

Independent Assessment

by Alan Radding

IBM System z case study: Univar

A fast, low cost path to business intelligence

Executive Summary

Univar USA, a largely mainframe shop, long wanted a modern enterprise business intelligence (BI) platform. Instead, it made do with a decade-old data warehouse running on Oracle and IBM Series P-series servers.

Budget constraints, a deepening economic turndown, and other priorities precluded undertaking the kind of top-to-bottom overhaul of the organization's BI strategy. Those needing information would have to be satisfied with whatever they could coax out of the Oracle data warehouse.

Following a data seminar the company received an invitation from IBM to participate in the InfoSphere Warehouse for System Z beta program. It jumped at the chance to take advantage of InfoSphere's cubing capabilities that allowed it to transform its System z DB2 data on the fly. What if BI users could access interesting slices of DB2 data without having to go through the Oracle data warehouse?

When the InfoSphere beta program ended Univar applied some of the modeling techniques they had explored with Data Architect, combined with Crystal reports and Business Objects, both of which were already deployed against Oracle, to craft an effective tactical BI strategy using the System z and DB2 data. The functionality isn't as rich as could be obtained through InfoSphere Warehouse for System Z and its cubing services, but it is an affordable tactical BI solution they could have fast.

This case study shows how an organization can leverage the System z and DB2 to deliver meaningful BI quickly and inexpensively. It is a lesson for every organization wanting to do more despite constrained resources.

Challenge—BI when resources are constrained

An application modernization effort started in 2003 led Univar to migrate from VSAM to DB2 on the System z. From a budget standpoint, it just made sense to invest in the System z than to replace it. The organization also had a decade-old Oracle-based data warehouse long overdue for modernization.

At that same time the architecture team led by Kevin Campbell, the company's manager of application and data architecture, really wanted to modernize and re-host the data warehouse. Other more compelling business cases trumped what was perceived as IT plumbing.

Management instead was focused on growing the company and by 2007 had completed the acquisition of their 4th largest US competitor. Hot on the heels of this significant acquisition was implementation of the company's CICS/DB2 ERP system for a sister company in Canada. Now Univar needed to provide some form of analytical capabilities to the new Canadian implementation, but significant data model changes ruled out a quick and easy extension of the existing information warehouse. With a long-term BI roadmap project yet to start there was a need for a more tactical approach.

The company ran an aging but proven ERP system. Home-grown upgrades extended the life of the ERP system but complaints persisted about the difficulty of getting good information out. With an enterprise solution some way in the future Campbell's team was left to piece something together on a shoestring using resources it already had or could get easily.

Background—Univar

According to the company, Univar began in Seattle in 1924 and for much of its early life did business as Van Waters & Rogers, including expansion into Canada in 1950. In 1986 the US organization purchased McKesson Chemical and subsequently acquired Ellis & Everard in 2001. The Ellis acquisition solidified Univar USA's leadership position in the United States. With the 2007 acquisition

of CHEMCENTRAL Corporation, Univar USA expanded yet again. Today Univar Corporation operates a network of over 170 distribution centers throughout North America, Europe and China.

Univar, a global leader in the industrial chemical business, provides chemical wholesaling by purchasing bulk chemicals in truck, rail car, and tank car volumes and breaking them down for customers who purchase in smaller quantities. In addition, Univar provides specialized services such as chemical blending and packaging, labeling, and just-in-time delivery. It also lends its expertise to help customers comply with increasingly complex environmental and regulatory requirements.

Finally, Univar's broad knowledge of chemicals and hazardous materials handling combined with its transportation and warehousing resources, allows the company to provide another valuable service: ChemCare waste management. With the aid of selected partners in the waste disposal business, these waste materials are safely transported to licensed treatment, storage, and disposal facilities.

BI on the System z

BI is a natural on the System z because that's where most of the important production and transaction data live and where most of the processing happens. At least that was the case with the forerunners of BI.

