Modular Computing: The New Enterprise Computing Model

Industry
Cross Industry

Business Challenge
Making enterprise data centers more agile, easier to manage, and less costly

Technology Solution
The Egenera BladeFrame* System

Server Processing Power
Intel® Xeon™ Processor MP

SOLUTION ARCHITECTS

egenera™

intel®

IT's Challenge

In the past three years, the world has changed for Information Technology groups. In the late 1990s, the predominant problem was deploying equipment and software quickly enough to keep up with demand for computing. While the tech sector boomed on Wall Street, money was no object. IT budgets swelled and the numbers of computers in data centers grew exponentially.

Now, in the early 2000s, the picture is very different. IT budgets are flat to down, yet business demand for IT services continues to escalate. This combination of more demand and constrained budgets has compelled IT groups to consider new approaches to IT infrastructure, approaches that offer more flexibility and lower cost of ownership.

The common theme is cost cutting. In today's world, profits come less easily than in the 1990s. Competitors are more experienced, and competition is more intense. Corporations that trim costs while providing great service will prevail over those that can't.

IT plays a major role in this competitive situation. As competition becomes more intense, so does the pressure on IT to cut costs and boost contribution. Now more than ever, large corporations are using their computing assets as tools to pull ahead of the competition.

The January 13, 2003, issue of Time Magazine* provides a great example of how IT contributes in new ways. Executives at a big-box retailer were considering dropping a particular brand of chicken from the shelves because sales volume was poor. Then the retailer's data miners found that customers who bought that brand of chicken also bought large amounts of other merchandise. The chicken stayed.

Data mining, online transactions, and other new computing demands require collecting and processing enormous amounts of data. Still, IT departments are expected to keep up, even with budgets flat to down. The bottom line is that IT will be doing more with less.

Modular Computing can slash costs in IT infrastructure. It enables IT groups to consolidate equipment, conserving expensive real estate. It offers the opportunity to migrate applications from expensive proprietary platforms to more open, powerful, and manageable systems.

Doing More With Less

To keep up with computing demand while operating within restricted budgets, IT must find ways to optimally use computing resources and reduce people costs. There are many areas for improvement.

Solution Blueprint:
Modular Computing: The New Enterprise Computing Model

Solution Provider:
Egenera

Proof Points:
A survey of hundreds of executives in Global 3500 firms shows reduced spending in e-Business.[1] Of the executives surveyed, more than 70% forecast headcounts for internal IT, along with consulting and implementation services, being flat to down.

In a survey of 100 IT managers at multinational Fortune* 1000 companies, cost cutting is the top priority for 2003.[2]

Cost of Over-Provisioning

As data centers have moved toward servers and away from mainframes, IT has found that some mainframe capabilities weren't available on servers. A glaring example is that smaller servers were unable to rapidly obtain more processing power to accommodate peaks in computing demand.

As applications became more transactional, for example with customers entering information via the Web, these peaks in computing demand became more visible. During peak demand, customers saw their transactions slow down. In situations where these transactions affect the bottom line, as when customers enter purchases, prompt processing becomes vital to the business.

As the number of customers using Web services has increased, the peaks in computing demand became more intense and more frequent. Consequently, customers more frequently saw declines in performance.

Many data centers have ensured responsiveness to business requirements by over-provisioning—proactively sizing computing resources in anticipation of peak demand. In the world of traditional servers and legacy mainframes, over-provisioning makes sense. In fact, many advisory firms once recommended over-provisioning as a means of meeting peak demand.[3]

In the ideal, an alternative to over-provisioning is for IT to obtain additional resources and bring them online as they see demand increase. In practice, even after obtaining the hardware, setting it up and configuring the software can take weeks. Given the real-time nature of the changes in computing demand, deployment takes too long, so IT began relying on over-provisioning.

Over-provisioning has its own drawbacks. It leaves costly resources idle most of the time. CPU utilization in many large data centers ranges from 15 to 20 percent for non-mainframe servers, chiefly because of inability to rapidly reallocate unused resources during off-peak periods.[3] Too much capital is tied up in under-utilized resources.

