Volume 4, Issue 4, April 2014 ISSN: 2277 128X

International Journal of Advanced Research in Computer Science and Software Engineering

Research Paper

Available online at: www.ijarcsse.com

Resource Management in Cloud Infrastructure

Jomol Thankachan, Mr.C.Anandaraj(ME)

CSE,Anna University Chennai, TamilNadu

India

***Abstract:** The recent emergence of public cloud offerings, surge computing -outsourcing tasks from an internal data center to a cloud provider in times of heavy load- has become more accessible to a wide range of consumers. Deciding which workloads to outsource to what cloud provider in such a setting, however, is far from trivial. The objective of this decision is to maximize the utilization of the internal data center and to minimize the cost of running the outsourced tasks in the cloud, while fulfilling the applications' quality of service constraints. We examine this optimization problem in a multi-provider hybrid cloud setting with deadline-constrained and pre emptible but non-provider-migratable workloads that are characterized by memory, CPU and data transmission requirements. Linear programming is a general technique to tackle such an optimization problem.*

***Keywords---** Federated infrastructures, resource allocation, resource mapping, virtualization, cloud computing, quality of service*

I. INTRODUCTION

Cloud computing builds upon advances on virtualization and distributed computing to support cost-efficient usage of computing resources, emphasizing on resource scalability and on demand services. Moving away from traditional data-center oriented models, distributed clouds extend over a loosely coupled federated substrate, offering enhanced communication and computational services to target end-users with quality of service (QoS) requirements, as dictated by the future Internet vision. Toward facilitating the efficient realization of such networked computing environments, computing and networking resources need to be jointly treated and optimized. This requires delivery of user-driven sets of virtual resources, dynamically allocated to actual substrate resources within networked clouds, creating the need to revisit resource mapping algorithms and tailor them to a composite virtual resource mapping problem.

In this paper, toward providing a unified resource allocation framework for networked clouds, they first formulate the optimal networked cloud mapping problem as a mixed integer programming (MIP) problem, indicating objectives related to cost efficiency of the resource mapping procedure, while abiding by user requests for QoS-aware virtual resources.

In this paper, we study the virtual resource allocation problem for networked cloud environments, incorporating heterogeneous substrate resources, and provide an appropriate approximation approach to address the problem. Specifically, for the node mapping phase, we provide a MIP problem formulation capable of taking into consideration QoS requirements.

II. PROPOSED SYSTEM

By providing a unified resource allocation framework for networked clouds, first we will formulate the optimal networked cloud mapping problem as a mixed integer programming (MIP) problem, by indicating objectives related to cost efficiency of the resource mapping procedure, at the time of abiding by user requests for QoS-aware virtual resources. The main advantages are,

- Allows for a flexible, structured, and comparative performance evaluation.
- The capacity of physical resources can be multiplexed among requested resources allowing us to accommodate more requests
- It adopts a heuristic methodology for resource allocation.

III. SYSTEM DETAILS

We implement a system architecture for solving this problem.

