Convergence of Distributed Clouds, Grids and Their Management

WETICE2011 – Paris, June 27 – 29, 2011

Track Chair's Report

Dr. Rao Mikkilineni
IEEE Member
Kawa Objects Inc.,
Los Altos, California, USA
rao@kawaobjects.com

Dr. Giovanni Morana
DIEEI
University of Catania
Catania, Italy
giovanni.morana@dieei.unict.it

***Abstract*— The stated objective of the first workshop on Collaboration and Cloud Computing" in WETICE 2009 was "to analyze current trends in Cloud Computing and identify long-term research themes and facilitate collaboration in future research in the field that will ultimately enable global advancements in the field that are not dictated or driven by the prototypical short term profit driven motives of a particular corporate entity."**

This report describes the progress made within the span of three conferences which helped develop a new approach to the convergence of distributed clouds, grids and their management. This track presents a new computing model and its implementation which resulted directly from the collaboration of the two workshops, Collaboration and Cloud Computing Workshop and the Emerging Technologies for Next-Generation grids sponsored under the Aegis of WETICE. In this paper we present some of the key issues that still need to be addressed to improve the efficiencies and utilize the new generation of many-core servers that are transforming the information technology landscape. We hope WETICE2011 will continue the tradition to forge new collaborations and lead to innovation without a short term agenda.

Keywords-component; Cloud Computing; grid computing; Distributed Intelligent Managed Element Networks; Distributed Services Management; Services Virtualization; Parallel Computing;Many-core Servers

I. INTRODUCTION

According to Prahlad and Krishnan [1], the traditional sources of advantage — access to technology, labor and capital — are no longer unique differentiators for most firms. They say that the "new source of competitive differentiation may lie in the internal capacity to reconfigure resources in real-time." It is hardly an exaggeration, today, to say that the communications, collaboration and commerce at the speed of light, have dramatically altered both the consumer and business economics. There are three major business drivers that are forcing new approaches to meet the global demand to deliver distributed services:

1. The consumer appetite for new generation of Internet hosted applications such as Twitter, Facebook, You Tube, Hulu, Nertflix, and Animoto are driving the need for service providers to address the need to support massive scalability and address wild fluctuations in demand.
2. Enterprises are starting to deploy their own private clouds in order to protect their proprietary applications but still leverage the benefits of virtualization and in some cases use hybrid clouds that combine both private and public clouds. The massive scalability and the economics of multi-tenancy through increased sharing of resources are forcing enterprises to re-architect their data centers to leverage globally distributed resources.
3. Escalating power, space and energy costs are driving enterprise data centers to consolidation and technology refresh.

While cost savings, the need for massive scaling, multi-tenancy with infrastructure consolidation and requirements for high availability, performance and security are driving the requirements for new approaches to distributed systems design, there are three key technology trends that are forcing the re-engineering of current services development, delivery and assurance infrastructure:

1. ***Many-core Technology:*** The hardware upheaval caused by the new generation of processor chip, consisting of dozens to hundreds of cores with parts of each core that can be dedicated to specific purposes like network management, graphics, encryption, and decryption along with a majority of cores dedicated for application programs, is forcing the reexamination of operating systems, programming paradigms and software architectures. In addition, hardware assisted virtualization made available in large servers to small

desktop, laptop and mobile computing platforms makes it possible to deploy software independent of physical infrastructure with features such as “live migration.” These hardware advances provide a unique opportunity for a new class of distributed systems that can leverage distributed computing, network and storage clouds transcending physical and geographical boundaries and constrained only by the application latency requirements, workload variations and business priorities. The decoupling of business workflow execution from the underlying hardware infrastructure also demands the decoupling of infrastructure management from the workflow management itself and offers an opportunity for a new generation of distributed systems design that exploits parallelism, distribution and scaling of resource utilization.

2. ***High performance servers with hundreds of multicore processors coupled with high bandwidth networks:*** They are altering the current data center landscape by providing consolidation of unparalleled proportion. A large server containing hundreds of multicores with a DS3 pipe can eliminate a host of current generation special-purpose networking equipment and the associated management complexity if the processors inside the server can be programmed to scale and exploit the parallelism. In addition, high bandwidth access to mini-clouds aggregating virtualized desktops, laptops and mobile devices offers seamless connectivity and interoperability thus enabling a new class of distributed systems that span across multiple clouds.

