Dynamically Generating and Orchestrating Virtual Machines on the Grid

by

Timothy H. Ward

Research Proposal

Submitted by Timothy H. Ward

in partial fulfillment of the Requirements for the Degree of

Bachelor of Software Engineering with Honours (2770)

Supervisor: David Abramson and Wojtek Goscinski

Clayton School of Information Technology

Monash University

May, 2008

Contents

1 Introduction **2**

2 Research Context **4**

3 Research Plan and Methods **6**

3.1 Research Methods 6

3.2 Proposed Thesis Chapter Headings 6

3.3 Timetable 7

3.4 Special Facilities Required 8

4 Contribution of the Proposed Research **9**

1 Introduction

Grid computing has opened up possibilities for e-Scientists to conduct and collaborate on computer intensive experiments which would have once been infeasible[5]. The next generation of large scale experiments for E-Science requires access to large scale computing resources and data storage, and grid computing has made this possible. High-performance computing experiments can now be run without requiring direct access to a single super-computer.

These experiments are often represented by scientific workflows[19], which involve the break down of an experiment into logical and ordered components that may be responsible for collecting and processing data. These components may be dependent or a dependency for other components within the workflow. These scientific workflows have simplified the process of designing and executing scientific experiments.

Due to the nature of the computing landscape, grids commonly consist of heterogeneous resources; every resource on a grid can potentially have different physical characteristics and a different configuration. For an e-Scientist to successfully use the full potential of a grid they must tailor their experiment to run on all or a subset of these resources. In most cases an e-Scientist may have some experience in software development. However, their main concern is in their field of research. For e-scientists, the process of developing and deploying software across a range of platforms, configurations and organisational boundaries is challenging[6, 7].

One approach to reducing the effort required for developing grid applications is using virtualisation to abstract resource characteristics and allow e-Scientists to define their own run-time environment for an experiment application[3]. Using this approach removes potential application development issues such as portability from the e-Scientist's responsibility. This can be achieved by using platform virtual machines, which emulate a complete machine including its hardware, operating system, and software. However this method still poses some problems for e-Scientists as the configuration of such environments can be time consuming and requires knowledge of operating systems concepts and system administration.

As mentioned, the process of creating and configuring these environments is tedious. Environments initially need to be configured with the base requirements for the experiment such as an operating system, software libraries, and other application dependencies. The experiment application then needs to be configured and installed within the environment. Experiment data then needs to be sourced and passed into the environment either from the experiment repository and/or being streamed from another experiment application. The experiment application then needs to be executed within the environment and be monitored to ensure that progress is made. Once the experiment is completed the experiment results need to be passed back to the experiment repository and/or passed on to another experiment application. The environment then needs to be cleaned up to allow the releasing of the underlying grid resource. These experiment applications are usually incorporated into a scientific workflow and as such this process of configuring environments needs to be repeated multiple times.

This project aims to look at how virtual machine environments can be automatically and dynamically created as work units and how they could be incorporated into a workflow. This will include investigating the issues related with the development and deployment of software on to virtual machines, and how the virtual machines can be deployed, maintained, and orchestrated across the grid. There by, simplifying the effort required by e-Scientists to conduct high-performance computing experiments on grid infrastructure as well as in the hope of a greater adoption of grid computing.

2 Research Context

Attempts have been made to standardise and make grid computing more accessible[4, 17]. However, even with such toolkits and standards, developing for grid computing still remains a challenge[6, 7]. In response, virtualisation is being applied as a solution to this problem[3]. Virtual machines have been successfully applied to grid computing, using both application level virtualisation and platform level virtualisation[9, 12, 1].

Application virtual machines that make use of .Net/Java virtual machine technologies are already implemented as middleware across different grid implementations[8]. These implementations allow e-Scientists to develop portable applications which can be executed across a range of environments. In some cases these environments support legacy code to a certain degree[8]. However, they do not give complete control for the e-Scientist to completely specify their experiments run-time environment; this includes controlling the underlying operating system and other legacy application dependencies.

