

Informix Accelerates Analytic Integration into OLTP

Reference Code: IT014-002611

Publication Date: 16 Jul 2012

Author: Madan Sheina

OVUM VIEW

Summary

Informix's database technology, now under IBM's wing, still garners a fiercely loyal customer base. IBM has been working steadily on its roadmap to extend the core Informix database kernel to deliver a scalable, high-performance platform for building and running high-end enterprise data warehouses (EDWs). The company has also released Informix Warehouse Accelerator (IWA), which integrates with Informix databases to offload and run business intelligence (BI) and complex analytic data queries at a consistently fast pace. As IWA continues to mature, the question is whether IBM can break the Informix brand out of its traditional niche of online transactional processing (OLTP) and embedded applications, and stand tall in a fiercely competitive world of high-performance analytics. Technically the IWA-Informix combination possesses all the necessary capabilities. But what is its appeal to organizations that have not already invested heavily in Informix for OLTP processing?

ANALYSIS

Speed is everything for IWA

IWA is the fruition of IBM's "Blink" research project to develop a scalable query engine that consistently responds to BI and analytic queries against a data warehouse without the need for a complex "performance layer" of indexes, materialized views, and pre-computation. From a product perspective it is best described as the data warehousing edition of the Informix 11.7 database. It is an analytically-aware and pre-tuned software-only appliance that is designed to be used as an add-on to an Informix database server. Building on heritage Informix capabilities – XPS (for massively parallel processing) and Red Brick (for star schema datamarts/warehouses) – IWA provides a fully workload optimized system for rapid query processing throughput. A high level of performance is achieved using multi-core parallelism, in-memory processing, and columnar data representation (i.e. there is no need to maintain aggregate tables).

IWA slices operational data from Informix OLTP databases and creates a highly compressed, in-memory analytic datamart that can be accessed concurrently by multiple users firing off analytic queries without fear of performance degradation. IWA executes queries serially, one query after another, but using full parallelism for each single query. Based on IBM's claims, the performance gains are impressive; most reference IWA customers (mainly large government agencies and retailers) tend to see gains of between 90x and 300x faster than alternatives.

IWA is tightly coupled with the Informix database server, and is typically deployed as a component of the Informix Ultimate Warehouse Edition – which also bundles the Informix 11.7 database server, the IBM Smart Analytics Optimizer Studio (an integrated Eclipse-based interface for configuring and managing the “accelerator” component and defining/deploying datamarts from the Informix database to IWA), and a Storage Optimization Feature (i.e. compression scheme). There is also a Growth Warehouse Edition that is aimed at mid-size companies that want to start within a restricted configuration (16 cores, 48GB of RAM) but want to safeguard future deployment scalability.

Informix customers that want “in-line” analytics will benefit most

Companies that stand to benefit most from IWA's speed are Informix's legacy customer base. The tight links between IWA and transaction-oriented Informix systems offer the ability to integrate analytic business insights into operational processes.

Logically, this means bringing analytics closer to transactional data. Architecturally this is achieved by balancing query processing workloads across two servers processing in tandem: a primary HA Informix server for OLTP workloads where data is frequently changed, inserted or deleted; and a secondary server that deals with OLAP workloads on data that change only very little or not at all. That said, IWA can also work with an Informix database server that supports only a data warehouse database. Supporting both operational and analytic workloads with the same system is a relatively unique idea that is also being pursued by rivals such as Oracle. Ovum believes that IBM now has a strong story to tell with IWA.

The software-only design of IWA is also well suited to flexible cloud deployment, particularly when combined with Informix technologies such as High Availability Data Replication (HDR), Enterprise Replication (ER), MACH11 (multi-node active clusters for high availability), and Flexible Grid. It is difficult to build Oracle RAC or a DB2 pureScale architecture in, for example, an Amazon Cloud since they require expensive hardware components such as InfiniBand switches.

IWA expands data warehousing options for IBM customers

IWA sits in the middle of a spectrum of EDW options that IBM now offers to its customers. On the extreme left is the Netezza appliance. Center-left is the IBM Smart Analytics system, which is positioned as a flexible, integrated system. On the extreme right is the flagship custom IBM InfoSphere Warehouse. IWA sits center-right as one of two data warehouse accelerators that IBM develops; the

other being DB2 Analytics Accelerator for zOS mainframe systems. The spectrum runs the gamut from simplicity (an appliance) to flexibility (custom build).

Arguably much of IBM's EDW spotlight (and marketing effort) has focused on the Netezza and InfoSphere platforms. But over the past two years it has also started to promote IWA as an alternative option aimed at providing analytics to an OLTP environment, even though it can run as a separate warehouse. IWA certainly provides an opportunity for existing Informix customers to engage in industrial-strength and high-performance analytics, and do both OLTP and OLAP processing on the same platform. Hence IWA will appeal greatly to die-hard Informix customers who use it predominately for OLTP-based and embedded database applications. This represents a substantial market. For example, Informix continues to be the embedded database of choice for applications such as Cisco's Call Manager.

But it is not just the prospect of having a single platform vendor for OLTP and data warehousing that attracts. Deep integration and transparency between the two database environments is perhaps more compelling, effectively providing companies with a way to use Informix as a single platform for both OLTP processing (for example, supporting a web storefront) and OLAP processing (for example, analyzing retail data) while retaining service-level expectations for both applications.

To sustain this level of performance IWA incorporates technical features such as frequency partitioning, where the system looks for the most frequently occurring data values in columns and encodes those values with the least number of bits that can adequately represent the data; and predicate evaluation directly on compressed data to reduce the amount of data transferred from memory to the CPU. Informix 11.7 also has features significant for IWA such as: interval partitioning (to support time-cyclic analysis based on when new data flows into the data warehouse); synchronized data refreshes between IWA and Informix environments; and support for high-availability cluster configurations.

IBM also continues to improve integration between IWA and Informix environments with its Cognos front-end BI toolset, including a native content store on the Informix database and SQL merge capabilities. DBAs will also appreciate IWA's pre-optimized, "no-knobs" structure which eliminates many of the mundane and resource-draining tasks associated with tuning and tweaking complex queries – i.e. you don't have to worry about maintaining indexes, materialized query tables, etc.

APPENDIX

Author

Madan Sheina, Lead Analyst, Information Management

madan.sheina@ovum.com

Further reading

IBM blends analytics and operations at runtime (July 2012)

Near-realtime analytics in banking: delivering on old promises? (June 2012)

Disclaimer

All Rights Reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior permission of the publisher, Ovum (an Informa business).

The facts of this report are believed to be correct at the time of publication but cannot be guaranteed. Please note that the findings, conclusions, and recommendations that Ovum delivers will be based on information gathered in good faith from both primary and secondary sources, whose accuracy we are not always in a position to guarantee. As such Ovum can accept no liability whatever for actions taken based on any information that may subsequently prove to be incorrect.