \frac{de(t)}{dt}$$

The first component of this controller, is known as the *proportional component*, and it adjust the result in *direct proportion* of the delta between the desired performance level and the observed performance level. The second component computes the integral of the delta between the desired SLO performance and the observed performance, adjusting the result with respect to the *historical observations* relative to the delta. The last component computes the derivative of the delta between the desired SLO performance and the observed performance, attempting to *anticipate* the upcoming performance level. Three coefficients are specified to adjust the weight of each of those components, and regulate the behavior of the controller. It is then possible for the application deployer to only specified the desired utilization level for each types of nodes, and this autonomic controller will adjust the resources to respond to the fluctuations in the workload autonomously.

Ultimately, using task queues as an integral part of this architecture and leveraging decentralized autonomic controllers, we are able to provide autonomic scalability and efficient load balancing. The load-balancing component is considered to be *application-dependent*, because there are multiple factors to take into account and it poses a multi-dimensional optimization problem which is not trivially solved. This optimization problem yields different solutions in different contexts, and different applications requires different solutions. Thus, we provide some functionalities as intrinsic features of our architecture, but we leave the more sophisticated functionalities to be provided by the application developer.

4.4.5 Security

The Security component is concerned with providing means to establish/entertain *secure communication channels* between the different nodes, but also between the end-user of the application and the node responsible of handling the requests. Another major concern is access control, including *authentication* of the end-user and providing *multi-tenancy*. We present each of these concerns, and explain how they are accounted for in this API.

The first concern is to provide means for secure communication between the nodes. It can be ensured using standardized technologies such as communicating over TCP/IP

using Secure Sockets Layer (SSL). SSL certificates are usually negotiated between a web server and a client. In the context of this architecture the web server is (*usually*) hosted on the application deployer node and the client is a candidate node. Consequently, a candidate node would negotiate its certificate with the application deployer to establish a secure communication channel between them. Upon successful negotiations, the node and the application deployer are able to communicate securely using the SSL protocol.

The communications between the end-user and the web server can also be secured using standardized protocols, by communicating using HyperText Transfer Protocol Secure (HTTPS), which consists of using SSL on top of HTTP. The negotiation for SSL certificates would be done between the web server and the end-user. We use well established standard technologies (SSL, HTTPS, etc.), as would any other web application when required to communicate over unsecured networks, such as the Internet.

Access control is another important security concern, controlling access to resources based on various schemes, using *authorization* mechanisms and *authentication* mechanisms. We can identify two basic types of accesses in our architecture: *access from node to node* and *access from end-user to the web application*.

The first type of access is illustrated using this example. If we introduce a *Byzantine* node, that is a node for which the requests or responses are incorrect, either because of an error in the computations or by malicious intent, and it requests to drop all the tables in the database. How do we differentiate between this request being erroneous or malicious in intent, and a legitimate request to drop all the tables? There exists multiple schemes and strategies to enforce various level of access control [Dar14] [Gol10], and depending on the context in which the application is deployed it can vary substantially. Thus, we implement a pragmatic strategy for access control, where only the application deployer is allowed to perform administrative tasks. This is known as Mandatory Access Control (MAC) [NO95]. The identity verification is done using a simple challenge-response between the nodes to assess the identity.

The second access type, refers to the authentication of the end-users interfacing with the application deployed. Authentication, as a service for applications, *can* be provided using Single-Sign-On (SSO) technologies. For example, incorporating hooks to the Google SSO API provides the ability to the application deployer to offload the authentication responsibilities, using a *delegation protocol* such as OAuth 2.0 [JH12], to a trusted third-party. Upon authentication, the user will receive a *OAuth* token, also known as a *bearer token*. When interacting with the application, the user presents this token to demonstrate its identity as certified by the trusted third-party. These tokens are valid for a certain period of time, after which the user is required to obtain a new token. The user will

also be able to maintain a session using this token. For more information on Google's SSO API, see [Dev15].

Authentication can be application dependent. Some prefer resorting to a 3rd-party to provide authentication, whereas others may prefer holding the credential in a database and enforcing authentication themselves. It really depends on the security requirements of the application, and thus we leave the responsibility to the application developer to provide control over this type of access.

This leads to the final portion of the security component, *multi-tenancy*, consisting of the ability to provide parallel multi-user support for an application without providing each user with a dedicated instance of the application. Consequently, seasoned web developers will find using this architecture very intuitive, because instead of forcing sophisticated design patterns, resulting in a steeper learning curve, we favor the use of *sessions* to achieve multi-tenancy. By using one session per user, it is possible to encompass all the information specific to this user in a self-contained web primitive.

Finally, we provide secure inter-node communication via SSL, whereas secure communication between the end-user and the web application via HTTPS. The access control scheme provided is minimal and only the application deploying node has the ability to perform administrative tasks. We provide multi-tenancy by maintaining parallel user-sessions, which are provided using a SSO service from a trusted third-party or maintained by the web server.

4.4.6 Application Deployment and Management

This component encompasses all that pertains to the administrative tools, used to manage an application. Due to the distributed nature of our architecture, we are faced with a challenge foreign to current PaaS providers and it is related to *source-code distribution*. Every other service in this component, can be commonly found in most of the PaaS providers, notably *application configuration* and *various monitoring capabilities*.

The distribution of the source-code of the application is a security concern. Because of the distributed nature of the infrastructure it entails that a contributing node trusts the application deploying node and the source-code distributed. Sophisticated schema exist to *reduce* the trust-level assumed between parties, with regard to source-code distribution, in a distributed context.

One example consists of creating *bundles*, or atomic units, of code and data, which can then be distributed and executed [DKMyC04]. This deployment framework provides authentication and authorization mechanisms, and certifying capabilities, creating a trust

environment where source code can be distributed and executed securely. An undesirable consequence, is that it forces a complete deployment infrastructure into the architecture. Applications are required to be designed respecting a set of guidelines, leaking into the development process and impeding the flexibility of the architecture.

Instead, we propose a way of providing similar guarantees, that ultimately relies on the users judgment and effectively removes any trust factor. Because we provide an open-platform, we apply the same open-source principles⁸ to the application it hosts by adopting a *white-box* approach. The workflow for source-code distribution can be summarized as follows:

- Application deployer stores the source-code into a public or private repository supporting any version control (git, mercurial, svn).
- Upon contracting a node for contribution, the application deployer provides the location of the application repository to this node, and any credentials if necessary.
- The contributing node will then be notified of its content, by presenting the source-code in a text editor or browser.
- The contributing node is then asked whether or not it accepts to execute this piece of code (in a container).
- Finally, depending on the response, the code will or will not be downloaded into the container (and all of its dependencies).

This workflow is designed to work in a fully untrusted environment. In a trusted or semi-trusted environment (cluster or private network), we can simply turn this option off and upon contracting a node for contribution the source-code is directly downloaded into the container, implicitly assuming that the user accepts to execute its content.

The CSP investigated offers application configuration services using a web portal that lets the user adjust the different configurable parameters of their application. The configurable parameters varies depending on the CSP and are largely dependent on the underlying proprietary technologies used. We offer *application configuration* using configuration files. The configurable parameters includes: minimum number of nodes, performance policies, type of environment deployed, any communication primitives and any security primitives. Before the creation of a node instance, its configuration file is read

⁸By adopting a similar perspective to the Free Software Foundation, who promotes free software and open source development, to ensure maximal transparency [inc14a].

and interpreted. This configuration mechanism is easily extensible and could compensate for any other requirements.

Monitoring is an important feedback mechanism for the application deployer, it is primarily used to anticipate possible bottlenecks or to track the resource consumption of the application. We provide monitoring capabilities by implementing a heartbeat mechanism, in which every node periodically sends information for a collection of (user-defined) dynamic attributes. Then the information is collected and aggregated, presenting the result as system-wide monitoring information to the application deployer. This mechanism provides a sufficient perspective of the application state, with respect to the load-balancing and scalability mechanism presented earlier.

In summary, we provide a flexible workflow for code distribution adopting a *white-box* approach. We provide the ability to statically configure an application through various parameters. We provide monitoring capabilities using a heartbeat mechanism that generates system-wide information, using locally available information from the participants.

As future work, we intend to provide dynamic configuration capabilities, in which the nodes will be able to modify the configuration parameters dynamically, when applicable. Providing a visual representation of the data monitored, using graphs to represent the topology and various metrics, would also be desirable from a user experience point of view.

4.5 Component Interaction

In this section we present a way to reason about developing applications that is central to this architecture, representing the main interactions between the different components of a distributed application.

When creating an application using this architecture, it is fundamental to reason about the problem at hand in terms of the request-response cycle, that is inherent to web applications. This enables the developer to understand how to formulate a problem into a compatible solution, using the different constructs available with this architecture. We can reason about a distributed application as follows:

1. **(Incoming Web Request)** A user send a request to the web server.
2. **(Generate HTTP Request for App.)** Web server receives the incoming request, then formulates a HTTP POST request using the content of the original request, and sends it.

3. **(Incoming RESTful API Request)** A node receives a request from the web server, using a web protocol.
4. **(Create Task)** It extracts the information from the request, creates a task and queue it up.
5. **(Task Available)** Upon queuing the task, it becomes available for offloading. It is then dispatched to any available nodes.
6. **(Task Completion)** Upon completing the task, the node sends the results to the dispatcher of the task.
7. **(Task Returns)** Upon receiving the results, the original dispatcher of the task uses a web protocol to formulate a HTTP POST request containing the results, and sends it back to the web server.
8. **(Result Returns)** Finally, the web server receives the results and formulates the appropriate response to present the information back to the user.

Writing applications for this architecture forces an event-driven programming model, resulting in applications designed using the Event-Driven Architecture (EDA) approach. For a distributed application, we can identify four logical event flow layers of any EDA, as presented in [Mic06]:

- **Event Generator** is the source of (all) the events, it is the web server in the case of web applications.
- **Event Channel** is the medium used to transport events from the generator(s) to the event processing engine(s). In our architecture, the event channel corresponds to a web protocol, where the incoming events from the generator are used to derive the task(s) to be sent to the event processing engine(s). At a higher-level, we could say that the application deployer itself is the event channel, but more precisely the web protocol does the translation from raw events into events that can be processed using this architecture.
- **Event Processing** is done by the *event processing engines*, taking the appropriate actions in response to the events. In our architecture, the contributing nodes are the *engines* and they process task(s), which are a derived form of events.
- **Downstream Event-Driven Activity** is the downstream activity initiated by an event, and it occurs upon processing the event. In our architecture, it consists

of returning the results of the completed task to the node that dispatched it, and can trigger a series of collateral events.

It is possible to distinguish between different types of event flow processing, either **simple**, **stream** or **complex**. Some simpler problems impose only *one* event generator, corresponding to *simple event flow processing*. Whereas, more complex problems incur *multiple* event generators or introduce *continuous streams* of events, corresponding to *complex event flow processing* and *stream event flow processing* respectively [Mic06].

Using this kind of reasoning, it is possible to effectively separate the presentation logic from the business logic. We now have a clear understanding of the role of both types of nodes: the application deployer will contain the *event generator* and will be responsible to dispatch the events using the *event channel*. It will also use the event channel to handle any *downstream event-driven activities*. Whereas the contributing nodes will only serve as *event processing engines*. This rationalization of our architecture is crucial to ensure: proper implementation of the problem domain, an appropriate separation of the presentation logic from the application logic, and the avoidance of possible design mistakes that could impede scalability.

4.6 Discussion

In this section we present a discussion of all the layers, and how they contribute in fulfilling the requirements of both the evaluation framework, as presented in Section 3.1 and the requirements of this thesis, as presented in Section 1.2. The overall rating of the system is summarized in Table 3.1, alongside the two related projects.

4.6.1 Collaborative Peer-to-Peer System Framework Implementation

Section 3.1, presents a framework that illustrates the essential functionalities required for a peer-to-peer collaborative system to scale, by mitigating the complexities inherent to these type of systems. This subsection aims at presenting the different mechanisms implemented in our architecture to provide such functionalities.

If we recall, the framework is composed of 7 key phases relating to the life-cycle of a participant in a peer-to-peer collaborative systems, which were: *Advertise*, *Discover*, *Select*, *Match*, *Bind*, *Use*, and *Release*.

The first two phases, **Advertise** and **Discover**, which relates to the advertisement and discovery of resources using a specification, are implemented as a best-effort mech-

anism, as presented in Figure 4.7. By this we mean that the newly arrived nodes create a *Resource Specification (RS)* using a common template, and it can be represented as follows:

$$RS(node_i) = [IPaddress, Port, SA, DA] \text{ where}$$

$$SA = \{StaticAttribute_1, \dots, StaticAttribute_n\} \text{ and}$$

$$DA = \{DynamicAttribute_1, \dots, DynamicAttribute_n\}$$

Note that the set of *Static Attributes* and *Dynamic Attributes* are extensible to accommodate any desirable attributes, and are initially empty. Each application deployer publish any desired attributes to a common repository (or in the context of a DHT, append the values to a specific key used by all nodes to construct their resource specification simply known as the *template*).

Upon creating their resource specifications, the nodes gather the list of application deployers using the *Ring*, and periodically send candidacy messages, containing their resource specification. We say best-effort, because the nodes send repeatedly their candidacy messages, until an application deployer contracts them for contribution.

Other solutions propose to publish the resource specifications to a repository where it is possible for the consumer to query the repository for the most relevant resources [BJ13]. These solutions are adequate if the repository is protected against possible storage attacks against the resource specifications, this is not the case for DHTs [UPS11]. Our mechanism enables us to mitigate the possible storage attacks, without having to resort to centralized management of the repository, because we do not persist the information.

The **Select** phase relates to the selection mechanism offered to possible consumers to query the resource specifications. In this architecture, in order to select a node, an application deployer must open a TCP/IP server connection on a specific port, and evaluate the upcoming candidacy messages individually by examining its content. As a matter of fact, this selection mechanism resort to a *publish-subscribe* messaging pattern, where the nodes are the publishers and the application deployers are the subscribers. The application deployers are then allowed to subscribe only to a *meaningful* subset of the attributes published as part of a resource specification, representing a *topic*. The nodes use the concept of topics to publish the various attributes that compose the resource specifications. Then, application deployers will do a tentative selection, notifying the resources that it considers them as candidate resources. The selection is completed only after all the required resources, to host the application, are *tentatively selected*. This is a blocking operation and can be deferred to a background thread, and automated by

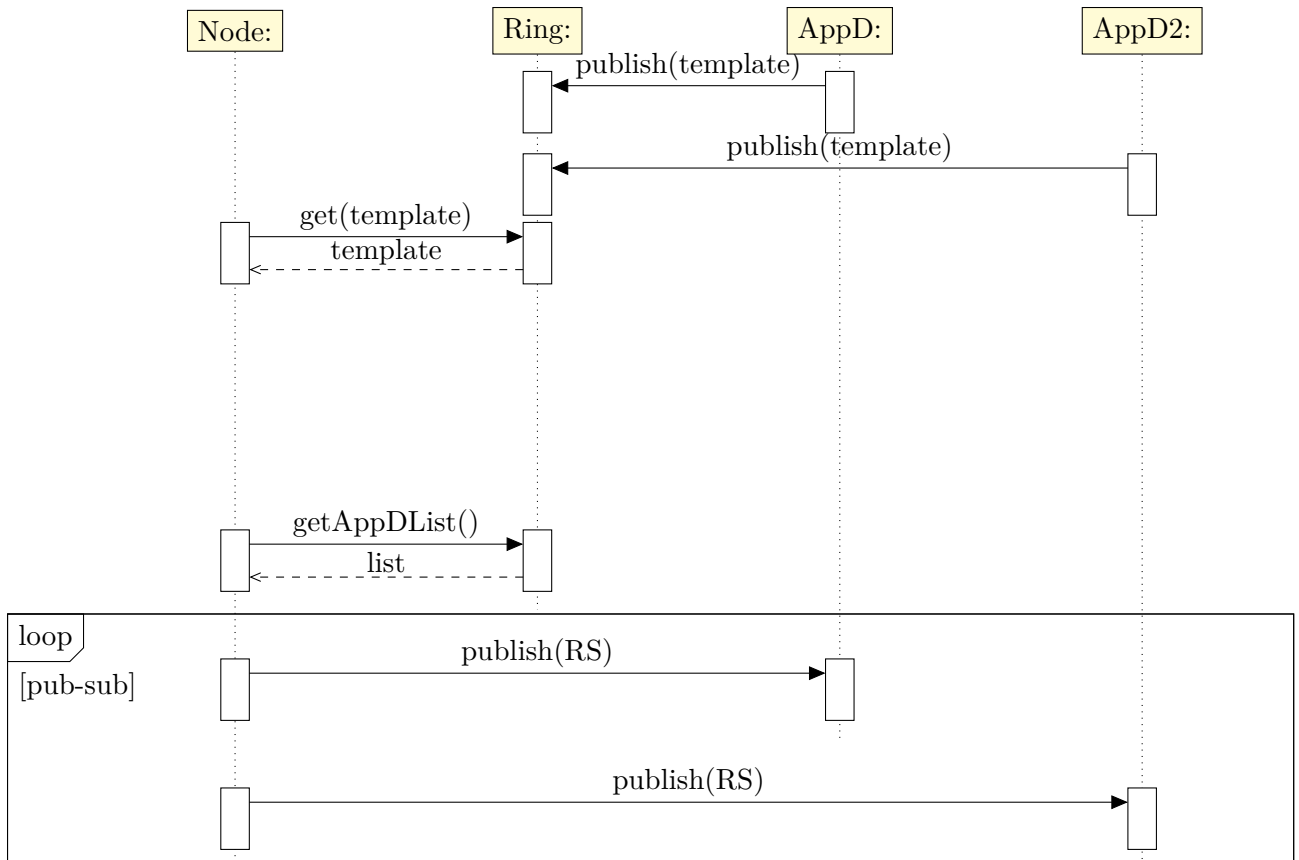


Figure 4.7: Advertise and Discover Phases Sequence Diagram.

defining a selection policy. This policy contains the desired values for the static attributes, and the intervals of desired values for the dynamic attributes.

