


Realtime
publishers

Simplified High-Performance Computing for Everyone

The Essentials Series

sponsored by



Don Jones

Introduction to Realtime Publishers

by Don Jones, Series Editor

For several years now, Realtime has produced dozens and dozens of high-quality books that just happen to be delivered in electronic format—at no cost to you, the reader. We’ve made this unique publishing model work through the generous support and cooperation of our sponsors, who agree to bear each book’s production expenses for the benefit of our readers.

Although we’ve always offered our publications to you for free, don’t think for a moment that quality is anything less than our top priority. My job is to make sure that our books are as good as—and in most cases better than—any printed book that would cost you \$40 or more. Our electronic publishing model offers several advantages over printed books: You receive chapters literally as fast as our authors produce them (hence the “realtime” aspect of our model), and we can update chapters to reflect the latest changes in technology.

I want to point out that our books are by no means paid advertisements or white papers. We’re an independent publishing company, and an important aspect of my job is to make sure that our authors are free to voice their expertise and opinions without reservation or restriction. We maintain complete editorial control of our publications, and I’m proud that we’ve produced so many quality books over the past years.

I want to extend an invitation to visit us at <http://nexus.realtimepublishers.com>, especially if you’ve received this publication from a friend or colleague. We have a wide variety of additional books on a range of topics, and you’re sure to find something that’s of interest to you—and it won’t cost you a thing. We hope you’ll continue to come to Realtime for your educational needs far into the future.

Until then, enjoy.

Don Jones

Introduction to Realtime Publishers..... i

Simplified High-Performance Computing for Everyone..... 1

 What Is HPC? 1

 Assessing Your HPC Needs 1

 Other Components You'll Need 3

 Don't Forget the Storage..... 4

 Getting the Right Software 4

 Making It All Work Together 4

HPC: Ready for Your Organization..... 5

Copyright Statement

© 2011 Realtime Publishers. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtime Publishers (the “Materials”) and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtime Publishers its web site sponsors. In no event shall Realtime Publishers or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtime Publishers and the Realtime Publishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtime Publishers, please contact us via e-mail at info@realtimepublishers.com.

Simplified High-Performance Computing for Everyone

High-Performance Computing (HPC) is one of the most interesting developments in the information technology (IT) industry in the past few years. It's also one of the most-overlooked. Essentially, HPC enables an organization to harness massive computing power by distributing workloads across numerous relatively-inexpensive computers. Rather than investing millions in a supercomputer, you simply aggregate the power of multiple PC-based servers—often “commodity” servers—built using standardized components and architectures. Adding computing power is as easy as adding new servers to the HPC cluster.

HPC is helping organizations of all kinds churn through incredible amounts of data much faster. Life sciences companies use HPC for genetic and molecular modeling. Businesses use HPC to forecast complex sales trends, make supply chain predictions, and spot business opportunities. Educational institutions use HPC to power research projects involving massive amounts of data. HPC can be sized for project workgroups, departments within a larger organization, or an entire organization. You can even size HPC to form massive “compute clouds” capable of supporting multi-organizational projects.

In the past, building HPC clusters has involved a lot of complexity. Today, however, the HPC revolution has put this kind of computing power into the hands of nearly any organization. This revolution combines powerful, less-complex hardware with commercial operating systems (OSs)—including HPC-specific variants of Linux and Microsoft Windows—for simpler HPC configurations that still deliver an incredible amount of power.

What Is HPC?

HPC solutions consist of *clusters* of computers. The idea is simple: Take a massive workload that would require far too long for a normal computer to process, break that workload into chunks, and distribute those chunks across a set, or *cluster*, of computers. With multiple computers working on the problem at once, the problem is solved faster. A central coordinator, often called the *master node*, distributes workload to the computers working on the actual problem—the cluster's *compute nodes*.

Assessing Your HPC Needs

HPC isn't a one-size-fits-all computing model. HPC solutions are often custom-built for specific purposes, although many vendors have started offering pre-built solutions designed for a range of common requirements. When considering an HPC solution, you will have to start by understanding your needs.

