

IBM Cognos BI for Linux on System Z

Unmatched performance
and scalability



By Dave Beulke and Associates
October 2009

Numius, an IBM business partner, recently performed a proof-of-concept (POC) for improving and lowering the costs of one of its client's business intelligence systems. The client's distributed business intelligence environment was facing degrading performance as query workloads and user populations increased. The POC demonstrated the advantages of porting the client's distributed system to a centralized Linux on System Z environment. The results of this POC port demonstrated-

- The significant scalability improvements when consolidating a distributed set of servers in to a single, System z environment
- The ability to deliver the same functionality provided by the multiple distributed servers
- The capabilities inherent within System Z, including the security, reliability, and availability strengths of System Z platform for business intelligence workloads
- The ability to deliver the performance and scalability necessary to support very large complex query workloads and potentially thousands of concurrent Cognos BI users


This white paper discusses the client environment, the POC architecture, its query workload, and the best practices employed for achieving the scalability for large query workloads and large user populations.

The client environment and POC architecture

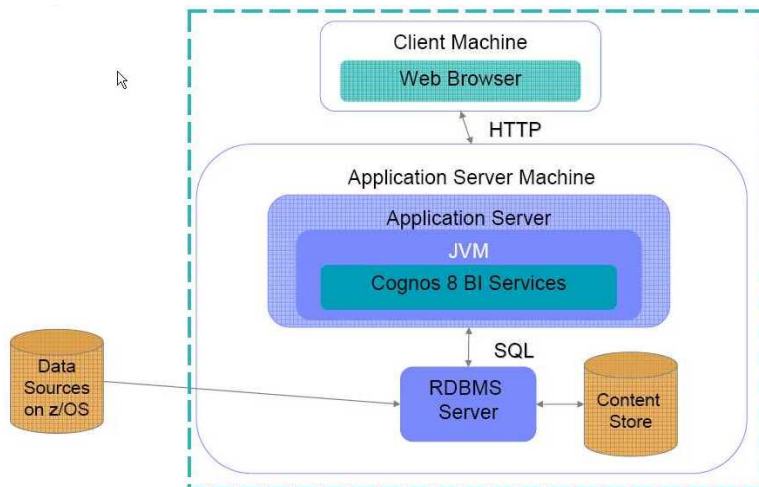
The Numius client provides administrative and financial data services to approximately 10,000 Belgian companies, mainly in the retail sector. Using Oracle Forms and file input, Cognos Powercubes and nightly reports experienced degrading performance as the number of clients and queries increased. The client wished to expand its individualized service to corporations and business-to-business clients, however the performance issues were prompting the client to explore outsourcing the entire infrastructure to improve operations.

The Numius team, consulting with their client, leveraged the IBM Montpellier Customer Center to understand, demonstrate and validate the capabilities of IBMs server, storage and software technology. The Montpellier Customer Center has three different environments for testing and benchmarking customer workloads, a z/OS environment, a Linux on System Z environment and an Oracle RAC environment. These environments are available for benchmarking any client's environment and analyzing IBM's technology for any POC situation.

The client's original distributed server configuration, shown below, included a number of intel based servers that could be centralized on a single System Z server for the POC.



The Montpellier Customer Center consolidated these different server workloads, by leveraging Cognos BI for Linux on System Z, with DB2 Version 9 for z/OS to server as the centralized database. Three stand-alone web client workstations were used to submit the client’s workload. The logical perspective of that environment is below.




The Scalability of the new environment

The workload and the simulation tests leveraged the client's own application and data to replicate the client's existing end-user environment.. To accurately duplicate the workload, the web clients submitted batch loads to the Cognos BI layer that queried the data on DB2 9, producing a PDF document of the resulting report saved to the content store.

Multiple workload tests were executed to understand the scalability and performance characteristics, and determine the point where the productivity peaked. The queries were complex and simulated a variety of Cognos workloads including four different reports, 130 named users, and a number of scalability situations.

For the POC analysis, four reports were developed to create realistic workload situations. These included a small report using a simple database query, a large report using a complex database query, a small report executing an OLAP database query, and a large report using a complex OLAP database query.



One hundred (100) different workload variants of these four reports were used in the performance testing. These workloads ran in parallel lane groups to simulate large numbers of concurrent end users. The parallel lanes channeled and simulated the concurrent workloads on the different Cognos web, Cognos services, and database components within the POC system.

