

Software Architecture Definition for On-demand Cloud Provisioning*

Clovis Chapman
Department of Computer
Science, UCL
Gower Street
London, UK
c.chapman@cs.ucl.ac.uk

Wolfgang Emmerich
Department of Computer
Science, UCL
Gower Street
London, UK
we@acm.org

Fermín Galán Márquez
Telefónica I+D
Emilio Vargas 6
Madrid, Spain
fermin@tid.es

Stuart Clayman
Department of Electrical
Engineering, UCL
Gower Street
London, UK
s.clayman@ee.ucl.ac.uk

Alex Galis
Department of Electrical
Engineering, UCL
Gower Street
London, UK
a.galis@ee.ucl.ac.uk

ABSTRACT

Cloud computing [22] is a promising paradigm for the provisioning of IT services. Cloud computing infrastructures, such as those offered by the RESERVOIR project, aim to facilitate the deployment, management and execution of services across multiple physical locations in a seamless manner. In order for service providers to meet their quality of service objectives, it is important to examine how software architectures can be described to take full advantage of the capabilities introduced by such platforms. When dealing with software systems involving numerous loosely coupled components, architectural constraints need to be made explicit to ensure continuous operation when allocating and migrating services from one host in the Cloud to another. In addition, the need for optimising resources and minimising over-provisioning requires service providers to control the dynamic adjustment of capacity throughout the entire service lifecycle. We discuss the implications for software architecture definitions of distributed applications that are to be deployed on Clouds. In particular, we identify novel primitives to support service elasticity, co-location and other requirements, propose language abstractions for these primitives and define their behavioural semantics precisely by establishing constraints on the relationship between architecture definitions and Cloud management infrastructures using a model denotational approach in order to derive appropriate service management cycles.

1. INTRODUCTION

Until recently, operating systems managed the allocation of physical resources, such as CPU time, main memory, disk space and

*This research has been partly funded by the RESERVOIR EU FP7 Project through Grant Number 215605.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

HPDC 2010 21-25 June, Chicago, USA
Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

network bandwidth to applications. Virtualisation infrastructures, such as Xen [7] and VMWare [19] are changing this by introducing a layer of abstraction known as a *hypervisor*. A hypervisor runs on top of physical hardware, allocating resources to isolated execution environments known as *virtual machines*, which run their own individual virtualised operating system. Hypervisors manage the execution of these operating systems, booting, suspending or shutting down systems as required. Some hypervisors even support replication and migration of virtual machines without stopping the virtualised operating system.

It turns out that the separation between resource provision and operating systems introduced by virtualisation technologies is a key enabler for Cloud computing¹. Compute Clouds provide the ability to lease computational resources at short notice, on either a subscription or pay-per-use model and without the need for any capital expenditure into hardware. A further advantage is that the unit cost of operating a server in a large server farm is lower than in small data centres. Examples of compute Clouds are Amazon's Elastic Compute Cloud (EC2) [1] or IBM's Blue Cloud [2]. Organisations wishing to use computational resources provided by these Clouds supply virtual machine images that are then executed by the hypervisors running in the Cloud, which allocate physical resources to virtualised operating systems and control their execution.

With an increasing number of providers seeking to migrate services to the cloud in order to save on deployment costs, cater for rapid growth or generally relieve themselves from the responsibility of provisioning the infrastructural resources needed to support the service, whether related to power, bandwidth, software or hardware [6], there is a crucial need to ensure that a same service quality can be retained when relying upon clouds while generally delivering on the promise of lowering costs by minimising overprovisioning through efficient upscaling and downscaling of services.

In this paper we review the implications of the emergence of virtualisation and compute Clouds for software engineers in general and software architects in particular. We find that software architectures need to be described differently if they are to be deployed into a Cloud. The reason is that scalability, availability, reliability, ease of deployment and total cost of ownership are quality attributes that

¹Specifically, virtualisation is an enabler for Infrastructure-as-a-Service (IaaS) Clouds, the type on which this paper is focused. For a description of the different types of Cloud computing see [22]

need to be achieved by a software architecture and these are critically dependent upon how hardware resources are provided. Virtualisation in general and compute Clouds in particular provide a wealth of new primitives that software architects can exploit to improve the way their architectures deliver these quality attributes.

The principal contribution of this paper is a discussion of architecture definition for distributed applications that are to be deployed on compute Clouds. The key insight is that the architecture description needs to be reified at run-time so that it can be used by the Cloud computing infrastructure in order to implement, monitor and preserve architectural quality goals. This requires several advances over the state of the art in software architecture. Architectural constraints need to be made explicit so that the Cloud infrastructure can obey these when it allocates, activates, replicates, migrates and deactivates virtual machines that host components of the architecture. The architecture also needs to describe how and when it responds to load variations and faults; we propose the concept of elasticity rules for architecture definitions so that the Cloud computing infrastructure can replicate components and provide additional resources as demand grows or components become unavailable. Finally, there must be an understanding of how a management cycle for a service deployed on a Cloud can be derived from a description of these constraints and we demonstrate how this can be achieved using a *model driven approach*.

As such the overall contributions of the paper are as follows: we identify a list of requirements and constraints that a provider must describe when deploying and hosting a multi-component application on a cloud. We present a language for the description of such requirements that builds on existing standards and introduces new abstractions such as the ability to specify on demand scaling. We define an architecture for a cloud infrastructure to support these abstractions and specify clear behavioural semantics for our language with respect to this architecture using a model denotational approach. Finally we evaluate our language primitives experimentally with a distributed computational chemistry application.

This paper is further structured as follows: In Section 2, we present the background to this research and in particular the primitives that are now provided by modern virtualisation and Cloud computing technologies in general and the infrastructure that has been developed by the RESERVOIR project in particular. We then use a motivating example in Section 3 to argue why architecture definition for Cloud computing differs from more conventional deployments. In Section 4, we describe our novel abstractions, such as co-location constraints and elasticity rules for describing architectures that are to be deployed in a Cloud. In Section 5 we describe the experimental evaluation of our approach by means of a computational chemistry application that we have deployed in a computational Cloud. We discuss related work in Section 6 and conclude the paper in Section 7.

2. BACKGROUND

The *Resources and Services Virtualization without Barriers* (RESERVOIR) project is a European Seventh Framework Programme (FP7) project which aims to promote through standardisation an open architecture for federated Cloud computing. It defines an open framework of modular components and APIs that enable a wide range of configurations within the Cloud space, focusing on Infrastructure as a Service (IaaS) Clouds [22].

The RESERVOIR architecture aims to satisfy the vision of service oriented computing by distinguishing and addressing the needs of *Service Providers*, who understand the operation of particular businesses and offer suitable *Service* applications, and *Infrastructure Providers*, who lease computational resources in the form of

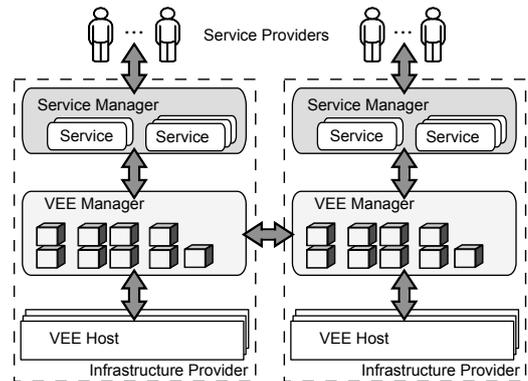


Figure 1: RESERVOIR architecture

a Cloud computing infrastructure. This infrastructure provides the mechanisms to deploy and manage self contained services, consisting of a set of software components, on behalf of a service provider.

The Cloud computing abstractions supported by this infrastructure and the architecture we have used to implement these abstractions are described in detail in [16]. The architecture is shown in Figure 1. A partial implementation of the architecture (the VEEM layer), which we have used in our experimental evaluation described below is publicly available from OpenNebula.org. It is also included from release 9.04 in the Ubuntu Linux distribution.

