InformationWeek Research & Reports

intelligent enterprise
www.intelligententerprise.com

October 17, 2008

Data Strategy

Scaling The Data Warehouse

Data warehouses aren't just exploding in size, they're also supporting more users and increasingly complex queries, all in shorter time frames. Here's how to make sure yours is ready to scale.

By Richard Winter

LGR TELECOMMUNICATIONS has a 310-TB Oracle data warehouse that's used daily by 2,500 people at one of its telecom carrier clients. The warehouse powers an LGR service, called CDRlive, that gives its carrier customers access to call data records. It's updated round the clock, in near-real time, and is available for query 24 hours a day, 365 days a year.

"There are no batch jobs," says Hannes van Rooyen, chief architect at LGR, which supplies data warehouse software and services to the telecom industry. "Instead, as many as 13 billion records a day are added, and an equal number are dropped in an online update process that runs concurrently with user queries."

The data warehouse keeps more than a petabyte of disks spinning and has grown by a factor of 10 during the last four years. It's expected to at least double in the coming year.

Most companies still don't hold hundreds of terabytes of data, but they're up against the same data warehouse problems that face LGR—soaring data volume, more users, complicated queries, and fast-changing information. Throw in a growing number of vendor options and it's time for companies to re-evaluate their data warehouse strategies.

The new generation of data warehouses looks a lot like LGR's: growing at an extraordinary pace, in multiple dimensions, and supporting critical business processes that must react quickly to

ALSO IN THIS REPORT

events around the company. Whether your company has 250 GB or 250 TB of data, you're likely facing the same questions: Do we have the right architecture? Is it on the right platform? Is the warehouse about to run out of headroom? What will it take to service new users? How do we move from batch loading to continuous update? And with technology changing so rapidly, how do we know we're on the right system?

All the answers loop back to managing scalability. Getting control of scalability might mean embracing the highly parallel processing and scale-out architectures long offered by Teradata and IBM and elements of which are now emerging in new products from Oracle and Microsoft (see story, p. 3). Or it might just require more effective management of existing data warehouse practices, including quantifying requirements, measuring alternative solutions, and acting earlier on potential problems.

MULTIPLE DIMENSIONS OF SCALABILITY

The convergence of three key trends is driving the ever-expanding scalability challenges facing data warehouse managers. The first is well known: Data volumes are increasing rapidly. The largest data warehouses are tripling every two years, according to WinterCorp surveys.

That's about how fast LGR's data warehouse is growing; it will approach 3 PB in 2012. Hundreds of other data warehouses, including those run by retail, health care, and financial services companies, also will reach petabyte scale in the next few years and thousands will surpass 100 TB. In many cases, competitive pressures are driving businesses to capture more data in hopes they can better analyze, understand, acquire, and retain the most valuable customers.

Data warehouses also are getting more time sensitive. The extraordinary velocity of the data in LGR's warehouse is a case in point: Billions of records pour in throughout the day, loaded into the database within minutes, and acted on almost immediately. If a mobile phone customer calls in because he had a bad experience, "we want to see exactly what happened, what calls were dropped, what tower was involved, and so forth, while they're still on the phone," van Rooyen says. "At the same time, you want the customer service person to know the customer's history." Problems are resolved faster, customers get better service, and "the business works better all around," he says.

High-velocity use of data—also called "operational business intelligence"—isn't a new concept. Teradata identified it several years ago as "tactical data warehousing," and IBM's Dynamic Data Warehousing pushes a similar notion of "right time" data. But the business pressure to provide such capability is rising.

Tactical data warehousing facilitates the moment-by-moment decisions employees must make. Many of these decisions are similar and repetitive: What should I offer this customer?