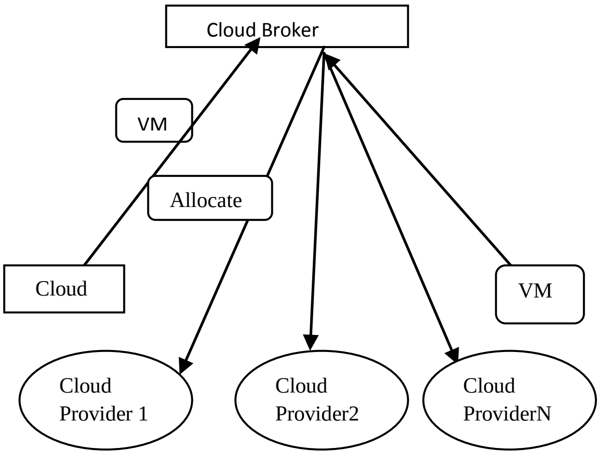

Fig. 1. System architecture

3.1Provision Provider

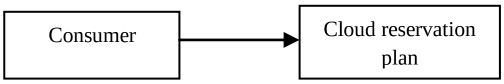

There are three provisioning phases: reservation, expending, and on-demand phases. These phases with their actions perform in different points of time (or events) as follows. First in the reservation phase, without knowing the consumer's actual demand, the cloud broker provisions resources with reservation plan in advance. In the expending phase, the price and demand are realized, and the reserved resources can be utilized. A cloud provider can offer the consumer two provisioning plans, i.e., reservation and/or on-demand plans. For planning, the cloud broker considers the reservation plans medium- to long-term planning, since the plan has to be subscribed in advance and the plan can significantly reduce the total provisioning cost. In contrast, the broker considers the on-demand plan as short term planning, since the on-demand plan can be purchased anytime for short period of time

3.2Node mapping

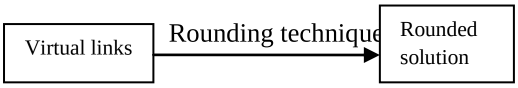

In this node mapping we are taking into consideration of virtual links as demands. Due to the nature of the MIP problem presented, the optimal fractional solution is computed for the problem's linear programming relaxation of the integer variables, which can provide a solution at least as good as the integer one. The relaxed problem can be solved by any suitable linear programming method, in polynomial time (e.g., CPLEX dual simplex routine). A rounding technique is applied to obtain the integer solution of the aforementioned relaxed MIP problem. Randomized rounding for LP relaxations was introduced for multi-commodity routing problems, where the fractional values contained in the optimal LP solution were treated as probabilities.

3.3Virtual link allocation

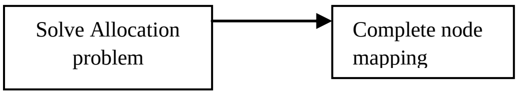

Once the aforementioned node mapping procedure has been successfully completed, link mapping is achieved by solving the multi-commodity flow allocation problem allowing traffic bifurcation. Alternatively, a shortest path algorithm can be applied to restrict each flow to a single path.

3.4Mixed Integer Programming

The proposed MIP algorithm will facilitate the adoption of cloud computing of the users as it can reduce the cost of using computing resource significantly. This MIP algorithm is proposed to minimize the total cost for provisioning resources in a certain time period. To make an optimal decision, the demand uncertainty from cloud consumer side and price uncertainty from cloud providers are taken into account to adjust the trade-off between on-demand and oversubscribed costs.

IV. Federated Infrastructures

Federated customers are those that often find themselves with too many cooks in the kitchen. In addition, business and IT may considerably differ in opinion about the value of integration, consolidation, and collaboration. As such, federated infrastructures push the limits of flexibility on Microsoft software, often negating best practices in favor of protecting fiefdoms, or more importantly, jobs. Federated infrastructures are quickly identified by their unnecessary complexity, customer pride in that complexity, split responsibility among groups that oppose one another, and a lack of service level agreements or accountability.

With such deployments, however, come problems and this blog allows me to share those experiences with you in hopes of reducing your own times-to-resolution.

V. Resource Allocation

Resource allocation is the distribution of resources – usually financial - among competing groups of people or programs. When we talk about allocation of funds for healthcare, we need to consider three distinct levels of decision-making.

Level 1: Allocating resources to healthcare versus other social needs.
Level 2: Allocating resources within the healthcare sector.
Level 3: Allocating resources among individual patients.

An Example of Resource Allocation

Let's consider an example: A community receives a gift of $100,000 from a wealthy donor to spend on healthcare, education and housing. The funds can be distributed among the three areas or dedicated to a single area, such as healthcare.

Level 1 : At this level, community members consider how to distribute the funds among one, two or three of the competing programs. For example, should the funding be split in three equal portions or should one program, possibly under-funded in the past, get all or most of the money?

