

Database Connectivity: Vendor Selection Criteria

An Application Architect's Guide

Executive Summary

In the first part of our *An Application Architect's Guide* series, *Data Connectivity for Critical Business Systems*, we described the integral role that the database connectivity layer provides in critical business applications. These critical business systems have zero tolerance for delays or errors relating to accessing, processing, and storing data. Inefficient or unreliable data connectivity can lead to poor performance, availability, and scalability, and to data integrity issues. Application issues can directly impact the revenue, cost, and risk structures of the organization.

This paper outlines the data connectivity options available to software architects. It should be used as a guide to ensure that the database connectivity options selected at design time meet the requirements of production systems.

Key Highlights

- Your database connectivity choice plays a critical role in the overall success of an application, greatly affecting production performance, reliability, and scalability.
- Choosing the right data connectivity option at design time can dramatically reduce the time spent developing, testing, and deploying critical applications. Architects are best served by making a pro-active, conscious decision regarding their data connectivity solution.
- The default connectivity options provided by database vendors frequently fall short in terms of quality, reliability, and performance requirements of critical systems. This results in additional application risk while driving up the overall development and maintenance costs of the application.
- Open source projects are appealing because they have no upfront cost. However, a variety of technical, legal, and support issues make open source alternatives impractical for critical production systems.
- Independent data connectivity vendors are the preferred choice for architects designing systems with significant requirements for performance, reliability, and interoperability.

Database Vendor-Provided Drivers/Data Providers

All relational databases are shipped with drivers that provide basic database connectivity. Database vendor-provided components provide a minimal level of functionality and performance needed to use the database engine. Because they are included with the database, vendor-provided components often are seen as the most cost-effective choice for new application development. However, these components have several critical limitations, which can have a significant impact on project success, and development and support costs.

At the heart of these limitations is one key fact: database vendors are in business to develop and market database engines. They distribute drivers because they need to provide developers a way to connect to *their* database. Bearing this in mind, application architects should look closely at the following attributes of the drivers provided by the database vendors:

- Application interoperability
- Performance/scalability
- Quality/reliability
- Feature support

Application Interoperability

Database vendor-provided components are typically designed to work with a specific version of a specific database. Database vendors generally support multiple versions of their drivers, which correspond to the supported releases of their database engines. For example, an application developed against an Oracle-provided driver will *not* work with an IBM DB2 database. It also may not work with more than one version of Oracle. Because many applications need to work with a variety of databases and database versions, this creates a problem for application developers, testers, and support personnel.

Additionally, because database vendors are in business to sell more database licenses, they have a strong incentive to make it difficult for customers to switch to another vendor's database engine. In short, they want to create "lock-in." To achieve this, database vendors often implement proprietary extensions to data connectivity standards. Using proprietary extensions for data connectivity eliminates the benefits of using a standard API and reduces the flexibility of the application.

Performance/Scalability

Database vendors focus on the performance and scalability of the database engine itself, not on their drivers and data providers.

Quality/Reliability

Not surprisingly, database vendors invest the majority of their testing and QA efforts on their database engines, not on their database connectivity components. Database connectivity components are usually only tested for

specific functionality within a specific database version. As a result, quality suffers. Poor quality in the database connectivity component can cause databases to return unpredictable results or cause certain functionality to simply not work. Vendors are unlikely to delay the release of the database engine because of bugs in the database connectivity component.

Feature Support

As an architect, you might think that database vendors would offer the most feature-rich database connectivity components available because of their knowledge of the database engine. However, this is not the case; database vendors typically focus on adding features to the database engine. When we review the capabilities included in the data connectivity specifications, the lack of features provided by the data connectivity components offered by database vendors is glaringly obvious. In most cases, they lack significant features that are defined in the latest specifications. This means that critical systems that require important features such as key transactional capabilities, resource pooling features, and so on cannot use this functionality.

Open Source Drivers/Providers

Several open source communities are developing and distributing database connectivity components on an ad hoc basis. Unlike commercially supported or foundation-backed initiatives such as Apache, Linux, and JBoss, open source database connectivity components lack a model of viable development, technical support, and legal indemnification. The initial cost advantage associated with these open source solutions is overwhelmed by additional development, support, and legal costs.

Open source solutions have the following key limitations:

- Software quality
- Technical support
- Legal risk

Software Quality

The data connectivity components developed by the open source communities depend on a small number of volunteer developers to build and test software. Due to this, functional depth, software quality, and performance capabilities of the open source solutions are extremely suspect.

