

*We've got people to get it done—nearly 6,000 impassioned employees who wake up every day and think about data warehousing and analytics. We have a great set of solutions and a tremendous roster of more than 850 customers. It's why we are the technology and market leader. Now that the new Teradata era has begun, I see a really bright future for all of us.*

---- Mike Koehler, CEO of Teradata, in Teradata Magazine, December, 2008.

## Teradata Reborn

### Introduction

On September 30, 2007, Teradata Corporation (NYSE: TDC) became an independent company - a spin-off from NCR Corporation (NYSE: NCR). That spin-off instantly became the largest organization with a business strategy focused solely on data warehousing and enterprise analytics. After just six short months of its independent existence, the company had total assets of \$1.3 billion, total shareholders' equity of \$645 million, no long term debt and strong liquidity. In addition, Teradata had more than 850 corporate customers and more than 2000 implementations in organizations including: 90% of the top 10 global telecommunications companies, 70% of the top 10 global airlines, 60% of the top 10 global transportation logistics companies, 50% of the top 10 global retailers and 50% of the top 10 global commercial and savings banks. The company's 2008 financial statements are included in Exhibit One. By the end of 2008, revenue from the company's services stood at \$913 million, and revenue from products was \$849 million. Investments in Research and Development (R&D) for 2007 and 2008 were \$126 million and \$108 million, respectively. On February 12, 2009, Reuters quoted CEO Mike Koehler as stating, "Teradata finished the year with a strong fourth quarter and, overall, had a successful and very productive first year as an independent company. In 2008, we extended our technology lead, launched a new innovative family of purpose-built data warehouse appliances and added sales territories to broaden our market coverage."<sup>1</sup> In that same article, the 2009 outlook was characterized as one of uncertainty given the global economic environment and its potential impact on Teradata's current and potential customers. However, it was stated that R&D funding was anticipated to increase above expenditures incurred in 2008.

---

*This case was prepared by Michael Goul, Professor, W. P. Carey School of Business at Arizona State University. Mary Gros, Teradata Director of Business Relations, provided substantial context for the case and arranged for numerous interviews with key executives, customers, business partners and other academics in related disciplines. This material is free for use for educational purposes. Copyright © 2009 by Michael Goul and the Arizona Board of Regents.*

Teradata's products and services provide organizations with the hardware, software, best practice methods, logical data models and professional services needed to capture, store and analyze data in order to gain insights and make informed decisions. Including the time spent while still with NCR, Teradata has been in the data warehouse business for more than 25 years.<sup>2</sup> The organization currently employs over 5900 associates in over 40 different countries. Company offerings support customer organizations in their efforts to conduct data intensive research and development in order to both discover new insights - and to deliver both strategically and operationally on those insights. Teradata's products and services support their customers' harnessing of data-intensive insights to make smarter, faster decisions in areas ranging from customer management to supply and demand chain management.

Industry analysts agree that the data warehousing market of the millennium's first decade is changing dramatically, and that an independent Teradata will soon be called upon to meet significant challenges. This view was best summarized by Ventana Research in a posting on October 14, 2008 just after the annual Teradata Partners Users Conference:

"For today, your data warehouse investment is under scrutiny to ensure you get the best bang for the buck. As the cost for memory, storage, and computing power decreases, the pressure on technology providers like Teradata to lower prices continues to increase. Can Teradata escape the perception that they only address the largest of data warehouses and are very expensive to manage? Can they once again grow their US market or have to depend on other geographies to carry them forward? Can they escape the rumors and potential opportunity for SAP or others to purchase them? There are new choices ahead for Teradata and much different than last year and higher stakes to their future and their envied position in the data warehouse market."<sup>3</sup>

### **Business Intelligence Marketspace**

Shortly after Teradata started its operations as an independent company in 2007, there was an explosion of mergers, acquisitions and new strategic alliances throughout the enterprise, database, and data warehousing marketspace. That space had quickly come to be known as 'Business Intelligence' and was commonly referred to by the acronym 'BI.' By November 14, 2007, the space had changed dramatically. IBM had purchased Cognos for \$5 billion, Oracle had purchased Hyperion for \$3 billion and SAP had acquired Business Objects for \$6.8 billion. Subsequent to the merger and acquisition frenzy, many analysts were scratching their heads wondering what was to become of the companies that weren't involved in all of the buzz surrounding the big enterprise players' acquisitions. On one hand, the mergers could usher in an era of stagnation as the complexities of managing multi-business alignments often require significant investments of time and energy. On the other hand, many analysts suspected that with some of the major players out of the game until the mergers could be completed, the smaller players would emphasize innovation at a break-neck pace. Following is an example commentary:

"Business intelligence is increasingly seen as core to enterprise strategy. To stay competitive, SAP, Oracle and IBM needed business intelligence offerings to complete their portfolio, but recognized that they would never catch up if they created those offerings organically. So they have extended through acquisition. Business Objects,

Hyperion and Cognos—all strong companies—allowed the larger software vendors to gain market presence and a large customer base quickly. The deals were natural extensions. The consolidation itself is simply the natural maturation of the business intelligence market. The ERP market underwent the same consolidation. Although the attention is placed on [the big BI players on the] upper end, there's a whole bunch of little guys who are really focusing on specific types of issues—analytics, open source and real-time data access. Innovation will continue..."<sup>4</sup>

Teradata adheres to what is called a two-sided platform strategy. The phrase ‘two-sided’ refers to the relationships Teradata manages: one with customers (who purchase and use Teradata products) and another with third-party companies called complementors. Scholars describe companies that successfully manage two-sided platform strategies as those where, “platform leaders create an industry ecosystem greater than the sum of its parts.”<sup>5</sup> In contrast to companies following a strategy that relies on purchasing or merging with complementors to provide better products and services to customers, Teradata must carefully manage its collaborations with each partner. Critical decisions are made in terms of what innovations to share and collaboratively improve upon (and with whom), what products and services to limit competition around, and they must decide how to handle situations such as when a former collaborating partner becomes a direct competitor.

