THE CLIPPER GROUP
Navigator™
Navigating Information Technology Horizons℠
Published Since 1993 Report #TCG2006006 January 24, 2006

System z9 Adds *zIIP to* Ally with *DB2 on z/OS* to Better Serve the Onslaught of Business Data

Analyst: Anne MacFarland

Management Summary

The perspective of business is constantly changing. What was needed for survival last year is now, often, insufficient. The need for a single view of the customer, the partner, and other significant entities that impinge on a business' strategy and tactics has become critical. This single view requires that the onslaught of *data pouring in* through email, POS devices, and instruments and the *data boiling up* from research, business units, customers, and joint initiatives with partners be handled effectively. Predictive analysis is used to determine tactical positioning - demanding more analysis of more information in less time. As old business fiefdoms have been called to participate more broadly in the enterprise and the results of their analyses must be shared more broadly, because understanding the changes in corporate strategy is required throughout an organization, if it is to respond coherently to opportunity. If it is only partially responsive, misaligned processes will leave unrealized revenue to the competition and the enterprise will be a shadow of what it could have been.

At an operational level, too, the world impinges on old procedures. In design and development, more iterations and more input make better decisions, but this intensity and inclusiveness can cause data and data management issues. The new cross-enterprise initiatives reveal opportunity, but they also cause risk, for without the self-centeredness of the old fiefdoms, internal data protection must be stronger. Identity awareness and as role-based security are needed – and they must be imposed in a way that does not slow to a crawl the sharing of data.

These trends find their most imperative incarnation in the very large enterprise – whose value nets are global and whose needs for agility spur a need to know that frequently transcends organizational structure. **Enterprise data must be secure, consistent and correct - yet contextually rich. Access to data must be both timely and secure. Management of the data must be effective, yet not transcend the skim-coat budget slice allotted to administration.**

Most of these very large enterprises have built their capabilities on relational databases, which can aggregate very large amounts of information to satisfy the quick-time heartbeat of business decision-making. IBM just beefed up *DB2* on *z/OS* and added an *Information Integration Processor* (called *zIIP*) to *System z9* hardware. Together with the richness of z/OS, these support an innovative approach to the use of business data. For more of the details, read on.

IN THIS ISSUE

The Problems of Timely Data Service in a Large Scale Domain

While the mushrooming of unstructured data concerns storage-centric folks and the nature of what was said in that content concerns the lawyers, it is the traditional structures of journals, account books, and electronic spreadsheets and databases that have been the backbone of business processes for centuries. The multi-dimensional data displays allowed by relational databases (and, now, object databases with search) have supported the growth of data warehouses. These data structures amass and integrate a broad array of business data and present it to applications for the analysis needed to derive (in theory) the comprehensive knowledge needed to choose effective strategies to address markets.

The problem has always been one of time. Structured query language (SQL) has given us a common (well, mostly common – more on that later) way to find information, but there have been limits on how broadly it can be applied. Information systems are logical but literal, and awareness of divergent data types and incompatible schema cannot be achieved without some kind of transformation. Without such transformation, queries are either stymied to a stand-still by such information, or unaware of it. The latter is preferable for IT managers, but can seriously impair the quality of the decisions, if the overlooked entries contain key information.

The need to transform[1] information process to make it broadly useful presents the business with some draconian alternatives.

1. Standardize on key metrics and business indicators, and only use the information that fits. This makes for bad decisions.
2. Transform everything. This makes for decisions based on *what was*, not *what is*.

There is, however, a third alternative.

3. **Make the transformation more atomic – more precipitated - so that information can be left in place (less costly and more preserving of context) and translated as needed.**

IBM's Strategy

Some years ago, IBM decided that this third approach – a federated approach – was needed to complement the traditions of a data warehouse. Like the aerobic exercises that complement weight training, **IBM's federated diversification strategy would meet more fully the business need to use information for both long term strength and the more immediate quick-twitch responsiveness to threat and opportunity.**

IBM made a series of acquisitions to broaden the structures of information they could manage, including unstructured information, forms, product information, and master data. IBM enriched their content management product with new forms of search and discovery (*Web Fountain*). IBM open-sourced an acquired Informix *Cloudscape* DBMS as Derby, providing a new low point of entry to their data management environment, and transformed *DB2* into *UDB*, a platform-independent DBMS that could run on everything from mainframes to embedded devices.

IBM built out its information integration capabilities, both in DB2 and in *WebSphere*. The integration of the code base of the various IBM Software business units, including DB2, has allowed data integration capabilities to be positioned opportunistically where needed. IBM also worked significantly on a number of XML standards bodies and on the concept of gateways for quick access to heterogeneous sources of data. DB2 evolved to use the common SQL queries needed to access data from other relational databases, and all the middleware to do so safely and securely.