RESERVOIR relies on generally available virtualisation products, such as Xen [7] or VMWare [19]. The lowest layer of the RESERVOIR architecture is the *Virtual Execution Environment Host* (VEEH). It provides plugins for different hypervisors and enables the upper layers of the architecture to interact with heterogeneous virtualisation products. The layer above is the *Virtual Execution Environment Manager* (VEEM), which implements the key abstractions needed for Cloud computing. A VEEM controls the activation of virtualised operating systems, migration, replication and de-activation. A VEEM typically controls multiple VEEHs within one site. The key differentiator from other Cloud computing infrastructure is RESERVOIR's ability to federate across different sites, which might be implementing different virtualisation products. This is achieved by cross-site interactions between multiple different VEEMs operating on behalf of different Cloud computing providers. This supports replication of virtual machines to other locations for example for business continuity purposes. The highest level of abstraction in the RESERVOIR architecture is the *Service Manager*. While the VEEM allocates services according to a given placement policy, it is the Service Manager that interfaces with the Service Provider and ensures that requirements (e.g. resource allocation requested) are correctly enforced. The Service Manager also performs other service management tasks, such as accounting and billing of service usage.

3. MOTIVATION

How is a software development project now going to use Cloud computing infrastructures, such as the one provided by RESERVOIR? We aim to answer this question and then derive the main research questions for this paper using a running example, an enterprise resource planning system, such as those provided by SAP [3]. A high-level architecture of an SAP system is illustrated in Figure 2. SAP ERP systems have a multi-tiered software architecture with a relational database layer. On top of the database is an appli-

cation layer that has a *Central Instance*, which provides a number of centralised services, such as synchronisation, registration and spooling capabilities, and generally serves as a database gateway. Moreover SAP applications have a number of *Dialog Instances*, which are application servers responsible for handling business logic and generating HTTP-based dialogues that are shown in a browser. A *Web Dispatcher* may be used to balance workloads between multiple dialog instances. To date SAP systems are hosted in substantial data centres, but organisations (playing the role of Service Provider) in the future might wish to deploy it in a compute Cloud in order to avoid the significant capital investment associated with the construction of a data centre.

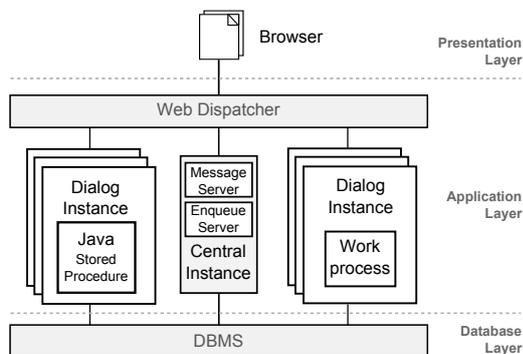


Figure 2: SAP three-tiered architecture

If the SAP system is handed over to a Cloud computing provider that uses an infrastructure based on the RESERVOIR architecture, a number of architectural constraints of the SAP system will need to be obeyed by the Cloud. For example, the Central Instance will frequently make queries to the database and in a typical SAP configuration the Central Instance and the database need to be co-located on the same LAN subnet. When the VEEM allocates the virtual machines for the Central Instance and the DBMS to particular VEEHs, it will have to respect this co-location constraint. Another architectural constraint is that the Central Instance can not be replicated in any SAP system. Dialog Instances, on the other hand are replicated to accommodate growing demand. Therefore these architectural constraints have to be expressed by the SAP provider and made available to the Service Manager so that they can obey these constraints at run-time.

This requires the software architecture to be expressed in terms of services, their relationship and their resource requirements. While a traditional SAP deployment would require that spare capacity is retained to deal with peaks in demand, Cloud computing introduces the ability to minimise overprovisioning by tying resource provision directly to the needs of the application. Resource requirements, in this case dictated by the number of Dialog Instances required to handle the current load, can be scaled dynamically and adjusted to maximise cost savings. This requires a means of describing the state of the system and rules for adjustments.

4. ARCHITECTURE DEFINITION

4.1 Requirements

Based on an understanding of the underlying hosting infrastructure described in Section 2, and the example provided in Section 3, we can break down the core issues that must be defined when deploying and running software systems on a Cloud computing infrastruc-

ture into the following high-level requirements, in order to provide a suitable definition for the terminology employed in later sections:

MDL1 Software composition: *A software system may be composed of one or more loosely coupled components, which may have differing resource (e.g. CPU, memory) and software (e.g. operating system, libraries, disk image) requirements.*

The components of the multi-layered SAP system, the Web Dispatcher, Central Instance, Dialog Instance and DBMS, will have varying hardware and software requirements, but will nevertheless be required to be managed jointly. We can expect for example the DBMS service to be very I/O and memory intensive and with large storage requirements. In contrast, the Dialog Instances may be more processor intensive, and hardware requirements may be adjusted accordingly.

MDL2 Network topology: *The system may require a specific network topology both to interconnect components of the system and communicate with external systems.*

With respect to the SAP system, the Web Dispatcher should provide an external interface and internal components should be at the very least interconnected, though external access may not necessarily be required.

MDL3 Capacity adjustment: *Hardware requirements may evolve during the lifetime of the system according to workload, time or other application-level variables.*

In order to deal with potential increases in requests, it may be necessary to deploy additional Dialog Instances in order to facilitate load balancing and ensure a certain level of performance.

MDL4 Dependencies: *Deployment and un-deployment dependencies may exist between components.*

The order in which components of an SAP system are started or stopped may affect the overall operation of the system. The DBMS and Central Instance components, serving as the backbone of the system, should be active before individual Dialog Instances.

MDL5 Location constraints: *Constraints on the distribution of service components across physical locations may exist.*

Federation of Clouds is key to enabling scalable provisioning of services. However along with the ability to seamlessly deploy services across multiple physical and administrative domains comes a need to allow service providers to control the “spread” of the application by defining clear constraints on the distribution of services across sites. These constraints can be of a technical nature (e.g. deploy certain components on a same host) or administrative (e.g. avoid un-trusted locations). Though we have, for example, established that the Central Instance and DBMS should be located on a same (virtual) network, a service provider may wish to minimise latency by ensuring proximity.

MDL6 Customisation: *Components may be dependent on configuration parameters not known until deployment.*

When deploying multiple instances of a same component, certain application-level parameters may be instance specific. As such it may be necessary to customise individual instances upon their creation and deployment. Dialog Instances may for example require the IP addresses of the Central Instance and DBMS to be provided, if this information is not known at pre-deployment time (e.g. dynamic IP allocation via DHCP).

In order to automate the management of a software system on a Cloud infrastructure it is necessary for a service provider to com-

municate both the software system stack (OS, middleware, application, configuration, and data) providing self contained services in the form of a virtualised image (addressing requirement MDL1) and a description of these requirements in the form of a *Service Definition Manifest* (addressing requirements MDL2-MDL6). The manifest therefore serves as a contract between service and infrastructure providers regarding the correct provisioning of a service. It hence reifies key architectural constraints and invariants at run-time so that they can be used by the Cloud.

To define manifests, we require a declarative language whose syntax should be sufficiently flexible to cater for a general purpose service provisioning environment, and provide the necessary abstractions to describe capacity and operational requirements of the software architecture both at deployment time and throughout the entire lifecycle.

We rely in our implementation on the *Open Virtualisation Format* (OVF) [5], a DMTF standard backed by VMWare and XenSource which aims to offer a portable packaging mechanism in a portable and platform neutral way. Building on open standards facilitates interoperability particularly in the context of federation and eases compatibility with existing services and tools. In addition it ensures that as cloud technology matures, continued compliance with the standard avoids vendor lock-in and potential deployment on newer platforms. OVF hence serves as a building block for our manifest, and provides the syntax for the description of virtual disks, networks, resource requirements and other issues related to dependencies or customisation. However, OVF (as other service description languages for existing virtualisation technologies) primarily caters for the initial distribution and deployment of fixed size services [5], which does not by itself fully realise the vision of Cloud computing.

Indeed, Clouds differ from traditional software deployment in many ways. Beyond the impact of virtualisation on multi-component architectures, existing deployment mechanisms are typically one-way “channels” where a service is configured and deployed according to an initial deployment descriptor. There is no feedback mechanism to communicate specific state, parameters and other information from a deployed service back to the infrastructure to adapt the execution environment dynamically. The manifest should enable the automation of provisioning and management through template based provisioning, where the service manifest is used as a template for easily provisioning instances of the application, and support for resource consumption control.

