

MiCADO - Towards a microservice-based cloud application-level dynamic orchestrator

Hannu Visti ¹, Tamas Kiss ^{Corresp., 1}, Gabor Terstyanszky ¹, Gregoire Gesmier ¹, Stephen Winter ¹

¹ Centre for Parallel Computing, University of Westminster, London, United Kingdom

Corresponding Author: Tamas Kiss
Email address: kisst@wmin.ac.uk

In order to satisfy end-user requirements, many scientific and commercial applications require access to dynamically adjustable infrastructure resources. Cloud computing has the potential to provide these dynamic capabilities. However, utilising these capabilities from application code is not trivial and requires application developers to understand low-level technical details of clouds. This paper investigates how a generic framework can be developed that supports the dynamic orchestration of cloud applications both at deployment and at run-time. The advantages and challenges of designing such framework based on microservices is analysed, and a generic framework, called MiCADO - (Microservices-based Cloud Application-level Dynamic Orchestrator) is proposed. A first prototype implementation of MiCADO to support data intensive commercial web applications is also presented.

MiCADO – Towards a Microservice-based Cloud Application-level Dynamic Orchestrator

Hannu Visti, Tamas Kiss, Gabor Terstyanszky, Gregoire Gesmier, Stephen Winter

Centre for Parallel Computing, University of Westminster, London, UK

(H.Visti, T.Kiss, G.Z.Terstyanszky,G.Gesmier, S.C.Winter)@westminster.ac.uk

1 **Abstract**—In order to satisfy end-user requirements, many scientific and commercial applications require access to dynamically
2 adjustable infrastructure resources. Cloud computing has the potential to provide these dynamic capabilities. However, utilising these
3 capabilities from application code is not trivial and requires application developers to understand low-level technical details of clouds. This
4 paper investigates how a generic framework can be developed that supports the dynamic orchestration of cloud applications both at
5 deployment and at run-time. The advantages and challenges of designing such framework based on microservices is analysed, and a generic
6 framework, called MiCADO – (Microservices-based Cloud Application-level Dynamic Orchestrator) is proposed. A first prototype
7 implementation of MiCADO to support data intensive commercial web applications is also presented.

8 I. INTRODUCTION

9 Many scientific and commercial applications require access to computation, data or network resources based on dynamically changing
10 requirements. Applications running on distributed computing infrastructures, such as grids or clouds, typically fall into this category.
11 End-users can access these applications via desktop or web-based high-level user interfaces, such as science gateways. When
12 executing applications or accessing services via high-level user environments, users and providers both require these applications or
13 services to dynamically adjust to fluctuations in demand and serve end-users at required quality of service and speed, and also at
14 optimized cost. The challenge of developing such dynamically adaptable applications without requiring application developers to deal
15 with low level details of the underlying distributed computing infrastructure is the topic of this paper.

16 One of well-documented benefits of cloud computing [1] is its ability to supply a variable amount of resources (computational power,
17 storage, network capacity), which can scale dynamically up and down, forming the supply side of figure 1. On the demand side, we
18 can see applications that are likely to be formed of one or more services. Services can be either in-house developed or provided by
19 external suppliers or open-source communities. Services could also be shared between applications.

20 Services consume *baseline resources* that can be defined as resources consumed by the component in idle state. *Variable resources*
21 are consumed when the component performs its duties. This can include heavy computation or storage based on application use,
22 which can vary significantly based on the nature of the application.

23 Overall resource demand is the sum of baseline and variable components, and the presence of a variable component makes the overall
24 resource demand variable as well. IaaS (Infrastructure as a Service) clouds are elastic and have the ability to supply a variable amount
25 of resources. However, applications need to be specifically programmed in order to utilise this elasticity and dynamically vary the
26 amount of resources provided. Customising dynamic provision for each application individually for a specific cloud environment is
27 costly to do. Our ultimate aim is to implement a generic service or layer that provides this functionality for any application
28 automatically, and in a cloud resource agnostic way.

29 As shown in Figure 1, cloud service providers are indifferent to the resource needs of applications since they have no means of
30 predicting application capacity demand behaviour. In practice, the operator of an application requests cloud resources based on static
31 predictive estimates (typically, worst-case estimates), but after being commissioned, these resources remain static without operator
32 intervention. If demand exceeds supply at a given point of time, applications do not function within their required parameters. If
33 supply exceeds demand, resources are wasted, which generally has a cost impact.

34 This work investigates possibilities of replacing manually adjusted supply of cloud services with an automatically adjusted supply, as
35 also illustrated in Figure 1. The aim is to create a framework, where automatically adjusted cloud service supply can be arranged,
36 based on application component demands. Such framework would allow cloud application developers to build cost and performance
37 optimization mechanisms into their application code through a high-level API.

38 The suggested solution is based on microservices and their dynamic orchestration in a cloud computing environment. The rest of this
39 paper defines a generic microservices-based architecture for application-level cloud orchestration, called MiCADO (Microservices-
40 based Cloud Application level Dynamic Orchestrator), and describes its first prototype implementation utilizing container-based open
41 source cloud technologies.