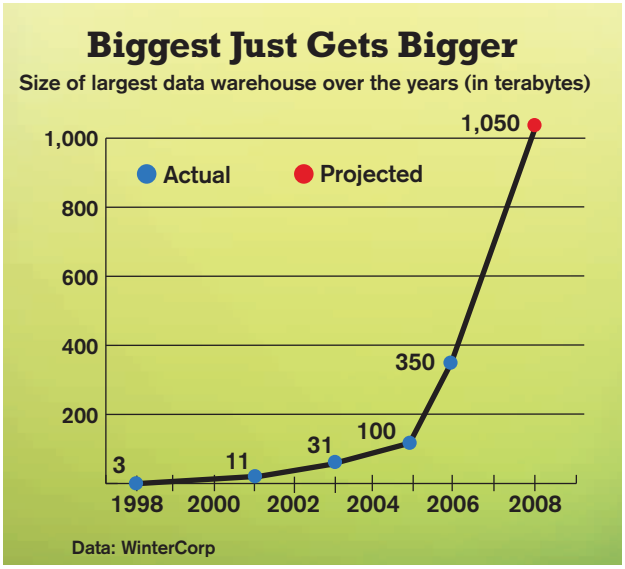

How do I treat this unexpected shipment that just turned up? Businesses that can make such decisions in a systematic way, informed by up-to-the-minute data, find they produce significantly better results.

Operational BI has big implications for data warehouse scalability. It results in larger user populations; more frequent, time-sensitive interactions; a need for fresher data; and support of business processes that can't tolerate downtime.

The third trend is rising complexity in data, queries, workloads, and analysis, all of which amplifies scale. When data warehouses are doing only simple things, such as predictable updating and straightforward reporting, they can grow without creating fundamentally new problems. But when they have to respond interactively to complex and unpredictable queries—perhaps performing large, complex joins, aggregations, sorts, and calculations on trillions of records—the requirements have truly escalated.

Many modern data warehouses perform complex queries, analyses, and reports. They also operate on more complex schemas than in the past, with thousands of tables, hundreds of thousands of columns, and a complex web of data relationships.

EXTREME MULTIDIMENSIONAL GROWTH

There are few better illustrations of the multidimensional growth phenomenon than eBay. About 85% of the queries run on the company's

Microsoft, Oracle Join The Scale-Out Crowd

ORACLE AND MICROsoft, in a bid to land midsize data warehouse customers, are pitching new products aimed at a "scale-out" option—running large data warehouses on clusters of small, low-cost servers.

By bringing products for highly parallel architectures to midmarket users, two of the largest database vendors are acknowledging there are multiple dimensions to scalability. Companies can have as little as a terabyte of data but use complex queries or schemas, or have lots of people accessing the data. Such users often find they need a scale-out architecture.

Oracle last month rolled out the HP Oracle Exadata Storage Server and the HP Oracle Database Machine, both designed to raise performance for data warehouse queries. Oracle's products use the Exadata storage cell as a building block, relying on low-cost Hewlett-Packard hardware and intelligent Oracle software to off-load database processing to the storage tier and increase disk I/O bandwidth. The performance version of an Exadata storage cell will store 1 TB of user data and deliver 1 GBps of raw I/O bandwidth.

The effective bandwidth in processing a query can actually be much greater than 1 GBps per cell because of compression and database operations, such as filtering and projections, performed within the storage cell. This lets Oracle data warehouses offer significantly higher performance, while requiring less space, power, and cooling; they also cost less compared with conventional storage arrays.

MICROSOFT BUYS IN

At Microsoft's Business Intelligence Conference last week, the company said it will integrate the technology it acquired as part of its purchase of data warehouse appliance vendor DATAllegro with Microsoft SQL Server. The first products are expected in 2010.

Before this move, Microsoft focused on growing data warehouses via scaling up; customers would buy larger SMP servers when they needed a bigger warehouse. This approach has advantages in operational simplicity, but it imposes a ceiling on capacity. Microsoft still touts the scale-up option, but the DATAllegro technology adds a scale-out option.

The scale-out approach isn't new to large-scale data warehousing. Teradata has used it since 1984, IBM since the mid-'90s, and Oracle for nearly 10 years with RAC and now grid computing. HP Neoview and many data warehouse appliance startups emerging this decade are using it, too. In addition to reducing hardware costs, a good scale-out architecture promises modular capacity and potentially little or no disruption for upgrades.

With these latest announcements, Oracle and Microsoft are aiming to capture larger data warehouse deployments. Of course, like all highly parallel architectures, theirs will have limitations and bottlenecks. And they'll have to prove they're as good as those from vendors that started out with a highly parallel approach. **—RICHARD WINTER**

data warehouse are "exploratory in nature," says Oliver Ratzesberger, eBay's senior director of architecture and operations. They come from end users, with no opportunity for a database administrator to apply a tuning tool to them. "The queries hit the engine, and it has to handle them," Ratzesberger says.