DB2 for z/OS Version 8 was released in 2004 and was the start of a significant DB2 evolution. Advances in on-line schema evolution continue to give improvements to already industry-leading continuous availability, while simplifying administrative processes. Scalability continues to improve through support of two-dimensional clustering with more effective I/O as well as more efficient use of memory. Major progress allowed better support of Java, XK\ML, and WebSphere workloads. Data warehousing support is optimized by the ability to support MQTs (Materialized Query Tables). The move toward a more standards-compliant SQL provided more congruence with other database management systems, broadened the addressable domains of data, and enabled the transformations of business data. **Nonetheless, there remained two significant challenges:**

- **The need for real-time access to insights, and**

[1] People transform information very well but they do not facilitate its broad use.

- **The need for comprehensive rationalization and coordination of business use of information as an integrated whole.**

To accomplish these functions, IBM has now endowed its mainframe *System z9-109* environment with a special *zIIP*[2] data-serving engine (i.e., a special-purpose segregated processing complex), and a new version of the DB2 - *UDB for z/OS* – to be available later this year, which has been recomposed to be more performant and more fully capable of addressing and managing all the data sources that infest the modern enterprise. The zIIP alone (which works with the current version of DB2 on z/OS) will expand how you do data services, but the new capabilities coming in the next version of DB2 for z/OS will let you think differently about how you structure your enterprise's data domain.

The New Version of DB2 on z/OS

DB2 version 8 on z/OS started to recentralize inefficient distributed processes back into z/OS and added the multi-level security there. The coming version of DB2 on z/OS continues the recentralization of process control to z/OS. It retains some distributed processes as legacies, but moves more and more to a *hub and spoke* or *star* incarnation for greater efficiency. Having a hub and spoke architecture allows a single view of data wherever it sits, and a single point of coordination of data use. The hub (DB2 on z/OS with the assist of the zIIP processor) is doing all of the load balancing and locking of this extended environment in a centralized way.

Security

Timely access to a vast domain of data requires a scalably-architected, atomically-implemented security strategy. The number of potential internal threats grows as the domain of data spreads to encompass competing business units. z/OS encryption options and multi-level security[3] allow this larger domain to be shared safely by many different user groups.

In the next version of DB2, IBM will add stronger role-based security for database administrators and other authorized users[4]. There will no longer be a default of trust in internal networks and topologies. Instead, DB2 will support a granular security for specific physical and logical connections. IBM continues to build security into the environment from the outset by promulgating its z/OS *System Integrity Programming Standard*. z/OS system integrity has been part of the DB2 architecture since its first release. These enhancements, taken as a whole, basically beef up the security beyond the traditional system level to address an entire distributed environment from development and throughout its operational life.

Data Aggregation

In IT systems, data aggregation and sharing can sap performance. For a long time, DB2 on z/OS has automatic data compression that improves performance. The mainframe's heritage of shell environments (*VM* is decades old) allows a near-linear scalability of DB2 cluster nodes when run in that environment, while full security is maintained. Now, the recentralization of database processes to be managed by z/OS increases their performance, and the coming version of DB2 for z/OS includes optimistic lock support and time stamping in the DB2 Engine.

XML Developments

With its *Viper* announcement in 2005, DB2 provided native support for XML – the ability to have a hybrid XML–relational database. Previously, DB2, like its competitors, had shredded XML into the rows and columns of the relational structure – something that was time consuming and lost the coherence of the object as a whole asset. With Viper, DB2 Stores XML as a *parsable tree*. The hybrid database engine can parse SQL, XQuery, and Xpath. It also features a bilingual query compiler, XML path and values indexes, and filtering.

The next version of DB2 for z/OS will change the way the mainframe does XML processing for all three alternatives. These integrated capabilities, and the DB2 database in the *WebSphere Application Server*, will allow datasets managed by DB2 on z/OS to be managed as targets of WebSphere SOA applications.

SQL Commonality

DB2's ANSI-compliant SQL (available in the

[2] *zIIP* stands for *z9 Integrated Information Processor*.

[3] z/OS enables row-level security in DB2. This and the new security on authorized users (see footnote 5 below) makes auditing and troubleshooting more effective and regulatory compliance a simpler process.

[4] The permissions for authorized users used to be all-or nothing: Now they can specify location and time of day/week for all databases DB2 manages, with better trace and track capabilities.

next version for z/OS) will be 95%[5] or more congruent with the SQL used by other major databases, such as those from Oracle and Microsoft. This makes migration and porting of data and processes between different databases easier. Furthermore, natively-stored SQL procedures will be supported by DB2 on z/OS, making it more congruent with DB2 on other platforms.

