International Journal of Electronics Communication and Computer Engineering
Volume 2, Issue 2, ISSN (Online): 2249–071X, ISSN (Print): 2278–4209

Multi-Site Virtual Cluster Management Mechanism for Grid Computing Environments

Mr. Kailash Patidar[1] Prof. Ravindra Gupta[2] Prof. Gajendra Singh[3]
Mr. Ajay Mishra[4] Mrs. Smita Sharma[5]

***Abstract* - Grid computing is a term referring to the combination of computer resources from multiple administrative domains to reach a common goal. The grid can be thought of as a distributed system with non-interactive workloads that involve a large number of files. What distinguishes grid computing from conventional high performance computing systems such as cluster computing is that grids tend to be more loosely coupled, heterogeneous, and geographically dispersed. Although a grid can be dedicated to a specialized application, it is more common that a single grid will be used for a variety of different purposes. Grids are often constructed with the aid of general-purpose grid software libraries known as middleware. Grid computing is a promising technology for utilizing distributed computing resources seamlessly. However, it is still difficult for end-users to deploy and manage their own computing environments into large number of distributed locations easily and rapidly. Recent grid technologies do not address the problems involved in distributed deployment and management of grid-enabled system programs; they do not allow end-users to obtain computing resources dynamically from different locations and organizations, and to build their computing systems with them easily and scalably. Therefore, we propose a user-oriented, distributed deployment and management mechanism for grid computing systems, based on multi-site virtual clusters. It enables end-users to easily configure grid computing environments on distributed computing resources through their own large-scale clusters of virtual machines. It is composed of inter-domain resource control protocols, resource virtualization technologies including virtual machines and virtual private networks, and scalable virtual node management mechanisms.**

***Key words* - Cluster Management, Grid Computing, Virtual Network's, WANS.**

I. Introduction

In grid computing environments, users want to access distributed computing resources located in geographically separated locations, and administrated by different organizations, seamlessly, recent grid technologies enable users to utilize remote resources through service interfaces provided by grid-enabled middleware and application programs. These middleware suites and application programs for computing resource collaboration need to be installed and configured appropriately in advance of being accessed by users. Only privileged users can perform system-wide deployment and configurations of such software programs.

For massively distributed resources of different organizations, however, this "manpower intensive" deployment task becomes a bottleneck hindering rapid and dynamic collaboration over diverse computing resources by end-users. It is difficult to deploy software programs with their appropriate settings in ever changing heterogeneous resource environments, in an efficient manner for all kinds of computing resource usage. This problem is caused by administrator oriented management approaches to resource allocation and software deployments, which means that end-users cannot decide how to combine distributed resources on demand and how to fully configure them for their purposes. In this paper, we propose a user-oriented, distributed deployment and management mechanism for grid computing systems, which is based on multi site virtual clusters; allowing users to incorporate distributed computing resources into their own single virtual clusters over WAN. Our mechanism is composed of inter-domain resource control protocols, resource virtualization technologies including virtual machines (VMs) and virtual private networks (VPNs), and scalable virtual node management mechanisms. It allows greater autonomy for End-users in obtaining and utilizing distributed computing resources from differ- ENT organizations.

A multi-site virtual cluster is composed of distributed VMs in different locations, which are created by end-users via the inter domain resource control protocols. The virtualization mechanisms available from resource providers create VMs at the request of end-users, and the allocated VMs are interconnected via VPNs for single cluster views. Its virtual node management mechanisms allow users to deploy operating systems and grid applications in their distributed VMs rapidly, and to customize each virtual node flexibly, with the minimum management cost. The contribution of this paper is to address system deployment and management problems in large-scale distributed computing environments, which have not yet been tackled by both grid and virtualization technologies. It presents novel system deployment and management mechanisms that enable end-users to build their own distributed computing environments easily and rapidly. It also shows the validity of the proposed mechanisms in WAN environments through verification by our experiments.

II. Related Work

Although existing grid technologies provide remote access mechanisms for distributed resources, they do not address the distributed deployment and management of grid-enabled programs by end-users. Virtualization is the key to creating isolated computing environments for users. Recent virtualization research and technologies, however,

International Journal of Electronics Communication and Computer Engineering
Volume 2, Issue 2, ISSN (Online): 2249–071X, ISSN (Print): 2278–4209

have not addressed distributed deployment and management problems.

