

WHITE PAPER

Don't Upgrade Blind: Critical Issues to Consider Before Executing a DB2 Upgrade



AUTHOR:

Craig Mullins

ABSTRACT:

Change is a fact of life, and DBMS products change quite rapidly. A typical release cycle for DBMS software is 18 to 24 months for major releases, with constant bug fixes and maintenance updates delivered in between major releases. Indeed, keeping DBMS software up-to-date can be a challenging job.

The fact is that after installing and using DB2 for Linux, UNIX and Windows for a while, your organization will inevitably have to decide whether or not to upgrade to a new version or release of the DBMS. How and when you approach upgrading DB2 should rest on multiple factors including the features in the new version, organization's risk tolerance, and other important criteria. In this white paper, we'll cover the important factors you need to consider before taking on the challenge.

content

Upgrading DB2 – a Version or a Release?	1
Understanding the Latest DB2 Versions	1
DB2 Editions	2
Understand the New Features	3
Complexity of the DB2 Environment	3
Creating an Upgrade Strategy	4
Skipping Releases	5
The Danger of Being Out of Support	5
DBA Staff Skill Set	5
Platform Support	6
Supporting Software	6
Migration Verification	6
Fallback Planning	6
Summary	6
About Datavail	7

Upgrading DB2 – A Version or a Release?

IBM makes a distinction between a “version” and a “release” of DB2. A new version of DB2 is a major undertaking, with many changes and new features. A release typically has fewer changes and not as many new features.

For example, moving from DB2 Version 9.5 to Version 9.7 is a release upgrade, which is typically easier to manage than a version upgrade, such as moving from DB2 Version 9.7 to DB2 Version 10.1. Usually significant functionality is added for version upgrades, less so for point releases. Nevertheless, upgrading from one point release to another can have just as many potential pitfalls as version upgrades. It depends on the nature of the new features provided in each specific release. And each version or release upgrade will have its own set of new and deprecated features to be analyzed and managed.

A DB2 version upgrade can be thought of as a special case of a new installation. All the same procedures required of a new installation still apply to an upgrade: you must plan for appropriate resources, reconsider system parameters, and ensure that all supporting software is appropriately connected.

Another consideration to be managed when upgrading (as opposed to initial install) is planning for existing users and applications. Care must be taken to minimize the impact and downtime experienced by DB2 applications during the upgrade process. And any impact must be broadcast to impacted business units and the timing of the upgrade negotiated with business requirements in mind.

Furthermore, any additional software that works with DB2 (such as purchased applications, DBA tools, utilities, and so on) must be verified to be compatible with the new version or release. Therefore, upgrading can be a tricky and difficult task.

The issues and concerns discussed in this white paper pertain to upgrading to a new release and to a new version.

Understanding the Latest DB2 Versions

The latest release of DB2 is DB2 11.1 for Linux, Unix and Windows, and it was released for general availability in June 2016. The most significant improvements in DB2 11.1 focus on enhancements in BLU Acceleration and pureScale functionality. But there are many new features and improvements that should compel organizations to adopt it quickly.

HADR

A significant improvement for organizations that demand high availability is support for HADR in pureScale clustered DB2 databases. HADR, or high availability disaster recovery, can be used to mirror data from a primary pureScale database cluster to a second local or remote standby pureScale cluster. Should an emergency cause your entire primary cluster to go offline, the standby cluster will be fully synchronized and can take over the workload in seconds.

Online Fix Pack Update Support

An additional improvement in DB2 11.1 for pureScale capability is the delivery of online fix pack update support. You can apply DB2 11.1 fix packs to individual members running DB2 11.1 in a pureScale cluster without having to take the entire cluster offline. By delivering HADR and online fix pack maintenance operations for DB2 pureScale clusters, you can improve the availability of your most critical systems and eliminate costly downtime for your end users.

