Coherent Cluster Performance Management

by David Freund

As Information Technology has grown to become a vital part of almost all business processes, so has the need for reliable, predictable service. This has resulted in a need for ensuring enough computing power is available, and for ensuring that the service can continue even if major IT components fail. Clustering technology has evolved as a way to address both of these needs, scalability and availability, but not without a cost – increased complexity. Ensuring smooth system, storage and application fail-over and fallback – all without users noticing – is challenging by itself. But service availability is not enough; performance levels must also be assured for the end-user. This is particularly true in an Internet and E-Business world, where a company's competition can literally be a click away. Performance problems can lead to loss of money, customers, market share, and perhaps even the entire business.

Making the most effective use of a cluster, whether it's based on Microsoft Cluster Server for Windows NT, a UNIX clustering technology, Compaq OpenVMS clustering, or any other clustering technology, requires an understanding of details such as how data is shared, how multi-node communications and data accesses are performed and synchronized, what overhead is added by the clustering services, as well as how much capacity is consumed by the application services the cluster supports. Such a comprehensive understanding is gained in one or a combination of ways, including: reading books and manuals specific to the clustering technologies being used, sending IT managers to cluster management and internals courses, using outside cluster performance experts, and by using tools designed to help the tool user manage the cluster as a single entity.

In real-world deployments, clusters have typically been more dynamic and complex after installation and use than they appeared when they were first purchased. Application behaviors, their interactions with each other, their interactions with the operating system, and a variety of cluster services, can become quite intricate. Software updates from operating system, cluster, and application software vendors come regularly. Application loads and even specific feature usage can change on a moment's notice. In fact, the relationships among these components and their performance impact can change quickly and surprise not only the IT managers, but also the software and hardware vendors themselves.

The Challenge: Clustering Interrelationships

Performance management tools have typically been strong at managing a single component or technology a layer at a time. They focus on a single vendor's hardware, an operating system, a database, an application, specific types of network devices, or specific storage subsystems. However, they don't provide much insight into how these components or layers might interact with one another.

The best approach to cluster performance management, whether it's for planning purposes, or for solving a performance problem that is taking place right now, is to assume that no person or tool can possibly have a complete understanding of all the performance interrelationships for the current system in real-time. What's needed, then, is a systematic way to discover these relationships, across all the software and hardware components, and across all of the nodes in the cluster.

Network managers have been faced with a similar problem. When a certain critical network component would fail, their network monitoring software would get hundreds of alarms – or events – each claiming something different was wrong. Event correlation, finding a reciprocal or complementary relationship among the events, emerged as a "must have" feature. Based on the timing of the events, and knowledge of the network topology, one can rapidly pinpoint the root cause of all the events.

In managing clusters, event correlation is also useful, particularly when handling application or subsystem failures. But event correlation is limited to just that – events. In basic performance management, one can define upper and lower thresholds for various performance metrics and generate events when those thresholds are violated. But this still does not help unless one of the thresholds violated is for a resource that is the root cause of the problem. And it would be unreasonable to try to put meaningful thresholds on every possible system and application resource metric.

The Solution: Automated Performance Data Correlation

What's needed, then, is performance *data* correlation – finding relationships among the metrics themselves. This requires good instrumentation – gathering all the relevant performance data – and efficient storage of that data over a sufficient period of time. When an undesired performance event occurs, data correlation can be used to determine what was going on, not only at the instant the event occurred, but also over a period of time leading up to the event. To make this truly effective, such correlation should be on data gathered from all of the technology layers – hardware, operating systems, cluster services, databases, and applications – and from all members of the cluster.

Most IT managers so far have been forced to perform this type of correlation manually. Some performance management products provide the ability to "overlay" performance graphs (like placing one transparency over another on an overhead projector) to allow the manager to attempt to find a correlation visually, but this still assumes that the manager knows what metrics to overlay. In other words, he or she has to know exactly what to look for – a nearly impossible task, in almost all cases.

The good news is that there are tools available that automatically perform this type of correlation using logic-driven software, which is by far the most effective method both in time and thoroughness. Performing this correlation on volumes of data also requires of very efficient storage and retrieval method. To date, only stream-lined file structures have yielded acceptable levels of performance to provide feedback in real-time, as relational and object-oriented databases have not proven to be well suited to this type of task.

The Next Step: Relate Cluster Performance to IT Infrastructure

To then gain further control over cluster performance, the IT manager should be able to relate performance of the cluster and the applications to the rest of the IT infrastructure supporting the applications. This includes such things as storage subsystems (including RAID arrays and Storage Area Networks), web servers, and so on. Being able to correlate performance data across this entire structure provides the ability to automatically discover potentially invisible interrelationships that affect performance, and to do something about them in an immediate and productive way.

Adding a rules-based tuning advice engine to the above capability, which can be customized by vendors and IT managers, allows the performance management software to monitor the environment, watch out for problems that are starting to occur, or have occurred, and automatically notify the responsible party not only that something "went wrong," but also what they can do about it. This enables a company to delegate the more mundane performance monitoring and trouble-shooting tasks to the general IT staff, and provide more time for the staff skilled in performance issues to handle more challenging tasks.

Such software needs to be easily deployed, be able to generate analysis reports both automatically and on-demand, and to provide a single station from which all of this can be viewed. Ideally, a web-based interface is available; if it is, however, there must be security features that prevent unwanted access to the data. Finally, its ability to integrate with other system and network management tools, provides a deep performance analysis extension to the overall management capability.

Given the dynamic nature of cluster technology, applications, and the way they used in business today, the only reliable way to ensure a good understanding of performance relationships among cluster components is to have this type of systematic way to automatically discover them. Whether a company IT staff member or an external service provider performs the management of the systems or applications, having performance data correlation, rapid root-cause analysis, rules based tuning advice, and a single console view is an essential component to successfully implement clustering technology for maximum performance and availability.

David Freund is the Director of Product Management at FORTEL Inc. Based in Fremont, CA, FORTEL is an innovative supplier of performance management software and solutions for E-Business.