From a quality perspective, the open source alternatives:

- Have a limited set of resources allocated to testing
- Do not leverage a rigorous or transparent model for testing
- Fall short in terms of testing certification

Although you could argue that these drivers and providers provide sufficient performance and scalability for small-scale deployments, understand that they lack a sophisticated architecture, performance tuning, and benchmark infrastructure, as well as high-scale production experience needed for today's critical systems.

To compensate for these shortcomings, the development organization responsible for critical systems is forced to expend additional development, testing, and deployment resources.

Technical Support

Reliable technical support is a basic requirement when working with software from an external source. In today's hyper-competitive business environment, reliable, multi-channel (phone, email, and web), 24x7x365 support is imperative for mission-critical applications. If reliable technical support is not available, the IT organization assumes a significant level of risk because even simple technical support issues can turn into major project delays, lost developer time, and application downtime.

Open source database connectivity solutions are not backed by organized technical support. At best, these projects provide online forums where developers can submit questions and bug reports. The developer has no assurance that a response will be provided, and no contract is in place that regulates the response time of the resolution. The organization must use its own development resources to solve issues relating to installation, configuration, performance, and reliability.

Legal Risk

Organizations relying on open source solutions for data connectivity assume a significant level of legal risk because they assume total responsibility for any potential software infringement. The legal risk is compounded by the fact that an amorphous group of open source developers contribute code with little oversight to copyright or patent infringement. Legal indemnification, which is commonly provided by commercial vendors, is unavailable. Without legal indemnification, the organization using the open source software, not the open source developers, is the potential target of any legal action.

In addition, the complexity of the various open source licenses requires that organizations be extremely diligent about how the drivers are deployed. If organizations are not careful, they run the potential of "infecting" their proprietary code, which forces their own intellectual property to be subject to the terms of the open source license.

Third-Party Vendors of Drivers/Providers

When performance, scalability, portability, and quality are important, third-party database connectivity components are the best choice for critical system deployments. Third-party components typically support required features without forcing lock-in to a specific database or database version. Within the category of third-party components, developers and IT managers should look closely at the following factors before selecting a vendor:

- Product comprehensiveness
- Production proven
- Technical support
- Technical leadership
- Corporate focus and strength

Product Comprehensiveness

The critical success of a database connectivity solution depends on:

- Robust functionality
- Timely support for the latest data connectivity specification versions
- Database and platform coverage

If a database connectivity component lacks functionality, the application capabilities are limited unless the application developers compensate for the missing functionality. To make sure that you do not hinder the effectiveness of your application, it makes sense to rely on a vendor that supports the newest functionality in the data connectivity specifications in a timely fashion and provides additional functionality outside of the specification, such as security features and developer productivity tools.

To successfully support database independence, a third-party vendor needs to offer a wide range of support for databases, database versions, connectivity APIs, and platforms. Otherwise, the advantage of database independence is lost. To see significant value, look for a third-party vendor who can service the widest range of requirements.

Production Proven

It's critical that you, as developers and development managers, look closely at third-party database connectivity solutions and the companies behind them. The key is to find a vendor whose products go through rigorous testing and have been deployed in thousands of environments.

For critical systems, an architect must select the vendor that leads the industry in terms of quality. This means relying on a vendor that consistently meets the certification process for connectivity standards, uses the most thorough test suite to validate connectivity components, and provides transparency in terms of testing results.

Another critical success factor relates to the number and quality of production implementations. It is important to select the vendor with the most proven track

record in terms of large-scale, mission-critical application success. Analyzing the number of independent software vendors that bet their business on a particular vendor's data connectivity components is a particularly effective method for selecting a vendor.

Technical Support

Technical support is key throughout the lifecycle of a critical system; it is not enough to simply provide support for the production environment. Project costs can easily escalate if timely, high-quality support is not available during the initial selection or pilot, design and development, testing and deployment, and ongoing maintenance phases of the project. Developers must carefully assess the availability of support (24x7x365), the commitment to service-level agreements, the overall quality and backing by the vendor's development team, coverage by locale and language, as well as the various modes of interaction provided by the third-party connectivity vendor (phone, email, fax, online forums and knowledgebase).

To ensure project success, it is highly advisable to select a vendor that is consistently recognized for its support organization.

