

(GOOD)

BI!

CRUEL
WORLD



{ Business intelligence vendors are struggling to expand their reach. Our prescription: Go real time, for real users, and integrate with existing practices more effectively—or risk falling into irrelevance }

By DAVID STODDER

LET THE POLITICIANS DEBATE AD NAUSEAM whether isolationism works in the global arena. The argument's been settled on the business-infrastructure front: No information system is an island, and he who depends on outdated data loses market share. Business intelligence vendors must release products that mesh with these realities if they expect to expand their empires.

And the territory is certainly there for the grabbing: A poll released last month by our sister publication INFORMATIONWEEK showed nearly half of the 500 IT professionals surveyed plan to increase spending on software for viewing and analyzing business information from 2006 levels (nwc.com/go/0402bi-spending).

Forrester says BI platform revenues will reach \$7.3 billion by 2008, and CIOs surveyed by Gartner identified BI as their No. 2 technology priority last year, up from No. 10 in 2004. Despite this, analysts have long puzzled over the relatively low penetration rate of BI tools, generally pegged at less than 20 percent of potential customers.

Why the disconnect? After all, BI suites provide the platforms from which critical data is aggregated, searched, presented and analyzed. Sure, it's a complex process. Data must be pulled from disparate sources, such as ERP, order entry and inventory management. But up-to-date information is the lifeblood of busi-

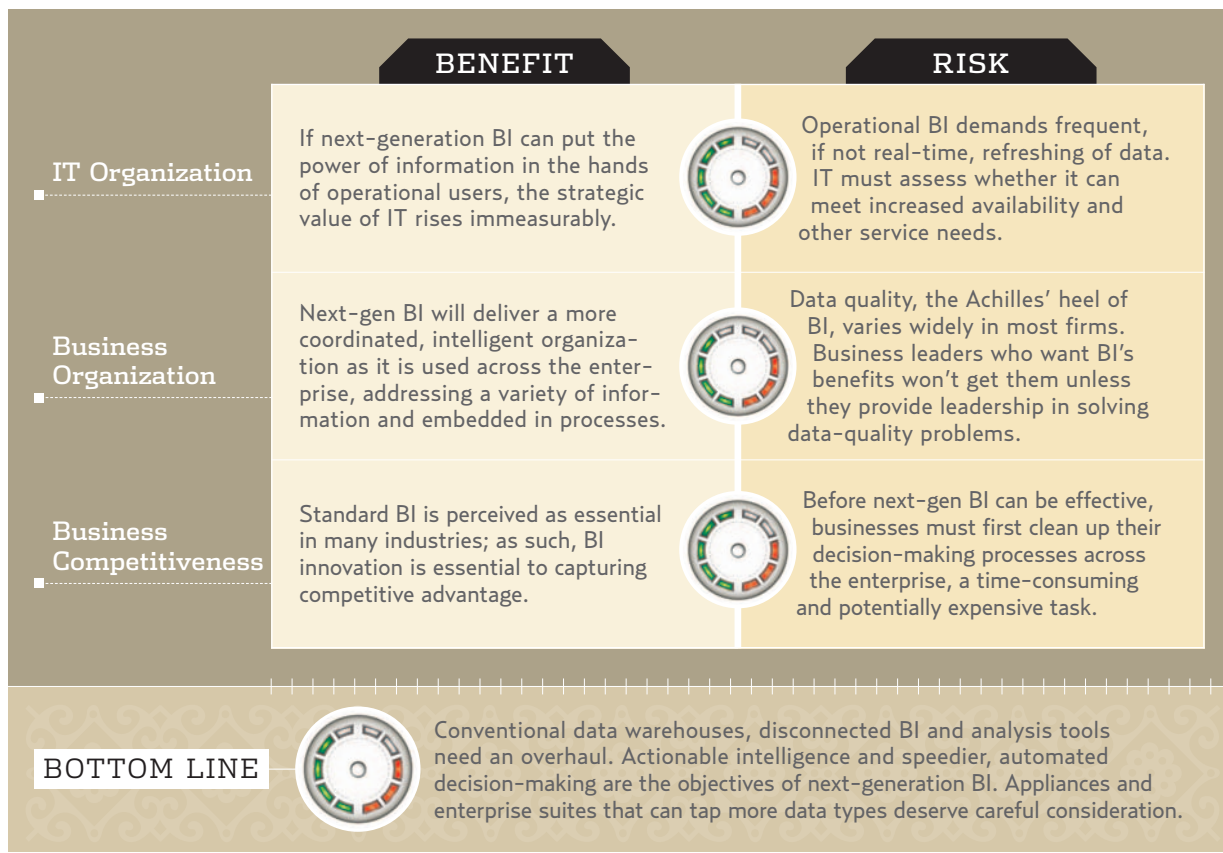
ness. How can four out of five not be buying in?