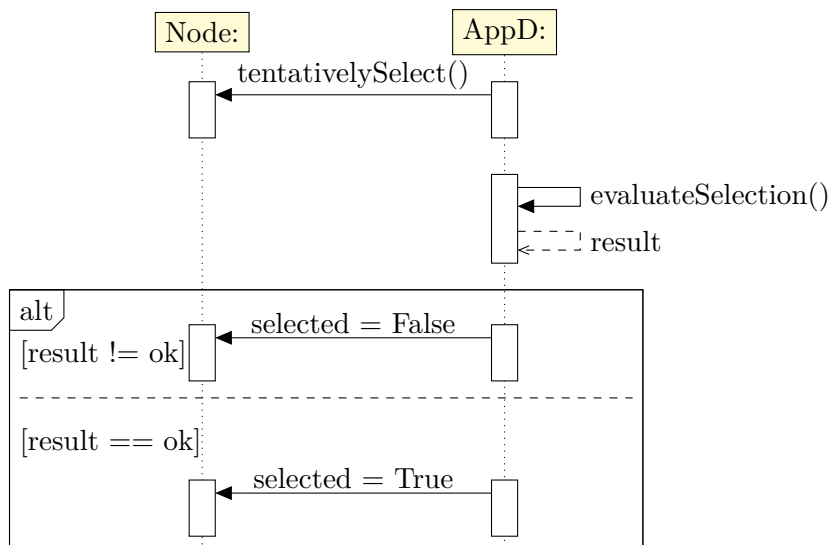


Figure 4.8: Selection and Match Phases Sequence Diagram.

The next phase is **Match**, it encompasses the ability to formally specify the inter-resource relationship requirements and to enforce them on the selection of resources, as shown in Figure 4.8. As a consequence of using a best-effort mechanism, it is possible to define these relationships as an integral part of the selection policy. For example, once a group of resources is *tentatively selected*, it is *then* possible to enforce these inter-resource relationship requirements, and repeat the process until the *tentative selection* fulfills all the requirements in the selection policy.

The **Bind** phase, is also best-effort, because of the mechanism in the *Select* phase. The priority for an application to contract a specific node is determined with respect to the time of the response to a candidacy message. In other words, applications which responses were received first are given priority, by the contributing node, over applications which responses were received later.

The **Use** phase states that it should be possible to use the resources contracted in the previous phase to execute the tasks pertaining to the application. It is accomplished by sending tasks to be executed after the enrollment process (selection, binding, handshaking, initialization of the node) has completed.

The **Release** phase is concerned with providing the capability to release a resource after contributing, either because its SLA prescribes it or because the workload has diminished to the point that this resource is no longer needed. Depending on the SLA between the application deployer and the contributing node (encompassed in a policy), release will happen if possible. Releasing a contributing node consists of a node leaving its current fellowship and returning to the ring, to advertise its resources again.

Future work includes providing more extensive functionalities for each of the phases. We do not include extend these functionalities, because the mechanisms described above are sufficient to operate the system efficiently.

4.6.2 Research Requirements

We now present how the research requirements were fulfilled in our architecture, and present a very high-level recapitulation of the various layers.

Requirement 1 states that a collection of heterogeneous devices can use this architecture to deploy multiple applications simultaneously. It is satisfied by the design of the Network layer, because we provide multi-application support using the *Ring* and we do not impose any restrictions on possible contributors, as long as they meet the light virtualization requirement of being able to host a container, which inherently all Linux

OS are capable of.

Requirement 2 states that this architecture should not impose third-parties to provide any of its services, and this architecture should be self-contained. It is partially satisfied by the Application layer, because it does not introduce any third-parties⁹ to provide any of the services described in the API. Then, we satisfy the remaining portion of this requirement, by providing the desired self-containment properties using light virtualization.

Requirement 3 states that this architecture should not introduce any single point of failure and should be fault-tolerant to ubiquitous failures of different participants. The *security* component¹⁰ of the API specification satisfies a portion of this requirement, by providing secure communication channels and enforcing MAC to regulate and control the administrative tasks. The *Ring* satisfies the remaining portion of this requirement, because it provides the ability to create and maintain a decentralized structured overlay network to connect the nodes.

Requirement 4 states that no special or dedicated equipment should be necessary to participate in this system. It is satisfied because there are no explicit or implicit restrictions on the intended hardware to be used with this architecture. Consequently, recycling the currently available hardware is the most cost-efficient and eco-efficient, *ceteris paribus*, alternative to participate in this system.

Sub-Requirement 4.1 states that the memory footprint should be sufficiently small to allow lower-end devices to be an integral part of this architecture. It is satisfied to the extent that it reduces the overhead induced by the virtualization technology to a *smaller* footprint.

Requirement 5 states that this system should provide *dynamic membership* capabilities as well as scalability to the applications that are deployed. This requirement is satisfied using task queues as a foundational construct of this system, but also by providing dynamic membership capabilities through the use of autonomic controllers to regulate the load on the resources.

We can make the following observations concerning the *Network layer*. Using the *Ring* and the *Fellowship* abstractions, we are able to separate the concerns of public-resource pooling and resource provisioning. Furthermore, we are able to provide fully isolated multi-application support using the *Fellowships*. Subsequently, the *Network layer* can be

⁹We do admit that web hosting has to be outsourced, but as we stated earlier in Section 4.4.3, it remains an open-problem because of the underlying infrastructure of the Internet. For a more complete presentation of the problem see Section 7.2.

¹⁰In no way is this component providing *absolute* security to this architecture, and more attack vectors need to be analyzed to account for a realistic threat model.

easily adapted to different environments (cluster, grid, or open Internet) without incurring any major changes, but by simply selecting which underlying networking primitive is a better fit. As an example, when operating in a highly dynamic environment (Internet), using a unstructured overlay network might be better to cope with the higher churn rate, but when operating in a moderately dynamic environment (shared-cluster) using the current DHT implementation would be more than sufficient to cope with the moderate churn rate.

The *Virtualization layer* uses lightweight virtualization technologies to abstract the resources from their underlying physical specificities, providing isolation guarantees similar to full virtualization technology. We can define the environment of an application and all its dependencies using a single configuration file, easing the portability, and creating a fully self-contained deployable construct. We are aware of the security concerns related to lightweight virtualization technologies, especially because it has direct access to the Linux Kernel's system call interfaces. This means that if a vulnerability exists in one of the system call interface, then it is possible to exploit it from within a container. This can be addressed by applying security profiles, such as Linux's *seccomp*, to restrict the breadth of accessible interfaces to a selection of safe and tested system calls. Docker applies this approach to augment the security of its containers and encourage the developer to be aware of these possible exploits [Inc15b].

The *Application layer* provides a minimal API specification for distributed computing platforms, including PaaS platforms, by providing the essential components required to develop distributed applications and an extensible API for all the other non-essential components.

Chapter 5

Implementation

In this chapter we present an implementation of the architecture proposed in the previous chapter. In order to provide context, we present the technologies used to implement the architecture and how they influenced the design, when applicable. We then present the central constructs of this architecture, followed by the algorithms and workflows used to provide the underlying logic composing this architecture. Finally, we present a very simple proof of concept that illustrates the orchestration of these constructs and algorithms.

5.1 Technology Used

We have used different technologies in concert to provide the ability to develop applications that can be deployed using this architecture.

The code was written using Python and the Twisted Event-Driven Networking Framework [Lef14]. We adopted an *event-driven programming model*, because it provides the ability to reason about problems from a non-sequential perspective, which is useful when implementing distributed applications.

Event-driven programming is conceptually single-threaded, although it supports multi-threading it is offered as an extra feature and it is not required. The primary advantage of using the event-driven programming model lies in its use of a single control thread to interleave multiple tasks, resulting in an opportunistic execution scheme. Whereas, a single-threaded program executes tasks sequentially, forcing the program to pause its execution at every blocking operation, until their completion. This sequential programming model effectively augments the total execution time, by introducing gaps in the timeline spent waiting (idling) for a blocking operation to return. Thus, the event-driven

programming model reduces the complexity inherent to the development of distributed applications, focusing on the possible events and their appropriate responses. Whereas the multi-threaded programming model, consists of using a collection of threads to accomplish different tasks, in parallel, and requires complex concurrency and synchronization mechanisms.

Ultimately, the event-driven programming model benefits from parallel execution of tasks, similarly to the multi-threaded programming model, and it combines the simplicity of the single-threaded programming model by resorting to a single control thread. Figure 5.1 illustrates the distinction between these three programming models.

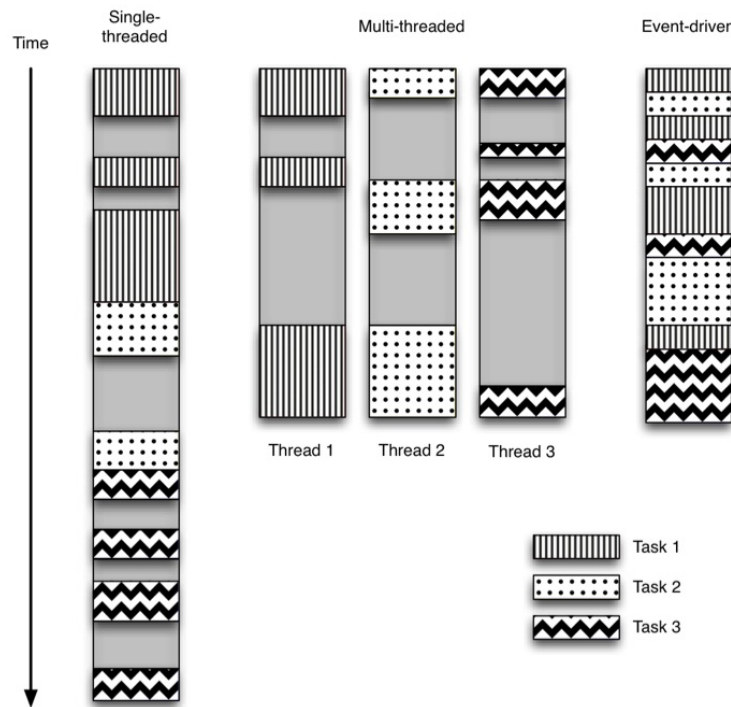


Figure 5.1: Comparison between Programming Models. [MF13]

Twisted leverages the event-driven programming model by inter-leaving multiple blocking operations and responding to consequential events, such as completion or failure, using callbacks. A *callback* is a function to call upon the occurrence of an event to handle the results. Twisted allows the developer to specify a callback to handle the successful completion of an event, and a callback to handle the failure of an event; both are encompassed into the *Deferred* abstraction, as shown in Figure 5.2. Twisted orchestrates the callbacks, contained in different deferreds, by registering them with an event-loop, called the *Reactor*. This event-loop controls the thread of execution, enabling operations to block, and keeping track of which callback is associated with which operation. Upon

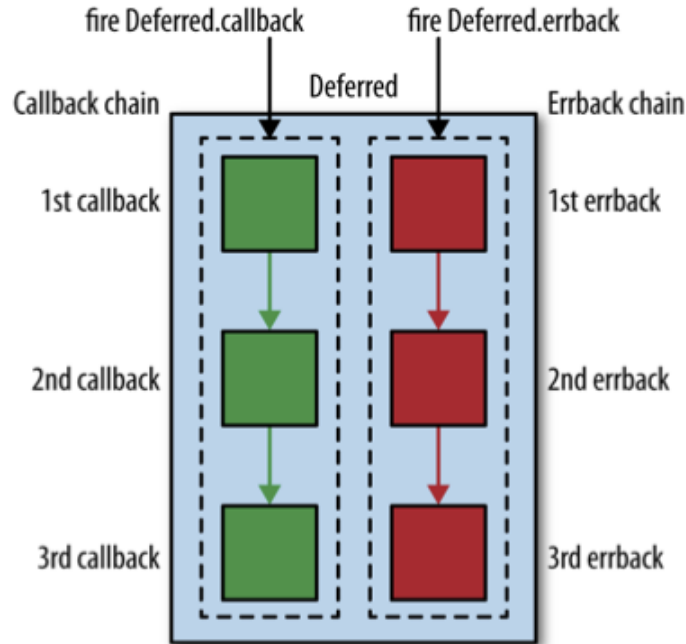


Figure 5.2: Graphical Representation of a Deferred. [MF13]

completion (or unblocking) of an operation the event-loop is notified and it passes the result to the corresponding callback, to be processed. By chaining callbacks, it is possible to achieve very complex asynchronous behavior in a cooperative fashion which is at the heart of the philosophy of the Twisted framework. The concept of deferreds influenced the design of the constructs composing this implementation of our architecture and their interactions among each other.

The overlay-network used for to *The Ring*, is a Python implementation of the Kademlia DHT using *Twisted* [Mul14]. Our decision was based on the fact that it was implemented using *Twisted*, and it could be easily integrated to our architecture incurring minimal disruptions. Another benefit, is to have uniformity in the programming-model used throughout the architecture, by focusing on the event-driven programming model.

For the web server, we used CherryPy, a minimal Python Object-Oriented Web Framework [Tea14], because of its simplicity and illustrative capabilities. It enables us to use a web server, without any bloating features and focuses strictly on what is necessary, easing the development of prototypes and proof of concepts. We can substitute this web framework with a more complete web framework, incurring only minimal changes by using our clearly defined interfaces.

To provide virtualization capabilities to the nodes we used Docker Containers [Hyk15].

Docker provides self-containing capabilities to each node, by specifying the dependencies inherent to this architecture and the dependencies of the applications developed, using *DockerFiles*. They act as configuration files, prescribing the requirements of the environment, all the dependencies, and the quantification of the resources available to execute the containers.

Using these technologies we were able to quickly and efficiently create prototypes, while constructing a solid foundation for our architecture. This foundation is the result of creating well-defined interfaces between the various components, accentuating the modularity and maximizing the extensibility of the architecture.

5.2 Constructs

In this section we present the fundamental elements of our architecture, namely the **Task**, the **ApplicationNode**, the **Ring** and the **Fellowships**. We present each of the elements and discuss their practical implementations.

5.2.1 Task

There is a need to represent a series of consequential actions resulting from an incoming request, into an easily distributable and self-contained entity. This entity *must* provide the ability to return a response to a requester, without any ambiguity regarding the originator of the request in the presence of large amount of requests.

Tasks are heavily inspired by Twisted's concept of *Deferreds*, because they make similar promises. As we have shown earlier, a *Deferred* promises to eventually return from a blocking operation with a result, and to apply a processing logic, in the form of callbacks, to this result. Similarly, a task embodies this promise of eventual completion and provides the ability to specify the result processing logic. As a matter of fact, *Deferreds* are used to chain the various processing steps for any given task. Upon the creation of a task, a *Deferred* is attached to it, consisting of the processing logic that must be applied when this task is dispatched to a node. The principal function required to process this unit of work, is the first callback to be chained in the callback chain of the *Deferred*. Then, any subsequent post-processing function will be chained to that same callback chain, representing the complete sequence of operations. A task corresponds to an atomic unit of work in this system. Analogous to the database atomic transaction property, a task can be in one of two states, processed or not processed. Task(s) are fully self-contained and stateless in the sense that any task can be dispatched to any node,

without having to synchronize the states of the nodes, receiving the task is sufficient to execute it, and carry out the corresponding sequence of operations. We distinguish between two types of tasks, a *Worker Task* and a *Data Task*. The former consists of **ANY** type of computational task, whereas the latter consists of task that pertains to persistent data (either storing or retrieving data).

Task objects are simple constructs containing the parameters, the module and the function names related to a unit of work. Once created they are dynamically linked to the module provided and the corresponding function, shown here:

```

# Task(operation)
new_task = Task("DB_STORE")
new_task.type = Task.Data
new_task._data = payload

# create(module_name,
#         func_name, data)
# where data represents the
# task data, but includes also
# any post-processing jobs as a
# list of: [module_name,
#          func_name], data
new_task.create('data_process',
                'save_file',
                [['worker_process', 'retrieveData'], filename,
                'data\binary',
                self.db_name])

```

A **task function** takes 2 parameters: a *task object* and a *list of arguments*, where the former is the instance of the task itself, and the latter is any parameters that are necessary to execute this function. Thus, a normal function can be refactored into a task function simply by modifying the signature of the function to take only the *task object* and a list representation of all the (current) parameters; and include logic to extract the parameters from the list.

Here is an example of a simple arbitrary function:

```
def arbitraryFunc(op1, op2):  
    return op1 + op2
```

And here is this function re-factored to be a task function:

```
def arbitraryFuncTask(taskObject, params):  
    op1 = params[0]  
    op2 = params[1]  
  
    taskObject.results = op1 + op2  
    taskObject.completed = True  
  
    return taskObject.completed
```

Developing applications by defining the business logic into self-contained units of work, helps mitigate a large portion of the complexity inherent to distributed systems, but imposes a fairly strict programming model on the developer. Such a programming model may be difficult to abide to in some cases, notably where the services provided by an application cannot be easily parallelized or exhibit inherent serial properties.

5.2.2 ApplicationNode

This construct represents any participating node in the network, and distinguishes between application deployers, and contributing nodes.

Application Deploying Nodes are the nodes contracting the contributing nodes in order to host an application using this system. Hosting, in this context refers to the ability to provide to a nodes the appropriate processing logic (according to their role) and to dispatch the incoming workload to the nodes accordingly. The *Application Deployer* is responsible for translating the incoming requests into tasks and for providing tasks to the contributors. Then, once a response is formulated as the result of executing a series of tasks, the *Application Deployer* returns the response to the originator of the request.

On the other hand, **Contributing Nodes** process the tasks assigned by the *Application Deployer*. Similarly to the types of tasks, such a node can adopt one of two roles: *Data Node* or *Worker Node*. The former is responsible for processing Data Task(s) and the latter is responsible for processing Worker Task(s).

5.2.3 The Ring

This abstraction was presented in Section 4.2.1, and provides the ability to connect all the participants in a public environment through a well-defined interface. It also enables the *application deployers* to find potential *contributing nodes*, and it provides a public repository for the common resource specification template. Using a DHT, we are able to connect all the nodes together in a public environment, the Internet. This construct serves as a public meeting space for contributing nodes and application deploying nodes. In order for nodes to interact with(in) the Ring, we have defined a network interface that declares the following functions:

- **bootstrap()**: defines the bootstrapping mechanism for the Ring allowing to configure any newly arriving node.
- **connect()**: defines the connection procedure once a node has been configured, in order to join the Ring.
- **set()**: defines how to store a value in the Ring.
- **get()**: defines how to retrieve a value from the Ring.