EBay's data warehouse contains about 5 PB of disk storage distributed over primary and secondary systems, both running Teradata. The secondary system for disaster recovery is located about 1,000 miles from the primary one. Each system has a complete copy of the company's core data, organized as an enterprise data warehouse. Both copies are updated every 15 minutes, round the clock, and are continuously active servicing queries.

There are more than 5,000 users and about 10 million queries each day. The daily update volume ranges from 10 billion to 15 billion records per day. Thousands of tables are involved, and queries range from simple lookups to complex analyses that run for hours. The system is constantly managing a mixed workload with different service-level objectives for each of the various classes of work.

Given the scale of the system, the growth rates are even more remarkable: The number of users grew 25% last year, the number of queries doubled, and the size of the system has at least doubled each of the last four years.

EBay's experience shows how data warehouses don't just grow in quantity of stored data. They also expand in several dimensions at once, including data volume, number of users, query volume, data latency, and data and query complexity. Decisions on architecture and spending must take into account the likely growth of all these dimensions.

FIVE-STEP PROGRAM

To be clear, don't try to preach "multiple dimensions of growth" to business unit managers. They see scalability as simply the ability to buy

EBay Turns Data Marts Into A Service

EBAY HAS INSTITUTED a utility computing model to better manage the growth of its data warehouse that Oliver Ratzesberger, senior director of architecture and operations, refers to as "analytics as a service."

This service lets authorized eBay employees access a virtual slice of the main data warehouse server where they can store and analyze their own data sets—either in isolation or in combination with core data in the enterprise data warehouse. EBay's virtual private data marts have been quite successful—hundreds have been created, with 50 to 100 in operation at any one time. They've eliminated the company's need for new physical data marts that cost an estimated $1 million apiece and require the full-time attention of several skilled employees to provision.

Virtual marts are often used only for a few days or weeks, so system resources are quickly reclaimed. Users typically introduce less than a terabyte of new data, which they often want to analyze in conjunction with the data in the enterprise data warehouse. If these projects were implemented as separate physical data marts, the required core data would probably be extracted to the data mart, swelling its size, requiring a way to keep replicated data up to date, and multiplying cost and complexity in other ways.

EBay's analytics as a service is a way for people to do "agile prototyping," Ratzesberger says. "They can do experiments quickly and succeed or fail quickly and inexpensively." This helps the company move faster to find and exploit opportunities in connection with Web site optimization, fraud detection, and revenue generation.

When an analytic environment is needed for more than 90 days, the data warehouse team explores whether the user's data ought to be incorporated into the enterprise model.

—RICHARD WINTER

systems—including data warehouses—without unusual worries about growth. They expect data warehouse growth won't cause a disproportionate increase in costs, unreasonable disruptions in business activities, or big hits to performance. Oh, and never run out of headroom.

Sound daunting? Here's a five-step approach to deal with extraordinary data warehouse growth and meet business expectations for scalability:

1. Develop quantitative requirements. Use a systematic, measurement-based engineering process to document quantitative requirements. They should include working estimates of the size and macro structure of the database and workload, service-level objectives, and operating schedule. These key inputs provide much of the information required to develop a physical database and evaluate alternatives.

The database's macro structure covers the likely size and structure of the largest tables, the likely set of the most heavily used relationships, and the likely distribution of data values of the most significant columns. The macro structure of the workload covers the 10 to 25 query or transaction types expected to account

Don't Do This: 7 Gotchas Of Scalability

1. Wait until the system is built to test for scalability There's always a temptation to wait too long before doing performance and scalability testing. The classic trap is waiting until the system is ready to go into production. Sure, the test is realistic because you can run the actual database and application, but if you discover something wrong, it's often too late to do anything about it. Test for scalability before you're committed.

2. Live with vague expectations Database people think that if no requirements are established, no one can prove they failed. In reality, it's often worse in this situation; management assumes that the system will meet all expectations, so the system is never good enough. You're much better off setting realistic expectations that can be met.

3. Skip requirements Users often don't know what the requirements are. You have to help them visualize a new business process and the requirements for supporting it. Only then can you develop valid usage scenarios and engineering requirements.