The answer may be that too many BI platforms are mired in historical analysis across siloed back ends. The action is in a new, more integrated world of dynamic, real-time information that's emerging from Microsoft, Oracle and other BI vendors. These suites empower real users—not IT pros drafted into duty—and let them draw valuable data from processes, events and other sources beyond conventional data warehouses.

Their vehicle for display? Real-time dashboards that process up-to-the minute information and present it for immediate use and analysis. But that level of integration carries risks: Data sharing within BI brings up serious security, compliance and privacy concerns. And then there are the turf wars, as departments, employees and business partners scramble to protect data—their prime intellectual asset—from internal and external competitors.

Still, a new vision of real-time, networked intelligence is possible thanks to affordable computing and storage platforms. These, combined with advances in leveraging BAM (business activity monitoring) to track strategic business objectives and cost-effective data-warehousing appliances, should help vendors extend the use of BI throughout the business world. Now, will they rise to the challenge? And will IT buy in?

IMPACT ASSESSMENT: NEXT-GENERATION BI



UNITED BI NATION

The list of BI players is a veritable who's who: Big application and systems providers—including IBM, Microsoft, Oracle, SAP and most recently, Hewlett-Packard—are trumpeting BI, analytics and data warehousing and are challenging in markets once owned by pure-plays such as Business Objects, Cognos, Hyperion, Information Builders, MicroStrategy and SAS. In addition, systems integrators and resellers are clamoring for BI expertise to buff up their services.

In mid-February SAP acquired Pilot Software,

which helped define the term *business intelligence* and now specializes in performance-management applications based on its BI/analytics platform, PilotWorks. More recently, in early March, Oracle made an offer to buy Hyperion. If the deal goes through, Oracle will muscle up its portfolio with the addition of Hyperion's Essbase OLAP server and BI tools. More importantly, however, Oracle will assume the market lead for performance management tools and financial applications. Market analysts predict more mergers and acquisitions to come.

SURRENDER THE SPREADSHEETS

THE ENTERPRISE BI INDUSTRY has grown to its present semi-lofty stature by convincing enterprises that tools independent of their data sources, such as ERP, provide for easier, more consistent and better decision-making. It asserts that BI software draws on shared resources and provides richer context, specialized for turning data into actionable intelligence, and points out that leveraging data for BI has not been the core competency of transaction-oriented application or database providers. And indeed, Oracle and SAP are taking BI more seriously than they did before: Not only are they expanding their tool portfolios, they're dropping more BI, data warehousing and analytical functionality into their service-oriented middleware platforms, such as Fusion and NetWeaver. Meanwhile, IBM, Oracle, Microsoft and Teradata are pushing more BI, data-mining and other analysis functions into their database engines to support more complex user queries. It's all very integrated and strategic.

So why are simple Excel spreadsheets still the BI tool of choice for many?

This speaks directly to the very human desire to maintain control over the coin of the business realm, knowledge. We understand the resistance, but IT must educate users that depending too heavily on spreadsheets can drag down the business as managers argue over who has the correct data. In exchange for less personal control over information, enterprise BI promises to enrich the user experience with shared access to better-quality information. To wean users off Excel, consider tools from Microsoft, Hyperion's SmartView and SAP's Duet that are tightening the BI link between desktop and enterprise applications.

Still, the hard truth is that to stay competitive, businesses must do more than share data effectively; they must clean up their decision-making processes across the enterprise. This is a daunting, expensive task: BI, spreadsheets and departmental data marts

added to myriad applications with their proprietary logic and rules all contribute to the "silo" problem.