Application Deployers collectively define the Resource Specification (RS) template and include any desirable attributes by appending them to the template stored in a public repository. This public repository corresponds to the *template key* stored in the DHT using *set()*. Contributing Nodes are then capable of advertising their resources according to the template, by retrieving it from the DHT using *get()*. Upon contracting all the necessary resources, the collection of nodes (including the application deployer) will form a *Fellowship*.

We have decided to use a DHT, because of its storage capabilities. Although, as we have presented in the previous chapter, we could supplant the DHT with any overlay network of our choice to satisfy any (other) networking requirements, such as better response to very dynamic networking environment by using an unstructured overlay network [LCP⁺05]. This could be done utilizing the network interface we have defined, and by redefining the mandatory functions.

5.2.4 The Fellowships

This construct provides the ability to connect the selected nodes with the application deployer in a private networking environment. It ensures privacy for the data transmitted, but also enforces any security policy between the nodes.

Currently we use a *whitelisting* mechanism to provide a private environment, meaning that only the nodes contained in the list defined by the application deployer are allowed to participate in this environment. This list is the result of selecting the resources, and publishing the list of these resources under the application deployer's corresponding key in the DHT. Participating nodes are then able to retrieve this list, and use as a provisional measure for validating incoming communications. More elaborate security schemes can be devised to satisfy a variety of security requirements, such as implementing a public-key infrastructure using a DHT [LLC11]. Due to time constraints, we have not implemented this construct as a stand-alone overlay network as prescribed in Chapter 4, and instead uses the whitelisting approach to security. As future work we would implement Fellowships using a protocol similar to T-Man [JMB09], to provide an overlay network; and we would explore more complete security schemes.

5.2.5 Conclusions

By using these constructs we enforce the adoption of a task oriented programming model from the application developer's perspective, which is fully compatible with the event-driven programming model used to construct this architecture. We are required to have centralized public point of access, due to the underlying Internet infrastructure, which force us to distinguish between two types of nodes, Application Deploying nodes and Contributing nodes. We are providing a flexible public networking platform using a well defined interface, that provides the ability to change the peer-to-peer networking primitive of the platform effortlessly. Finally, we are also providing a private networking environment, in which the nodes hosting the deployed applications can communicate between them, using a whitelisting mechanism.

5.3 Workflows and Protocols

In this section we present the procedure to initialize a node, and the protocol necessary to interact and be part of this architecture. We present the initialization workflow of a node, and illustrate the differences between the workflows of an Application Developer's node and a Contributing node. Then we present the protocol used by the nodes to communicate within a Fellowship. We conclude with the presentation of the web protocol interfacing with the web server, followed by a brief discussion of the error-handling capabilities of the various protocols.

5.3.1 Initializing a Node

A node follows a specific initialization procedure, which varies slightly depending on the type of the node. We distinguish between the Contributing node's workflow and the Application Deployer's node workflow.

A contributing node iterates through this progression of steps, forming a sequential workflow:

1. Node creation.
2. Connect to the Ring.
3. Retrieve Resource Specification Template.
4. Initiate Advertisement mechanism, and wait until selected.
5. Start communication with the Fellowship.

Whereas the Application Deployer's node workflow contains additional steps:

1. Node creation.
2. Connect to the Ring.
3. Read configuration file and update Resource Specification Template.
4. Initiate resource Seeking mechanism, and select adequate nodes.
5. Publish the list of nodes selected, to the Ring.
6. Start the web server, and start receiving requests.
7. Start communication with the Fellowship.

Using these two workflows we are able to create a node and successfully join the system. We are also able to deploy applications, by selecting the nodes that satisfies our requirements and joining them together to form a Fellowship.

5.3.2 Fellowship's Protocol

The communications inside a *Fellowship*, are conducted according to the **Fellowship's Protocol**. This protocol defines 6 different actions, that can be taken. Similar to client-server protocols, here we depict the contributing node as being a client and the application

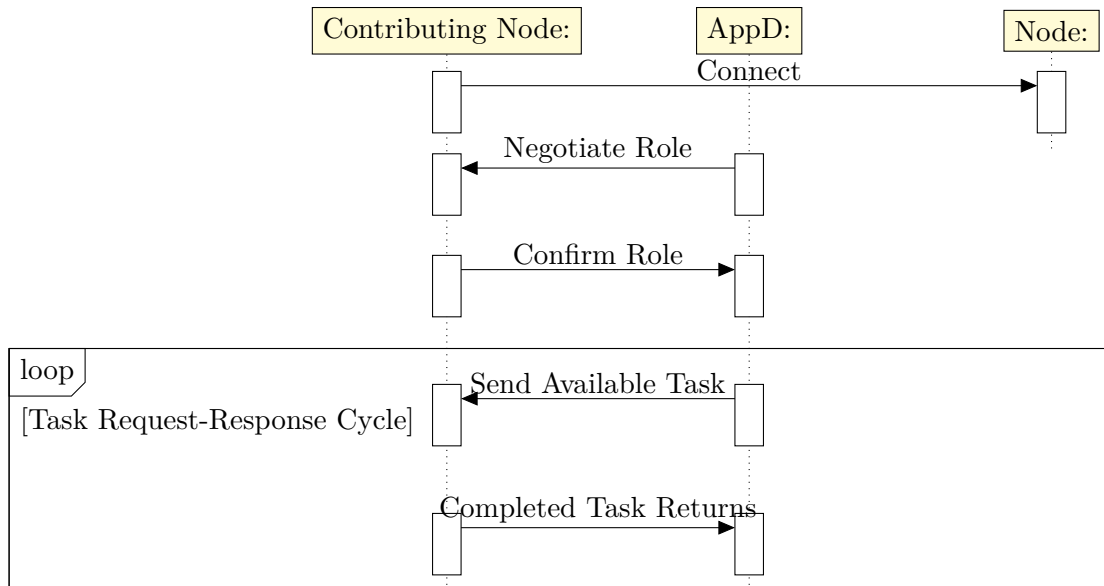


Figure 5.3: The Fellowship Protocol

deploying node as being a server¹.

As illustrated in Figure 5.3, a *Contributing Node* **connects** to any node in the *Fellowship*. After successfully connecting to the network, it will **negotiate** a role with the *Application Deployer*, either this node will become a *Data Node* or a *Worker Node*. Then, after **confirming** its role, the *Contributing Node* will begin the **Task Request-Response Cycle** and it will become an active node of the *Fellowship*.

5.3.3 Web Protocol

The **Web Protocol** is used to translate incoming web requests into tasks, by implementing a RESTful API between the application and the web server, which enables complete separation from the web hosting and processing facilities.

Translation of web requests into tasks, consists of receiving a web request and then the web server generates the appropriate RESTful request to be sent to the application. For all intents and purposes, any web request is translated into a POST request containing the information about the task, such as which module and function to call, and the arguments to call it with. Upon receiving the request from the web server, the application creates a task using the information contained in the request. When a task is completed, the result is retrieved by the application deployer node, who formulates a HTTP POST

¹The application deploying node is depicted as a server since it hosts the web server.

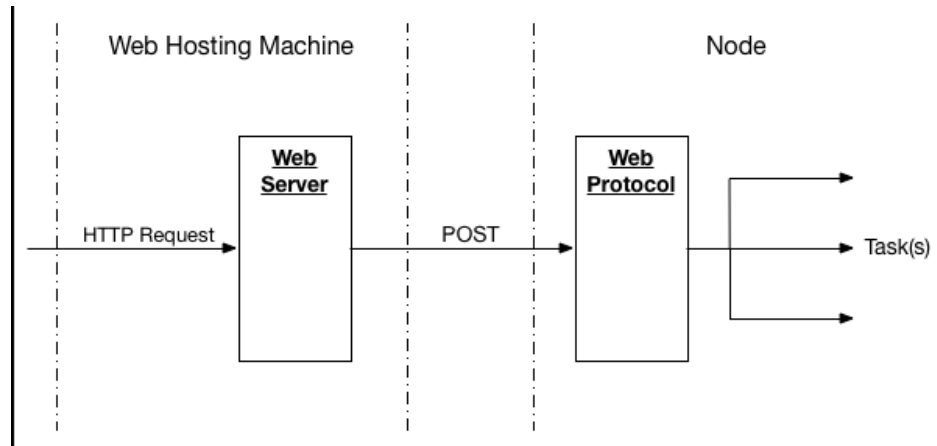


Figure 5.4: Web Request Translation into Task(s).

request using the web protocol and sends it to the web server. The web server then is responsible to present the information back to the originator of the initial request.

The rationale behind this, is to completely decouple the presentation logic from the application logic, thus maximizing extensibility and modularity. Any web framework can be used with this architecture, and it can be hosted on any web service provider, as long as it is able to emit/receive HTTP requests, it is compatible. It is possible to have more elaborate web hosting scenarios, where multiple web servers are spun and they all emit HTTP requests to a subset of nodes of an application to be translated into tasks. Meticulous synchronization is required to ensure the responses are sent back to the corresponding requesters in that case, but by using session primitives it is trivial to locate the originating web server, and consequently the originator of the request.

This component is crucial to the claims of extensibility that this architecture makes. It is also important in order to achieve scalability for the application deployed, where web hosting can be an important bottleneck. Lastly, it provides the foundation to implement robust web applications, and to provide fault-tolerant mechanisms, such as redundant web servers, without incurring any major changes.

5.3.4 Error-Handling

In the various protocols presented above, we've omitted a crucial component, that is: error-handling, which varies slightly from one protocol to another. This system being implemented as a prototype, it does not cover the handling of every single error-case nor of all the edge-cases, but rather covers the handling of the most obvious errors. As part of the second iteration of the development lifecycle of this system, the protocols will be implemented as finite state machines, but for illustrative purposes these verbose

protocols are sufficient to demonstrate the error-handling capabilities.

The first protocol covers the steps to initialize a node, be it a consumer or a contributor, in both cases the obvious potential errors are: connection failures between consumer/contributor and vice-versa. The connection failures could be relative to a connection timeout, or a connection drop. To circumscribe these possible errors, both parties define a timeout value and a maximum number of retries. Using the timeout value a consumer or contributor can determine whether an error occurred, similarly for the maximum number of retries, after a number of unsuccessful exchanges, consumer or contributors can then deem this destination node to be in an error-state. Once a node is deemed to be in an error-state, if it is the application deployer then it is necessary to bring this node back to an error-free state in order to proceed with the deployment of the application, whereas if it is a contributing node, then it is simply ignored and the application deployer attempts to communicate with a different contributing node.

The *Fellowship's Protocol* on the other hand is a bit more complex in terms of error-handling, because it is also vulnerable to connection failures, but also to computation failures. In this case computation failures refer to a completed task with erroneous results, or an uncompleted task. The former can be handled by performing sanity checks on the parameters, on the intermediate results and final results. All of which can be easily integrated in the application development workflow using Twisted's Deferreds, since every deferred encompasses 2 chains of callbacks: **callbacks** and **errbacks**, see Figure 5.2. The latter relates to the return of an uncompleted task, or the non-return of a completed task. In the case of the return of an uncompleted task, it is possible to handle this error on the application deployer side by enforcing mandatory sanity check(s) for all the completed tasks, and if the check(s) fails simply re-queue the task in a best-effort fashion. If task does not return in a timely manner (as prescribed by the application deployer), then it is considered incomplete and it is re-queued.

Similarly for the *Web Protocol*, using Twisted's Deferreds, it is possible to define the appropriate behavior of the node to handle any error-cases. Again, this discussion about error-handling is a preliminary discussion that paves the way to an exhaustive error-handling analysis, which takes part in the next iteration of the development lifecycle of this system.

5.4 Proof of Concept: Calculator

In this section we present our very first proof of concept application using this architecture. We demonstrate how to create a sensible solution, by reasoning about the problem

at hand in terms of an Event-Driven Architecture (EDA) and illustrate the interactions between the different components. We present the role of each type of nodes, their responsibilities, and present conclusions related to this proof of concept.

5.4.1 Overview

The application is a simple web-based calculator. It takes two operands, and using a button signifying the operator to be applied, computes the result using the resources available to the application (nodes). Four different operators are implemented: the addition (+), the subtraction (-), the multiplication (*) and the division (/).

Implementing this application, we make the following assumptions:

- Web Hosting is the responsibility of the application deployer, thus it will host the web server.
- No data is persistent, thus it will not require any Data Node(s) (or database).
- At least one Worker Node(s). The more nodes available, the more concurrent requests can be satisfied.

We have implemented the web server, as mentioned previously, using CherryPy. We present a minimalistic web interface that contains 2 fields for each possible operators, 8 fields in total and 4 buttons, as illustrated in Figure 5.5.

Cloud Computation Execution -- Yet Another Calculator!

Use this page to input the number you wish to compute, then let the cloud do it for you!

<input type="text" value="8"/>	<input type="text" value="10"/>	<input type="button" value="Add [+]"/>
<input type="text" value="8"/>	<input type="text" value="10"/>	<input type="button" value="Subtract [-]"/>
<input type="text" value="8"/>	<input type="text" value="10"/>	<input type="button" value="Multiply [*]"/>
<input type="text" value="8"/>	<input type="text" value="10"/>	<input type="button" value="Divide [/]"/>

Figure 5.5: Calculator's Web Interface

And the results are presented using a simple string, consisting of: *"The result is :"* followed by the result of the computation.

5.4.2 Application Deploying Node

In this subsection we present the implementation of the Application Deploying Node for the Calculator application.

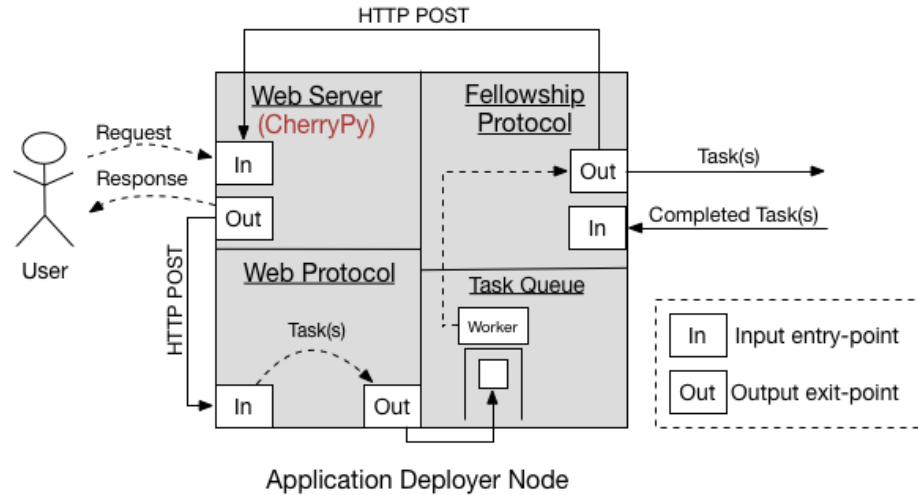


Figure 5.6: Application Deployer Node as part of the Calculator Application.

This node encapsulate the sole event generator and the event channel of this application. When an event is generated (from the web server), it is passed to the event channel (web protocol) where it is translated into a task and then queued. Using the Fellowship protocol, this node will dispatch the task(s) to the available nodes and collect the result(s) of the completed task(s). The results are extracted from the completed task(s), and are presented to the end-user, using a HTTP POST, as a downstream event-driven activity emerging from the processing of the event, as shown in Figure 5.6.

The application deploying node is responsible for the task queues in this specific application, because we host the web server on that same node and there are little to no value in distributing task queues among several nodes in such a simplistic example.

5.4.3 Contributing Node

In this subsection we present the implementation of a *Contributing Node*, and then the logic necessary to process the tasks.

A contributing node receives a task from the application deployer node, executes it, and returns the completed task to the application deployer node. All the communication between the application deployer node and the contributing node is performed using the Fellowship protocol.

In the context of this application, we have defined the logic to process the various tasks into a module called *worker process*. It encapsulates all the logic necessary to perform any operation on the two operands, and facilitates the addition of new operations, by simply adding a new corresponding functions. It makes for a clear separation between the

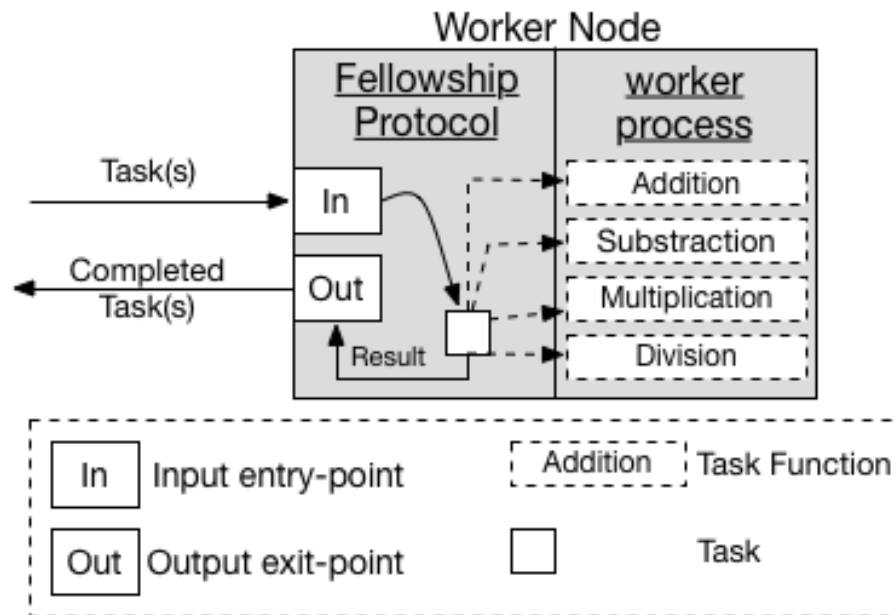


Figure 5.7: Contributing Node as part of the Calculator Application.

application-dependent code and the architecture related code, and it augments portability of the source code.

In order to execute a task, the *contributing node* will call the appropriate function inside the module and pass it any parameters necessary using a list of parameters. The *contributing node* will append the results to the task object, passed in parameters, and reply to the application deployer with the completed task, as shown in Figure 5.7.

5.5 Conclusion

We can observe, in Figure 5.8, the rudimentary workflow of this architecture: receive a request, translate it into a task, dispatch the task, execute the task, return the results, and respond back to the originator of request. It exemplify the request-response cycle perfectly, corresponding to an Event-Driven Architecture. Resorting to this way of reasoning is intuitive when developing web-applications. It is a central characteristic of this system. This way of reasoning facilitates the development of new applications, but also considerably reduces the effort required to translate an existing monolithic application into a distributed application, as we will see in the following chapter.

