META Group

META Delta

5 June 2003 Delta 2259

Performance Monitoring Versus Performance Analysis

Operations Strategies, Service Management Strategies

Glenn O'Donnell

The infrastructure and application management market is rich with tools to monitor performance metrics and report them in graphical form. Although this function is useful, many tools lack the ability to perform broader data analysis. The next generation of performance monitors is emerging and provides additional insight through advanced methods such as statistical analysis, long-term pattern recognition, and multivariant composite processing. This added insight enables better understanding of performance issues, which can be used for problem identification and performance optimization.

Performance monitoring is a relatively mature cornerstone of infrastructure and application management (IAM). However, most monitoring tools have progressed little beyond data collection and presentation. Advances in user interfaces, threshold exceptions, real-time visualization, and ad hoc reporting are welcome evolutionary steps to basic reporting, but such features are quickly becoming commodities in IAM products. A few vendors are expanding (e.g., Opnet, Mercury Interactive, ProactiveNet, Entuity, BMC) or emerging (e.g., Fidelia, Quantiva, Network Physics) to provide the next step in the evolution of performance analysis products. Performance analysis is a sign of a mature environment, and the activity itself is proactive — seeking patterns in data, looking for areas to enhance, and tuning systems prior to additional problems. Because more than 90% of collected data is never used, this deeper automated analysis is more efficient. Many organizations disable data collection in servers and applications because of this data neglect.

Few IT organizations (5%-7% of Global 2000 [G2000]) currently perform automated, detailed analysis of performance data. Most generate reports with simple sorting to identify infrastructure and applications with worst-performance levels. By 2005/06, deeper analysis will finally augment simple monitoring in the majority (50%-60%) of G2000 IT organizations. Simple monitoring will never vanish completely, but the shift to deeper analysis is undeniable. Comprehensive exploitation of end-to-end relationships for highly automated analysis will be a difficult goal through 2010, though significant progress is already underway (see SMS Delta 1146).

Service-level management (SLM) is a prime directive of IT organizations, and companies are investing in SLM tools even with tight budgets. Measurement of services (i.e., application performance monitors, end-user monitors) is preferred over complicated assembly of composite component metrics, though many SLM tools merely report direct measurements without further analysis. As with any performance metrics, many of the same mathematical algorithms can be applied to SLM. The distinction between traditional infrastructure performance monitors and SLM tools is fading; the two are becoming intractably entwined into unified products.

The most positive outcomes of this unification are the common analysis engine and the ability to map relationships across data for multivariant composite analysis (see SMS Deltas 1044 and 1146). New insight is being gleaned from patterns in data streams and composite databases on relationships. Advanced mathematical methods rooted in such fields as statistics, linear algebra, and operations research consume large data sets and produce performance anomalies, trends, and systematic events not possible with simple reporting.

In most situations, absolute replacement of the organization's performance tools is difficult or impossible and, in some cases, unnecessary. This can create a dilemma when incumbent tools are acceptable for basic data collection and reporting but not for analysis.

META Trend: Through 2003-05, real-time status reporting will be demanded in context (e.g., performance versus the SLA — not element level) and include extended metrics (e.g., financial analysis, mean time to restore service) that drive abstraction technology. Through 2005, tools targeting service-level "everything" (reporting, verification, management, etc. — SLx) will remain immature and fragmented. During 2006/07, SLx technology will begin to mature (demanding XML-based data integration), supported by common metrics.

208 Harbor Drive • Stamford, CT 06912-0061 • (203) 973-6700 • Fax (203) 359-8066 • metagroup.com

We recommend organizations investigate adjunct tools (e.g., from ProactiveNet or SAS IT Resource Management) to perform additional processing on this existing data.

Substantial effort has been expended on performance reporting for human interpretation, but manually reviewing reports is time-consuming, wasteful, and error-prone. Much of the effort has been around data visualization and not necessarily data simplification. Companies should automate the analysis wherever possible. In effect, this automation is a correlation phase that consumes raw performance data and produces more relevant performance anomalies to trigger action. If a performance metric is safely within exception thresholds, there is little operational need to focus on that data source. This automated analysis filters out irrelevant data and enables the user to more easily focus on anomalies.

We recommend organizations distribute the processing of performance data to reduce the load on the central management server. As distributed processing matures, localized analysis can be more robust and timely (see SMS Delta 1145). Management appliances will prove to be useful packaging models for such distributed processing, along with other management functions such as localized discovery and event correlation (see SMS Delta 1054).

We also recommend organizations incorporate deeper analysis of performance data to bolster capacity management efforts (see SMS Delta 1098). Such analysis is central to the capacity management process, and any analysis models (i.e., relationship-based, statistical, or otherwise) used in performance management can be leveraged in capacity management or vice versa. In addition, the data reduction of the analyzed performance information is a more relevant input to the capacity management process.

Performance in context is necessary to understand business impact. The context can be a relationship to the specific business process. It can be temporal (i.e., performance thresholds are ignored during off-business hours), or it can be a more complex relationship to a specific technology issue (e.g., switched Ethernet near 100% utilization may be acceptable despite conventional wisdom). The point is that static performance thresholds for unacceptable behavior are useless in a vacuum. The relationship of the performance values is what makes a situation notable. Organizations must establish these relationships for successful performance analysis.

Much of this analysis is still immature, so an important intermediate step is to deploy performance monitors to enable flexible interactive navigation of the collected data. This is especially important when analyzing related data sets (e.g., network traffic flows showing raw performance metrics, traffic sources, and applications). Although there is little automated analysis in such products, the interface for relating this data and altering the relationship chain is a useful aid to manual analysis.

Many performance monitors are limited to reporting short-term data sets (e.g., several minutes, a few hours). Better monitors will also perform long-term data reduction to coarser samples (e.g., daily or weekly averages) to allow for more efficient reporting of long-term data (e.g., monthly, yearly). Although this is helpful, it involves little actual analysis. However, the long-term domain presents an opportunity for an important analytical technique. Long-term pattern recognition can identify chronic performance anomalies that can be difficult or impossible to identify with simple performance monitors. These patterns can sometimes be visible from reports (e.g., a periodic spike each Friday at 2:00 pm), or they can be obscured inside apparently normal data. Advanced mathematical methods can extract these hidden anomalies, though such functions are still emerging from research efforts and will evolve slowly within commercial products.

Bottom Line

We recommend organizations investigate advanced analytics to enhance performance monitoring and enable more efficient operational processes and process integration. Analysis products are only now emerging and will continually develop as algorithms and processing systems evolve.

Business Impact: Automating IT performance analysis increases operational efficiency to reduce costs and improve services.