CLOUD COMPUTING: NETWORKING AND COMMUNICATION CHALLENGES

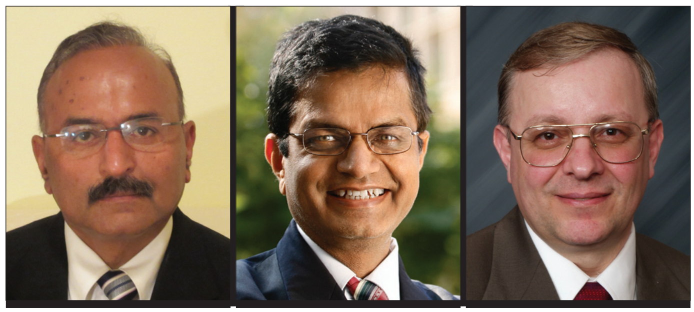

Amitabh Mishra *Raj Jain* *Arjan Durresi*

Over the past few years, cloud computing has rapidly emerged as a widely accepted computing paradigm built around core concepts such as on-demand computing resources, elastic scaling, elimination of up-front capital and operational expenses, and establishing a pay-as-you-go business model for computing and information technology services. With the widespread adoption of virtualization, service oriented architectures, and utility computing there has been a significant development in the creation of cloud support structures to deliver IT services within QoS bounds, service level agreements, and security and privacy requirements.

While cloud data centers are dominated by the server and infrastructure costs followed by networking and power, it turns out that networking and systems innovations are the key to the success of the cloud. The capital cost of networking gear for data centers is a significant portion of the cost of networking and is concentrated primarily in switches and routers, and load balancers. The remaining networking costs are concentrated in wide area networking, such as *peering*, *data center links*, and *regional back-haul* facilities needed to reach wide area network interconnection sites [1]. The value of the wide area network is shared across the data centers, and costs vary with industry dynamics (e.g., with tariffs), and are sensitive to site selection. Clever design of peering and transit strategies combined with optimal placement of micro and mega data centers therefore have a role to play in reducing network costs which can be further reduced by optimizing the network usage through better design of the services themselves, and better partitioning of their functionality. For example, with micro centers built out close to users, the latency of response can be reduced, but under the threat of substantial increases in wide area network costs. Networking has a role in data partitioning and replication, which requires better methods for design and management of traffic across the network of data centers, as well as better algorithms to map users to data centers.

Significant advancements in virtualization technologies have led to the development of large clouds that exist today. Virtualization creates several networking challenges that arise at the data link (layer 2) and network (layer 3) layers, which must be overcome by the networking gear (switches and routers) used in creating cloud computing infrastructure. Cloud computing networks must contend with a large number of attached devices consisting of physical and virtual devices, a large number of independent subnetworks, collocated software components belonging to different applications, and automated creation, deletion, and live migration of virtual machines — all of which may possibly come from different vendors. In order to be able to create a true multivendor cloud infrastructure that supports resource pooling the standardization of components is of utmost importance.

In the Call for Papers for this feature topic we solicited contributions on "networking and communication challenges" that need to be addressed for the long-term success of cloud computing. We kept our scope wider to consider all aspects of the infrastructure, which included computing centers, data centers, the cloud network, and the supported end-user services. The challenges related to the architecture, performance, reliability, security, maintainability, and virtualization were all within the scope of this issue. From the numerous submissions we have accepted four manuscripts for this special issue.

The first article in this feature topic, "Connecting Through Clouds: Open Standards and Proprietary Protocols for Data Center Networking," discusses the standardization protocols for switches and routers used in clouds at layers 2 and 3. It also presents a few proprietary protocols used in equipment manufactured by commercial vendors. One such standardized protocol is Spanning Tree Protocol (STP), which is a layer 2 switching protocol that calculates a loop-free single-path tree structure for the entire network. STP, which works well in classical Ethernet, suffers from several limitations when deployed in the cloud such as:

- Reduction in aggregate bandwidth as a result of blocking of redundant paths.
- Scalability.
- Path isolation.
- Support for multiple applications — multiple tenancy
- The need to discover a new path if a node or a link fails on a given path adds latency of several seconds to minutes, causing disruptions to virtual machine migrations.

