



Breaking the Curse of Distributed Environments

By Don Dugdale

As decentralized computing comes to carry the load of mission-critical information systems, IS managers have to make them perform well. Here are some tips for tuning those networks.

Distributed computing naturally creates headaches—"challenges" in vendor-speak—for those who have to keep each site and the whole network running. This should not surprise anyone who follows the progress of technology. Jonathan Eunice, director of the client/server program for IT consultants Illuminata in Nashua, NH, uses an analogy to the past to explain distributed computing. "In the early years of this century, when electrical lines were coming in, it was a messy infrastructure," he says. "If you look at pictures of New York City in 1915, every street was like the inside of some electrical rat's nest. That's typical of first-generation technology; you put wires everywhere." Until people developed a recognized body of knowledge on how to handle the electrical infrastructure, that messy wiring caused many fires and electrocutions. Then, over time, the infrastructure was refined.

"It took 10 to 20 years and became safe and effective," says Eunice. "That's what's happening in networked computing. We're about 10 years into this phenomenon on a broad basis. Until two or three years ago, we were still putting up more wires and tying them together. The experience level has become broad, and people coming from the mainframe world or the PC world have seen the benefits and are understanding how to put it together."

As heterogeneous, distributed computing has matured, IS managers have sought to raise the performance levels of their multivendor, multitier, geographically diverse systems to those of their host-centered systems housed in a data center—with mostly frustrating results. "Performance is bad," says Chet Geschickter, director of research for Hurwitz Consulting in Newton, MA. "Distributed computing doesn't have to degrade performance, but it can tend to."

How finely IS managers can tune their systems and how soon they can do it will be key determinants of computing's future direction. "Most everybody agrees that client/server as a business technology will rise or fall in the next five years, based on the strength of the tools that are offered to support it," says Ted Collins, vice president of Viatch Lab, a division of Platinum Technology in Eagan, MN. Although there may be basic agreement on that prediction, there's less consensus on how far the supporting technology has come. Collins estimates that it's no further along than the automobile industry was in the 1920s.

Others see distributed performance management at about the same point as where the IBM MVS mainframe operating system was in its midlife; that is, advanced enough to do a pretty good job but not quite mature. "Whenever new technology is introduced, technology comes first and manageability second," says Doug



McBride, manager of technical marketing for Hewlett-Packard's resource and performance management group in Roseville, CA. "People focus on implementation and getting things to work—availability and functionality. Then they start to get their hands around it."

With multivendor systems, networking, applications and database components proliferating throughout the enterprise, no one questions that distributed computing can be a strange and hard-to-manage milieu, requiring a more complex set of tools than the mainframe environment. Although the sophistication of those tools is expected to improve markedly over the next two years, reliable performance tuning is possible now. Here are some ways to realize it.

Identifying Tools

A variety of tools exists for performance monitoring, metering, comparison and prediction. IS managers need to know what tools are there and how they function at the system, network, application and database levels. "It's pretty much within the last year that we've seen these tools mature from a network-centric, up-down focus to an infrastructure focus that gives you qualitative as well as quantitative measures on the behavior of the infrastructure," McBride says. "Once you start to roll them out and use them, the distributed system is as controllable, for the most part, as the mainframe was in its mid-lifetime."

However, because the function of systems management tools varies so widely, expertise on exactly what's currently available is rare. "People need more diligent understanding of the tools offerings and the different vendors," Geschickter says. Rather than working with a single vendor, as was common in the mainframe-centered world, an enterprise will have to pick key partners and several vendors to obtain the best set of tools for system management and control. For example, Geschickter says that one global oil and gas company picked three key partners—a platform infrastructure provider, a development tools vendor and a middleware vendor. "You'll have many vendors and a few partners," he says. "That selection

process includes not only technology but vendor support."

