

MySQL 5.0 Vs SQL Server 2005

A dilemma mired in common myths, pitfalls and stray decisions

Written by

Vishwenath S. Kizhapandal
iLink Systems, Inc

The release of SQL Server 2005 and MySQL 5.0 continues to fan the debate on the suitable database development platform of choice. There are well established camps on both fronts armed with several studies, performance benchmarks, feature comparisons and counter arguments. These camps and their arguments fall anywhere in the spectrum ranging from the open source fanatics to the commercially vested interest groups. To a large extent, the already “enlightened” crowd and the plethora of information on both sides have contributed to the dilemma in the minds of the IT decision makers. While it is futile to preach to the enlightened, the well intentioned and neutral decision makers are often led astray due to the common myths, pitfalls and focus on the wrong decision points.

The purpose of this paper is not an elaborate feature by feature comparison of SQL Server 2005 and MySQL 5.0. There are several online comparison matrices that are readily available that compare and contrast the features. In fact in most cases, such comparison and technical differentiation are primarily of academic interest and the technology selection is primarily governed by the common myths. Hence, some of the more common myths are first discussed. However, certain technical factors that might affect the platform choice for enterprise systems are briefly highlighted in the second section. In the final section we recommend a more business centric approach to the selection process. More often the decision makers focus on licensing cost even before they have a product roadmap or deployment plan. Irrespective of the cost incurred or saved, the lack of a roadmap or a plan is more harmful to the business solution than any technology, licensing or cost savings. Instead, having a well thought out development, scaling and product life cycle plan would enable the decision maker to better evaluate the competing platforms and make an informed choice.

Common Myths

Let us first visit the common myths.

SQL Server is expensive while MySQL is free

MySQL is open source software that is indeed distributed under the GPL license. But this version is uncertified, unsupported and is generally not considered ready for production use. This is sometimes generally used to discredit MySQL as well. However, MySQL AB the company behind the development of MySQL does provide certified software with commercial licenses for production and enterprise use purposes with a technical help desk for production support. These licenses and support can cost anywhere between approx. \$500 to \$5000 per server based on a tiered pricing model and cost of ownership now will depend of your hardware configuration and support level requirements.

With the launch of SQL Server 2005, Microsoft has also launched SQL Server 2005 Express edition that replaces the older MSDE. This version can be distributed freely under the distribution agreement. Microsoft has a range of pricing and version options on SQL Server 2005 which caters to a low cost of acquisition approach that can be seamlessly scaled with business and volume growth. SQL Server 2005 is actually free for “development use only” . And SQL Server Developer Edition can be purchased for approx \$50. The Standard and Enterprise Edition are priced using two models, per Machine or per CPU. Depending on the pricing model the price can vary from \$1400 to \$13,000.

SQL Server Express Edition is a very limited edition

Obviously Microsoft has imposed limitations on SQL Server Express, but it is not limited in terms of SQL Server features. An application developed using any edition of SQL Server can be deployed and distributed to run on SQL Server Express seamlessly. The Advanced Services edition also comes with a full featured management studio and even includes the Business Intelligence developer studio. The Toolkit Edition provides the necessary development tools if Visual Studio 2005 is not already available to the developer. It also overcomes several performance bottlenecks in MSDE (no workload throttles on > 5 concurrent tasks, no connection limits etc.). The Express edition does have some limitations (< 4GB and one processor usage) in terms of database size and processor usage but it might suffice a range of applications such as data driven websites, data collection clients in distributed systems etc. So if you are looking to just have data collection functionality in your field applications or want to use it as a data store in your data driven websites it should very well fit the bill. As long as you do not store large graphics content in your database, one would be surprised how much data can be stored in 4 GB. That alone is worth the exercise to size databases with similar requirements. iLink Systems developed a distributed vault management application for one of its clients. The distributed clients store their data in a SQL Server Express database which is synchronized with the central servers. The development was done using the Developer edition of SQL Server and the database was seamlessly deployed in Express edition for distribution along with the client installs.

MySQL lacks key RDBMS support

MySQL did not support several key features of RDBMS by design in the prior versions. But support for cursors, views, triggers and stored procedures have been implemented in MySQL 5.0. Referential integrity is also possible with the implementation of foreign keys support. However, these new features in the latest release have not fully proven itself in a large production install base and the features have not been fully implemented in all the different underlying data stores in MySQL.

MySQL is so simple to use, faster to develop and administer without additional resource costs

This is a more complicated myth with several variations, some of which are contradictory and even counter intuitive. In general it is a corollary claim to the primary myth that it is free and readily available. Without questioning the larger wisdom of saving labor costs on a DBA or assuming that programmers always write best queries or the lack of need for data modeling it can safely be said that SQL Server 2005, Visual Studio 2005 and the .Net framework definitely provides an equivalent if not superior set of tools to develop, administer and manage the database and applications.