Cloud computing with virtualization, such as Amazon EC2 [1], is a web service through which users create VMs and customize them for their applications. However, because its target customers are web service providers and not grid application users, it does not have the functionality to manage large-scale distributed VMs from different organizations. Virtual Workspace [2] is a virtualized execution environment based on the Globus Toolkit framework. It allows clients to create VMs on remote resource providers through the WSRF protocols. However, it lacks an integration mechanism for VMs allocated in different locations. Our research aims to allow users to manage large-scale distributed VMs in an easy and flexible manner, as if they were in a single physical cluster; providing the distributed computing infrastructure through which end-users and different organizations can collaborate easily by using easy-to-use multi-site virtual clusters. Planet Lab [3] is a distributed computing test bed for large-scale network applications. Users can run their network programs on VMs distributed at many sites over the Internet. Since it focuses only on academic network research communities, it does not address deployment and management of grid computing systems for end-users.

III. Multi-site Virtual Clusters

A multi-site virtual cluster is an advanced cluster virtualization concept for greater scalability and flexibility, as well as sufficient usability. Beyond physical hardware limitations at one site, it aims to integrate distributed computational resources into a single cluster form by exploiting recent broad bandwidth backbone advances in WAN environments. A conceptual overview is illustrated in Figure 1, and its advantages are summarized as follows.

Scalability: A virtual cluster overlaying multiple sites can be composed of a greater number of computational nodes than can single-site hardware resources. This breaks hardware resource limitations at a single site, and integrates distributed computer resources as virtual clusters.

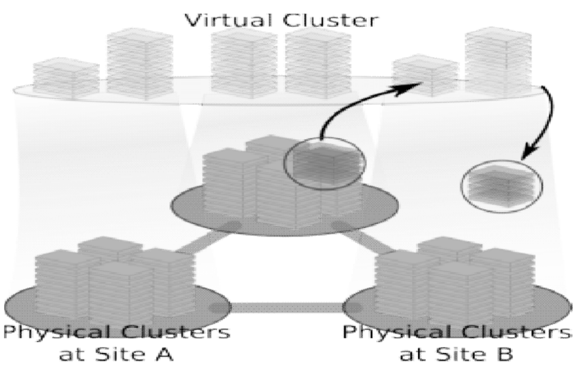

Fig.1. multi-site virtual clusters

In recent years, broadband networks over 1Gbps have been widely deployed between large-scale data centers and computer centers, and end-to-end reservations of network resources are becoming possible in academic and commercial backbone networks. These networks enable virtual clusters to be extended to other remote sites in a scalable manner.

Flexibility: A virtual cluster can be dynamically expanded according to user requests, by adding more and more virtual machines from remote sites. It is also possible to shrink it by releasing virtual machines. The virtual nodes belonging to a virtual cluster can be substitutable with other virtual machines of remote sites in a fairly transparent manner; for instance, some virtual nodes can be moved to other sites for hardware maintenance or power saving of physical clusters.

Usability: Even though a virtual cluster consists of distributed virtual machines located at remote sites, it can be seen as a single system image "cluster" for users and applications. Its underlying networking topology is appropriately hidden from the inside of the virtual cluster. Remote computer resources are integrated into a uniform computing environment over its underlying heterogeneity, so that existing cluster applications can be applied to distributed environments with smaller efforts. A user basically operates a virtual cluster in the same manner as a physical one, using unmodified programs designed for physical clusters.

IV. The Concept of Multi-site Virtual Clusters

We propose the concept of a multi-site virtual cluster, which is a single cluster of distributed VMs allocated in different locations. It is an easy-to-use computing environment, through which users can rapidly deploy their customized software programs to a large number of virtual nodes, and easily manage them as a single cluster.

The requirements of multi-site support are summarized as follows: First, the multi-site extension needs to have compatibility with existing components and transparency for the inside of a hosting virtual cluster. Our single-site virtual clusters are basically designed to be scalable, flexible, and usable as much as possible for available hardware resources. We consider this basic design architecture is also applicable to multi-site virtual clusters. Its multi-site support should be achieved by a straightforward extension of the existing design; scalability, flexibility, and usability are improved in a transparent manner for users, applications, and cluster management systems. Users should be able to utilize multi-site virtual clusters seamlessly by using existing programs for physical clusters.

Second, multi-site virtual clusters should be allocated rapidly, and their internal operating systems and applications should be also installed and configured quickly,

International Journal of Electronics Communication and Computer Engineering
Volume 2, Issue 2, ISSN (Online): 2249–071X, ISSN (Print): 2278–4209

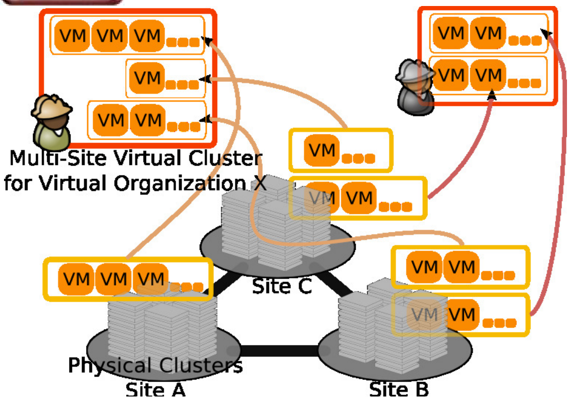

Fig. 2. User's view of multi-site virtual clusters

with the minimum manual configurations. These installations and configurations need to be performed dynamically over networks for additional virtual machines, so that users can fully customize their virtual clusters anytime. Both efficient deployment and full customization must be supported in this mechanism, even for remote nodes over network latencies.