To reduce capital costs, IT needs an alternative to over-provisioning—a means of reallocating resources in minutes rather than weeks to accommodate intense peaks in demand for an application.

Cost of High Availability

As transaction processing applications have become more common, more applications have been deemed mission critical—capable of severely affecting business when they slow down or stop running altogether. Hence the growing need for high availability.

However, high availability traditionally comes with a high price. Redundant equipment is expensive to buy, maintain, and manage. Additional software licenses, clustering software, and the professional services needed to implement a traditional configuration for high availability can cost more than the initial hardware. As a result, many IT groups continue to rely on expensive mainframes or RISC servers that use costly switched redundant connections to provide high availability.

Data centers need high availability, but they don't need added expense. They need high availability on equipment that costs less, eliminates the need for extra software and professional services, and automates management.

Cost of Too Many People Doing Low-Level Tasks

Labor is the largest expense associated with IT. According to Giga Information Group, labor represents 46% of IT budgets.[4] Finding ways to move administrators from low-level tasks to more-productive tasks would greatly improve an IT department's ROI.

Excessive Server Management

Consider a data center with 1000 application servers. Each class of server has its own management and provisioning process. To support these servers, IT needs experts for each class of server. In addition to their unique knowledge, these experts have many redundant skills.

If server management could be simplified, many of these experts could be shifted to tasks with higher ROI than managing servers.

Excessive Deployment Expense

Installing and configuring hardware and software takes much more administrator time than one would expect. According to Giga Information Group, "Management of most large collections of servers is a manually intensive process. Highly automated management of servers, particularly the deployment of applications and operating system images, is more the exception than the rule…."[3] Moving an application from one server to another is a delicate task requiring days for a skilled administrator.

Complex deployment also contributes to stranded resources, discussed below.

Excessive Cable Management

A full rack of traditional servers can need over 200 cables to provide the redundant connections necessary for high availability. Such large numbers of cables complicate cable management. Giga Information Group says that, in large data centers that have many reconfigurations, system administrators can spend up to 25 percent of their time managing cables.[3]

IT needs a means of spending less time on cables.

Cost of Stranded Resources

Closely related to over-provisioning is the dilemma that causes stranded resources. For example, suppose that demand for an application crests, then declines over a period of months. Three factors make data center management reluctant to harvest computing resources associated with the application:

- The cost of the administrative time spent removing the resource from the first application and reconfiguring it for the second.
- The risk of destabilizing the declining application. Removing some of the hardware used to process an application is complex. Without extreme attention to detail, it's possible to cause the application to fail.
- The possibility that demand for the declining application may return after the resources have been reassigned. Should demand return, another costly and risky harvest and reallocation would begin.

The resulting stranded resources remain unused, prematurely forcing IT groups to buy equipment to deploy new applications and upgrade existing ones.

Enterprise Modular Computing

Many research firms have come to the conclusion that enterprise computing must change. Gartner Group envisions policy-based computing.[5] Forrester Research envisions Organic IT.[6] Giga Information Group envisions modularity and virtualization.[3]

Intel is taking a leadership position in this movement. Intel expects Modular Computing to play a major role in enterprise computing.

A New Computing Paradigm

Modular Computing relies on a new paradigm for computers. Modular Computing draws elements from pools of computing resources—processing, storage, and networking. Together, these resources become a virtual server, a computer that can be assigned to run one or more applications. However, unlike a traditional server, when demand for an application changes, virtual servers can dynamically be repurposed, in just minutes.

A virtual server is logically integrated rather than physically integrated. This distinction is essential for enabling potent management of the resources. A control module, running Modular Computing software, manages the creation of virtual servers and facilitates real-time allocation and deallocation of resources.

