

# gridcentric

## SIMPLIFYING CLUSTER MANAGEMENT

11/12/2009

Harness the power of cloud computing

GridCentric brings the future of cloud computing to your organization. The GridCentric platform enables instantaneous on-demand sharing of high-performance computing resources, simplifies the application view of your storage substrate and makes cluster administration a one-click task. With GridCentric, organizations reduce IT expenditures by achieving significantly higher utilization of their physical infrastructure, and improved sharing of their data assets.

# Simplifying cluster management

HARNESS THE POWER OF CLOUD COMPUTING


## VIRTUALIZATION IS KEY

Virtualization technology enables a powerful abstraction, the Virtual Machine (VM), which encapsulates an entire “computer” – including the operating system and all application code and data. By effectively separating the physical hardware from the software, VMs enable a more powerful and flexible administration model. Thanks to advanced techniques such as paravirtualization and strong support from hardware vendors, VMs are now able to achieve application performance on par with bare-metal execution.

## Secure Sharing and Serialization

Fundamental benefits of virtualization are secure isolation and resource multiplexing: two VMs can execute simultaneously on the same physical hardware just as securely as on separate hardware. This is in stark contrast to typical cluster or server installations based on common Operating Systems (OSs): applications are isolated through OS interfaces that enforce a weak notion of isolation, are mired with vulnerabilities, and are prone to leaking information or crashing applications. The stronger isolation provided by virtualization guarantees that Bob's VM cannot attack Alice's VM, Bob's application cannot crash the system along with Alice's application, and that if Bob demands a carefully tuned configuration it won't conflict with Alice's.

The VM abstraction also allows a running “computer” to be serialized: it may be frozen in time and its state saved to a file from which it can be unfrozen later. Running VMs may be quickly moved from one physical machine to another without requiring a shutdown.




## THE GRIDCENTRIC VISION

GridCentric is a software platform that brings the power of cloud computing to an organization's internal High Performance Computing (HPC) infrastructure. GridCentric uses virtualization to make the hardware fungible, breaking the binding of software to hardware resources. Software stacks, encapsulated as serialized VMs, can arbitrarily run on any machine. All of an organization's hardware resources can be combined into a pool, grid, or cloud upon which user applications can be arbitrarily “swapped in” on-demand. No machine is left out of the pool for special projects or because only specific machines have certain software and libraries available.

## Challenges of VM-based Clouds

Organizing physical resources as a pool and swapping in VMs on-demand presents unique challenges. A significant obstacle is the size of the VM's footprint for typical modern deployments – after all, VMs encapsulate an entire software stack, including the OS and a virtual disk. For instance, a Windows Vista VM with Office 2007 installed can easily top 16GBs; a fresh install of Linux (Fedora 10) requires at least 5GBs. Massive VM footprints are at odds with a fundamental bottleneck in any cluster – the network. In order to maximize application throughput, most network bandwidth must be reserved for application data, not swapping in VM images. This limits the utility of current virtualization-driven cloud/cluster/grid computing offerings that must move all the state for a VM before application execution can be resumed, potentially adding several minutes of latency. Moreover, the new VM must be explicitly programmed to coordinate with other worker VMs and accomplish distributed tasks as typically required for HPC workloads. As a result, significant efforts must be expended by researchers and developers to port their programs to current virtualization-driven cluster computing offerings.



