5th International DMTF Academic Alliance Workshop on Systems and Virtualization Management: Standards and the Cloud (SVM 2011)

Poster Session Abstracts

Mearns, Leaney, Parakhine, Verchere, and Debenham [1]

Cloud computing continues to expand in both utility and coverage with growing public cloud infrastructures and hybrid service clouds offering a diverse range of services. Along with the cloud services are the options associated with the telecommunication service. These combined management issues have introduced greater complexity in the provisioning and performance management of cloud services. Choosing and associating all these services, including provisioning and managing their performance is currently left to the user. The user is often locked into one provider and their services. This single provider environment puts a limit on the scalability of any global applications. Researchers in both the cloud computing and telecommunications are moving towards a dynamic market based service structure, where services can be purchased from the various providers at need.

We propose extending the management strategies for this complex environment. Firstly, we aim to cover the totality of modern complex services, managing both the connectivity and virtual infrastructure. Secondly, we aim to accept responsibility for the complex service in the open marketplace, where responsibility is composed primarily of the judgement of risk and the management of resilience. We are working towards a bundled service provider agent architecture, which can negotiate on a combined cloud / telecommunication open service market, and which will eventually help to optimise the utilisation of the providers infrastructure while reducing the risk of failure to users through total service management.

An advantage of this approach, is that while the bundled service agent remains active, there is no need for multiply contracted providers to be utilised until needed. Another advantage of this approach is the ability to isolate the problems and take appropriate action, resolving the problems in the appropriate part of the service, rather than trying to re-provision the whole complex service. Further the bundled service provider would be able to manage the problems arising in dependent services as well. We present the specification, design and simulation of the bundled service provider and its agents in a marketplace environment.

Textor [2]

I. MOTIVATION

In the domain of IT management, numerous models, pro- tocols and tools have been developed. Notable models in- clude the OSI network management model (CMIS/CMIP) and the still widely used simple network management protocol (SNMP). A more recent approach to specify a comprehensive IT management model is the Common Information Model (CIM [1]), a widely recognized Distributed Management Task Force (DMTF) standard. The more complex an environment gets, the higher the need for comprehensive, highly automated IT management. To achieve this long-term goal, the various available sources of information need to be combined. How- ever, syntactic translations from one model to another are often not sufficient, as the same concept can be expressed in different ways in two different domains [2]. For this reason, several researchers have proposed to use ontologies to unambiguously and comprehensively model IT environments. An ontology is a formal representation of a set of concepts within a domain and the relationships between those concepts. It can also contain behaviour rules and instance data that represents concrete entities. Thus, ontologies are a possible formal base for automated IT management.

II. RELATED WORK

To allow the application of ontologies to IT management, a suitable domain model is required. In [3], CIM has been proposed for this purpose, as it is an actively used, maintained and freely available standard. As CIM is also the foundation for other models, such as the storage-related standard SMI- S (Storage Management Initiative Specification) of the SNIA (Storage Networking Industry Association, [4]) and is used in the DMTF SMASH specification (Systems Management Architecture for Server Hardware, [5]), the approach presented here is also applicable in environments where one of these is employed.

The authors in [6] point out that CIM is usable for inferring properties about distributed systems, but is a semi-formal ontology with limited support for knowledge interoperability and aggregation, as well as reasoning. This firstly means that support in CIM for the specification of formal rules is very limited. Secondly, the interrelation

of the specified information with knowledge from other domains (which are usually not specified in CIM) is not easily possible in a way that supports
automatic management, e.g. business processes that refer to physical or logical IT systems. For the approach presented here, the first point is of primary importance, while in the long run, the connection to information models from different domains (or different views of the environment in a business) will become more relevant.

III. APPROACH

In this work, we construct an IT-management domain ontology in the Web Ontology Language (OWL, [7]) by transforming the CIM schema into OWL. This way, the rich and extendable model of the CIM schema can be used, while simultaneously a formal knowledge base is created that not only contains the domain model, but can also express rules and incorporate instance data, i.e. all the real entities of the managed system. The ontology can be enriched with rules formulated in the Semantic Web Rule Language (SWRL, [8]). SWRL rules can be directly embedded in OWL ontologies. That way, a comprehensive and formal description is formed that can be used as the basis for automated management. Details about the transformation of CIM to OWL are described in [9].
Based on the ontology that contains the domain model and rules, a management system is constructed. This management system loads the CIM ontology and then starts to read runtime information from the managed system by querying a CIM Ob- ject Monitor (CIMOM). Returned information is transformed into OWL instance data which matches the ontology entities of the previously loaded CIM ontology. As the ontology combines model and instance data, the SWRL rules can refer to both structural and instance data. A semantic reasoner component evaluates the ontology and the rules contained within, and can trigger reactions based on the evaluation results, thus creating a feedback loop to the managed system.