If, for example, you currently mail a giant catalog to all customers quarterly, and you want instead to do 100 targeted mailings of specialty catalogs each going to about 2% of your customers, then hold two facilitated discussions with stakeholders. First, talk about what the new mailing process will be and how it will get carried out 25 times as often each year. Second, explore the information capabilities needed to support the process. Then work out usage scenarios and develop the necessary workloads, service levels, and other requirements. Don't skip identifying requirements, or you'll end up back at pitfall No. 2.

4. Skip risk analysis Once you develop requirements, identify, test, and manage the risks that emerge.

5. Accept flimsy "proofs" Beware of salespeople taking over the definition of the proof. Never let the vendor define the test to be performed. If you don't have the expertise in-house, get a consultant with experience defining benchmark specifications for testing complex data management systems. Your test has to capture the key challenges of scale and performance.

6. Underestimate growth rates Knowing this year's requirements isn't enough. Architectural and platform decisions will take awhile to implement and longer to change. Project requirements out two to three years, at least—better to have a projection that gets revised than shoot in the dark. And don't assume that the data and business growth rates are the same. Data and workloads tend to grow faster than the related business because data gets used more intensively as the business gains momentum.

7. Ignore any dimension of scalability Data size is the dimension easiest to measure, but the workload, data complexity, query complexity, availability, and data latency dimensions are nearly as important. They can all drive configuration size and determine whether you're on the right platform. Take them all into account. —RICHARD WINTER

for the main performance challenges and their expected frequency.

When coming up with these estimates, the key is to get them in the ballpark—absolute accuracy is far less important than scale. Getting the scale right lets you understand whether you're building a passenger car, an 18-wheeler, or a freight train. Don't settle on this too soon: Document a set of numbers, talk through the estimates with stakeholders, and then use them in the management process as well as architectural and engineering decision making.

2. Forecast long-term needs. Within a few years, your data warehouse could be several times larger than it is today. To estimate long-term requirements, consider factors such as new applications, new or expanded subject areas, additional levels of data detail, as well as new users, tools, and data sources. The engineering requirements should define how a system will grow along each of the dimensions of scalability.

How To Manage Scale

- **DEVELOP** quantitative requirements
- **FORECAST** long-term needs
- **IDENTIFY** critical risks
- **MEASURE** potential solutions against the requirements and risks
- **MANAGE** areas where the data warehouse doesn't meet requirements

Don't just extrapolate existing growth rates, since they don't reflect changes in technologies and practices that might support major new opportunities. In retail, data scale increased dramatically first when point-of-sale and then when Web clickstream data were added to the data warehouse. In the supply chain, the next big leap in scale will come if there's full deployment of RFID. Extrapolating from past trends might grossly understate the impact of future needs.

3. Identify the critical risks. The process of documenting requirements—with vendors, user groups, reference companies, consultants—should raise the big risks: "We lose money if we can't load that data in time," or "We're dead ducks if we go down on any of the big weekends."

You'll spot some yourself—like the engineering problems no one has a convincing solution for, or the recurring queries everyone knows are complex and time sensitive, but no one can say how much time they'll take.

But not all engineering requirements are equally important; focus on the ones critical to business objectives. In a fraud-detection application, it may be critical to get the data into the database within minutes or seconds of receipt—no matter what the circumstances. That may be fairly simple to do except during peak hours, yet those are the exact times it's most critical to spot fraud because a lot of money is being spent. So ingesting data quickly during peak hours becomes a critical factor. In other areas, response time may be important, such as with customer-facing queries. If a moderately complex query happens while a customer is talking with a call center agent, and thus has a desired 2-second window to complete, it could become a risk.

It may be easy to show that the requirements will be met initially, when data volumes are small and usage is light, but what happens in the second year, when volumes skyrocket? The trick is focusing on engineering with two characteristics: There's no proof that the targets can be met, and missing the targets causes major pain to the business. These are the critical risks.

4. Measure solutions against targets. This step is key: Measure the solutions to a critical risk against the requirements today and as a company reaches the projections developed in the second step.

For this step to work, be realistic about scale and complexity. Don't cheat on the dimensions of scalability by, say, running a few simple queries, one at a time, on a 5% sample of the data. Instead, run a realistic simulation of the workload against a realistic, full-scale database, and take into account how the application is likely to evolve over the next three years.

