

# VMLab: Infrastructure to Support Desktop Virtualization Experiments for Research and Education

**Prasad Calyam**

The Ohio State University  
[pcalyam@oar.net](mailto:pcalyam@oar.net)

**Alex Berryman**

The Ohio State University  
[berryman@oar.net](mailto:berryman@oar.net)

**Albert Lai**

The Ohio State University  
[albert.lai@osumc.edu](mailto:albert.lai@osumc.edu)

**Matthew Honigford**

VMware  
[mhonigford@vmware.com](mailto:mhonigford@vmware.com)

## Abstract

In terms of convenience and cost-savings, user communities have benefited from transitioning to virtual desktop clouds (VDCs) that are accessible via thin-clients, moving away from dedicated hardware and software in “traditional desktops”. Allocating and managing VDC resources in a scalable and cost-effective manner poses unique challenges to cloud service providers. User workload profiles in VDCs are bursty, such as in daily desktop startup, or when a user switches between text and graphics-intensive applications. Also, the user quality of experience (QoE) of thin-clients is highly sensitive to network health variations within the Internet.

To address the challenges associated with developing scalable VDCs with satisfactory thin-client user QoE, we developed a “VMLab” infrastructure for supporting desktop virtualization experiments in research and educational user communities. This paper describes our efforts in using VMLab infrastructure to support the following:

- Desktop virtualization sandboxes for system administrators and educators
- Research and development activities relating to VDC resource allocation and thin-client performance benchmarking
- Virtual desktops for classroom lab user trials involving faculty and students
- Evaluation of the feasibility to deploy computationally intensive interactive applications in virtual desktops, such as remote volume visualization
- Educational laboratory course curriculum development involving desktop virtualization exercises

## I. Motivation and Significance

Today, common user applications such as email, photos, videos, and file storage are supported at Internet-scale by cloud platforms, including HP Cloud Assure, Google Mail, and Amazon S3. Even academia increasingly is adopting cloud infrastructures and related research themes to support scientific research and education communities, such as the National Science Foundation Cluster Explorer (NSF CluE) and the Department of Energy’s (DOE) Magellan project. The next frontier for these user communities is to transition traditional distributed desktops with dedicated hardware and software installations into virtual desktop clouds (VDCs) that are accessible via thin-clients.

Moreover, in the not so distant future, we can envisage home users signing up for virtual desktops (VDs) with a VDC service provider providing Desktop-as-a-Service (DaaS) as a utility. With such a utility service, a thin-client such as a settop box can be shipped to a residential user to access a VD in a manner similar to what we have today for other common computing and communication needs, such as VoIP and IPTV. The settop box can be connected to television monitors or computer monitors, and multiple residential users can have their own unique login through this box to personalized VDs.

The drivers for transitioning traditional desktops to VDCs are obvious in terms of user convenience and cost-savings:

- Easier management of desktop support in terms of operating system, application and security upgrades
- Reduction in the number of underutilized distributed desktops unnecessarily consuming power
- Wider access to applications and data by mobile users

Allocating and managing VDC resources in a scalable and cost-effective manner poses unique challenges for service providers. User workload profiles in VDCs are bursty, such as in daily desktop startup, or when a user switches between text and graphics-intensive applications. Also, the user quality of experience (QoE) of thin-clients is highly sensitive to network health variations within the Internet. Unfortunately, existing solutions focus mainly on managing server-side resources based on utility functions of CPU and memory loads [1–4] and do not consider network health and thin-client user QoE. There is surprisingly little work being done [5–6] on resource adaptation coupled with measurement of network health and user QoE. Investigations such as [6] and [7] highlight the need to incorporate network health and user QoE factors into VDC resource allocation decisions.