Modular Computing

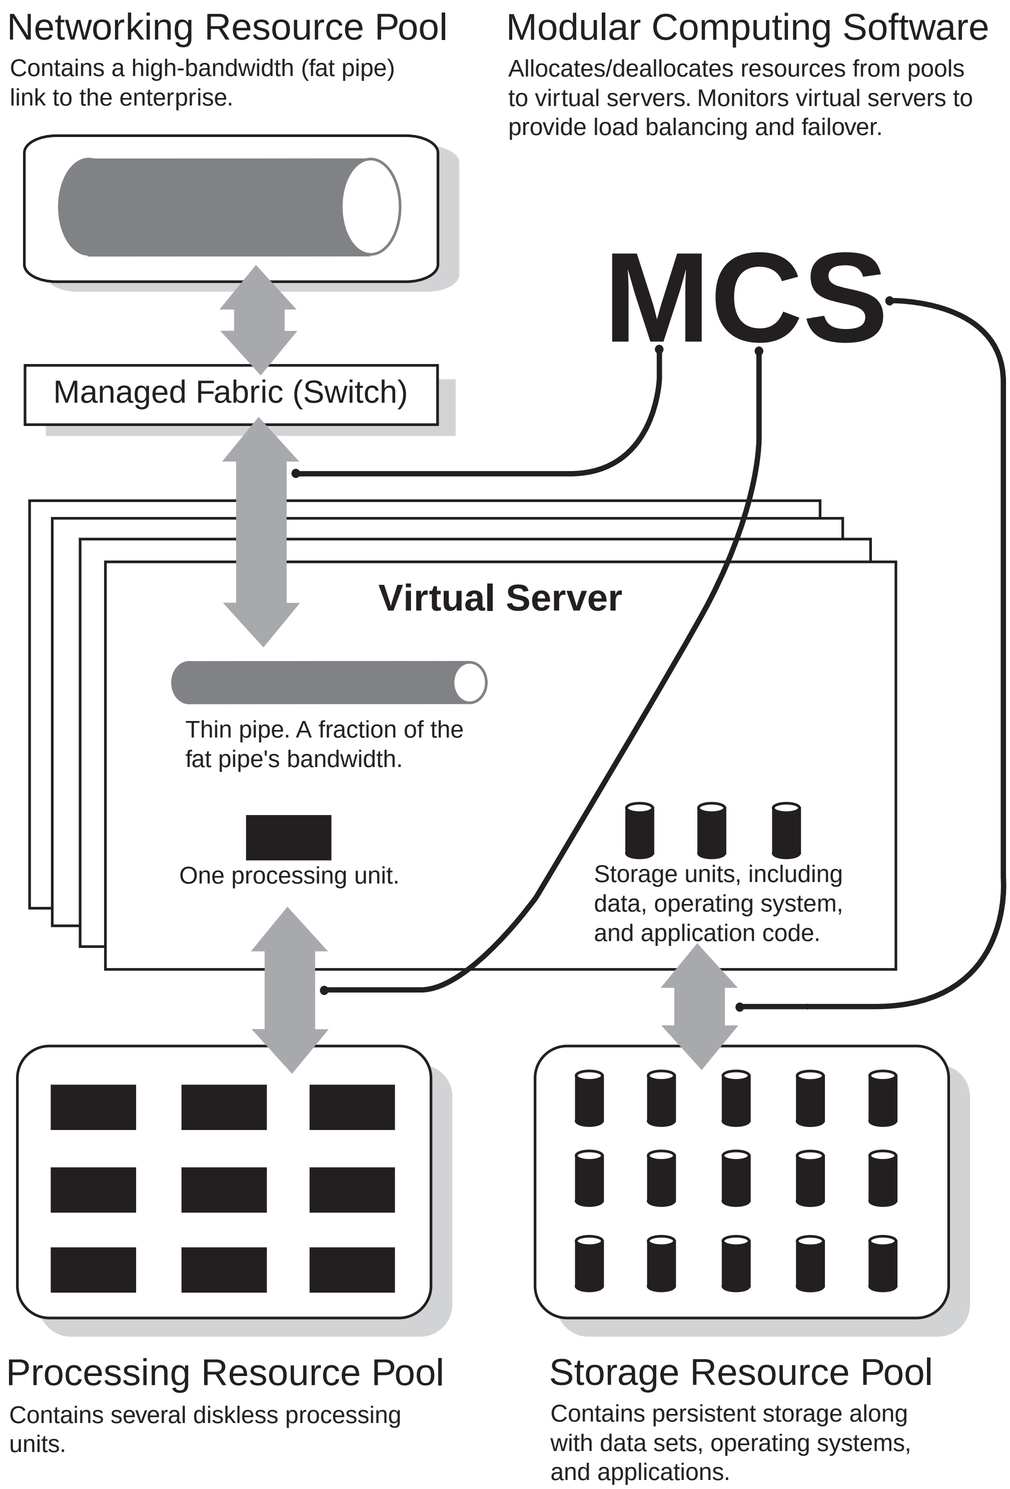

Processing Resource

Modular Computing's processing resource is based on Intel Architecture (IA) processors because of their superior price for performance across all business and technical workloads. As the Web site for the Transaction Processing Performance Council (www.tpc.org) shows, IA processors offer significantly better price/performance than proprietary processors in this business computing benchmark. The HPC and Web Server benchmarks also show superior price/performance for IA processors. Both of these benchmarks are on the Intel Web site at www.intel.com/products/benchmarks/server/index.htm

Intel server processors range from the 32-bit Intel Xeon™ processor MP with strong transaction and I/O processing capabilities, to the 64-bit Itanium® processor family with high-performance floating point execution.

Because of their robust capabilities and price for performance, servers based on IA processors are very popular.[7] Consequently, software developers have written hundreds of applications for servers based on IA processors. All this software provides IT groups with many choices, and competition between developers helps reduce software costs.

For the balance of this paper, "processing unit" refers to the smallest "chunk" of processing power that can be deployed from the processing resource pool. For example, if the processing resource pool consists of 4-way SMPs, a processing unit is a 4-way SMP.

Storage Resource

For Modular Computing, the storage resource should be a Storage Area Network (SAN) or network-attached storage (NAS).

Using SAN or NAS allows a computing facility to concentrate the storage in one physical location and obtain economies of scale. For example, mirroring, backup, and offsite archiving processes are much more cost effective on a SAN or with NAS than when applied to directly attached storage.

In addition, SAN or NAS allow the server's personality (operating system, application, and data) to be defined completely by the content of storage. The processing resource can be diskless and anonymous. This allows any processing unit to be assigned to any application, facilitating the dynamic nature of logical, rather than physical, connections. If storage were directly attached to the processing resource, the personality would follow the processing resource, making it less suitable for use with a different application.

If a virtual server consumes its storage resource, the Modular Computing software automatically allocates another unit of storage to the server.