Performance and benchmarking myths

There are numerous benchmarks on performance on both sides that add to the confusion. In general, MySQL is faster for “read only” or “read mostly” data access and SQL Server is a more scalable, high available database for high transactional volume. However, most benchmarking studies assume reference application or hardware configuration that might not remotely mimic the real world system to be built or might be beneficial to the internal implementations of a particular database engine. SQL Server 2005 also inherently has several advanced features that require proper tuning for performance. Hence, benchmark studies are not a primary indicator to make a decision.

Some Technical Differentiators

SQL Server 2005 has had a head start over MySQL 5.0 as a robust and more mature RDBMS with a stabilized set of advance features. Hence it is typically the preferred choice for enterprise applications that require some of the more advanced features such as replication, clustering, recovery and high security. MySQL 5.0 might suffice the needs of small to mid tier applications. However, the term “enterprise level” or “mid tier” is loosely defined. Some of the technical differentiators that might affect the selection process are addressed below

Performance & Platforms

MySQL performs well on a variety of platforms. Particularly, on UNIX based hardware platforms it can take advantage of some of the hardware specific features. However, it does sacrifice features for the performance gains. Some of the storage engines does not support referential integrity or ACID transactions or provide less data security. For read only or read mostly application (s), where these features do not matter MySQL might be ideal. SQL Server on the other hand is highly scalable and has the fastest performance on the Windows platform achieving the highest

TPC-C result ever recorded on Windows. However, SQL Server does have an advanced set of features that might sacrifice performance for better data integrity, recovery or clustering support if it is not specifically tuned for these scenarios.

Security

MySQL does provide some basic security mechanisms but lacks the robust security features offered by SQL Server. SQL Server 2005 is C2 compliant, hence for government applications will become the automatic choice over MySQL. SQL Server also provides native encryption support in the database using a combination of certificates, asymmetric and symmetric keys. MySQL does not have such provisions.

Replication

MySQL does offer a decent replication mechanism based on its binary logs. However, SQL Server provides more advanced replication options such as transaction or merge replication options. iLink Systems implemented a distributed application at one of its clients that required data stored on the web portal to be synchronized frequently with a central repository. Hence transactional replication was a requirement and SQL Server was the automatic database choice in this case.

Integrated development and Management environment

SQL Server 2005 has a more advanced integrated development, deployment and management platform due to its seamless integration with Visual Studio and the .Net CLR greatly reducing development and maintenance time and costs.

A Business Approach

The different myths surrounding the debate mires the decision making process when choosing between the two competing platforms. More often than not, the plethora of information leads to the intricate details of the two platforms with respect to their strengths and weaknesses. However, a more business driven approach to the selection process might view the problem in an entirely new light and expose some of the common pitfalls, the decision makers must avoid.

Costing Vs Planning

The first common pitfall that one must avoid is short term costing in lieu of long term planning. One of the most common lures is the need to save on licensing costs with MySQL. But the goal is to save licensing costs or reduce the total cost to build, own and operate. Typically, people have an end state architecture and deployment in mind and try to achieve or plan for it with their

current budget/revenues. Licensing costs is typically a lot smaller compared to corresponding labor costs to build and maintain any non trivial production system. In order to avoid this pitfall we need to first create and then evaluate the nature, purpose and life cycle of the application or system. This will enable the decision maker to not only have a better business plan but also leverage the SQL Server 2005 low cost of acquisition based pricing models that can grow with your business.

Database Vs Business Intelligence

What is the purpose of the data? What is the need in storing the data without analyzing it? SQL Server 2005 comes with Analysis Services, Integration and Reporting services. The Business Intelligence stack seamlessly integrates and interoperates with the .Net development platform. The Visual Studio environment even provides the ability to seamlessly step into the stored procedures from the application code while debugging. Even SQL Server Express edition has full support for reporting services and comes with a Business Intelligence Developer Studio. This integrated development environment will not only reduce development time but also enable the business to properly mine and use the information.

Enterprise Scalability and Migration Costs

SQL Server 2005 is a more robust, secure and scalable database for enterprise grade applications. For small or startup business with a good plan for growth, scaling will be a key issue in the technology selection. Though there are tools available to migrate from MySQL to SQL Server, it can still be time consuming and costly based on the size of the production application. iLink Systems is currently in the process of developing the next generation product on the .Net Platform for a large continuing education client . One of the key initiatives involved the conversion and migration from MySQL to a SQL Server 2005 database. The subsystems that constitute the client platform could have easily used SQL Server Express in which case the upgrade to SQL Server 2005 would have been seamless without a need for migration costs.

Conclusion

SQL Server 2005 is the ideal platform for the development of enterprise grade applications. The scalability, security, robustness, clustering and replication support along with the power of the Microsoft BI stack and the .Net platform interoperability makes it the better choice. For applications including data driven websites, distributed data collection clients, and prototypes; SQL Server Express might be a good fit for future seamless upgrades. MySQL might be a choice when a size limitation of SQL Server Express is an impediment.