Even while mergers and acquisitions made frequent headlines in media stories discussing the BI marketplace, another major market evolution was taking place – a shift to the offering of Data Warehouse (DW) appliances. A DW appliance is an:

“... integrated set of servers, storage, OS, DBMS and software specifically pre-installed and pre-optimized for data warehousing. Alternatively, the term is also used for similar software-only systems that purportedly are very easy to install on specific recommended hardware configurations. DW appliances provide solutions for the mid-to-large volume data warehouse market, offering low-cost performance most commonly on data volumes in the terabyte to petabyte range.”<sup>6</sup>

Companies like Greenplum, Kognitio, Vertica, EXOSOL, Netezza, Calpont, Dataupia, InfoBright and ParAccel entered the DW appliance market, and even Hewlett Packard launched the HP NeoView appliance and an Exadata appliance in partnership with Oracle. With increasing competition from the big enterprise players and new competition from DW appliance vendors, the pressure on Teradata to innovate to stay ahead of the competition mounted. Had the company anticipated this new competitive landscape by making earlier decisions and establishing important priorities around multiple streams of innovation?

### **Teradata’s Strategic Moves Just Prior to the NCR Spin-Off**

On September 18<sup>th</sup>, 2006 – about two years before the NCR spin-off of Teradata, a 64-bit Teradata parallel database version of the mainstay Teradata product for the SUSE Linux Enterprise Server 10 from Novell was announced. The original Teradata Relational Database Management System (RDBMS) of the late 1970’s and 1980’s was implemented on a proprietary operating system called TOS (Teradata Operating System). After the acquisition of Teradata/NCR by AT&T in the early 1990’s, the RDBMS

was ported to a version of the original open source UNIX source code (System V, Release 4, 'SVR4') and named as "MP-RAS." Later, NCR ported its Unix, a variant of AT&T's Unix System V, to X86 and X64 hardware that was running Microsoft Windows 2000 server or Windows 2003 – both operating in 32-bit mode. However, analysts reported that major Teradata customers were excited to reinforce the need for a move of Teradata to the 64-bit platform and Novell's SUSE Linux Enterprise Servers 10 (SLES 10). According to *The Linux Beacon*, speculation was that there might have been two customers involved; one runs what has been mentioned to be the largest data warehouse on the globe, and that company was likely Teradata's largest customer at the time.<sup>7</sup>

SUSE is an acronym that originally stood for, "Software und System Entwicklung" (Software and System Development), a German Unix dating back to 1992. It is pronounced "zoo-sah," but those who speak English often say "soo-say." There are SUSE Linux versions that are free, and others (including the Novell Enterprise version) that are licensed most often on a per-seat usage basis. Novell also has an open source version of SUSE Linux named openSUSE ([www.opensuse.org](http://www.opensuse.org)). As one has come to expect with open source software, this Linux version engages a large community that promotes the use of Linux everywhere. Project participants include those getting started with Linux, those who are reporting and addressing bugs, developing new packages and features, supporting cross-organization innovation, and all the while, community members are engaging and learning from each other. Teradata had a difficult choice to make. Should it invest in making the move to the SUSE Linux Enterprise platform or stay with its proprietary NCR version of Unix? If it did make the move, and there were few customers who would actually benefit, what was the real reason for this move?

At the time, Scott Gnau, Vice President and General Manager of Teradata Research and Development, announced his decision:

"With the addition of 64-bit SUSE Linux Enterprise to the list of operating systems, Teradata provides its data warehousing customers yet another choice. SUSE Linux Enterprise further enables the powerful Teradata enterprise data warehouse to integrate seamlessly with even the most complex and open IT architectures. "<sup>8</sup>

More recently, Gnau reflected on the decision as follows:

"Operating systems are not intended to be a core competency of Teradata, so we ported our products to Linux so that we could focus our research and development on what we do best – data warehousing. This was a decision that many people thought long and hard about. On one hand, we were making significant investments to move to Linux - and yet, we were maintaining the same functionality in our products. To some analysts, that seemed like a strange move. But now we are seeing the benefits. Both open source Linux experts and our Novell partners are at work innovating around the operating system infrastructure for Teradata products - and that enables Teradata to focus *our* innovation strategy on the areas where it matters most. We strive to invest our research and development expenditures in a balanced portfolio of open, proprietary, customer-driven, competition-driven and business partner-supported innovations. In addition, we strive to balance that portfolio on the dimensions of whether the potential of an innovation leads to incremental or architectural product enhancements - or even game-changing marketplace disruptions."<sup>9</sup>

Innovation has been at the very roots of Teradata's history. Teradata was born in a garage in Brentwood, CA in 1979 by Dave Hartke, Jerry Modes, Ira Moskatel, Phil Neches and Jack Shemer – all who shared the goal of building a massively parallel database system. In fact, Teradata actually introduced the first data warehousing appliance to the market in 1990, which, at the time, was referred to as a 'database machine.' The Teradata of the late 1970's and the 1980's was purchased (along with NCR) in December of 1991 by AT&T. Customer-driven innovations have long been a tradition at the company since the delivery of its first system to Wells Fargo Bank in 1983. For instance, the long-running annual Partners Users Conference, where Teradata users drive the agenda and actually decide on the program and the participants, is coordinated by a non-profit organization with a powerful steering committee composed of Teradata customers. In contrast to the other typical vendor conferences, the Partners Users Conference is not a marketing-driven event with a 'sales show' feel. Instead, Teradata is able to hear of problems and opportunities associated with its products and services during the conference, and they can identify specific innovation streams where there is the most potential to benefit the installed user base. Along with customer-driven innovation, Teradata has close partnerships with organizations including Microstrategy, SAS, and many others. Exhibit Two includes a list of the partners – each partner listing is a hyperlink where additional information about the partnership is available. Note that several competitors appear in this list – an interesting notion given the state of the marketplace. Exhibit Three provides a more comprehensive historic perspective to Teradata's evolution.