It is self-evident that any cloud platform's capability to support large user workloads is a function of both server-side desktop performance as well as remote user-perceived QoE. In other words, *a CSP can provision adequate CPU and memory resources to a VD in the cloud, but if the thin-client protocol configuration does not account for network health degradations and application context, the VD is unusable for the user.* Another real-world scenario that motivates intelligent resource allocation is the fact that: *CSPs today do not have frameworks and tools that can estimate how many concurrent VD requests can be handled on a given set of system and network resources within a data center such that resource utilization is maximized, and at worst, the minimum user QoE is guaranteed as negotiated in service-level agreements (SLAs).* Resource allocations without combined utility-directed information of system loads, network health, and thin-client user experience in VDC platforms inevitably results in costly guesswork and over-provisioning, even for as few as tens of users. Also, due to lack of tools to measure the user experience from the server-side of VDCs, management functions in VDCs, such as configuring thin-client protocol parameters, often are performed using guesswork, which in turn impacts user QoE.

## II. VMLab Infrastructure

### A. Resources and Setup

To address the research and development challenges in developing scalable VDCs with satisfactory thin-client user QoE, we developed a “VMLab” infrastructure [8] that supports desktop virtualization experiments for research and education user communities. Initially funded by the Ohio Board of Regents, VMLab now is supported by the VMware End-user Computing Group, VMware Academic Program, Dell Education Cloud Services, and the National Science Foundation (under award numbers NSF CNS-1050225 and NSF CNS-1205658).

In the current VMLab infrastructure, an IBM® BladeCenter® S Chassis acts as a VDC data center that can concurrently support up to approximately 50 VDs. The BladeCenter has two IBM HS22 Intel® blade servers each with two quad-core CPUs, 32 GB of RAM, and four network interface cards (NICs). The storage resource is approximately 9 TB of shared SAS storage. The client-side uses a mix of several physical thin-clients from IBM, HP, and Wyse.

A netem network emulator [18] on a Linux Kernel is used for laboratory experiments to introduce network latency and loss and constrain end-to-end available bandwidth between the client and server sides. The VMware View™ desktop virtualization solution is used primarily to provision resources and broker virtual desktops, and has prerequisites such as VMware vSphere®. A web portal (<http://vmlab.oar.net>) enables information sharing about VMLab resources, capabilities, and salient project results. The web portal also provides information for VMLab users to gain hands-on access to run desktop virtualization experiments in their own sandboxes.

Each sandbox in VMLab has a dedicated virtual network with a separate blade allocation, as well as storage and firewall resources. Resource provisioning is performed based on user requirements and experiment plans. VPNs are set up using the OpenVPN™ server [19] for WAN connections to avoid using public IP address for VDs, and to accept VPN connections from external IP addresses. A virtual pfSense® server is used as a firewall appliance to handle all traffic between VDs and the Internet. Firewall rules are set or modified to restrict access to certain ports and addresses based on user sandbox requirements. The hypervisor, Active Directory, web portal and other supporting infrastructure are hosted in individual virtual machines within the VMLab infrastructure.

For experimentation involving distributed, multi-data center VDCs with realistic settings, VMLab resources are augmented with additional data center and thin-client resources from the NSF-supported Global Environment for Network Innovations (GENI) [17] infrastructure. The GENI infrastructure is a federated cloud of system (Emulab [20], PlanetLab [21]) and network resources (Internet2® and National LambdaRail (NLR)) for controlled as well as real-world experiments. It also provides a sliceable Internet infrastructure with wide area network (WAN) programmability using OpenFlow technologies that enable the dynamic allocation and migration of virtual machines in experiment slices. A multi-domain test bed with extended VLAN connectivity is set up between two data centers: VMLab at The Ohio State University and Emulab at the University of Utah. Distributed GENI nodes located at several university campuses (Stanford University, Georgia Institute of Technology, University of Wisconsin, Rutgers University) are used as thin-client sites.