Level 2 : Assuming that healthcare gets a portion of the $100,000, the next decision community members face is how best to direct the spending among competing healthcare interests. Should most or all of the funds go to hospital care and medical equipment? What about the public education program that promotes healthy lifestyles and behaviors (like exercise or immunizations) that prevent disease? Or, community members could decide to spend the money to purchase health insurance for those who can't afford it.

Level 3 : The next level of decision making involves distributing the financial resources among individuals. Most communities have policies and guidelines to insure fairness in these situations.

Once the first problem is solved, we are faced with the second problem, how to produce. There are several technically possible ways a commodity can be made. A basic criterion used in deciding the best technique is that producers should avoid inefficient methods. Production is said to be insufficient when it is possible to reallocate resources and, as a result, produce more of at least one good without producing less of any other good. This information is powerful and helpful to beginners.

VI. Resource Mapping

This decade has brought with it several disturbing trends: large numbers of high school drop-outs; high-profile incidents involving youth in violence; and projections of labor shortages coupled with the rising and significant need among employers for workers with specialized talents. As a result, there is a growing interest among communities to invest in youth. Communities, including schools, need strategies to improve working relationships that prepare young people as they transition to adulthood. This support encompasses strategies that ensure high academic achievement, completion of high school, postsecondary enrolment, and competitive employment. Single programs are not comprehensive enough to meet the needs of all youth. As a result, communities that want more effective and efficient ways of providing these services are finding ways to better connect individual agency resources into a single system for youth. Many communities, states, and governments are utilizing mapping strategies. Community Resource Mapping facilitates the identification, alignment, and leveraging of community resources to improve the educational, employment, and developmental needs of all youth, including youth with disabilities.

Although there is no common definition for mapping, it is, generally, a methodology used to link community resources with an agreed upon vision, organizational goals, strategies, or expected outcomes. There are several principles that are unique to mapping efforts. First, mapping strategies focus on what is already present in the community. The idea is to build on the strengths within a community. Second, mapping is relationship-driven. Key to mapping efforts is the development of partnerships--a group of equals with a common interest working together over a sustained period of time to accomplish common goals. Third, mapping embraces the notion that to realize vision and meet goals, a community may have to work across programmatic and geographic boundaries. These principles provide the foundation for the mapping process. For example, a community vision might be to improve post school results for youth with disabilities. Some goals within this vision might include the reduction of dropout rates, or an increase in the number of students pursuing postsecondary education. Mapping can address curriculum, financial resources, human resources, policies and legislation, state academic standards, and youth and adult services related to these goals.

VII. Cloud Computing

Cloudcomputing is the use of computing resources (hardware and software) that are delivered as a service over a network (typically the Internet). The name comes from the use of a cloud-shaped symbol as an abstraction for the complex infrastructure it contains in system diagrams. Cloud computing entrusts remote services with a user's data, software and computation. The cloud also focuses on maximizing the effectiveness of the shared resources. Cloud resources are usually not only shared by multiple users but are also dynamically re-allocated per demand. This can work for allocating resources to users. For example, a cloud computer facility, which serves European users during European business hours with a specific application (e.g. email) while the same resources are getting reallocated and serve North American users during North America's business hours with another application (e.g. web server). This approach should maximize the use of computing powers thus reducing environmental damage as well since less power, air conditioning, rack space, etc. is required for a variety of functions.

The term "moving to cloud" also refers to an organization moving away from a traditional CAPEX model (buy the dedicated hardware and depreciate it over a period of time) to the OPEX model (use a shared cloud infrastructure and pay as you use it).

Proponents claim that cloud computing allows companies to avoid upfront infrastructure costs, and focus on projects that differentiate their businesses instead of infrastructure. Proponents also claim that cloud computing allows enterprises to get their applications up and running faster, with improved manageability and less maintenance, and enables IT to more rapidly adjust resources to meet fluctuating and unpredictable business demand.