Third, flexible relocation of virtual machines is required for management flexibility. Some virtual machines of a virtual cluster should be able to be moved to other host machines at remote sites; however the virtual cluster must continue to work properly during the move. Without this feature, it is difficult to maintain a large-scale virtual cluster distributed, across many sites consistently.

V. Design and Implementation

A multi-site virtual cluster is composed of virtual clusters at its master site and worker sites, and these virtual clusters at the sites are interconnected by Ethernet VPNs to be a single cluster. The cluster manager is modified to send/receive allocation requests of virtual clusters and also reserve network resources if possible. A virtual front end node of the local virtual cluster in the master site also works as the Rocks front end node of the multi-site virtual cluster; other virtual nodes at its worker sites are also installed and managed from the remote virtual front-end. Ethernet VPN services and transparent cache servers are incorporated into the virtual cluster management system as independent components for multi-site support.

We have developed a prototype implementation of the proposed multi-site virtual cluster system, which supports resource allocations and controls via the REST API. It has a portal Web site through which users can make reservations of multi-site virtual clusters and monitor their status. Through the portal Web site, users can automatically install an internal operating system optimized for system customization and deployment inside multi-site virtual clusters. These implementations are designed to be add-on packages of the NPACI Rocks cluster toolkit. They will be available under an open source license on our project page.

VI. Experiments

We performed experiments to show the validity of the proposed system. A primary issue is the deployment and management cost of widely distributed, multi-site virtual clusters. The proposed system has to enable users to quickly configure their customized systems into distributed VMs through the view of a single cluster. The experiments emulate a multi-site virtual cluster composed of VMs at two different locations. Its internal network is built with an Open VPN session linking them. It was captured during a demonstration creating multi-site virtual clusters from distributed computing resources at 5 remote sites. A user-created multi-site virtual cluster composed of 26 VMs is connected to the same network segment via VPNs. Six VMs were hosted on UCSD's cluster and others were at different locations in Japan. Four VPN connections were established for a private network over the Internet. Ethernet via a network emulator [8]. This adds network latencies to all traffic between the clusters to emulate a WAN environment.

The 134 VMs on Cluster B are reconfigured for the new settings of grid system programs from a frontend node on Cluster A. The VMs are reinstalled with a set of software packages with their new settings. All software packages are retrieved from the frontend node over the VPN session. A transparent package caching server, however, intercepts package requests and merges them to reduce network traffic over the VPN session.

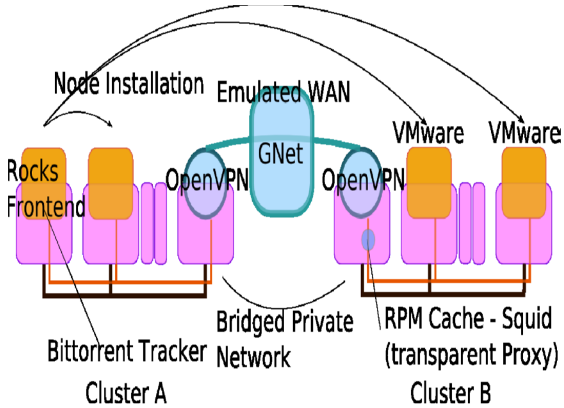

Fig.3. Experiment Setting

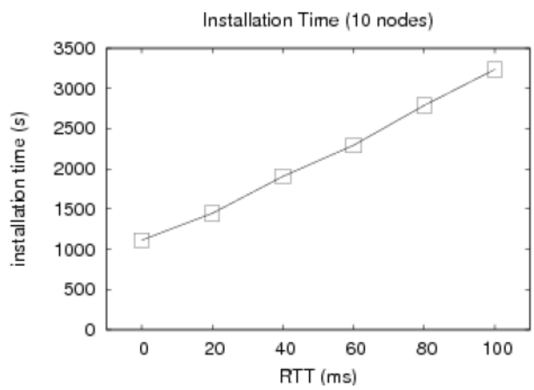

Figure 4: Installation time for 134 nodes under emulated network latencies

The reconfiguration was completed 800-1000 seconds faster by enabling a cache server. VPN traffic becomes much smaller than download traffic from the cache server. The total transferred data over the VPN is considered as a unique set of downloaded packages. In comparison with VPN traffic was far less than the case without caching. In the case where all packages are precached before reconfiguration, it repeatedly took 1300 seconds under various network latencies. VPN traffic was dramatically reduced due to the fact that there was no package transfer. It is reasonable to share common packages in a local cache repository at each location.