Again, IT can educate users that BI—especially as it shifts to meet performance-management objectives—will go far to help companies execute smarter and more efficient business processes, while also bringing order to complex, collaborative decision-making. Once an organization has done its housecleaning, it's easier to implement rules-based automation of processes that typically involve multiple decision-makers or departments, such as loan approval.

To support enterprise BI, IT must do its part to clean up the data integration layers and enterprise data warehouses in order to federate information silos into a shared resource. Decision-makers, automated process-management systems and Web services will benefit from a "single version of the truth" about customers, products and other objects of interest.

Metadata and master data-management systems, or those that identify and protect "gold copy" reference data, must link up with industry standards generally based on XML. A good example is the XBRL (Extensible Business Reporting Language), an emerging XML specification used to exchange financial information and backed by the U.S. Securities and Exchange Commission to formalize the reporting of quarterly and annual financial results.

XBRL adoption will require greater standardization of internal financial management. This, in turn, will depend on MDM and integration metadata to establish "single truths" among far-flung data sources in a typical global enterprise. These are not easy transitions for businesses that need to balance BI rigor with consideration of decision-makers who must protect their business creativity and innovation. If the tools are too hard to use, or don't live up to their promises, time doesn't wait—your users will make decisions without them.



If these commercial packages are out of your league, open-source BI is becoming a player, as are new data-warehouse appliances that promise analytical power at an affordable price. JasperSoft and Pentaho, two of the better-funded open-source start-ups, help develop BIRTs (BI and Reporting Tools) for Eclipse, the popular open-source development environment. This is no flash in the pan: BI veterans Actuate, Business Objects and Cognos are also BIRT developers.

Finally, BPM and rules-engine providers are in the hunt. BEA Systems with Fuego, as well as Fair Isaac, Savvion, TIBCO and WebMethods, all boast of activity monitoring, embedded analytics, decision management and business information integration in their products. These technologies will be critical to next-generation BI.

LET THE PEOPLE BI

Although BI tools are becoming somewhat easier to use, they remain a challenge for most nontechnical employees, especially if the suites require deep knowledge of underlying data models, schema and metadata. It takes

some training before an end user will know how to drill down and discover why, for example, a BI dashboard is showing that labor costs are high in comparison with sales in a particular store.

If BI is to enable quick, daily business decisions on issues such as allocating resources and inventory, users shouldn't be forced to waste time sifting through irrelevant data for answers. They certainly can't depend on IT to do their drilling.

BI and analytic application vendors will take the user's role and responsibilities into account. Business Objects, Microsoft, Oracle and others will further link BI into project lifecycle management systems to integrate a much more guided experience with established project methodologies.

What these vendors often overlook, though, is that executives and managers have always had methods of gathering internal information and applying it to their decisions. Maybe they had IT develop a proprietary application that paves the way by automating existing measures, or they might become spreadsheet jockeys them-

BI VENDORS TO WATCH

Enterprise BI platforms consolidate what was once a vast array of disparate tools for data access, analysis and reporting with data integration infrastructure technology, including administration, BI modeling, metadata management, portal integration and security. Operational BI pressures are pushing leading vendors to expand their technologies to support performance and process management; activity monitoring; and faster, more real-time information delivery. Vendors also are trying to keep pace with changing Web 2.0 trends in the user community that will favor BI services inside component mash-up applications, search and collaborative networks. Here's a snapshot of some key vendors to watch:

Business Objects	Regarded as the BI market leader, Business Objects has bought its way into the platform battle for performance management. The XI release gives Business Objects the benchmark enterprise BI suite, but its R&D is focused on making BI "ambient" through integration with search and collaborative, Web 2.0 technology.
Cartesis	As the big vendors sweep through, Cartesis is one of the more attractive pure-play performance management vendors left standing. Financial consolidation, planning and XBRL-based reporting put Cartesis on many shortlists—and in the running as a tempting acquisition target.
Cognos	A major enterprise BI platform player with Cognos 8, the company is expanding toward real-time, process-oriented BI with its acquisition of Celequest. A specialist in BAM, Celequest was developing a BI appliance. Mature performance-management tools also differentiate Cognos.
Endeca Technologies	When people don't know exactly what they're looking for, they search. Offering more focus than a typical Google search, Endeca takes a model-driven approach, bringing search into the BI realm.
Information Builders	This veteran BI provider offers WebFocus, a major operational reporting platform. IBI is a quiet leader in integrating search and BI. Visual workflow modeling and partnerships with specialist software providers give IBI a lift in process-oriented analytics and performance management.
Microsoft	With its market presence growing and its technology maturing, Microsoft is the biggest change agent in the BI industry. Its suite knits together services resident in SQL Server, Office and SharePoint for collaborative information delivery, reporting and analysis. Watch for the midyear introduction of PerformancePoint.
MicroStrategy	While now something of a niche player, MicroStrategy remains an innovator. The next release brings a richer user experience to the standard dashboard.
Oracle	Oracle's agreement to acquire Hyperion signals a major focus on performance management and financial analytic applications. Overall process-management strength gives Oracle the opportunity to lead in BI-process integration.
Pentaho	By packaging projects it sponsors in the open-source community, well-funded Pentaho pulled ahead of the pack in creating what proprietary vendors fear most: an open-source, enterprise-class BI and analytics platform.
QlikTech	With its graphical, in-memory approach to BI, the innovative QlikTech is pointing the way toward more powerful analytics riding on 64-bit platforms—without the usual throughput performance sacrifice.
SAP	While not quite enterprise BI, NetWeaver BI and BW are critical expressions of the company's unified, SOA foundation. BI Accelerator puts SAP squarely in the BI and data warehousing appliance markets. SAP's performance management is a work in progress, however.
SAS	Known for its deep analytics with data-mining prowess, SAS has made strides in establishing a respected enterprise BI platform. The company offers analytical depth in performance management as well.
Spotfire	Guided, in-memory analytics anchored in vertical industry depth will be important for process-oriented BI. Spotfire has built a leading-edge customer base focused on business objectives, not IT priorities.

selves. Some simply rely on instinct. Either way, old habits die hard. BI flourishes where executives, managers and business analysts are willing to challenge conventional wisdom, but not all organizations are that adventurous.

From a strategic standpoint, BI vendors are discovering that the best way to overcome entrenched practices and increase their presence is to infiltrate the application, service and process environment that operational workers encounter every day. In this capacity, BI and performance-management dashboards serve as components of corporate portals or other interfaces. Eighteen months ago we reported that Microsoft Excel was an integral part of the enterprise business intelligence strategy—and with good reason. Excel is flexible, provides both numeric and visual representations of data, and can be scripted to provide bidirectional communication that enables write-backs for forecasting or corrections to existing data (see “One Suite To Serve them All,” at nwc.com/go/O4O2bi).

Some BI vendors are still battling to pry users away from their spreadsheets, but most—including Microsoft—have accepted defeat gracefully and are working to bring BI’s benefits to those who prefer spreadsheets as their main tool and entry point. For more on Excel, see “Surrender the Spreadsheet,” page 60.

LAWS OF THE LAND

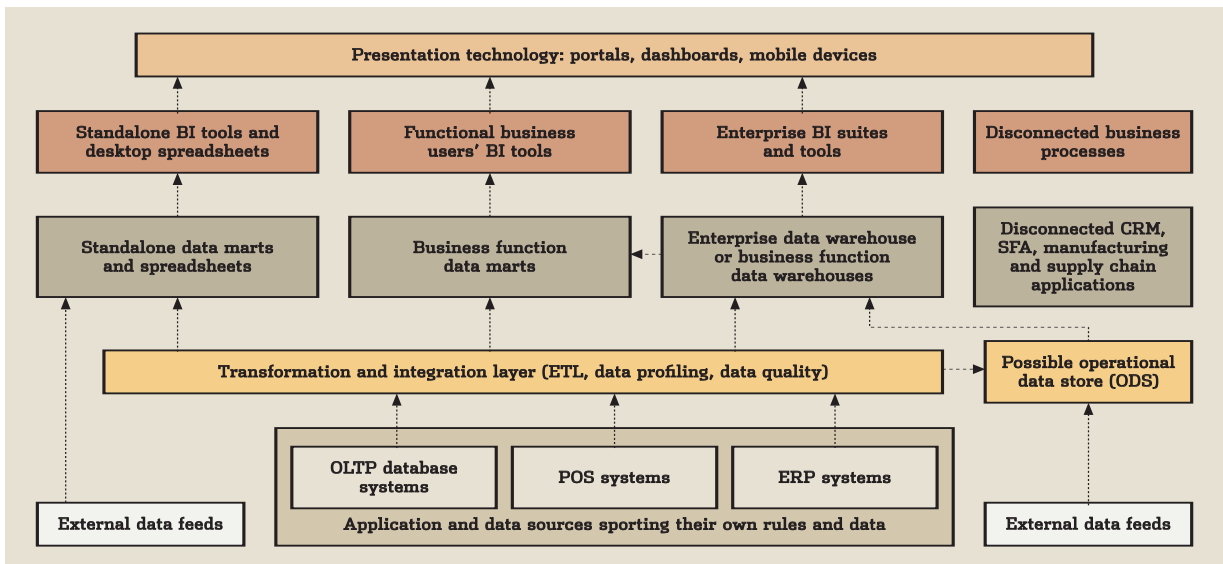
Inevitably, BI-generated aggregations and other intermediate analytical data stores house sensitive or proprietary information, meaning any BI implementation must address security and privacy concerns. Most BI vendors pay heed to database and network security best