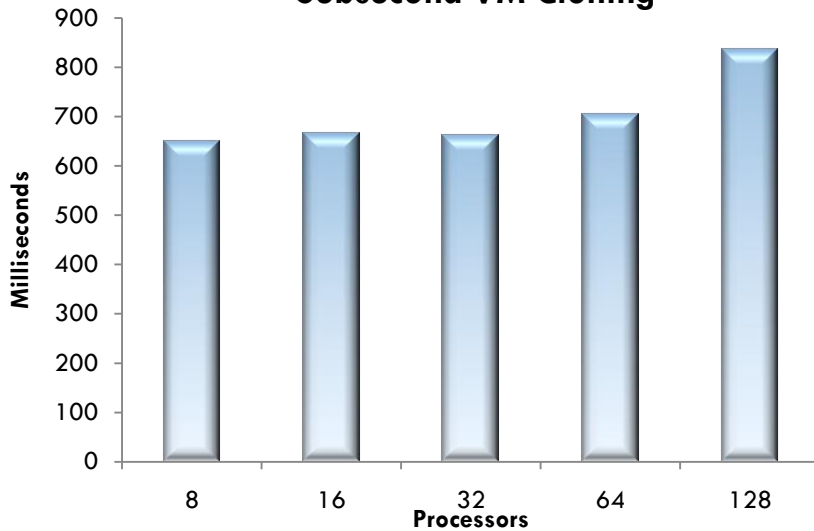
## THE GRIDCENTRIC SOLUTION

GridCentric enables an environment in which users can run completely arbitrary code, on-demand. A single application can rapidly expand its computing footprint to thousands of physical processors. This is done through a clone mechanism that allows a single VM to replicate itself into hundreds of instances *in mere seconds*. GridCentric performs this replication by efficiently propagating only the VM state that an application needs:

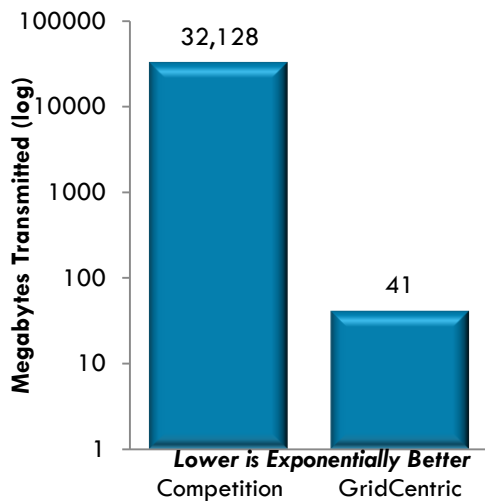
1. GridCentric sends VM state on-demand, ensuring the network only propagates state that is relevant. In contrast, transmitting a full VM will copy the entire OS, file system, and many system services that may not be needed when performing tasks like sequence alignment or stock pricing.
2. GridCentric multicasts VM state, leveraging the parallelism inherent in high-performance network hardware. For example, Gigabit Ethernet and Infiniband all provide native multicast facilities.
3. GridCentric avoids unnecessary propagation of state. Advanced heuristics detect memory used to store transient state like temporary I/O buffers, unused shared libraries, and unused allocation pools which can form a substantial component of a VM's footprint.

By combining these techniques, GridCentric can clone a running VM to hundreds of physical processors, spanning dozens of machines in a cluster, in less than a second. The cloned VMs can perform parallel HPC tasks such as rendering the frames of an animation, aligning short genomic reads, pricing stock options according to the Black-Scholes model, compiling a kernel, or building a sequence alignment tree.