Platform virtual machines abstract the entire computing resource by virtualising the underlying computer hardware thus emulating a complete computing environment in which the user actions and/or executions will not directly affect the underlying resource[16]. Platform virtual machines provide isolation, legacy-support, administrator privileges, resource control, and environment recovery[3]. Combined, these characteristics have the potential to provide a high level of control when conducting experiments in a grid environment. Furthermore, recent advances in virtualisation and virtual machines has led to performance overheads being dramatically decreased and has made it feasible to apply virtual machines to high-performance computing[13]. Virtual machines also allow the saving and restoration of machine states using checkpointing which can be used in speeding up of execution, such as removing the initial system booting times[14].

The integration of grid computing middleware with platform virtual machines has led to two major architectures. The first approach has led to the placing of grid middleware into the virtual machine. This approach allows grids to be implemented and deployed by running these virtual machines on the grid resources[9]. This is following the traditional steps of virtual appliances[15], where virtual machines are used to distribute software; in this case the appliance is the grid middleware. This architecture allows users to have access to a uniform set of resources. However, it does not allow e-Scientists to have the ability to completely define the environment for which their experiment will be executed in. The second approach is using existing grid infrastructure and using virtual machines as a work units for the execution of applications[14, 12, 1]. In this case the grid middleware is used for supporting the virtual machine deployment and execution, though it should be noted that the first approach can be used in conjunction with this method.

Virtual machine work units are often referred to as sandboxes as they allow users to customise the execution environment without compromising the resource[18]. Some techniques have been developed so that the virtualisation level is abstracted and the platform, process, and other virtualisation levels are not specific to a particular environment; rather the environment is a dynamic virtual environment[11] or virtual workspace[10].

Customisation of such environments is generally through a specification which is provided by the user. These may be passed to a service for auto-configuration[12, 1] or in some cases the user is provided a graphical user interface to the virtual machine[3]. VMPlant[12] is one example of a virtual machine factory that is responsible for creating virtual machines based on a configuration file supplied by the user. Most configurations provide the ability to specify the operating system, required network settings, application and application library dependencies, and user folder locations. These configurations can be specified using descriptions like directed acrylic graphs[12] or using XML schemas[10].

Creating these environments based on specifications can be achieved by different interactions with the virtual machines. For example VMPlant mounts CD-ROM ISO images and uses the auto-run to launch scripts for configuration[12], whereas other methods break environments into separate pre-defined modules (operating system, application, etc) and combine then based on the configuration[10].

Recently Amazon released a service for using its computing infrastructure. Amazon Elastic Compute Cloud (Amazon EC2)[2] provides developers with an application programming interface for dynamically creating and managing virtual machines on their large scale web infrastructure. This service can be utilised for grid computing by e-Scientists.

This project will look at these different methods for creating, configuring, deploying and orchestrating virtual machines across different grid infrastructures, such as Amazon EC2. This research will be done with the greater goal of developing a model which allows users to easily execute grid experiments using virtual machines.

3 Research Plan and Methods

3.1 Research Methods

This project will follow a iterative software development life-cycle (SDLC) model for building a software system that allows the creation of virtual machines as work units and the orchestration of these across grid infrastructure.

This will begin with looking at current virtual machine technologies looking at how they are implemented in terms of virtual machine images and whether the virtual machines will be emulated or virtualised. Considerations for development and deployment of grid software will also be analysed to determine what requirements are needed for the configuration of virtual machines.

After this, an overall design will be created based on the previous findings. This model will include a method of dynamically creating and configuring virtual machine images that meet the required environment as specified by an e-scientist. The project will also look how these work units can be incorporated into a workflow.

The resulting design will be implemented as a software system. The development will use a combination of languages; Perl and C++. The initial system will be developed on a Linux-based host but will be implemented following software engineering conventions to allow easy portability and maintenance.

The finalised implementation will then be evaluated to determine if the creation of virtual machine work units can be scaled on the grid, and whether it provides effective means of conducting high-performance computing across the grid.