To briefly recap BI mainframe history: in the days before relational databases, PCs, spreadsheets, GUI query tools, and multi-dimensional analytics the mainframe probably could be considered the original BI platform. It regularly generated business reports and managers who wanted something different could request custom reports. This typically entailed filling out a detailed form describing the information they wanted and how they wanted it. It would be passed to a programmer who would code the report. In a week or a month it would run as a mainframe job. Eventually a pile of paper would end up on the manager's desk.

This wasn't really BI as we think of it today. It certainly wasn't fast, and it wasn't iterative; there was no practical way to explore the data through a sequence of different queries. There wasn't even much managers could do with the data they received. In the pre-desktop computing days, the pre-VisiCalc days, managers couldn't even re-enter the data into spreadsheets. Maybe they could write FORTRAN programs if they were so inclined. Almost none were.

The advent of desktop computing, the PC, VisiCalc and all that followed ushered in what would become the modern BI era. In doing so, however, it would move BI away from the mainframe. Transactions still were processed and stored on the mainframe, production databases resided on the mainframe, even basic reporting remained on the mainframe, but custom reporting, data analysis, and what would eventually emerge as BI moved to

distributed platforms, mainly the PC. Armed with spreadsheets like VisiCalc, managers could do their own data analysis while business analysts created multidimensional analytical models.

With the rise of relational databases like DB2, BI should have remained on the mainframe, close to where the data was processed and managed. It wasn't until the advent of client/server computing and, more importantly, the emergence of SOA, however, did the mainframe begin to regain its role in BI. Finally, with the introduction of Rational's data modeling tools on the System z, IBM's acquisition of Cognos, and the introduction of new tools like InfoSphere Data Warehouse the System z could become a modern, competitive BI platform.

Univar pieced together its BI strategy on the System z using InfoSphere initially but also Crystal Reports, Business Objects, DB2, and Rational tools. It

Univar System z BI Components

System z9
zIIP
zAAP
2 IFLs
DB2 v9
InfoSphere Data Warehouse for System Z (beta)
IBM InfoSphere Data Studio
InfoSphere Data Architect

Crystal Reports
Business Objects

took almost three decades, but the mainframe now is back as an effective and competitive BI platform.

Benefits of BI on the System z

The benefits of BI are well documented: more insightful decision making, a more responsive organization, faster problem identification and resolution, avoidance of surprises, and more. These benefits are available from BI on any platform, so why use the System z.

Three reasons come immediately to mind:

1. Low TCO, which enables higher, faster ROI
2. The System z's traditional reliability, security, and manageability
3. It's where most of the data and application logic already reside

To these reasons, Univar can add several of its own. Its first two reasons mirror the benefits noted above: quick time to market and low incremental cost. The company already had DB2 and a System z, which was running its core business applications, so yes, the incremental cost of adding BI was low. Similarly, with the data and the z right there and nearly ready to roll, time to market was short. Univar could assemble the pieces and generate BI results quickly.

More specifically, as Campbell explains: the team already maintained a moment-in-time, read-only, operational data store that represented a snapshot of production data. All they had to do was refresh the data store using DB2 unload and load utilities. It turned out to be a trivial matter to include these databases in existing backup schedules, similarly for runstats, reorgs and other housekeeping—all of which typically are much quicker on the z than on distributed platforms, he reports

As for cost, the company experienced savings through the use of assist processors like the zIIP. Again, as Campbell explains: the incremental workload

would be zIIP eligible and probably not even cause a blip on Univar's MLC usage reports. The company needed no new hardware and no new software licenses.

Furthermore, since in Campbell's experience Workload Manager for z/OS (WLM) does a better job of maximizing utilization than alternative distributed virtualization offerings the team could safely introduce new BI workloads knowing they wouldn't swamp the higher priority demands yet still take advantage of as much of the z as is available. They considered this a significant advantage over any distributed platform, which typically struggle to support diverse workloads.

BI Implementation

Univar's tactical BI implementation began with a simple need: to add new measures to the data model used in the company's data warehouse. That brought a Univar team to an IBM Information on Demand conference where they saw InfoSphere Warehouse for System z and the associated BI modeling extensions to InfoSphere Data Architect. That suggested they could achieve a rich BI environment without having to build a new BI infrastructure from scratch.