IV. STATUS

A proof of concept was implemented that manages a Linux file system. File system information is acquired via the OpenPegasus CIMOM, and rules are evaluated that monitor file sizes. Rules include checks if sizes of certain files or directories excess a given threshold and are older than a given reference file. If that is the case, actions can be taken, such as moving or compressing folders.
In addition to the evaluation of rules at runtime, the ontology can be queried for runtime information using the SPARQL Query language [10]. In contrast to the CIM Query Language, this also allows querying facts that were reasoned by a semantic reasoner, and which were not explicitly given in the original model, for example, enumerate all files that are older than a given reference file.
Both the translation tool for the CIM to OWL translation and the runtime management system were implemented, and performance tests were carried out. The CIM schema with approximately 1400 classes is converted to about 100,000 OWL axioms. The as of yet unoptimized runtime system takes between 8 and 80 seconds on an Intel Core 2 Duo with 2 GHz and 2 GB RAM for one reasoning cycle, when the ontology contains the full CIM schema and up to 100,000 instances.
The translation tool for the CIM to OWL translation was not only used on the CIM schema, but also on the SMI-S 1.5.0 schema files, as well as the VMware CIM SMASH/Server Management API for ESX, which are both also specified in the CIM format.
Based on these results, next steps include performance optimizations and the application of the management system to automated storage management and/or automated virtual machine management together with a partner.

REFERENCES

[1] Distributed Management Task Force, “Common Information Model (CIM),” http://www.dmtf.org/standards/cim/.
[2] J. E. L. De Vergara, V. A. Villagra ́, and J. Berrocal, “Semantic Management: advantages of using an ontology-based management information meta-model,” in Proceedings of the HP OpenView University Association Ninth Plenary Workshop, 2002. [Online]. Available: http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.18. 7462
[3] J. E. L. De Vergara, V. A. Villagra ́, J. I. Asensio, and J. Berrocal, “Ontologies: Giving Semantics to Network Management Models,” IEEE Network, vol. 17, no. May/June, 2003.
[4] Storage Networking Industry Association, http://www.snia.org/.
[5] Distributed Management Task Force, “Systems Management Architec-
ture for Server Hardware (SMASH),” http://dmtf.org/standards/smash.
[6] S. Quirolgico, P. Assis, A. Westerinen, M. Baskey, and E. Stokes, “Toward a Formal Common Information Model Ontology,” pp. 11–21, 2004. [Online]. Available: http://www.springerlink.com/content/
a48nkmh3c4037xhe
[7] World Wide Web Consortium, “Web Ontology Language (OWL),” http:
//www.w3.org/2004/OWL/.
[8] ——, “Semantic Web Rule Language (SWRL),” http://www.w3.org/
Submission/SWRL/.
[9] A. Textor, J. Stynes, and R. Kroeger, “Transformation of the Common
Information Model to OWL,” in 10th International Conference on Web Engineering - ICWE 2010 Workshops, ser. Lecture Notes in Computer Science, vol. 6385. Springer Verlag, July 2010, pp. 163–174.
[10] World Wide Web Consortium, “SPARQL Query Language for RDF,” http://www.w3.org/TR/rdf- sparql- query/.

Dawoud and Takouna [3]

The rapid growth of E-Business and the frequent changes in sites contents pose the need for rapid and dynamic scaling of resources. Currently, cloud computing infrastructure enables agile and dynamic scalability of resources. However, current implementation of clouds' scalability is based on the Virtual Machine (coarse-grained) as a scaling unit which often leads to over-provisioning of resources. Hence, we propose Elastic VM (fine-grained) scaling architecture. It implements the scalability into the VM resources level. So, instead of scaling-out by running more VMs instances, our architecture scales-up the VM's resources themselves (e.g., number of cores and memory size) to cope with the workload demand. The theoretical comparison between the Elastic VM and the current Multi-instances scaling architectures (e.g., Amazon EC2 and GoGrid) shows that the Elastic VM scaling architecture is more able to maintain QoS metrics.Moreover, for an empirical comparison between the current scalability in the cloud (i.e., Multi-instances) and the Elastic VM scaling architectures, we locally implemented EC2 Amazon architectures that consist of Amazon Elastic Load Balancing and Amazon Auto Scaling. We considered the same parameters to be defined by both architectures. Our prototype is implemented into Xen virtualization environment.

