DB2 11 for z/OS: Unmatched Efficiency for Big Data and Analytics

by Julian Stuhler

Few IT professionals can have missed the big data phenomenon that has manifested itself in recent years. Industry publications and IT analysts have devoted a huge percentage of their output to the subject (creating a big data challenge all their own in the process). There can be little doubt that the advent of new technologies and methods of customer and business interaction have created unique challenges for organizations wishing to create actionable insight from very large amounts of unstructured data. Innovative tools and techniques have been developed to cope with these "big data" challenges (and indeed some of them are discussed in this paper, in the "Hadoop and Big Data Support" section).

However, beyond this somewhat narrow definition of big data, many organizations have been dealing with the challenges of processing, maintaining, and analyzing ever-increasing amounts of more traditionally structured data for many years. The inherent scalability and resilience of IBM® DB2® for z/OS® and the underlying System z® platform have proven to be a compelling combination for such applications, and IBM continues to invest in extending DB2's capabilities with each new release.

From transparent archiving to greater in-memory scalability through the use of 2 GB page frames, DB2 11 for z/OS, the latest release of IBM's flagship database, contains many new features specifically designed to help customers to address the challenges of managing traditional big data. A wealth of material exists on the technical changes within DB2 11, but finding descriptions of how those new features will improve your business results can be a challenge. The main body of this paper provides a high-level overview of the major new features from an IT executive's perspective, with emphasis on the underlying business value that DB2 11 can deliver.

This is the fourth paper in this series, with previous editions highlighting the business value offered by DB2 for z/OS V8.1, DB2 9 for z/OS, and DB2 10 for z/OS:

- *DB2 for z/OS 8.1: Driving Business Value* (J. Stuhler, Triton Consulting, 2004)
- *DB2 9 for z/OS: Data on Demand* (J. Stuhler, Triton Consulting, 2007)
- *DB2 10 for z/OS: A Smarter Database for a Smarter Planet* (J. Stuhler, Triton Consulting, 2010)

NOTE: Throughout the remainder of this document, all references to "DB2 9," "DB2 10," and "DB2 11" refer to the relevant release of IBM DB2 for z/OS.

DB2 11 for z/OS: The Database for Big Data and Analytics

In this section, we take a detailed look at the major features of DB2 11 for z/OS and see how many of IBM's most innovative enterprise customers plan to use them to deliver an enhanced IT service to the business. Many of these enhancements can deliver benefits "out of the box," with little or no effort required to begin exploiting them, reducing the time-to-value for a DB2 11 upgrade. See "DB2 11 New Features by Implementation Effort" (opposite) for a breakdown of the effort required to exploit each new feature.

This section is organized around the key DB2 11 themes:

- **Efficiency.** Reducing cost and improving productivity
- **Resilience.** Improving availability and data security
- **Business analytics.** Enhanced query and reporting

Efficiency

Even in the most favorable economic climate, businesses need to control costs and increase efficiency to improve their bottom line. In today's increasingly challenging business environment, this continues to be a key factor for the survival and success of enterprises of all sizes.

This section examines the major DB2 11 enhancements that are aimed at delivering the highest efficiency for core IT systems that rely on DB2, a key design objective for the new release. These features can help reduce ongoing operational costs, improve developer and DBA productivity, and enhance customer experience by increasing performance and delivering a more responsive application.

CPU Reductions

Most DB2 for z/OS customers operate on a CPU usage-based charging model, so any increases or decreases in the amount of CPU required to run DB2 applications can have a direct and very significant impact on overall operational costs.

Traditionally, IBM has tried to limit the additional CPU cost of adding new functionality into each release, keeping the net CPU impact below 5 percent. The move to a 64-bit computing platform in DB2 for z/OS Version 8 was an exception to this rule and introduced some significant processing overheads that resulted in many customers experiencing net CPU increases of 5 to 10 percent following the upgrade.