### **Innovating to Win**

Gartner Research has long touted Teradata's strength as a complete provider of data warehousing solutions because the company rounds out its "... physical warehouse products with strong consulting services and the widest variety of vertical solutions."<sup>10</sup> By 2004, Teradata had already secured a four year streak of being in the leadership position of Gartner's magic quadrant for Data Warehousing, and the recognition continues today. One reason is that the underlying technologies deployed by Teradata are noted for their regular ability to handle the largest organizations' data warehousing needs mainly due to the approach taken to use hardware nodes and virtual work units to process data needs in parallel. Another reason for its long term leadership has been a patented approach to 'logical data models' (LDMs) and a commitment to building high quality professional services staff capable of deploying those models to provide value to customers. LDMs for vertical solutions and professional services are both major elements of the factors emphasized by Gartner's evaluation.

#### ***On Categorizing Innovations***

When considering different types of innovations, it is appropriate to consider three different categories: innovations that provide short-term gains, referred to as incremental innovations, innovations that take existing technologies and reconfigure them for longer term business advantage, called architectural innovations, and innovations that revolutionize processes or require new operating principles in core subsystems that are referred to as discontinuous innovations.<sup>11</sup>

Teradata's technology platform is a result of major research and development in the manner by which the core subsystems of a database management system operate. First and foremost, the company's founders believed that two distinct orientations to database management systems would be the mainstay of

computerized business processing. While much of the market focused on operational databases (called On-line Transaction Processing or OLTP-based relational database management systems) that handle large volumes of business transactions predominantly involving updates and deletes, another need would be for dedicated database management systems for supporting decision making. In this decision support context, most operations would not be updates and deletions; rather, they would involve scanning, aggregating and drilling down into that granularity of data that serves to provide the solid evidence necessary to support decision making. This orientation to decision support-related operations is known in general as business intelligence, and the term data warehouse is typically used to distinguish platforms with this decision support orientation from the operational orientation of OLTP platforms. It should be noted that a data warehouse might or might not be based on a relational database management subsystem, i.e., some data warehouses actually rely on other types of data management subsystems.

The phrase ‘technology platform’ used above references a set of subsystems and interfaces that form a common structure from which a stream of derivative products can become efficiently developed and produced.<sup>12</sup> Not surprisingly, the phrase is used in the context of manufacturing many things ranging from automobiles to furniture. For example, if we compare software to a chair that has subsystems including the pedestal, stem, seat, armrest and backrest - then system interfaces include various fasteners, coasters and even user interfaces such as the height adjustment lever, seat cushions, etc. In the context of innovation then, an incremental innovation might improve the padding in an armrest or a seat cushion, an architectural innovation might be a type of chair where the backrest is flexible and moves in conjunction with the engagement of a leg rest, and a disruptive innovation might well be a ‘sports chair,’ a chair one can fold up, stow in an automobile and take to an event, unfold and then make use of it. Further, there is an important potential market advantage to a properly designed platform. There is the opportunity for that platform to become a standard basis for large-scale innovation. As scholars have point out:

Platform architecture allows a software company to become a conduit for other individuals or companies building modules based on that architecture. Strategically, the firm can foster the development of a multitude of derivative products created for different uses or market niches without having to bear the direct cost of programming labor... The power of the platform approach for software is similar to that for physical products: a software developer can provide a family of products without starting from zero every time... In most cases, a company’s improvements to its software platform are released as newer versions of an existing product line. In other instances, however, platform markets can open new markets for the same basic technology... Platforms make technical sense not only for software but also for the business.”<sup>13</sup>

### **The Teradata Platform: Born for Parallelism**

United States Patent Number 4,925,311 issued on May 15, 1990 to Teradata Corporation (before it became a part of NCR) describes a unique aspect of the Teradata platform. To put this patent into the context of the historical database market, it was about this time that parallel database systems were being considered as possibilities to displace mainframe database systems. At the time, it had become the norm that a business had to reinvest in every next generation data management system as the mainframe it was

operating on became overloaded. Historically, computing had followed Grosch's Law – a forerunner of Moore's Law. Grosch's Law has been summarized as follows:

“... computers present economies of scale: the more costly is the computer, the price-performance ratio linearly becomes better. This implies that low-cost computers cannot compete in the market. In the end, a few huge machines would serve all the world's computing needs.”<sup>14</sup>

While other vendor patents and other Teradata patents addressed intellectual property surrounding parallel processing, Teradata patent 4,925,311 details the notion of virtualized processors called Access Module Processors (AMPs) that are commissioned into groups to handle a particular workload. It also describes the de-commissioning of AMPs from these groups as work is completed. In this way, the AMPs are automatically commissioned and decommissioned based on dynamic workload needs, and they can be scheduled for other workload tasks as they are decommissioned in order to maintain a balanced workload. In the concluding sections of the patent's “Detailed Description of the Invention” section, the following is stated:

“Substantial advantages derive from this system and method, in that a great many tasks can be carried out concurrently and asynchronously, with the multiple processors being dynamically partitioned into appropriate groups, with group addresses being allocated and subsequently de-allocated as the tasks are completed. Any of the processors can establish such a group, and there is no time limit on the maintenance of the group. Once the group is in existence, however, the need for prefatory and confirmatory status messages is greatly diminished. In contrast to prior art, many partitions can be in effect simultaneously, and an individual processor can participate in many partitions simultaneously.”<sup>15</sup>

Similarly, the Teradata platform relies on virtualized processors called Parsing Engines (PEs) to decompose a query into manageable pieces, distribute the processing of those pieces by AMPs, and then coordinate the results from multiple AMPs' computations. Teradata's AMPs and PEs were optimized for a relational database, and they are independent of the hardware configuration. That means that the software can be ported to different hardware platforms as hardware technologies improve over time. Multiple AMPs and PEs can run on a single node where a node is a fundamental hardware entity (running an operating system) and configured with the Teradata platform. A multi-node configuration of the Teradata platform can manage very large-scale data warehouses. The way the Teradata platform connects with different node hardware and operating system configurations is through what are called Parallel Database Extensions (PDE's). These PDEs enable a layering insulating the platform from the operating system and enabling it to integrate within the Unix SVR4 MP-RAS, Microsoft Windows or SUSE Linux operating system environments. This also means that if there are underlying innovations to those operating systems, modified or new PDEs can be developed to take advantage of them – without having to recreate the Teradata platform. This is known as “platform independence.”

While AMPs and PEs facilitate parallelism, how do they work with memory? A substantial issue in database and data warehouses involves how memory is allocated to and/or shared between processors. For example, if memory is shared between processors, then there is communication overhead to ensure

that one processor isn't making changes to memory while another processor is trying to do the same thing. These conflicts are known as "interference," and they often arise in a shared memory parallelized platform. Another approach is referred to as a 'shared-nothing' approach (Exhibit 4a). This means that memory and is owned by a processor. Shared nothing platforms, like Teradata's, minimize overhead associated with communications and controls to guarantee proper resource sharing protocols, but appropriate methods for managing the coordination of the processors (e.g., AMPs and PEs) are required, and this is where Teradata's intellectual property has stood the test of time.

The inventors of the Teradata platform made an inherent assumption about computing that represented what at the time was a major gamble - it was a 'make or break' opportunity for the small company. That assumption was that high-performance, low-cost commodity components for the platform would be available at a time when commodity components had low performance and were of low quality. Some analysts cited this issue as critical to the industry's move from proprietary, closed systems to an open systems philosophy. An important open systems assumption was that other companies or individuals could build subsystems that might operate on top of, under or inside of the platform. This enables the platform to become the standard for more wide-scale innovation. Another major assumption was that a large set of processors could efficiently communicate across a network (originally called Teradata's Y-Net and now upgraded and referred to as BY-Net) for purposes of coordination. Decomposing data management processes into AMPs and PEs, relying on message exchanges to coordinate those AMPs and PEs, banking on commoditized low cost, high quality components, and counting on an open systems advantage in terms of innovation were all part of Teradata's biggest bet. In fact, Teradata was one of the first to wholeheartedly adhere to an "Intel Inside" philosophy, and many now say, "With Intel's fortunes, so went Teradata's fortunes." In the *New York Times* on September 18, 1990, prior to NCR's acquisition of Teradata, L.W. Fisher wrote:

"Key to the development of [NCR's] new products was Philip M. Neches, 38, a co-founder of the Teradata Corporation in El Segundo, CA, whom NCR hired last year as senior vice president and chief scientific officer. NCR executives joke that Mr. Neches has turned the company into a 106-year-old start-up business. Analysts note that his hair is a lot shorter these days than when he was hired. Teradata is a pioneer in the use of parallel processing, a technology until recently limited to scientific applications and supercomputers for commercial tasks like managing databases. In parallel processing, tasks can be performed more quickly by dividing them up among multiple computer processors operating in parallel. In Teradata's machine, and NCR's new line, these are standard microprocessors made by the Intel Corporation, which keeps costs down."<sup>16</sup>

Teradata's AMPs are where all of the work takes place in the Relational Database Management System (RDBMS). Once configured, queries and many complex database administration processes are accomplished by engaging multiple AMPs simultaneously (Exhibit 4b). In the earlier versions of the Teradata RDBMS, each AMP was actually a physical processor. Later versions 'virtualized' the AMPs, and many could exist on what is referred to as a single Symmetric Multiprocessor (SMP) node. Most multi-processor systems today use an SMP architecture. Even in the case of multi-core processors, the SMP architecture refers to the cores and typically regards each core as a distinct processor. Many SMP nodes can be organized to work together to form a single Massively Parallel Processing (MPP) system.

Obviously, a single SMP node will not handle as large of a computational load as an MPP system, but the MPP system is more complex and involves additional hardware and more software sophistication.