The business model, IT as a service (ITaaS), is used by in-house, enterprise IT organizations that offer any or all of the above services.

Using software as a service, users also rent application software and databases. The cloud providers manage the infrastructure and platforms on which the applications run. End users access cloud-based applications through a web browser or a light-weight desktop or mobile app while the business software and user's data are stored on servers at a remote location. Proponents claim that cloud computing allows enterprises to get their applications up and running faster, with improved manageability and less maintenance, and enables IT to more rapidly adjust resources to meet fluctuating and unpredictable business demand.

Characteristics

- **Agility** improves with users' ability to re-provision technological infrastructure resources.
- **Application programming interface** (API) accessibility to software that enables machines to interact with cloud software in the same way the user interface facilitates interaction between humans and computers. Cloud computing systems typically use REST-based APIs.
- **Cost** is claimed to be reduced and in a public cloud delivery model capital expenditure is converted to operational expenditure.

  This is purported to lower barriers to entry, as infrastructure is typically provided by a third-party and does not need to be purchased for one-time or infrequent intensive computing tasks. Pricing on a utility computing basis is fine-grained with usage-based options and fewer IT skills are required for implementation (in-house).The e-FISCAL project's state of the art repository contains several articles looking into cost aspects in more detail, most of them concluding that costs savings depend on the type of activities supported and the type of infrastructure available in-house.
- **Device and location independence** enable users to access systems using a web browser regardless of their location or what device they are using (e.g., PC, mobile phone). As infrastructure is off-site (typically provided by a third-party) and accessed via the Internet, users can connect from anywhere.
- **Virtualization** technology allows servers and storage devices to be shared and utilization be increased. Applications can be easily migrated from one physical server to another.

Multitenancy enables sharing of resources and costs across a large pool of users thus allowing for:

Centralization of infrastructure in locations with lower costs (such as real estate, electricity, etc.)`

Peak-load capacity increases (users need not engineer for highest possible load-levels)

Utilization and efficiency improvements for systems that are often only 10–20% utilized.

- **Reliability** is improved if multiple redundant sites are used, which makes well-designed cloud computing suitable for business continuity and disaster recovery.[30]
- **Scalability and elasticity** via dynamic ("on-demand") provisioning of resources on a fine-grained, self-service basis near real-time, without users having to engineer for peak loads.
- **Performance** is monitored, and consistent and loosely coupled architectures are constructed using web services as the system interface.
- **Security** could improve due to centralization of data, increased security-focused resources, etc., but concerns can persist about loss of control over certain sensitive data, and the lack of security for stored kernels. Security is often as good as or better than other traditional systems, in part because providers are able to devote resources to solving security issues that many customers cannot afford.
- **Maintenance** of cloud computing applications is easier, because they do not need to be installed on each user's computer and can be accessed from different places.

VIII. Virtualization

Virtualization (or **virtualisation**) is the creation of a virtual (rather than actual) version of something, such as a hardware platform, operating system (OS), storage device, or network resources.

While a **physical computer** in the classical sense is clearly a complete and actual machine, both subjectively (from the user's point of view) and objectively (from the hardware system administrator's point of view), a **virtual**

machine is subjectively a complete machine (or very close), but objectively merely a set of files and running programs on an actual, physical machine (which the user need not necessarily be aware of).

Virtualization can be viewed as part of an overall trend in enterprise IT that includes autonomic computing, a scenario in which the IT environment will be able to manage itself based on perceived activity, and utility computing, in which computer processing power is seen as a utility that clients can pay for only as needed.

The usual goal of virtualization is to centralize administrative tasks while improving scalability and overall hardware-resource utilization. With virtualization, several operating systems can be run in parallel on a single central processing unit (CPU). This parallelism tends to reduce overhead costs and differs from multitasking, which involves running several programs on the same OS.