### B. Users and Activities

Over the last three years, the VMLab infrastructure has supported:

- Desktop virtualization sandboxes for system administrators and educators [8]
- Research and development activities relating to VDC resource allocation and thin-client performance benchmarking [9]–[13]
- Virtual desktops for classroom lab user trials involving faculty and students [14] [15]
- Evaluation of the feasibility to deploy computationally-intensive interactive applications, such as remote volume visualization, in virtual desktops [14] [16]
- Educational laboratory course curriculum development involving desktop virtualization exercises [13]

The desktop virtualization sandboxes were set up for several campus system administrators in Ohio, Michigan, and Texas for a variety of experiments involving VMware Virtual Desktop Infrastructure (VDI) technologies, web portal and electronic lab notebook staging, and thin-client video streaming performance testing. Educators in three departments (Dept. of Chemistry, Dept. of Industrial and Systems Engineering, Small Animal Imaging Shared Resource (SAISR)) at The Ohio State University (OSU), as well as in Polymer Ohio have experimented with VMLab resources to set up VDs for classroom labs. The classroom lab applications within VDs ranged from common applications such as Microsoft Word and Microsoft Windows Media Player to remote volume visualization applications, such as surgical simulation and polymer injection-flow modeling, that are computationally-intensive, have massive datasets, and are highly interactive in nature.

In addition, the Ohio Board of Regents CIO Advisory Board Members recently sponsored a VDPilot project, a feasibility study of a VDC for classroom labs. This VMLab related study leverages universities' pre-existing high-speed access to the OARnet network and to national networks such as Internet2 and NLR in order to assess the user QoE of accessing desktops remotely compared to physically going to a computing lab, as well as analyzing the challenges and cost savings due to shared resources amongst collaborating institutions.

Development of novel, dynamic VDC resource allocation schemes and an OpenFlow controller are ongoing, integrating them into a VDC-Sim simulator. The VDC-Sim can act as a *cloud resource broker* to “control and manage routing flows”, as well as “measure and monitor user QoE delivery” in a “Run Simulation” mode, or can actually interact with VDC components in a GENI slice in a “Run Experiment” mode. VDC-Sim is being adapted to develop a graduate course curriculum involving desktop virtualization exercises in VMLab-GENI infrastructures through collaboration with the Department of Computer Science at Purdue University. These exercises include a study of resource allocation schemes, and comparing and optimizing thin-client protocol performance.

Lastly, a project involving underserved communities is being led by researchers in the Department of Computer Science and Engineering at Ohio State University. Here, researchers are working to equip the Linden community around Columbus, Ohio with VDCs as part of their emerging STEM education and other critical community support programs that can benefit from virtual desktop access.

### III. Exemplar Use Cases, Experiments, and Results

#### A. Thin-client Performance Benchmarking Toolkit Development

Clearly, instrumentation and measurement are needed on the server and thin-client sides to gather performance data for making the best resource adaptations in VDC platforms around system loads, network health, and thin-client user QoE. VMLab resources are being used to develop a novel virtual desktop benchmarking toolkit, called VDBench, [9] to create application and user group profiles based on CPU, memory, and network bandwidth measurements.

The current VDBench prototype can measure user QoE of atomic and aggregate application tasks in terms of interactive response times or timeliness metrics such as application launch time, web page download time, “Save As...” task time. Tasks are executed via different thin-client protocol configurations, such as RDP, RGS, and PCoIP, under synthetic system loads and network health impairments. It uses the concept of “marker packets” to correlate thin-client user events with server-side resource performance events in packet captures. It also leverages measurements from built-in memory management techniques in VMware® ESX®, such as ballooning under heavy loads [22], and earlier research on slow-motion benchmarking of thin-client performance under varying network health conditions [23]. Figure 1 shows the VDBench Java client prototype. The software can run on Windows and Linux platforms, and has capabilities for NIC selection for test initiation as well as interactions with the benchmarking engine for reporting test results. Enhancements to the VDBench Java client prototype are ongoing, including the ability to install and configure the software to run on physical desktops or commercial Windows or Linux operating system embedded thin-clients.