For example, many organizations will want to start by selecting an HPC OS, and they'll often select the one that corresponds to their existing server OS investment. Organizations already using Microsoft Windows Server, for example, might prefer to power their HPC solution with Microsoft Windows HPC Server. Linux-centric organizations might opt for Red Hat Enterprise Linux.

A key consideration will be density. HPC solutions, like any collection of servers, need space, so you'll want to make sure you're using your space wisely. HPC solutions are often measured by the number of processors they contain, rather than the number of servers. Specialized servers designed to support HPC can pack a large number of processors into a small space; by combining racks of these servers, you can build an HPC solution that maximizes data center efficiency. Or you can choose blade servers that offer high processor density combined with central power management and other features. You might also opt for less processor density and greater expansion options by using traditional rack-mounted servers.

Tip

It used to be that blade servers were the highest-density avenue to building an HPC cluster. Today, specially-designed 1U rack-mounted servers can contain two multi-core processors, creating higher density while still offering internal expansion capabilities. These do have some limitations on their ability to quickly move massive amounts of data to the local network, so an alternative is a 2U server that provides a bit more room in the server chassis for higher-speed components.

Density is extremely workload dependent, and a good HPC vendor can offer you invaluable advice on what to choose. For example, some HPC applications still require a great deal of computing power to process the individual "chunks" that are distributed across the compute nodes. In those cases, compute nodes may need four or more sockets, something that isn't usually possible in a smaller 1U or 2U chassis. Larger chassis, such as a traditional 4U rack mount server, can also offer lower power consumption.

The next balancing act is between performance and energy efficiency. In the past, HPC solutions were often built with a complete focus on computing power, with little regard for electrical power. Today, HPC vendors are more sensitive to environmental and budgetary concerns, and usually offer HPC solutions that can maximize energy efficiency, create a balance between performance and efficiency, or maximize performance when that's what the customer needs. Your energy/performance decision will inform not only the selection of processors inside your HPC cluster nodes but also the amount and type of memory installed within them.

HPC vendors can typically tune their cluster configurations to specific types of workload. These workload types are often represented by specific benchmarks, such as fluid dynamics, computer-aided engineering, general-purpose, and the like.

Finally, you'll need an estimate of how many gigaflops (GFLOPS) your HPC solution will need to support. A *flop*, which stands for *floating point operations per second*, is a measurement of computing performance. The exact application you plan to use with your HPC solution will inform how many billion operations per second you will need to support, and that number will tell your HPC vendor how many nodes your solution will need to contain. For example, a general-purpose HPC solution that can sustain 10,000 GFLOPS (that's ten petaflops per second) in a balance energy/performance configuration might require around 250 cluster nodes. Significantly smaller HPC solutions are more common, as few HPC applications really require that kind of power. For super-sized applications, of course, even larger HPC solutions can be created.

Other Components You'll Need

HPC solutions also require extensive interconnect networks, which must support the bandwidth requirements of the HPC applications that the solution will run. These interconnects serve to connect each HPC node with its neighbors, and to connect all the nodes to the controlling servers, logon workstations, and other HPC components. These interconnects are a core challenge of building HPC clusters, as the bandwidth of the interconnects needs to be sufficient for the cluster's workload.

Master, or *login*, nodes are an essential part of an HPC solution. These often run *pre-* and *post-processing applications*, which create the data input for jobs sent to the HPC nodes and analyze the final output of the computing operations. These nodes may also help create visualizations of data, helping human beings interpret it and utilize it more effectively. High-powered workstations are typically used for this task, providing more processing and graphics manipulation capabilities than a traditional personal computer offers. Interestingly, newer servers—especially those designed for HPC uses—are beginning to include those same high-end graphics capabilities. These higher-end graphics processing units (GPUs) are providing those servers with additional specialized compute power, which can be extremely beneficial for certain HPC workloads.