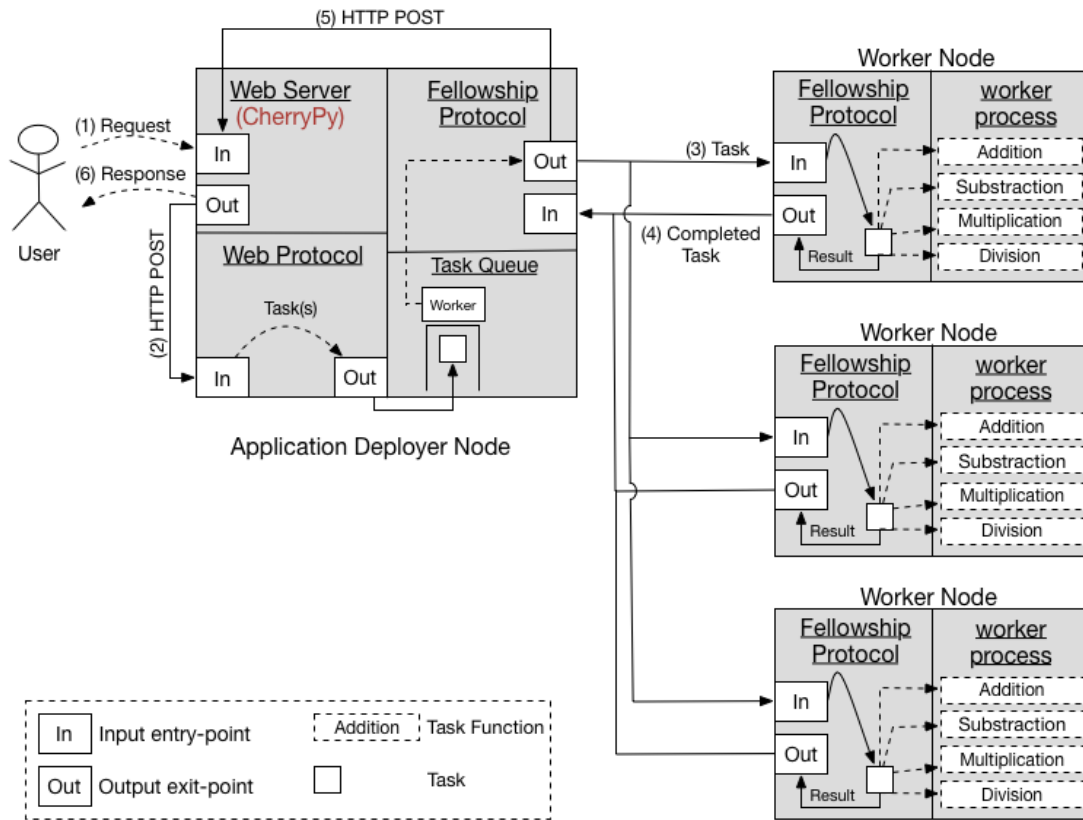


Figure 5.8: Calculator Application Architecture Overview.

Chapter 6

Case Study: Multi-Document Text Summarization using Genetic Algorithms

In the previous chapter, we have introduced the technologies used and the fundamental constructs of this architecture. Then, we have presented as a proof of concept, a calculator, in terms of our architecture and its implementation. In this chapter, we are presenting an example case study for a more computationally intensive application. Namely we present an adaptation¹ of a system for *Multi-Document Text Summarization using Genetic Algorithms* [QHH08].

In the following sections, we present the problem at hand and illustrate how a genetic algorithm is capable of addressing it. We present how we extend the existing work on text summarization using Genetic Algorithm (GA), in order to process multiple documents. We then use an analysis-based approach to translate this problem into a distributed web application. This analysis will serve as the foundation to devise an algorithm that is compatible with an event driven architecture. Finally, we present the characteristics and roles of each type of nodes, and we present a discussion of this implementation.

¹By no means, is this an efficient multi-document text summarization system. But it embodies a computationally intensive highly parallelizable problem which is appealing to showcase this architecture.

6.1 Problem at Hand

In order to provide some context², we present what *Automatic Summarization* is in general. Then we present the characteristics of *Multiple-Document Text Summarization*, followed by a high-level explanation on how to use *Genetic Algorithms* to solve this problem.

This section demonstrates how to transform a monolithic serial program into a distributed web-based application, including all the steps required to rationalize the transformation.

6.1.1 Automatic Summarization

To understand how to automate the summarization of documents, we must understand what constitutes a summary. A summary can be defined as: the product of a reductive process to include salient statements of a source text, by selecting or abstracting the original statements, forming a condensed representation of the argumentations and conclusions drawn [J⁺99].

To generate a summary, one of the techniques that can be used relies on the *extraction* of linguistic units from a document to represent the information conveyed in it. This technique is known as *extraction-based summarization*. Another popular technique that can be used to generate a summary, consists of *abstracting* the sections from a source text and generate statements, using natural language generation techniques, that convey the same information in a condensed form. This technique is known as *abstraction-based summarization*.

For this case study, we are interested in extraction-based summarization, more specifically we are interested in a technique called *Sentence Extraction*. Generally, this technique consists of scoring sentences in a text for a set of given metrics, and apply statistical heuristics to filter superfluous or not-so relevant sentences.

Up until now, we have defined a summary has being derived from a single source text, but what if we are presented with a set of documents on a single topic, can we summarize the aggregation of all these texts into a single summary? Doing this will provide the ability to aggregate all the different perspectives of each text, but also to reveal the overlapping perspectives concerning a singular topic. This variation of automatic sum-

²*Caveat lector*: extensive literature has been written on both, Automatic Summarization and Evolutionary Algorithms, and may be consulted for more information: see [MM99] as a starting point for Automatic Summarization, and see [ES08] as an introduction to Evolutionary Algorithms.

marization corresponds to *Multi-Document Text Summarization (MDTS)* [GMCK00].

6.1.2 Genetic Algorithms

The problem of creating a summary, using sentence extraction, can be defined as the problem of extracting the most salient sentences of a source text, and use them to generate a summary of a given length. We can interpret this problem as an optimization problem and apply an evolutionary algorithm to search the solution space for near-optimal solutions, as proposed by the authors of [QHH08] for single document text summarization.

A genetic algorithm, is defined as an adaptative heuristic search algorithm inspired by biological evolution processes, such as natural selection, mutations, cross-overs, etc. [ES08]. A very high-level overview of the outline of a simple GA is presented in Algorithm 1. There are several considerations to take into account while using GAs,

Algorithm 1 Genetic Algorithm: An Overview

```

1: procedure GA
2:   Initialize population:
3:     population = randomPopulation()
4:   Evaluate Population:
5:     fitness = fitnessFunction(population)
6:   while fitness <= desired fitness do
7:     Evaluate Population.
8:     Select parents for next generation.
9:     Perform Crossovers on the parents selected.
10:    Perform Mutations on the resulting population.

```

especially on how to represent an individual in the population, which type of crossovers should be applied, enforcing elitism or not in the selection process, the type of mutations to apply and how to formulate a meaningful fitness function. But this digression is not relevant in the context of this thesis and thus for more information see [ES08].

6.1.3 Extensions Proposed for MDTS

We have proposed extensions to the work of [QHH08], in order to provide multi-document summarization capabilities. But first, we present a high-level overview of their procedure, and then we present our proposed extensions.

Algorithm 2 is conceptually intuitive. The document is represented using a directed acyclic graph to preserve the order of the sentences in the text. In this graph, the sentences are the vertices and the similarity between two sentences is represented by a

Algorithm 2 Automatic Text Summarization using GA: An Overview

```
1: procedure ATS-GA
2:   Represent document as a Directed Acyclic Graph.
3:   Compute similarity metrics and weight of the sentences.
4:   Initialize population:
5:   population = randomPopulation()
6:   Evaluate Population:
7:   fitness = fitnessFunction(population)
8:   while fitness <= desired fitness do
9:     Evaluate Population.
10:    Select parents for next generation.
11:    Perform Crossovers on the parents selected.
12:    Perform Mutations on the resulting population.
13:  Extract the summary from the graph.
```

directed edge from one sentence to another, where the direction of the edge corresponds to the order of the sentences, ensuring a continuous progression in the traversal of the graph. The problem then becomes finding a path that traverses the graph, containing a number of vertices that corresponds exactly to the desired summary length, maximizing the fitness function. Some difficulties can be encountered when tuning the parameters of the genetic algorithm, but it is a problem inherent to GAs in general and not to this specific problem of automatic summarization.

We propose to use this methodology integrally, and execute it in parallel on multiple documents sharing a common topic in order to generate a collection of summaries. This collection of summaries is then used as a corpus of salient sentences. Instead of using the GA to find the final summary, we use a redundancy reduction function to remove excessively similar sentences and to achieve the desired summary length. Thus, our new and updated algorithm, able to summarize multiple document on a single topic is presented in Algorithm 3.

The rationale for this extension is the following: given that the genetic algorithms finds a near-optimal summary for each document independently, then by collecting all of these summaries together we have the most salient sentences (in relation to the topic) across all documents. Then, we remove the number of less salient sentences for which the redundancy score was the highest. The score was computed with respect to the number of overlapping words between two sentences, and the closeness to the topic.

Algorithm 3 Multi-Document Text Summarization using GA: An Overview

```

1: procedure MDTS-GA
2:   for each document in collection of documents do
3:     Represent document as a Directed Acyclic Graph.
4:     Compute similarity metrics and weight of the sentences.
5: Initialize population:
6:   population = randomPopulation()
7: Evaluate Population:
8:   fitness = fitnessFunction(population)
9:   while fitness <= desired fitness do
10:    Evaluate Population.
11:    Select parents for next generation.
12:    Perform Crossovers on the parents selected.
13:    Perform Mutations on the resulting population.
14:  while len(summary) > desired summary length do
15:    Apply redundancy reduction algorithm on collection of summaries.
16:  Extract the summary from the graph.

```

6.2 Translation of the Problem for this Architecture

In this section, we analyze the MDTS system, and translate it into a distributed application for our architecture. First, we state the assumptions, then the node requirements. Then we identify the logical event flow layers, and present the different possible event flows. Finally, we present the resulting workflow of the application.

To start our analysis, we need to make the following assumptions:

- We need to persist (some) data. At least one *Data Node* is required.
- We require an arbitrary number of computational resources. At least one *Worker Node* is required.
- Provide support for multi-tenancy.

Based on the assumptions above we can draft out the node requirements for a minimal implementation of this case study: at least one Application Deploying Node, one Data Node, and one Worker Node.

We can distinguish the two following alternatives: (1) The resulting EDA contains the sole event generator, the web server; and (2) The resulting EDA contains two event generators, the web server and the data node(s). The latter alternative distinguishes between a HTTP request received from the web server, resulting in a data task to store

its content, and the worker task created as a result of the completion of the data task. Whereas, the former doesn't, and rather it proposes to leave these subsequent tasks as part of the downstream event-driven activities.

The second alternative conflates the definition bestowed upon an event in this architecture with a task or self-contained unit of work. In this architecture, an event incurs the creation of a series of tasks representing the sum total of the work required to account for this event. In the context of web applications, we can draw the following analogy between the event and the request issued by the end-user: A request is received by the web server and it must formulate a response to this request, either by executing a number of operations or retrieving information from a data source. In this analogy the fashion in which the response is formulated is irrelevant to the end-user, and it can be perceived as a *black-box*, similarly for an event from the EDA's point of view. Everything that happens inside this black-box, as a result from this event, is irrelevant in terms of the EDA. All that matters is that the action(s) taken in response to this events yields the desired results, whether this chain of action is composed of an arbitrary number of links it is still initiated by an event, and ultimately concludes by formulating the appropriate system-response to this event. We can now position the concept of downstream **event-driven** activities as the links in the chain of actions undertook to formulate a response to this event, and thus these *activities* cannot originate from an event generator.

We can now appreciate why the first possible translation fits naturally within the context of this architecture, because an event corresponds to a user request, or the initiator of the chain of activities to be performed in order to formulate the appropriate system-response to this request, and upon completion no other activities are generated for this specific event, marking its termination. In more complex examples involving a multi-layered SOA, it may be justified to further the categorization of events with exterior events and interior events for this architecture. An exterior event is generated from an event-source that is at the edge of the application, i.e., a web server. Whereas an interior event originates from one service provider to another inside the SOA.

As a rule of thumb, one should always try to trace the computational steps of any chain of activities and identify its origins. If the originator of the chain of activities is a request from the end user, then it is safe to assume that any of the subsequent links in this chain are downstream event-driven activities. Conversely, if the current step is the originator, and tracing it back would take you outside the system, then it is safe to assume that this is an event. Using this approach it is trivial to reason about what should be an *event generator*, an *event channel*, an *event processing engine* and a *downstream event-driven activity*.

Now, we can identify the logical event flow layers according to our rationalization:

- **Event Generator(s):** in this case the initial events are generated by the web server (incoming request(s) from the user(s)).
- **Event Channel(s):** the web protocol is the event channel.
- **Event Processing Engine(s):** events communicated through the web protocol (event channel) are processed by a data node.
- **Downstream Event-Driven Activities:** once a data node completes the persistence of the data relative to the user's request, the application logic is applied as downstream event-driven activities. This triggers the creation of a series of tasks that will carry the execution of the application, involving data and worker nodes, and finally returning the result to the web server.

We can distinguish between three event flows from these logical event flow layers: (1) *Incoming Request Event Flow*, (2) *Processing a Request Event Flow* and (3) *Returning a Response Event Flow*.

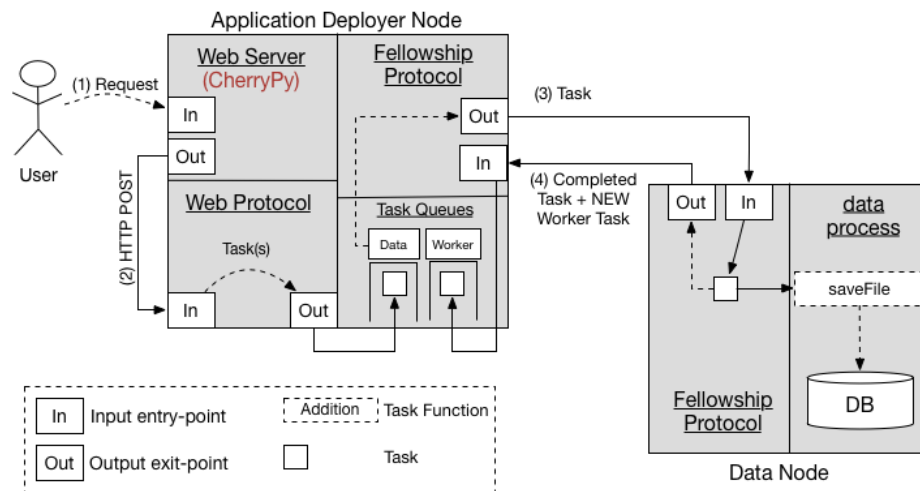


Figure 6.1: Event Flow: Incoming Request.

The **Incoming Request Event Flow**, as shown in Figure 6.1, embodies the logic of receiving a request from a user, persisting the data (documents) of this request in the database, and queuing up a processing task for this data.

The **Processing a Request Event Flow**, as shown in Figure 6.2, corresponds to processing a task related to a document, which implies retrieving the data (documents)

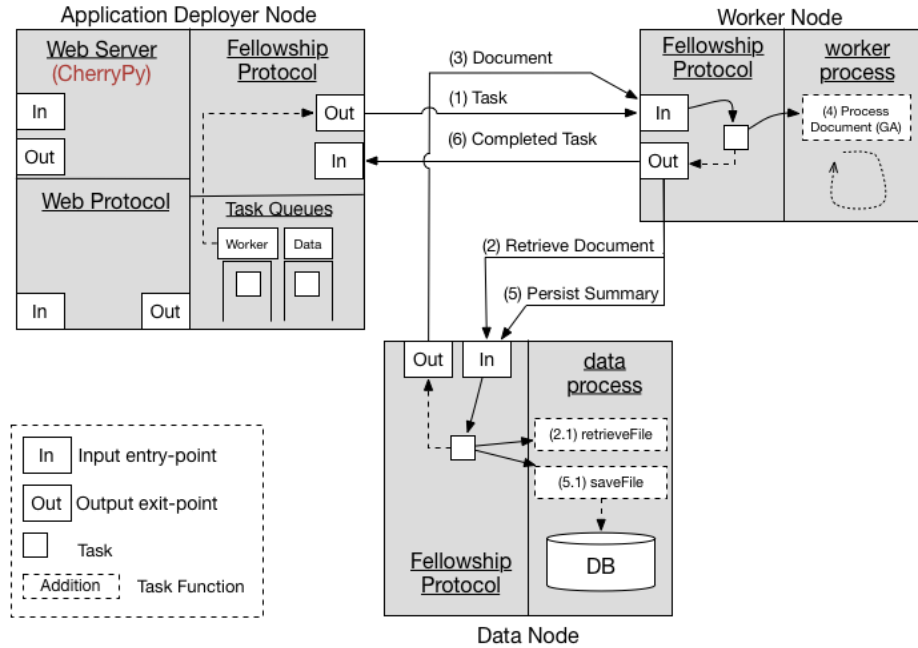


Figure 6.2: Event Flow: Processing a Request.

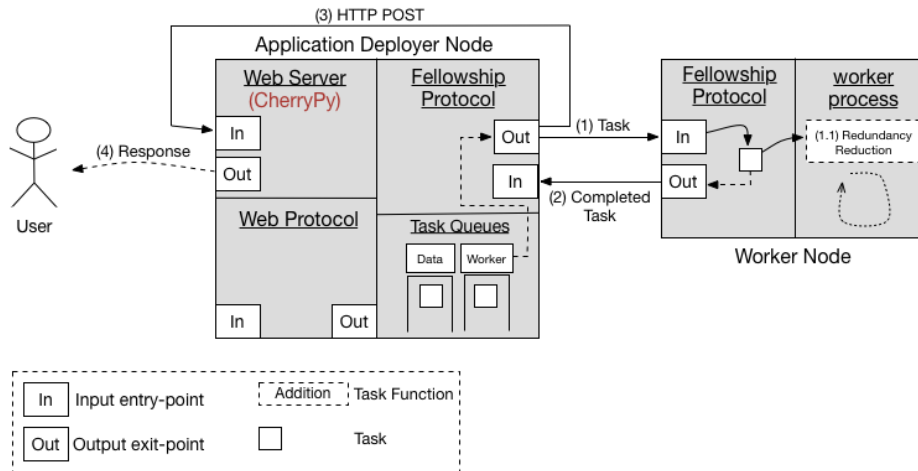


Figure 6.3: Event Flow: Returning a Response.

from the database, processing it using the genetic algorithm, and persisting the results back into the database.