DB2 9 for z/OS helped to redress the balance somewhat by delivering modest CPU improvements for many large customers, but the advent of DB2 10 completely changed the picture. IBM delivered the most aggressive performance improvements of any DB2 release in the past 20 years, with many customers seeing net CPU savings of 5 to 10 percent or more in their traditional DB2 online transaction processing (OLTP) workload without any application changes being required.[1] Unsurprisingly, these savings proved to be very popular and are consistently quoted as being one of the major reasons for customers to upgrade to DB2 10.

DB2 11 New Features by Implementation Effort

One of the most compelling features of DB2 11 is the number of enhancements that can deliver business benefit with little or no change being required to existing applications. The lists below categorize the covered DB2 11 features in this paper according to the amount of effort required to exploit them:

Minor Implementation Effort – Immediate. These features are available immediately after upgrading to DB2 11, with no database or application changes required. A REBIND may be required.

Minor Implementation Effort – Deferred. These features do not require any database or application changes but will be available only after the DB2 system has been placed in New Function Mode.

Significant Database/System Changes Required. These features require some changes to be made to DB2 objects and structures (typically by the DBA), but no application changes. These changes are typically quicker and less expensive to implement/test than application changes.

Significant Application Changes Required. These enhancements require some degree of application change in order to implement and will therefore be the most expensive to implement and test.

Minor Implementation Effort – Immediate

CPU reductions
Application compatibility
pureXML enhancements
Optimizer and query performance improvements
Data sharing performance enhancements
Enhanced dynamic schema change (some features)
BIND/REBIND enhancements

Minor Implementation Effort – Deferred

zEC12 exploitation (also requires DB2 to be running on a zEC12-class server)
Temporal data enhancements
Utility enhancements
Enhanced dynamic schema change (some features)

Significant Database/System Changes Required

Java stored procedure enhancements
Extended log record addressing
Security enhancements

Significant Application Changes Required

Transparent archiving
Global variables
Variable arrays
SQL aggregation improvements
Hadoop and big data support

IBM has further developed the CPU reduction theme within DB2 11, with initial savings of up to 5 percent expected for customers running simple OLTP workloads. Significantly higher savings are possible for complex OLTP and query workloads, as discussed below. Because these improvements are due to internal DB2 code optimization, they are available in DB2 11 Conversion Mode, without the need for any application changes. Additional CPU savings are possible once customers begin to use some of the other DB2 11 enhancements that require application change, as described elsewhere in this section.

Some workloads will benefit more than others from the performance enhancements offered by DB2 11. Figure 1 breaks down the anticipated CPU savings by workload type.

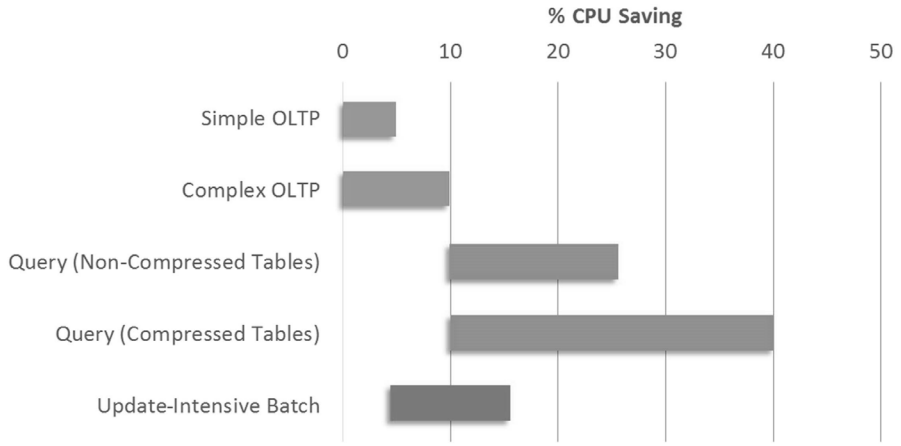

Figure 1: DB2 11 CM vs. DB2 10 NFM – Expected CPU savings by workload type

The most significant benefits are expected to be seen within query workloads. Complex reporting queries can see up to 25 percent savings for uncompressed tables and up to 40 percent for queries on compressed tables. Reporting queries with heavy sort processing may also see additional DB2 CPU savings.