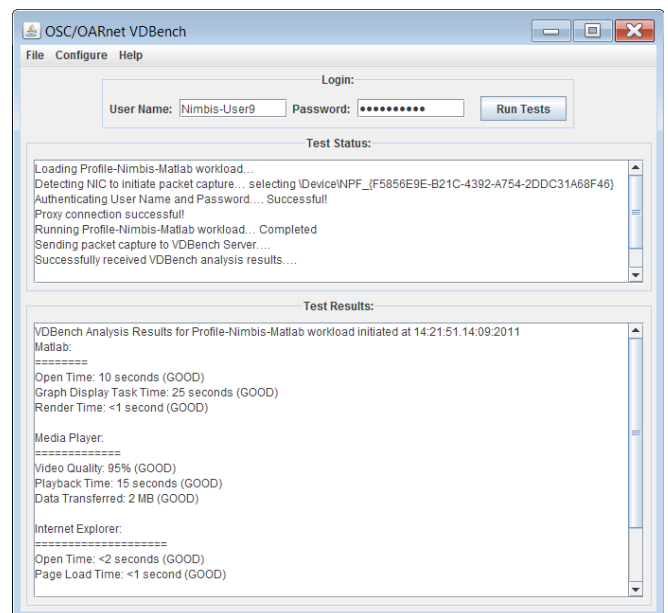


Figure 1. VDBench Java Client Prototype User Interface

#### B. Utility-directed Resource Allocation Scheme Development

In another set of salient research activities, application and user group profiles obtained through VDBench in VMLab are being used to develop utility-directed resource allocation schemes for VDCs at Internet scale. More specifically, we developed a utility-directed resource allocation model (U-RAM) [10] that uses offline benchmarking-based utility functions of system, network, and human components to dynamically (online) create and place VDs in resource pools at distributed data centers, while optimizing resource allocations along timeliness and coding efficiency quality dimensions. We showed how this resource allocation problem approximates to a binary integer problem whose solution is NP-hard.

To solve this problem, we proposed an iterative algorithm with fast convergence that uses combined utility-directed decision schemes based on Kuhn-Tucker optimality conditions [24]. The ultimate optimization objective was to allocate resources (CPU, memory, network bandwidth) to all VD such that the global utility is maximized under the constraint that each VD at least meets its minimum quality requirement along timeliness and coding efficiency dimensions.

To assess the VDC scalability that can be achieved by U-RAM provisioning, simulations were conducted to compare U-RAM performance with other resource allocation models:

- Fixed RAM (F-RAM), where each VD is over provisioned, something that is common in today's cloud platforms due to a lack of system and network awareness
- Network-aware RAM (N-RAM), where allocation is aware of required network resources yet over provisions system resources (RAM and CPU) due to a lack of system awareness information
- System-aware RAM (S-RAM), where allocation is the opposite of N-RAM
- Greedy RAM (G-RAM), where allocation is aware of system and network resource requirements based purely on conservative rule-of-thumb information rather than the objective profiling used by U-RAM

Several data center sites were considered, assuming each site had 64 GB of RAM, a 100 Mbps duplex network bandwidth interface, and a scalable number of 2 GHz CPU cores. Several factors were varied during the simulation runs, including the number of data center sites, the number of CPU cores at each site, and the type of desktop pools to which incoming VD requests belonged. The simulation results clearly showed that U-RAM outperforms other schemes by supporting more VDs per core and allowing a greater number of user connections to the VDC with satisfactory user QoE.

In addition, U-RAM and F-RAM are implemented in the VMLab-GENI test bed, as illustrated in Figure 2 [11]. Using Matlab-based animation of a horse point-cloud as the thin-client application, we demonstrated that U-RAM provides improved performance and increased scalability in comparison to F-RAM under realistic settings.