We hence need to add a number of abstractions to OVF, the primary being **elasticity specification** in the form of rules allowing conditions related to the state and operation of the service, such as application level workload, and associated actions to follow should these conditions be met, **application domain description**, which allow the state of the application to be described in the form of monitorable application level parameters and **placement and co-location constraints**, which identify sites that should be favoured or avoided when selecting a location for a service.

In previous work [11], we have discussed a number of additional extensions to the OVF syntax to support Clouds, including attribute and section changes to incorporate support for service components IDs in elastics arrays, cross virtual machines reference, IP dynamic addresses and elasticity rules and bounds. However, a syntactic definition of a deployment descriptor only forms part of what is necessary to ensure these requirements are met with respect to the underlying Cloud computing infrastructure.

Indeed, there must exist a clear understanding of how we derive from the language used to express the requirements of the Service Provider a *management cycle*, which will consist of several actions

being taken throughout the lifetime of a service to ensure a certain service quality being obtained. Using the RESERVOIR framework as a reference, and examining specifically issues related to dynamic capacity adjustment and service deployment, we now describe how the *behavioural semantics* for our manifest language are described and how they guide the operation of underlying Cloud components.

Focusing specifically on elasticity and application domain description, as well as service deployment, we refine and extend in this paper our OVF based service definition language syntax to incorporate these abstractions.

4.2 Manifest Language Definition

In this section, we describe the overall approach undertaken to define and provide support for the *Manifest Language*. This is achieved through the specification of three complementary facets of the language: the abstract syntax, the well-formedness rules, and the behavioural semantics. The abstract syntax of the manifest language is modelled using the *Essential Meta-Object Facility* (EMOF), an OMG standard part of the Model Driven Architecture initiative [14] for describing the structure of meta-data, and embedded within an object-oriented model of the RESERVOIR architecture. Because the manifest describes the way in which a RESERVOIR based infrastructure should provision a service application, the semantics of the language can be expressed in the model denotational style to define semantics that we introduced in [18] as constraints between the abstract syntax and domain elements that model the operation of Cloud infrastructure components. These constraints are formally defined using the *Object Constraint Language* (OCL) [15], a language for describing consistency properties, providing the static and behavioural semantics of the language. In this manner the language serves to constrain the behaviour of the underlying infrastructure, ensuring the correct provisioning of the software system services.

The motivations for this approach are two-fold: firstly by modelling the syntax of the manifest language as an EMOF model, we seek to express the language in a way that is independent of any specific implementation platform. Components of a cloud infrastructure such as RESERVOIR may rely on a number of different concrete languages, whether implementation languages (Java, C++, etc.), higher-level “meta” languages (HUTN, XML, etc.), or even differing standards (WS-Agreement, OVF, etc.). A higher level of abstraction ensures that we free ourselves from implementation specific concerns, and allows seamless and automated transitions between platform specific models as required by components.

Secondly, providing a clear semantic definition of the manifest using OCL allows us to identify functional characteristics that service management components should present in order to support capabilities such as application based service elasticity, again irrespective of the implementation platform. As such the definitions presented in this paper extend beyond the scope of RESERVOIR or any specific cloud infrastructure, instead providing a clear understanding of expected provisioning behaviours, with respect to identified and required component interfaces.

Finally, we may also consider that clear semantics ensure that we limit ambiguities when it comes to interpretation of the manifest. This is of crucial importance where financial liabilities may exist; a formal understanding of the nature of the service being provided is required in order to ensure that the service is provisioned as expected by both parties, and in a way that both can evaluate to be correct, either through run-time monitoring capabilities or historical logs. We achieve this by tying the specification of the manifest to the underlying model of the Cloud infrastructure.

4.2.1 Abstract syntax

The abstract syntax of the manifest describes the core elements of the language and their accompanying attributes. The core syntax relies upon, as previously stated, OVF [5]. The OVF descriptor is an XML-based document composed of three main parts: description of the files included in the overall architecture (disks, ISO images, etc.), meta-data for all virtual machines included, and a description of the different virtual machine systems. The description is structured into various “Sections”. Focusing primarily on the most relevant, the `<DiskSection>` describes virtual disks, `<NetworkSection>` provides information regarding logical networks, `<VirtualHardwareSection>` describes hardware resource requirements of service components and `<StartupSection>` defines the virtual machine booting sequence.

Incorporating the OVF standard ensures that we tackle several of the requirements identified in Section 4.1, providing the manifest language with a syntactic model for the expression of physical resource requirements and hardware configuration issues. We introduce new abstractions in the form of extensions to the standard rather than create new independent specifications. OVF is extensible by design and doing so ensures continued compatibility with existing OVF-based systems.

We model these extensions using EMOF. EMOF models are very similar to UML class diagrams, in that they describe classes, the data they contain and their relationships, but are at a higher level of abstraction: they describe the constructs, rules and constraints of a model. As such, EMOF is typically used to define the syntax of languages.

Application description language.

Reliance on Cloud computing introduces the opportunity to minimise overprovisioning through run-time reconfiguration of a service, effectively limiting resource consumption to only what is currently required by the application. However, when dealing with rapid changes in service context and load, timely adjustments may be necessary to meet service level obligations which cannot be met by human administrators. In such a case, it may be necessary to automate the process of requesting additional resources or releasing existing resources to minimise costs.

This automated scaling of service capacity to support potential variations in load and demand can be implemented in numerous ways. Application providers may implement such scaling at the application level, relying on an exposed interface of the Cloud computing infrastructure to issue specific reconfiguration requests when appropriate. Alternatively, they may have a desire to keep the application design free of infrastructure specific constraints and opt instead to delegate such concerns to the infrastructure itself. With a sufficient level of transparency at the application level for workload conditions to be identified, and through the specification of clear rules associating these conditions with specific actions to undertake, the Cloud computing infrastructure can handle dynamic capacity adjustment on behalf of the service provider.

It is the latter approach that we have chosen to adopt in the context of RESERVOIR. By providing a syntax and framework for the definition and support of elasticity rules, we can ensure the dynamic management of a wide range of services with little to no modification for execution on a Cloud. With respect to the syntax, we can identify the two following subsets of the language that would be required to describe such elasticity: service providers must first be able to describe the application state as a collection of *Key Performance Indicators* (KPIs), and the means via which they are obtained in the manifest. These will serve as a basis for the formulation of the rules themselves, described in the following subsection.

Because we do not want the manifest language to be tied to any specific service architecture or design, it is necessary to decouple the KPI descriptions from the application domain via the use of an *Application Description Language* (ADL). Though it is possible to build elasticity conditions based purely on infrastructure level performance indicators, this may prove limiting. Indeed, the disk space, memory or CPU load may not accurately reflect the current needs of the application, as will be seen in the evaluation. This language will allow the description of services in terms of components, parameters of interest and their monitoring requirements.

Alongside the syntactic requirements, a suitable monitoring framework must exist. A service provider is expected to expose parameters of interest through local *Monitoring Agents*, responsible for gathering suitable application level measurements and communicating these to the service management infrastructure. Though communication protocols with the underlying framework are outside the scope of the manifest language, there must exist a correlation between the events generated by the monitors and the KPIs described in the manifest. This is modelled in Figure 3.