BLU Acceleration Support

Support for BLU Acceleration in DPF environments for MPP scale out is another important new capability in DB2 11.1. For those unacquainted with DPF, it stands for Data Partitioning Feature, and is used to promote scalability by partitioning a database across multiple servers (or within a large SMP server.) By supporting BLU for DPF, DB2 11.1 can improve query performance for very large, partitioned applications.

With BLU MPP support, you can choose to distribute data using a hash key or a random distribution. Random distribution will try to evenly distribute data across the partitions.

Many SQL improvements are delivered with DB2 11.1. More types of SQL and optimization can be used in the BLU columnar engine, including native columnar OLAP functions (and others.) In terms of SQL compatibility, DB2 11.1 delivers improved compatibility with SQL on IBM PureData for Analytics, Oracle Database, and PostgreSQL.

Data Protection

Additional new availability and data protection capabilities have been added to DB2 11.1. On the security front, DB2 11.1 delivers the ability to use enterprise key management systems for storage of native encryption master keys. This means that you no longer need a local key store. DB2 11.1 also delivers encryption capability for the HADR replication stream using TLS.

Advanced Recovery Feature

Another compelling new addition is the DB2 Advanced Recovery Feature, which is a bundle of database backup, recovery, and data extraction tools that can help you improve availability, mitigate risk, and optimize administrative tasks.

Finally, for DB2 11.1, it is important to keep in mind the licensing and support changes that were made (as outlined earlier in this document.)

Although DB2 11.1 is the most recent version available, many organizations have yet to upgrade to DB2 Version 10.5, which has been available since June 2013. So if you are at an earlier DB2 release level it makes sense to learn the compelling features of DB2 V10.5 as well as DB2 11.1. So with that in mind, let's discuss a few of the most important new capabilities.

DB2 Version 10.5 New Capabilities

The most significant new feature of DB2 Version 10.5 is actually a bunch of great new functionalities which is collectively referred to as BLU Acceleration. The first thing you need to know about BLU is that it can help speed up query processing, especially for analytic processing. Users have experienced queries running up to 74 times faster when using features of DB2 with BLU than before.

BLU Acceleration

So what is BLU? The first aspect of BLU is the addition of a column store capability to DB2. Column organized tables are physically stored as sections of columns rather than as rows of data. By doing so, data warehouse queries, analytics, and other types of ad-hoc queries where aggregates are computed over large numbers of similar data items can be optimized.

« Actionable Compression

But BLU Acceleration is not just a column store. IBM had delivered three additional capabilities and improvements with BLU Acceleration. The first is called "actionable compression," which can deliver up to 10x storage space savings. Some customers have experienced as much as 90-95% data compression for their large data warehouse tables. But why is it called "actionable?" There are two key ingredients that make the compression actionable. There are (1) new algorithms enabling many predicates to be evaluated without having to decompress and (2) the most frequently occurring values are compressed the most, thereby saving the greatest level of storage space.

« SIMD

The second new feature of BLU Acceleration comes via the exploitation of the SIMD (Single Instruction Multiple Data) capabilities of modern CPUs. The basic idea behind SIMD is the ability for a single instruction to be able to act upon multiple items at the same time, which obviously can speed up processing.

« Data Skipping Technology

And finally, BLU Acceleration adds data skipping technology. You can probably guess what this does, but let's explain it a little bit anyway. The basic idea

is to skip over data that is not required in order to deliver an answer set for a query. Metadata is stored for sets of data records that can be accessed by DB2 to determine whether that particular set of data holds anything of interest. If not, it can be skipped over.

Other DB2 Version 10.5 Features

In addition to the significant features delivered with BLU Acceleration, DB2 Version 10.5 provides many other new features. For example, the Workload Manager (WLM) has been improved. WLM can better minimize contention and manage workload for analytics workloads. Additionally, DB2 Version 10.5 adds capabilities to automatically reclaim storage space when data is deleted from extents, as well as offering up a new constraint option to specify uniqueness but not enforce it with an index.