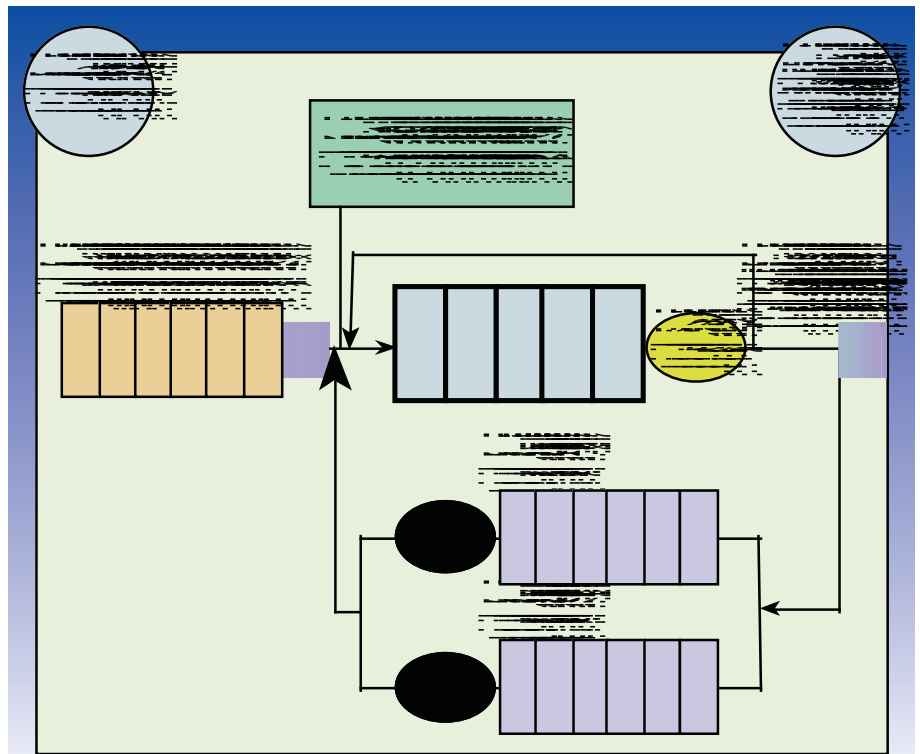
No single tools vendor is considered likely to meet all possible needs. Some systems management platforms such as HP's OpenView, IBM's NetView/6000, Sun Microsystems' SunNet Manager and Spectrum from Cabletron of Rochester, NH, are focused on the network. Others have a strong mainframe flavor, including Computer Associates' CA-Unicenter. The various products of Platinum Technology and OpenVision of Pleasanton, CA, are amalgamations of separately developed technologies bought by the parent companies and marketed as unified product lines. Still others, like the Tivoli Management Environment (now owned by IBM) and Power Center from Innovative Software of Englewood, CO, are attempts to create a single, comprehensive client/server management system. The trick is to know which is best for a given situation.

Emerging Trends

Organizations moving from a mainframe environment or a shared file and print

environment are in for a shock: the need to handle diverse computing resources as nodes on a network. The mainframe itself has to be viewed as another node. Therefore, configuration and performance of the network have become important elements in total system performance and are more complicated than simply relying on IBM's System Network Architecture (SNA). For example, managers have to understand the Simple Network Management Protocol (SNMP) and its applications. SNMP is a public domain protocol for managing multivendor devices and networks developed by the Internet Engineering Task Force (IETF). Also currently under development are management information bases (MIBs), documents that contain information allowing various vendors' products to address the components and objects on a network in the same way. These incomplete solutions help in network monitoring and are basic to understanding distributed systems.

Performance monitoring and workload distribution are considered basic to keeping a distributed system tuned. "Per-



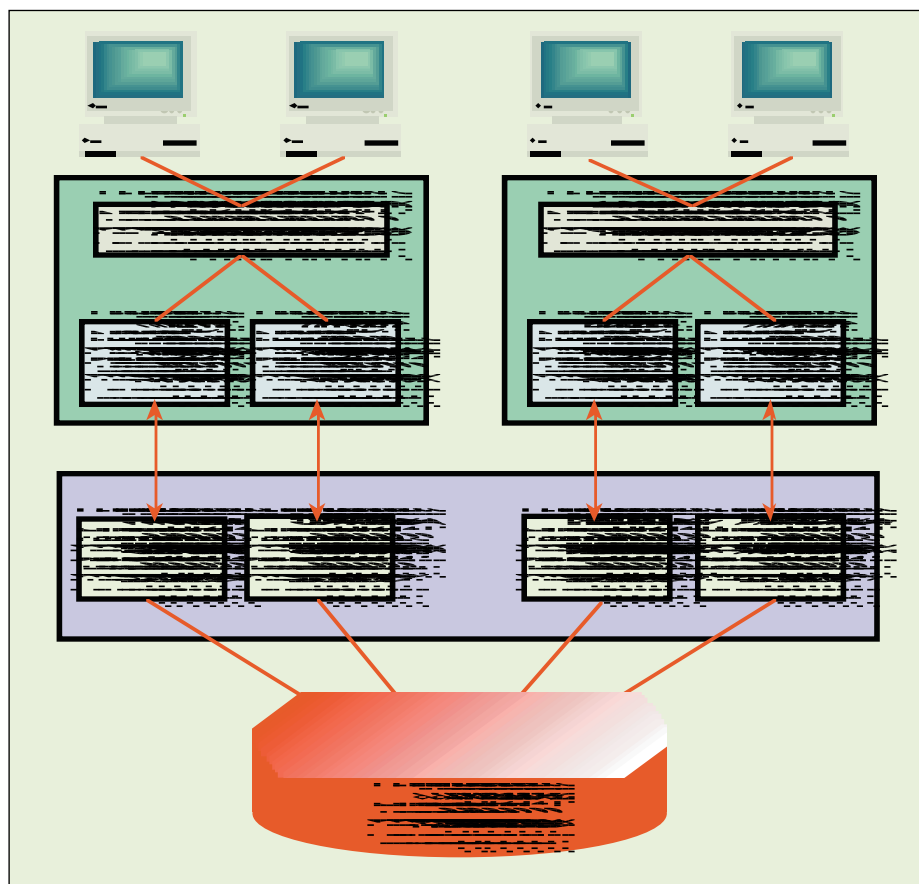
Tracking system response time in application-level management may require setting up a queueing mode. It's important to know how much time is spent waiting in each queue. This diagram illustrates the components in which an application may access resources.