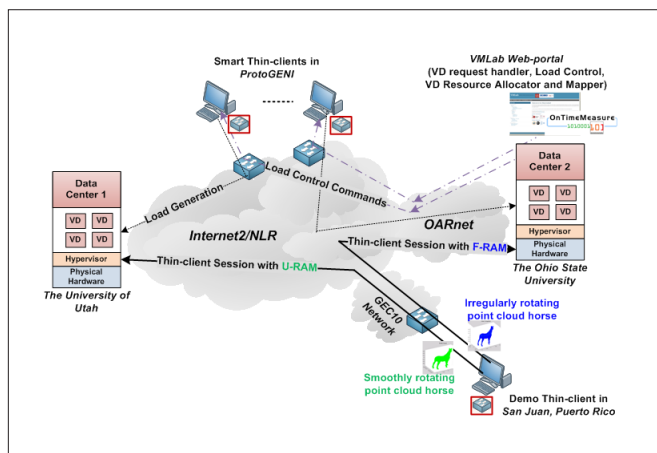


Figure 2. VMLab-GENI provisioning experiment to compare U-RAM and F-RAM schemes

The U-RAM work is being extended for provisioning VDs by investigating salient problems with subsequent placement of VDs across distributed data centers [12]. Placement decisions are influenced by session latency, load balancing, and operation cost constraints. In addition, placement decisions need to be changed over time for proactive defragmentation of resources for improved performance and scalability, as well as reactive VD migrations for increased resilience and sustained availability. Proactive defragmentation of resources is performed using global optimization schemes to overcome the resource fragmentation problem in VDCs that results from placements being done opportunistically to reduce user wait times for initial VD access. We refer to opportunistic placements as those that are performed using local schemes that use high-level information about resource status in data centers.

Over time, resource fragmentation due to careless packing of VDs on resources and changing application workloads leads to the “tetris effect” that decreases scalability (VDs per core) and performance (user QoE), thereby affecting the VDC Net-Utility. In contrast, reactive VD migrations are triggered by cyber attack or planned maintenance events, and should be performed in a manner that does not drastically affect the VDC Net-Utility. Not all VD migrations suggested by proactive or reactive schemes generate positive benefit in VDC Net-utility, since VD migration is an expensive and disruptive process. Therefore, we model the cost of migration and normalize it to utility of VDs, and migrate only the VDs (positive pairs) that generate positive Net-benefit in the VDC.

### C. VDPilot: Virtual Classroom Lab User Trials

Providing access to expensive, computational software such as Matlab® and SPSS has always been a logistical and licensing challenge for professors who want to train their students with industry-standard software. Although universities have labs with pre-licensed versions of the software available, lab access for some students is inconvenient. Furthermore, many students need pervasive access to the software and have trouble obtaining a license and installing the software correctly on their home computers. Professors who want to manage lab exercises, assignments, and exams use e-mail to send and receive large files, and are limited in their ability to access and assist in the work-in-progress of students.

To address these problems, the Ohio Board of Regents CIO Advisory Board Members recently commissioned a VDPilot feasibility study for hosting virtual desktops and shared storage for classroom labs within the Ohio-based university system. The study was initiated to investigate the use of federated shared infrastructure resources that would simplify classroom lab computing for faculty and students, and reduce costs for universities.

As part of the VDPilot study, VMLab was reconfigured to support subjective testing for approximately 50 faculty and students with secure remote access to lab software using thin-clients over the Internet [15]. User trials were conducted with professors and students, as well as some IT administrators, who were asked to compare going to a physical lab versus using the remote thin-clients while performing tasks in the virtual desktops using applications such as Microsoft Excel®, Matlab, SPSS, Windows Media Player, and Internet Explorer®.

Figure 3 shows the VDPilot survey (screenshot) that participants completed after following subjective testing instructions provided to them through the VDPilot web portal.

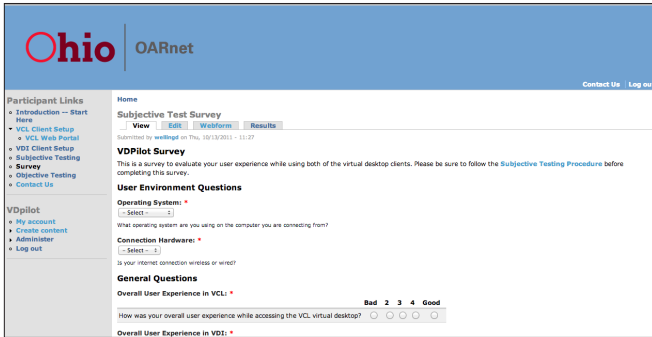


Figure 3. VDPilot survey participants completed after subjective testing

The survey results showed that 50 percent of participants found the virtual desktop user experience to be comparable to their home computer’s user experience, while 17 percent could not decide which user experience was better. Interestingly, 8 percent found the virtual desktop user experience to be better than their home computer user experience, particularly in the case of resource-intensive applications such as SPSS. Quotes recorded from faculty and students indicated they liked the virtual computer lab access in the pilot project in its current form. Two of the professors who participated in the study were eager to have their students use the pilot project test bed immediately as part their ongoing course offering. This confirms there is a real and current need for hosting virtual desktops and shared storage for classroom labs at universities.