Virtualization, in computing, refers the act of creating a virtual (rather than actual) version of something, including but not limited to a virtual hardware platform, operating system (OS), storage device, or network resources.

IX. QoS (Quality of Service)

QoS (Quality of Service) refers to a broad collection of networking technologies and techniques. The goal of QoS is to provide guarantees on the ability of a network to deliver predictable results. Elements of network performance within the scope of QoS often include availability (uptime), bandwidth (throughput), latency (delay), and error rate.

QoS involves prioritization of network traffic. QoS can be targeted at a network interface, toward a given server or router's performance, or in terms of specific applications. A network monitoring system must typically be deployed as part of QoS, to insure that networks are performing at the desired level.

QoS can be improved with traffic shaping techniques such as packet prioritization, application classification and queuing at congestion points. Using the Internet's Resource Reservation Protocol (RSVP), packets passing through a gateway host can be expedited based on policy and reservation criteria arranged in advance. Using ATM, which also lets a company or user preselect a level of quality in terms of service, QoS can be measured and guaranteed in terms of the average delay at a gateway, the variation in delay in a group of cells (cells are 53-byte transmission units), cell losses, and the transmission error rate.

Quality of Service (QoS) for networks is an industry-wide set of standards and mechanisms for ensuring high-quality performance for critical applications. By using QoS mechanisms, network administrators can use existing resources efficiently and ensure the required level of service without reactively expanding or over-provisioning their networks.

Traditionally, the concept of quality in networks meant that all network traffic was treated equally. The result was that all network traffic received the network's best effort, with no guarantees for reliability, delay, variation in delay, or other performance characteristics. With best-effort delivery service, however, a single bandwidth-intensive application can result in poor or unacceptable performance for all applications.

Quality of Service (QoS) parameters are a key factor in the roll-out of new technology. ETSI works on several QoS specifications and has been particularly active in Interoperability events on speech quality. QoS parameters are increasing in importance as networks become interconnected and a large number of operators and providers interact to deliver communications.

X. Resource optimization

Resource optimization is the set of processes and methods to match the available resources (human, machinery, financial) with the needs of the organization in order to achieve established goals. Optimization consists in achieving desired results within a set timeframe and budget with minimum usage of the resources themselves.

Intelligent Management sees the structure of an organization as a network of projects which cut across company "functions", in contrast with the hierarchical view of a company divided up into silos unable to recognize precise patterns and rules of interdependencies.

When we manage a company as a network of projects we must be able to allocate the resources available in the most efficient way possible, always bearing in mind that we have to achieve the global goal.

An efficient use of resources to carry out a project requires us to:

- Have a shared vision of the global goal to be achieved (remove unnecessary protection from individual tasks)
- Eliminate multitasking (increased effectiveness in the tasks)
- Identify the constraint (the **critical chain**) and protect it with a buffer of time (thus protecting the project from **variation**)
- Carefully manage the operational phases of the project (capitalize on time gained)
- Carry out a statistical analysis of the project buffer consumption using Statistical Process Control (**SPC**) (more effective project management)

11.Performance Evaluation

11.1Networked Cloud Mapping

Resource mapping determines the allocation of physical resources(substrate nodes,links,and paths) to the networked cloud request.Resource allocation does not change for the lease from the cloud provider while substrate resources are released upon request expiration.Request mapping is comprised of node assignment and link assignment.Specifically,node assignment is denoted as ;

$$M^N : N^V \rightarrow N^S$$

$$where\ M^N(n^V) \in V_a^S, n^V \in V_a^V \subseteq N^V.$$

u^S	$\mathbf{c}(u^S)$	$C(u^S)$
a'^S	$[8\ 16GB\ 50TB]$	$[8\ 16GB\ 50TB]$
b'^S	$[4\ 8GB\ 10TB]$	$[2\ 6GB\ 9.6TB]$
c'^S	$[8\ 16GB\ 50TB]$	$[8\ 16GB\ 50TB]$
d'^S	[15]	[14]
e'^S	[15]	[14]
f'^S	[15]	[14]

u^S	v^S	$bw(u^S, v^S)$	$BW(u^S, v^S)$
a'^S	d'^S	5 Gbps	5 Gbps
d'^S	e'^S	5 Gbps	4 Gbps
e'^S	b'^S	5 Gbps	4 Gbps
e'^S	f'^S	5 Gbps	4 Gbps
f'^S	c'^S	5 Gbps	5 Gbps