For Teradata's system, and for complex BI systems in general, the main unit of work to be carried out is a query. Parallelizing processes associated with serving a query in a way that saves time and costs is therefore a key objective. While Teradata's AMPs-based architecture provides many parallelism opportunities, there are also other innovations to support even more parallelism. Before a system is populated with data, that system is configured to consist of a given number of AMPs. The Teradata database relies on hash-partitioning the data across all of the AMPs that have been configured for a particular system. This means that each AMP is associated with a subset of the complete database. Therefore, when a query is processed (e.g., a join, an aggregation, a table scan, projection, a selection or an index scan), the work is spread to the AMPs associated with the applicable data. Each AMP performs its operations on data independent of the data associated with other AMPs. The Teradata RDBMS coordinates the processing and aggregates the query results. An additional opportunity for parallelism is to dynamically seek occasions during the processing of the query to do what is called pipelining. A query is decomposed into steps in Teradata's system, and each step may require a large chunk of work. In some cases, a task within that chunk of work can be started before the entire step is completed. For example, while a table is being scanned, rows with specified properties can be pipelined out to another task such as a join operation. In short, pipelining takes advantage of opportunities to increase parallelism when chunks of work are such that a second operation can jump off from a first operation before that first operation is complete. Multi-step or bushy parallelism is another unique and innovative way that Teradata's RDBMS can enhance query execution. The term 'bushy' is used in reference to join trees that can be constructed by parsing a query. As the term implies, a bushy tree means one where there is more balance than in a tree where sequential operations on relations is required. Additional parallelism opportunities are also made available in the Teradata system. For example, for groups of queries, the system provides an extension called 'Multi-Statement Request' to enable the bundling of those queries. This facilitates the optimization of sub-expressions shared by queries within the bundle; they are executed once, and the results are shared to other relevant queries in the bundle. This process is known as sub-expression elimination.

The bottom line to the benefits of parallelism in the Teradata platform is essentially how well it achieves two goals relevant to data warehousing: speed-up and scale-up. Speed-up is measured by the ratio:

$$\text{Speed-up} = \text{big\_system\_elapsed\_time} / \text{small\_system\_elapsed\_time}$$

where a fixed job is executed in both environments, and the size of the system is reflected in the number of nodes/processors it requires. So, the Speed-up is said to be linear if an N times as large (or an N times more expensive system) yields a Speed-up of N. In contrast, Scale-up measures the impact of growing the system size. It is defined as the "ability of an N-times larger system to perform an N-times larger job in the same elapsed time as the original system." This is reflected in the following ratio:

$$\text{Scale-up} = \text{small\_system\_elapsed\_time\_on\_small\_problem} / \text{big\_system\_elapsed\_time\_on\_big\_problem}$$

If the Scale-up ratio evaluates to 1, then Scale-up is said to be linear; if it is close to 1, it is said to be near-linear. Note that the ratios are especially useful in benchmarking and comparing systems where the growth of data is expected, and the need for rapid decision support is required. There are three major impediments to linear/near-linear Speed-up and linear /near-linear Scale-up. First is the start up time - the time it takes to start a parallel operation. If it takes significant time to start up thousands of processors, the start-up time will degrade performance in contrast to a non-parallelized system. Similarly, if there is significant interference such as when processors have to vie for access to shared memory, then parallelized solutions degrade in contrast to non-parallelized solutions. In terms of these dimensions, DeWitt and Gray characterized the Teradata solution in 1992 as follows:

“Teradata quietly pioneered many of the ideas we have presented here [in this research paper]. Since 1978 they have been building shared-nothing, highly parallel SQL systems based on commodity microprocessors, disks and memory. Teradata systems may have over a thousand processors and many thousands of disks. Teradata has installed many systems containing over one hundred processors and hundreds of disks. These systems demonstrate near-linear Speed-up and Scale-up on relational queries, and far exceed the speed of traditional mainframes in their ability to process large (terabyte) databases.”<sup>17</sup>

While the technical aspects of the Teradata platform have their roots in novel approaches to parallelism, Jim Dietz, Teradata Product Marketing Manager for Platforms, recently summarized these beginnings and brought them up to date in terms of the company’s current products and services:

“The scale-up and speed-up properties inherent to our original Teradata platform have been keys to achieving sustainable competitive advantage. The pioneers of our original platform had the innovative foresight to bet the farm on a future whereby our data warehousing software can continue to evolve by taking advantage of innovations to commodity microprocessors, memory and disks. Our R&D leaders even adopted a focused innovation strategy that kept our eyes on the data warehouse ball. This kept us from getting tangled in an operating systems quagmire, and it sent the message that Teradata *invited* partner and customer innovations to underneath, on top of and even inside the platform (through commodity components). If you could visit the labs of Teradata today, you’d see how we assemble components into various configurations, put those configurations through their hoops, and carefully benchmark scale-up and speed-up against our competitors’ systems. I call these configurations ‘technology bundles,’ and our R&D surrounding them has been key to the recent launch of the Teradata platform family of products and services in what we call the ‘Accelerate’ Program. So, these technology bundles are just one piece of the R&D puzzle – we also have a history of success in logical data model innovation, and our patents in that area are indicative of how we’ve approached the linkage between logical and physical data management.”<sup>18</sup>

### ***A Logical Starting Point and a Roadmap Leading to an Enterprise Data Warehouse***

Creating a new data warehouse solution is a complex process that requires a focus on business objectives. Teradata has even patented a visual modeling tool to help simplify the process and make it more efficient and effective: “System and Method for Planning and Implementing a Data Warehouse Solution” [ US

7,092,968 B1, August 15, 2006].<sup>19</sup> It is important to note that one company in an industry that is following a particular business strategy may need different data in their warehouse than another that has adopted a different business strategy (e.g., low-cost vs. focus vs. differentiation strategy). Similarly, organizations in one industry will need to manage different data than organizations in another industry (e.g., healthcare vs. manufacturing). It is no wonder that the time required to create a data warehouse solution from scratch for a particular business is typically measured in terms of years – capturing the business requirements and translating them into a working, complex data warehouse is a serious undertaking. It's also common that such an effort could require the skills and expertise of between fifty to one-hundred trained data modelers.

