New Technologies for Top CIO Challenges

By John Ladley

Overview

CIOs face shifting priorities and business issues every year. The list of challenges gets longer and longer. According to a recent survey by CIO Magazine (November 4, 2004), the top priorities for CIO's in 2005 are:

- *Aligning IT and business goals*
- *Increasing business efficiency through IT-enabled process improvement*
- *Creating competitive advantage through IT*
- *Improving internal customer (user) satisfaction*
- *Controlling IT costs*
- *Enabling/enhancing knowledge management/leveraging intellectual assets*
- *Staff development/leadership, developing business skills within IT*
- *Improving project management discipline*
- *Measuring and communicating bottom-line impact from IT*
- *Ensuring regulatory compliance*
- *Ensuring privacy of customer and employee data*
- *Implementing mechanisms for IT governance, portfolio management, etc.*

Many of these priorities are internal IT improvements or technology-agnostic, but several of them can be explicitly addressed through technology. This paper will explore the applicability of a new technology—data warehouse appliances—for those technology-oriented priorities. For the sake of this discussion, it is appropriate to group them into three scenarios: 1) Managing the total cost of ownership of IT, 2) Improving processes, and 3) Aligning with and supporting business goals.

Evolving Data Warehouses

There are numerous approaches and technologies available to CIOs to address business problems. One of the solutions within information technology which has dominated the CIO's radar for the last ten years, is the "data warehouse." However, the data warehouse has changed from the static historical collections of data to a wide range of potential solutions. Two drivers of changing data warehouse technology have been lowering latencies and increasing data volumes. The response to these issues has been the appearance of a technology known as data warehouse appliances.

Data warehouse appliances (DWA) are specialized data base/hardware hybrids (pre-integrated hardware and software) that can act as the primary platform for a data warehouse, or act as a back-end supplement to an existing data warehouse.

The various vendors in this space all have slightly different views and metaphors, but in general, the DWA works as a server that is specialized for high speed query capability. Data warehouse appliances (such as DATAllegro and Netezza) present a relevant alternative with a variety of potential applications. Their prime purpose is more speed, more data, and lower costs. However, a technical set of "whiz-bang" characteristics are not going to sell the concept. Business solutions are necessary as data warehouse technology matures. To best illustrate the role of this technology, we will examine how data warehouse appliances could be used in the aforementioned scenarios facing CIOs today.

Scenario 1 - Managing the Total Cost of Ownership of IT (TCO)

According to the survey, CIOs will be focused on controlling costs and IT architectures in the coming year. This means not only reducing and managing expenditures, but gaining efficacies through technology. In addition, CIOs must extract as much leverage as possible from existing data and knowledge.

The maturation of the data warehouse and data warehouse technology has brought scrutiny on the cost of operating and owning these environments. A sustainable data warehouse manages the ongoing total cost of ownership and constantly measures the benefit to the business. There have been instances of the data warehouse being turned off due to cost issues. Often, data warehouse environments become bloated report factories. When the number of so-called data warehouse developers (aka report writers) are added into the end-user training, license under utilization, constant data base appliance tweaking, and on-going development demands, the costs of ownership of large centralized data warehouse environments must be constantly scrutinized to ensure on-going return on investment. Given the price of upgrades to hardware, software licenses, and personnel cost, expanding the data warehouse environment becomes potentially as costly as the original project.

A common remark heard over the past 10 years in regard to specialized data warehouse technology is that it is "proprietary." The implication is that the technology will be non-standard within a particular IT shop, and therefore raise the total cost of ownership. Specialized resources will be needed, and the CIO will be at the mercy of the vendor. So-called integrated offerings of business information and ETL tools have also fallen short and best of breed products are still flourishing in these categories.

The Truth About Proprietary
Experience has shown the assumption about proprietary technology to be untrue. Each of the supposed advantages of using "standard" technologies really does not hold water.

1. Leverage existing staff - all the main stream RDBMS products (UDB, Oracle, etc.) require special training for data warehouse utilization. Actual experience shows very little overlap between database administrator support of operational and data warehouse data bases.
2. Leverage existing platforms – Actual practice has demonstrated several challenges when shops try to operate a data warehouse instance in the same environment as the operational system. Back plane contention, network overload, and security are all issues that push data warehouses to separate platforms anyway. At this point, licensing costs increase when the new server is introduced.
3. Leverage vendor relationships – Unless there are rigid negotiations, common off-the-shelf data warehouse offerings from mainstream vendors are priced separately. Only through aggressive negotiations have IT shops seen leverage by using data warehouse products from their operational data base vendors.

The reality of using a DWA is that existing staff and platforms are leveraged, and there is minimal impact on the existing cost structure. In this manner, enhancements to the core warehouse functionality can actually be executed for less TCO.

This approach also supports managing TCO through leveraging existing information management capital. The data warehouse appliance offers data warehouse environments a means to incrementally improve performance, capacity and content without considerable upheaval to the original environment. Often, this means a simpler business case as well.

Scenario 2 - Aligning with and Supporting Business Goals

Alignment to and enabling of the business has been a continuous challenge for CIOs. Far too often, strategic projects like a data warehouse become operated or funded by one area, and end up being part of the very silos the data warehouse was designed to prevent or cure.