Also, keep in mind that different DB2 editions support different combinations of new functionality. For example, WLM is not offered in DB2 Express-C, adaptive compression requires one of the Server Editions or the Developer edition, and column-organized table support is not offered for DB2 Express-C or the non-Advanced Server Editions. Before planning your upgrade be sure that your targeted edition of DB2 actually delivers the new functionality required by your applications and systems.

Of course, these are just the highlights of the latest versions of DB2, but you can see how there is much to help you improve your applications and administration tasks.

DB2 Editions

A further complicating factor for upgrades is the edition of DB2 that you are using. Currently, there are nine (9) different editions of DB2 11.1 for Linux, UNIX and Windows from which to choose. Each edition is focused on supporting different types of processing and offers different functionality and features. Differentiating which features exist in which edition can be a difficult task to perform. Essentially, though, there are editions for larger businesses, for smaller businesses, for developers, and for personal use. Further complicating things, IBM at times will change the naming convention and number of editions offered.

The editions are as follows:

- DB2 Advanced Enterprise Server Edition
- DB2 Enterprise Server Edition
- DB2 Advanced Workgroup Server Edition
- DB2 Workgroup Server Edition
- DB2 Direct Advanced Edition
- DB2 Direct Standard Edition

- DB2 for Big Data
- DB2 Developer Edition
- DB2 Express-C

Although all DB2 editions share the same code base, certain technical features, resources, and tool sets are limited based on the edition that you choose. As you work your way down the list of editions outlined above, each edition will have a different set of features. These limitations typically do not impact the way your applications operate using the database, but impact issues like performance, administration, and availability. For example, adaptive compression is available in the high end server editions but not in the workgroup edition. Of course, the pricing of each edition varies based on the features and capabilities that are supported.

The most significant new edition added for DB2 11.1 is DB2 for Big Data, which combines the powerful relational engine of DB2 with IBM's BigInsights platform for big data and analytics. Organizations might choose this edition to enhance portability when data needs to be shifted between relational systems and Hadoop.

The other two new editions are the DB2 Direct Advanced Edition and DB2 Direct Standard Edition. These two editions deliver the same functionality as DB2 Advanced Enterprise Server Edition and the DB2 Workgroup Server Edition respectively, but for digital delivery.

When editions change between DB2 versions, your upgrade plans must include an analysis of the editions that are supported. If the edition that you are currently using is no longer supported, you must match your requirements against the editions to be supported in the new version. You can either choose to match your current requirements as closely as possible, or use the upgrade as an opportunity to add functionality by choosing a more feature-laden edition. You can track the features of each edition on-line on IBM's website.

Keep in mind that moving to a new edition of DB2 will impact the cost of editions with more availability and functionality, and features will have a higher price.

Understand the New Features

Perhaps the biggest factor in determining when and how to upgrade to a new DB2 release is the functionality supported by the new release. Tightly coupled to functionality is the inherent complexity involved in supporting and administering new features.

It is more difficult to delay an upgrade if application developers (or even DBAs) are clamoring for new DB2 features. If DB2 functionality can minimize the cost and effort of application development, the DBA group will

feel pressure to migrate swiftly to the new release. An additional factor that will coerce rapid adoption of a new release is when DB2 issues are resolved in the new release (instead of through regular maintenance fixes.)

Regardless of a new release's "bells and whistles," certain administration and implementation details must be addressed before upgrading. The DBA group must ensure that standards are modified to include the new features, educate developers and users as to how new features work and should be used, and prepare the infrastructure to support the new DBMS functionality.


The types of changes required to support the new functionality must be factored into the upgrade strategy. When the DBMS vendor makes changes to internal structures, data page layouts, or address spaces, the risks of upgrading are greater. Additional testing is warranted in these situations to ensure that database utilities, DBA tools, and data extraction and movement tools still work with the revised internal structures.

Complexity of the DB2 Environment

The more complex your database environment is, the more difficult it will be to upgrade to a new DB2 release. Consider the size of your environment. The greater the number of database servers, instances, applications, and users, the greater the complexity. Additional concerns include the type of applications being supported. A DB2 upgrade is easier to implement if only simple, batch-oriented applications are involved. As the complexity and availability requirements of the applications increase, the difficulty of upgrading also increases.