In the process known as data modeling, a common first step is to create a business-relevant 'logical' data model. The logical data model (LDM) synthesizes what is important to the business, and by using a variety of representational methods and tools, data modelers collaborate with business decision makers to ensure that requirements for a new data warehouse are being properly articulated and captured. Since the ultimate data warehouse solution will require the services of a variety of experts in different business discipline and technical areas, the complexity of creating an LDM can be high. An inadequate LDM can doom an entire data warehousing solution to failure from its onset. Teradata approaches the logical data modeling effort in a unique way. Their approach begins with a perspective that focuses on what is called an "Enterprise Data Warehouse" or EDW, for short. An EDW strives to provide an integrated view of data that is consistent across the organization, i.e., a 'single source of the truth' – with the ability to provide many views on top of that single source. An EDW that is also capable putting right-time, vital information into the hands of front-line decision makers is called an 'active' EDW. Active EDWs represent the most sophisticated data warehouse ideals, and they are also among the most complex to construct – yet they can have the highest bottom line payoff. However, not every organization needs an active EDW, nor will all organizations begin their data warehousing journey with an enterprise-wide solution. Some organizations want to focus their initial data warehouse investment on a particular core business area, unit or a functional area. In other words, some organizations need initial logical data models that map only to the business processes they are targeting for decision making improvements.

Teradata's historical success has been in the area of active EDW, and in fact, the phrase itself is attributed to Teradata advancements related to near real-time capabilities harnessed within the company's EDW platform. Teradata applies the trademark "Active Enterprise Intelligence" to its vision for real-time business intelligence:

"An Active Enterprise Intelligence business environment is built on the foundation of a central data warehouse and delivers detailed visibility to a company's operational front-line employees, partners, and customers in real time. Active Enterprise Intelligence systems provide instant insight to customer profiles, supply chain disruptions, inventory status, fraud patterns, package and passenger routing, labor surges, or out-of-stock conditions – in fractions of a second. Active Data Warehouse (ADW) is the term for platforms which operate at the Active Enterprise Intelligence level of performance. Some business decisions are just too time-sensitive, and the stakes too high, to rely entirely on day-old analytics. "For the most time-critical decisions, enterprises require 'really urgent' analytics," said James Kobielus, a senior analyst at Forrester Research, in his August 11 report *Really Urgent Analytics: The Sweet Spot for Real-Time Data*

*Warehousing.* He added that "information and knowledge management professionals should familiarize themselves with the various approaches for adapting their enterprise data warehouse for real-time requirements." The term 'Active Data Warehousing' was coined in 1999 when a Teradata Magazine article described how a few customers were using data warehouse insights in operational processes. The approach driven by Teradata customers and engineers is today known by terms like 'pervasive BI,' 'operational BI,' and 'real-time BI.'" <sup>20</sup>

Teradata's Chief Technology Officer of Research and Development, Todd Walter, recently reinforced the significance that Teradata's Active Enterprise Intelligence vision plays in the recently released Teradata 13.0 product. With significant changes that are transparent to users, Teradata engineers virtualized the storage subsystem architecture, redesigned the management functions and introduced automation of replication services (replications are easier to manage and maintain) and the Multi-system Manager (reports on the status of sub-systems, data and applications, and it sends alerts if operational thresholds are reached). Walter wrote:

"Virtualization, management, ultra-high availability, functionality and performance are all offered through Teradata 13.0. Together, the changes provide the technology to meet the service levels required for Active Enterprise Intelligence implementations while maintaining the high performance that is valued in the Teradata system. With the installation of Teradata 13.0, organizations can further optimize their data warehouse investment."<sup>21</sup>

Teradata's innovations centering on LDMs are founded on their historical active EDW perspectives as enabled through their scalable, high performance parallelized platform. Because of these innovations - and through the lens of the company's many experiences in implementing successful EDWs - Teradata has developed broad capabilities and patented methodologies to realize what it calls its 'data warehousing roadmap.' Within that roadmap are building blocks that support organizations making the right choices from the beginning - and that enhance a direct migration path should a customer's initial data warehousing project (one with a focus on certain business processes, units or functional areas) expand to a broader active EDW journey down the line. This means that while LDMs used within the roadmap will support a particular focus (e.g., manufacturing, financial services, etc.), those seemingly independent focuses are actually highly integrated parts of what Teradata already possesses as its broader enterprise-wide LDM nucleus. This modularity and capability to synthesize are keys to the company's efficiency in scaling-up those customers that opt to start their data warehouse journey with a particular focus, but realize that an enterprise solution represents an important next step.

US Patent 7,359,906 dated April 15, 2008 exemplifies the company's continued march to speed-up the deployment of a Teradata LDM to address the needs and nuances of its individual customers - particularly with respect to linking together LDMs in a fast way to match hybrid business models that are becoming more common. The innovations described in this patent substantially improve the efficiency in customizing LDMs for Teradata customers. The patent states:

"A logical data model design methodology utilizing shared subject areas provides for more effective new LDM development through re-use of common elements and quicker deployment of horizontal applications across all industries. Shared subject areas enable

Professional Services (PS) consultants who work with multiple LDMs to leverage knowledge across LDMs and facilitate customer/PS combining of LDMs. Shared subject areas represent “configurable” LDM components that facilitate the development of hybrid business models. For example, a communications company that has retail store outlets to sell wireless services crosses the retail and communications LDMs. A component-based architecture allows configuration of LDMs for more complex business models.”<sup>22</sup>

Proprietary innovations surrounding the Teradata EDW roadmaps and LDMs have and will remain part of the Teradata culture. The LDM modular philosophy, along with incremental and managed LDM improvements, have remained a central thrust for Teradata as demonstrated in the following comment by Barb Vostmyer, Teradata Senior Manager for the Manufacturing Industry:

“The Teradata LDM nucleus is a single model with all industry atomic data prepared for an international implementation. The LDM is the starting point for a roadmap which can enable our customers to realize a single source of the truth. These LDMs require relatively little rework to custom fit to a customer's specific needs or to change when there is growth to that customer's enterprise data warehouse vision. Our Professional Services staff has experience customizing the LDM with customers in specific industries across the globe. Teradata keeps improving our LDM products so that they continue to include the latest data such as RFID tracking, and we keep them consistent with emerging industry standards such as a Basel II mapping for our financial services LDM.”<sup>23</sup>

While the LDMs are a target for continuous innovation, so too is the way those LDMs serve to provide a focus for efficiently customizing a customer roadmap. In other words, the LDMs are leveraged to speed-up the articulation and formulation of a client's requirements in a way that is customized to their particular strategy, data needs and data warehousing objectives. The manner by which the roadmap is co-produced with the customer is also supported by toolsets and applications. Teradata employees have contributed significant innovations to the roadmap process and applications, further exemplifying the innovation culture at the company. For example, Barb Vostmyer, quoted above, discussed an invention she was a part of as follows:

““To support a client engagement, we can leverage each focused LDM within our LDM nucleus. We have a new roadmap application that makes use of patent-pending intellectual property that our team helped to invent. Imagine you are a client contemplating what Key Performance Indicators (KPIs) that you really need to track in order to make the most important decisions in your company. Wouldn't you want to know what data you'd need in your data warehouse to support the right time calculation and display of those KPIs so they're in the right peoples' hands who can take action? Our newest roadmap tool does just that. Companies can create scenarios and we can efficiently determine their needs based on the scenario they select. We also can immediately begin work on going through the relevant LDMs with the company's key people to get them customized and move them along. Further, if they opt for a few KPIs to begin with, then when they want more, it becomes a matter of scaling-up to the additional data needs they'll have for their EDW, and due to our LDM nucleus approach,

the upward migration is much more effective. Our Teradata parallel platform, how we manage LDMs, and how well our overall vision of what an active EDW ought to do for a company have enabled us to bundle technical capabilities with rapidly customizable logical data model needs. These capabilities are central to the newest versions of what has come to be known as the Teradata platform family and the Accelerate Program.”<sup>24</sup>

## **Business Intelligence and Customer Value Streams**

The most important objective of a data warehouse investment is to support better organizational decision making. In US Patent 6,553,366 titled, “Analytic Logical Data Model,” Teradata inventors captured the spirit of innovation thrusts being reported in the data warehousing area even today. A simple, yet omniscient set of statements in this patent overview are relevant to an area called *in-database processing*:

“The present invention provides a relational database management system (RDBMS) that supports data mining operations of relational databases. In essence, advanced analytic processing capabilities for data mining applications are placed where they belong, i.e., close to the data. Moreover, the results of these analytics processing capabilities can be made to persist within the database or can be exported from the database.”<sup>25</sup>

Data warehouses built on RDBMS involve enormous amounts of data that are stored and processed. Data mining and analytics involve processes of identifying and interpreting patterns from those data warehouses. Discovered patterns can then be leveraged to support decisions. For example, if a data warehouse has a great deal of customer and sales data for an e-commerce catalog sales organization, then mining past purchasing patterns can provide support for decisions surrounding who best to target for special offers, new catalog sales, etc. Similarly, customer marketing campaigns extrapolate from data warehouses to help target and even personalize specific promotions through helping to identify segments. Mining for patterns can also provide insight on ‘customer churn,’ such as when potential customers may switch cell phone companies or upgrade their contracts. In supply chain management, data mining supports forecasting and risk decision making. In insurance and other industries, data mining is also an important ingredient for innovation. For example, an insurance company may design insurance plans through data mining. At eBay, even ‘analytics as a service’ for innovation activities is being pioneered. eBay runs the world’s largest commercialized data warehouse environment at more than 5 petabytes, and that data warehouse is powered by Teradata products. eBay research groups can, via self-service, leverage the Teradata platform’s virtual systems to do research and development activities on subsets of the EDW. Over 50 business units at eBay run such test environments on the EDW at a given time, and each test span is typically 90 days. This type of virtual laboratory exemplifies support for rapid experimentation using data mining and analytics, and it reinforces fast learning and thereby improved time to market.

Different industries, organizations and business strategies imply that there is no ‘one size fits all’ approach to harnessing the potential of data warehouses, logical data models, data mining and analytics. How, then, can a company like Teradata make the best choices for innovating, designing and targeting its products and services? In 1995, James Martin in *The Great Transformation* coined the phrase ‘strategic value streams.’<sup>26</sup> In essence, a company manages many value streams; each is an “end-to-end collection

of activities that creates a result for a customer.” That ‘customer’ may be the ultimate consumer or beneficiary, or it may be an ‘internal customer’ of the value stream, but the value stream has a clear goal: “... to satisfy or delight the customer.” A strategic value stream is more than an organization’s core competency; it is a unique capability that enables a corporation to move faster or better than another company - and it is one that competing firms cannot emulate easily. Value streams are prioritized in companies, and those not performing up to par can be targeted for reinvention. Of course, the value streams within a company are connected to relevant external entities as well as other within-enterprise value streams. The broad discipline of business intelligence, comprised of data warehousing, LDMs, data mining and analytics, targets the reinvention and improvement of enterprise value streams.