3.2 Proposed Thesis Chapter Headings

1. Introduction
2. Virtual Machines and Grid Computing
   (a) Virtual Machines
   (b) Grid Computing
   (c) Grid Computing Using Virtual Machines
3. Design of Virtual Machine Work Units
   (a) Configuration
   (b) Execution
   (c) Data Management
   (d) Workflow Incorporation

4. Implementation of Virtual Machine Work Units
    (a) Configuration
    (b) Execution
    (c) Data Management
    (d) Workflow Incorporation
5. Evaluation of Virtual Machine Work Units and Workflow Incorporation
6. Conclusion and Future Work
7. Bibliography
8. Appendices
    (a) Appendix A Data
    (b) Appendix B Programs

3.3 Timetable

Date	Week #	Activity	Deliverable
7/3	S1, W2/13	Commence Literature Review	
18/4	S1, W7/13	Commence Prototype Implementation	
25/4	S1, W8/13	Research Proposal Draft Writing Commencement	
28/4	S1, W9/13	Research Proposal Draft to Supervisor	
2/5	S1, W9/13	Research Proposal Finalisation	
9/5	S1, W10/13	Literature Review Draft Writing Commencement	*Research Proposal
23/5	S1, W12/13	Literature Review Draft to Supervisor	
30/5	S1, W13/13	Literature Review Draft Finalisation, Commence Design	
3/6	MY, W1/6	Interim Presentation Preparation	
6/6	MY, W1/6	Thesis Draft Writing Commencement, Thesis Chapter 2 Writing Commencement	*Literature Review
10/6	MY, W2/6	*Interim Presentation	
13/6	MY, W2/6	Commence Implementation	
4/7	MY, W5/6	Thesis Chapter 3 Writing Commencement	
5/9	S2, W8/13	Thesis Chapter 1 Writing Commencement	
8/10	S2, W12/13	Final Presentation Preparation	
10/10	S2, W12/13	Thesis Draft to Supervisor	
17/10	S2, W13/13	Thesis Finalisation	
22/10	EY, W1/2	*Final Presentation	
31/10	EY, W2/2		*Thesis

Key
S1 - Semester 1 **S2** - Semester 2
MS1 - Mid-Semester 1 **MS2** - Mid Semester 2
MY - Mid Year **EY** - End Year
***** - Milestone

Recurring Timetable

Occurrence	Begin Date	End Date	Description
Weekly	22/2	31/10	Supervisor Meeting - David Abramson and Wojtek Goscinski

3.4 Special Facilities Required

This project may require access to a grid infrastructure for testing and evaluating the effectiveness of the project implementation, though this is yet to be determined if required. For all other components the facilities provided by Monash University are sufficient for completion of this project.

4 Contribution of the Proposed Research

Most virtual machine grid solutions provide a means to use virtual machines to abstract grid resources but still provide no means for the automatic creation of virtual machine as work units. This project will investigate issues related with the automatic creation and deployment of virtual machines on to the grid as work units and how they can be incorporated into a workflow.

By abstracting resources and providing automatic environment creation will allow e-Scientists to concentrate on the core of their work, their experiment. This may also provide e-Scientists without the necessary technical skills to be able to begin and use grid technology in there high-performance computing experiments.

References

[1] Sumalatha Adabala, Vineet Chadha, Puneet Chawla, Renato Figueiredo, José Fortes, Ivan Krsul, Andrea Matsunaga, Mauricio Tsugawa, Jian Zhang, Ming Zhao, Liping Zhu, and Xiaomin Zhu. From virtualized resources to virtual computing grids: the in-vigo system. *Future Gener. Comput. Syst.*, 21(6):896–909, 2005.

[2] Amazon. Amazon elastic compute cloud (amazon ec2), May 2008. `http://www.amazon.com/gp/browse.html?node=201590011`.

[3] Renato J. Figueiredo, Peter A. Dinda, , and José Fortes. A case for grid computing on virtual machines. In *23rd IEEE International Conference on Distributed Computing Systems (ICDCS'03)*, page p. 550, 2003.