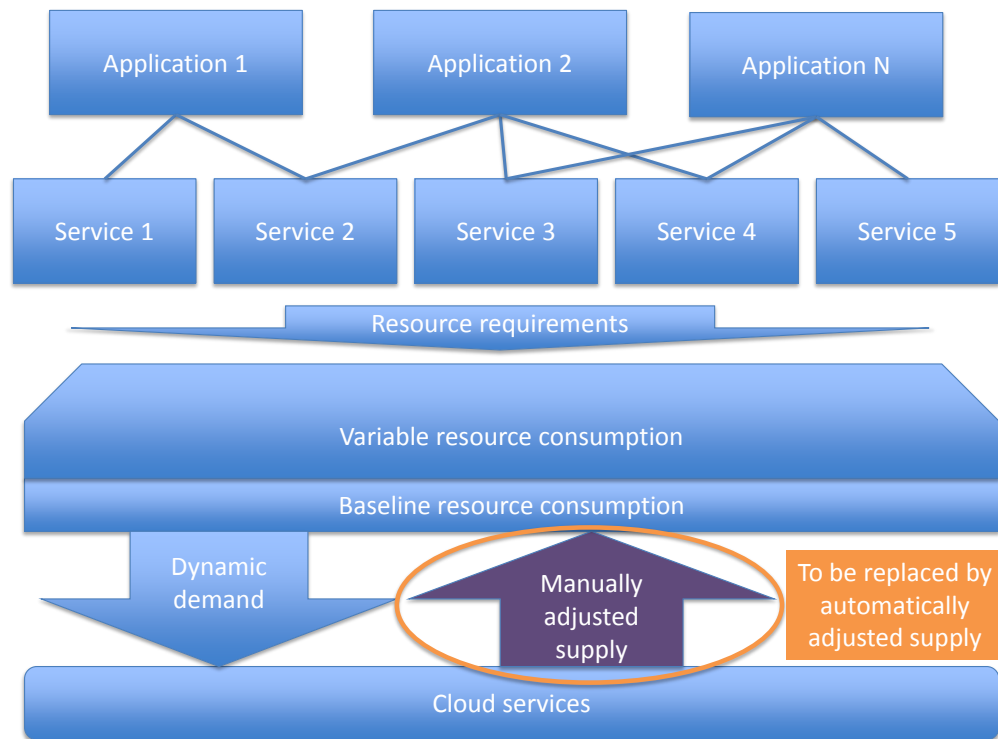


Fig. 1: Resource demand and supply of cloud applications

II. RELATED WORK

The problem of application-level orchestration has been recognized and a number of solutions have been designed and implemented. Marpaung et al. [2] discuss Altocumulus, AppScale, Cloudify and mOSAIC. Altocumulus focuses on deploying web applications to variety of public clouds, which limits its usability in private or hybrid clouds. It does not provide monitoring or dynamic changes to services. AppScale is an open-source product that supports execution of Google Application Engine applications and therefore restricted to this particular technology. mOSAIC provides a set of APIs to application developers to tackle cloud deployment issues. The limitation of mOSAIC is the implementation of these APIs; the application developer needs to integrate these to application components.

Cloudify [3] combines a TOSCA editor with deployment and orchestrator. It provides access to multiple clouds and a complete framework to describe microservices and execute them either in Docker containers or on cloud metal. Cloudify also provides dynamic service upscaling and downscaling based on microservice dependent parameters, for example number of transactions, number of threads, etc. Cloudify does not provide a container portability framework, nor do its metrics span dockerised microservices and the cloud metal executing them.

Pham, Tchana et al. discuss the problem of distributed applications [4]. They focus on application orchestration that can span several different clouds. They recognize the need to deliver microservice-specific parameters, for example port numbers and IP addresses, to other microservices, and provide a description language framework to do this. However, this solution does not support service discovery tools and dynamic relocation of microservices.

Amazon CloudFormation [5] provides to system administration developers an easier way for the collection, creation and management of related AWS resources through templates. These templates describe the AWS resources and associated dependencies. After the deployment of AWS resources, CloudFormation ensures the start of services in the correct order.

OpenTosca [6] [7] provides an open source ecosystem for the OASIS Topology and Orchestration Specification for Cloud Applications developed by Stuttgart University. OpenTosca is divided into three parts: a TOSCA runtime environment (OpenTosca container), a graphical modelling TOSCA tool (Winery) and a self-service portal for the application available in the container (Vinothek).

Although effective in particular narrow circumstances, none of these systems provide a fully-automated, cloud-agnostic solution for all, or even a wide range of applications and for wide variety of clouds.

74 III. OPPORTUNITIES AND CHALLENGES OF MICROSERVICE-BASED ARCHITECTURE DESIGN IN CLOUDS

75 In order to support application-level orchestration and dynamic resource provision in clouds, a microservices-based architecture is
76 proposed. This section highlights the motivations behind this approach and collects challenges that the architecture design needs to
77 overcome.

78 Software design based around microservices, in contrast to a traditional monolithic architecture, is a relatively new concept, and the
79 literature does not yet have an agreed definition for microservices. According to Newman “Microservices are small, autonomous
80 services that work together”. His definition includes further characteristics, chiefly “focused on doing one thing well” and being
81 “autonomous” [8]. Balalaie et al write “Microservices is [*sic*] a new architectural style [9] that aims to realize software systems as a
82 package of small services, each deployable on a different platform, and running in its own process while communicating through
83 lightweight mechanisms like RESTFull APIs.”

84 Both Newman and Balalaie et al. see microservices mainly as an agile development concept. A small, focused team can develop one
85 component of an application independently and deploy new versions without changes to other components. In this paper, we focus on
86 deployment of microservices-based architecture, where connectivity between microservices becomes challenging.

87 Cloud computing is a natural platform for microservices. Microservices achieve decoupling of independent components from a
88 monolithic application. Clouds enable execution and resource allocation of these independent components based on their specific
89 needs. One microservice might require a lot of storage while another could be CPU-intensive. Cloud execution offers the possibility to
90 optimise resource allocation, and thus resource cost, dynamically. The alternative would be to allocate a monolithic infrastructure, the
91 size of which would be large enough for a worst-case requirements scenario. However, for most of the time, the worst-case scenario
92 does not prevail and allocated resources of the monolithic infrastructure are wasted.