3. ***Current generation of software development environments***: Both on web based delivery platforms and mobile broadband delivery platforms, the new development environments have radically transformed the way business and consumer applications are developed and delivered. They have unleashed a new generation of service developers outside the conventional business application domains. Web based or mobile platform based applications can be developed with such ease that we have a proliferation of social networks, games and educational services made available on a massive scale to consumer market. This has created a need for new application architectures to meet the massive scaling and low cost constraints dictated by the consumer market. In the past, services development was in the purview of specialized programmers and IT development organizations. With the new tools and platforms providing ubiquitous access to millions of developers, to develop trillions of services on worldwide service delivery platforms has become simpler. The service developers can now focus on their domain and not on programming detail. The domain workflows can be translated to executing services and deployed on web services platforms today much faster than application development in IT organizations that used to consume months and years in the past. The massive scale of services development and deployment also has necessitated the need for services management to meet the mass scale usage, workload fluctuations, latency constraints imposed by the speed of light in delivering services across the globe and the business priorities.

In the first workshop, the papers discussed the state of the art of cloud computing and a lively discussion followed resulting in the following questions:

- Is the cloud just an XaaS stack or is there more to it?
- Is Unified Computing simply bundling servers, network and storage in blades with virtualization?
- Are lessons from the past i.e. from telecommunications and from the Internet, relevant to enabling massively scalable and globally interoperable clouds?
- How can infrastructure hardware vendors accelerate cloud deployment by including service enabling features in their hardware?
- Is the current trend in throwing multiple Operating Systems inside the server and including another networking abstraction layer that bridges them, the right model or should we look at a new network centric operating system that allows dynamic composition of distributed physical computing resources based on latency tolerance of services consuming the logical resources?
- Can we accelerate the creation of computing clouds through fresh ideas such as the concept of a virtual infrastructure fabric, a management services fabric and a business services fabric?

In the 2010 edition of WETICE, the chairs of the cloud computing workshop and the emerging technologies for next generation Grids recognized the synergy of the two workshops and agreed to hold a joint session. This resulted in a very lively discussion about computing models and the need for a unified computing theory. One of the proposals of a computing model called the DIME (distributed intelligent managed element) network computing model piqued the interest of two of the participants and has resulted in three contributions in this conference. The first track on convergence of distributed computing clouds, Grids and their management in WETICE 2011 is devoted to further the collaboration of experts in cloud computing, high performance computing, Grid computing and operating system disciplines. The track is divided into two themes:

1. Many-core processors and their impact on computing models
2. Convergence of distributed cloud-based services, grid-based services and their management

In Section II, we comment on the DIME network computing model and related papers. In Section III, we summarize three long papers dealing with clouds, grids and their management and in Section IV, we introduce three short papers. We sincerely hope that the current track will create further interesting collaboration and advance the art and science of distributed computing transcending the short term commercial interests. The upheaval in hardware innovation if matched by an appropriate software innovation will radically transform the current data centers and cause in our view a major transformation of current computing and communications industries. We believe that the evolution toward the transformation of the data centers from their current role of being just server, networking, and storage hosting centers to service switching centers with telecom grade trust is just around the corner.

II. COMPUTING MODELS AND MANY-CORE SERVERS

The DIME computing Model [2, 3] discussed in WETICE2010 is derived from the genetic computing model[1] and supports the genetic transactions of replication, repair, recombination and reconfiguration. The DIME network architecture (DNA) allows us to implement a managed workflow as a set of tasks, arranged or organized in a directed acyclic graph (DAG) and executed by a managed network of distributed computing elements (DIMEs). These tasks, depending on user requirements are programmed and executed as loadable modules in each DIME. The first paper by Morana and Mikkilineni describes a prototype demonstration of the DIME computing model on top of LINUX operating system. It decomposes each service into two parts 1) a service regulator and 2) a service package and implements a parallel computing workflow and a control workflow to orchestrate the end-to-end service workflow (Figure 1).