[4] Ian Foster and Carl Kesselman. Globus: A metacomputing infrastructure toolkit. *The International Journal of Supercomputer Applications and High Performance Computing*, 11(2):115–128, Summer 1997.

[5] Ian Foster and Carl Kesselman, editors. *The grid: blueprint for a new computing infrastructure*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1999.

[6] Wojtek Goscinski and David Abramson. Distributed ant: A system to support application deployment in the grid. In *GRID '04: Proceedings of the Fifth IEEE/ACM International Workshop on Grid Computing*, pages 436–443, Washington, DC, USA, 2004. IEEE Computer Society.

[7] Wojtek Goscinski and David Abramson. Application deployment over heterogeneous grids using distributed ant. In *E-SCIENCE '05: Proceedings of the First International Conference on e-Science and Grid Computing*, pages 361–368, Washington, DC, USA, 2005. IEEE Computer Society.

[8] Wojtek Goscinski and David Abramson. Motor: A virtual machine for high performance computing. In *High Performance Distributed Computing, 2006 15th IEEE International Symposium on*, pages p. 171–182, Melbourne, Australia, 2006.

[9] Joomla! Grid appliances, May 2008. `http://www.grid-appliance.org/`.

[10] K. Keahey, I. Foster, T. Freeman, and X. Zhang. Virtual workspaces: Achieving quality of service and quality of life in the grid. *Sci. Program.*, 13(4):265–275, 2005.

[11] Katarzyna Keahey, Karl Doering, and Ian Foster. From sandbox to playground: Dynamic virtual environments in the grid. In *GRID '04: Proceedings of the Fifth IEEE/ACM International Workshop on Grid Computing*, pages 34–42, Washington, DC, USA, 2004. IEEE Computer Society.

[12] Ivan Krsul, Arijit Ganguly, Jian Zhang, Jose A. B. Fortes, and Renato J. Figueiredo. Vmplants: Providing and managing virtual machine execution environments for grid computing. In *SC '04: Proceedings of the 2004 ACM/IEEE conference on Supercomputing*, page 7, Washington, DC, USA, 2004. IEEE Computer Society.

[13] C. Macdonell and P. Lu. Pragmatics of virtual machines for high-performance computing: A quantitative study of basic overheads. In *2007 High Performance Computing & Simulation Conference (HPCS'07)*, Prague, Czech Republic, June 4-6 2007.

[14] Sriya Santhanam, Pradheep Elango, Andrea Arpaci-Dusseau, and Miron Livny. Deploying virtual machines as sandboxes for the grid. In *WORLDS'05: Proceedings of the 2nd conference on Real, Large Distributed Systems*, pages 2–2, Berkeley, CA, USA, 2005. USENIX Association.

[15] Constantine Sapuntzakis, David Brumley, Ramesh Chandra, Nickolai Zeldovich, Jim Chow, Monica S. Lam, and Mendel Rosenblum. Virtual appliances for deploying and maintaining software. In *LISA '03: Proceedings of the 17th USENIX conference on System administration*, pages 181–194, Berkeley, CA, USA, 2003. USENIX Association.

[16] James E. Smith and Ravi Nair, editors. *Virtual Machines: Versatile Platforms for Systems and Processes*. Elservier Inc., San Francisco, CA, USA, 2005.

[17] Todd Tannenbaum, Derek Wright, Karen Miller, and Miron Livny. Condor: a distributed job scheduler. pages 307–350, 2002.

[18] David Isaac Wolinsky, Abhishek Agrawal, P. Oscar Boykin, Justin R. Davis, Arijit Ganguly, Vladimir Paramygin, Y. Peter Sheng, and Renato J. Figueiredo. On the design of virtual machine sandboxes for distributed computing in wide-area overlays of virtual workstations. In *VTDC '06: Proceedings of the 2nd International Workshop on Virtualization Technology in Distributed Computing*, page 8, Washington, DC, USA, 2006. IEEE Computer Society.

[19] Jia Yu and Rajkumar Buyya. A taxonomy of scientific workflow systems for grid computing. *SIGMOD Rec.*, 34(3):44–49, 2005.