



All Nodes Are Not Created Equal — Thinking Differently About the Grid Nodes

Analysts: Anne MacFarland and Mike Kahn

Management Summary

Many believe that grid architectures don't make sense for commercial applications and that commercial SMP servers have no place in a grid architecture. Both presumptions are wrong, but you need to think differently about using grids for the wider commercial needs of the enterprise.

A typical grid deployment strives to deploy physical assets against workload and data retrieval needs both opportunistically (using whatever capacity is available) and optimally (in terms of resource utilization). It provides a way for an enterprise to use its assets more flexibly and more comprehensively. This is also exactly what commercial enterprises are seeking, but **grid deployments for commercial data processing need more capabilities than traditionally occur in academic and scientific grid computing environments.** Commercial enterprises may need to use the grid pervasively (on a continuing, rather than an as-available basis), to include some of their existing heterogeneous data center assets, and to set their own terms (for security, manageability, integration, and coordinated application functionality). **Using an architecture developed, largely, as a narrow-use paradigm to deploy a wide variety of concurrent and interrelated commercial applications presents many challenges.** For many of these workloads, parallelization may be impossible or does not help speed processing, and input/output may be more important to throughput than straight-out processing power.

It is tempting, and common, to think of a grid as a physical-level manifestation of PC-like nodes connected across networked distances. **It is better to think about the functionality that a grid can provide the enterprise, in terms of work to be processed and data to be delivered.** If you take this approach, you become less concerned about the physical characteristics of the servers and more concerned about having enough power in the grid to get the work accomplished or data delivered - in the required time frame, i.e., when it is needed. **This is an important shift in thinking - from a focus on sharing and utilization of physical assets to a focus on achieving multiple goals for work completion in a cost-effective manner.**

This reorientation clears the way to think about including - in the grid - larger symmetric multi-processor (SMP) servers that are the mainstay of enterprise data centers for many key applications. Typically, these large servers can be partitioned to run multiple workloads at the same time. This may be disconcerting to those who think of grid assets as a collection of distinctly-separate, singularly-identifiable processing engines - but partitions are a well-established, virtual way to deliver a variably-powered set of processing to an application. **So think about a grid node - at a higher level - as a vehicle for delivering a given amount of processing power because it makes no difference whether is it physically-defined or virtually-partitioned, except possibly in the total cost of ownership.**

The SMP server can bring to the grid a richness of capabilities, because it can be defined - on the fly - as one or more virtual nodes of varying capabilities, allowing it to execute many different kinds of workloads, flexibly and opportunistically. Deploying an SMP server in a grid as a virtual set of nodes (or mega-node) may be more cost-effective than the many small physical nodes that may be required to get all of the work done, when it needs to be done. Read on to understand the many ways a mega-node can work in an enterprise grid. **Remember, it's all about achieving collective results, in a most flexible way, at an advantageous cost.** It is not about perpetuating last decade's preconceived notions.

IN THIS ISSUE

- Expanding the Concept of the Node ... 2
- Fitting the Grid to the Enterprise..... 3
- Rethinking the Grid..... 4
- The Virtues of a Mega-Node 4
- Conclusion..... 5

Expanding the Concept of the Node

Servers, switches, appliances, and storage arrays have so many overlapping functionalities that it becomes productive to think of them all as nodes with particular functional attributes. Virtualization techniques allow nodes to be aggregated - or pooled - for simplified management via clustering (for servers), RAID (for storage) and alternate pathing (for switches and ports) and workload scheduling (workload managers, multiplexing, and grid, etc.). **Any vehicle that allows allocation of resources to a node without the node necessarily possessing identical physical characteristics is *virtualization*.**

Many larger physical nodes have a divisible capacity in the form of ports, zones, LUNs, and in the case of servers, partitions (via hardware or software, native or appended). The ability to disaggregate and re-aggregate at multiple levels turns a simple topography of physical assets into a much more flexible environment. To date, grids have not taken advantage of this flexibility.