Location of the database servers also affects the release upgrade strategy. Effectively planning and deploying a DB2 upgrade across multiple database servers at various locations supporting different lines of business is difficult. It is likely that an upgrade strategy will involve periods of supporting multiple versions of DB2 at different locations and for different applications. Supporting different versions in production should be avoided if possible, but that is not always possible to achieve.

Finally, the complexity of the applications that access your databases must be considered. The more complex your applications are, the more difficult it will be to ensure their continuing uninterrupted functionality when DB2 is upgraded.



Creating an Upgrade Strategy

In a complex, heterogeneous, distributed database environment, a coherent upgrade strategy is essential. Failure to plan your upgrade can result in improper and inefficient adoption of new features, performance degradation of new and existing applications, and downtime.

Benefits

Upgrading to a new DBMS release offers both rewards and risks. The following are some of the benefits of moving to a new release.

- Developers can benefit from new features and functionality delivered only in the new release. If development requires a new feature, or can simply benefit from a new feature, program development time can be reduced or made more cost-effective.
- For purchased applications, the application vendor may require a specific DB2 version or release for specific versions of its application to enable specific functionality within the application.
- New DB2 releases usually deliver enhanced performance and availability features that can optimize existing applications. Sometimes a new DBMS release is required to scale applications to support additional users or larger amounts of data.
- IBM often provides better support and responds to problems faster for a new release of their software. This is because vendors are loath to allow bad publicity about bugs in a new and heavily promoted version of their products. It is also true because IBM has a vested interest in moving customers to newer versions of DB2 in order to reduce its support costs.
- There may be cost-savings that accrue by upgrading to a new DB2 release. Older versions of DB2 that have reached their “end of support” date can be costly to manage. Support can be purchased for unsupported versions, but usually at a very high cost.
- Production migration to a new version of DB2 will align the test and production database environments, thereby providing a consistent environment for development and implementation. If a new release is running in the test environment for too long, database administration and application development tasks become more difficult because the test databases will operate differently than the production databases.

Risks

However, an effective DBMS upgrade strategy must balance the benefits against the risks of upgrading to arrive at the best timeline for migrating to a new DBMS version or release.

The risks of upgrading to a new DBMS release include the following:

- A DBMS upgrade usually involves some level of disruption to business operations. At a minimum, databases may not be available while the DBMS is being upgraded. This can result in downtime and lost business opportunities if the upgrade occurs during normal business hours (or if there is no planned downtime.) Clustered database implementations may permit some database availability while individual database clusters are migrated to the new version.
- Other disruptions can occur, such as having to convert database structures or discovering that previously supported features were removed from the new release (thereby causing application errors). Delays to application implementation timelines are another possibility.
- The cost of an upgrade can be a significant barrier to release migration. First, the cost of the new version or release must be budgeted for (price increases for a new DBMS version vary but can amount to as much as 10% to 25%.) The upgrade cost must also factor in the costs of planning, installing, testing, and deploying not just the DBMS but also any applications using databases. Finally, be sure to include the cost of any new resources (such as memory, storage, additional CPUs) required to use the new features delivered by the new DBMS version. It is important to meticulously examine the specifications of any new version or release. Sometimes features and functionality supported by past releases can be removed for a new release, thereby necessitating the acquisition of additional software (and expense) to replace the lost capabilities.
- As indicated earlier, a new version of DB2 is likely to come with performance gains. But keep in mind that when SQL optimization techniques change, it is possible that a new version of DB2 will generate SQL access paths that perform worse than before. DBAs must implement a rigorous testing process to ensure that new access paths are helping, not harming, application performance. When performance suffers, application code may need to be changed—a very costly and time consuming endeavor. A rigorous test process should be able to catch most of the access path changes in the test environment.
- New DB2 releases may cause features and syntax to be deprecated that are being used in existing applications. When this occurs the applications must be modified before migrating to the new release can proceed.
- Supporting software products may lack immediate support for a new DB2 release. Supporting software includes the operating system, transaction processors, message queues, purchased application, DBA tools, development tools, and query and reporting software.