practices; however, newer systems are getting tougher, using guided analytics and role-based access to protect data sources. Master data management (sharing a master data set among disparate IT systems and groups) and other infrastructure-management capabilities are rapidly becoming enterprise BI product differentiators.

Of course, the hammer that will really pound vendors into shape is regulatory compliance. The Sarbanes-Oxley Act, for example, outlines “need-to-know” access to data. For compliance with Basel II, HIPAA and other regulations, companies need better audit trails, better data quality and improved security. They also must be able to track what information and analysis contributed to a business decision. BI and data warehousing can help improve accountability and stewardship of information resources. However, BI and analytics activity will put database and application managers under added pressure to safeguard data sources from suspicious queries, aggregations, loading and other activity.

Regulatory compliance is pushing some BI vendors specializing in performance management, such as Cognos and Hyperion, to expand their roles beyond passive, strategic business analysis. Through partnering with search technology providers, BI’s reach is broadening into unstructured data found in Office documents, e-mail, voicemail and other assets. Rather than focusing on discrete transactions, BI must be able to comprehend events amid a sea of data to address fraud detection, security surveillance, and new business activities such as RFID tracking and algorithmic trading in financial services. Previously, BI has kept its focus on the past,

THE CONVENTIONAL BI STACK



More like a kludge! The primary orientation is toward the report: that is, gaining a comprehensive view of historical data for further query and analysis. Although enterprise BI has made inroads, standalone systems are still prevalent. This **creates a chaotic system** that has trouble producing a “single view of the truth” about customers, products and other objects of interest.

primarily on relational, transactional data. With help from search, XML and event processing technology, the blinders are starting to lift, and none too soon.

THE X FACTOR

XML is alive and well, and popular for content management as well as providing the basis for industry standards such as XBRL (Extensible Business Reporting Language). This is presenting data-management challenges that ultimately will impact the BI establishment. XML data comes into systems and is shared primarily through documents and forms. Should the data be “shredded” into relational rows and columns? Or is it necessary to retain the document context?

IBM, Microsoft and Oracle, and new vendors such as Ipedo, Ixiasoft and Mark Logic are marketing “pure” XML database technology that retains context. These vendors, plus middleware providers such as Progress Software’s DataDirect, support the emerging XQuery standard, which will serve a role similar to that of SQL. This will allow querying, manipulation and processing of XML data as easily as an Oracle or SQL Server database is queried today.

Most BI tools arrived on the scene to replace manual SQL coding. Query engines will have to upgrade to accommodate XQuery and XML, or use partner technologies from DataDirect or other middleware providers to work with XML-based sources. Data warehouses that bring in XML data and content also will have to adjust to new usage patterns that will differ from what designers anticipated with traditionally structured data. XML database technology, whether integrated with the big vendors’ relational systems or

offered by pure-play providers, will be important to broadening BI into the XML realm.

TIME IS THE ESSENCE

When people first started talking about real-time BI, skeptics didn’t hesitate to chime in. Do sales managers, or even most executives, really need up-to-the-second numbers? Or are vendors just trying to sell us more expensive stuff?