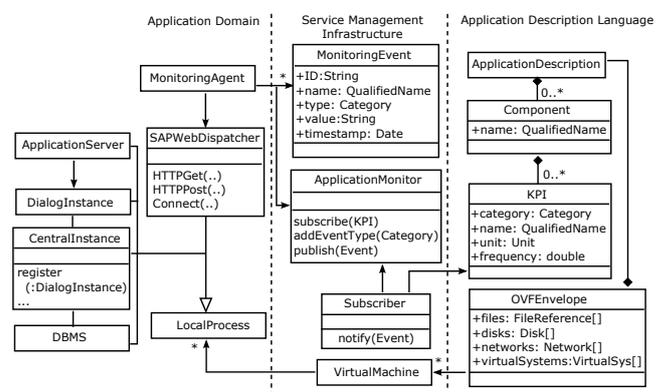


Figure 3: Application Description Language

Based on our running example, the figure exemplifies the relationship between the ADL, the RESERVOIR application-level monitoring infrastructure, and the application domain. The syntax of the ADL consists of one or more named components, with a number of associated KPIs. These KPIs are identified using appropriate qualified names (e.g. `com.sap.webdispatcher.kpis.sessions`), that will allow the underlying infrastructure to identify corresponding events obtained from an application level monitor and forward these to subscribers responsible for the enforcement of elasticity rules.

We are concerned in this example with the number of simultaneous web sessions managed by the web dispatcher, as there is a proportional relationship between resource requirements and sessions. The number of simultaneous sessions will be used as a basis for scaling the number of Dialog Instances. However, directly monitoring the traffic to and from the web dispatcher would be impossible, as SAP uses proprietary protocols. The SAP system can nonetheless report the number of sessions provided an appropriate query is formulated. The monitoring agent would be responsible for such queries and forwarding obtained responses, bridging the gap between application and monitoring infrastructure.

KPI qualified names would be considered global within the scope of a service. If there exists a need to distinguish the KPI measurements produced by multiple instances of a same component, this is achieved by using distinct qualified names. Monitoring agents can, for example, include instance IDs in the qualified name.

The structure of the qualified name itself would not fall within the scope of the manifest specification. Instances of an application service as a whole however would be considered distinct. At the implementation level, KPIs published within a network are tagged with a particular service identifier, and rules, covered below, will also be associated with this same identifier. Multiple instances of an application service would hence operate independently.

Elasticity Rules.

With respect to the rule syntax, we adopt an *Event-Condition-Action* approach to rule specification. This is a widely adopted model for rule definition, adopted for example in active databases and rule engines, and suited in this instance. Based on monitoring events obtained from the infrastructure, particular actions from the VEEM are to be requested when certain conditions relating to these events hold true. This requires rules to be expressed with respect to the interface of the underlying VEEM and monitoring events.

A representation of the elasticity rules based on a general rule-base model and their relationship to monitoring events and the Cloud infrastructure is illustrated in Figure 4. The syntax specifies conditions, based on monitoring events at the application layer or otherwise, which would lead to specified actions, based on a set of operations presented by the VEEM. The operations, modelled on the OpenNebula framework capabilities will involve the submission, shutdown, migration, reconfiguration, etc. of VMs and should be invoked within a particular time frame. The conditions are expressed using a collection of nested expressions and may involve numerical values, arithmetic and boolean operations, and values of monitoring elements obtained. The relationship between KPIs specified in the manifest and these events has been described in the previous section. The elasticity rules will be supervised by the Cloud infrastructure at the Service Manager layer during the running of the software system and it is expected that a rule interpreter will receive events from the infrastructure or application monitors and trigger such operations accordingly.

With respect to the example, this language enables us to express that virtual machines with new Dialog Instances should be created as the number of user sessions maintained by the SAP web dispatcher grows in order to handle the increased load. A concrete example of an elasticity rule will be provided in Section 5.

It is worth briefly discussing the subject of time management. The service provider controls the timeliness of the response in multiple ways. Firstly the rate at which monitoring events are sent by the application level monitor is entirely dictated by the application and this should be balanced against expected response time to avoid duplicate responses. Secondly service providers can specify a time frame within which particular actions should take place, as described above. Finally, the current time can be introduced as a monitorable parameter if necessary.

Additionally service providers may prefer expressing conditions regarding a series of measurements within a time frame rather than focusing on single events. We may be concerned here with the average number of active sessions in a window in order to limit the impact of strong fluctuations. While the language is extensible and presents the opportunity to provide such functionality, and we are currently working on the ability to specify a time series and operations related to that time series (mean, minimum, maximum, etc.), this can be achieved by aggregating measurements at the application level, with the monitoring agent performing such tasks.

Elasticity rules can be a powerful tool to express capacity constraints. The structure is kept purposely simple: not intended as a substitute for a programming language, elasticity rules only aim to inform the Cloud infrastructure of the corrective process to be

undertaken. Auto-scaling is not a form of capacity planning but it aids in introducing a certain degree of flexibility in resource allocation which ensures that strong and often unexpected variations in demand can be met. In general, more complex relationships between performance indicators can be computed at the application level, before being forwarded to the service manager.

4.2.2 Semantic Definition

We examine in this section the dynamic semantics of the manifest language as OCL constraints on the relationship between the syntactic model of the manifest and the infrastructure and application domains. Dynamic semantics are concerned with deployment and run-time operation. These will specify behavioural constraints that the system should adhere to during its execution.

The question of how and when we verify that these constraints hold true during the provisioning of a service should be discussed briefly. Defined invariants should be true at any moment in time, however it is not feasible in practice to continuously check for this. Instead it is preferable to tie the verification to monitoring events or specific actions, such as a new deployment. Another question to be posed is what should be done when an evaluation of the state system does not fit the specified constraints. This will depend on the context: an exception may occur, or an operation should be invoked to trigger some corrective action, as would be the case with elasticity rules.

Service Deployment.

As the manifest is processed by the various independent components of the Service Manager to generate a deployment descriptor for submission to the VEEM, it becomes important to ensure that the final product, which may be expressed using a different syntax, is still in line with the requirements of the service provider. In the case of RESERVOIR, the VEEM would introduce it's own deployment template. Using OpenNebula as a reference implementation of a VEEM, the deployment template relied upon by the system is roughly based on a Xen configuration file. The association between manifest and deployment template is illustrated in Figure 5.

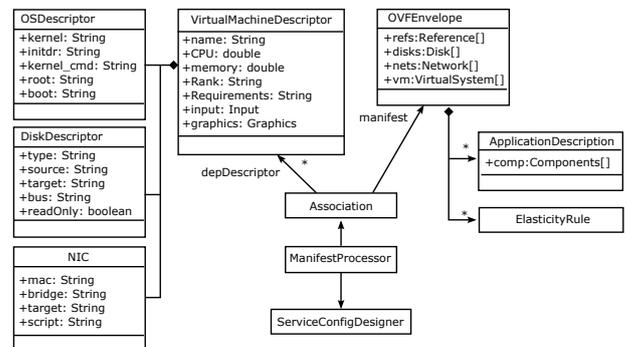


Figure 5: Service Manifest and deployment descriptor

It is presumed that the `ManifestProcessor` will be responsible for parsing the manifest and generating one or more deployment templates accordingly. The `ServiceConfigAnalyzer` may be used to further optimise the placement with regards to the multiple sites at which it may be deployed, though the manifest specification is not concerned with this. It is only necessary to ensure that the optimisation process respects certain constraints regarding resource requirements. This is a *design by contract* approach [23]. We are not

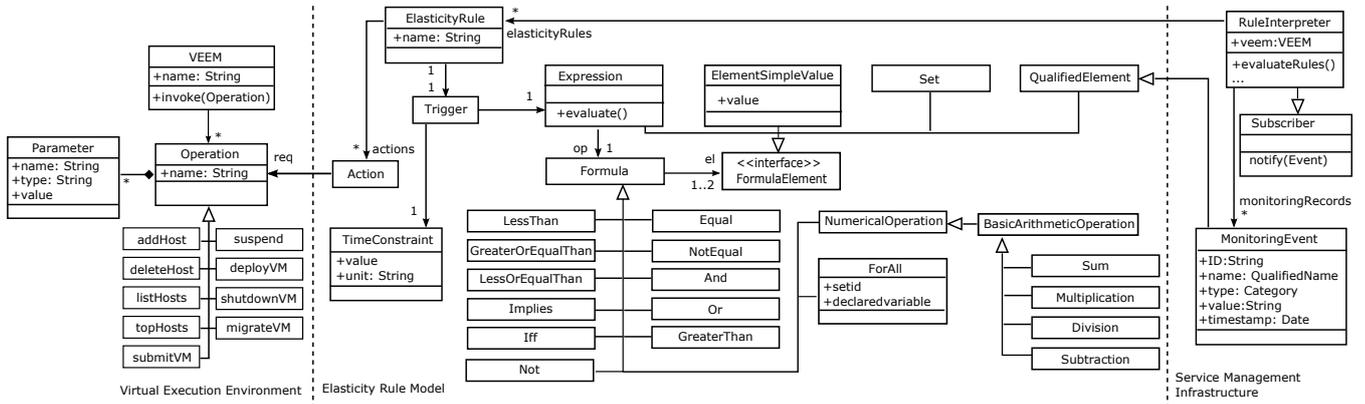


Figure 4: Elasticity Rules Syntax

concerned with the actual transformation process, but rather that the final product, i.e. the deployment descriptor, respects certain constraints. These can be expressed in OCL as follows:

```

context Association
inv:
manifest.vm -> forAll ( v |
depdescriptor . exists ( d |
d . name = v . id &&
d . memory = v . virtualhardware . memory &&
d . disk . source =
(manifest . refs . file ->asSet()->
select (id = v . id)->first () . href
...
)