Often, the original data warehouse morphs from its original intent (static, time variant, analytical) into a lower latency, pseudo-operational application. When the original data warehouse architecture can no longer support business needs or direction, it is often more prudent to add on structures that can be fed from the main data warehouse, but can be used to merge other source data, or deal with new latency requirements. The appliance can act as a vehicle for exploration, special reporting, reduced latency, or sophisticated mining.

This ability to respond rapidly, yet still maintain the core structures is also more competitive. Decreasing latencies are usually brought on by the need for business to respond more quickly. Increasing data volumes also tend to be promulgated by business changes. For example, moving from static, market-based analysis to real time response to web clicks.

One of the higher ranked challenges to CIOs in 2005 is the management and support of compliance issues. Business pressures related to Sarbanes Oxley (Sarbox), Basel II, etc. have placed additional burdens on IT shops. Most IT efforts are related to installing tactical controls, or installing re-hashed internal control products.

Compliance issues first and foremost promote data volume challenges. Sarbox, for example, promotes storage of documents that are scanned or of less structured than traditional "row and column" data. Use of a data warehouse appliance will offer not only an affordable common storage point of critical documents, but the population process enables documentation of precise, compliant data controls.

IT shops have been told that Sarbox compliance is a wonderful opportunity to accomplish great architectural things. In reality, the business wants quick visibility to show regulators that they are in compliance without spending more than they must to achieve that.

The reality is that long term management of information to maintain compliance is low on corporate radar. The data warehouse appliance can have a roll here. In a compliance scenario, the data warehouse appliance can act as a flexible add-on integration point. Rather than muddying-up an existing data warehouse or enterprise requirements planning package (e.g. SAP) with the burden of integrating data for Sarbox, the appliance can act as an add-on facility.

Scenario 3 - Improving Processes

Business efficiency improvement has always been a great motivator for IT. At the same time, CIOs continuously strive to improve the perception of IT to their internal customers. These two priorities go hand-in-hand.

For example, the data warehouse appliance acts as an integration point of financial data from disparate sources. Many very large, multi-national companies are exploring this type of integration strategy at this time. Consolidations, timing issues, and currency changes create an environment where even the most savvy business processes cannot keep up with changes to data. The DWA offers assistance here as well whereby the myriad of currency variations and final timing issues can be offloaded to a platform that can restate the information at any time without overloading existing data warehouses. This permits a single consistent view of financial data that is not only useful for enterprise performance measurement, but drastically reduces financial processing cycles.

Risks

The aforementioned advantages to data warehouse appliances need to be balanced against potential risks. Foremost is the use of a data warehouse appliance as an "instant warehouse." While departmental usage of these appliances is certainly a possibility, they are subject to all of the rules of data warehouse best practices. There must be a reasoned development of a businesses aligned strategy and business performance. The dangers of data quality and transformation lurk in the background even with these products. Data must be used to add value to an enterprise. Therefore, business processes must explicitly be supported.

The New Data Warehouse

Most of this paper has focused on sustaining an existing data warehouse. For those organizations that are just starting a data warehouse, these appliances offer a legitimate alternative as a stand-alone solution, especially when large amounts of data are suddenly confronted. Examples include scenarios where statistical analyses of massive numbers of events or massive data clicks are required. (e.g. process control from manufacturing or genetic research)

Additionally, these products embody 20 years of lessons learned in handling large complex data sets. The major vendors in this space have created DBMS engines designed specifically for the task of data warehousing.

Do We Need One?

Remember the value proposition of these products – lower total cost of ownership, less database administration overhead, and extraordinary performance increases. These benefits translate into a checklist of technical challenges, that, if indicated, may mean a data warehouse appliance is worthy of exploration.

Performance
Are decreasing latency requirements challenging your environment? If these requirements can be isolated into a sub-set of the data warehouse, or are focused on one portion of the business community, then adding a data warehouse appliance to an exiting data warehouse offers a means to address latency issues without disrupting the current architecture. If the product integrates well with your existing data warehouse database management systems, then this is even better.

If query performance and volume issues are challenging your data warehouse, again, the data warehouse appliance offers a relatively efficient means to address these specific issues. The addition of huge new subject areas is not uncommon, and frequently requires re-engineering large portions of a data warehouse.

Cost of Ownership
If the data warehouse is becoming hard to justify due to high overhead in terms of support, i.e. too many database administrators, these products offer intelligent data engines that actually reduce the need for additional database administrators. Again looking at the addition of a large subject area, the data warehouse appliance offers an alternative if the production maintenance processes of the current data warehouse are difficult to upgrade.

Summary

Data warehouse appliances have added another effective tool to the many components that are available to CIOs. They offer a valid alternative to updating entrenched data warehouse architectures. They also offer cost effective back-end processing to lower costs, and even offer an interesting alternative for data new warehouse projects. The challenges that CIOs will be addressing over the next few years strongly indicate the need for alternative, creative application of technology to meet business needs.

John Ladley is an internationally known practitioner and a popular speaker on information and knowledge management. John is widely published and has several regular columns. Prior to founding KI Solutions, John was Senior Program Director of Data Warehouse strategies and a Research Fellow at Meta Group. Mr. Ladley is an authority on information architectures, business performance measurement architectures, knowledge management, collaborative applications, and information resource management. He can be reached at jladley@kisolves.com. KI Solutions is now a part of Navigant Consulting.

Thank you for reading this complimentary paper provided by DATAllegro.
For more information on this topic or DATAllegro products, please contact us at (949) 330-7690, or visit *www.datallegro.com*.