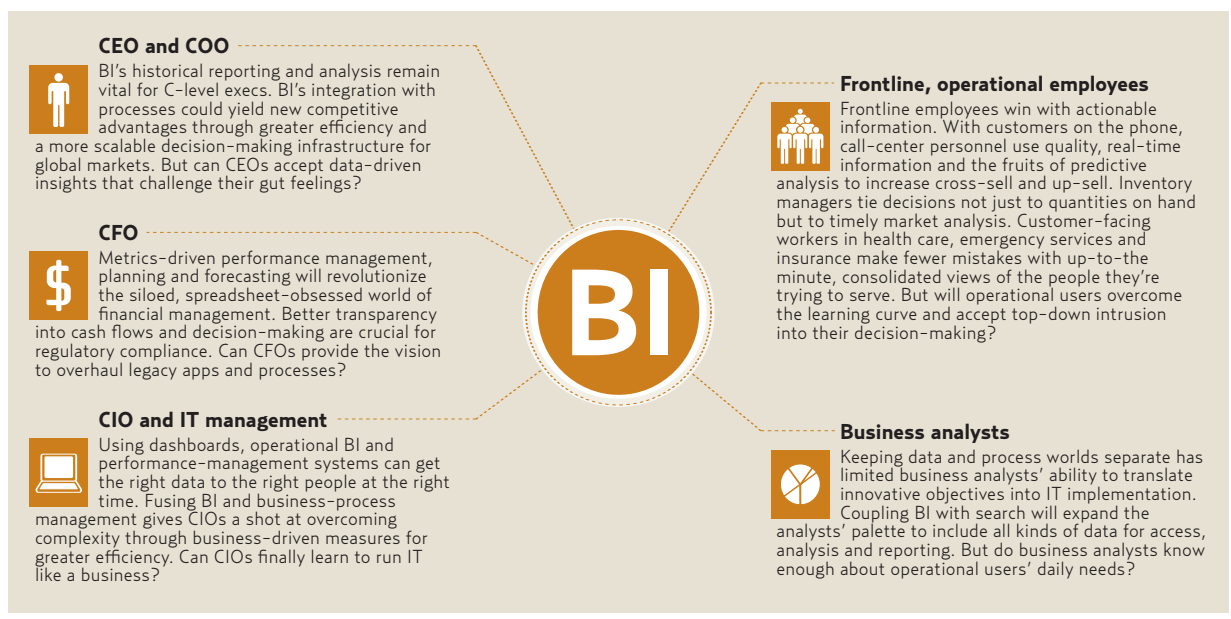
BI applications in a variety of areas, such as pricing, manufacturing, yield management, fraud detection, supply chain planning and call-center management, can help decrease the latency between data capture and when it’s available for analysis. That can only be good for business.

This demand for real-time processing hit the data warehousing community where it hurt. Call it collateral damage: As data volumes have grown, so has the difficulty involved in loading massive, sometimes multiterabyte, extracts into enterprise data warehouses. Most systems need to go offline for periods of time before the updated state is finished and BI tools can have at the data.

Moreover, data warehouses have been gradually moving from monthly to weekly update schedules. More than one daily update is currently considered “right-time” in many shops—that is, close enough to real-time. The result? As BI plays a more operational role, downtime leaves businesses vulnerable.

To speed up the process, companies are turning to active data-warehouse approaches that include trickle feeding the system smaller updates rather than updating in one big lump. Others are changing their architectures to include middle-tier servers that function much like an application or transaction server to manage data feeds and play query traffic cop.

WHO REAPS NEXT-GEN BI’S BENEFITS?



Enterprise data warehouse specialist Teradata, now in the process of splitting off from parent company NCR, cites customers such as Continental Airlines and Harrah's Entertainment that use active concepts with centralized warehouses to drive real-time customer intelligence, letting them make offers or knowledgeably respond to customers' behavior when they're on the line or playing at one of its gaming establishments.

Cost pressures are also driving interest in data-warehousing appliances from DATAlegro, HP, Netezza and Sun Microsystems' Greenplum. Teradata and Oracle, too, are close to marketing data-warehousing appliances—generally hardware, software and (in Netezza's case) storage bundles that commoditize massively parallel, shared-nothing computing. Shared-nothing computing is a distributed computing architecture whereby data and processing are located on multiple machines with no centrally located data store or processor. This leads to easily scalable systems that have, in theory, no single point of failure.