93 As discussed earlier, microservices provide APIs to enable communication with them, and rely on APIs of other microservices to
94 access other services, on which they are depending. A typical API connection is a TCP (Transmission Control Protocol) socket. To set
95 up a TCP connection, the connecting computer needs two parameters: the IP address of the server host and a port number of the
96 particular service on that host. Many port numbers are defined either officially in RFCs (Request For Comments), or by convention,
97 and even in case of a completely new service, the port number would remain static within the particular architecture. The problem
98 arises from dynamic IP address allocation strategy employed by the cloud. To form an API connection, the client application needs to
99 know the dynamically assigned IP address of the server host.

00 In a microservices-based architecture commissioned on cloud, the IP address information may change during the life-cycle of the
01 application. In a microservices-based model, microservices can and will be developed independently. Moreover, cloud computing
02 provides dynamic allocation of hardware resources. From performance and financial optimisation, it would be occasionally necessary
03 to allocate more capacity by migrating a microservice to a more powerful platform, or to consolidate less regularly used microservices
04 to a single platform and shutting down some cloud capacity. By doing this, the IP address of the server host would change.

05 An additional challenge arises if multiple microservices providing the exact same service are deployed on the same host. If the service
06 API is provided in an established port, only one of the microservices could occupy this, and the others would need to be remapped.
07 This would normally require changes to a configuration file. For example, if a single cloud instance hosted three MySQL
08 implementations, only one of them could use the default port 3306. Running the other two would require configuration file changes,
09 including additional changes to configuration files and service start-up scripts to assign a different location to data and configuration
10 files.

11 One workaround would be to ensure that only one microservice of a kind could run on each cloud instance. This would remove the
12 need to reconfigure port numbers and file locations, but it might waste resources if new cloud instances would be needed just for this
13 purpose. Moreover, this would require building an additional coordination layer to keep track of running cloud instances and currently
14 running microservices, to be used to allocate new microservices to cloud instances where such microservices do not currently run. It
15 would also be entirely possible for two different services to occupy the same network port. While not common, there is no mechanism
16 or policy that would prevent this, especially if microservices came from different sources and their development were independent.
17 Thus, the aforementioned mapping would need to consider network services as well as file system locations to ensure there is no
18 overlap between two microservices designed to share the same cloud instance. In Section V we will show such steps are unnecessary.

19 Obtaining benefits from the dynamic nature of cloud requires an understanding of current and/or predicted resource usage, and the
20 possibility to scale infrastructure up or down on demand. A mechanism is needed to analyse usage, allocate microservices to cloud
21 instances best suited to serve them, and allow optimisation based on given parameters, such as performance or cost. Microservices can
22 set different requirements to physical cloud instances and resources based on the nature of the service. For example, an authentication
23 service would require little CPU power or storage, an HTTP application server would require considerable CPU power but little
24 storage, and a database server could require both powerful computation *and* storage capabilities.

25 Setting up a static microservices based architecture on cloud would be possible manually, by designing the correct launch order of
26 services, and configuring manually IP addresses of the already started services to those depending on them. If more capacity were
27 needed or any changes to the infrastructure were initiated by changes in cloud services provisioning, an operator would need to
28 reconfigure microservices API information manually. This could in many cases lead to suboptimal service performance or waste, as

29 from cost and workload management viewpoint it might be easier to react to changed requirements only when absolutely necessary.
30 Automated monitoring and control could thus enhance benefits enabled by the use of cloud and microservices.

31 The above detailed challenges can be summarized as follows:

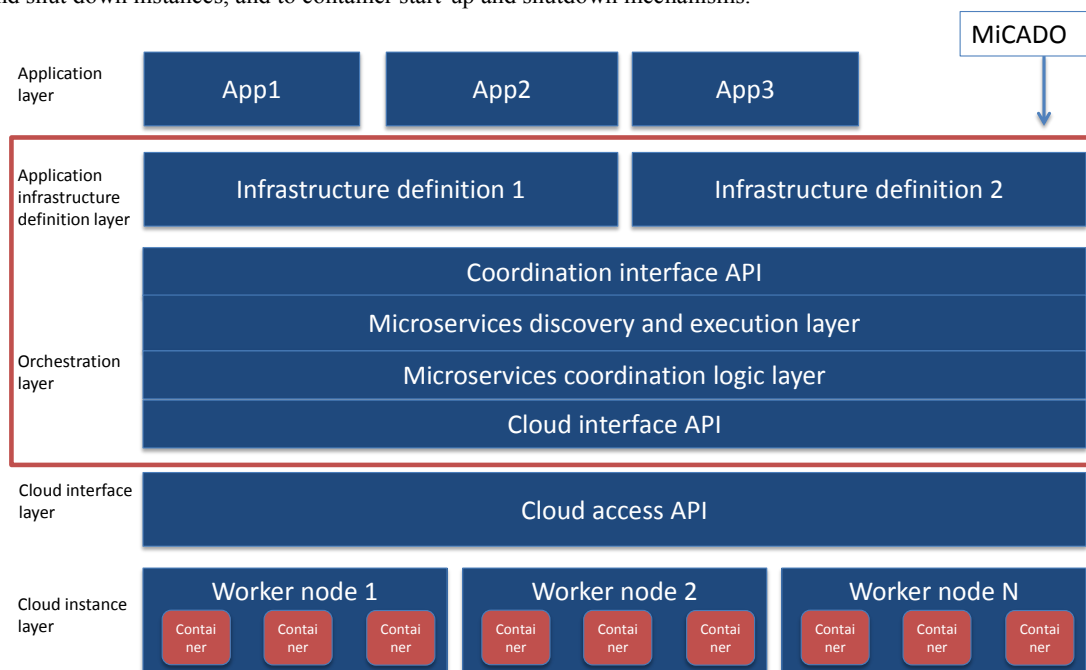
- 32 C1. Discovery of IP addresses and port numbers of running microservices.
- 33 C2. Ability to run several microservices of the same kind on a single cloud instance.
- 34 C3. Allocation of microservices to cloud instances based on available resources.
- 35 C4. Restriction mechanism to ensure microservices can only be allocated to cloud instances that have the functionality and
- 36 capabilities of serving them.
- 37 C5. Ability to dynamically scale up and down. If cloud infrastructure usage exceeds given parameters, the condition must be detected
- 38 and more cloud resources need to be allocated.