5. Manage the gaps. Realistic analysis and testing often reveal that the intended data warehouse won't meet all the requirements. If so, address the issue with stakeholders before it becomes a problem. By measuring the alternatives, you can enter the discussion with real data on the options. Can users accept the 4-second response that is feasible under the current budget? Or would they increase the budget by 50% to get a 2-second response? Should we stay with the company standard platform, which has never been used with more than 10 TB of data, or take 90 days to evaluate other options, now that we understand we're likely to have 100 TB of data within 18 months?

A systematic engineering approach puts you in control, providing options with known outcomes and trade-offs as data warehouse requirements increase rapidly in six formidable dimensions. Where you have higher risks, you have analyzed, measured, and set up fallback plans. You can discuss the trade-offs and options with stakeholders and prepare them for the likely outcomes. This is a much better approach than the oft-used "forge ahead and hope" approach to managing data warehouse scalability (see story, p. 5).

ULTIMATE SCALABILITY

The new technology trend designed to deal with multidimensional data warehouse growth is toward highly parallel architectures. The HP Oracle Exadata Storage Server, announced last month, is designed to keep data flowing to and from more disks at once, increasing the pace at which I/O-intensive tasks can be performed. And Microsoft has just revealed that it will incorporate the DATAllegro technology acquired earlier this year into the next release of SQL Server, thereby increasing both I/O bandwidth and processor parallelism. Almost everyone is moving to exploit lower-cost hardware. Though big symmetric multiprocessor servers aren't about to disappear, there's an ever greater emphasis on scale-out architectures.

In the 1990s, conventional wisdom had it that massively parallel processing would never be more than a niche architecture, used for extreme requirements at the margins. But MPP has become reliable, manageable, and affordable—and suddenly it seems that nearly everyone is hungry for scalability. So highly parallel architectures—whether you call them MPP, cluster, or something else—have become part of the mainstream.

A lot of data warehouse practitioners are struggling with the changes brought on by rising data warehouse scale and rapidly evolving architectures. The most important thing to remember is that business problems aren't solved by buying new hardware or introducing new architectures. They're solved by determining the requirements of a solution and then implementing systems that meet those requirements.

To do that, follow these three recommendations in any data warehouse development project: Introduce a systematic management process to

deal with the scalability problems. Avoid the seven gotchas of scalability management. Emphasize quantitative requirements and use measurements or tests at every stage of the development life cycle. With a systematic approach, you will meet business expectations and have a scalable data warehouse with long-term business value.

Richard Winter is the president and founder of WinterCorp, a consulting firm focused on large-scale data management. In addition to advising companies in industries including retail, health care, financial services, and distribution, WinterCorp provides consulting services to vendors including Hewlett-Packard, IBM, Microsoft, Netezza, Oracle, and Teradata.

Performance Upgrade

Teradata Rolls Out Peta-Scale Appliance

By Doug Henschen

NOT EVERY COMPANY NEEDS AN APPLIANCE that can scale up to 50 petabytes, but who would say no to database upgrades that claim to improve performance by as much as 30%? Promising something for everyone, a new peta-scale appliance and the Teradata 13 database upgrade were the headliners at this week's Partners Teradata User Group Meeting in Las Vegas. Teradata also demonstrated a prototype appliance of the future using solid-state disk (SSD) drives that promise both faster performance and lower power consumption.

Teradata's new Extreme Data Appliance 1550 is designed for high-data-scale applications that are characterized by focused queries, departmental scale, and not-so-time-sensitive querying. Examples include Web site clickstream analysis, multiyear regulatory compliance, manufacturing processing and testing, RFID-product movement, and cell phone network usage. Many of these apps were heretofore viewed as impractical, or they were relegated to server farms and flat-file processing. The 1550 combines 1-TB hard drives with Intel quad-core nodes and built-in data protection software to offer extreme storage density, starting at a list price of $16,500 per terabyte.

"This gives many people an option to get [vast sets of data] into a relational format and start doing analytics on it for discovery and huge data sifting," says Scott Gnau, Teradata's chief development officer. "It's for applications in which performance isn't extremely important, but ease of use and ease of integration will be important."