Gigabit Ethernet, InfiniBand and proprietary interconnects let nodes communicate and distribute the work. The appliances also feature chip-level optimization and specialized disk I/O—long BI and data warehousing's chief performance bottleneck—so the machines can focus solely on the kind of I/O associated with data warehousing and analytics. Feel the wind in your spreadsheets.

Parallel, shared-nothing architectures have forever been considered more scalable and faster for analytics

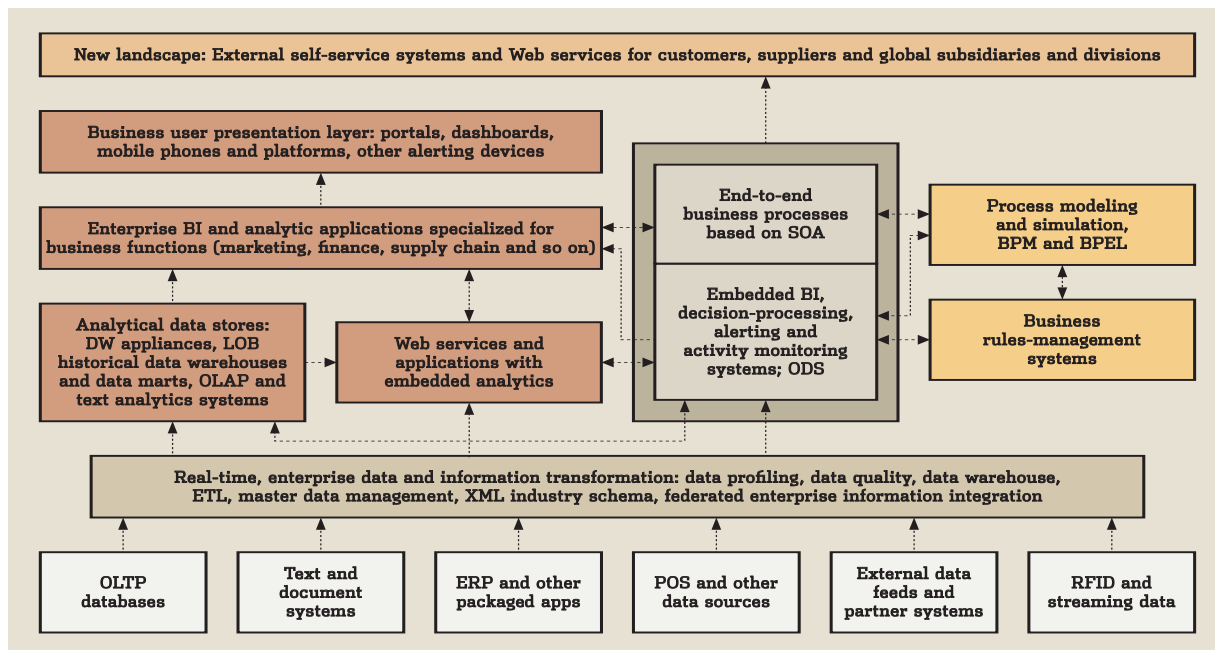
and data warehousing, but only Teradata has had true success with it in the market—and Teradata is expensive. Data-warehousing appliances are bringing the architecture into the mainstream, putting price pressure on Teradata and offering an alternative to the standard, shared memory or disk approaches offered by Oracle and Microsoft. IBM stakes out the middle ground by offering some database systems with parallel, shared-nothing architectures.

As packages, these systems avoid saddling administrators with tuning the entire stack to meet time-sensitive pressures and ever-growing data quantities. However, there are downsides. The appliances are customized for a specific purpose; they are not meant for mixed workloads, such as OLTP plus complex data querying, or even analytical workloads that require specific tuning to meet unusual objectives.

This could be a showstopper if "unusual" means an innovation that adds up to a major competitive business advantage. However, they may be just right for high-performance data marts—that is, data warehouses or analytic engines built specifically for one department or purpose, such as tracking online customer shopping. Attracted by their potential affordability and easier maintenance, midmarket firms might also jump on the data-warehousing appliance bandwagon.