39 This work proposes an architecture that helps mitigate the aforementioned challenges. The approach is also agnostic to chosen
40 applications, cloud providers, and microservices needed to provide applications.

41 IV. MiCADO – THE PROPOSED SOLUTION

42 The proposed solution places microservices in lightweight virtualisation containers in worker nodes. These containers can be hosted
43 on any of the different worker nodes, and one node can run one or more containers. An orchestration and coordination mechanism is
44 required to enable service discovery and performance management. The overall MiCADO (Microservices-based Cloud Application
45 level Dynamic Orchestrator) architecture is illustrated in Figure 2. The solution takes the following overall approaches to challenges
46 identified in section III:

- 47 C1. To address dynamic discovery of IP addresses and ports, use a service discovery framework. Several such frameworks exist, for
48 example Consul [10], Zookeeper [11] and Etcd [12].
- 49 C2. To enable deployment of several microservices of the same kind on the same instance, use kernel namespace based lightweight
50 virtualisation solution to run microservices in containers. These solutions (for example Docker [13] built on Linux containers)
51 provide separation of application files and allow port mapping functionality to masquerade standard microservice API ports to
52 dynamically allocated ones.
- 53 C3. To support the allocation of microservices to cloud instances based on available resources, form a cluster of worker nodes that is
54 aware of the capacity and status of each node, and is able to make service allocation decisions based on a documented logic.
- 55 C4. The chosen clustering mechanism must be aware of the hardware configuration of cloud instances, and allow constraints to be set
56 to limit automatic allocation decisions (C3) to those nodes that have the requisite physical capabilities, for example a public IP
57 address, to serve a particular microservice.
- 58 C5. To support automatic scaling up and down, the cluster needs to have an alert mechanism, to help cluster management detect
59 nodes and services that exceed thresholds or nodes that appear to be underutilised. Decision-making logic needs to be
60 programmed based on this information or obtained within a software component. The logic needs to have interfaces to clouds to
61 start and shut down instances, and to container start-up and shutdown mechanisms.



62

63

Fig 2: Generic MiCADO architecture

64 Based on the above described generic approaches to handle the five identified challenges, the following layered architecture is
65 proposed for the application level orchestration of cloud applications. The layers in the box entitled MiCADO represent the actual
66 orchestration architecture. Layers below MiCADO provide access to cloud resources, while the top layer represents actual
67 applications to be optimized. The layers in Figure 2 are described from bottom to top.

- 68 1) **Cloud infrastructure layer.** This layer contains cloud instances, which in turn run containers that execute actual microservices.
69 One instance can run one or more containers.
- 70 2) **Cloud interface layer.** This is a set of APIs that provides means to launch and shut down cloud instances. There can be one or
71 more cloud interfaces to support multiple clouds. Either native interfaces of targeted clouds can be applied (e.g. EC2) or generic
72 cloud access layers that provide access to multiple heterogeneous clouds.
- 73 3) **Microservices orchestration layer.** This is divided into four sub-layers.
 - 74 a) *Cloud interface API.* This layer is needed to abstract cloud access from the layers above. Cloud access APIs can be complex
75 interfaces, as they typically cater for a large number of services provided by the cloud provider. On the other hand, the
76 microservices execution and coordination logic layers (see 3b and 3c) only need to shut down and start instances. Abstracting
77 this to a cloud interface API simplifies the implementation of the aforementioned layers, and equally, if new Cloud access
78 APIs are implemented, only this layer needs to change.
 - 79 b) *Microservices coordination logic layer.* With large infrastructures, and to reap the benefits from cloud-based execution, it
80 becomes necessary to understand how the current execution environment is performing. Information needs to be gathered
81 and processed. If bottlenecks are detected or the currently running infrastructure appears underutilised, it may be necessary to
82 either launch or shut down cloud instances, and possibly move microservices from one physical worker node to another.
 - 83 c) *Microservices discovery and execution layer.* This layer manages the execution of microservices and keeps track of services
84 running. Execution management combines both start-up and shut-down of microservices. Service management gathers
85 information about currently running services. Information gathered includes service name, IP address and port where the
86 service is reachable, and optional service tags to help in service coordination.
 - 87 d) *Coordination interface API.* This layer provides access to orchestration control and decouples the orchestration layer from
88 the application infrastructure definition. This set of APIs enables application developers to utilise the dynamic orchestration
89 capabilities of the underlying layer and support the convenient development of dynamically and automatically scalable
90 cloud-based applications.
- 91 4) **Application infrastructure definition layer.** This forms the basis for creating a functional application infrastructure. At this
92 level, software components and their requirements as well as their interconnectivity are defined. This layer does not contain any
93 application-specific data. For example, to provide HTTP based services, this layer can define that to provide this functionality, a
94 MySQL database, Apache HTTP server and Nginx proxy server are needed, and that Nginx needs connection to Apache, which
95 in turn needs connection to MySQL. As the infrastructure is agnostic to the actual application using it, this definition can be
96 shared with any application that requires such an environment.
- 97 5) **Application layer.** This layer contains actual application code and data to make an incarnation of a defined application
98 infrastructure (4) function in such a way that the desired functionality is achieved. For example, this layer could populate a
99 database with initial data, and configure an HTTP server with both look and feel and application logic.

