



# A Taxonomy of Metrics for Hosted Databases

Journal of Internet Banking and Commerce, April 2006, vol. 11, no.1  
(<http://www.arraydev.com/commerce/jibc/>)

---

By Jordan Shropshire, Doctoral Student/Mississippi State University

Email: [jds372@msstate.edu](mailto:jds372@msstate.edu)

Jordan Shropshire is currently a doctoral student of Management Information Systems at Mississippi State University, USA. He has a undergraduate degree in Business Administration from the University of Florida. His research interests center around Information Security, Diffusion of Innovations, and Technology Adoption.

---

## Abstract

The past three years has seen exponential growth in the number of organizations who have elected to entrust core information technology functions to application service providers. Of particular interest is the outsourcing of critical systems such as corporate databases. Major banks and financial service firms are contracting with third party organizations, sometimes overseas, for their database needs. These sophisticated contracts require careful supervision by both parties. Due to the complexities of web- based applications and the complicated nature of databases, an entire class of software suites has been developed to measure the quality of service the database is providing. This article investigates the performance metrics which have evolved to satisfy this need and describes a taxonomy of performance metrics for hosted databases.

---

## Introduction

Enterprise information management has changed a great deal over the three decades. Throughout the seventies and eighties, many organizations developed proprietary systems in-house (Wustenhoff, 2002). By the mid-nineties, the advent of the client-server paradigm shifted attention towards enterprise resource planning systems, and their associated applications. The majority of these systems were purchased from software vendors such as SAP or Oracle. Organizations are now in the midst of yet another IT paradigm shift, one oriented towards outsourced services, such as managed databases or client relationship management systems (Callaghan, 2002). Companies are entrusting some or all of their information requirements to third-party organizations such as Salesforce.com. These organizations are known as managed service providers because they fulfill the information needs of their clients by providing them access to enterprise applications which have been configured to meet their requirements (McKean, 2003). Data storage and processing, not actual software, is the provided service.

As the paradigm has shifted towards larger and more complicated agreements between

companies and IT vendors, contracts have become invariably complex. These agreements between customer and service provider are known as service level agreements (Wustenhoff, 2002). Oversight of the execution of these contracts is equally challenging. The need for adequate performance metrics and means of appraisal is quite salient. This statement may be evidenced by the number of IT firms who now sell service level management software (Schwartz, 2004). However, the extent to which these programs adequately manage and assess the service level agreement is unknown.

The purpose of this article is to explore key issues related to performance metrics for hosted databases. A review analysis of relevant literature on service level agreements and service level management is conducted in the next section. A taxonomy of performance metrics for hosted databases is then developed and explained.

## Background

Before proceeding to study performance metrics for hosted services, it will be necessary to provide some background information. In order to present a more complete picture of the current trend, this section includes a review of IT service contracts.

The hosted services paradigm presents complexities not previous encountered in the organizational computing domain. These complexities stem from two areas, changes in information technology, and changes in the management of information technology (Clyman, 2004). Enterprise reliance on web-based applications is a relatively new phenomenon; there is a lack of accepted performance measures (Schwartz, 2004). In addition, the increased reliance on IT vendors as providers of enterprise information storage and processing has shifted attention to the need for improvements in the client-vendor relationship. A crucial innovation which addresses these issues is the service level agreement.

The service level agreement (SLA) was originally adopted by IT vendors as a competitive advantage; customers were more comfortable dealing with vendors which guaranteed a specific quality of service (Sage, 2001). Before its adoption by IT vendors, the SLA was probably most associated with telecommunications service providers. Local and long distance phone companies have long included quality of service agreements in contracts. Cellular service companies include specific performance metrics in their SLAs. One of the most common metrics wireless providers use is the number of calls dropped. For example, some carriers, such as Alltel, include a SLA stipulation of one free minute for each call dropped.

• Technical support	• Storage capacity and server resources
• Billing procedures	• Customer data privacy protection
• Escalation of unresolved issues	• Data integrity
• Emergency maintenance	• Data retention
• Hardware replacement	• Customer responsibilities
• Network availability	• Total penalty limits
• Network monitoring	• Web server monitoring
• Web server availability	• Network and server availability

**Table 1: Typical components of a web-hosting service level agreement**

It seems intuitive that the means for evaluating hosted services would draw from telecommunications service providers. Some of these companies provide additional services such as web-hosting. Often, those desiring an internet presence would simply contract with their internet service provider to add web-hosting to the services which they are already subscribed (Wustenhoff, 2002). The ISP describes the quality of the web-hosting service through a comprehensive service level agreement which typically treats many aspects of the proposed service (see Table 1). When ISPs began offering additional services, such as email hosting and spam filtering, the SLA was most likely adapted to fit the conditions unique to those services.