Technical Leadership

Data connectivity is both critical and complicated. Data connectivity depends on knowledge of connectivity standards as well as the underlying database engines and protocols used to communicate with them. You should look for a vendor that has strong relationships with the top database vendors and significant ties to major standards committees. This ensures that the vendor has the knowledge and relationships to be quick to market with fully functional products. It is also critical that a vendor employs engineers, testers, customer support professionals, and development managers that are experts in data connectivity.

Corporate Focus & Strength

Architects should consider the overall corporate focus and strength of vendors when selecting a database connectivity solution. Significant corporate backing ensures that the data connectivity solution is supported with the requisite enhancements as the needs of the critical system evolve over time. Selecting a vendor with limited staying power can cripple the application or force the development organization to go through a time-consuming and costly replacement project.

As an architect, it is highly advisable that you select a vendor that is 100% focused on database connectivity. Many vendors sell data connectivity solutions as an afterthought, or they include them as an isolated offering within a disparate product line. These vendors simply won't provide the commitment necessary for a critical system deployment.

DataDirect Comparison Summary

DataDirect is focused on data access, enabling you to quickly create applications that meet the performance, reliability, and security requirements of today's critical systems. DataDirect provides an unparalleled range of data source support using standards-based interfaces such as ODBC, JDBC, ADO.NET, and XQuery. In addition to the corporate market, DataDirect is production proven in over 450 products from leading independent software vendors. The following table compares DataDirect's product offering to other third-party vendor options. For detailed information concerning these product and company differentiators, please visit <http://www.datadirect.com>.

Selection Criteria	DataDirect Technologies	Other Third-Party Vendors
Product Comprehensiveness	<ul style="list-style-type: none"> • Unrivaled product breadth, including being current with specifications and additional features that complement the specifications. • Unmatched coverage across APIs (JDBC, ODBC, ADO.NET), databases (Oracle, Microsoft SQL Server, DB2, Sybase, Informix, and more) and operating systems (Windows, UNIX, Linux, iSeries, z/OS). 	<ul style="list-style-type: none"> • Product breadth typically lags behind the specifications and complementary functionality is minimal or non-existent. • Support for APIs is often limited to a single API and database combination. Support for non-Windows, UNIX is minimal.
Production Proven	<ul style="list-style-type: none"> • Pervasive coverage throughout the industry, production proven in countless environments. • Quality proven through specification certification, largest independently owned test suite. • Huge customer base – proven in over 450 products from independent software vendors. 	<ul style="list-style-type: none"> • Production support is limited in large-scale, mission-critical deployments. • Products typically fail some portion of the certification test suite, internal testing process limited. • Customer base can be deceptively large based on number of project-based deployments.
Technical Support	<ul style="list-style-type: none"> • Unrivaled leader and four-time consecutive winner of the Omega NorthFace Scoreboard Award. • 24x7x365 phone support for all major countries. • Multi-channel support via phone, fax, email, web. • Largest industry knowledgebase including technical forums, 1800+ documents. 	<ul style="list-style-type: none"> • Support organizations not considered world class. • Phone support is limited to standard office hours for North America. • Email is typically the only alternative channel. • Web sites are inadequate for mission-critical support.
Technical Leadership	<ul style="list-style-type: none"> • Industry-trusted specification leader for JDBC, ODBC, ANSI SQL, and XQuery. • A trusted relationship with database vendors provides DataDirect with insight into the wire protocol for each database. 	<ul style="list-style-type: none"> • Other vendors do not participate in the standards process, limiting visibility into connectivity design. • Weak relationship with database vendors forces vendors to reverse engineer wire protocol design.
Corporate Focus & Strength	<ul style="list-style-type: none"> • Industry leader for 15+ years, backed by financially sound Progress Corporation. • 100% focused on database connectivity. 	<ul style="list-style-type: none"> • Other vendors have minimal corporate backing. • Connectivity solution is an isolated point solution within a disparate product line.

Summary

Database connectivity choices have material impact on the cost, performance, and risk profile of critical systems. The right choice can ensure that your application avoids downtime and performance issues while satisfying service-level commitments and keeping your users satisfied. The wrong choice can lead to application failure, disgruntled users, and damage to your reputation.

Architects must carefully assess the functionality, technical support, and the performance and software quality characteristics of each option to ensure the success of critical systems. This guide details the various database connectivity options that are available and provides information about the factors that influence the selection of the option that is best for your critical system.

For over 15 years, DataDirect Technologies has been the industry leader in data access technologies. Learn more about DataDirect and download free evaluation software at www.datadirect.com.