#### D. Remote Interactive Volume Visualization for Researchers

VMLab resources have been used in experiments for the evaluation and support of a Remote Interactive Volume Visualization Infrastructure for Researchers (RIVVIR) [14] [16] to serve an increasing user base in the Small Animal Imaging Shared Resource (SAISR) at the Ohio State University and Polymer Ohio communities. RIVVIR provides an environment in which users can access VD that host computationally-intensive interactive applications and their related massive data sets.

Given the growing trend of users of data-intensive remote volume visualization applications that deal with gigabyte to petabyte sized data sets, it is impractical to carry or download these data sets and run computational analyses. Users of such applications in communities such as the ones supported by SAISR and Polymer Ohio inevitably must use VDs that have high-performance computing (HPC) capabilities at the data location and high-speed intermediate networks. Moreover, VDs allow a visual interpretation of large data sets, a powerful medium that can foster science and engineering innovations. Further, RIVVIR allows users to access their computationally-intensive interactive applications using handheld devices. They can jointly collaborate on the application steering with researchers at other institutions for visualization and analytics tasks related to their research and development efforts.

In our RIVVIR development efforts, we are interested in exploration that is beyond the classical thin-client model. We are exploring hybrid computing models and advanced multimedia stream content processing schemes, where execution is apportioned between thin-clients and the back-end server. For high-motion session output or computationally-intensive rendering such as 3D, we are investigating thin-client protocol optimizations that leverage increased server processing power to improve frame rate. The session switches to a general thin-client protocol configuration for low-motion video and other routine rendering.

In addition, we are evaluating the potential of caching repetitive video blocks, such as desktop backgrounds and menu items, to reduce server-side processing, bandwidth consumption, and interaction delays. Furthermore, there has been a trend lately in the use of thin-clients with high-resolution displays and significant computing power, especially with the latest generation of Apple iPads and similar products. To support these emerging thin-client platform applications, we are exploring related hybrid computing issues, where computing is distributed between the client and server ends, depending upon application context.

Figure 4 shows the RIVVIR configurations within VMLab for SAISR and Polymer Ohio community users for high-end demonstrations [14] [16]. A remote SAISR or Polymer Ohio user accesses a VD application similar to other users of typical applications, such as Microsoft Word or Internet Explorer, while their computationally-intensive interactive applications in the back-end rely on HPC infrastructure at the Ohio Supercomputer Center (OSC).

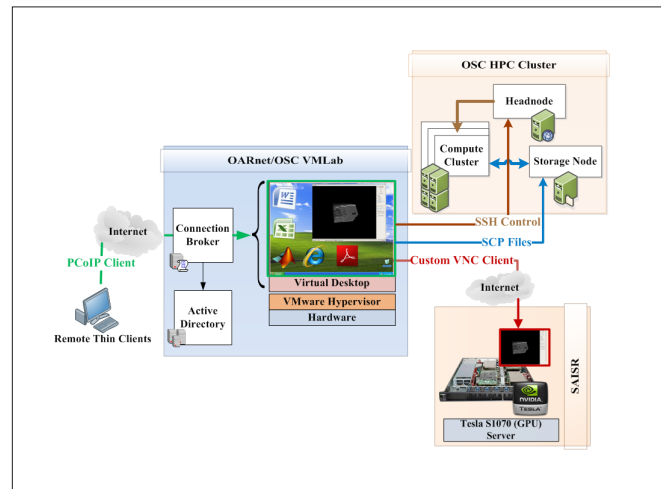


Figure 4. RIVVIR configurations for SAISR and Polymer Ohio community users

There are several challenges in configuring VDCs to support such applications due to their computational, storage, and network resource requirements for optimal user QoE. As in the case of SAISR users, custom applications need to be deployed that may rely on legacy thin-client protocols such as VNC, which are unsuitable for large delay or lossy networks that are common on the commercial Internet, as shown from our previous studies [25] [26]. We plan to