To overcome the limitations of STP, Multiple Spanning Tree Protocol (MSTP) and Link Aggregation Group (LAG, IEEE 802.3ad) protocols have been standardized. The extension of LAG is called multichassis link aggregation (MC-LAG), which creates loop-free topology, allows dual homing, and works with existing management and multicast protocols. It has been extensively deployed. Equal Cost Multi-Pathing (ECMP) is also a standardized layer 3 protocol, which can be adapted to cloud computing due to its ability to create multiple load balanced paths that can provide variable bandwidths depending on the needs of the applications. However, one of the limitations of ECMP related to the live migration of virtual machines (VMs) has been fixed in other standardized protocols such as Transport Interconnection of Lots of Links (TRILL) and Shortest Path Bridging (SPB).

TRILL is an Internet Engineering Task Force (IETF) standard based on link state routing between routing bridges (RBridges) that computes shortest paths on a hop-by-hop basis between switches. TRILL offers several benefits such as scalability, discovery of loop-free multiple paths, and efficient bandwidth utilization. SPB is a layer 2 standard (IEEE 802.1ad) that fulfills the same requirements as TRILL but uses a slightly different approach to provide simpler data center virtualization by separating the connectivity services layer from the physical layer, and making endpoints fully aware of the entire path. OpenFlow is a new industry standard that achieves virtualization using software defined networking (SDN). It supports features such as packet flows, topology change, QoS, firewalls, statistical analysis of data streams, and network management. Besides the standardized protocols, this article also discusses a few proprietary protocols that are vendor specific. For example, Cisco uses a layer two protocol called FabricPath which is TRILL like but has a better performance when it comes to creation of a large number of virtual machines and MAC addresses. Similarly, Virtual Cluster Switching (VCS) and QFabric are used by Brocade and Juniper Networks, respectively.

Without virtualization it is hard to imagine if cloud computing could have emerged as a new computing paradigm. Virtualization hinges on the management of virtual machines (VMs) to achieve optimum utilization of cloud resources. The second article of this special issue, "Dynamic Resource Management Using Virtual Machine Migration," discusses the importance of VM migration in clouds. Virtual machine migration is needed in clouds to:
- Conserve energy, memory, and bandwidth resources
- Solve load imbalance among servers to meet application performance requirements
- Achieve server consolidation, which advocates use of a few servers with higher loads supporting a workload than on many lightly loaded servers

Consolidation also reduces server sprawl. One of the beauties of VMs is that they can be migrated from one server to the other live even while executing applications, and thus can mitigate server hotspots as and when they arise.

A cloud is a very dynamic resource center where new VMs are created as jobs arrive, and old machines are removed as jobs complete their demands. Management of resources therefore needs to be done dynamically using heuristics-based algorithms that determine the right time to migrate a VM to a specific server such that certain thresholds are satisfied. The thresholds for live VM migration, for example, can include factors such as reductions in:
- Frequency of migrations
- The cost of migration
- Energy consumption
- Service level agreement violations
- Resource utilizations

The article describes in sufficient detail the current heuristics-based VM migration algorithms that have been proposed in the literature to achieve server consolidation, load balancing, and hotspot mitigation.