100 V. MiCADO REFERENCE IMPLEMENTATION

101 A prototype of the proposed MiCADO architecture has been implemented via a combination of open-source tools and in-house
102 developed extensions, as is illustrated in Figure 3. Although it would be possible to develop a platform to implement MiCADO
103 functionality from scratch, reusing and utilizing existing open-source components has significantly speeded up the implementation
104 process. As a restriction, this first MiCADO prototype does not include Cloud interface and Coordination interface APIs. The
105 Coordination interface API is replaced by a simple command line, and the Cloud interface API is represented by direct calls towards
106 the cloud access layer (in this case the CloudBroker Platform [14] [15]). Moreover, instead of a generic approach towards
107 infrastructure definition, a concrete architecture has been realized implementing a real-life case-study. Applying these limitations, a
108 proof-of-concept implementation has been completed to justify the validity of the proposed approach when handling the five
109 challenges detailed in section III. For infrastructure management and orchestration the following tools have been selected to
110 implement the previously described MiCADO layers.

111 **Microservices discovery layer** was implemented based on Consul [10]. Consul is an open-source service discovery tool that also
112 includes health check and alerts functionality. An additional component, Registrator [16] is used on worker nodes to register
113 information on running containers to Consul.

114 Consul agents can work either in a server or serf role [17]. Servers keep track of events happening in the network and notify other
115 servers of changes, while serfs create a server connection and maintain a list of backup servers. Consul servers create network
116 overhead by communicating between themselves using Consensus protocol [18]. In large environments, most nodes should be
117 configured as serfs to minimise this traffic. In the proof of concept example, the master and all nodes are configured as servers, to
118 keep node configurations identical.

119 When a Consul agent starts, it needs to connect to an existing network. This can be provided either by a specific Consul server set in
120 *bootstrap mode*, by connection to a node or set of nodes already running Consul, or by connection to Atlas service [19]. Atlas is a
121 service that registers running Consul infrastructures based on a secret token. When a Consul agent configured to use Atlas starts, it

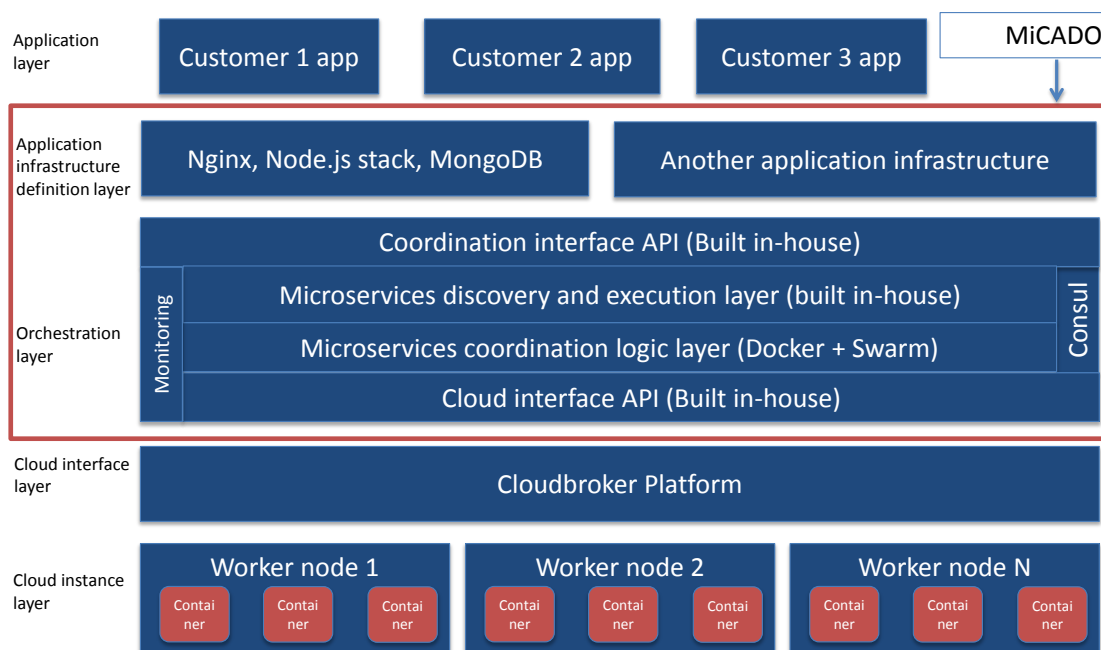
22 registers itself with the service and queries for other nodes already present. The service is provided via a REST API, making it easy to
 23 use from firewalled networks. The first node registering does not receive anything. When the second node joins the infrastructure, it
 24 registers itself and receives in response the details of the first node, allowing it to initialise connection to it. When a server finds one
 25 Consul server, it communicates with it and shares information about the entire network. Thus, a Consul network does not need to be
 26 described in nodes, only the bootstrap information needs to be present.