Database Management Efficiency

Other new capabilities coming in DB2 for z/OS include an enhanced ability to clone tables[6], partition by growth[7], table append, default databases and table spaces, and automatic unique indexes. All these enhancements decrease the management complexity (and costs) of large-enterprise structured data environments. There will be a GUI-based Optimization Service Center that captures workload information and suggests optimization strategies – all part of the autonomics that curtail the routines of administration

The zIIP Processor for z9

zIIP does for database operations what *zAAP* did for *Java*.[8] As with zAAP, DB2 on z/OS is deployed in the same logical partition as z/OS (remember, other partitions may be running under Linux). z/OS then coordinates with DB2 to redirect requests for data serving cycles to the zIIP data processor, which speeds the time to completion. There is a limitation on the number of zIIPs that can be deployed in a system – there must be a general-purpose processor for every zIIP, just as there has to be a general-purpose processor for every zAAP. But that same general processor can be counted for both a zIIP and a zAAP, if your workloads can benefit from both Java and data serving offloads. For many enterprises dependent on their databases and J2EE applications, this will be the case. As with zAAP, there are no z/OS or other operating software charges for DB2 work done on a zIIP. zIIP will work on System z9 109 running z/OS 1.6 and later with DB2 for z/OS v8.

When configured with zIIP data processor(s), zAAP Java processor(s), and z/OS to fully the administer the offload to both specialty engines, the mainframe steps firmly into its role as *enterprise hub* that can harness the power of the information that pervades the enterprise. This hub can coordinate the work distributed throughout the enterprise infrastructure and can become the point of reconciliation of the activities of disparate, heterogeneous systems.

The highly-evolved, stabile, and comprehensive capabilities of the mainframe environment - with its long heritage of resource sharing – and the z/OS operating system - capable of managing the arbitrations needed to work deftly and quickly with heterogeneous data sources - made z9 the obviously most-appropriate host for these new data-sharing capabilities. **Now, IBM builds out the notions of scalability, availability, and security in the specific context of data sharing to make the mainframe not just a wonder of technology but also specifically capable of addressing the timely information analysis challenges of today's enterprise**.

Conclusion

The equivalent of a space-time continuum that business information requires is basically the three dimensions of *data capture*, *data access*, and *data management*, together with an accelerant to make navigating those dimensions happen in a timely way. The combination of the next version of DB2 on z/OS and the zIIP processor gives a rich package of tools to meet this challenge.

If you don't want to navigate the seas of opportunity armed only with an incoherent idea of whom you used to be, you will want to spruce up how you leverage your data, which may mean rethinking how you serve it. IBM's new version of DB2 on z/OS will not be out until later in 2006, but zIIP will work with z9 and DB2 for z/OS V.8, and it is not too soon to start thinking about how you can reorganize and rationalize data serving for your enterprise.

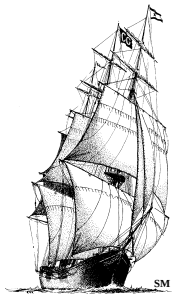

---

[5] The incongruent SQL is for legacy processes seldom used in these larger, heterogeneous domains.

[6] This is helpful to Web-based applications where quick switching between dual-mirrored tables is a key element of performance.

[7] Instead of setting key ranges, administrators can set policy thresholds.

[8] Fort more on zAAP, see **The Clipper Group Navigator** dated April 7, 2004, entitled *zSeries Zips through Java with zAAP*, which is available at http://www.clipper.com/research/TCG2004030.pdf.

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

Anne MacFarland **is Director of Data Strategies and Information Solutions for The Clipper Group.** Ms. MacFarland specializes in strategic business solutions offered by enterprise systems, software, and storage vendors, in trends in enterprise systems and networks, and in explaining these trends and the underlying technologies in simple business terms. She joined The Clipper Group after a long career in library systems, business archives, consulting, research, and freelance writing. Ms. MacFarland earned a Bachelor of Arts degree from Cornell University, where she was a College Scholar, and a Masters of Library Science from Southern Connecticut State University.

- ***Reach Anne MacFarland via e-mail at Anne.MacFarland@clipper.com or at 781-235-0085 Ext. 128. (Please dial "128" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, **The Clipper Group Explorer**, **The Clipper Group Observer**, **The Clipper Group** *Captain's Log*, **The Clipper Group Voyager**, and *"clipper.com"* are trademarks of The Clipper Group, Inc., and the clipper ship drawings, "*Navigating Information Technology Horizons*", and "*teraproductivity*" are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.