Traditional OLTP workloads are also likely to benefit from the efficiency enhancements in DB2 11. Savings of up to 5 percent are expected for simple OLTP,[2] with reductions of up to 10 percent for more complex transactions. Finally, update-intensive batch workloads may enjoy CPU reductions of 5 to 15 percent.

Figures 2 and 3 depict some actual observed CPU reductions for sample workloads, run as part of IBM's internal performance testing for the new release. These figures are broadly in line with the high-level expectations detailed above.

The overall out-of-the-box CPU savings within DB2 11 are expected to be one of the major factors supporting the business case for upgrading to the new release.

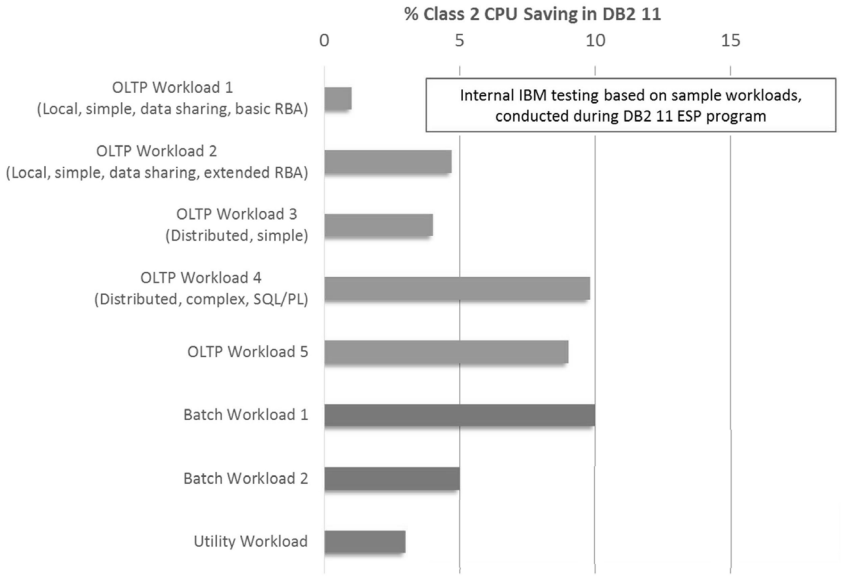

Figure 2: DB2 11 CM vs. DB2 10 NFM – Sample OLTP/batch CPU savings

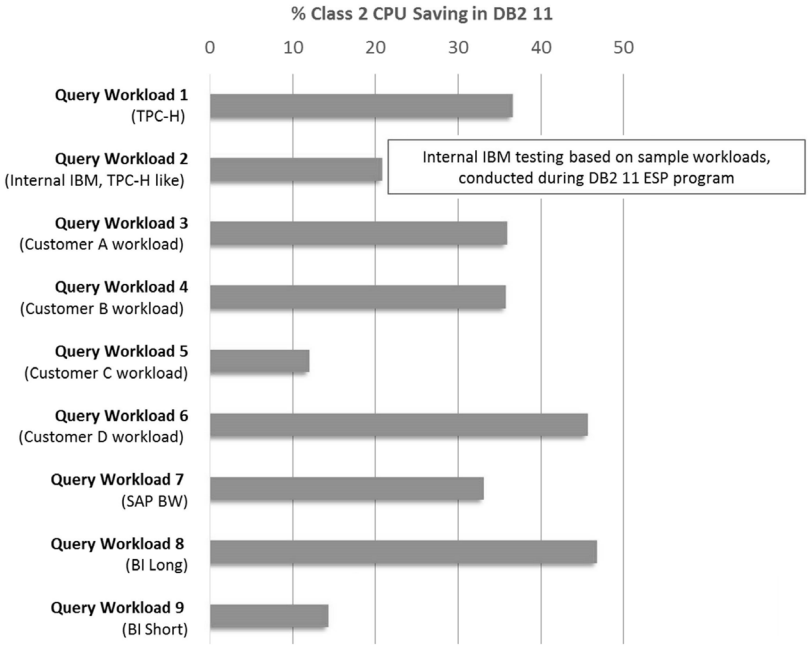

Figure 3: DB2 11 CM vs. DB2 10 NFM – Sample query CPU savings

zEC12 Exploitation

In August 2012, IBM announced the latest-generation IBM zEnterprise® EC12 (zEC12) enterprise servers, with up to 101 configurable processors per server, each running at an industry-leading 