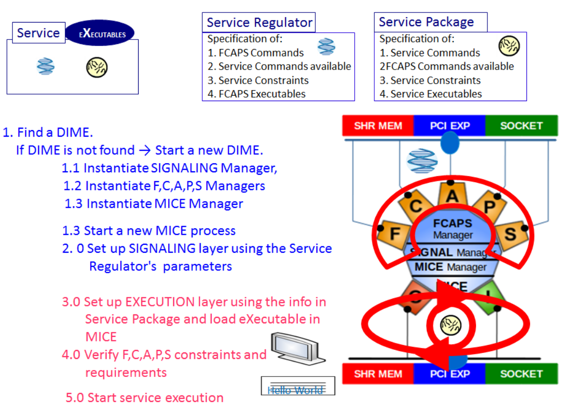

Figure 1: The DIME computing model incorporates self-management of computing elements and network-aware signaling capability to collaborate as a group to implement a managed workflow.

Self-management of each computing element, network-aware signaling abstractions and both node level and network level fault, configuration, accounting, performance and security (FCAPS) management using parallelism distinguish this model from von Neumann stored program control computing model. While the paper describes a small application as an example, it provides the proof of concept demonstration of the various features that define the model. It claims to provide shared memory, PCIexpress and Socket based TCP/IP communication between various DIMEs based on latency requirements and available resources. This offers an interesting solution to communication between many Linux images deployed today in many-core servers to deliver web services[2]. Only time will decide its usefulness in many-core servers.

The second paper by Mikkilineni and Seyler [3] describes a new operating system specially designed for multi-core servers utilizing the DIME computing model. It also briefly reviews the current efforts that attempt to address many of the issues with current operating systems when deployed in many-core servers. The operating system called Parallax distinguishes itself by supporting signaling and FCAPS management at the operating system level and is implemented in assembler language for efficiency to work currently with Intel Xeon multi-core chips. It demonstrates replication, dynamic reconfiguration, and repair of applications deployed in multiple many-core servers as a DIME network without the need for a Hypervisor. It seems to offer an alternative to current virtualization methods for "live-migration". While the operating system is in its nascent form, it offers a new approach that is different from current research efforts in operating systems. As the paper points out, "The history of the evolution of current OSs is filled with lessons on wasted billions (does anyone remember Multics or OS2?), unmet expectations (who would have thought UNIX, the original System V, would vanish), surprise winners (Windows and Linux), and stealthy survivors (Mach in a Mac)." Only time will decide its usefulness.

The third paper by Celesti and Tusa [5] related to the computing model describes how the security can be implemented leveraging the parallelism offered by the DIME computing model. The dynamism offered by the signaling mechanism is fully exploited by the authors to assure application specific security both at the node level and at the network level to provide transaction security that spans across multiple distributed infrastructures.

III. THE CONVERGENCE OF COLUD COMPUTING AND GRID COMPUTING

Resource optimization strategies are very important in both grid and cloud utilization and the fourth paper presents CHASE, an autonomic engine designed to optimize the scheduling of virtual machines in a cloud environment. The prototype of CHASE is tested and evaluated on PerfCloud, an environment for IaaS provision based on cloud and grid

[1] As Mitchell Waldrop explains in his book on Complexity [4], "the DNA residing in a cell's nucleus was not just a blue-print for the cell – a catalog of how to make this protein or that protein. DNA was actually the foreman in charge of construction. In effect, was a kind of molecular-scale computer that directed how the cell was to build itself and repair itself and interact with the outside world."

[2] While many-core servers provide space and energy cost savings, current operating systems cannot support more than few tens of cores in a single OS image and the multiple images in the same enclosure communicate with each other over socket communication with TCP/IP protocol.

integration, and on Cloud@Home, a Seti@home-like application based on BOINC and volunteer-computing.

In the fifth paper, the authors make a detailed analysis on the use of "small world" effect for building an effective algorithm for resources finding in distributed environments which are most commonly used with SOA – both P2P Grid and Cloud environments.