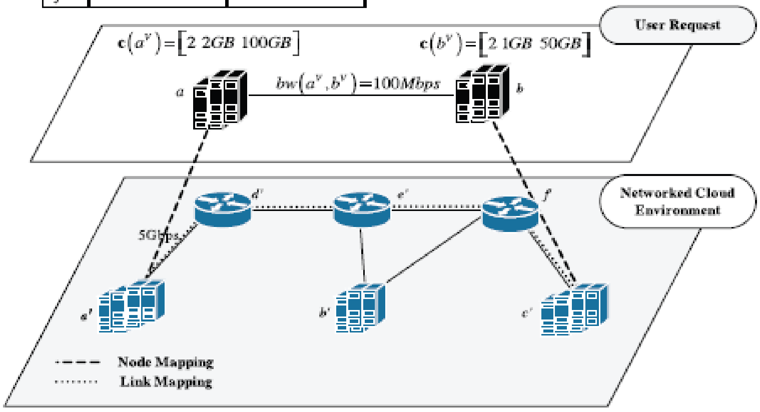

Fig. 1. Networked cloud environment and request mapping.

$$M^N(n^V) = M^N(m^V) \iff n^V \equiv m^V.$$

For a virtual node $n^V \in V_a^V$ to be mapped to substrate node $n^S \in V_a^S$ each requested capacity $i \in I$ must not exceed the remaining capacity C_i of the substrate node n^S. That is,

$$c_i(n^V) \leq C_i(n^S)$$

$$C_i(n^S) = c_i(n^S) - \sum_{\substack{\forall m^V, where \\ M^N(m^V)=n^S}} c_i(m^V).$$

Similarly, the bandwidth capacity of the virtual link is subject to,

$$bw(n^V, m^V) \leq \sum_{\forall P^S \in M^E(n^V, m^V)} bw(P^S)$$

$$bw(P^S) = \min_{\forall (u^S, v^S) \in P^S} BW(u^S, v^S)$$

$$BW(u^S, v^S) = bw(u^S, v^S) - \sum_{\substack{\forall (j^V, k^V), (u^S, v^S) \\ \in M^E(j^V, k^V)}} bw(j^V, k^V).$$

Effective resource optimization requires a certain rigor, consistency and agreement on processes. The following five steps can help you make the most of this optimization framework.

Five steps to resource optimization

Step 1: Define the objective to reflect organizational mission and strategy

The resource optimization model must reflect not only the well-defined, often narrow departmental objectives but also the objectives that are most important to the organization as a whole. There also needs to be an understanding of how activities will support these objectives, and how success or failure will be measured.

Step 2: Get executive buy-in and foster accountability

Putting the "best" choice for each decision variable into action requires accountability and commitment from implementers and executives.

Step 3 : Define the conceptual resource optimization model

To define the model, you first need to determine what input data is available. The cleaner and more accurate the data, the better. The more historical depth and relevance, the better. Next, identify variables that can actually be changed and decisions that can realistically be made in this organization within the given time frame.

Step 4. Formulate the resource optimization model

This step is the translation of your conceptual model into an analytic model with more rigor and detail, represented in mathematical terms.

Step 1 : Implement and update the model Using analytical software such as SAS, build and implement the model. Its output can provide recommendations as to the best values of the decision variables to support the objective, given the constraints and data available.