In addition to the incumbent internet service companies, new organizations were formed to compete only on hosted services. The service level agreements in the current era of enterprise information systems relate more to the businesses which offer hosted services than the traditional internet service providers (McKean, 2003). These businesses are sometimes referred to as application

service providers (ASP). ASPs are often some derivative of a traditional enterprise software company (Kresch, 2006).

The service level agreements for a hosted service such as a web-enabled database will be much more complex than an SLA for standard website hosting (Schwartz and Cowley, 2004). The arrangement between the organizations must treat issues associated with all hosted services, such as uptime, availability, and maintenance, as well as features unique to the hosted database, such as data integrity, information security, query capability, and database interface customization.

## Methodology

The purpose of this research is to identify the performance metrics which are used to assess the quality of service of a hosted database. In monitoring a hosted database, these metrics are usually collected and analyzed by a software application acquired by the service provider (Lamonica, 2000). This class of software suites may be referred to as SLM, SLA, or QoS software. The first step in the methodology is to identify these applications and determine what they are capable of monitoring.

Using several popular internet search engines and phrases associated with hosted databases, a number of applications which are capable of monitoring hosted databases were identified. For the purposes of this study, the search process continued until twenty relevant programs were identified. Potential applications were screened to ensure that (1) they had the capabilities of monitoring network services, and (2) they possessed the ability to monitor databases.

Following initial screening, a second round of investigation was conducted in order to identify the monitoring capabilities of each system. The purpose of this exercise was to examine the various means by which an SLM application could monitor the performance of a hosted database. For the detailed investigation, as well as initial screening, information about the applications was gathered from the website of its parent company. White papers, technical briefs, webinars, product manuals and demonstration videos were all accessed in order to assess the capabilities of each suite. Only information related to the program of interest was included in the study, and no third-party reviews or external websites were included.

## Results

Once the monitoring capabilities of each program were identified, the next phase in the research program consisted of organizing the capabilities into a structure which could be used as a taxonomic guide. Each metric identified in the procedure described above was organized into the most appropriate category. Category assignment decisions were based on the attributes of the measures, and the stages in the hosted-service model they apply to (see Table 2). This recursive technique has been consistently used for developing and populating taxonomies.

The results of the grouping exercise include the creation of five categories: client, network, server, application, and support metrics. Each of the newly identified categories is explained and assigned meaning in order to increase the utility of taxonomy. The descriptions provided for each category are a summary of the combined capabilities of the related performance metrics.

The first category in the taxonomy is client metrics. These performance metrics are used to measure end-to-end performance of the hosted database. Only six of the twenty service level management software suites had performance measures for clients, and most were simply identified as client. They provide end-user response times, transactions per second, hits per second, and other user status measures. They are used to gauge the level of user activity, and can be used to predicate usage levels on servers, networks, and hosted applications. These metrics are useful for determining when to supply extra capacity. Roughly ten percent of the metrics identified in the study could be classified as client-side metrics.

Client	Network	Server	Application	Support
Client Monitoring-7	Availability-4 Bandwidth-2	Memory-4 CPU Load-5 Hard Disk Space-3 Event Monitoring-3 Server Availability-7 Server Security-1	Availability-10 Query Processes-4 Database Monitoring-2 Response Time-4 Transaction Processing-3	Help Desk-3 Change Management-1 Financial Impact-1
7 10.29%	8 11.76%	23 33.82%	25 36.76%	5 7.35%

**Table 2: Towards taxonomy of hosted database metrics**

Network measures are the second component in the taxonomy. They represent 11.76 percent of the total measures identified in the analysis. The principal measures for networks are bandwidth and availability. The bandwidth measures represent the speed at which data can flow through a network. Due to the data-intensive nature of databases, the high speed data transmission is critical component for success. The availability measures indicate whether the network is online or offline. They allow for the isolation and resolution of telecommunications systems hardware issues, such as faulty switches, routers, or cabling.