Using the parallel lanes method, the POC simulated concurrency of large numbers of users maintained over extended periods of time, in order to test the overall scalability of the POC environment. Using this method, groups of 10, 21, 51 and 100 parallel lanes of concurrent users executing 100 variants of the four different reports, were channeled through the system. These different groupings used randomized report sizes to simulate up to the 130 named users within the test system. While the client's original distributed server environment struggled with 12 end users, the POC demonstrated the ability to easily handle 130 users, representing a ten times (10x) improvement over the client's original environment, on a single foot print.

The POC also completed workload tests to demonstrate the benefits of DB2 database parallelism and the System Z workload manager software (WLM). The WLM software offers a functionally rich, sophisticated environment capable of balancing all the virtualized workloads and their resource consumption. WLM balances the allocation of resources to all portions of the workload on a priority basis, while driving higher utilizations of all available CPU resources across the various Cognos virtualized components. Leveraging parallelism and WLM capabilities further improved the overall throughput, scalability, and overall response time of the Cognos solution.

During this POC, the randomized workload demonstrated that WLM was able to maintain efficient work flow in the system, preventing the large queries from clogging up the system, while allowing the fast, concurrent execution of small queries in the environment.

This capability was especially important, as the client's original distributed Oracle/server configuration was unable to service these different workloads concurrently and found it difficult to provide adequate facilities to balance resources across the different workload types.


Within the POC, the Cognos configuration and architecture executed and completed 400 large and small queries while the original client distributed Oracle/server configuration was dominated by only a single large query. The resulting throughput was 400 times better on System Z. Additionally, some reports that were never able to complete on the original client's distributed architecture, executed easily on the System Z POC solution.

Additional tests were conducted utilizing WLM discretionary service capabilities to exploit the specialty ZIIP processors within the System Z environment, which lowered the general processor requirements for the workload. This is extremely beneficial in reducing the overall cost of running the workload on System z, as these specialty ZIIP processors are not included in the MSU software-pricing model charges and let the workload fully utilize their capacity without affecting the core price of the software licenses.

Phase 1 of WLM testing:

In the first test, the WLM was established without its discretionary service mode. The results showed excellent response times with the shortest elapsed time to execute the 100 channeled groups of four reports. During the execution of the workload, all the general and ZIIP specialty processors were utilized at 100% at times, while CPU utilization peaks and valleys reflected the different size reports and different query workloads on the different components of the POC configuration.

- **Result without z/OS WLM Discretionary Service Goal**
- Output is measured at the end of every two minute slot during the test.
- We have no report failures (400/400 successful).
- In all tests, we have a peak at the start, followed by a "wait time".




In the next phase of the POC, WLM was set up with discretionary service mode, the workload showed excellent response times, and more overall elapsed time to execute the same 100-channelled groups of four reports. This elongated elapsed time was because the WLM discretionary service mode shifted the workload to use 100% of the ZIIP specialty processors before utilizing the general processors. Successfully minimizing the general processor usage results in lowering the MSU software- pricing model charges, while maximizing the use of these specialty ZIIP processors without affecting overall IT costs.

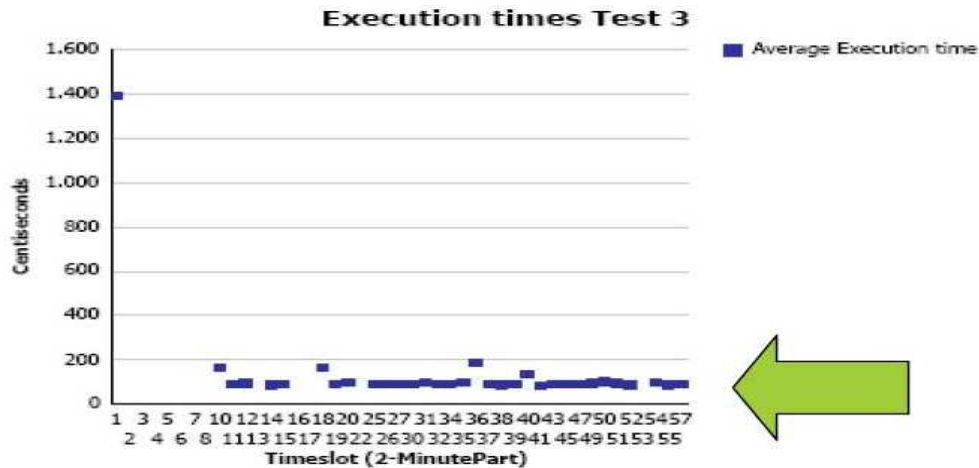
These WLM tests demonstrated that as the workload increased the response time became more regular. So regardless of the amount of queries going into the system, the end user response time is consistent, even as the workload on the server became larger. By delaying long running queries, while prioritizing the shorter more critical queries, WLM allows for the smart utilization of resources within the system. This is a huge advantage over the client's original distributed Oracle/server configuration since a single large query would hold up or clog processing of all other queries in the workload.