The last paper presents a Java classloader, entirely based on web services, called WSBCL (Web Services Based Classloader). WSBCL allows loading Java classes in systems with connectivity restrictions (firewall and NAT), low computational power (mobile phones, set-top boxes, etc.) and different platforms/programming languages. In particular, the authors describes the architecture and the implementation of WSBCL classloader in the Grid Anywhere, a middleware for grid computing that aims at allowing many kinds of devices to be used to build a distributed desktop grid computing environment.

IV. SHORT PAPERS

WETICE prides itself in encouraging graduate students and researchers bringing new ideas that have not yet been accepted by the main stream or authors who would like to present yet-incomplete research to get feedback from experts to present their ideas in the form of short papers. It is hoped that the following discussions will result in some new research papers in the future. The DIME computing model and resulting papers presented in this session are the results from such a discussion in the last two conferences.

In this track, three short papers are presented:

1. The first short paper presents the "Sam Dog," a Java Sandbox which allows a flexible security policy specification using its hierarchical/cascading rules, model and context attribute/value pairs. It is useful especially when "unknown" applications have to be executed. It defines and enforces the Access Control Policies adopting access rules organized by three elements: Domain Security Manager (DSM), entity and task. Since the implementation has been made using independent modules, the Sam Dog Policy Manager can be easily used to manage the access control in many different distributed environments, including clouds and desktop grids.

2. The second short paper "Managing Dependability in Distributed System" proposes a new strategy for managing dependability issues on distributed environments. This strategy, exploits the concepts of "centrality" and "clustering" taken from the graph theory, and is based on the analysis of the deployment of the applications on the resources and on the most used "distributed transactions" for obtaining an allocation/scheduling policy to improve the robustness of the applications in distributed environments.

3. The final short paper "An improvement of a different approach for medical image storage" presents a newer version of the HDF5 paradigm for storing and retrieving medical images within a PACS (Picture Archiving and Communication System). In particular, the authors suggest the adoption of two additional services, the C-Store and C-Find that improve the efficiency of the whole system.

V. CONCLUSION

The website http://wetice.org claims "what sets WETICE apart from larger conferences is that the conference tracks are kept small enough to promote fruitful discussions on the latest technology developments, directions, problems, and requirements." We believe that this had an influence in the evolution of the collaboration and computing clouds and the emerging technologies for next generation grids workshops that allowed some new ideas to be introduced and evolved in a collaboration that spanned across multiple countries. We hope the current track on the convergence of distributed clouds, grids and their management will continue the tradition with lively and unconstrained discussion of emerging business and technology trends. As von Neumann [5] pointed out "The basic principle of dealing with malfunctions in nature is to make their effect as unimportant as possible and to apply correctives, if they are necessary at all, at leisure. In our dealings with artificial automata, on the other hand, we require an immediate diagnosis. Therefore, we are trying to arrange the automata in such a manner that errors will become as conspicuous as possible, and intervention and correction follow immediately." Comparing the computing machines and living organisms, he points out that the computing machines are not as fault tolerant as the living organisms. He goes on to say "It's very likely that on the basis of philosophy that every error has to be caught, explained, and corrected, a system of the complexity of the living organism would not run for a millisecond." We have a long way to go from current computing architectures to architectures that mimic living systems. The many-core computing architectures and parallel distributed and scalable computing models are a first step in that direction.

ACKNOWLEDGMENT

The Chair persons thank many of the reviewers and the WETICE organizing committee for their support.

REFERENCES

[1] Prahalad C.K, and Krishnan M. S, "The New Age of Innovation", McGrawa Hill, New York, 2008 p7

[2] Mikkilineni, R "Is the Network-centric Computing Paradigm for Multicore, the Next Big Thing?" Retrieved July 22, 2010, from Convergence of Distributed Clouds, Grids and Their Management: http://computingclouds.wordpress.com

[3] Rao Mikkilineni and Ian Seyler, "Parallax – A New Operating System for Scalable, Distributed, and Parallel Computing", The 7th International Workshop on Systems Management Techniques, Processes, and Services, Anchorage, Alaska, May 2011

[4] Mitchell Waldrop, M., "Complexity: The Emerging Science at the Edge of Order and Chaos", Penguin Books, London, 1992, p218

[5] Neumann, J. v., General and Logical Theory of Automata. edited and compiled by William Aspray and Arthur Burks, MIT Press, 1987, p408