Today's computational grids¹ are being used for massive tasks that can be subdivided and run in parallel in academic or scientific computing environments, using dedicated machines or idle-cycle harvesting schemes. While faster throughput is the goal, "fast" for these massive, typically long-running workloads is not the same as the need to instantly process a business transaction. Non-computational data grids have been developed to distribute content, particularly large files.² **Both computational and data grids typically involve dedicated homogeneous herds of similar small servers called "nodes".** (See Glossary on the next page.) **In these nodes, the level of granularity is the entire server.**

The ability to partition a server to run more than a single application is not usually considered in a grid environment. Either the

applications were large and subject to parallelization (think "scientific") and thus consumed all of the available processing power, or you just didn't care that some of the processing power wasn't being used, because it was too hard to consider or the node was deemed to be "cheap". Often, many of these nodes reside in the same rack or in adjacent racks. Think of a room filled with identical rack-mounted servers, with 10 or more per rack. Data centers are full of these today. **These nodes become a grid when they are networked together under some kind of grid management software.** These nodes can be distant as well. The network and grid manager make them look like they are a pool of servers, although distance can be a factor in performance.

More Node Varieties

With blade servers, the computing industry is moving toward a more optimized and manageable configuration that places many nodes in a single box.³ **These multi-node blade servers still deliver each blade separately to the grid pool.** This is in contrast to symmetric multi-processing (SMP) servers that can allocate more than one processor (and, usually, shared access to memory) to execute a large application. **Multi-node servers (whether blade or SMP) can also be seen as a *grid-in-a-box*,** with dozens or hundreds of processors working under the control of a grid manager that resides inside.

While it may be harder to visualize, consider a *mega-node* that consists of several of these multi-node servers within the same cabinet, usually by carving out or *partitioning* a larger SMP server into smaller *domains*, each capable of executing isolated work on these virtual multi-nodes. This is akin to having multiple grids in the same box, which we call a *multi-grid*.

¹ See *Computational Grids - Server Consolidation For A Distributed, On-Demand World* in **The Clipper Group Explorer** dated June 21, 2002 at <http://www.clipper.com/research/TCG2002021.pdf>.

² See *SGI Grids Accelerate Business Process* in **The Clipper Group Navigator** dated September 24, 2002, at <http://www.clipper.com/research/TCG2002034.pdf>.

³ See *IBM Blade Center - A Glimpse at the Future of Computing* in **The Clipper Group Navigator** at <http://www.clipper.com/research/TCG2002038.pdf> dated October 4, 2002.

Clipper's Glossary of Grid Terminology

Node – *An independent (think “stand alone”) computing system, capable of executing a single instance of an operating environment across one or more processors.*

Allocated or Virtual Node – *A slice of processing capacity defined without reference to its underlying hardware.*

Multi-node – *A set of independent nodes that exist in the same rack ore cabinet (think blade server- or SMP server) possibly with some shared architecture and/or management, which is transparent to each internal node*

- This multi-node could be parceled out as to different workloads – or it could work like a grid-in-a-box. As a grid-in-a-box, it could use indigenous partition flexibility and workload management to further optimize the immediate environment for the workload at hand.

Mega-node – *A computer system capable of hosting several-to-many multi-nodes simultaneously. It can be a multiprocessor with soft and hard partitions – or virtual machines that further fractionalize the partitions*

- A mega-node can be a grid itself and be attached to a grid, sort of a *grid of grids.*

Grid – *A collection of nodes interlinked by a network over which workloads are deployed in order to optimize access to available resources*

Virtual or logical grid – *The virtual nodes available on the Grid, regardless of their physical underpinnings*

Multi-grid – *A network of connected grids*

Grid manager – *A piece of software or set of protocols (like Globus) that facilitate aggregation, reservation, scheduling and use of grid computing resources*

Mixed grid – *A grid that contains different kinds of environments (servers and images). Think “variety of combinations of chip-based platforms and operating systems” like Intel/Windows, Intel/Linux, and SPARC/Solaris.*

So there are many shapes that a grid or multi-grid may take, once we open our minds to new thinking. **It is not about how the grid is physically implemented at the box or processor level, but how the application sees the server to which it is delivered for execution. It is not about the physical grid, which can have many manifestations, it is really about a logical or virtual grid.**

Fitting the Grid to the Enterprise

The benefits of grid architecture come at different levels. **Because of their inherently distributed governance, grids offer the opportunity to optimize opportunistically as unused assets become available.**

- At a physical level, grids can optimize the use of resources.
- At a logical level (say, in a data grid), they can optimize access to information.
- At an application level, they can provide a highly-available and flexible deployment infrastructure.
- At a data center level, grid protocols can link data centers functionally, giving another layer of resilience.