By efficiently sharing the general processors across a number of the architectural components within this implementation, and by efficiently leveraging the specialty ZIIP processors, the System Z configuration was able to maximize the overall throughput, scalability and response time for the workloads. Offloading significant amounts of the processing demands from the general CPU to the specialized ZIIP processors reduced the overall end-user response times while improving the consistency and scalability of the workloads.

- As the load on the system increases through a higher degree of parallelism, the degree of variation in report execution times is reduced.
- On a system that is only scarcely loaded, report execution times may go extremely fast (no other processing going on) or slower (accidental overlap with complex other processing). On a heavily loaded system, average times may go up, but the behavior becomes more regular.



This study further reinforced the overall configuration of the System Z platform for BI demonstrating that as the workload increased, the WLM facilities, the architecture of the Cognos solution, and the off load of work to the ZIIP processors, delivered a consistent average query response time of less than two seconds for the large number of parallel users.



Within this POC, the workload naturally consumed peak I/O and CPU when the largest number of users were executing the largest reports. During these peak processing times, the study showed that the System Z WLM software properly distributed the workload to its specialty ZIIP processors.

In addition, as the workloads grew on System Z, the results showed their overall response times were linear and consistent no matter how big the database was, how many users were concurrently using the system, or the type of query workload that was executing in the system. In the POC test, the IBM Cognos 8 BI solution demonstrated throughput that was consistent and scalable.

The Numius POC test documents that the System Z platform delivers linear scalability and query performance when doubling the computing resources available to the Cognos system. The unique capabilities of the System Z features effectively utilize resources to ensure that response times and throughput performance are consistent and linear as the workloads increase.

Summary

This Numius POC demonstrates that IBM Cognos 8 Business Intelligence on Linux for System Z is an optimal solution for BI and DW workloads. Its linear performance and scalability when scaling from small to large data warehouses, and its ability to support large numbers of users is unmatched by any other platform or solution available.

These tests show that the Cognos solution can execute a mixed workload ranging from simple to very complex BI reports with excellent response time. They show that the Cognos solution can take full advantage of System Z's inherent strengths and WLM capabilities to deliver excellent response times for all users, while completing many different queries without interference or delays from complex or long running requests.

- System Z can “virtualize” potentially hundreds of distributed BI server environments

- Distributed servers can be consolidated and scale within the Linux for System Z environment without any major changes
- System Z WLM software can balance and leverage its specialty ZIIP and general-purpose CPU resources to maximize performance
- System Z specialty ZIIP CPUs minimize general purpose CPU usage, thereby lowering software licensing costs
- IBM Cognos 8 Business Intelligence on Linux for System Z can perform better than distributed server architectures and can scale linearly as demand grows, to support large numbers of users and large complex query workloads

Finally, the Cognos architecture, with its tiers of distributed components is a perfect fit for the “virtualized” System Z environment. The IBM Cognos 8 Business Intelligence on Linux for System Z provides the best overall solution for any company to build, consolidate or operate their BI DW environment because it provides the reliability, performance and linear scalability at a competitive cost of ownership.

References

Introduction to IBM Cognos 8 BI on Linux for System Z
Internal IBM paper available from your IBM Cognos representative

50 TB Data Warehouse Benchmark on IBM System Z
www.redbooks.ibm.com/abstracts/sg247674.html

IBM and Numius:
Performance and scalability results for data warehousing and BI on System Z
www-01.ibm.com/software/os/systemz/telecon/jun2/

About the Author

Dave Beulke and Associates, a division of Pragmatic Solutions, Inc is an independent technology research, analysis, education and consulting organization specializing in architecture analysis, database performance, data warehouses designs, and internet architectures.

Other research and analysis conducted by Dave Beulke can be found at www.davebeulke.com.

Dave Beulke and Associates
3213 Duke Street
Suite 805
Alexandria, Virginia, 22314
(703) 798-3283
Dave@DaveBeulke.com