formance monitor probes can help you gather and quickly synthesize and respond to performance data," Geschickter says. Ideally, a workload balancing tool will collect statistics on how the workload is processed, then use it to spread the processing so no resource is overwhelmed.

As distributed systems mature, trends in system management will determine effective approaches to performance tuning. One current trend is toward application-level management rather than system-focused management. This area also has no all-in-one solution. "Nobody today gives you end-to-end performance monitoring," says Steven Foote, director of product marketing and integration for OpenVision.

Some products now available monitor performance on the application level. Among them are EcoNet and EcoTools from Compuware of Farmington Hills, MI; Measureware from HP; and Patrol from BMC Software of Houston. The idea is to monitor and control the system's response to a single transaction within an application—for example, to track the beginning and end points of a catalog order—through an application programming interface (API) call from a client-based application. "The biggest need I see right now is for the capability to see how long it takes to do a mission-critical business transaction," Geschickter says. "That's how long the customer representative or the service person is going to be sitting at their screen getting mad at their computers."

Once that transaction time is known, an application dependency stack could be used to "drill down" into the various levels of the system itself (the user interface, tools, network, database, operations system and hardware) to determine where hangups are occurring and how system response might be improved. "Rather than just focus on a particular piece of the infrastructure, like a network segment on the LAN or how an operating system reports hardware resources being consumed, we're starting to see tools that look at the bigger view—how the infrastructure is being used by the application," McBride explains. To get that level of control, the application itself has to take a role in its own manageability. "To try to understand what end users need," he says, "you can have your application report



System response time in a distributed environment can be divided between waiting time and service time as a transaction is processed through the operating system, client, database and network.

Source: BGS Systems

back the responsiveness that it's actually giving." That, in turn, allows a system manager to directly improve the quality of service to the end user.

The trend toward application-level management is being facilitated by the IETF's attempt to define an application-level MIB under SNMP and by the Desktop Management Interface (DMI), a spin-off from the IETF that deals with the interface between desktop applications and management tools. "It will probably take through this calendar year before there is a specification that people can feel comfortable about using in order to instrument their applications for manageability," says Wayne Morris, open systems product marketing manager for BMC, one of the companies involved in standards development. "Pragmatic solutions that groups of vendors get together on will solve the problem."

A corollary to the development of this kind of management is the decreasing emphasis on "point solutions" (tools that analyze only one level of the system).

"The people that are having trouble today are still using those point solutions," McBride says. "They don't even know where to start looking [for a problem], especially if it's a multitiered, highly partitioned application that's spread across the IT infrastructure."

Another trend, according to Geschickter, is toward Unix-oriented distributed systems and network technology in place of products with MVS origins. "As you move toward a distributed environment, you need a distributed solution," he says. "We expect that those technologies will start to take the central role from the more mainframe-oriented approaches."

That need for a distributed solution is also taking IT away from the concept of "islands of management," which resulted from the decentralization of IS, both by geography and by platform. "In order to implement an enterprise management solution, a company is faced with having VMS people agree with their HP-UX people, with their Windows NT people, with

(CONTINUED ON PAGE 34)

This page
intentionally left
blank

(CONTINUED FROM PAGE 32)

their Sun Solaris people," says Jay Whitney, president and chief technical officer of Innovative Software. "When they all agree on a solution, they're going to pray that it works together. A new mind-set in IS is to not seek that island concept."