Just as data-warehousing appliances could clean up the gathering and preparing of data, BI appliances hold promise in simplifying the "downstream" OLAP

THE NEW INTEGRATED, PROCESS-DRIVEN BI



The BI stack provides actionable information for continuous improvement of end-to-end business processes built on SOA. Rules and process-management systems work intimately with BI and data transformation systems.

for use in exception reporting, alerting and other operational BI chores.

Celequest, acquired in January by Cognos, made a name for itself by offering tools for BAM, which focuses on detecting patterns in processes, events and other activities. In late 2006, the company began to market an in-memory, behind-the-firewall appliance for operational BI and performance management, with SAP installations as a target market. It will be interesting to watch how Cognos, already prominent in performance management, plays its hand with the Celequest technology—and how competitors, including application providers Oracle and SAP, will respond.

Whether appliances or not, BI tools and analytic applications are racing to exploit 64-bit processing platforms that can support big, in-memory and data-intensive applications. These platforms will let users rely less on IT and perform more extensive and complex queries without having to retrieve data from disk. Appian and QlikTech are the most prominent specialized vendors in this realm, while SAS, long known for its deep analytics, leads the pack of established BI platform providers.

PARADISE BY THE DASHBOARD

If empowering business users to measure, monitor and manage business performance is the end game for BI, then performance management is an area where vendors must excel.

Not surprisingly then, performance management dashboard development has been the rage in BI, in some cases letting organizations dust off shelfware and finally apply the tools effectively. A dashboard, frequently in tandem with business-scorecard applications, presents information in a visually powerful way, often with gauges, metrics and graphs. The interface draws on more comprehensive performance-management engines that link into underlying applications, BI systems, data warehouses and other sources. Performance management has established a foothold in the CFO's office, where it aids in planning, budgeting, forecasting and consolidation.

All major BI and ERP players, plus specialists such as OutlookSoft and Spotfire, are chasing this market. However, none has achieved breakout success in making performance management essential to operational activities beyond finance. Some BI vendors are focusing on making performance management part of robust analytical applications specialized for certain industries or functions. However, as Cognos' Celequest acquisition shows, BI vendors realize it will take more than what they have in their traditional technology bases to succeed. Operational dashboards must tap timely information drawn directly from processes and event streams to be fully relevant. And for that, we need BAM.

UP A REAL-TIME NOTCH

BPM (business process management) vendors such as BEA with Fuego, IBM with Filenet, MetaStorm, Oracle and Savvion have long provided BAM functionality as a way of enabling continuous improvement of processes, often to meet Six Sigma or other best-practices requirements. As automated processes proliferate, it takes BAM and BI functionality to help companies stay ahead of the complexity they're busily creating and be able to monitor and measure whether their processes are fulfilling strategic business objectives.

Many organizations would love to bring together real-time BAM capabilities with their enterprise BI platform's facilities for long-range planning and analysis. It's the best of both worlds: Processes could be optimized continuously, with input coming from throughout the enterprise through operational BI systems. Linked more tightly with processes, conventional BI analysis and reporting won't lag so far behind—a situation that prevents companies from recognizing trends and patterns indicating problems in customer processes or supply chains until it's too late. Through dashboards and mobile devices, alerts and BAM notification systems can enable key personnel and systems to take action.

With BI and BPM more closely aligned, organizations can apply predictive analytics to anticipate problems, even use rules engines from the likes of Fair Isaac, ILOG and Pegasystems to automate remedial decision-making. There's also room at this intersection for innovative start-ups. In February, Altosoft launched Insight, which uses adapters to pull data from processes and other sources, perform predictive analysis and then use simulations to show the potential effect of changes to processes.

This BI-process synergy fits with larger organizational objectives to drive BI implementation through business models. In other words, rather than start with the data requirements, a model-driven approach looks first at what needs to be accomplished in sales, finance, product design or other business areas. Business modeling requires a richer metadata and semantic integration layer than most data warehouses currently have. However, with the adoption of XML-based schemas, business process modeling languages, such as BPEL, and service-oriented architectures could accelerate the evolution of information integration to support model-driven BI. ■



DAVID STODDER, /// now VP and research director for information management at Ventana Research, was previously editorial director and editor in chief of INTELLIGENT ENTERPRISE, a CMP Technology publication. He has been covering BI and related IT topics for nearly 20 years. Write to him at dstodder@msn.com. Post a comment or question on this story at nwc.com/go/ask.html.