

Taking the Blame out of Performance Management

InfoManagement Direct, January 2002

Chuck DeLouis

One of IBM's TV commercials shows a group of people sitting around their company's conference table on a Sunday afternoon. Everyone's pointing fingers at everyone else and at their IT vendors – the network provider, the service manufacturer, the software supplier and so on – because a catastrophic failure just shut down the company's online business. The reason the commercial is so effective is because the situation it portrays is so common. What often passes for performance management is really a blame game – the results of which can be incredibly destructive – both in terms of the damage done to relationships and the lost business that results until a corrective solution is found.

The easiest thing to do when this game is played is to blame those who are playing it. But to do so misses the real lesson: it is the IT environment itself, especially in the marketing data management arena, which sets people up to fail at IT performance management. Just making people smarter will not necessarily lower the threat. The context must also change. Otherwise the game is about survival of the fittest, not survival of the business.

The Threat Environment

Performance management is especially challenging given today's technical topography. Calling the IT environment multitiered really doesn't do it justice. In addition to layers, there are also numerous fissures (or silos), the sheer number of which makes a defensive view of the landscape almost unavoidable. People are constantly bombarded "over the horizon" with problems they perceive as coming from outside their specific areas. What's worse is that their own performance often gets criticized because of the effect these external problems are having on their internal capabilities. It would be bad enough if a slow disk array were dragging down database response times. But what really hurts is when management blames that degraded database on the database administrator. The result is a very real "us versus them" competition, which makes performance problems more likely to occur rather than less.

[Integrating Existing Applications](#)
Any technology application or integration solution must be able to interoperate with existing applications. This is why Open Source data integration is a key [paper](#).

Written by Randy Cronk
greatwriting.com

The effects of these silos are magnified in market data management by the very ad hoc nature of the work. Whether the application is a customer-facing storefront, a service-desk CRM tool or an internal data analysis and campaign management application – the impact on IT is often

quite unpredictable. It is in the nature of data warehousing, for example, that users are frequently trying to correlate what were previously unrelated pieces of data. Very often marketers will then build campaigns that leverage these new relationships. The upshot is that different pieces of the technical infrastructure must interoperate in previously unanticipated ways, causing perhaps unanticipated performance problems. An application that wasn't I/O bound yesterday might be I/O bound today – not because the application has changed, but because it generates a different SQL statement that perhaps uses network bandwidth or storage resources or some other resource in an adverse way.

The effect on the user (and the business) is equally frustrating: intermittent and unexplained breakdowns in the ability to use data. The results are especially insidious because analysts who wait too long for answers often find themselves unable to ask good follow-up questions of the data. That means the business mission is blunted. Fault isolation is also hampered because events leading up to the failures are often difficult to recreate. People don't remember exactly what they did (not that tracing IT problems is something users ought to become involved with anyway). Nor is improbable that "canned" diagnostics, i.e., those that rely on synthetic rather than historical data, will recreate these events either. The fact that a specific resource "checks out" as okay suggest to the resource owner that the performance problems most likely come from somewhere else. Of course, all the other silos run their own resource-specific tools, too. None of them takes into account the end-user's experience because none provides end-to-end visibility of the infrastructure. Yes, this is the experience that ultimately counts most.

Marketing Analytics

The technical topography won't change anytime soon. That doesn't mean that the infrastructure can't be healed, i.e., that the silos and tiers can't operate as a coherent whole greater than the sum of its parts, at least with respect to performance management. The key ingredient is one that should resonate well with market data management specialists. Just as marketers need to funnel data from many different sources to a central point for analysis, so too do IT managers looking to improve infrastructure performance. A marketing analyst, for example, might take credit bureau data slice it by specific demographics, correlate it against a database of mortgage holders living in certain ZIP codes and apply previous campaign results to derive probable buying behavior for given target segments. In a multisegmented marketing universe, no one would expect to conduct a campaign off a single data set. The same logic applies to managing performance in a multisegmented IT infrastructure.

Suppose a marketing analyst experiences slow response times while working with a specific set of files. Ideally, an IT analyst would be able to look at a console and see the relative impact that this research activity is having on various parts of the IT infrastructure from one end of that infrastructure to the other. Thresholds are "relative" in two dimensions – first, with respect to what is defines as a "normal" level of performance for that particular resource, i.e., disk I/O; second, with respect to other resources. In other words, not just which resources are running "in the red zone," but also which ones are running more in the red zone (or yellow) than the others. That is possible, of course, only if thresholds have already been put in place for every link in the performance chain: disk arrays, networks, database systems, application servers,

Web servers and so on. This, in turn, implies that there is some trend data collection going on to see just what “normal” means for each of these resources in this infrastructure. Another benefit of trend data collection is the ability to see where the trend lines are moving, and how fast, in order to foresee where potential future problems might occur.

But seeing what impact a user’s activity is having on a specific resources is only one step. The next step is to see what impact these resources are having on each other and why. Correlating activity in one area with activity in another is key. If we can see, for example, that a file or set of files are performing poorly, we should be able to go to the database and see what activity in the database corresponds to the suspect activity in the files. Database activity can then be traced back to application activity, for example, an errant SQL statement. Suppose there is a Web server in front of that database which, in turn, is causing the specific SQL statement to get executed in the database which, in turn, is causing the poor file performance.

If, on the other hand, we were to just look at storage, we might only see the obvious: that some files have bad I/O. In that case, we might have simply tuned the file system – perhaps scouring the disk for unauthorized MP3 or MPEG files or perhaps asking our storage vendor to take a look at its allocation algorithms. What would be missing would be a view of the problem from the application’s perspective – i.e., how resources affect one another as they are actually being consumed. By correlating data across resources, we attack the problem where it lives: perhaps changing the way SQL gets generated or avoiding a suspect SQL.

A Cycle of Business Improvement Performance

Performance is much easier to manage once there is a view of performance data that does two things: 1) captures historical trends against resource thresholds; and 2) correlates real (as opposed to synthetic) activity in one resource with activity in another. This type of visibility is a function of the performance management technology installed. The key ingredient is the ability to reach into disparate parts of the architecture and correlate disparate data. That ability results primarily from the technology’s own architecture – which precisely conforms to the infrastructure on which it resides. Distributed resource-specific performance monitors are required, but can only work if they pool their information in a common warehouse for access by a common set of analysis tools.

This changes the perceptual landscape and the effect is remarkable. Because it is now easier to take the large view, that is exactly what people tend to do. It’s both empowering and interesting to see how everything impacts everything else. A self-perpetuating cycle of business performance improvement begins. Greater visibility begets greater cooperation which begets greater visibility. Rather than obstruct performance improvement, the environment now encourages it. Best of all, there is less time spent managing and more time spent performing.

Chuck DeLouis is vice president of Professional Services for Precise Software Solutions, a company that focuses on optimizing business through application performance management. DeLouis has extensive technology and MIS experience in large-scale data warehousing, database design and modeling. He can be reached at cdelouis@precise.com.

For more information on related topics, visit the following channels:

- DW Administration, Mgmt., Performance
- Enterprise Information Portal (EIP)

©2010 Information Management and SourceMedia, Inc. All rights reserved.

SourceMedia is an Investcorp company.

Use, duplication, or sale of this service, or data contained herein, is strictly prohibited.