I.PIs: Karsten Schwan, Ada Gavrilovska, Matt Wolf, Greg Eisenhauer, Georgia Institute of Technology, Atlanta GA
schwan@cc.gatech.edu, ada@cc.gatech.edu, mwolf@cc.gatech.edu, eisen@cc.gatech.edu

Our current research touches on the following topics, with topics of primary interest being 8. and 10:
1.Cloud Architectures and Systems
3.Data Portability, Consistency, and Management
8.Cloud Self-Monitoring and Autonomic Control
10.Green Clouds
11.Cloud Test-Beds

Cloud management (monitoring and control) is being investigated in joint efforts with HP Labs researchers, with Yahoo, Intel, and more recently, with VMware. One outcome is a software infrastructure that combines monitoring and analysis – monalytics – to perform both continuously and at scale. On that basis, we are also creating new resource management methods. Monalytics and active management are at the scale of 1000s of VMs in the Georgia Tech 'GreenIT' facility where we can investigate both the performance and the energy use implications of running large cloud applications. As a full member of the OpenCirrus cloud initiative, as part of which we will be one of the sites to offer remote cloud access to other OpenCirrus members, we can experiment with multi-data center settings, where we are also working with Engineering at Georgia Tech to construct campus-wide cloud systems used for teaching (we already have two such infrastructures on campus, one in CS, one in Engineering). Finally, we work with researchers at Emory Univ. to investigate data-intensive cloud applications for applications in cancer research. Some of our proposed research in cloud management is explained on Page 2 below. The remainder of this abstract describes past work and other ongoing efforts.

1. Monalytics is defined in ICAC 2010. Novel black box behavior detection methods realized in monalytics appear in NOMS 2010, HotCloud 2010, with additional papers under submission.
2. We have developed and evaluated novel resource management methods for a wide variety of cloud applications and systems software, including for the HDFS file system [Amur – SOCC 2010], in part in joint work with researchers at CMU and at Intel, Pgh.
3. Cloud resource management methods for cloud infrastructures must smoothly interact with management methods already used in the virtualization layer (e.g., DRS and DPM for VMware's ESXServer). Toward that end, we are also composing a larger scale cloud computing mix comprised of a number of popular benchmarks, in experimental work using VMware's vCloud. This work is joint with VMware.
4. We will provide management software for the OpenCirrus cloud infrastructure, in joint work with Intel and HP, and in that same context, we are working to characterize the energy usage of virtualized applications (e.g., VM power metering), with first results described in [Gavrilovska - Wiosca 2010], and we are working with faculty in Mechanical Engineering at Georgia Tech to improve data center energy usage for virtualized applications, with a first approach published in [Nathuji – HotPower 2008]. The data captured from experiments in the GreenIT data center will be placed into an open repository made available externally.
5. Research to improve virtualization technology uses the Xen hypervisor on modern multicore processors, e.g., to enable heterogeneity (considering platforms with both general purpose and say, graphics processors – in joint work with HP and in part funded by NSF; considering performance asymmetric platforms in joint work with Intel. We have improved platform-level I/O performance isolation, described in [Kesavan - SOCC 2010], with extensions to cluster systems ongoing.

II.Managing Large-Scale Systems – A JIT Approach to Understanding and Managing Utility Cloud Data Centers

Challenges in managing data center systems include (i) *dealing with their continuing increases in scale* in terms of the numbers of entities and resources being managed, (ii) *understanding and acting upon new application-level requirements* like efficient access to large data sets (e.g., ClickStream data with Hadoop applications) or operation across multiple data center partitions or sites, and (iii) *runtime management in the presence of increasingly dynamic application behavior* with new usage paradigms like cloud computing.

We propose research that takes a multi-pronged approach to making advances concerning (i)-(iii).

(i) Automated, online flexible monitoring and analysis - two avenues for dealing with scale: (1) automate the methods for acquiring and analyzing monitoring data, to continuously maintain a high level of understanding about current application and data center behavior, and (2) make methods flexible in terms of their ability to operate at dynamically determined scopes with consequently variable overheads, e.g., by at any one time, monitoring and evaluating only those system or application elements that currently require attention.

(ii) Requirements as application performance states – since applications/systems differ in their requirements and because those can change over time, many are actively managed using built-in methods typically not visible to the hypervisor. Such methods attempt to `move applications/systems toward desired performance states'. Our monitoring and management techniques will exploit this fact by (1) leveraging 'set points' that encode these states provided by applications/systems while at the same time, (2) using black box methods that attempt to determine such states or state changes using system- and hardware-level performance counters.

(iii) Continuous management over varying time scales - much like web companies that are creating streaming mapreduce implementations to analyze data in real-time, this will also be the case for hypervisor-level monitoring, analysis, and management. (1) Management methods will operate at different time scales, to match levels of dynamics while containing overheads. (2) Methods continually try to move systems into acceptable states rather than trying to obtain an 'optimal' state immediately.

Concrete use cases will drive the proposed research, using installations of VMWare's ESXServer and of Xen in small-scale data center systems at Georgia Tech (an older 3000 core set of machines, plus a new 1000 core set of machines, as well as remote machines accessible to us due to our OpenCirrus membership and our collaboration with a bio-medical research group at Emory University). Two concrete ongoing research tasks are described next.

I.We propose to extend VMware's vCloud infrastructure (in joint work with VMware) to develop a 'resource/host trading' model across different management domains (vClusters): (1) by extending our initial approach to move from a model focused on improving infrastructure utilization to a future model that also takes into account application requirements, (2) by considering other resources, like I/O, in addition to compute and memory. This work will also benefit from our links with other companies and groups that provide us with usage data (e.g., with Travelport in Atlanta, with Yahoo, with HP).

II.We will develop new VM migration methods that operate over time to continuously assess benefits derived from such active management in lieu of costs. One such method considers other data center resources, such as their limited bi-section bandwidth. Initial results show that it is possible to dynamically identify the 'cohorts' of VMs that must be managed. VM migration methods, overhead vs. benefit assessment, and continuous monitoring and analysis for these purposes are subjects of future proposed work. We will also consider thermal/cooling issues when identifying suitable VM placements and migrations, exploiting the GreenIT data center at Georgia Tech. `