experiment with various reconfigurations of the enhanced VMLab infrastructure to study optimizations in thin-client protocols, such as tunneling and network-awareness, that can deliver satisfactory remote volume visualization user QoE. In the case of tunneling experiments, working through campus firewalls is a challenge. Additionally, we plan to address challenges in tunneling legacy protocols within the latest thin-client protocols, such as PCoIP to support remote users with large delay or lossy networks between thin-clients and servers.

## IV. Conclusion

Our VMLab deployment efforts have provided valuable operations experience to support a wide variety of research and education use cases relating to desktop virtualization experimentation. These experiences are helping significantly the engineering and operations of production services for desktop virtualization at the Ohio Supercomputer Center and OARnet, and new service models for research and education are being developed for user communities. Developers of the GENI community within infrastructure groups and instrumentation and measurement services groups have developed new capabilities by supporting unique experiment requirements of our VDC Future Internet application. We expect future desktop virtualization experimenters to benefit from these advancements.

The number of researchers and educators desiring to use VMLab resources is increasing rapidly, and several new initiatives are looking to use VMLab to scale to a large number of actual users in classroom labs and underserved communities. Our ongoing research and development projects on VDC resource allocation and thin-client performance benchmarking are ramping up for more extensive experiments. As a result, the VMLab infrastructure is being enhanced to support the concurrent provisioning up to 150 VDs. In addition, more than 20 physical thin-clients are being deployed. These units can be shipped to end-user sites for VDC experiments with a diverse group of geographically distributed users. This will enable us to validate successful DaaS offerings, where service providers can observe performance and control the end-to-end components to consistently meet SLAs and deploy an economically viable service delivery model.

We believe research and development outcomes from VMLab enable the realization of the next frontier—one that transforms end-user computing capabilities in the future Internet and enables society to derive benefits from computer and network virtualization. As a result, VMLab infrastructure enhancements are focused on collecting valuable real-world data sets to gain a better understanding of workload profiles for diverse applications and user groups in VDCs. This is the kind of understanding, missing in today's research literature, that could, in turn, fuel the development of new frameworks and tools to allocate and manage resources for improved performance, increased scalability, and cost-effective VDCs.

## Acknowledgment

This material is based upon work supported by the VMware Academic Program and the National Science Foundation under award numbers CNS-1050225 and CNS-1205658. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of VMware or the National Science Foundation.