Networking Resource

The networking resources should be a high-speed network (gigabit Ethernet, for example) accessed through a high-speed switch. This should provide access to both the LAN and, if needed, the Internet.

Just as storage resources are flexibly allocated to meet computing demand, networking resources must be scalable so bandwidth does not hinder performance.

Modular Computing Software

The Modular Computing software (MCS) is a vital part of Modular Computing. It obtains resources from resource pools and aggregates them into virtual servers. It also provides an interface for administrators. Running on a control module, it can oversee several virtual servers.

The MCS monitors and facilitates activity of each virtual server. For example, when demand for an application increases enough that existing virtual servers are unable to keep up, the MCS spawns another to process the same application. Similarly, when demand falls below a threshold set by an administrator, the MCS disbands a virtual server and returns its resources to the resource pools.

Depending on the nature of the MCS, the allocating and deallocating of resources can happen automatically, using parameters set by the administrator. Or the administrator can control it directly.

The MCS also monitors the health of each virtual server. If a resource fails, the MCS disassociates the failed resource from the virtual server, allocates replacement, then informs the administrator about the status of the failed resource.

Fewer Cables to Manage

To facilitate expansion and maintenance, the processing and networking resources, along with the control modules, could be mounted in the same rack. If this rack provides a high-speed interconnect, it can reduce the number of cables from more than two hundred to a mere handful. Two cables from redundant switches replace the NIC cables for all the servers. All the virtual servers access the storage resource through just two cables. And a forest of KVM cables is eliminated by providing an administrator interface across the network.

Benefits of Modular Computing

Modular Computing increases agility, while reducing equipment and people costs.

Increased Agility

Changes in computing demand need no longer cause panic. The Modular Computing software (MCS) can monitor the status of virtual servers in real time. As demand for an application changes, the MCS can adjust the number of virtual servers to match, in minutes instead of weeks.