```

This OCL description is a sample of what is required to establish a relationship between manifest and deployment descriptor. Here, we describe that there should be at least one deployment descriptor generated for every virtual system described in the manifest definition that has the same identifier and memory requirements. The full OCL specification contains a full mapping of attributes of our manifest language to that of a VEEM deployment descriptor.

Service Elasticity.

Similarly, we can specify the expected outcome of elasticity rule enforcement with respect to both the syntax of the manifest and the underlying RESERVOIR components. OCL operations are *side effect free*, as in they do not alter the state of the system. Nevertheless they can be used to verify that the dynamic capacity adjustments have indeed taken place when elasticity rule conditions have been met, using the *post* context.

This is described in OCL as follows:

```

— Collect monitoring records upon notification
context RuleInterpreter :: notify (e : Event)
post : monitoringRecords =
monitoringRecords @pre ->append (e)
— Evaluate elasticity rules and check adjustment
context RuleInterpreter :: evaluateRules ()
post : elasticityRules ->forAll (er |
if self . evaluate (er . expr) > 0 then
er . actions ->forAll (a | veem ^ invoke (a . req))
else true
endif)
— Query simple type value
context RuleInterpreter :: evaluate (el :
ElementSimpleType) : Real
post : result = el . value
— Obtain latest value for monitoring record
— with specific qualified name
context RuleInterpreter :: evaluate (qel :
QualifiedElement) : Real
post :

```

```

if monitoringRecords ->select (name=qe . name)
->last()->exists () then
result = monitoringRecords
->select (name=qe . name)
->last (). value
else result = qel . default
endif
— Evaluate expressions
— Defined as post in order to use recursion
context RuleInterpreter :: evaluate (expr :
Expression) : Real
post :
if expr . op . oclIsTypeOf (GreaterThan) then
if self . evaluate ((op . el ->first ()) >
self . evaluate ((op . el ->last ()) then
result = 1
else result = 0
end if
else
....
endif
endif

```

This OCL pseudo-code is only a subset of the complete OCL specification that aims to illustrate how we can specify the correct execution of elasticity rules with respect to the rule syntax. The code is split into a number of individual segments. The first simply states that monitoring events obtained are expected to be collected as records for the purpose of later evaluation. The second element states that if the conditional element of one of the elasticity rules is found to be true, then particular actions should have been invoked. To reiterate, the operations are side effect free, implying that no processing of any kind will take place in an OCL statement. Instead we only check that there has been communication with the VEEM to invoke certain operations if the conditions described hold true. How it is implemented is then left to developers.

The final segments relate to the evaluation of the conditions themselves. The `RuleInterpreter::evaluate(qel: QualifiedElement)` describes that upon evaluation of the rules, values for key performance indicators described in the document are obtained from the monitoring records, by examining the latest monitoring event with matching qualified name. This defines the relationship between KPIs and monitoring events. This asynchronous model is chosen because the Cloud infrastructure does not control application level monitoring agents. As there is no guarantee over how often monitoring information is provided, and rules may involve measurements from several services, it is for the implementation to determine when the rules should be checked to fit within particular timing constraints rather than tying checks to the reception of any specific monitoring event. Finally the last segment illustrates the recursive evaluation of expressions based on the type

of formula selected by the service provider.

4.2.3 Concrete syntax

While the specification of our manifest language is kept free of implementation concerns, the model-denotational approach adopted here provides a basis for automatically deriving concrete human or machine readable representations of the language that can be manipulated by the end-user or processed directly by the RESERVOIR based infrastructure. Moreover, beyond creating and editing the manifest itself, the syntax and accompanying semantics can be used as input for a generative programming tool to automate the generation of applications to control the service provisioning process.

In practice, the RESERVOIR architecture may be implemented using a wide range of programming languages and existing technologies. The semantic definition described in this paper will generally serve as an important software engineering artefact, guiding the design and development of components. However, the potential for errors to occur during the provisioning process always exists, due to implementation or a failure to correctly interpret the specification of our language. We can assist in identifying and flagging such errors by *programmatically* generating monitoring instruments which will validate run-time constraints previously described in Section 4.2.2.

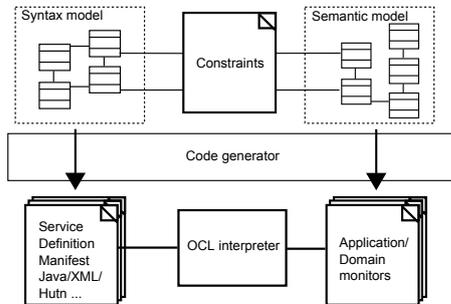


Figure 6: Programmatic generation of monitoring instruments

The process by which this is achieved is illustrated in Figure 6. In previous work, we have developed the UCL-MDA tools [17], a graphical framework for the manipulation of EMOF and OCL implemented as a plug-in for the Eclipse IDE. The framework relies on existing standards for the transformation of EMOF models and OCL into code, such as the Java Metadata Interface (JMI) standard and OCL interpreters, and is available at uclmda.sourceforge.net.

We have extended the framework for the purpose of this work. Our extensions introduce the ability to create, edit and validate manifests describing services to be deployed on a RESERVOIR based infrastructure. Element attribute values are input via the graphical interface in accordance with the structure of the language. Infrastructure related attributes and configuration values may be included in order to verify that OCL constraints are correctly maintained. This may be used amongst other things to verify that deployment descriptors generated by the infrastructure fit within the requirements specified by the manifest as covered.

Via the interface, users can additionally request the creation of stand-alone monitoring instruments in Java capable of interaction with our implementation of the RESERVOIR framework. These are currently of two forms. The first is simply responsible for gathering and reporting the values of specific KPIs described in the manifest. The second will validate the correct enforcement of elasticity rules by evaluating incoming monitoring events and ver-

ifying where appropriate that suitable adjustment operations were invoked by matching entries and time frames in infrastructural logs. The framework also allows the generation of custom stubs which the service provider may use as a basis for the development by the service provider of monitoring agents, handling issues such as communication protocols, measurement labelling and packaging, and providing a control interface to manage frequency and operation. This would have to be supplemented with appropriate probes responsible for the application level collection of measurements.

Java code is generated from a combination of data obtained from the specification, element values input by the user and Java templates, the latter being used to bridge the gap between the abstract model of the infrastructure and the actual implementation. As previously discussed, issues such as communication channels for the distribution of monitoring events fall outside the scope of the manifest language specification. Templates provide the necessary code to gather KPI measurements or parse infrastructure logs and pass this information to OCL interpreters.

Our implementation of the RESERVOIR monitoring framework itself is described in [16]. Data sources, such as application level monitoring agents, encapsulate one or more probes responsible for the collection and publication of specific attributes and values. Communication between producers and consumers of data is divided between multiple communication planes: the control plane, to manage the execution of infrastructure level probes, such as turning them on and off and reconfiguration, the information plane, to facilitate the exchange of related meta-data, and the data plane, for the actual communication of KPI measurements. Multiple transport mechanisms may be plugged in. In order to minimise the number of connections established between end points a number of solutions are in place, including the potential use of IP multicasting as a transport mechanism, and intermediate data aggregation points, which will be responsible for collecting data packets and processing these if needed to produce new performance indicators.