The Elastic VM architecture implementation implies modifying the VM's kernel to accept on-the-fly adjustment of the VM resources without restarting the operating system or interrupting the service. Furthermore, the hypervisor is extended with interfaces allow changing resources automatically by our resources provisioning controllers. Experimental results show that the Elastic VM scaling architecture is able to reduce scaling overhead, maintain a high throughput, mitigate Service Level Objectives (SLOs) violation, and enable a simple scalability to broader range of applications including databases.

Carlson [4]

A new computing paradigm is quickly emerging called Cloud Computing. The actor enablers of this model have embraced the need to provide interoperability between enterprise computing and cloud service providers.

Virtualization technology and the evolution from discrete software packages to complete systems that can be created and deployed as a collection of virtual images is becoming the primary focus for delivering and managing software solutions into enterprise customers today. As these customers look to also take advantage of cloud computing, standards are needed to enable interactions between private clouds within enterprises, and between private and public cloud providers, and between separate public cloud providers to exploit this emerging business model.

This poster will show how the DMTF Cloud Management Working Group has focused on addressing the management interfaces between the cloud service consumer / developer and the cloud service provider. The poster outlines both the model and protocol bindings of a new DMTF standard for interoperable Infrastructure as a Service (IaaS) management. The poster shows how the standard addresses common use cases for creating cloud resources and managing the lifecycle and operation of virtual machines and their associated storage and networking resources.

The protocol binding for this version of the standard includes the RESTful use of HTTP. Standard HTTP verbs such as POST, GET, PUT and DELETE are used with standard mime types for each of the resources in the model. Protocol security is addressed by referencing existing internet security standards. Both JSON and XML renderings of the message body are shown, and examples are used to illustrate typical usage.

Chillarón and López De Vergara [5]

Current trend in IT services is that users, companies and large corporations access computing resources through the network in what is named Cloud Computing, enabling a better utilization of hardware and reducing costs. In this type of infrastructures, it is common to deploy virtual machines to deal with the users' requirements. Then, it is also important to provide the necessary computing power when needed, allocating the load of each virtual machine in a server that can handle it. However, this can be a challenging task, because this load is very variable, and many virtual machines are sharing the same host.

To solve this problem, we have designed and implemented an autonomic system to manage these virtual machines in the following way: all servers monitor their performance as a whole and also taking into account each virtual machine running on them. This information is shared among all the servers in the system. When one server is overloaded, it looks for the best servers in the system where to offload those virtual machines, trying to relieve this problem. For this, each host predicts what is going to be the load of the servers once the virtual machines have been migrated, choosing the right server and the right virtual machine to migrate at each time. The system follows the Monitor-Analysis-Plan-Execute loop, fulfilling the self-configuration, self-healing, self-optimization, and self-protection rules. In this poster we present the architecture of the developed system, and provide the results of the tests we have done, showing the autonomic behavior that regulates the servers load.

Fernández, Cordero Ordoñez, Somavilla, Rodriguez, Corchero, Tarrafeta and Galan [6]

The poster proposed here is related with the article "Distributed Virtual Scenarios over Multi-host Linux Environments" already accepted for the SVM 2011 workshop. The poster will mainly present the characteristics of Virtual Networks over LinuX (VNX), a tool designed for the deployment of virtual network scenarios over clusters of LinuX machines.

The poster will mention the general utility of virtual network scenarios deployment tools, quoting other tools or initiatives similar to VNX. Later, it will present the main characteristic and design decisions taken in VNX, the improvements achieved over the previous version of the tool (VNUML) and the technologies, virtualization platforms and standards used. Special attention will be paid to the virtual machine autoconfiguration mechanisms used in VNX, based on the ideas of dynamically created cdroms offered to virtual machines that were proposed in OVF.

Moreover, some of the experiments made with VNX tool to create complex network scenarios in the context of academic network laboratories and research projects will be shown. In particular, the complex experimentation scenario presented in the paper, which resembles an ISP which offers VPN services based on MPLS to two companies made of 6 different sites, will be shown. That scenario, made of 44 virtual machines deployed over two test environments (a multiprocessor server and a cluster of 6 personal computers) has been used to test the validity of the tool and to get some interesting results.

Finally, the poster could be complemented with simple demos of the VNX tool that could be made with one or two notebooks. The only requirement for that demo would be to have a small table near the place the poster will be located.

Lamers [7]

The System Virtualization, Partitioning and Clustering (SVPC) Work Group is developing DMTF standards for the management and packaging of virtual systems. This includes the deployment, discovery, configuration, resource allocation and active management of virtual computer systems. The work group is also developing specifications for the packaging and distribution of virtual appliances composed of one or more virtual computer systems. The first generation work has been published. In the second generation work currently under way enables improved scalability and the capability of handling data center level collections of virtual computer systems.