This real-time load balancing prevents applications from slowing down for long periods. The users of the application don't suffer lengthy response times associated with overloaded servers.

Equipment failure no longer takes applications offline. When the MCS detects a failure in the equipment allocated to a virtual server, the MCS logically swaps out the failed equipment, replacing it with resources from the pool within minutes. Applications keep running.

Because this failover capability is automatic and fast, it enables administrators to extend failover coverage beyond mission-critical applications to all applications running in the Modular Computing environment.

Suppose demand grows for many applications, threatening to regularly consume all of one of the resource pools. Rather than purchasing an expensive, traditional server, IT purchases only the resources needed (processing units, storage units, or network capacity) and adds them to the resource pools. The MCS takes care of deployment in minutes as demand fluctuates. Compare this to the weeks it takes to set up a new traditional server.

Reduced People Costs

With traditional, physically integrated servers, equipment failure often means an administrator needs to visit the rack immediately to make replacements. Each such visit is time consuming and costly. Rack visits become rare with Modular Computing.

The Modular Computing software (MCS) acts automatically. It uses parameters set by administrators to govern resource distribution. Once the administrator has set the parameters, the software can balance loads or invoke failover procedures without human intervention, in minutes.

In addition, because adding resources to a Modular Computing environment is so easy, substantially less administrator time is spent on configuration and setup.

Management, too, becomes easier. *All* applications running in a Modular Computing environment are monitored by the Modular Computing software. Compare this to a collection of disparate, physically integrated servers, where each server class needs unique management tools. Reducing the number of management tools means fewer specialized experts.

Consequently, IT management can move people from administrative duties to activities with higher ROI, such as planning or application development.

Reduced Equipment Costs

All applications running in a Modular Computing environment share the same resource pools. In other words, the entire collection of virtual servers draws load-balancing or failover resources from the same resource pools. In contrast, with traditional computing, *each* mission-critical application needs spare equipment standing by for failover or load balancing.

With Modular Computing, a little spare resource protects all applications. Because less resource can do the job, utilization of resources is higher.

A related benefit of Modular Computing is the absence of stranded resources. The MCS harvests under-utilized resources automatically.

Modular Computing helps IT do more with less. By increasing utilization of computing resources, Modular Computing holds down capital expenditures. By freeing administrators from tasks such as load balancing and deployment of hardware, it makes them available for other tasks, with higher ROI.

So when does Modular Computing become reality? It's available now. The Egenera BladeFrame* System has been shipping since late 2001.

The Egenera BladeFrame* System

The Egenera BladeFrame system consists of Modular Computing software, connections for SAN or NAS and IP networking, and as many as 24 virtual servers based on Intel processors.

The BladeFrame provides a pool of up to 96 Intel processors, deployable entirely through software, with no physical intervention. The system components are listed in the following table.

Component	Description
Processing Blade*	2-way or 4-way, diskless, symmetric multiprocessors (SMPs) using Intel processors. Each virtual server uses one Processing Blade. The BladeFrame* system can contain as many as 24 Processing Blades.
Control Blade*	This is the control module for the BladeFrame system. It runs the Modular Computing software and provides security for the Processing Blades. To ensure high availability, each BladeFrame system has two Control Blades.
Switch Blade*	This is the networking resource for the BladeFrame system. It provides communication with the SAN or NAS and the IP network. To ensure high availability, each BladeFrame system has two Switch Blades.
BladePlane*	High-speed interconnect. Enables communication between components within the BladeFrame system.
PAN Manager*	Modular Computing software (MCS) to configure virtual servers and govern failover and load balancing. Administrators can use the browser-based interface or can write scripts to provide control through a command-line interface.

The system resides in a 24x30x84-inch chassis containing a redundant BladePlane, two Control Blades, two Switch Blades, and up to 24 Processing Blades.

The BladeFrame system is a processing resource for the data center. The Processing Blades are diskless, accessing the data center's storage area network (SAN) or network attached storage (NAS) for storage resources, software, and data.

Separating processing resource from storage lets the processing resource remain anonymous—not permanently dedicated to any particular application(s). Anonymity facilitates reallocating Processing Blades between the processing resource pool and virtual servers. Egenera calls this diskless architecture a Processing Area Network, or PAN, and the management software is called PAN Manager*.