The tool hence serves the following purposes: firstly, it allows users to specify and manipulate manifests. Secondly it allows the generation of code allowing the service provider to verify the correct provisioning of a service at run-time according to the semantics of the language. Finally it provides the means of interfacing a service with the RESERVOIR monitoring architecture.

5. EXPERIMENTAL EVALUATION

In the evaluation of our work we aim to prove the following hypothesis: provided that an architecture definition correctly specifies requirements and elasticity rules, and that the Cloud computing infrastructure obeys the constraints identified in the semantic definition, then the *quality of service* that can be obtained from a Cloud computing infrastructure should be equivalent to that obtained were the application hosted on dedicated resources. In addition, through the specification of elasticity rules, providers can considerably reduce expenditure by minimising over-provisioning.

We will demonstrate this hypothesis by deploying a production level service on an infrastructure consisting of the RESERVOIR stack and a custom implementation of the Service Manager. This Service Manager incorporates monitors that validate the specified constraints. The selected service is a grid based application responsible for the computational prediction of organic crystal structures from the chemical diagram [10].

The application operates according to a predefined workflow involving multiple web based services and Fortran programs. Up to 7200 executions of these programs may be required to run, as batch jobs, in both sequential and parallel form, to compute various subsets of the prediction. Web services are used to collect inputs from

a user, coordinate the execution of the jobs, process and display results, and generally orchestrate the overall workflow. The actual execution of batch jobs is handled by Condor [21], a job scheduling and resource management system, which maintains a queue of jobs and manages their parallel execution on multiple nodes of a cluster.

This case study provides many interesting challenges when deployed on a Cloud computing infrastructure such as RESERVOIR. Firstly, the application consists of a number of different components with very different resource requirements, which are to be managed jointly. Secondly, the resource requirements of the services will vary during the lifetime of the application. Indeed, as jobs are created, the number of cluster nodes required to execute them will vary. Our goal in relying upon a Cloud computing infrastructure will be to create a *virtualised cluster*, enabling the size of the cluster to dynamically grow and contract according to load.

For this evaluation, we will compare the quality of service, i.e. the duration required to complete the prediction) when executing this workflow on a dedicated cluster, compared to a Cloud computing infrastructure that provides support for our abstractions. We are not concerned here with the overhead of hypervisors such as Xen, which are well documented [8]. Instead we are concerned with evaluating the costs of dynamically adjusting the resource provisioning during the application lifecycle and determining whether an appropriate level of service can still be obtained.

5.1 Testbed Architecture

5.1.1 Service components

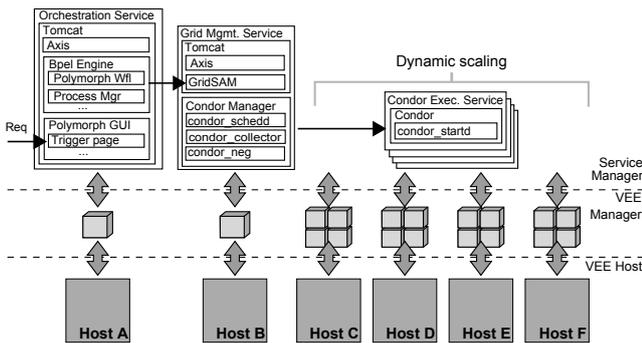


Figure 7: Testbed architecture

The testbed we use is illustrated in Figure 7. Three main types of service components can be distinguished. The *Orchestration Service* is a web based server responsible for managing the overall execution of the application. It presents an HTTP front end enabling users to trigger predictions from a web page, with various input parameters of their choice. The Business Process Execution Language (BPEL) [13], is used to coordinate the overall execution of the polymorph search, relying on external services to generate batch jobs, submit the jobs for execution, process the results and trigger new computations if required.

The *Grid Management Service* is responsible for coordinating the execution of batch jobs. It presents a web service based interface for the submission of jobs. Requests are authenticated, processed and delegated to a Condor scheduler, which will maintain a queue of jobs and manage their execution on a collection of available remote execution nodes. It will match jobs to execution nodes according to workload and other characteristics (CPU, memory, etc.). Once a target node has been selected it will transfer binary and input files over and remotely monitor the execution of the job.

The last type of component is the *Condor Execution Service*, which runs the necessary daemons to act as a Condor execution node. These daemons will advertise the node as an available resource on which jobs can be run, receive job details from the scheduler and run the jobs as local processes. Each node runs only a single job at a time and upon completion of the job transfers the output back to the scheduler, and advertises itself as available.

5.1.2 Deployment

Packaged as individual virtual machines encapsulating operating system and other necessary software components, the three components are deployed on the RESERVOIR-based infrastructure. The associated manifest describes the capacity requirements of each component, including CPU and memory requirements, references to the image files, starting order (based on service components dependencies), elasticity rules and customisation parameters. For the purpose of the experiment, the Orchestration and Grid Management Services will be allocated a fixed set of resources, with only a single instance of each being required. The Condor execution service however will be replicated as necessary, in order to provide an appropriate cluster size for the parallel execution of multiple jobs.

The elasticity rules will tie the number of required Condor execution service instances to the number of jobs in queue as presented by the Condor scheduler. This enables us to dynamically deploy new execution service instances as the number of jobs awaiting execution increases. Similarly as the number of jobs in the queue decreases it is no longer necessary to use the resources to maintain a large collection of execution nodes, and hosts in the Cloud can be released accordingly. This is expressed as follows in the manifest using an XML concrete syntax, which conforms to the abstract syntax described in Section 4.2.1. We use a similar elasticity rule for downsizing allocated capacity as the queue size shrinks.

```
<ElasticityRule name="AdjustClusterSizeUp">
  <Trigger>
    <TimeConstraint unit="ms">5000</TimeConstraint>
    <Expression>
      (@uk.ucl.condor.schedd.queueSize /
      (@uk.ucl.condor.exec.instances.size + 1) > 4) &&
      (@uk.ucl.condor.exec.instances.size < 16 )
    </Expression>
  </Trigger>
  <Action run=
    "deployVM(uk.ucl.condor.exec.ref)"/>
</ElasticityRule>
```

The elasticity rules will refer to key performance indicators that would have to be declared within the context of the application structure. This will be expressed as follows:

```
<ApplicationDescription
  name="polymorphGridApp">
  <Component name="GridMgmtService" ovf:id="GM">
    <KeyPerformanceIndicator category="Agent" type="int">
      <Frequency unit="s">30</Frequency>
      <QName>uk.ucl.condor.schedd.queueSize</QName>
    </KeyPerformanceIndicator>
  </Component>
  ...
</ApplicationDescription>
```

All components and KPIs will need to be declared in this manner. This enables the infrastructure to monitor KPI measurements being published by specific components and associate them to the declared rules according to the previously stated semantics. In this particular instance we are specifying that a monitoring agent associated with the Grid Management Service will publish measurements under the `uk.ucl.condor.schedd.queueSize` qualified name every 30 seconds as integers.

The overall management process can hence be described as follows: upon submission of the manifest, the Service Manager, which

is responsible for managing the joint allocation of service components and service elasticity, will parse and validate the document, generating suitable individual deployment descriptors to be submitted to the VEEM beginning with the Orchestration and Grid Management components. The VEEM will use these deployment descriptors to select a suitable physical host from the pool of known resources. These resources are running appropriate hypervisor technology, in this case the Xen virtualisation system, to provide a virtualised hardware layer from which the creation of new virtual machines can be requested. Upon deployment, the disk image is replicated and the guest operating system is booted with the appropriate virtual hardware and network configuration.

When the Grid Management component is operational, a monitoring agent, as described in Section 4.2.1, will begin the process of monitoring the queue length and broadcast the number of jobs in the queue on a regular basis (every 30 seconds) under the selected qualified name (`uk.ucl.condor.schedd.queue_size`). These monitoring events, combined with appropriate service identifier information, will be recorded by the rule interpreter component of the Service Manager to enforce elasticity rules. When conditions regarding the queue length are met (i.e. there are more than 4 idle jobs in the queue), the Service Manager will request the deployment of an additional Condor Execution component instances. Similarly, when the number of jobs in queue falls below the selected threshold, it will request the deallocation of new virtual instances.