The third article of this feature topic, "Scheduling in Hybrid Clouds," deals with the problem of scheduling in hybrid clouds and surveys some of the schedulers that have appeared in the literature. In addition, the authors also examine the impact of available bandwidth between a private cloud and the public cloud on the performance of a few scheduling algorithms to decide which client tasks should be executed on the public cloud. Performance results computed using three different schedulers confirm the importance of available bandwidth between the private and public clouds. The fourth article of this special issue, "Toward Cloud Ready Transport Network," discusses the need for evolving the existing transport networks to provide high-bandwidth connectivity to data centers on demand. To be able to provision such services, the authors propose that the transport networks must be capable of supporting automatic network configuration, adaptive bandwidth allocation, and multilayer network control driven by cross-layer optimization of available resources. The article suggests that, fortunately, several standard interfaces for control/management plane already exist. These interfaces can facilitate cross-layer optimization to improve the utilization of network resources, and reduction in infrastructure costs.

We would like to thank all authors who submitted manuscripts to this special issue and the reviewers for their wisdom in helping us select the four articles that are included here. This issue would not have been possible without the support of Editor-in-Chief Dr. Steve Gorshe, and assistance from the publication staff, particularly Joseph Milizzo and Jennifer Porcello.

REFERENCES

[1] A. Greenberg *et al.*, "The Cost of a Cloud: Research Problems in Data Center Networks," *ACM SIGCOMM Comp. Commun. Rev.*, vol. 39, no. 1, Jan. 2009, pp. 68–73.

BIOGRAPHIES

ARJAN DURRESI [SM] (durresi@acm.org) is currently an associate professor of computer and information science at Indiana University/Purdue University Indianapolis. He was a senior software analyst at Telesoft Inc. in Rome, Italy; then a research scientist of computer and information sciences at Ohio State University, Columbus; and an assistant professor of computer science at Louisiana State University in Baton Rouge. He has authored over 70 articles in journal and more than 160 articles in conferences in the fields of networking and security. He is the founder of several international workshops, including Heterogeneous Wireless Networks in 2005, Advances in Information Security in 2007, Bio and Intelligent Computing in 2008, and Trustworthy Computing in 2010.

RAJ JAIN [F] is a Fellow of ACM, a winner of the ACM SIGCOMM Test of Time award, CDAC-ACCS Foundation Award 2009, Hind Rattan 2011 award, and ranks among the top 70 in Citeseer's list of Most Cited Authors in Computer Science. He is currently a professor of computer science and engineering at Washington University in St. Louis, Missouri. Previously, he was one of the co-founders of Nayna Networks, Inc., a next generation telecommunications systems company in San Jose, California. He was a senior consulting engineer at Digital Equipment Corporation in Littleton, Massachusetts, and then a professor of computer and information sciences at Ohio State University in Columbus. He is the author of *The Art of Computer Systems Performance Analysis*, which won the 1991 Best-Advanced How-to Book, Systems award from the Computer Press Association. Further information about him including all his publications can be found at http://www.cse.wustl.edu/~jain/index.html.

AMITABH MISHRA [SM] (amitabh@cs.jhu.edu) is a faculty in the Computer Science Department at Johns Hopkins University in Baltimore, Maryland. His current research is in the area of cloud computing, data analytics, dynamic spectrum management, and data network security and forensics. In the past he has worked on the cross-layer design optimization of sensor networking protocols, media access control algorithms for cellular ad hoc interworking, systems for critical infrastructure protection, and intrusion detection in mobile ad hoc networks. His research has been sponsored by NSA, DARPA, NSF, NASA, Raytheon, BAE, and the U.S. Army. In the past, he was a member of technical staff with Lucent Technologies — Bell Laboratories in Naperville, Illinois, where his focus was on architecture and performance of communication applications running on the 5ESS switch. GPRS, CDMA2000, and UMTS were a few of the areas he worked on while with Bell Laboratories. He obtained his M.Eng. and Ph.D. in 1982 and 1985 in electrical engineering from McGill University, and an M.S. in computer science in 1996 from the University of Illinois at Urbana-Champaign. He is the author of *Security and Quality of Service in Wireless Ad Hoc Networks* (Cambridge University Press, 2007) and a Technical Editor of *IEEE Communications Magazine*.