Instead, the team saw that they could model data at the transaction level and publish it to InfoSphere's cubing server running on zLinux. That would let them leverage the existing DB2 data on System z. InfoSphere Warehouse for System Z at that time was available as part of a beta program so they could evaluate its functionality and the viability of the approach without expenditure.

To test the approach, the BI team first put an area of the Oracle data warehouse on a tiny LPAR on the z. The results surpassed their benchmarks and were quite easy to implement. From that, the team concluded that the System z was a good candidate for revamping their BI approach. It would involve tools they already had, skills they already knew, and avoid much of the ETL effort that bogged down previous attempts.

Working from overnight point-in-time snapshots of the day's transaction data, the team could avoid remodeling the data, instead simply altering it logically to use in their star schemas. In effect, they built BI models on top of the data

while leaving the core transactions alone. InfoSphere Cubing Server would handle whatever transformation was required in memory enabling very fast model changes.

Fast forward ahead a few months. The InfoSphere beta program is over and Univar, under budget constraints, deferred any BI purchases until the conclusion of a strategic roadmap project. However, managers had acquired a taste for basic BI on the System z.

Without the modeling and in-memory cubing features of InfoSphere Warehouse the team instead implemented much of their transformation and business logic with DB2 as views. These views were then published to Univar's Crystal Reports developers whose reports executed directly against DB2 ODS data on the z.

Results

In its somewhat unconventional, low budget approach, the Univar BI team achieved surprisingly effective tactical BI by defining a logical model on top of the transaction data and performing light ETL on the fly. Certainly users wouldn't experience the rich cubing possible through InfoSphere Warehouse, and they could only go back 18 months, the amount of history available in DB2. Most users, however, only wanted to look back a few weeks or months anyway.

The team initially expected most users to access the data through Excel where managers could take advantage of pivot tables to create multi-dimensional views. Instead, more users were opting for Business Objects as the front end. This worked but required iterative querying to get the multi-dimensional views they would otherwise get from InfoSphere Warehouse Cubing Server.

What Univar built was a poor-man's version of BI. They achieved it because they could leverage the existing investment in the System z and DB2 and the capabilities each brought to the effort. They also could take advantage of software licenses Univar already had. The result was effective tactical BI with only minimal incremental expensive and effort.

Lessons Learned

The Univar BI experience teaches a number of lessons that are especially valuable when resources are constrained. Despite a limited budget, the team was able to deliver more than adequate BI and avoid the need for a massive and costly revamping of the organization's data warehouse and BI infrastructure.

Along the way it learned the following:

- Take advantage of learning opportunities such as IBM Information on Demand conference
- Prepare to fully leverage existing assets, especially the System z and DB2
- Learn from users about the data they need and how they use it
- Continue to demonstrate the value of BI to management

Maybe when the economy turns around and as management gets a taste for the value BI can deliver, they may be more willing to provide the resources for tools like InfoSphere, which is all that is still needed.

Independent Assessment analysis

BI belongs on the System z. That's where the data lives and that where it is processed and stored. Although BI started on the mainframe, it shifted decades ago to distributed platforms. IBM's latest investments in the System z environment—everything from Linux on z, SOA, InfoSphere, and more—makes possible a shift back to BI on System z.

The Univar BI team's effort is innovative and should provide a model for other organizations experiencing resource constraints. What they did that others can learn from is very simple, they:

- Educated themselves about BI at minimal cost
- Made the most of what they could do with existing tools
- Fully leveraged the resources they did have,
- Understood their data and what could be done with it within their constraints

- Watched how their users used the data and made appropriate adjustments

There are IT teams everywhere that can take these lessons to heart. And when the economy improves, these will still be worthwhile lessons.

About Independent Assessment

Independent Assessment (<http://www.independentassessment.com>) is the IT and business assessment, analysis, and writing service of Alan Radding, an independent business and IT analyst/writer for over 20 years. It provides independent ROI and TCO analysis, competitive assessment and positioning reports, case studies, white papers, and Web content.

Independent Assessment publishes *dancingdinosaur*, the independent blog covering the System z, <http://dancingdinosaur.wordpress.com>