The second largest category in the taxonomy is that of the server performance monitors. The general purpose of these metrics is to determine the extent to which the server is being utilized. Too much strain of the server system will lead to lagging application performance. There are many, many unique metrics to analyze various aspects of the server. Memory monitors measure the amount of logical memory being used by the server. CPU utilization monitors report the percentage of the central processing unit that is being utilized. Hard disk space is an especially important metric for hosted databases because they rely heavily on stable, long-term memory. Server event monitors include user logons, number of client requests, and number of errors processed. Server availability is a composite measure of the server hardware components. Every application suite included in the study had several server measures. In total, 33.82 percent of all the metrics fall into this category.

The category containing the most metrics was applications. These contained over thirty six percent of all the measures used by the service level management software. Application metrics are concerned with the availability and adequacy of the required service. All of the software programs included in the analysis had features specifically design to monitor and evaluate the performance of database applications. Of all the metrics, availability was the most common. Availability measures are meant to indicate the extent to which the application is being used, and thus not available. This is mainly indicated by the number of non-idle application threads being currently being processed. Query processing is an indicant specific to database information storage and retrieval processes. Other statistics include the total number of logins started per seconds, logouts per second, and number of users connected to the system. Because of the data-intensive nature of databases, it is important to manage the applications memory utilization. Associated measures include lock memory, target server memory, and maximum workplace memory.

The final category in the taxonomy is support. Metrics in this category are primarily concerned with the administrative aspects of the hosted service. Although this category only contains seven percent of the total metrics, it includes measure of several important aspects of the provided service. Furthermore, support measures might be included in the service level agreement. The support metrics observed in the software suites include helpdesk, change management, and financial performance. Helpdesk operations are an integral part of hosted services. Metrics may include the number of tickets processed, total tickets, and unresolved tickets. Change management measures are usually composite metrics. They are used to guide service improvement projects. Because helpdesk staffing expenses can quickly escalate, there is much incentive to calculate financial measures and try to improve helpdesk efficiency. Such metrics might include cost per ticket, cost per client, or cost per application hosted.

## Conclusion

An analysis of the performance metrics of twenty service level management software suites was conducted in this research. Each of the software suites included in the study possessed monitor capabilities specific to hosted databases. These programs are used to monitor the quality of service

being provided by the organization hosting the database. It is important to understand how the provided metrics relate to the service level. Because the software suites provide are many performance metrics for assessing the quality of service being provided, a system for organizing and understanding them is critical. This paper presented a step towards that goal by providing a taxonomy of performance metrics.

## References

- Callaghan, D. (2002). "CRM providers go custom." eWeek December 23, 2002: 15.
- Clyman, J. (2004). CRM on demand: Have it your way. PC Magazine. July: 62-63.
- Ferguson, R. (2004). "Hosted CRM choices emerge." eWeek November 1, 2004: 15-17.
- Ferguson, R. (2005). "Hosted services gain ground." eWeek February 7, 2005: 26-27.
- Institute, I. G. (2002). COBIT 4.0. Rolling Meadows, IL, IT Governance Institute: 1-209.
- Knorr, E. (2005). "Software as a service." Infoworld.com: 40-42.
- Koblentz, E. (2000). "Hosted Services brings FAQs to life." eWeek September 4, 2000: 40.
- Kresch, M. (2006). "Advantages of hosted services." Business Communications Review February 6, 2006: 64-65.
- Lamonica, M. (2000). "Handicapping an IT industry battle royale." Infoworld.com December 18, 2000.
- McKean, K. (2003). "The rise of ASPs highlights what enterprise software is all about." Infoworld.com: 6.
- Montalbano, E. (2006). "Verio airs Windows hosting services." NetworkWorld.com February 27, 2006: 12.
- Rapoza, J. (2003). "IBM package gets e-commerce package right." eWeek September 15, 2003: 63-66.
- Sage (2001). Service Level Agreement Requirements: A Qualitative Assessment, Sage Research: 1 - 292.
- Schwartz, E. (2004). "Salesforce.com pushes integration." Infoworld.com November 8, 2004: 24.
- Schwartz, E. and S. Cowley (2004). "Hosted CRM extends its reach." Infoworld.com December 3, 2004: 2.
- Wustenhoff, E. (2002). Service Level Management in the Data Center. Sun Blueprints. Santa Clara, CA: 1-15.