## References

- [1] D. Gmach, S. Krompass, A. Scholz, M. Wimmer, A. Kemper, "Adaptive Quality of Service Management for Enterprise Services", *ACM Transactions on the Web*, Vol. 2, No. 8, Pages 1-46, 2008.
- [2] P. Padala, K. Shin, et. al., "Adaptive Control of Virtualized Resources in Utility Computing Environments", *Proc. of the 2nd ACM SIGOPS/EuroSys*, 2007.
- [3] B. Urgaonkar, P. Shenoy, et. al., "Agile Dynamic Provisioning of Multi-Tier Internet Applications", *ACM Transactions on Autonomous and Adaptive Systems*, Vol. 3, No. 1, Pages 1-39, 2008.
- [4] H. Van, F. Tran, J. Menaud, "Autonomic Virtual Resource Management for Service Hosting Platforms", *Proc. of ICSE Workshop on Software Engineering Challenges of Cloud Computing*, 2009.
- [5] K. Beaty, A. Kochut, H. Shaikh, "Desktop to Cloud Transformation Planning", *Proc. of IEEE IPDPS*, 2009.
- [6] N. Zeldovich, R. Chandra, "Interactive Performance Measurement with VNCplay", *Proc. of USENIX Annual Technical Conference*, 2005.
- [7] J. Rhee, A. Kochut, K. Beaty, "DeskBench: Flexible Virtual Desktop Benchmarking Toolkit", *Proc. of Integrated Management (IM)*, 2009.
- [8] P. Calyam, A. Berryman, A. Lai, R. Ramnath, "VMLab Testbed for Desktop Virtualization to Support Research and Education", *MERIT Desktop Virtualization Summit*, 2010 - <http://vmlab.oar.net>.
- [9] A. Berryman, P. Calyam, A. Lai, M. Honigford, "VDBench: A Benchmarking Toolkit for Thin-client based Virtual Desktop Environments", *IEEE Conference on Cloud Computing Technology and Science (CloudCom)*, 2010.
- [10] P. Calyam, R. Patali, A. Berryman, A. Lai, R. Ramnath, "Utility-directed Resource Allocation in Virtual Desktop Clouds", *Elsevier Computer Networks Journal (COMNET)*, 2011.
- [11] P. Calyam, M. Sridharan, Y. Xiao, K. Zhu, A. Berryman, R. Patali, "Enabling Performance Intelligence for Application Adaptation in the Future Internet", *Journal of Communications and Networks (JCN)*, 2011.
- [12] M. Sridharan, P. Calyam, A. Venkataraman, A. Berryman, "Defragmentation of Resources in Virtual Desktop Clouds for Cost-Aware Utility-Optimal Allocation", *IEEE Conference on Utility and Cloud Computing (UCC)*, 2011.

- [13] P. Calyam, A. Venkataraman, A. Berryman, M. Faerman, "Experiences from Virtual Desktop Cloud Experiments in GENI", *GENI Research & Educational Experiment Workshop (GREE)*, 2012.
- [14] P. Calyam, D. Stredney, A. Lai, A. Berryman, K. Powell, "Using Desktop Virtualization to access advanced Educational Software", *Internet2/ESCC Joint Techs*, Columbus, OH, 2010.
- [15] P. Calyam, A. Berryman, D. Welling, S. Mohan, R. Ramnath, J. Ramnathan, "VDPilot: Feasibility Study of Hosting Virtual Desktops for Classroom Labs within a Federated University System", *International Journal of Cloud Computing*, 2012.
- [16] A. Lai, D. Stredney, P. Calyam, B. Hittle, T. Kerwin, D. Reed, K. Powell, "Remote Interactive Volume Visualization Infrastructure for Researchers", *AMIA Annual Symposium*, 2011.
- [17] Global Environment for Network Innovations - <http://www.geni.net>
- [18] The Linux Foundation Netem - <http://www.linuxfoundation.org/collaborate/workgroups/networking/netem>
- [19] OpenVPN - <http://openvpn.net/>
- [20] Emulab/ProtoGENI Infrastructure - <http://www.emulab.net>
- [21] PlanetLab- GENI - <http://groups.geni.net/geni/wiki/PlanetLab>
- [22] C. Waldspurger, "Memory Resource Management in VMware ESX Server", *ACM SIGOPS Operating Systems Review*, Vol. 36, Pages 181 - 194, 2002.
- [23] A. Lai, J. Nieh, "On The Performance Of Wide-Area Thin-Client Computing", *ACM Transactions on Computer Systems*, Vol. 24, No. 2, Pages 175-209, 2006.
- [24] R. Rajkumar, C. Lee, J. Lehoczky, D. Slewlorek, "A Resource Allocation Model for QoS Management", *Proc. of IEEE RTSS*, 1997.
- [25] P. Calyam, A. Kalash, A. Krishnamurthy, G. Renkes, "A Human-and-Network Aware Encoding Adaptation Scheme for Remote Desktop Access", *Proc. of IEEE MMSP*, 2009.
- [26] P. Calyam, A. Kalash, R. Gopalan, S. Gopalan, A. Krishnamurthy, "RICE: A Reliable and Efficient Remote Instrumentation Collaboration Environment", *Journal of Advances in Multimedia's special issue on Multimedia Immersive Technologies and Networking*, 2008.