27 Registrator is started on all worker nodes. Registrator is a helper process to Consul that registers all running Docker containers to it.
 28 Registrator runs in a Docker container itself, and it can be started either locally or via the Swarm launch mechanism. The roles of
 29 Docker and Swarm will be explained in detail under the Microservices coordination logic layer.

30 **Microservices coordination logic layer** was implemented based on Docker and Swarm [20]. Swarm is a clustering mechanism built
 31 on Docker that is aware of worker nodes and their current workload, and is able to allocate new containers to the node currently least
 32 used. Docker and Swarm allow the capability of giving worker nodes “tags” to enable the use of constraints. For example, giving a
 33 different tag to nodes with a lot of disk space allows allocation of database containers to these particular nodes, while worker nodes
 34 tagged for proxy use will have a public IP address. To fulfil requirement C4, worker nodes are tagged with their roles. The tags used
 35 in the current application scenario are *proxy*, *application* and *database* (see Application infrastructure definition layer later on). This
 36 tag is set in the Docker start-up configuration file.

37 **Cloud Access API layer** is represented by calls to the Cloudbroker Platform. The Cloudbroker Platform is a production level tool
 38 developed by a CloudBroker GmbH that allows interfacing to various commercial IaaS clouds, for example Amazon and
 39 CloudSigma, and also to private cloud infrastructures based on Openstack, OpenNebula and Eucalyptus [21]. An installation of the
 40 platform that is operated by the CloudSME European Project [22] was utilized in our experiments. Using the CloudBroker Platform,
 41 even this very first MiCADO implementation is capable of interfacing with a large variety of clouds without further cloud interface
 42 plug-in development.

43 For **cloud instances**, an Ubuntu 14.04 image was prepared with Docker, Consul, Swarm and Registrator installed. The CloudBroker
 44 platform provides a mechanism for preparing such images to be later launched in participating clouds. In addition to pre-installed
 45 applications, the cloud image needs a start-up script to initialise and launch the above mentioned tools. Consul and Docker create,
 46 upon initial start, random identifiers to distinguish between nodes in the same cluster. When a cloud instance launches the first time, it
 47 needs to remove these identifiers to force initialisation.

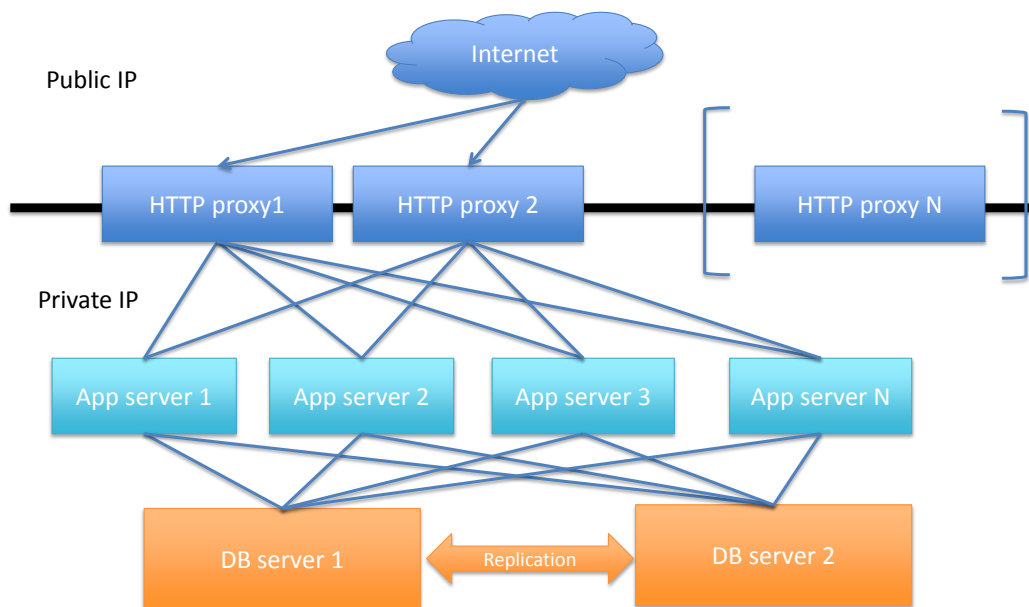


48
49 **Fig. 3:** MiCADO reference implementation

50 **Application infrastructure definition layer** currently implements one particular topology based on a specific set of use-cases
 51 provided by Outlandish LLP [24]. Outlandish is a small employee-owned digital agency specialising in middleware, usability, search
 52 and scalable data applications. Outlandish develops and hosts various web applications for multiple corporate clients. Usage estimates
 53 vary greatly between applications and customers, and in some cases, usage activity can be expected to vary greatly based on external
 54 triggers, such as calendar, time of day, outside events and news. The company wishes to serve its clients to agreed quality of service
 55 but also tries to minimise its costs on the provision side. Therefore, a cloud-based application deployment and hosting environment
 56 that is capable of making such optimisation choices dynamically at run-time is highly desired.

57 A typical Outlandish application is formed of three components: proxy servers with public IP address, application servers and
 58 database servers, as illustrated in Figure 4. The concrete implementation applies Nginx reverse proxy servers, Node.js stacks for
 59 application servers, and MongoDB databases. This set-up is generic for several Outlandish-developed and -hosted applications and as
 60 a consequence, several application scenarios and clients can be served based on this application pattern.