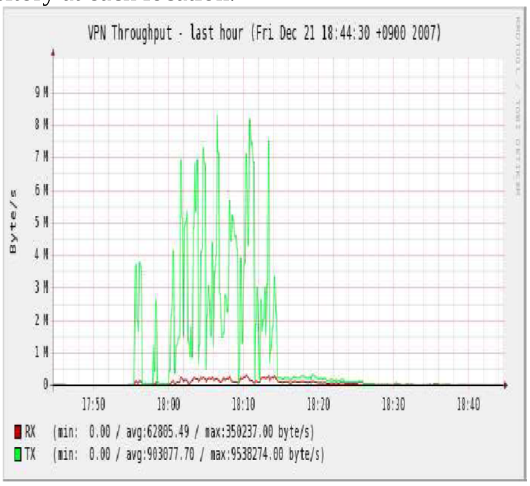

Figure 5: VPN throughput with package caching

These results clarify that our cluster configuration framework designed for multi-site virtual clusters resolves the major issues of distributed deployment and management of large-scale VMs. Distributed VMs in remote locations are rapidly configured by our efficient package caching mechanism, and also easily.

VII. Conclusions

We proposed a user-oriented, distributed deployment and management mechanism for grid computing systems, based on multi-site virtual clusters. It enables end-users to create large-scale clusters composed of VMs at different locations, through which they can easily and rapidly configure their grid system programs distributed over different locations. Our mechanism is composed of inter-domain resource control protocols, resource virtualization technologies including Vmsa and VPNs, and scalable virtual node management mechanisms. Our prototype implementation enables users to rapidly build multi-site virtual clusters over the Internet, and to easily configure their internal systems through the view of a single cluster. The experiments showed that our cluster management framework allowed rapid configuration of large-scale VMs under WAN latencies.

Acknowledgment

This supports resource allocations and controls via the REST API.

References

1. Amazon Elastic Compute Cloud. http://aws.amazon.com/ec2

2. Zhang, X., Freeman, T., Keahey, K., Foster, I., Scheftner, D.: Virtual clusters for grid communities. In: Proceedings of the Sixth IEEE International Symposium on Cluster Computing and the Grid (CCGRID2006), IEEE Computer Society (2006)

3. Peterson, L., Anderson, T., Culler, D., Roscoe, T.: A blueprint for introducing disruptive technology into the Internet. ACM Computer Communication Review 33(1) (2003) 59–64

4. Foster, I., Kesselman, C., Tsudik, G., Tuecke, S.: security architecture for computational grids. In: Proceedings of the Fifth ACM Conference on Computer and Communications Security (CCS'98), ACM Press (1998) 83–92

5. Open VPN. http://openvpn.net/

6. Papadopoulos, P.M., Katz, M.J., Bruno, G.: NPACI Rocks: Tools and techniques for easily deploying manageable Linux clusters. In: Proceedings of Cluster 2001: IEEE International Conference on Cluster Computing, IEEE Computer Society (2001)
258–267

7. Squid: Optimizing Web Delivery. http://www.squid-cache.org/

8. Kodama, Y., Kudoh, T., Takano, R., Sato, H., Tate be, O., Sekiguchi, S.: GNET-1: Gigabit Ethernet network test bed. In: Proceedings of Cluster 2004: IEEE International Conference on Cluster Computing, IEEE Computer Society (2004) 185–192

9. Nakada, H., Yokoi, T., Ebara, T., Tanimura, Y., Ogawa, H., Sekiguchi, S.: The design and implementation of a virtual cluster management system. In: Proceedings of the First IEEE/IFIP International Workshop on End-to-end Virtualization and Grid Management (EVGM2007), Multicon Verlag (2007) 185–192

10. Mr. Kailash Patidar[1] Prof. Ravindra Gupta[2] Prof. Gajendra Singh[3] Prof. Surendra mishra[4] The Design and Implementation of a Virtual Cluster Management System national. In: Proceedings of the First national conference in "BIST" Bhopal on End-to-end Virtualization.

[11] NISHIMURA, H., MARUYAMA, N., AND MATSUOKA, S. Virtual clusters on the fly - fast, scalable, and flexible installation. In Proceedings of the Seventh IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2007) (2007), IEEE ComputerSociety