The actual physical resources which are managed by the RESERVOIR infrastructure used in this experiment consist of a collection of six servers, each of them presenting a Quad-Core AMD Opteron(tm) Processor 2347 HE CPU and 8 GBs of RAM and with shared storage via NFS. OpenNebula v1.2, as the VEEM implementation, is used to manage the deployment of virtual machines on these resources according to the requirements specified by a Service Manager.

Both the Orchestration and Grid Management components will be allocated the equivalent of a single physical host each, due to heavy memory requirements, and up to 4 Condor Execution components may be deployed on a single physical host, limiting the maximum cluster size to 16 nodes. This mapping however is transparent to the Service Manager, Service Provider and application.

5.1.3 Metrics

It is also important to briefly describe the characteristics of the overall application workflow, in order to determine appropriate metrics for the experiment. Our primary indicator of quality of service is the overall *turn around time* of a prediction. The turn around time can be defined as the amount of time elapsed between the moment a client user requests a search to the moment results are displayed on the web page. As previously stated, the overall process combines functionality from a number of different Fortran programs into a larger workflow. Based on our selected input, two long running jobs will first be submitted, followed by an additional set of 200 jobs being spawned with each completion to further refine the input. We must also take into account the additional processing time involved in orchestrating the service and gathering outputs.

Another important metric to consider is that of resource usage. The goal of service elasticity is to reduce expenditures by allowing Service Providers to minimise overprovisioning. While the actual financial costs will be dependent on the business models employed by Cloud infrastructure providers, we can at the very least rely upon resource usage as an indicator of cost.

5.1.4 Experiment results

	Dedicated Environment	Cloud Infrastructure
Search turn around time (s)	8605	9220
Complete shutdown time (s)	N/A	9574
Average execution nodes for run until shutdown	16 N/A	10.49 10.42
Percentage differences		
Resource usage saving		34.46%
Extra run time (jobs)		7.15%

Table 1: Experiment Results

We compare turn-around time and resource usage obtained on our Cloud infrastructure with elasticity support with that obtained in an environment with dedicated physical resources. The objective is to verify that there are no strong variations in turn around time, but a significant reduction in resource usage. The results are illustrated in Figure 8. The number of queued jobs is plotted against the number of Condor execution instances deployed. Both charts show large increases in queued jobs as the first long running jobs complete and the larger sets are submitted. In addition, the first chart represents the execution of the application in a dedicated environment and shows a set of 16 continuously allocated execution nodes. The second chart represents the execution of the application with elasticity support, shows the increase in the number of allocated nodes as jobs in queue increases, and a complete deallocation as these jobs complete. The overall turn around time and resource usage obtained is described in Table 1.

As we can see from the results, a 7.15% increase in turn around time occurs. As there is little difference in execution times in the individual batches of jobs on either the dedicated or virtual cluster, the increase in turn around time comes primarily from the additional time that is taken to create and deploy new instances of the Condor execution service as jobs are added in the queue. This can be verified in Figure 8, where a small delay can be observed between increases in the number of jobs in queue, and the increase in Condor execution services. The overhead incurred is due to the deployment process, which will involve duplicating the disk image of the service, deploying it on a local hypervisor, and booting the virtual machine, and the registration process, which is the additional time required for the service to become fully operational as the running daemons register themselves with the grid management service. There exists ways of reducing this overhead independently of the Cloud computing infrastructure, at the expense of resource usage, such as relying on pre-existing images to avoid replication.

A 10 minute increase of time in can however be constituted as reasonable considering the overall time frame of a search, which is well over 2 hours. This is particularly true as we consider the overall resource usage savings. Indeed as can be seen in the table, with respect to execution nodes, the overall resource usage decreases by 34.46% by relying on service elasticity. This is because the totality of the execution nodes are not required for the initial bulk of the run, where only 2 jobs are to be run. It is only in the second stage that more nodes are required to handle the additional jobs.

Of course the savings here are only considered in the context of the run itself. If we consider the overall use of the application over the course of a randomly selected week on a fully dedicated environment where resources are continuously available, even more significant cost savings will exist. Examining logs of searches conducted during this period, and based on cost savings obtained here, we have estimated that overall resource consumption would drop by 69.18%, due to the fact that searches are not run continuously; no searches were run on two days of the week, and searches, though of varying size, were run only over a portion of the day, leaving re-

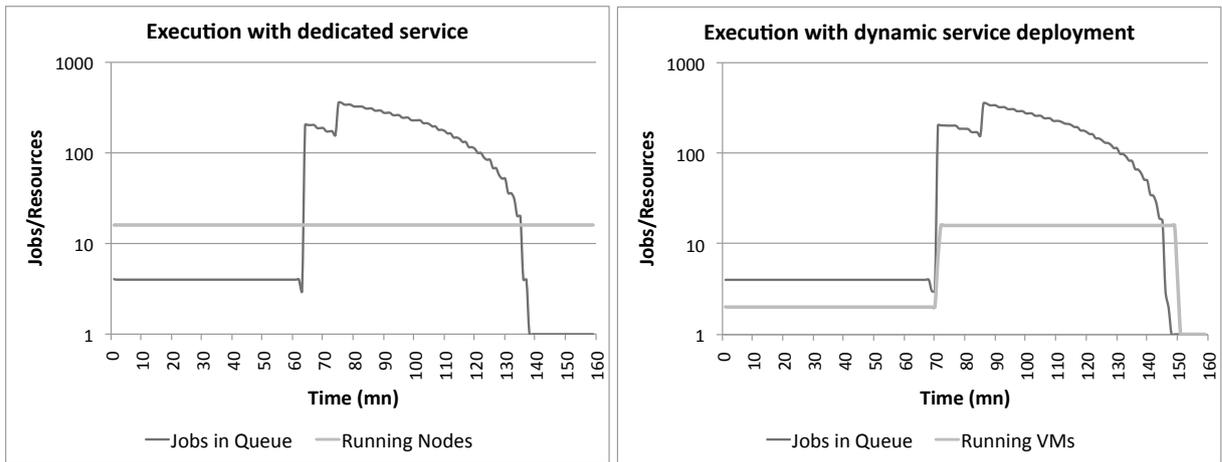


Figure 8: Job Submission and Resource Availability

sources unused for considerable amounts of time.

6. RELATED WORK

Much of this work builds on the foundation previously established with SLAs, where we used a model denotational approach to specify service level agreements for web-based application services [18]. In this paper, we have aimed to broaden the approach to encapsulate Cloud computing primitives and environment, providing a specification for a manifest language describing software architecture, physical requirements, constraints and elasticity rules.

In addition, it is worth examining research developments related to service virtualisation, grid computing and component based software architecture description languages. With respect to virtual environment deployment descriptions, the manifest language proposed here builds upon the Open Virtualisation Format (OVF) [5], whose limitations have already been discussed.

There exists a number of software architecture description languages which serve as the run-time configuration and deployment of component based software systems. xArch [9], for example, is a standard, extensible XML-based representation for software architectures. It does not in itself provide the support for physical resource requirement descriptions that may be relevant in a virtualised environment. Moreover, at the application level, we are primarily concerned with observable parameters of interest.

The CDDLM Component Model [20] outlines the requirements for creating a deployment object responsible for the lifecycle of a deployed resource with focus on grid services. Each deployment object is defined using the CDL language. The model also defines the rules for managing the interaction of objects with the CDDML deployment API. Though the type of deployment object is not suited to virtual machine management, the relationship between objects and the deployment API can be compared to our approach we have undertaken here, providing a semantic definition for the CDL language. However in our case the relationship between domain and syntactic models is performed at a higher level of abstraction, relying on OCL to provide behavioural constraints. Our specification is hence free of implementation specific concerns.