61 The currently implemented MiCADO prototype successfully deploys the above described generic architecture for several Outlandish
 62 applications into Docker containers, clusters these containers based on Swarm, and provides the required service discovery
 63 mechanism via Consul and Registrator. Outlandish application developers need only provide the application code and the
 64 corresponding Docker configuration files. The optimised deployment to available cloud instances is managed by MiCADO.
 65 Therefore, the implemented solution successfully addresses challenges C1-C3 (IP address and port number discovery, running same
 66 kinds of microservices on a cloud instance, and allocating microservices to cloud instances based on available resources), and also by
 67 tagging the resources it provides mechanisms to support C4 (allocate microservices to specific instances that have the capabilities to
 68 serve these). However, please note that this current prototype does not offer run-time monitoring and automated optimisation
 69 capabilities (C5) that is part of our future work.



70
 71 **Fig. 4:** Web-based application infrastructure example implemented for the Outlandish use-cases

72 VI. DEPLOYING AND MANAGING APPLICATIONS WITH MiCADO

73 Based on the above described first proof-of-concept reference implementation of the MiCADO framework, a short overview on how
 74 this implementation supports the deployment and management of cloud applications is provided in this section. Please note that the
 75 current implementation supports only command-line access. However, the objective of our future research is to develop a set of well-
 76 defined APIs that can be conveniently embedded into desktop or web-based science gateway solutions.

77 Before an application developer can utilize MiCADO services, certain infrastructure components need to be deployed and set-up by
 78 an administrator. A specific set of scripts has been developed [25] that can be used to create and deploy the necessary virtual machine
 79 images on the CloudBroker Platform. MiCADO differentiates between two types of virtual machines: manager and worker nodes. The
 80 manager node runs permanently until the infrastructure is required and can be used to deploy or terminate microservices running in
 81 containers, or cloud instances. Worker nodes are hosting various Docker containers running different microservices (e.g. applications,
 82 databases or proxies, as illustrated in figure 4). Both manager and worker node images include Docker, Consul and Swarm.

83 In order to create a new MiCADO installation, the administrator needs to launch a manager node instance through the CloudBroker
 84 platform in the target cloud. Please note that as the CB platform is used at the cloud access layer, the solution is cloud-agnostic, and
 85 the manager node can be launched in any targeted cloud in which an image has been created. Once the manager node is launched,
 86 application developers can log-in to this instance to manage their applications formed of multiple microservices.

87 Application developers first need to create a Docker image of the specific microservice they want to deploy (for this process no
 88 specific support is provided currently in MiCADO). This Docker image is then uploaded to a private Docker registry operated by the
 89 CloudSME project. In order to launch a new Docker container with a required image, application developers need to execute a Python
 90 script from [25] providing the name of the image and the type of node (e.g. node with public IP address for a Proxy server) as
 91 parameters. The container will be launched to the least utilised node automatically by Swarm. When a new container is launched,
 92 MiCADO checks whether there is enough memory available on current cloud instances to run the container. If there is not enough

93 memory available, then a new instance is launched automatically, and the launch of the container will be held until the new cloud
94 instance is up and running on the infrastructure.

95 The current MiCADO prototype also supports automated downscaling of the infrastructure at run-time. There is a cron job running
96 continuously to determine if the infrastructure needs to be downscaled. This cron job will first establish which instance on the
97 infrastructure is the least utilized. It will then list all the running containers on the least utilized instance and get their memory usage.
98 The cron job will check if the memory usage on the instance is less than the availability on the rest of the infrastructure. If there is
99 enough memory available on the other running instances then it will check if each container can be relocated on another instance. If
00 the containers can be relocated then these are launched on the other instances, the containers on the least utilized instance are stopped
01 and removed, and finally the least utilized instance is stopped and removed from the infrastructure. With this functionality MiCADO
02 provides some basic capabilities to optimize resource utilization and therefore save unnecessary expenses for application providers.

03 VII. CONCLUSIONS AND FUTURE WORK

04 This paper collected and analysed challenges towards a microservices-based implementation of a dynamic application level cloud
05 orchestrator called MiCADO. Moreover, it presented a first proof-of-concept implementation of the architecture that successfully
06 addresses four out of the five identified challenges, and also facilitates a future implementation to tackle the fifth challenge. The
07 MiCADO concept succeeds in decoupling application level logic from cloud orchestration logic, while preserving the elastic and
08 dynamic nature of the cloud. The experiment shows it is possible to narrow the gap between IaaS and SaaS models.

09 While the solution was tested on the infrastructure described earlier, there is no fundamental reason why the principles could not be
10 applied to other operating systems, monitoring tools, service discovery tools and cloud APIs. The most significant limitations are in
11 security, performance measurement and downscaling. No work has been done on infrastructure level security. Network overhead
12 incurred by Consul and Swarm has not been investigated, nor has the CPU overhead caused by Docker. Consul has a built-in
13 mechanism to reduce network overhead by designating only a few nodes as servers and the rest as serfs. Therefore, more work would
14 be needed to define the optimal configuration. Moreover, performance measurement has been done on a most basic level, to prove a
15 concept instead of providing a ready solution.

16 As this paper only analysed the MiCADO concept and presented a proof of concept implementation, future work is rather versatile.
17 On the one hand it should evolve around measurement and optimization logic of dynamic scalability. As this is the core of the
18 MiCADO platform, making this more versatile and intelligent without sacrificing usability is a challenge. Another aspect includes
19 infrastructure definition reusability and sharing. This could lead to a definition of application topology and description templates that
20 are applicable for a wide range of application scenarios. For such purposes we are looking at existing standardized description
21 formats, such as TOSCA (Topology and Orchestration Specification for Cloud Applications) by OASIS [23]. Another angle for
22 future work is information security related to service up/downscaling and portability, as this has not been touched at all.