After weighing the benefits of upgrading against the risks of a new DB2 release, the DBA group must create an upgrade plan that works for the organization. Sometimes the decision will be to upgrade immediately upon availability, but often there is a lag between the general availability of a new release and its widespread adoption. Your organization style will come into play when deciding on an upgrade timeline.

Organization Style

Every organization displays characteristics that reveal its style when it comes to adopting new products and technologies. Industry analysts at Gartner, Inc. have ranked organizations into three distinct groups labeled types A, B, and C. A type-A enterprise is technology-driven and, as such, is more likely to risk using new and unproven technologies to try to gain a competitive advantage. A type-B organization is less willing to take risks but will adopt new technologies once others have shaken out the bugs. Finally, a type-C enterprise, very conscious of cost and averse to risk, will lag behind the majority when it comes to migrating to new technology.

Only type-A organizations should plan on moving aggressively to new DB2 releases immediately upon availability if the new features of the release will deliver advantages to the company. Type-C enterprises should adopt a very conservative strategy to ensure that the DB2 release is stable and well-tested by types A and B companies first. Type-B organizations will fall somewhere between types A and C: Almost never upgrading immediately, the type-B company will adopt the new release after the earliest users have shaken out the biggest problems, but well before type-C enterprises.

Skipping Releases

When the risks of a new release outweigh the benefits, some organizations may decide to skip an interim release if doing so does not impact a future upgrade. For example, a good number of DB2 users migrated directly to DB2 Version 10.1 from DB2 Version 9.5, skipping DB2 Version 9.7.

When skipping interim releases, be careful to understand and prepare for the interim release, as well as the destination release. That means taking care to evaluate the prerequisites for both releases, as well as understanding the new features being made available for both releases.

Although a multiple release upgrade takes more time, you will be able to more effectively control when and how you migrate to new DB2 releases instead of moving steadily to every new release that is delivered by IBM. For example, if you are still running DB2 Version 9.5, which will no longer be supported as of April 30, 2018, you might choose to move to DB2 Version 10.1 or 10.5 instead of DB2 Version 9.7, which would be the immediate next version after 9.5.

The Danger of Being Out of Support

As new releases are introduced, DBMS vendors retire older releases and no longer support them. The length of time that IBM will support an old release of DB2 must be factored into your DB2 release migration strategy. If you decide to continue running a DB2 release in production that is no longer supported, IBM will not be able to resolve problems that arise because you are no longer paying for support.

Some type-C (and perhaps type-B) organizations try to get by running a DB2 version that is no longer supported by IBM. This might, at first blush, seem to be a less risky strategy than upgrading to a new version, but this actually can be a much more troublesome tactic.

Unless you negotiate a high-cost, limited support contract from IBM, any bugs or errors in DB2 can bring your organization's applications to a standstill. Running unsupported DB2 software is not a wise course of action if your organization is conducting mission-critical services using DB2. If you absolutely must continue using a retired version of DB2 (for business or application issues) be sure to investigate IBM's policies for unsupported releases as well as the related costs.

Furthermore, as software ages you run the risk of new hardware and other updated software not working with the old software. This can force an organization to stop upgrading to any newer technologies. As your hardware and software ages, if it does not fail it may slowdown, thereby impacting your applications and your business. And you also risk losing DBAs and developers not interested in working with obsolete technologies.

With these considerations in mind, standard operating procedure should be to make sure that you are running only supported versions of DB2 in your organization. IBM documents DB2 end of support dates on their website.

DBA Staff Skill Set

Upgrading DB2 is easier if your DBA staff is highly skilled or experienced. The risk of an upgrade increases as the skills of the DBA staff decrease. If the DBAs at your organization are not highly skilled, or have never migrated DB2 to a new release, consider augmenting your DBA staff with consultants or DBA services such as Datavail's, for the upgrade. Deploying skilled DBAs will ensure that your upgrade goes as smoothly as possible.