The general approach to dynamic and automated provisioning may also be compared to the self-managing computing systems associated with autonomic computing research [12]. While our approach to elasticity is *explicit*, in that providers define appropriate scaling rules based on an event condition action model, we

have laid a foundation for further methods to be developed relying on predictive and autonomic mechanisms to anticipate future allocation changes and further minimise over-provisioning, providing monitoring channels and a rule based framework for the dynamic management of services.

Finally it is important to examine current developments in production level Cloud environments, such as Amazon’s EC2 offering. In particular, auto-scaling has been introduced by Amazon to allow allocated capacity to be automatically scaled according to conditions defined by a service provider. These conditions are defined based on observed resource utilisation, such as CPU utilisation, network activity or disk utilisation. Whilst the approach laid out in this paper can be used to define elasticity rules based on such metrics, this can prove limiting. With respect to the evaluation, the need to increase the cluster size cannot be identified through these metrics as we require an understanding of the scheduling process. The ability to describe and monitor application state is crucial if we wish to correctly anticipate demand.

In addition, the focus of our paper has primarily been on Infrastructure-as-a-Service Clouds. Nevertheless it is still important to briefly discuss the relevance of this work with respect to Platform-as-a-Service (PaaS) clouds such as Windows Azure [4]. PaaS clouds provide an additional level of abstraction over IaaS clouds, providing a runtime environment for the execution of application code and a set of additional software services, such as communication protocols, access control, persistence, etc. Windows Azure allows services to be described as distributed entities: clients can specify the interfaces exposed by services, communication end points, channels and roles (web or worker) and different hardware requirements may be allocated. However a need to control the management, distribution and lifecycle of multi-component systems still exists, though with the added benefit of application specific operations being more readily exposed to the infrastructure. How this is implemented will be tied to the specifics of the platform itself but we do believe there is potential to adapt many aspects of our approach to the platform specific interfaces and tools.

7. CONCLUSION AND FUTURE WORK

In this paper, we have proposed an abstract syntax and semantic definition for a service manifest language which builds on the OVF standard and enables service requirements, deployment constraints (placement, co-location and startup/stopping order) and elasticity

rules to be expressed. We believe that clear behavioural semantics are of paramount importance to meet quality goals of both the Service and Infrastructure provider. Our model-driven approach aims to strengthen the design of the RESERVOIR stack, identifying functional capabilities that should be present and constraints the system should observe. It has also served to drive the implementation of the RESERVOIR monitoring framework.

We have shown experimentally that the implementation of these concepts is feasible and that a complete architecture definition that uses our manifest syntax can enable Cloud computing infrastructure to realise significant savings in resource usage, with little impact on overall quality of service. Given that a Cloud can deliver a near similar quality of service as a dedicated computing resource, Clouds then have a substantial number of advantages. Firstly application providers do not need to embark on capital expenditure and instead can lease the infrastructure when they need it. Secondly, because the elasticity rules enables the application provider to flexibly expand and shrink their resource demands so they only pay the resources that they actually need. Finally, the Cloud provider can plan its capacity more accurately because it knows the resource demands of the applications it provides.

While Cloud computing is still a relatively new paradigm, and as such changes in standards, infrastructural capabilities and component APIs are inevitable, defining the software architecture of services hosted on a cloud with respect to the capabilities of the underlying infrastructure is key to optimizing resource usage. This allows us to bridge the gap between application and infrastructure and provides the means for providers to retain some control over the management process.

Our work paves the way towards quality of service aware service provisioning. In future work, we aim to develop appropriate syntax and semantics for resource provisioning service level agreements. Building upon the approach laid out here, we aim to provide a framework for the automated monitoring and protection of service level obligations based on defined semantic constraints.

8. REFERENCES

- [1] Amazon Elastic Compute Cloud (Amazon EC2). [Online] <http://aws.amazon.com/ec2>.
- [2] IBM blue cloud. [Online] <http://www-03.ibm.com/press/us/en/pressrelease/22613.wss>.
- [3] SAP Enterprise Resource Planning. [Online] <http://www.sap.com/solutions/business-suite/erp/index.epx>.
- [4] Windows Azure Platform. [Online] <http://www.microsoft.com/windowsazure/>.
- [5] Open Virtualization Format Specification. Specification DSP0243 v1.0.0, Distributed Management Task Force, 2009.
- [6] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. H. Katz, A. Konwinski, G. Lee, D. A. Patterson, A. Rabkin, I. Stoica, and M. Zaharia. Above the clouds: A berkeley view of cloud computing. Technical Report UCB/EECS-2009-28, EECS Department, University of California, Berkeley, Feb 2009.
- [7] P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt, and A. Warfield. Xen and the art of virtualization. In *SOSP '03: Proceedings of the nineteenth ACM symposium on Operating systems principles*, pages 164–177, New York, NY, USA, 2003. ACM Press.
- [8] L. Cherkasova and R. Gardner. Measuring CPU overhead for I/O processing in the Xen virtual machine monitor. In *Proceedings of the USENIX annual technical conference*, pages 387–390, 2005.
- [9] E. M. Dashofy, A. van der Hoek, and R. N. Taylor. An infrastructure for the rapid development of xml-based architecture description languages. In *ICSE '02: Proceedings of the 24th International Conference on Software Engineering*, pages 266–276, New York, NY, USA, 2002. ACM.
- [10] W. Emmerich, B. Butchart, L. Chen, B. Wassermann, and S. L. Price. Grid Service Orchestration using the Business Process Execution Language (BPEL). *Journal of Grid Computing*, 3(3-4):283–304, 2005.
- [11] F. Galán, A. Sampaio, L. Rodero-Merino, I. Loy, V. Gil, L. M. Vaquero, and M. Wusthoff. Service specification in cloud environments based on extensions to open standards. In *Fourth International Conference on COMMunication System softWARE and middlewaRE (COMSWARE 2009)*, June 2009.
- [12] A. Ganek and T. Corbi. The dawning of the autonomic computing era. *IBM Systems Journal*, 42(1):5–18, 2003.
- [13] OASIS. Web Service Business Process Execution Language Version 2.0 Specification, april 2007. OASIS standard.
- [14] Object Management Group. Meta Object Facility Core Specification 2.0, OMG Document, formal/2006-01-01, 2006.
- [15] Object Management Group. Object Constraint Language (OCL) 2.0, OMG Document, formal/2006-05-01, 2006.
- [16] B. Rochwerger, A. Galis, E. Levy, J. Cáceres, D. Breitgand, Y. Wolfsthal, I. Llorente, M. Wusthoff, R. Montero, and E. Elmroth. RESERVOIR: Management Technologies and Requirements for Next Generation Service Oriented Infrastructures. In *The 11th IFIP/IEEE International Symposium on Integrated Management,(New York, USA)*, pages 1–5, 2009.
- [17] J. Skene and W. Emmerich. Engineering runtime requirements-monitoring systems using MDA technologies. *Lecture notes in computer science*, 3705:319, 2005.
- [18] J. Skene, D. D. Lamanna, and W. Emmerich. Precise service level agreements. In *ICSE '04: Proceedings of the 26th International Conference on Software Engineering*, pages 179–188, Washington, DC, USA, 2004. IEEE Computer Society.
- [19] J. Sugarman, G. Venkitachalam, and B.-H. Lim. Virtualizing IO/Devices on VMware Workstation's Hosted Virtual Machine Monitor. In *Proc. of the 2001 USENIX Annual Technical Conference*, Boston, Mass, June 2001. Usenix.
- [20] J. Tatemura. CDDL Configuration Description Language Specification 1.0. Technical report, Open Grid Forum, 2006.
- [21] D. Thain, T. Tannenbaum, and M. Livny. Condor and the Grid. In F. Berman, G. Fox, and T. Hey, editors, *Grid Computing: Making the Global Infrastructure a Reality*. John Wiley & Sons Inc., December 2002.
- [22] L. M. Vaquero, L. Rodero-Merino, J. Cáceres, and M. Lindner. A break in the clouds: towards a cloud definition. *SIGCOMM Comput. Commun. Rev.*, 39(1):50–55, 2009.
- [23] J. Warmer and A. Kleppe. *The Object Constraint Language: Getting Your Models Ready for MDA*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2003.