The grid infrastructure template is a very good fit with today's need to link processes together on multiple levels. Both are information processing structures created to deal with complex environments. Managing the environment as a whole (to the benefit of many), as well as the sum of its parts, is crucial to both grid and traditional data center operations. The challenge of bringing grid architecture to the commercial enterprise is partially due to a mismatch between the original “distributed” approach of grid philosophy and the traditionally optimized (i.e., integrated) workload engines of the enterprise data center. **Providing access to more (as in incremental processing in the grid) is only as valuable as the more can be deftly managed and used in the context of the enterprise. At the same time, grid protocols are a way to make single-purpose large servers more accessible and useful to the IT environment as a whole.**

The following factors complicate the situation.

- Enterprise environments have multiple platforms and multiple vintages of servers. Despite open standards, servers on different platforms have different ways of doing things, and are limited in their functional cooperation.
- Despite efforts to achieve simplicity, most enterprises still have a need to do a wide variety of tasks, sometimes accessing the same sources of data.
- Often, enterprise workloads cannot be split into parallel parts.
- The need for interactive access to “real-time” data has raised synchronization issues with the replication approach to pervasive access.
- Core applications may have different patterns of processing and pause, depending on which modules are used. These may be affected by other applications providing access via Web Services.
- Different industries have different bottlenecks, and companies have unique ways of dealing with them, which is often a key part of what makes them “tick” operationally.

The commercial enterprise needs a grid that can handle interactive workloads, and it needs to leverage the capabilities of its assets in place. This argues for a *mixed grid* of heterogeneous servers of different capacities – one that allows an enterprise to deploy workloads that need consistency across homogeneous server domains, but that also allows use of distinctive nodes for workloads congenial to them.

Rethinking the Grid

You need to switch your thinking from the physical assets and how much you pay for them, both at purchase and over time, to focus instead on the workloads and the levels of service (response time, process urgency) they require. By doing so, you get a way to manage IT that is much more in synch with the business process that is your mission to support. This is service-centric computing. It is not completely alien - the basic variables are still cost and time.

Make your inventory an inventory of functionality. With virtual machines technology, you can divide even uniprocessor boxes into several discrete virtual servers. With that in mind, rethink what kind of processing you need and how you may most economically achieve it.

With adequate bandwidth, what is the difference between 100 virtual nodes in a single box and 100 nodes of equal functionality on a grid? You have the same aggregation, service negotiation (though the points and tools of manageability may be different, and the costs may differ). There is little difference until you need more elasticity. Then a grid of partitions can take advantage of the ability to reallocate more processing capability, or memory space on the fly. For this you need a *mega-node*.

The Virtues of a Mega-node

The more dynamic your business processes, the more you will care about the elasticity and manageability that mega-nodes provide.

- **Capacity partitioning**, sometimes dynamic, for both hardware and software, with prioritization policies to arbitrate the reallocation, allows an enterprise to match resources with the task at hand.
- **Use of virtualization software** for sub-processor partitioning increases the granularity of processing that can be assigned to a workload.
- **The ability to schedule and manage workloads** within the node allows localized optimization of the different patterns of processor use or multiple workloads.
- **Clustering and internal diagnostics** can sense processor malfunction and reallocate workloads accordingly, delivering greater resiliency.
- **Open standards** compliance is necessary for communications and manageability.
- **Capacity on demand** satisfies the need for invocable additional capacity and provides a resilience beyond the elasticity of dynamic resource allocation. In automotive terms, it is the spare tire, not the steering and suspension.