The **Returning a Response Event Flow**, as shown in Figure 6.3, consists of monitoring the database for all the results to be persisted using a database trigger³ and then creating and dispatching a task to process the results by applying the redundancy reduction algorithm. Upon completion of the task, the results are sent back to the web

³This is not shown in the figure, the worker task has already been created and queued in the figure.

server and it presents them to the end-user.

Based on this analysis we devise a general workflow that encompasses all the interactions between the nodes and the end-user concerning a single request, as shown in Figure 6.4. Multitenancy is provided through the use of sessions, and we use these session primitives to organize the data in the database for each user.

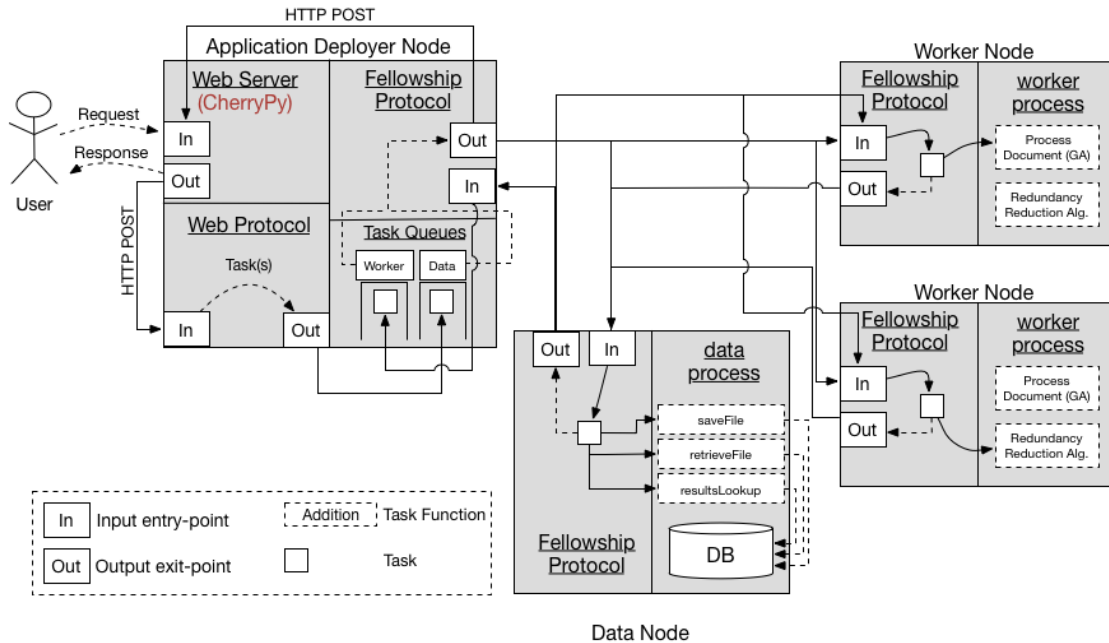


Figure 6.4: MDTs: General Workflow.

6.3 Implementation Details

In this section we present the implementation details for each type of nodes. And we outline the specificities of this particular application in the context of our architecture.

6.3.1 Application Deploying Node

The application deploying node is responsible for receiving the requests from the users and translating them into initial data tasks. It is also responsible for the task queues, because we host the web server on that same node and there are little to no value in distributing task queues among several nodes in this example. Because we are using data node(s) we need an additional task queue, thus we have one task queue for the worker tasks and one for the data tasks. Ultimately, the application deploying node is

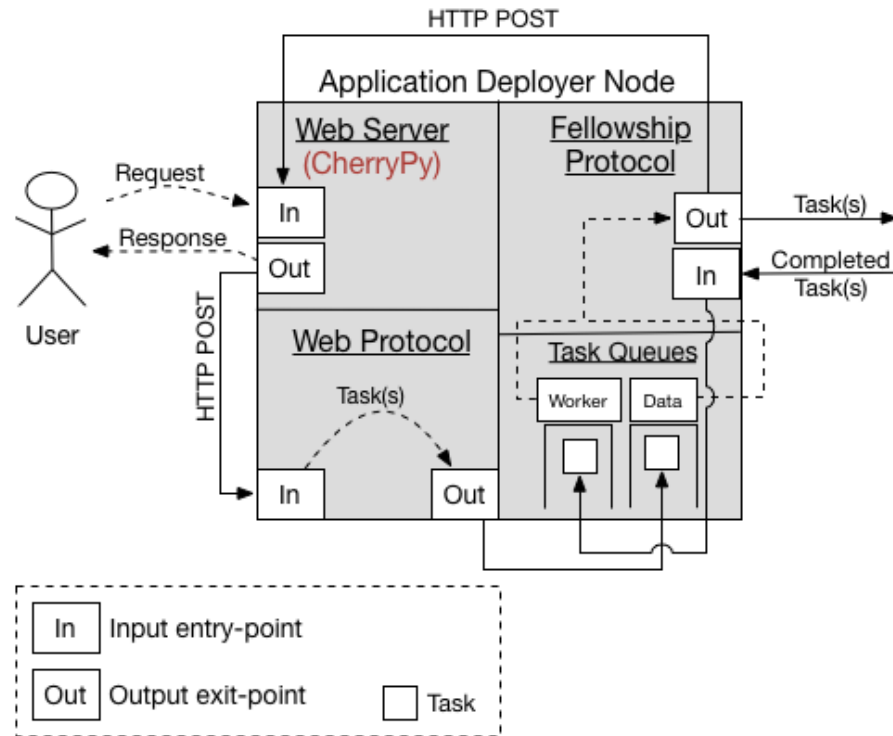


Figure 6.5: Application Deployer Node as part of the MDTs Application.

responsible for dispatching any tasks contained in its queues, to any available nodes, as shown in Figure 6.5. Another characteristic of this implementation of the application deploying node, is the task creation logic resulting from a completed data task. As part of the core logic of this architecture, when a data task completes the data node can specify a post-processing task to be scheduled as the result of this data task. We leverage this capability to generate the downstream event-driven chain of activities, representing the application logic. We have implemented the web server, as mentioned previously, using CherryPy. In Figure 6.6, we present a minimalistic web interface, where the users can upload their files, specify the different parameters of the genetic algorithm and the summarization parameters.

The results are presented using a simple string, consisting of: *"The result is :"* followed by the topic of the documents and the resulting summary. There is nothing inherently different between this version of the application deploying node and the version presented in the previous chapter, see Section 5.4.2, aside from the addition of a data task queue. It still hosts the web server and it is still responsible for the interactions between the contributing nodes and the end-users.

Multi-Document Text Summarization using Genetic Algorithm!

Use this page to upload a collection of (similar) documents to generate a text summary.

Filename: No file chosen

Genetic Algorithm: Fitness Function Parameters

Cohesion Factor Proportion:
 Topic Relation Factor Proportion:
 Readability Factor Proportion:

Genetic Algorithm Parameters

Population Size:
 Maximum Number of Generations:

Summarization Parameters

Desired Summary Length:

Figure 6.6: MDTs Application's Web Interface

6.3.2 Data Node

Data nodes are responsible for persisting the data of the application, as shown in Figure 6.7, which is achieved by using a relational database, such as *RethinkDB* [Ope15].

We grouped all of the database functionalities into a single module named *data process*, and thus if we would want to replace the implementation of the database, only some of the functions need to be redefined in this module. It defines a *DataProcess* instance, corresponding to a long-standing process created for the database server instance. The *DataProcess* instance contains all the information necessary to access the database, such as the client driver port and the authentication key. But it also provides the facilities to start and stop the database server instance gracefully, through eponymous functions.

We created a simple database schema, to represent the data that needs to be persisted, and we make a distinction between two types of data: *files*, shown in Figure 6.8, and *results*, shown in Figure 6.9.

The task functionalities are implemented in the format defined in Section 5.2.1, using a list of parameters and passing it alongside the task object. We defined three different functions corresponding to the actions to be performed on the database.

The first function is *saveFile*, it consists of extracting a file's content and name, from a task object and storing it in the database. Each user has its own database, allowing to store very large sets of related documents. If it is the first file to be stored, the database and the table corresponding to the *files* table schema are created. Upon completion of the task containing the *saveFile* function, a *worker* task is created to retrieve the file

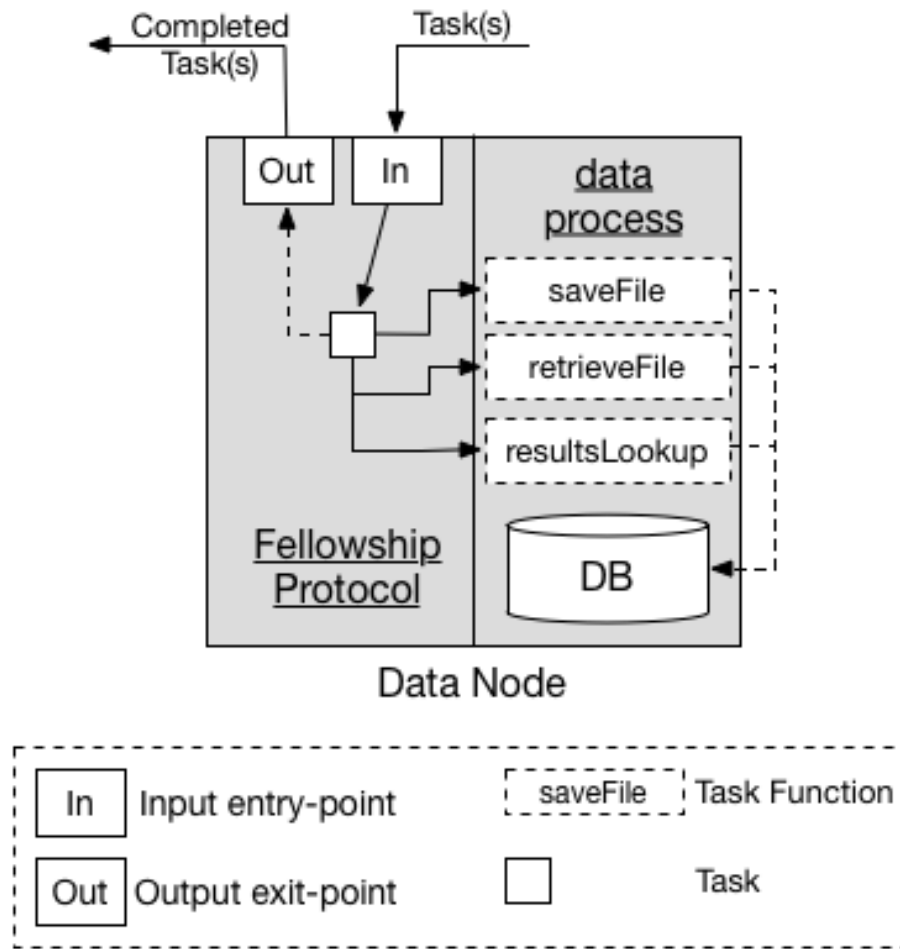


Figure 6.7: Data Node as part of the MDTS Application.




Figure 6.8: Files table schema.

from the database, to process it (generate a summary), and to commit the results back to the database.

The second function is *retrieveFile* and it consists of retrieving a file from the database, for a worker node, to be processed. The function logic formulates a query to retrieve a

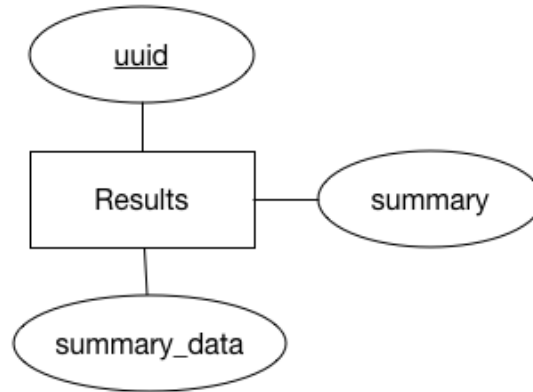


Figure 6.9: Results table schema.

file from the appropriate database, i.e., the document prescribed in the worker task.

The third function is *resultsLookup* and it consists of creating a trigger that monitors the changes in the *results* table. It can be done using the primitives offered by *RethinkDB*, such as the *changes()* function, which returns a stream of all the changes made to a table or database [Ope15]. When all the results have been written to the database, the data node creates a task indicating that the results are ready to be processed.

It is possible to distribute the database across a cluster of data nodes, and we can leverage the clustering capabilities of the DBMS to do it. Sharding and Replication are automatically taken care of by the DBMS, but can be configured to cater to different requirements [Ope15]. To leverage the clustering capabilities, it is only required to specify the argument *-bind all* when starting the database server instance of the first data node, then any other data node can join the cluster by specifying that argument and the *-join* argument with the IP address of the first data node. Thus, to augment the number of data nodes of **ANY** application, simply specify the corresponding arguments when instantiating the data nodes.

Data nodes are simply database servers, augmented to function within this architecture. Their functionalities oscillate around the Search-Create-Replace-Update-Delete (SCRUD) operations. Augmenting availability of the data can be done effortlessly, on-the-fly by recruiting new nodes and creating additional database server instances using the proper arguments.

6.3.3 Worker Node

Worker nodes are the central components to the business logic of any application hosted using this system, because they form the computational resources of the system. The

re-factoring effort required to transform any function into a task function, is minimal.

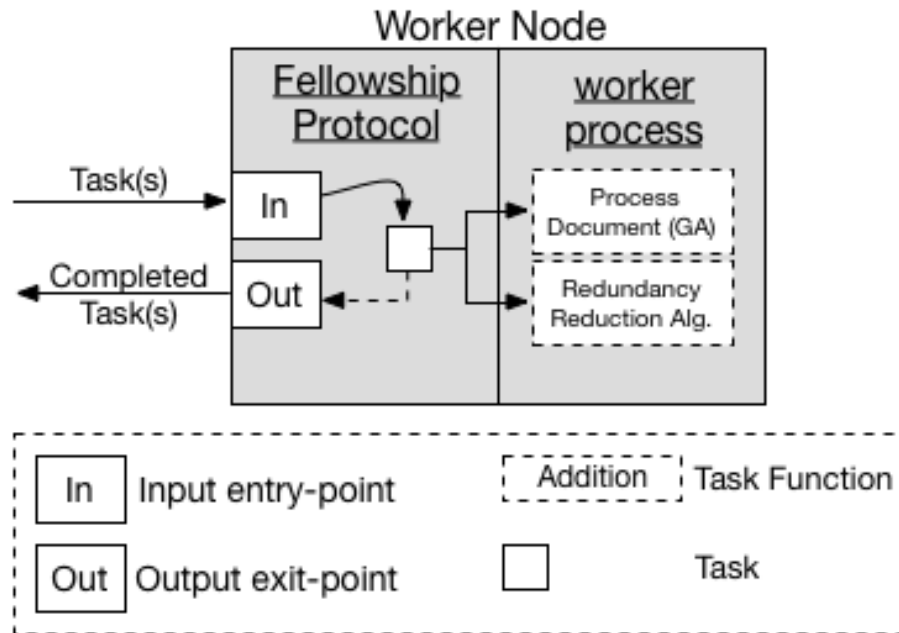


Figure 6.10: Worker Node as part of the MDTs Application.

Similarly to the Calculator example, we have encapsulated all the application logic into one single module, named *worker process*. As shown in Figure 6.10, this module exposes two task functionalities⁴, again using the format defined in Section 5.2.1.

The first function, *retrieveData*, as its name implies is responsible for retrieving the data from the database, to create a summary, and to write the results back into the database. Our original summarization code required only one modification, and it was to tailor the code to process only one document (instead of a directory). Ultimately, this function contains the logic to retrieve a document from the database, and calls each of the summarization functions in their respective order.

The second function, *retrieveDataPostProcessing*, retrieves the results from the database, and applies the redundancy reduction algorithm, then sends the resulting summary back to application deploying node to be presented back to the end-user.

⁴In all the figures containing a worker node in this case study, we changed the task function names to better portray their functionalities and thus the function name for *Process Document (GA)* in the module corresponds to *retrieveData*. Similarly, the function name for *Redundancy Reduction Alg.* in the module corresponds to *retrieveDataPostProcessing*.

6.4 Experimentation

In this section, we present the experimental setup used to test and develop the MDTS application for this architecture. We conclude with a discussion of our evaluation and its limitations.

Our experimental setup consists of:

- One laptop hosting two VMs, where the three are connected using the local loopback adapter.
 - The laptop is an Apple MacBook Pro (15-inch, 2.53GHz, 8GB of RAM, Mid 2009) running OS X Yosemite (10.10.2).
 - We use Oracle VM VirtualBox version 4.3.12 to host the VMs on the laptop.
 - Both VMs are identical, both were allotted 512MB of base memory, 12Mb of video memory, 8 GB SSDs and are running Lubuntu 14.04.

As per our initial assumptions, we must have at least one application deployer node, one worker node, and one data node. The VMs correspond to the contributing nodes, and the laptop corresponds to the application deployer node. The contributing node’s copy of the application differs from the application deployer copy, because they are not allowed to instantiate (themselves as) application deploying nodes. The application deployer is running a version of this application on MacOS X, whereas the contributing nodes are running their applications inside a Docker container. We used **valgrind** instrumentation framework, more precisely *massif*, a tool to generate the heap memory profile of applications, to profile the application running outside a container, when summarizing 9 documents.

n	time(i)	total(B)	useful-heap(B)	extra-heap(B)	stacks(B)
0	0	0	0	0	0
1	160,800,345	2,789,560	2,722,880	66,680	0
2	319,964,709	4,559,048	4,464,383	94,665	0
3	432,207,257	4,655,904	4,567,228	88,676	0
4	580,771,287	6,908,216	6,740,605	167,611	0
5	711,952,664	9,685,952	9,479,920	206,032	0
6	800,083,101	10,210,536	9,999,586	210,950	0
7	934,281,944	15,176,552	14,942,006	234,546	0
8	1,111,714,521	16,942,128	16,704,676	237,452	0

Table 6.1: First 8 Snapshots: Heap Memory Usage Peaking at Snapshot 8.