Hlavacs and Treutner [8]

The ability to live migrate virtual machines (VMs) between physical servers without any perceivable service interruption is pivotal for building more energy efficient Cloud Computing infrastructures in the future. Nevertheless, energy efficiency is not worth the effort if quality metrics (e.g., QoS, QoE) are severely decreased by, e.g., dynamic consolidation using live migration. We identify the most significant utilization metrics to predict the service level during live migrations for a web server scenario. We show important correlations, give reasons and draw conclusions for systems using live migration for yielding higher energy efficiency. We also give reasons for extending the current hypervisors' capabilities regarding VM utilization collection and reporting.

We present the effects of live migration on service levels for different workload scenarios. In particular, we demonstrate that live migration should be done preventively. This anticipates disproportional high service level degradation due to live migration. We examine the most important utilization metrics for predicting the service level by both stepwise and exhaustive regression. As a result, we can explain 90% of the service level variance during live migration with a single variable, using more variables yields 95%.

Shah and Slaight [9]

Platform Management Components Intercommunications (PMCI) is a Working Group (WG) within

DMTF. This poster will provide an overview of PMCI standards to address “inside the box” communication and functional interfaces between the components of the platform management subsystem such as management controllers, BIOS, and intelligent management devices. This poster will cover PMCI scope, work areas, published specifications, and the most recent specification development efforts including the following important technologies: Management Component Transport Protocol (MCTP), Platform Level Data Model (PLDM), and Network Controller Sideband Interface (NC-SI).

MCTP is a physical medium independent transport protocol designed to support communications between different platform management subsystem components. PLDM is a data model targeted for efficient communications among platform management subsystem components for accessing low-level platform inventory, monitoring, control, event generation and logging, and data/parameters transfer functions including SMBIOS data transfer and BIOS control/configuration. NC-SI is a sideband communication interface to enable the exchange of management data between the Management Controller and Network Controller.

Hilland [10]

This poster will provide a technical overview of SMASH & DASH management initiatives, including their differences and their similarities. The poster will cover implementation requirements, including the protocol requirements for the SM CLP and WS-Management and a brief discussion of the Management profiles for SMASH/DASH.

References

[1] H. Mearns, J. Leaney, A. Parakhine, D. Verchere, and J. Debenham. “Comprehensive Cloud Management via an Open Marketplace”. 2011 5th International DMTF Academic Alliance Workshop on Systems and Virtualization Management (SVM 2011).

[2] A. Textor. “A CIM-based Ontology for Semantic IT-Management”. 2011 5th International DMTF Academic Alliance Workshop on Systems and Virtualization Management (SVM 2011).

[3] W. Dawoud and I. Takouna. “Elastic VM: A Fine-grained Scalability Architecture for Virtualized Environments”. 2011 5th International DMTF Academic Alliance Workshop on Systems and Virtualization Management (SVM 2011).

[4] M. Carlson. “DMTF's Infrastructure Cloud Management Interface”. 2011 5th International DMTF Academic Alliance Workshop on Systems and Virtualization Management (SVM 2011).

[5] A. Chillarón and J. E. López De Vergara. “Design and Implementation of an Autonomic Management System for Virtual Machines in a Cloud Environment”. 2011 5th International DMTF Academic Alliance Workshop on Systems and Virtualization Management (SVM 2011).

[6] D. Fernández, A. Cordero Ordoñez, J. Somavilla, J. A. Rodriguez, A. Corchero, L. Tarrafeta and F. Galan. “Distributed Virtual Scenarios over Multi-host Linux Environments: Virtual Networks over LinuX (VNX)”. 2011 5th International DMTF Academic Alliance Workshop on Systems and Virtualization Management (SVM 2011).

[7] L. Lamers. "DMTF SVPC Work Group". 2011 5th International DMTF Academic Alliance Workshop on Systems and Virtualization Management (SVM 2011).

[8] H. Hlavacs and T. Treutner. "Predicting Web Service Levels During VM Live Migrations". 2011 5th International DMTF Academic Alliance Workshop on Systems and Virtualization Management (SVM 2011).

[9] H. Shah and T. Slaight. "Platform Management Components Intercommunications: Technology Update". 2011 5th International DMTF Academic Alliance Workshop on Systems and Virtualization Management (SVM 2011).

[10] J. Hilland. "SMASH & DASH Overview". 2011 5th International DMTF Academic Alliance Workshop on Systems and Virtualization Management (SVM 2011).