This PAN architecture facilitates efficient use of processing resources. As demand for a particular application declines, PAN Manager software reduces the number of virtual servers assigned to that application, making their resources available for other applications. PAN Manager shifts resources automatically, in minutes. By rapidly distributing resources to where they are needed, PAN architecture eliminates costly over-provisioning.

Should a piece of equipment fail, PAN Manager detects the failure, notifies the administrator, and allocates a replacement resource, all within minutes.

The BladeFrame system greatly reduces cable count. With traditional architecture, each single-processor server can require more than five cables, without providing redundancy. With the BladeFrame, as many as 96 Intel processors can be redundantly connected to the storage and IP networks with as few as four cables. This huge reduction in cables saves many error-prone hours during installation, while offering fewer failure points and increased density of servers. By reducing cable count, the BladeFrame contributes to higher reliability (because of fewer failure points), more efficient use of administrators (by saving cabling time), and less stranded equipment (by simplifying harvesting and redeployment).

Administrators can easily add additional Processing Blades to a BladeFrame system. The administrator simply slides it into the chassis. The Control Blade detects the new Processing Blade and automatically places it in the processing resource pool. Adding and deploying a new Processing Blade takes just several minutes. Because storage is separate from the Processing Blades, adding Processing Blades requires no manual installation of software.

Many benefits of the BladeFrame system derive from the Egenera PAN Manager software, which provides a single control point for monitoring and allocating both physical and logical resources. For example, administrators create pServers, a logical equivalent of a conventional server, by associating individual Processing Blades with storage and IP network capacity.

Similarly, administrators can create several logical PANs (LPANs) on a single Blade Frame. LPANs have physically distinct, secure resources even though they exist on the same BladeFrame. They do not communicate with each other. LPANs allow administrators to run applications for disparate customers (such as individual corporate divisions) on the same BladeFrame. And administrators can establish a pool of reserve Processing Blades shared by LPANs to provide incremental capacity and/or automatic failover. The ability to share a pool of Processing Blades slashes the cost of high availability and boosts processor utilization.

Using PAN Manager software, administrators can rapidly adjust logical configurations to service changing demand. Tasks that were once physical and required weeks (such as reconfiguring a cluster of servers to securely accommodate a new customer) are now accomplished through software in minutes.

The hardware and software modules of the BladeFrame system work together to provide automation and rapid, flexible deployment. The BladeFrame system saves administrator time associated with cable management and other deployment issues. It automates harvesting and reassigning resources, while slashing the cost of high availability.

Stretching the IT Dollar

Modular Computing replaces the physical connections between computing resources with logical connections. Because the connections are logical, Modular Computing software can monitor and control how virtual servers use resources.

This software-based monitoring and controlling enables automated resource management, where the software continuously redistributes resources according to parameters provided by an administrator. This, along with simplified server management and reduced cable count, means large collections of servers need fewer administrators.

Modular Computing uses small amounts of spare resource to provide failover and load balancing for all applications running in the Modular Computing environment. This eliminates stranded resources, boosts resource utilization, and holds down capital expense.

Because Modular Computing is built on IA processors, it offers better price for performance and ensures a broader choice of software vendors and software. A larger selection of software can speed application development, and competition between software vendors can hold down development costs.

Modular Computing is a concept for the future, but it is available now, in products shipping today. It is already proving itself by saving money for IT.

---

[1] Forrester Research, Inc., *The New Computing Utility*, February 2002.

[2] Goldman Sachs, *IT Spending Survey*, November 29, 2002.

[3] Giga Information Group, Inc., *The Future of the Data Center—Modularity and Virtualization*, May 8, 2002.

[4] Giga Information Group, Inc., *Budgeting for IT—Average Spending Report*, August 23, 2002.

[5] Gartner, Inc., *The Evolution Toward Policy-Based Computing Services*, October 2002.

[6] Forrester Research, Inc., *Organic IT*, April 2002.

[7] IDC, *Q2 Server Tracker* indicates that IA processors hold 87% of the server market segment share as of September 2002.