Figure 6.11: MDT5 Application Heap Memory Profile.

Figure 6.11 shows the graphical representation of the memory profile, generated using the `ms_print` tool from the **valgrind** framework. We can observe that this application heap memory profile peaks at 16.94MB, shown in Table 6.1, and this is used as a baseline to profile the memory usage of this application. We say baseline in this context, because it is the smallest memory footprint for this application independently of the virtualization technology used. Next, we identify how much memory is required to run this application inside a container⁵ using the `free` Unix command to capture the amount of memory used before and after the creation of the container, shown in Table 6.2 and Table 6.3 respectively.

We can now compute the delta between the memory used before and after, or $489096KB - 481064KB = 8032KB$, then we extrapolate this result, *naively*, to deduce that it requires $8032KB$, or around 8MB of RAM to create and execute a idling Docker container of Ubuntu 14.04. Thus, our guesstimation of the memory footprint of this application

⁵Docker was wrote using the emerging Golang (or Go) language, and it is not trivial to generate memory profiles for Golang binaries because of the immaturity of the language. No tools are available for memory profiling, thus we improvise this evaluation scheme, which consists of finding the heap memory profile of the *bare* application and add the amount of memory used to create and run the container to provide a very high-level guesstimation of the actual memory footprint.

	total	used	free	shared	buffers	cached
Memory (KB)	501832	481064	20768	3020	76564	161628
-/+ buffers/cache (KB)	242872	258960				
Swap (KB)	522236	348	521888			

Table 6.2: Memory Statistics Before Container Creation and Execution.

	total	used	free	shared	buffers	cached
Memory (KB)	501832	489096	12736	3180	76672	162052
-/+ buffers/cache (KB)	250372	251460				
Swap (KB)	522236	348	521888			

Table 6.3: Memory Statistics After Container Creation with Active Container.

running in a Ubuntu 14.04 Docker container is computed by adding the peak value of the heap memory profile to the memory required to create and execute the container, and is equal to $8032KB + 16942KB = 24974KB$.

Profiling the application footprint using this improvised scheme has its limitations, and we resort to this methodology to provide a glimpse of the *potential* memory footprint in a real application. Optimizations to the core of the architecture framework could definitely be made to augment the performance, and effectively reduce the memory footprint. Further experimentations and tests are required to provide a more accurate memory profile. One interesting test would remove the use of Docker altogether and resort to the Linux containerization primitives directly, then we could profile the memory footprint of minimal light virtualization technologies, effectively removing Docker’s overhead and approach our $16942KB$ baseline.

6.5 Conclusions

By reasoning using the event-driven architecture approach, we were able to circumscribe stateful dependencies in the event flows and mitigate them, enforcing the task properties of being stateless and persisting data to the database. This is a central aspect of this approach used to translate monolithic applications into distributed applications in accordance to our architecture. It also provides the tools necessary to divide the application logic into meaningful atomic units of work. We have presented a way of dividing an existing application into tasks, but it is not the only way to do so. As a matter of fact, a developer has the ability to divide the application into many meaningful task, spanning multiple modules, because tasks instances are dynamically linked to their functionalities. Thus, it is possible to use different module names and function names for the

functionalities of the Worker nodes and the Data nodes.

All the components of this case study were written using the Python programming language. The amount of modifications and coding required to translate the application was very small, only 99 additional lines of code to transform the summarization logic into worker task functionalities and 164 lines of code for the `DataProcess` class and its (data) task functionalities. The *Web Protocol* was written in 113 lines of code, and the web server was written in 145 lines of code.

The *Fellowship Protocol* is application independent, and is written in 514 lines of code. Similarly for the core components of this system totalizing 556 lines of code, including all the logic pertaining to the nodes and the network interface to the *Ring*.

These numbers include the comments, and it showcases how minimal is the translation from a monolithic application to a distributed application, incurring a total of 521 lines of code for the application-dependent components.

Using a primitive experimentation environment, we profiled this application with a baseline heap memory footprint of $16942KB$ and deduced that *at least 8MB* of additional memory was required to execute this application in a container. Further experimentations are required, and due to the immaturity of the toolset for Golang binaries it may be necessary to revert to LxC, or to use directly the containerization primitives to gather an accurate memory and performance profile of this application, but also of the architecture in general.

Chapter 7

Discussion

In this chapter we present various musings about our architecture and the problems we have encountered. We initiate the discussion by presenting the positioning of this project amongst the wide landscape of distributed computing platforms. Then conclude by presenting various open problems that were encountered along this research path, including peer-to-peer networking primitives for collaborative systems, the state of the current Internet infrastructure, and ethical problems inherent to public collaborative systems.

7.1 Positioning of the System

In this section we present the positioning of this system amidst the sea of different distributed computing platform and systems. It is necessary, because the purpose of many of these systems gets conflated together, and it provokes a severe misunderstanding of what they provide. This is especially true with cloud computing.

Cloud computing has become a *buzzword* nowadays, which implies that it becomes the go-to solution for all problems. Even today, a majority of professionals are still kept in the dark amidst all these promises and features, that very few are capable of providing a concise definition of what cloud computing entails [JM14] [CS14]. We don't blame them. We agree that the information circulating around about this technology is, confusing, at best.

On the other hand, many researchers have tried to devise a proper ontology for this *new* paradigm, see [YBDS08]. This attempt made us realize that the definition of cloud computing can be rationalized, and with the help of documents like [MG11] it is possible to clearly define cloud computing. But if we attempt to complete this already expansive definition, by defining what the end-product of cloud computing is, this term *cloud*

computing becomes an *umbrella* term and can be used to define any distributed system [Cit12] [WR15].

Thus, we present this system as being part of a sub-class of cloud computing, referred to as *volunteer cloud computing*, offering a PaaS-*like* computing platform. Notice how we *italicized* the preposition *like*, we put emphasis on the fact that this system is **not** aimed at providing a volunteer cloud computing PaaS service model, but rather a computing platform akin to it. What sets them apart, is the environment in which they are intended to operate.

Our system is meant to operate in a semi-trusted environment, in which the contributors (usually) knows (in)directly the application deployers, and thus there exist some sort of trust relationship between the two. For example, if we refer to a medium to large size secondary school, such an establishment could be the host of upwards to 200-300 computers. These computers are used during the day for pedagogical purposes, but remains idling or off at night. The school IT administrator, could then use this system to provide these 200-300 computers to academics from a local university wishing to run scientific experiments. This example prescribes a semi-trusted environment, because the academics are in a trust relationship with the IT administrator. We qualify this environment as *semi-trusted*, because both consumers and producers trust each others, but are required to communicate using untrusted networks, such as the Internet.

Another example, in the context of local communities, could be a chess club, amongst any other clubs. Given that the club contains a small community of members, in the range of 100-150 members, this system could be used to host their web application. Such a web application could contain the statistics about the players, recorded games, event announcements, and any other desirable features to chess players. Then a subset of the community, at the very least 10 members, could contribute their household computing resources to host their web application, and to serve their community without incurring any extra cost ¹. Two or three nodes could host the database, to provide fault-tolerance in terms of reliability, availability and consistency. Whereas one node could host the web server, and delegate two other nodes to take over in case of failures, exposing only 3 IP addresses. These nodes are not required to host a domain, and could simply connect directly using the IP addresses, but for convenience sake's lets assume that they register a domain name to one of these IP addresses. The remaining 3-4 nodes can be used as worker nodes, to respond to incoming requests from the web server. Again, this is a *semi-trusted* environment because the producers and consumers have a pre-existing trust

¹Given that their computers are usually turned on, if not then one should factor in the cost of electricity.

relationship, but operates in a public environment. The only cost incurred by this chess club would be the domain registration, and it is optional.

Consequently, we emphasize that this computing platform caters to a specific niche of potential users, but it empowers them to recycle currently available resources and to re-factor these resources with a new purpose. All of this, while maintaining a focus on providing an intuitive and rational way to port such applications, using the event-driven programming model and the Python programming language.

7.2 Problems Encountered

We have encountered many problems while conducting this research effort, and we present the problems that remains open and require further research.

Peer-to-Peer Networking Primitives and Collaborative Systems

The first problem we have encountered was related to the peer-to-peer networking primitives and the collaborative systems, more precisely how well suited are the current networking primitives to be used in very-large scale collaborative systems. Networking primitive in this context relates to overlay networks connecting all the peers in a collaborative system.

One publication debated this question and the authors of this publication concluded that there are no peer-to-peer networking primitive, out of the 12 investigated, that caters to all the requirements of a peer-to-peer collaborative system, or in other words a peer-to-peer networking primitive that addresses the resource collaboration problem [BJ13]. We have presented these requirements, in Section 3.1, they encompass the lifecycle of a contributing resource in a collaborative system, and present an adequate starting point for the development of such systems. Out of the 12 peer-to-peer networking primitives investigated, there were 4 unstructured overlay network-based solutions and 8 structured overlay network-based solutions. All of these solutions were mutually exclusive with respect to their structure and functionalities. The authors are currently working on a solution that would account for all their requirements and address the resource collaboration problem, which has yet to be published.

The major difficulty outlined, was the mutually exclusive characteristics of some requirements. Such as the ability to perform deterministic queries and to perform MADQ, which seems to be completely incompatible because they occur in different networking primitives exclusively. Similarly for mutable object support and the deterministic querying capabilities of structured primitives. Faced with this problem, we decided to create

an architecture focused on extensibility, to be able to supplant the underlying networking primitive at any desirable point in time, whether because a new and more efficient networking primitive has been devised or simply to port this system to a different environment.

Collaborative systems imposes several design constraints, as a consequence of the environment in which these systems operate, but also because of the collaborative nature of these systems, which requires trust between the participants. We think that this is why this problem is difficult, but also why it is interesting. This operating environment is paradoxical, because it requires the participants to trust other participants to just the right extent, for it to be still beneficial to offload the workload to the participants.

Consequently, novel peer-to-peer networking primitives are required for this specific class of distributing systems, if not to fulfill all the requirements [BJ13] but to cater to this semi-trusted computing environment, emerging out of peer-to-peer collaborative systems.

Collaborative web hosting

As we already mentioned in Section 4.4.3, the current Internet infrastructure is not adequate for collaborative web hosting. Furthermore, the authors of [AB14], present an account of these deficiencies and propose a solution to mitigate them.

They outline the fact that these deficiencies arise from two characteristics inherent to peer-to-peer systems. First, is the volatility of the resources and the dynamic nature of the networking infrastructure in peer-to-peer system, whereas the current Internet infrastructure is stable and *somewhat* static. The second characteristic emerges for the constant migration of the data among peers to ensure availability and consistency, meaning that one peer may have some data for some time, and then it goes offline, it no longer hosts data and it delegates (indirectly) to another peer.

The *naming* scheme employed by the Internet infrastructure is the first deficiency presented, it refers to the Domain Naming System (DNS). By definition DNS, operates by resolving an IP address from a URL, using a map of URLs and IP addresses. This is not adequate in the dynamic environment of peer-to-peer systems, because there are no guarantee that the IP address associated with a specific URL corresponds to a live node, possibly attempting to access this inaccessible resource.

The second deficiency presented relates to *searching and indexing* resources. It is traditionally done using a centralized search engine that provides the service to locate any resource on the web. Whereas, in peer-to-peer systems indexing is (usually) distributed and done on a voluntary basis.

The last deficiency presented relates to the *content availability*. The current web

infrastructure relies on web servers to be online continuously. Whereas in peer-to-peer systems, the nodes join and leave the system constantly. It is not easy to provide similar availability guarantees using peer-to-peer systems and it is even worst in fully decentralized systems. These guarantees are essential for the usability of a collaborative web hosting infrastructure.

This provides a glimpse into the limitations of the current web infrastructure with regards to collaborative web hosting, and again for a full and comprehensive account followed by a candidate solution please refer to [AB14].

Ultimately, if an application needs to be accessible from the web, it must commit to at least one static IP address for the DNS to include this application as a contactable resource on the web. This effectively induces a single point of failure into this web-based resource, namely its IP address, exposing the resource to potential Denial of Service types of attacks. This showcases the immaturity of the Internet as a platform, and how initial design decisions still transpires to this day by exposing the limitation of the design. Although, we must acknowledge that collaborative web hosting is *possible*, using technologies such as Dynamic DNS (DDNS), it is not as intuitive and it reduces usability. DDNS also requires the web server to communicate with a 3rd-party using a client to update the corresponding DNS entries, and thus it may incur a monthly fee for this service. As a result it is more sensible when designing a system such as this one, to leave these details to the application developer's discretion.

Ethical use of Computational Resources

The last problem we have encountered is a criticism for providing open access to collections of computational resources.

One of the underlying factor that drove this research effort, was to explore the possibility for the consumer to emancipate himself from current cloud service providers and rather, leverage resources from the Internet as a community. We were naively inspired by this utopic vision of a computing platform that promotes: ecological awareness by recycling rather than consuming new resources, censorship-resistance by resorting to decentralized topologies and community by sharing *private* resources.

Upon further investigation, and helpful comments from other researchers we were faced with the reality that this vision was in fact utopic, because of which we reposition ourselves to preserve our vision but in a more fitting context. We position our system to operate in a semi-trusted or semi-private environment, which is more realistic and provides the liberty to enact our vision. By resorting to whitelisting mechanisms we are able to limit access to the system to only a subset of trusted peers. Again, as stated in Section 7.1, this results in a semi-trusted environment, and reduces the risk

of malicious nodes harnessing the computational resources to commit malicious, even illegal, actions. One example, applicable in a untrusted environment without whitelisting mechanisms, is a node that contracts large amount of contributing nodes and rather than using these resources to deploy legitimate application, this malicious node uses the computational resources to raise a bot-net army to perform Distributed Denial-of-Service (DDoS) attacks.

Thus, the question raised by this anecdote gravitates around the capacity for the ethical use of freely available computational resources in an environment such as the Internet. This relates to a field known as *cyberethics*, being highly debated in the early days of the Internet, some guidelines were devised such as [NWG89] or [Bar92]. The guidelines states that wasting computational resources, disrupting the functioning of the Internet, gaining unauthorized access to resources, or violating privacy of the users are all deemed unethical.

Now, after quarter-century of continuously presenting similar ethical guidelines to the users, it seems that the state of the *open* Internet is more unethical than ever. What gives?

This question is more akin to a philosophical inquiry, and its root cause as well as its potential solutions digress too much for the scope of this thesis. Nonetheless, the implications and consequences of this question are a real impeding factor when researching *and* developing Internet related technologies. For more information on cyberethics, see [Spi10] or [Tav10].

Chapter 8

Conclusion

This thesis presented a distributed computing platform, similar to the PaaS computing platform offered by cloud service providers, which is composed of volunteered resources rather than dedicated resources. We have illustrated how to provide such a platform without introducing any additional resources, but rather by simply recycling the resources that are available. As a consequence, the computing platform we have proposed is well suited for a variety of devices including lower-end computational resources.

8.1 Requirements Fulfilled

In the context of this research effort, we focused our attention on two sets of requirements simultaneously: (1) The *evaluation requirements*, see Section 3.1; and (2) The *functional requirements*, see Section 1.2. The former set of requirements originates from [BJ13], and contain the essential phases for a contributing resource in a collaborative system. Whereas the latter set of requirements, are the functional requirements that we devised for the intended application of this system.

We consider the five requirements to be fulfilled as follows:

Requirement 1 Was completely fulfilled by devising the *Ring* abstraction, which provides the means to connect the participants and to manage the application deployment. It effectively supports the deployment of multiple applications simultaneously.

Requirement 2 Was fulfilled, because the API proposed for the *Application layer* does not introduce and/or force any third-party to provide any of the services it encompasses. Opting for light virtualization, we enforced the self-containment property, not only to the system-level but also to the application-level.

Requirement 3 The *security* component of the proposed API covers very basic threat models common to all web-based applications, including access control, authentication and secure communication channels. Also, by choosing a DHT as the networking primitive for the *Ring* abstraction, we obtain a decentralized peer-to-peer networking topology.

Requirement 4 By design no extra hardware is required, rather it encourages to recycle the resources available and enables lower-end and legacy resources to participate. More specifically, we reduced the overhead incurred by traditional (full) virtualization technologies by using light virtualization technologies, such as containers.

Requirement 5 Dynamic membership is designed to be supported by this application, using the autonomous PID controllers, but still requires to be implemented. We provide scalability to the applications hosted using this architecture, through these PID controllers and because this system is built around asynchronous task queues.

We consider to have addressed the *evaluation requirements*, if not exhaustively, at least sufficiently to ensure the proper functioning of this system as a *collaborative system*. We demonstrate this claim as follows:

- We provide a way to represent resources based on a set of desirable attributes, and means to **Advertise** and **Discover** these resources using the *Ring* abstraction.
- We use a two-step best-effort mechanism to **Select** the most relevant resources based on a consumer-defined query using a pub/sub messaging pattern. Then, in the second step we provide means to evaluate the inter-resource relationships of the tentative selection, effectively addressing the concerns expressed in the **Match** phase.
- Using acknowledgments in the *Select* mechanism, we provide this system with the ability to **Bind** the resources to a tentative consumer, thereby mitigating any potential concurrent bindings of the same resources to two different applications.
- Resources are **Used** to execute tasks immediately after the initialization process is completed.
- **Release** of the resources, is performed autonomously using Proportional-Integral-Derivative controllers that continuously monitor the workload and performance of each type of nodes in an application, and they release any superfluous resources.

8.2 Contributions

In this thesis we introduced four contributions which are novel to the best of our knowledge, see Section 1.3.

Fully-Decentralized Collaborative Web Computing Platform

Compared to the Cloud@Home system, which provides a full-fledged volunteer cloud computing infrastructure and a business-model to entice contributions, we proposed to provide a fully-decentralized web computing platform, focusing on the PaaS service model. Instead of devising a business model for a completely untrusted environment, we resort to a whitelisting mechanism operating in a *semi-trusted* environment.

The Peer-to-Peer Cloud System focuses on a decentralized IaaS volunteer cloud computing infrastructure, whereas our architecture focuses on a decentralized PaaS volunteer cloud computing infrastructure operating in a semi-trusted environment. Our API contains more features than the API they offered, but doesn't aim to provide the same service model.