### Subsecond VM Cloning

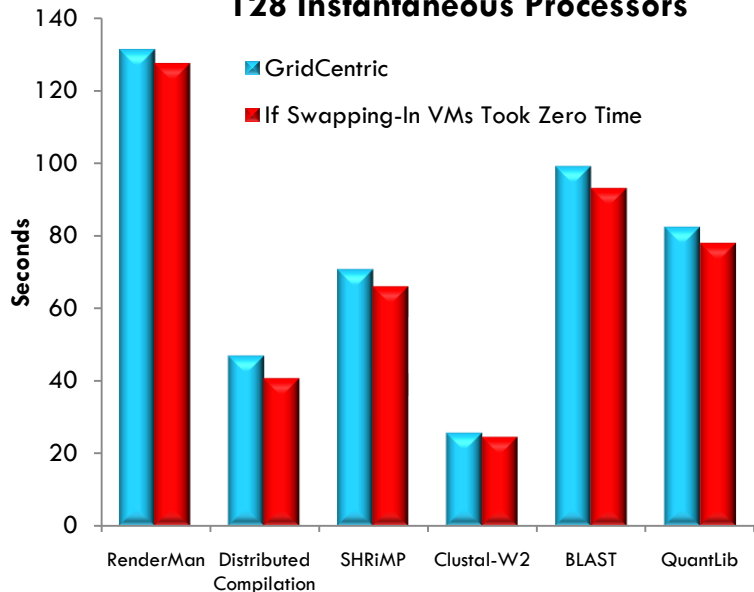


### Cloning a VM to 128 Processors Without Thrashing Your Network



The VMs performing these tasks are swapped-in by transmitting a few megabytes of state, three to four orders of magnitude more efficiently than any other VM-based cloud or cluster computing platform. While any of these platforms would add an overhead of several minutes to these HPC tasks, with GridCentric's unique technology you only pay a few seconds of overhead – in exchange for a whole new dimension in simplicity and cloud computing-based possibilities.

### 128 Instantaneous Processors



## Cloning VMs – a disruptive technology

GridCentric delivers technology that revolutionizes distributed and parallel cluster computing:

### An Intuitive Programming Model for Parallel Computing

The VM clone paradigm provides the ability to instantaneously clone a running VM into hundreds of identical replicas, thus allowing the user to transparently grow their allocation from a single VM to an entire cluster. This paradigm is rooted in decades of accumulated wisdom invested in UNIX programming patterns that use cloning facilities such as process forking or thread cloning. As a result, the VM clone model is a new, yet familiar and intuitive programming model for parallel computing.

### The Power of Thousands of Nodes with the Complexity of One

Cluster administrators no longer have to synchronize software stacks across hundreds of machines. The software is completely encapsulated in a VM image which is automatically replicated with the clones. As a result, security patches, configuration changes, and version updates need only be applied once to the original VM.

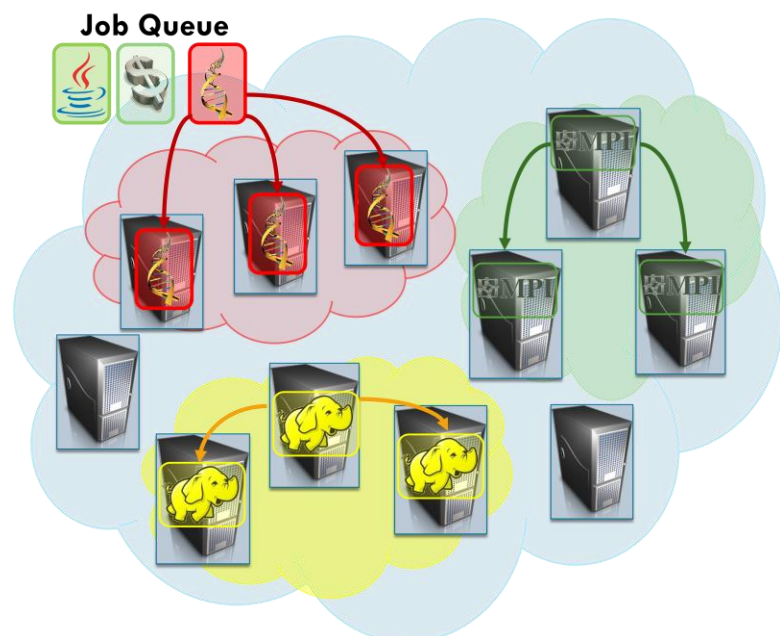
### Better Sharing of Pooled Resources

GridCentric enables practical encapsulation of jobs as VMs. Jobs can immediately span multiple physical hosts in the infrastructure, and execute leveraging fully-customized software environments. Once jobs have completed, all resources are immediately relinquished, and made available for other tasks. The same pool of physical machines is transparently shared, and leveraged with ease and agility.

## FEATURES FOR THE DATA CENTER

GridCentric integrates its technology with a number of staple cluster paradigms that simplify adoption.