23 VIII. ACKNOWLEDGEMENTS

24 This work was funded by the CloudSME Cloud-Based Simulation platform for Manufacturing and Engineering Project No. 608886
25 (FP7-2013-NMP-ICT-FOF). We also acknowledge the contribution of Outlandish LLP whose staff provided us with vital information
26 regarding their requirements and valuable feedback throughout the implementation.

27 REFERENCES

- 28 [1] Thomas, D. (2009). Cloud Computing — Benefits and Challenges! *The Journal of Object Technology*, 8(3), 37.
29 [2] Marpaung, Sain, & Hoon-Jae Lee. (2013). Survey on middleware systems in cloud computing integration. *Advanced Communication Technology (ICACT), 2013*
30 *15th International Conference on Advanced Communications Technology*, 709-712.
31 [3] Anon. (2016). *Cloudify* [online] Available from: <<http://getcloudify.org>> [Accessed 16/03/2016]
32 [4] Manh Pham, L., Tchana, A., Donsez, D., Zurczak, V., Gibello, P., & De Palma, N. (2015). An adaptable framework to deploy complex applications onto multi-
33 cloud platforms. *The Institute of Electrical and Electronics Engineers, Inc. (IEEE) Conference Proceedings*, 169-174.
34 [5] Anon. (2016). *AWS CloudFormation* [online] Available from: <<https://aws.amazon.com/cloudformation/>> [Accessed 16/03/2016]
35 [6] Anon. (2015). *OpenTosca* [online] Available from: <<http://www.iaas.uni-stuttgart.de/OpenTOSCA/#ecosystem>> [Accessed 16/03/2016]
36 [7] Binz, Breiter, Leyman, & Spatzier. (2012). Portable Cloud Services Using TOSCA. *Internet Computing, IEEE*, 16(3), 80-85.
37 [8] Newman, Sam. (2015). *Building Microservices*. O'Reilly.
38 [9] Balalaie, A., Heydamoori, A., & Jamshidi, P. (2015). Migrating to Cloud-Native Architectures Using Microservices: An Experience Report.
39 HashiCorp. (2016). *CONSUL by HashiCorp* [online] Available from: <<https://www.consul.io>> [Accessed 16/03/2016]
40 [11] Apache. (2016). *Apache Zookeeper* [online] Available from: <<https://zookeeper.apache.org>> [Accessed 16/03/2016]
41 [12] Coreos. (2016). *Etd* [online] Available from: <<https://coreos.com/etcd/>> [Accessed 16/03/2016]
42 [13] Anon. (2016). *Docker* [online] Available from <<https://www.docker.com>> [Accessed 16/03/2016]
43 [14] Anon. (2016). *CloudBroker Platform* [online] Available from: <<http://cloudbroker.com/platform>> [Accessed 15/03/2016]
44 [15] Taylor SJE, Kiss T, Terstyanszky G, Kacsuk P, Fantini N. (2014). Cloud Computing for Simulation in Manufacturing and Engineering: Introducing the
45 CloudSME Simulation Platform. *Proceedings of the 2014 Annual Simulation Symposium, ANSS '14, Society for Computer Simulation International: San Diego,*
46 *CA, USA, 2014*; 12:1–12:8.
47 [16] Anon. (2016). *Registrar* [online] Available from: <<http://gliderlabs.com/registrator/latest/>> [Accessed 16/03/2016]
48 [17] Anon. (2016). *Consul Architecture* [online] Available from: <<https://www.consul.io/docs/internals/architecture.html>> [Accessed 16/03/2016]
49 [18] Anon. (2016). *Consensus Protocol* [online] Available from: <<https://www.consul.io/docs/internals/consensus.html>> [Accessed 16/03/2016]
50 [19] Anon. (2016). *Atlas Integration* [online] Available from: <<https://www.consul.io/docs/guides/atlas.html>> [Accessed 16/03/2016]
51 [20] Anon. (2016). *Docker Swarm Overview* [online] Available from: <<https://docs.docker.com/swarm/overview/>> [Accessed 16/03/2016]
52 [21] Baset, Salman, A., Carey, Michael, & Hand, Steven. (2012). Open source cloud technologies. *Cloud Computing Proceedings of the Third ACM Symposium*, 1-2.

- 53 [22] Kiss, T, Dagdeviren, H, Taylor, JES, Anagnostou, A, Fantini, N. (2015). Business Models for Cloud Computing: Experiences from Developing Modeling &
54 Simulation as a Service Applications in Industry, *Proceedings of the 2015 Winter Simulation Conference*, Pages 2656-2667, ISBN: 978-1-4673-9741-4, IEEE
55 Press
- 56 [23] Oasis committee specification (2013). *Topology and Specification for Cloud Applications version 1.0* [online] Available from: <[http://docs.oasis-](http://docs.oasis-open.org/tosca/TOSCA/v1.0/os/TOSCA-v1.0-os.html)
57 [open.org/tosca/TOSCA/v1.0/os/TOSCA-v1.0-os.html](http://docs.oasis-open.org/tosca/TOSCA/v1.0/os/TOSCA-v1.0-os.html)> [Accessed 15/03/2016]
- 58 [24] Anon. (2016). Outlandish [online] Available from <<http://outlandish.com/>> [Accessed 19/03/2016]
- 59 [25] Anon. (2016). *MiCADO setup scripts* [online] Available from <<https://github.com/gregGes/DSCR-infrastructure>> [Accessed 08/05/2016]