5.5 GHz. In addition to an impressive list of general performance and capacity improvements over the previous-generation z196 enterprise servers, the zEC12 models include a number of features that DB2 11 will specifically exploit.

2 GB page frames. DB2 10 for z/OS introduced support for 1 MB "large page frames," an enhancement designed to reduce processing overheads for very large DB2 buffer pools by letting z/OS manage the underlying storage in fewer 1 MB pieces rather than many more 4 KB pieces (Figure 4).

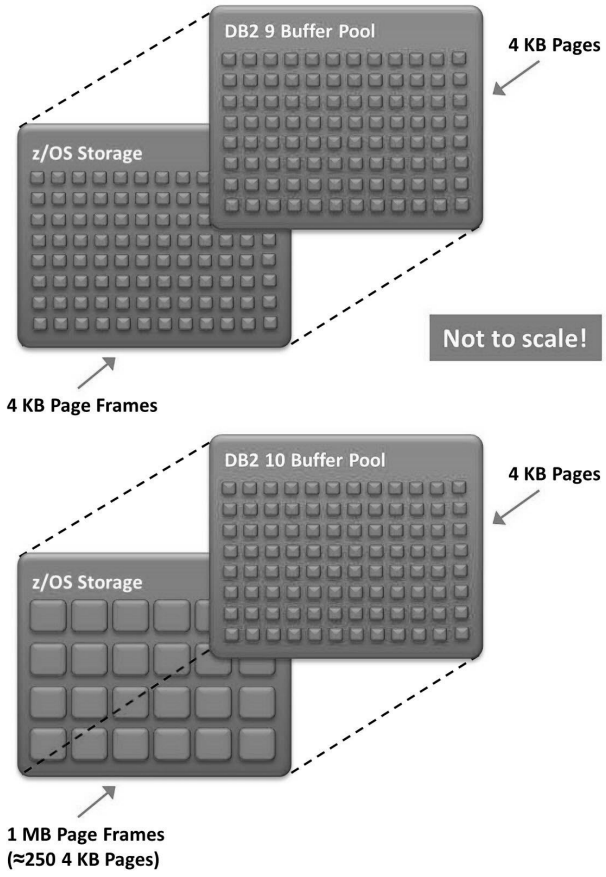

Figure 4: DB2 10 large page frame enhancement

Many customers with larger DB2 buffer pools were able to achieve CPU savings of up to 4 percent by exploiting this capability. However, as memory prices fall and workloads increase, DB2 buffer pools continue to increase in size, and the overheads of managing even the larger 1 MB page frames start to become significant.

In recognition of these trends, when running on an zEC12 server DB2 11 will support even larger 2 GB page frames, each of which will map onto more than half a million 4 KB pages (Figure 5).