**MPI:** GridCentric transparently supports the Message Passing Interface (MPI) middleware for parallel applications. The user only maintains a single VM with their unmodified MPI applications. An MPI job is started by cloning the VM into a virtual cluster spanning the desired number of processors. The MPI code then proceeds with execution, oblivious to any difference between the virtual cluster and a regular dedicated physical cluster. Upon completion of the job, the virtual cluster disappears.



**Grid Batch Queuing:** GridCentric transparently supports standard grid job scheduling. When a user queues a set of jobs in the scheduler, parallel or serial, their VM is cloned with the replicas running the jobs she placed in the queue. Their entire executing environment is transparently available to each job in the queue -- this environment could be radically different from that of fellow users. As the jobs finish executing, they place their results in a shared file-system, much like in a regular cluster setup, and the cloned VMs are gradually discarded, making room for the VMs encapsulating the jobs of other users.

**MapReduce:** Like the EC2 interface, Google's MapReduce data-parallel processing API is swiftly becoming a de facto standard, particularly through its open source incarnation Hadoop. GridCentric can transparently support MapReduce-like workflows, in similar fashion to its MPI support. GridCentric enhances MapReduce execution through its capability of immediately cloning the user VM to constitute the cluster of workers that will participate in the data-parallel computation.

## EXAMPLE USE CASES

**Day and Night:** In a financial institution the traders and the quantitative financiers have their own dedicated clusters and run different software. During the day, the traders' cluster performs automated trading and analysis. The cluster is sized with excess capacity to prevent peaks in activity, caused by a market event, from catastrophically degrading performance. The analysts pricing options, or the risk assessment group, have their own clusters. These clusters are lightly loaded during the day, yet crunching numbers at night while the traders' machines sit idle. Excess capacity in any cluster is rarely used, and unused cycles are never shared with the other groups. With GridCentric, all physical clusters can be unified for maximum potential peak availability, and virtual clusters with each group's configuration can be deployed and re-sized instantaneously to match the changing needs of the different groups. This allows the physical resources to be shared effectively by both groups; the institution may choose to either reduce the total computing resources or each group will have access to additional computing power.

**Open Data:** A typical genomics lab generates thousands of gigabytes of information from high-throughput sequencing devices. In order to perform analysis on a specific set of biological data, scientists have to download a large data set and acquire a cluster to crunch the numbers – something which is always impractical and often financially out of reach. With GridCentric, an organization can pool their resources to offer a cloud for internal and external consumption. Different data sets can be encapsulated in containers, with access controlled on a per-user basis. External users can be charged for time executing in the cloud and for data accessed. Users can install or develop vastly different applications performing analysis, assembly, annotation, or search. Using the MPI, grid batch queuing and Map-Reduce facilities, VMs can be dynamically and instantaneously cloned to hundreds of host to perform their parallel high-performance data intensive tasks. The users do not need to worry about setting up anything other than their private VMs. The infrastructure provisioning is as simple as defining data containers and policy, booting physical hosts, and taking care of the billing. Such a setup is currently difficult or impossible, but is capable of accelerating the speed of science and research.

## THE BEST OF BOTH WORLDS

GridCentric combines the benefits of virtualization and grid computing technologies. The cluster is simplified and virtualized without sacrificing performance.

	Virtualization	Grid Computing	GridCentric
Dynamic provisioning	✓		✓
Multi-group consolidation	✓		✓
Simultaneous multi-OS support	✓		✓
Live system migration	✓		✓
Transparent live snapshots	✓		✓
Enhanced power management	✓		✓
Instantaneous resource allocation		✓	✓
Scalable architecture		✓	✓
Cluster-wide resource accounting		✓	✓
Flexible allocation policies		✓	✓
Sub-second virtual cluster creation			✓
Multiple cluster harmonization			✓
Simple parallel programming APIs			✓

## CONTACT US

For more information, visit our website, send us an e-mail or give us a call.

web: <http://www.gridcentric.ca>

e-mail: [sales@gridcentric.ca](mailto:sales@gridcentric.ca)

phone: +1-888-365-GRID