Platform Support

When IBM unleashes a new release of DB2, it is possible that not every platform and operating system will be immediately supported. Be sure to engage with IBM to understand when your particular platform and operating system will be supported by the new version of DB2.

Supporting Software

Carefully consider the impact of a DB2 upgrade on any supporting software. Supporting software includes purchased applications, DBA tools, reporting and analysis tools, and query tools. Each software vendor will have a different timeframe for supporting and exploiting a new DB2 release.

Note that supporting and exploiting a new release are two different things. Some vendors carefully differentiate between being able to support versus exploit a new DB2 version or release. Software that supports a new release will continue to function the same as before DB2 was upgraded, but with no new capabilities. Therefore, if a DBA tool, for example, supports DB2 Version 10.5 it can provide all the services it did for the last release, as long as none of the new features of the new version are used. In contrast, a DBA tool that exploits a new version or release provides the requisite functionality to operate on the new features of the new DB2 release.

So, to use a concrete example, IBM added BLU Acceleration in DB2 Version 10.5, which includes the ability to create column-organized tables. A DBA tool can support DB2 Version 10.5 without operating on column-organized tables, but it must operate on them to exploit DB2 Version 10.5.

Prior to migrating to a new version of DB2, make sure you understand the difference between supporting and exploiting a new version, and get a schedule for both from your third-party tool vendors for the tools you use.

Migration Verification


The DBA should implement procedures—similar to those for a new installation—to verify that the DBMS release upgrade is satisfactory. Perform the same steps as with a brand-new DBMS install, but also test a representative sampling of your in-house applications to verify that the DBMS upgrade is working correctly and performing satisfactorily.

Fallback Planning

IBM will provide an installation guide for each new version of DB2 that outlines the new features of the release and describes the fallback procedures to return to a prior release. Be sure to review the fallback procedures and implement the procedures faithfully if being able to go back to a prior DB2 release is important. You may need to return to a previous release if the upgrade contains a bug, performance problems ensue, or other problems arise during or immediately after migration.

Summary

Keeping your DB2 environment up-to-date and viable requires planning and research to ensure that your organization's functionality and support requirements are being met. Be sure to allocate sufficient time and resources to keep your DB2 systems upgraded to an appropriate level for your organization's needs. If your DB2 team needs assistance or the benefit of having help to upgrade, consider Datavail's DB2 experts. They can help you get the job done.



Biography

Craig Mullins

Consultant at Mullins Consulting, Inc.



Craig S. Mullins is a data management strategist, researcher, and consultant working with Datavail and its DB2 practice to expand offerings and client base.

He is President and Principal Consultant at Mullins Consulting, Inc. and the publisher of The Database Site (thedatabasesite.com). He has three decades of experience in all facets of database management and has worked with DB2 on the mainframe since V1.

About Datavail

Datavail Corporation is the largest provider of data and database administration (DBA) services in North America, offering database design and architecture, administration and 24x7 support. The company specializes in Oracle, Oracle E-Business Suite, Microsoft SQL Server, MySQL, MongoDB, DB2 and SharePoint, and provides flexible on-site/off-site, onshore/offshore service delivery options to meet each customer's unique business needs.

Contact Us

General Inquiries:

877-722-8247

Fax Number: 303-469-2399

Email: info@datavail.com

Corporate Headquarters:

Datavail Corporation

11800 Ridge Parkway

Suite 125

Broomfield, CO 80021

Powai Office:

Datavail Infotech Pvt. Ltd

A-902, Supreme Business Park,

Hiranandani Gardens, Powai

Mumbai – 400076

Maharashtra

Bangalore Office:

Datavail Infotech Pvt. Ltd

Concept Business Park

#319/9, 1st floor, Block A

Hosur Main Road

Bommanahalli, Bangalore 560100

datAvail

DATABASE SERVICES

www.datavail.com | 877.634.9222