Our platform is decentralized to the extent that is possible within the current infrastructure of the Internet, because it is not possible to associate a URL to a non-static IP address, or simply impractical using the current Internet infrastructure. Consequently, every Web-based application built for this system will have at most one single point of failure, the web server. Thus, unless the application is not providing a point of access through the web and resorts to a dedicated application communicating directly, this problem remains open.

The computing platform is flexible enough to cater to a variety of applications, akin to what current cloud service providers offer in terms of PaaS. We have demonstrated this by implementing two proofs of concept, and have showed how intuitive it is to refactor applications for this architecture using the event-driven programming model.

Candidate Minimal API Specification for Computing Platforms and/or PaaS

The API devised, see Appendix A, is based on an investigation of the three major cloud service providers: *Google*, *Amazon* and *Microsoft*. It circumscribes the essential services and features provided in a computing platform and it is minimal, because it provides the basic building blocks for any web applications. Given that a *specific* feature is missing, it is easily extensible and can be incorporated into the API, as long as it is included as an application-specific dependency in the container configuration file. Thus, this system provides a computing platform that is minimal and does not impose any superfluous features, because it is intended to be used with lower-end computing resources.

Using Light Virtualization to Abstract and Isolate Contributors

To the best of our knowledge, we are the first to propose the use of *light virtualization*, or containers, as an atomic unit composing the computing platform and also in the context of PaaS. Since the start of this work, the technological landscape shifted and now containers are gaining in popularity, especially with the support of major partners like Microsoft.

Minimally Intrusive System Inducing Small Memory Footprint

Our system is minimally intrusive and intuitive, as shown in the case study of *Multi-Document Text Summarization using a Genetic Algorithm*. By resorting to *light virtualization* rather than full virtualization, we have reduced the overhead induced by the virtualization technologies to a minimum, while maintaining the desirable isolation and abstractive properties of virtualization. However this observation is largely informal¹, and has yet to be precisely quantified, but from our early experimentation we have found that the memory footprint of the container itself was around 8MB, which supports our intuition.

8.3 Future Work

Several research issues were identified while investigating the research problems of this thesis. We have identified five short-term research issues that pertain to the implementation of this system, and one long-term research issue pertaining to the environment of the system. We have presented some open problems that we have encountered while writing this thesis, see Section 7.2, but these problems are not specific to our system.

Short-term Research Issues

1. Implement the *T-Man* overlay network gossip-based protocol, [JMB09], to construct the overlay network topologies in the *Fellowships*, rather than connecting the nodes directly to each other.
2. Provide dynamic configuration support for nodes, enabling them to modify their configuration at run-time rather than through the modification of configuration files.
3. Implement PID controllers to provide dynamic membership capabilities, and test how well it operates under various workloads. But also research what kind of

¹See Section 2.4.2.

parameters are suitable for these autonomic controllers, given that in some cases it may not be trivial.

4. Investigate more elaborate access control schemes to evaluate how much change is required to adapt this system to an *enterprise* environment, for it to operate as an Enterprise Service Bus.
5. Perform exhaustive performance testing and accurate memory profiling of the architecture, in order to identify possible memory leaks and to augment the performance of the overall system. Additionally, performing these tests in a distributed environment rather than locally would provide a realistic performance assessment.

Long-term Research Issues

The primary long-term research issue that rises out of this academic inquiry, is how can we adapt (if possible) the current system to operate in a fully untrusted environment? Researching this issue will enlighten us on how to provide a truly public volunteer cloud computing infrastructure using a *decentralized* networking primitive. Multiple concerns arises in a fully untrusted environment, security being the most obvious one but more precisely how to prevent this system in such an environment to be used for malicious intents, such as raising a bot-net army? Thus security concerns are supplemented with ethical concerns in an untrusted environment, and is it possible to mitigate them using this design?

Appendix A

API Specification

Functions/ Parameters	Description
Database/Storage	
<code>create(authKey)</code>	This function is used to create a database instance using a specific authentication key. Returns a new DataProcess instance.
<code>stop()</code>	This function is used to stop the database server. Returns a boolean indicating the outcome of the function.
<code>start()</code>	This function is used to start the database server. Returns a boolean indicating the outcome of the function.
<code>run(query)</code>	This function is used to run a query on the database, sending it to the server. Returns the result of the query.
Network/Communication	
<code>connect(node,list-of-ips)</code>	This function is used to connect the <i>node</i> , bootstrapping it to the network using the <i>list-of-ips</i> , which contains a list of contactable nodes.
Continued on next page	

Continued from previous page

Functions/ Parameters	Description
	Returns a deferred when this function completes.
<code>set(key, value)</code>	This function is used to set a value for a corresponding key in the DHT. Returns a deferred when this function completes.
<code>get(key)</code>	This function returns the valued corresponding to the key passed in parameters, if it exists in the DHT. Returns a deferred with the result.
Load Balancing/Scalability	
workloadWN	Configuration file parameter for the percentage of the desired utilization of the <i>Worker Nodes</i>
workloadDN	Configuration file parameter for the percentage of the desired utilization of the <i>Data Nodes</i>
coeffP	Configuration file parameter for the coefficient of the P component of the PID controller.
coeffI	Configuration file parameter for the coefficient of the I component of the PID controller.
Continued on next page	

Continued from previous page

Functions/ Parameters	Description
coeffD	Configuration file parameter for the coefficient of the D component of the PID controller.
Security	
checkCredentials(<i>user</i> , <i>hashPW</i>)	This function takes a username and a hashed password (SHA-2), and verifies them against the credentials stored in the database. Returns a boolean indicating whether the credentials are verified or not.
Application Deployment/Management	
resourceSpecification	Configuration file parameters for the resource specification to be stored in the <i>template</i> . Example: cpuUsage, memUsage, cpuCapacity, memCapacity, hddCapacity, ...
Concluded	

Appendix B

Code Samples

Caveat lector: this implementation is still in pre-alpha stage, and some details were left out in order to achieve a working prototype and to produce two proofs of concepts. It includes, communicating directly between the two processes of the web server and the Web Protocol, rather than using HTTP POST requests, because these processes were residing on the same host.

B.1 Configuration Files

B.1.1 Architecture Configuration File

```
# resource specification: cpuUsage, memUsage, cpuCapacity, memCapacity, hddCapacity, ..
cpuUsage

# number of db_nodes
db_nodes 1

# NOT IMPLEMENTED – for PID Controllers
workloadWN 90
workloadDN 75
coeffP 75
coeffI 45
coeffD 55
```

B.1.2 Docker Configuration File

```
FROM ubuntu:14.04
MAINTAINER Dany Wilson <dwils098@gmail.com>

RUN apt-get update
RUN apt-get install -qy python-zmq
RUN apt-get install -qy python-pip
RUN apt-get install -qy python-dev
RUN apt-get install -qy python-gevent
RUN apt-get install -qy python-twisted
RUN pip install cherrypy

RUN apt-get update
RUN apt-get install -qy rethinkdb

# Rethinkdb process
EXPOSE 28015
# Rethinkdb admin console
EXPOSE 8080

# Create the /rethinkdb_data dir structure
RUN /usr/bin/rethinkdb create

ENTRYPOINT ["/usr/bin/rethinkdb"]

ADD mtds /scripts/mtds

# params: user_IP, DHT_port, appD, known_IP
RUN /scripts/mtds/mtds_core.py 192.168.56.1 5555 True 192.168.56.101
```

B.2 Calculator

B.2.1 Web Server

```
import cherrypy
class Calculator(object):

    @cherrypy.expose
    def index(self):
        return file('example/index.html')

    @cherrypy.expose
    def add(self, op_1, op_2, **params):
        print 'add',op_1, op_2
        import sys
        result = sys.stdin.readline()

        return result

    @cherrypy.expose
    def subtract(self, op_1, op_2, **params):
        print 'sub',op_1, op_2
        import sys
        result = sys.stdin.readline()

        return result

    @cherrypy.expose
    def multiply(self, op_1, op_2, **params):
        print 'mul',op_1, op_2
        import sys
        result = sys.stdin.readline()

        return result

    @cherrypy.expose
    def divide(self, op_1, op_2,**params):
        # send the operands to through the pipe (to the webProtocol)
        print 'div',op_1, op_2
```

```
import sys

# wait for the response...
result = sys.stdin.readline()

out = "The result is: " + result
return out

# start the web_server first
cherrypy.log.screen = None
# bind to all IPv4 interfaces
cherrypy.config.update({'server.socket_host': '0.0.0.0'})
cherrypy.tree.mount(Calculator())
cherrypy.engine.start()
cherrypy.engine.block()
```

B.2.2 Web Protocol

```
from twisted.internet import stdio
from twisted.protocols import basic

#----- Task Functionalities -----
def addition_function(task_obj, params):
    result = params[0] + params[1]

    task_obj.results = result
    return result

def subtract_function(task_obj, params):
    result = params[0] - params[1]

    task_obj.results = result
    return result

def multiply_function(task_obj, params):
    result = params[0] * params[1]

    task_obj.results = result
    return result

def division_function(task_obj, params):
    result = params[0] / params[1]

    task_obj.results = result
    return result

#----- End of Task Functionalities -----

# This class takes requests in from the webServer + process them into tasks
class WebProtocol(basic.LineReceiver):
    from os import linesep as delimiter

    def __init__(self, app_obj):
        self.app = app_obj

    def connectionMade(self):
```

```
self.app._netHandle._log.info('webProtocol --> connectionMade = SUCCESS!')
```

```
def lineReceived(self, line):
```

```
    self.app._netHandle._log.debug('WP --> '+line)
```

```
    result = ""
```

```
    operator,op_1, op_2 = line.split()
```

```
    #retrieve the current file name and strip extension
```

```
    import os
```

```
    module=os.path.basename(__file__).split('.', 1)[0]
```

```
    self.app._netHandle._log.debug(module)
```

```
    from node_netHandle import Task
```

```
    new_task = Task("WEB")
```

```
    new_task.type = Task.Worker
```

```
    if operator == 'add':
```

```
        new_task.create(module, 'addition_function', [int(op_1), int(op_2)])
```

```
    if operator == 'sub':
```

```
        new_task.create(module, 'subtract_function', [int(op_1), int(op_2)])
```

```
    if operator == 'mul':
```

```
        new_task.create(module, 'multiply_function', [int(op_1), int(op_2)])
```

```
    if operator == 'div':
```

```
        new_task.create(module, 'division_function', [int(op_1), int(op_2)])
```

```
    self.app.addTask(new_task)
```

```
    return result
```

```
        ' ' + file_name + ' ' + str(size))
    sys.stdout.write('\n')
    sys.stdout.write(output)
    count+=1

    sys.stdout.write(OP_RETRIEVE + ' ' +str(cherrypy.session.id) + ' ' + str(count))
    sys.stdout.write('\n')

    return 'Number of FILES: ' + str(count)

@cherrypy.expose
def upload(self,**kwargs):
    # data_file dict
    data_files={}
    data_bytes=""

    list_of_files = []
    for myFile in kwargs['myFiles']:
        return_str=""
        data=""
        while True:
            data_chunk = myFile.file.read(8192)
            if not data_chunk:
                break
            size += len(data)
            data+=data_chunk

        # save the file as bytes
        saved = open(myFile.filename, 'wb')
        saved.write(data)
        saved.close()

        # load bytes
        loaded = open(myFile.filename, 'rb')
        data_bytes = loaded.read()
```

```

        loaded.close()

        # to the db
        data_files[str(myFile.filename)] = data_bytes
        with open('post_'+str(myFile.filename), 'wb') as fh:
            fh.write(data_files[str(myFile.filename)])

        # call task creation function
        count = self.create_tasks(kwargs, data_files)

        # load bytes
        import sys
        result = sys.stdin.readline()

        return result

# start the web_server first
cherry.py.log.screen = None
# bind to all IPv4 interfaces
cherry.py.config.update({
    'tools.sessions.on' : True,
    'server.socket_host': '0.0.0.0'
})
cherry.py.tree.mount(MTDS())
cherry.py.engine.start()
cherry.py.engine.block()

```

B.3.2 Web Protocol

```

from twisted.internet import stdio
from twisted.protocols import basic

#----- Task Functionalities -----
def addition_function(task_obj, params):
    result = params[0] + params[1]

```

```

    task_obj.results = result
    return result

def subtract_function(task_obj, params):
    result = params[0] - params[1]

    task_obj.results = result
    return result

def multiply_function(task_obj, params):
    result = params[0] * params[1]

    task_obj.results = result
    return result

def division_function(task_obj, params):
    result = params[0] / params[1]

    task_obj.results = result
    return result
#----- End of Task Functionalities -----

# This class takes requests in from the webServer + process them into tasks
class WebProtocol(basic.LineReceiver):
    delimiter='\n'

    def __init__(self, app_obj):
        self.app = app_obj
        self.buffer = ""

    def connectionMade(self):
        self.app._netHandle._log.info('webProtocol --> connectionMade = SUCCESS!')

    def rawDataReceived(self, data):
        self.app._netHandle._log.debug(str(self.exp_len))

        self.buffer+=data

```

```

if len(self.buffer) >= self.exp_len:
    buf = self.buffer[:self.exp_len]
    rem = self.buffer[self.exp_len:]

    operator, filename, data_o = str(buf).split('-',2)

    from node_netHandle import Task

    new_task = Task("DB_STORE")
    new_task.type = Task.Data
    new_task._job_data = data_o
    new_task.create('data_process',
                    'save_file',
                    [filename, 'data\binary',
                     self.db_name,
                     ['worker_process','retrieveData']])

    # create a task to save the file, and as parameters
    # include any consequent actions. In this case we want
    # to generate a job for retrieving the data:
    #           from module = worker_process.py
    #           function    = retrieveData

    self.app.addTask(new_task)
    self.buffer = ""
    self.setLineMode(rem)
    return

def lineReceived(self, line):

    result = ""
    if 'store' in line:
        operator, session_id, filename,size = line.split()
        self.db_name = session_id
        self.exp_len = int(size)
        # change mode to receive file
        self.setRawMode()
    if 'ret' in line:

```

```
operator, session_id, num = line.split()
self.db_name = session_id
from node_netHandle import Task
new_task = Task("DB_RETRIEVE")
new_task.type = Task.Data
new_task._job_data = num

new_task.create('data_process',
               'result_lookup',
               [self.db_name, num,
               ['worker_process', 'retrieveData_postprocessing']])

self.app.addTask(new_task)
return result
```

Bibliography

- [AAC⁺11] Rocco Aversa, Marco Avvenuti, Antonio Cuomo, Beniamino Di Martino, Giuseppe Di Modica, Salvatore Distefano, Antonio Puliafito, Massimiliano Rak, Orazio Tomarchio, Alessio Vecchio, et al. The cloud@ home project: towards a new enhanced computing paradigm. In *Euro-Par 2010 Parallel Processing Workshops*, pages 555–562. Springer, 2011.
- [AB14] Reaz Ahmed and Raouf Boutaba. *Collaborative Web Hosting: Challenges and Research Directions*. Springer, 2014.
- [ABC⁺11] Rocco Aversa, Dario Bruneo, Antonio Cuomo, Beniamino Di Martino, Salvatore Distefano, Antonio Puliafito, Massimiliano Rak, Salvatore Venticinquè, and Umberto Villano. Cloud@ home: Performance management components. In *Euro-Par 2010 Parallel Processing Workshops*, pages 579–586. Springer, 2011.
- [ACK⁺02] David P Anderson, Jeff Cobb, Eric Korpela, Matt Lebofsky, and Dan Werthimer. Seti@ home: an experiment in public-resource computing. *Communications of the ACM*, 45(11):56–61, 2002.
- [AF06] David P Anderson and Gilles Fedak. The computational and storage potential of volunteer computing. In *Cluster Computing and the Grid, 2006. CCGRID 06. Sixth IEEE International Symposium on*, volume 1, pages 73–80. IEEE, 2006.
- [And03] David P Anderson. Public computing: Reconnecting people to science. In *Conference on Shared Knowledge and the Web*, pages 17–19, 2003.
- [And04] David P Anderson. Boinc: A system for public-resource computing and storage. In *Grid Computing, 2004. Proceedings. Fifth IEEE/ACM International Workshop on*, pages 4–10. IEEE, 2004.

- [APHB14] Maeva Antoine, Laurent Pellegrino, Fabrice Huet, and Françoise Baude. A generic api for load balancing in structured p2p systems. In *Computer Architecture and High Performance Computing Workshop (SBAC-PADW), 2014 International Symposium on*, pages 138–143. IEEE, 2014.
- [ASK15] Mohammadfazel Anjomshoa, Mazleena Salleh, and Maryam Pouryazdanpanah Kermani. A taxonomy and survey of distributed computing systems. *Journal of Applied Sciences*, 15(1), 2015.
- [Bar92] Ramon C Barquin. In pursuit of a ‘ten commandments’ for computer ethics. *Computer Ethics Institute*, 1992.
- [Bar01] David Barkai. *Peer-to-Peer Computing: technologies for sharing and collaborating on the net*. Intel Press, 2001.
- [BDF⁺03] Paul Barham, Boris Dragovic, Keir Fraser, Steven Hand, Tim Harris, Alex Ho, Rolf Neugebauer, Ian Pratt, and Andrew Warfield. Xen and the art of virtualization. *ACM SIGOPS Operating Systems Review*, 37(5):164–177, 2003.
- [BFH03] Fran Berman, Geoffrey Fox, and Anthony JG Hey. *Grid computing: making the global infrastructure a reality*, volume 2. John Wiley and sons, 2003.
- [BG09] Andrew Berns and Sukumar Ghosh. Dissecting self-* properties. In *Self-Adaptive and Self-Organizing Systems, 2009. SASO’09. Third IEEE International Conference on*, pages 10–19. IEEE, 2009.
- [Bir07] Ken Birman. The promise, and limitations, of gossip protocols. *ACM SIGOPS Operating Systems Review*, 41(5):8–13, 2007.
- [BJ12] HMN Dilum Bandara and Anura P Jayasumana. Evaluation of p2p resource discovery architectures using real-life multi-attribute resource and query characteristics. In *Consumer Communications and Networking Conference (CCNC), 2012 IEEE*, pages 634–639. IEEE, 2012.
- [BJ13] HMN Dilum Bandara and Anura P Jayasumana. Collaborative applications over peer-to-peer systems—challenges and solutions. *Peer-to-Peer Networking and Applications*, 6(3):257–276, 2013.