It all comes back to time and money and the characteristics of workloads you want to run. You may have processes involving real-time data access, where latency hurts and nearness to recipient matters. There may be other processes or parts of processes that can be split into parts and run in parallel to diminish the elapsed time. There may be lower priority processes that can be multiplexed with other critical workloads, but only when there is spare capacity. You will want to take advantage of all of these workload characteristics in your flexible enterprise grid environment.

If you are trying to reduce the costs to your users for computing, you can no longer afford to think just about the physical assets. However, you must implement there, because your applications may care about the processor, the server firmware, the middleware, the resource management, and the capabilities and resilience of the operating system in order to get the most bang for the buck from each node. What physical servers you chose to deploy in your grid will depend on your need for responsiveness, your concerns about security, the kinds of workloads that need to be run, and, of course, your ability to buy new capital assets. **The on-the-fly flexibility characteristics of multi-node and mega-node servers to work as nodes in your grid may be the ultimate determinant, because every unused cycle becomes a source of loss that cannot be recovered.**

The benefits of and need for enterprise grids go far beyond the economics of capacity utilization. An opportunity is knocking on the door of the beleaguered enterprise. **By changing your perspective from *managing physical assets and running applications on them* to *managing and deploying workloads on your collection of physical assets*, you will give the enterprise the needed flexibility to be responsive to today's rapidly evolving markets.**

Conclusion

Grid architecture gives new tools for consolidation and cost control to the IT infrastructure managers for commercial enterprise. More importantly, it gives new ways to optimize the performance of the applications that support the enterprise. **The inclusion of the mega-node in a grid allows core enterprise applications to participate in and add to the benefits of a grid architecture.**



About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

➤ **The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.**

About the Authors

Anne MacFarland is Director of Enterprise Architectures and Infrastructure Solutions for The Clipper Group. Ms. MacFarland specializes in strategic business solutions offered by enterprise systems, software, and storage vendors, in trends in enterprise systems and networks, and in explaining these trends and the underlying technologies in simple business terms. She joined The Clipper Group after a long career in library systems, business archives, consulting, research, and freelance writing. Ms. MacFarland earned a Bachelor of Arts degree from Cornell University, where she was a College Scholar, and a Masters of Library Science from Southern Connecticut State University.

➤ **Reach Anne MacFarland via e-mail at Anne.MacFarland@clipper.com or at 781-235-0085 Ext. 28. (Dial "1-28" when you hear the automated attendant.)**

Mike Kahn is Managing Director and a cofounder of The Clipper Group. Mr. Kahn is a thirty-year veteran of the computer industry. For the vendor community, Mr. Kahn specializes on strategic marketing issues, especially for new and costly technologies and services, competitive analysis and sales support. For the end-user community, he focuses on mission-critical information management decisions. Prior positions held by Mr. Kahn include: at International Data Corporation — Director of the Competitive Resource Center, Director of Consulting for the Software Research Group, and Director of the Systems Integration Program; President of Power Factor Corporation, a Boston-based electronics firm; at Honeywell Bull — Director of International Marketing and Support; at Honeywell Information Systems — Director of Marketing and Director of Strategy, Technology and Research; with Arthur D. Little, Inc. — a consultant specializing in database management systems and information resource management; and, for Intel Corporation, Mr. Kahn served in a variety of field and home office marketing management positions. Earlier, he founded and managed PRISM Associates of Ann Arbor, Michigan, a systems consulting firm specializing in data management products and applications. Mr. Kahn also managed a relational DBMS development group at The University of Michigan where he earned B.S.E. and M.S.E. degrees in industrial engineering.

➤ **Reach Mike Kahn via e-mail at Mike.Kahn@clipper.com or via phone at 781-235-0085 Ext. 21. (Dial "1-21" when you hear the automated attendant.)**

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log and "**clipper.com**" are trademarks of The Clipper Group, Inc., and the clipper ship drawings, "*Navigating Information Technology Horizons*", and "teraproductivity" are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.