Figure 3: Workforce distribution optimization model

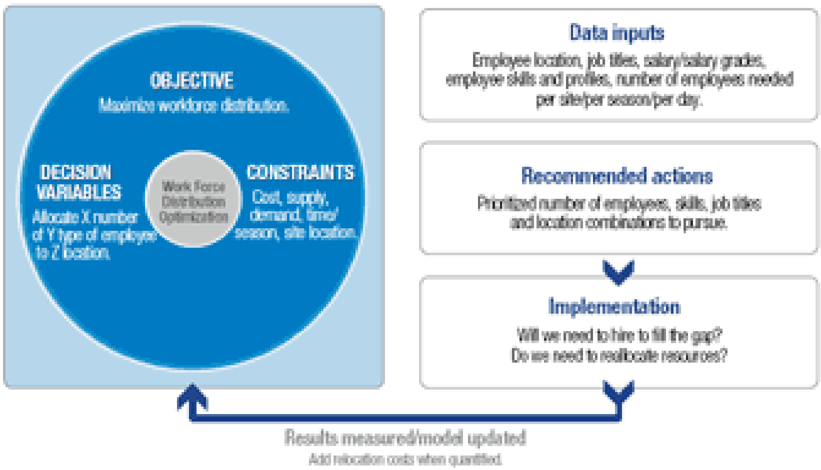

Figure 2: Marketing campaign optimization model

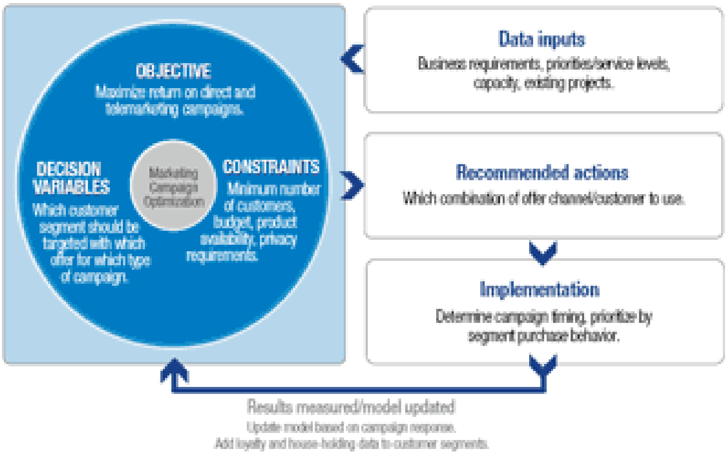

Test the optimization model for suitability. Commit to resource optimization Changing conditions will warrant corresponding changes in your resource optimization models. Periodically cycling through this five-step process will help organizations highlight areas to improve as they update their models to generate insights that continue to be relevant and valuable. A commitment to resource optimization will help to ensure that your organization remains focused and productive in an ever-changing competitive environment.

XI. CONCLUSION

In this paper, we study the virtual resource allocation problem for networked cloud environments, incorporating heterogeneous substrate resources, and provide an appropriate approximation approach to address the problem. Specifically, for the node mapping phase, we provide a MIP problem formulation capable of taking into consideration QoS requirements. Appropriate relaxation and application of a randomized rounding technique leads to a polynomial time solution. Following, link mapping is determined by solving the corresponding multi commodity flow problem. The proposed solution is compared against two well-known approaches on embedding virtual resource requests to a physical substrate. the proposed NCM approach overall outperforms other commonly applied algorithms. Specifically, NCM provides a tradeoff between G-SP and G-MCF in terms of acceptance ratio of NCM requests and number of hops on the substrate, per virtual link. At the same time, NCM manages to embed requests that generate more revenue, at a cost similar to G-MCF. An appropriate reconfiguration strategy has been also adopted to deal with the highly dynamic networked cloud environment.

REFERENCES

[1] R. Buyya, C.S. Yeo, and S. Venugopal, "Market- Oriented Cloud Computing: Vision, Hype, and Reality for Delivering it Services as Computing Utilities," Proc. IEEE Int'l Conf. High Performance Computing and Communications (HPCC '08), pp. 5-13, Sept. 2008.doi: 10.1109/HPCC.2008.172.