- [BMT12] Ozalp Babaoglu, Moreno Marzolla, and Michele Tamburini. Design and implementation of a p2p cloud system. In *Proceedings of the 27th Annual ACM Symposium on Applied Computing*, pages 412–417. ACM, 2012.
- [CDP10] Vincenzo D Cunsolo, Salvatore Distefano, and Antonio Puliafito. Cloud@ home on top of reservoir. In *Cloud Computing*, pages 41–56. Springer, 2010.
- [CDPS09a] Vincenzo D Cunsolo, Salvatore Distefano, Antonio Puliafito, and Marco Scarpa. Cloud@ home: Bridging the gap between volunteer and cloud computing. In *Emerging Intelligent Computing Technology and Applications*, pages 423–432. Springer, 2009.
- [CDPS09b] Vincenzo D Cunsolo, Salvatore Distefano, Antonio Puliafito, and Marco Scarpa. Volunteer computing and desktop cloud: The cloud@ home paradigm. In *Network Computing and Applications, 2009. NCA 2009. Eighth IEEE International Symposium on*, pages 134–139. IEEE, 2009.
- [CDPS10a] Vincenzo D Cunsolo, Salvatore Distefano, Antonio Puliafito, and Marco Scarpa. Applying software engineering principles for designing cloud@ home. In *Cluster, Cloud and Grid Computing (CCGrid), 2010 10th IEEE/ACM International Conference on*, pages 618–624. IEEE, 2010.
- [CDPS10b] Vincenzo D Cunsolo, Salvatore Distefano, Antonio Puliafito, and Marco Scarpa. From volunteer to cloud computing: cloud@ home. In *Proceedings of the 7th ACM international conference on Computing frontiers*, pages 103–104. ACM, 2010.
- [CDPS10c] Vincenzo D Cunsolo, Salvatore Distefano, Antonio Puliafito, and Marco Scarpa. Open and interoperable clouds: The cloud@ home way. In *Cloud Computing*, pages 93–111. Springer, 2010.
- [CDPS10d] Vincenzo D Cunsolo, Salvatore Distefano, Antonio Puliafito, and Marco Scarpa. Open and interoperable clouds: The cloud@ home way. In *Cloud Computing*, pages 93–111. Springer, 2010.
- [Chi08] David Chisnall. *The definitive guide to the xen hypervisor*. Pearson Education, 2008.
- [Cit12] Citrix. Most american confused by cloud computing according to national survey, 2012. Taken from

- <https://www.citrix.com/news/announcements/aug-2012/most-americans-confused-by-cloud-computing-according-to-national.html>, Consulted: June 4th 2015.
- [CS99] Peter H Carstensen and Kjeld Schmidt. Computer supported cooperative work: New challenges to systems design. In *In K. Itoh (Ed.), Handbook of Human Factors*. Citeseer, 1999.
- [CS14] O’Neil Software Christine Spisto. Cloud computing confusion – the root cause?, 2014. Taken from <https://www.oneilsoft.com/news-portal/blogs/780-cloud-computing-confusion-the-root-cause>, Consulted: June 4th 2015.
- [CW09] Abhishek Chandra and Jon Weissman. Nebulas: Using distributed voluntary resources to build clouds. In *Proceedings of the 2009 conference on Hot topics in cloud computing*, pages 2–2. USENIX Association, 2009.
- [Dar14] Sashank Dara. Multi-user protocols with access control for computational privacy in public clouds. *arXiv preprint arXiv:1406.1823*, 2014.
- [DCP10] Salvatore Distefano, Vincenzo D Cunsolo, and Antonio Puliafito. A taxonomic specification of cloud@ home. In *Advanced Intelligent Computing Theories and Applications. With Aspects of Artificial Intelligence*, pages 527–534. Springer, 2010.
- [Dev15] Google Developers. Google identity platform. <https://developers.google.com/identity/>, 2015. Accessed: 12-05-2015.
- [DFP11] Salvatore Distefano, Maria Fazio, and Antonio Puliafito. The cloud@ home resource management system. In *Utility and Cloud Computing (UCC), 2011 Fourth IEEE International Conference on*, pages 122–129. IEEE, 2011.
- [DKMyC04] Alan Dearle, Graham Kirby, Andrew McCarthy, and Juan Carlos Diaz y Carballo. A flexible and secure deployment framework for distributed applications. In *Component Deployment*, pages 219–233. Springer, 2004.
- [DIRB14] Jose De la Rosa and Kent Baxley. Lxc containers in ubuntu server 14.04 lts. <http://en.community.dell.com/techcenter/os-applications/w/wiki/6950.lxc-containers-in-ubuntu-server-14-04-lts>, June 2014. Accessed: 30-12-2014.

- [Dol00] Shlomi Dolev. *Self-stabilization*. MIT press, 2000.
- [DP80] IW Draffan and F Poole. *Distributed data bases*. CUP Archive, 1980.
- [DP12] Salvatore Distefano and Antonio Puliafito. Cloud@ home: Toward a volunteer cloud. *IT Professional*, 14(1):27–31, 2012.
- [DPR⁺11] Salvatore Distefano, Antonio Puliafito, Massimiliano Rak, Salvatore Venticinque, Umberto Villano, Antonio Cuomo, Giuseppe Di Modica, and Orazio Tomarchio. Qos management in cloud@ home infrastructures. In *Cyber-Enabled Distributed Computing and Knowledge Discovery (CyberC), 2011 International Conference on*, pages 190–197. IEEE, 2011.
- [EK07] Pavel Emelyanov and Kir Kolyshkin. Pid namespaces in the 2.6. 24 kernel. *LWN. net*, November, 2007.
- [ES08] Agoston E Eiben and James E Smith. *Introduction to evolutionary computing*. Springer Science & Business Media, 2008.
- [Fit11] Brad Fitzpatrick. Memcached: a distributed memory object caching system. *Memcached-a Distributed Memory Object Caching System*, 2011.
- [GGL03] Sanjay Ghemawat, Howard Gobioff, and Shun-Tak Leung. The google file system. In *ACM SIGOPS Operating Systems Review*, volume 37, pages 29–43. ACM, 2003.
- [GMCK00] Jade Goldstein, Vibhu Mittal, Jaime Carbonell, and Mark Kantrowitz. Multi-document summarization by sentence extraction. In *Proceedings of the 2000 NAACL-ANLP Workshop on Automatic summarization-Volume 4*, pages 40–48. Association for Computational Linguistics, 2000.
- [Gol10] Dieter Gollmann. Computer security. *Wiley Interdisciplinary Reviews: Computational Statistics*, 2(5):544–554, 2010.
- [GSL14] Ian Gergin, Bradley Simmons, and Marin Litoiu. A decentralized autonomic architecture for performance control in the cloud. In *Cloud Engineering (IC2E), 2014 IEEE International Conference on*, pages 574–579. IEEE, 2014.
- [GSRU07] Debanjan Ghosh, Raj Sharman, H Raghav Rao, and Shambhu Upadhyaya. Self-healing systems—survey and synthesis. *Decision Support Systems*, 42(4):2164–2185, 2007.

- [Hyk14] Solomon Hykes. Dockercon. <http://www.dockercon.com/>, 2014. Accessed: 30-12-2014.
- [Hyk15] Solomon Hykes. Docker. <http://docs.docker.com/>, 2015. Accessed: 16-02-2015.
- [inc14a] Free Software Foundation inc. Free software foundation – working together for free software. <http://www.fsf.org>, 2014. Accessed: 30-12-2014.
- [Inc14b] Google Inc. Google app engine documentation. <https://cloud.google.com/appengine/docs>, 2014. Accessed: 08-01-2015.
- [Inc14c] Microsoft Inc. What is microsoft azure? <https://azure.microsoft.com/en-us/overview/what-is-azure/>, 2014. Accessed: 30-12-2014.
- [Inc15a] Amazon Web Services Inc. Amazon web services. <http://aws.amazon.com/>, 2015. Accessed: 08-01-2015.
- [Inc15b] Docker Inc. Docker security, v1.6, 2015. Taken from <http://docs.docker.com/articles/security/>, Consulted: June 4th 2015.
- [J⁺99] K Sparck Jones et al. Automatic summarizing: factors and directions. *Advances in automatic text summarization*, pages 1–12, 1999.
- [JC10] Xing Jin and S-H Gary Chan. Unstructured peer-to-peer network architectures. In *Handbook of Peer-to-Peer Networking*, pages 117–142. Springer, 2010.
- [Jel] Márk Jelasity. Gossip-based protocols for large-scale distributed systems.
- [JH12] M Jones and Dick Hardt. The oauth 2.0 authorization framework: Bearer token usage. Technical report, RFC 6750, October, 2012.
- [JK06] Mark Jelasity and A-M Kermarrec. Ordered slicing of very large-scale overlay networks. In *Peer-to-Peer Computing, 2006. P2P 2006. Sixth IEEE International Conference on*, pages 117–124. IEEE, 2006.
- [JM14] Forbes Joe Meck Kendrick. What cloud computing customers want: Clarity, simplicity, support, 2014. Taken from <http://www.forbes.com/sites/joemckendrick/2014/07/19/what-cloud-computing-customers-want-clarity-simplicity-support/>, Consulted: June 4th 2015.

- [JMB05] Márk Jelasity, Alberto Montresor, and Ozalp Babaoglu. Gossip-based aggregation in large dynamic networks. *ACM Transactions on Computer Systems (TOCS)*, 23(3):219–252, 2005.
- [JMB09] Márk Jelasity, Alberto Montresor, and Ozalp Babaoglu. T-man: Gossip-based fast overlay topology construction. *Computer networks*, 53(13):2321–2339, 2009.
- [JOV05] Hosagrahar V Jagadish, Beng Chin Ooi, and Quang Hieu Vu. Baton: A balanced tree structure for peer-to-peer networks. In *Proceedings of the 31st international conference on Very large data bases*, pages 661–672. VLDB Endowment, 2005.
- [JVG⁺07] Márk Jelasity, Spyros Voulgaris, Rachid Guerraoui, Anne-Marie Kermarrec, and Maarten Van Steen. Gossip-based peer sampling. *ACM Transactions on Computer Systems (TOCS)*, 25(3):8, 2007.
- [KC03] Jeffrey O Kephart and David M Chess. The vision of autonomic computing. *Computer*, 36(1):41–50, 2003.
- [KC12] Nidhi Jain Kansal and Inderveer Chana. Existing load balancing techniques in cloud computing: a systematic review. *Journal of Information Systems and Communication*, 3(1):87–91, 2012.
- [KL03] Alexander Keller and Heiko Ludwig. The wsla framework: Specifying and monitoring service level agreements for web services. *Journal of Network and Systems Management*, 11(1):57–81, 2003.
- [LCP⁺05] Eng Keong Lua, Jon Crowcroft, Marcelo Pias, Ravi Sharma, Steven Lim, et al. A survey and comparison of peer-to-peer overlay network schemes. *IEEE Communications Surveys and Tutorials*, 7(1-4):72–93, 2005.
- [Lef14] Original Author: Glyph Lefkowitz. Twisted matrix labs. <https://www.twistedmatrix.com/trac>, 2014. Accessed: 30-12-2014.
- [Lew11] Brian Lewis. Separation of responsibilities., 2011. Taken from <http://mythoughtsonit.com/2011/04/infrastructure-as-a-service-platform-as-a-service-software-as-a-servicetake-a-look-at-the-management-stack>, Consulted: November 19th 2014.

- [Lin76] James G Linders. Distributed data bases. *Computers & Geosciences*, 2(3):293–297, 1976.
- [LLC11] Zhen Luo, Zhishu Li, and Biao Cai. A self-organized public-key certificate system in p2p network. *Journal of Networks*, 6(10):1437–1443, 2011.
- [Los15] Canonical Ltd. and open source. Linux containers. <https://linuxcontainers.org/>, 2015. Accessed: 12-05-2015.
- [Men11] Paul Menage. Linux kernel documentation/cgroups/cgroups.txt, 2011.
- [MF13] Jessica McKellar and Abe Fettig. Twisted network programming essentials. 2013.
- [MG11] Peter Mell and Tim Grance. The nist definition of cloud computing. x, 2011.
- [MGLPPJ13] Rubén Mondéjar, Pedro García-López, Carles Pairet, and Lluís Pamies-Juarez. Cloudsnap: A transparent infrastructure for decentralized web deployment using distributed interception. *Future Generation Computer Systems*, 29(1):370–380, 2013.
- [Mic06] Brenda M Michelson. Event-driven architecture overview. *Patricia Seybold Group*, 2, 2006.
- [MKL⁺02] Dejan S Milojevic, Vana Kalogeraki, Rajan Lukose, Kiran Nagaraja, Jim Pruyne, Bruno Richard, Sami Rollins, and Zhichen Xu. Peer-to-peer computing, 2002.
- [MM99] Inderjeet Mani and Mark T Maybury. *Advances in automatic text summarization*, volume 293. MIT Press, 1999.
- [MM02] Petar Maymounkov and David Mazieres. Kademia: A peer-to-peer information system based on the xor metric. In *Peer-to-Peer Systems*, pages 53–65. Springer, 2002.
- [Mul14] Brian Muller. Kademia: A dht in python. <http://findingscience.com/python/kademia/dht/2014/02/14/kademia:-a-dht-in-python.html>, 2014. Accessed: 30-12-2014.

- [NO95] Matunda Nyanchama and Sylvia L Osborn. Modeling mandatory access control in role-based security systems. In *DBSec*, pages 129–144. Citeseer, 1995.
- [NWG89] Internet Activities Board Network Working Group. Rfc 1087: Ethics of the internet, 1989.
- [Olk06] Timo Olkkonen. Generic authentication architecture. In *Security and Privacy in Pervasive Computing, Seminar on Network Security, Espoo*, 2006.
- [Ope15] Opensource. Rethinkdb: The open-source database for the realtime web. <http://rethinkdb.com/>, 2015. Accessed: 12-05-2015.
- [ÖV11] M Tamer Özsu and Patrick Valduriez. *Principles of distributed database systems*. Springer, 2011.
- [PD11] Harald Psailer and Schahram Dustdar. A survey on self-healing systems: approaches and systems. *Computing*, 91(1):43–73, 2011.
- [PG74] Gerald J Popok and Robert P Goldberg. Formal requirements for virtualizable third generation architectures. *Communications of the ACM*, 17(7):412–421, 1974.
- [PMRS14] Mathieu Pasquet, Francisco Maia, Etienne Rivière, and Valerio Schiavoni. Autonomous multi-dimensional slicing for large-scale distributed systems. In *Distributed Applications and Interoperable Systems*, pages 141–155. Springer, 2014.
- [QHH08] Vahed Qazvinian, Leila Sharif Hassanabadi, and Ramin Halavati. Summarising text with a genetic algorithm-based sentence extraction. *International Journal of Knowledge Management Studies*, 2(4):426–444, 2008.
- [RCL09] Bhaskar Prasad Rimal, Eunmi Choi, and Ian Lumb. A taxonomy and survey of cloud computing systems. In *INC, IMS and IDC, 2009. NCM'09. Fifth International Joint Conference on*, pages 44–51. Ieee, 2009.
- [RR08] Leonard Richardson and Sam Ruby. *RESTful web services*. " O'Reilly Media, Inc.", 2008.

- [RV11] Etienne Riviere and Spyros Voulgaris. Gossip-based networking for internet-scale distributed systems. In *E-Technologies: Transformation in a Connected World*, pages 253–284. Springer, 2011.
- [sal15] inc. salesforce.com. Crm and cloud computing to grow your business: Salesforce.com canada. <http://www.salesforce.com/ca/>, 2015. Accessed: 16-02-2015.
- [Sar10] Siamak Sarmady. A survey on peer-to-peer and dht. *arXiv preprint arXiv:1006.4708*, 2010.
- [SMJ00] Rick Sturm, Wayne Morris, and Mary Jander. *Foundations of service level management*, volume 13. Sams Indianapolis, IN, 2000.
- [SMK⁺01] Ion Stoica, Robert Morris, David Karger, M Frans Kaashoek, and Hari Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. *ACM SIGCOMM Computer Communication Review*, 31(4):149–160, 2001.
- [SPEW11] Amir Ali Semnanian, Jeffrey Pham, Burkhard Englert, and Xiaolong Wu. Virtualization technology and its impact on computer hardware architecture. In *Information Technology: New Generations (ITNG), 2011 Eighth International Conference on*, pages 719–724. IEEE, 2011.
- [Spi10] Richard Spinello. *Cyberethics: Morality and law in cyberspace*. Jones & Bartlett Learning, 2010.
- [Tav10] Herman T Tavani. *Ethics and technology: Controversies, questions, and strategies for ethical computing*. John Wiley & Sons, 2010.
- [Tav12] Djamshid Tavangarian. Virtual computing: The emperor’s new clothes? In *Software Service and Application Engineering*, pages 53–70. Springer, 2012.
- [Tea14] CherryPy Team. Cherrypy: A minimalist python web framework. <http://www.cherrypy.org>, 2014. Accessed: 30-12-2014.
- [UPS11] Guido Urdaneta, Guillaume Pierre, and Maarten Van Steen. A survey of dht security techniques. *ACM Computing Surveys (CSUR)*, 43(2):8, 2011.
- [Wat08] Jon Watson. Virtualbox: bits and bytes masquerading as machines. *Linux Journal*, 2008(166):1, 2008.

- [Wik15] Wikipedia. Seti@home wikipedia article. <http://en.wikipedia.org/wiki/SETI@home>, 2015. Accessed: 16-02-2015.
- [WR15] Cloud Tech News William Rabis. The confusion around hybrid: Cloud isn't for everyone, but everyone can use a bit of cloud, 2015. Taken from <http://www.cloudcomputing-news.net/news/2015/jun/01/confusion-around-hybrid-cloud-isnt-everyone-everyone-can-use-bit-cloud/>, Consulted: June 4th 2015.
- [YBDS08] Lamia Youseff, Maria Butrico, and Dilma Da Silva. Toward a unified ontology of cloud computing. In *Grid Computing Environments Workshop, 2008. GCE'08*, pages 1–10. IEEE, 2008.
- [ZCB10] Qi Zhang, Lu Cheng, and Raouf Boutaba. Cloud computing: state-of-the-art and research challenges. *Journal of internet services and applications*, 1(1):7–18, 2010.