In larger HPC solutions, it is common practice to separate the master and login functions. Login roles are still usually accomplished on a powerful workstation, but the master role—the node that distributes workload to the overall solution—is often moved to a dedicated server.

Don't forget that all the servers and workstations in your HPC solutions *are* still independent machines. They'll need to be managed, kept up-to-date on patches for OS and applications, and so forth. If a node experiences a hardware failure, you're going to want to know about it quickly so that corrective action can be taken. All of this means selecting servers that have robust built-in management capabilities. Remote management features are something we often look for in a standalone server and, for the same reasons, you will want them in your HPC nodes as well.

Don't Forget the Storage

Storage is a fundamental component of any HPC solution. Although it's easy to focus on raw computing power as you assess hardware decisions like processor selection, number of nodes, and so forth, the fact remains that the results of all that computing—not to mention intermediate data used to generate results—has to be stored somewhere. HPC solutions typically measure storage both in terms of total capacity (usually in terabytes) and raw throughput (usually in gigabytes per second). You not only must have *enough* storage but it must be the right *kind* of storage, capable of sending and receiving information quickly enough to let your HPC nodes sustain their computing pace.

Storage solutions designed to be paired with HPC systems are often built in the form of Storage Area Networks (SANs), using arrays of high-speed drives to provide the capacity you need at the speed your HPC solution will require. SAN connectivity is often built over fiber-optic networks or extremely fast copper (using 10Gb Ethernet, for example) to keep data moving quickly.

Getting the Right Software

In addition to an HPC-compatible OS, you will need a great deal of other software to make your HPC system work properly. Perhaps the most crucial components of the software stack are the hardware drivers and middleware that enable the OS to communicate with, and manage, the hardware. Poorly-written driver software can destroy an HPC solution's performance from the inside. Thus, it's *very* important that you select an HPC vendor that can provide well-written, thoroughly-tested, and proven drivers to ensure that the software doesn't become a weak point in the configuration.

You'll also need HPC-specific applications. You can't just drop any application onto an HPC solution; applications need to be designed to break down large workloads into small, independent pieces, which are then sent out to the HPC nodes for processing. Results are then assembled from the individual nodes. HPC is distributed computing at its finest, but it does require specialized applications capable of distributing their workload in this fashion.

Making It All Work Together

As you begin to plan for your HPC solution and evaluate vendors, there are a number of considerations to have in mind in addition to price, support options, general quality, and so forth. Chief amongst these is an independent validation of the hardware. HPC solutions are made up of a huge number of tightly-interconnected components. It's critical that every component—processors, memory, chipsets, motherboard busses, storage components, networks, and more—all be selected and tuned to suit each other. Spending extra for faster processors won't help if the HPC nodes are connected by an underperforming interlink network. In fact, one of the biggest hurdles to HPC in the past has been the difficulty in assembling the proper components.

HPC vendors and component manufacturers are aware of the problem, and many have taken steps to help. Intel, for example, offers the Intel Cluster Ready program, which is designed to ensure the interoperability of HPC components. Rather than trying to build your own HPC solution from off-the-shelf components, a Cluster Ready-certified system is designed and certified to function properly. Many HPC applications are also registered with the Intel Cluster Ready program; by running those applications on a Cluster Ready-certified HPC solution, you'll know that everything will work correctly together.

Aside from Intel's program, most HPC-savvy server vendors also maintain their own cluster certification and verification programs. In most cases, these programs enable the vendors to build and customize cluster solutions to each customer's specific workload while verifying the overall interoperability of every component within the final solution.

HPC: Ready for Your Organization

HPC clusters used to be the exclusive domain of large, research-focused businesses and educational institutions. Today, the availability of preconfigured, cluster-ready hardware configurations from HPC-aware vendors means that this kind of computing power can be leveraged by almost any organization. Whether you're solving scientific riddles or forecasting complex sales and supply scenarios, HPC clusters can help you find your answers faster.