Good Management

Experts agree that sound IT management practices are an integral part of performance tuning. Their suggestions take practical approaches to effective IT management. For instance, it is not necessary to discard mainframe-oriented disciplines in distributed computing. "Performance is a symptom of how the various capacities of the infrastructure are being managed," McBride says. "This is a science, believe it or not. It has been used in the mainframe world for years, and people are now trying to transition those techniques and methodologies into Unix and other operating environments."

Collins of Platinum points to other originally mainframe-based techniques, including centralized production control, capacity planning, backup and restore, and disaster recovery. "All those disciplines are relevant and necessary," he says.

It is also worthwhile to focus on how technology can serve IS management purposes. For example, observers suggest setting performance goals and then using monitoring tools to measure whether those goals are being met. Some advise implementing service-level agreements, as was done in a host-centered environment. IS agrees to provide certain services to the users within the enterprise; this arrangement can prevent misunderstanding. "That boils down to the concept of managing user expectations—managing against service levels that are required to run the business," McBride says.

Keeping the emphasis on business needs also helps. "If you build this stuff to provide service rather than just to provide technology, then less and less often you

find yourself getting into a situation where things don't work as they should," he says.

Users should understand that, during a technology transition, they may have to put up with glitches where complex networking is involved, says Tim Lee-Thorp, business development manager for Remedy, a tools vendor in Mountain View, CA. "People need to realize that this stuff just happens," he says. "A ripple goes through the network, and it's gone. I don't have time to go and mess with it."

Moving from a host-centered system in phases can help to avoid performance problems, suggests Julia Lockwood, director of product marketing for Unison Software, a performance tools vendor in Sunnyvale, CA. "Our successful customers are the ones who start with one application and go one step at a time," she says. "They haven't thrown away their mainframe or other legacy systems. The people who jump in all at once have a tremendous management issue."

Planning for system growth is another crucial task that should be examined up front. Make sure your tools can grow with your system, experts advise. "Whatever your solution is, be sure it scales," Collins says. "Unless it scales, you are setting the corporation up for inflexibility and frustrating the whole reason for going to open systems in the first place."

What Not to Do

Just as important as what to do about performance tuning is what not to do. Remembering that distributed computing is still maturing can help everyone manage their expectations. For example, don't be overly concerned with precision. "Trying to tune to the gnat's eyebrow is more trouble than it's worth," says Lee-Thorp of Remedy. "Give yourself some slack and get on to the next problem."

Says Whitney, whose company developed a distributed systems management platform, "We tend to think it's more important to have 100 percent availability in a computing platform than to squeak

the last percent of performance out of it. If a system is not available, the performance is zero."

In these days of downsizing and budget crunches, it still remains tempting to overspend to solve problems. But buying more hardware or hiring more personnel to deal with performance is an expensive solution. "We have had a tendency to throw application processors at the problem," says Pete Bates, principal of LPB Systems, a consultancy in Mill Valley, CA, and formerly CIO for Esprit, an apparel vendor. "That's one of the advantages of Unix. It's easy to add and subtract the capacity and move an application from one processor to another. The flexibility is a marvelous thing."

Other advisors caution against this temptation. "The consumers of these products are confused," Collins says. "They spend money because they're confused about how to go about getting the metrics."

Similarly, it's a mistake to substitute infrastructure for management. "If you understand this as a business and understand what your users require to get their work done and make the enterprise more competitive, then you know how to build the infrastructure and how to solve the problem," McBride says. "Many people are not quite sure how much to put where, who's going to use it and why."

In the end, tuning distributed systems for best performance will be inevitably more complicated and difficult than the same process on a mainframe with terminals attached. And it remains true that the ideal management tool has yet to be devised. "The caustic statement is that everything has been a year away for the last five to seven years," says Lee-Thorp. But available tools can begin to penetrate the complex nature of networking and distributed application processing. Using them with sound IT management practices, you can gain control of distributed systems. Perhaps within a couple of years, comprehensive management solutions may come to market, but most organizations cannot wait. **IT**

Don Dugdale is a technology writer based in San Jose, CA. He can be reached at dugdale@netgate.net.

Moving from a host-centered system in phases can help to avoid performance problems.

This page
intentionally left
blank