[2] "Creating Energy Efficient Data Centers," U.S. Dept. of Energy,May 2007.

[3] D. Breitgand, A. Epstein, and B. Rochwerger, “Resource Management Mechanisms to Support SLAs in IaaS Clouds” Achieving Federated and Self-Manageable Cloud Infrastructures: Theory and Practice, pp. 288- 307, IGI Global, 2012, doi:10.4018/978-1- 4666-1631-8.

[4] I.M. Lloriente, R.S. Montero, B. Sotomayor, B. Breitgand, and D. Maraschini, “On the Management of Virtual Machines for Cloud Infrastructures,” Cloud Computing: Principles and Paradigms,pp. 157-191, John Wiley & Sons,Jan.2011, doi:10.1145/1809049.1809052.

[5] X. Meng, C. Isci, J. Kephart, L. Zhang, E. Bouillet,andD.Pendarakis,“Efficient Resource Provisioning in Compute Clouds via VM Multiplexing,” Proc. Seventh Int’l Conf. Autonomic Computing,pp. 11-20, June 2010. doi:10.1145/1809049.1809052.

[6] I. Houidi, W. Louati, W.B. Ameur, and D. Zeghlache, “Virtual Network Provisioning AcrossMultipleSubstrateNetwork,”J. Computer Networks, vol. 55, no. 2, pp. 1011- 1023, 2011, dx.doi.org/10.1016/j.comnet.2010.12.011.

[7] M. Chowdhury, M.R. Rahman, and R. Boutaba, “ViNEYard: Virtual Network Embedding Algorithms With Coordinated Node and Link Mapping,” IEEE/ACM Trans. Networking, vol.20, no. 1,pp. 206-219, Feb. 2012, doi: 10.1109/TNET.2011.2159308.

[8] “The FEDERICA Project Website,” http://www.fp7-federica.eu, 2013.

[9] D. Andersen, “Theoretical Approaches To Node Assignment,”Unpublished Manuscript, http://www-2.cs.cmu.edu/~dga/papers/andersen-assign.ps, 2002.

[10] M. Yu, Y. Yi, J. Rexford, and M. Chiang, “Rethinking Virtual Network Embedding: Substrate Support for Path Splitting and Migration,” ACM SIGCOMM Computer Comm. Rev., vol.38, no. 2,pp. 17-29, Apr. 2008, doi:10.1145/1355734.1355737.

[11] W. Szeto, Y. Iraqi, and R. Boutaba, “A Multi-Commodity Flow Based Approach to Virtual Network Resource Allocation,” Proc.IEEE GLOBECOM ’03, vol. 6, pp. 3004-3008,Dec.2003, doi:10.1109/GLOCOM.2003.1258787.

[12] J. Fan and M.H. Ammar, “Dynamic Topology Configuration in Service Overlay Networks: A Study of Reconfiguration Policies,” Proc. IEEE INFOCOM ’06, pp. 1-12, Apr. 2006, doi:0.1109/INFOCOM.2006.139.

[13] J. Lu and J. Turner, “Efficient Mapping of Virtual Networks onto a Shared Substrate,” Technical Report WUCSE-2006-35, Washington Univ. St. Louis, 2006.

[14] Y. Zhu and M.H. Ammar, “Algorithms for Assigning Substrate Network Resources toVirtual Network Components,” Proc. IEEE INFOCOM ’06, pp. 1-12, Apr. 2006, doi:10.1109/INFOCOM. 2006.322.

[15] I. Houidi, W. Louati, and D. Zeghlache, “A Distributed Virtual Network Mapping Algorithm,” Proc. IEEE Int’l Conf. Comm.(ICC ’08), pp. 5634-5640, May 2008, doi:10.1109/ICC.2008.1056.