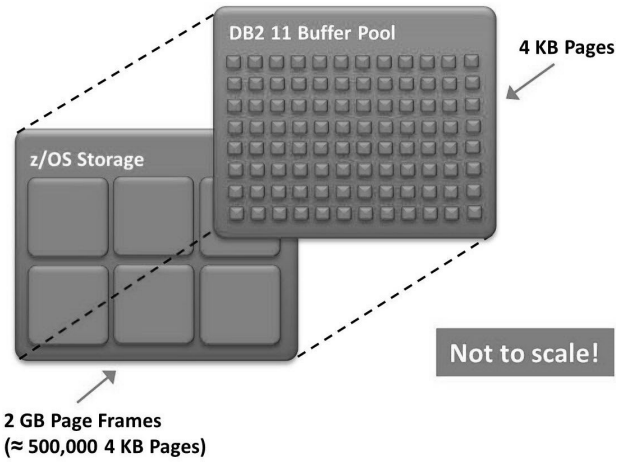

Figure 5: DB2 11 large page frame enhancement

Those customers using very large DB2 buffer pools will see further CPU reductions by moving to 2 GB page frames. Other sites may not have sufficiently large pools for 1 MB page frames to be a significant limitation today, but that situation will undoubtedly change in the future as buffer pool sizes continue to grow. By moving early to support 2 GB page frames, IBM has recognized and eliminated an important future scalability issue.

DB2 code using large page frames. As discussed in the previous section, DB2 10 and DB2 11 have exploited 1 MB and 2 GB large page frames to allow more efficient handling of large buffer pools. However, despite the extensive use of large memory objects in the past few releases of DB2, the storage used for DB2 code (as opposed to the data held in buffer pools) remained backed by standard 4 KB page frames.

DB2 11 is able to utilize large page frames for DB2 code objects and log output buffers, in addition to buffer pools. This enhancement reduces the z/OS overheads associated with DB2 code objects, lowering CPU consumption and operational costs. (Support for running DB2 code in large page frames requires z/OS 2.1.)

Application Compatibility

Many new releases of DB2 introduce enhancements or new features that require application and/or SQL code to be changed. These include additional SQL reserved words, changes to DB2 behavior or processing and even changes to SQL return codes. Although IBM tries to minimize these "incompatible changes," they cannot always be avoided. They may be required in order to ensure that DB2 adheres to evolving SQL standards, to support new functionality, or perhaps to address an earlier defect in the DB2 code.

A major part of planning for a new release is to analyze the impact of these incompatible changes and arrange for the necessary amendments to be made to DB2 application code so it will continue to work as designed under the new release. This situation poses some challenges for DB2 customers:

- Analysis of the impact of incompatible changes can be difficult, time consuming, and error-prone. Missing one or more of the required changes may result in application outages when DB2 is upgraded (or worse, the application may continue to work but return unexpected results).
- Finding the necessary resources to undertake any required remedial work (and scheduling the associated change slots) can be expensive and require significant elapsed time. All of the changes within a given subsystem or data sharing group must be completed before the upgrade can commence, so a lack of resources within a single application team could impact the upgrade schedule for the entire environment.

Figure 6 depicts these challenges.

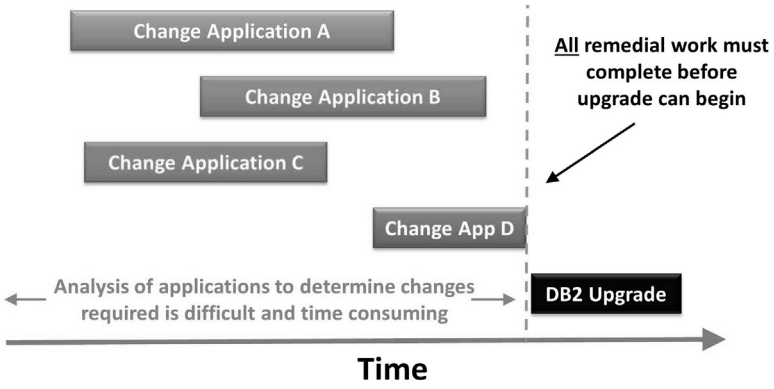

Figure 6: Application compatibility issues

To address these issues and allow customers to upgrade their DB2 systems with less effort and risk, IBM has introduced some new capabilities in DB2 11 for z/OS that remove the hard dependency on all remedial work being conducted before a version upgrade and allow the impact of incompatible changes to be more easily assessed. Figure 7 summarizes these enhancements.