R. Whittle and C. Myrick have summarized a set of generic value streams in reference to their work on enterprise business architectures.<sup>27</sup> Their approach’s generality enables a cross-industry, company and business strategy perspective. A value stream can be expressed in an ordered pair of phrases or keywords that address a transformation between particular states in the context of a set of business processes and events. For example, concept-to-development is a value stream that expresses: when a company has a concept, then a key set of business processes and events are associated with converting or otherwise operating on that intangible ‘concept’ to achieve the resulting state of ‘development.’ This implies that the ideas expressed in a conceptual form are operationalized in a more tangible state. Some of the value streams that are most relevant to today’s business intelligence market are as follows:

- *Prospect-to-Customer*: Taking a business prospect and transforming that prospect to an actual customer
- *Order-to-Cash*: Receiving an order and performing the necessary business processes to ultimately receive the cash for the order
- *Manufacturing-to-Distribution*: Once a good is manufactured, the processes and events necessary to achieve the state whereby the good has been transported to the proper location for sale or further elaboration
- *Request-to-Service*: Once a request is made of a service provider, the co-production associated with delivering the service to satisfy the request
- *Forecast-to-plan*: Once a forecast is made, those processes involved in making that forecast a part of an actionable plan
- *Requisition-to-Payables*: The business processes starting with a requisition and achieving the point where there is a payable, i.e., the procurement process
- *Awareness-to-Prevention*: Once a particular problem or issue has risen to the level where action is needed, those business processes necessary to prevent the problem.
- *Insight-to-action*: A significant insight leads to actions that leverage that insight.

Many of the value streams above are targets for business intelligence since there are many insights that can be gleaned from reviewing historical data to discover patterns for making process improvements and making better state transformation decisions. In fact, business intelligence can help to surface new business rules that can be applied directly to improve a value stream’s underlying processes. In addition, business intelligence can help to innovate new business processes to improve or redefine value streams – such is the case at eBay where virtual laboratories are set up for research and development personnel to test and evaluate their ideas<sup>28</sup>.

What value streams are most important for Teradata's customers? At the most generic level, the insight-to-action value stream is among the most important. Exhibit 5a is a figure eight model of this generic value stream. On the action loop, there are three steps: sense, assess and respond. On the insight loop (often called the strategic intelligence loop), there are four steps including: reporting, analyze and predict. A connecting arc flowing from the action loop to the insight loop is labeled, "plan." Consider an example of an e-commerce DVD sales portal. On the insight loop, *reporting* involves answering the question, "What happened?" For example, a decision maker might want to know what DVD sales were this week versus last week. For *analyzing*, the question becomes, "Why did it happen?" For example, in shopping carts with DVDs, what other items were purchased most? Next, *predicting* involves asking the question, "What will happen?" Here, we may want to know what will happen to kid's DVD sales after a 6% price increase. Finally, *plan* sets the stage for next steps, for example, "Let's figure out how to make the best outcome a reality." *Sense, assess* and *respond* often relate to what is referred to as 'operational intelligence.' The following perspective provides insight on this loop:

"What he [Darwin] theorized was that the fittest - or those that are best at adapting to their environments - will survive... You don't have to be the strongest - just the fastest to adapt to a changing environment. Transfer that idea to businesses and you get the idea that if your company is not the biggest but is more flexible and agile, then it can react faster than its bigger competitors to a moving marketplace -- eventually eclipsing and surpassing its slower moving "big brothers". To do this, your business paradigm must change from the "Make and Sell" model of the past (think Henry Ford) to the new "Sense and Respond" one. The company must be able to sense what markets and customers really want, that is -- to understand the patterns in buying behaviors, the trends in markets, the reactions of your competitors, etc. -- so you can respond more quickly and appropriately. This is where a well engineered business intelligence environment comes into play. The agile enterprise must be able to sense and respond quickly to opportunities and threats in its environment. How? The way to do this is to change from:

1. Sequential to synchronous information flows
2. Batch to right-time updates (the right data at the right time to make the right decision)
3. Chains to "hubs" in the company's process architecture
4. Enterprise automation to value network optimization

This last step -- value network optimization -- means recognizing that most business processes span beyond your company to include suppliers, partners and customers. Each link -- supplier, partner, customer, and employee -- must be able to operate and make good decisions at the right time which means having access to the right data." <sup>29</sup>

While there are likely thousands of value streams and organizations might define them differently - and there are a variety of approaches for identifying and prioritizing them, the underlying theme of insight-to-action is a key. Teradata's focus is on its customers' value streams, and innovations surrounding those customer needs are focal points for its own value streams. For example, Teradata's value streams have targeted innovations explicitly discussed in this case: Custom LDMs-to-Customizable LDMs (e.g., Patent 7092968), Slow relational DBMS-to-Scalable, high performance data warehouses (Patent 4925311) and Slow analytics-to-Fast analytics (e.g., Patent 6553366). That said, Stephen Brobst recently stated, "At Teradata, we won't just focus on *our* value streams over the next several years, it is very important for us to emphasize that our products and services are an integral part of the innovation streams *for our customers.*" In fact, Teradata provides customers with an overall roadmap for a data warehousing journey (Exhibit 5b) in terms of specific value stream questions.

In terms of Teradata's customers, there is another overarching theme to investments made on behalf of innovating, improving or reengineering value streams. That theme can be expressed as a value stream in its own right: *initiative-to-results*. This is one of the most important value streams related to investing in business intelligence, that is, when such an initiative is given the green light - how quickly and effectively can the actual value be realized? To answer that question, Teradata R & D must always have a focus on the next generation of potential products and services. For example, in a recent interview, CTO Brobst stated:

“Teradata has been in the data warehouse game for more than 25 years. One might think that after all this time the industry would have reached a status quo, [but] more and more companies are jumping into the data warehousing fray, and innovation is accelerating rather than slowing down. Four areas will be particularly interesting over the next several years and are likely to drive huge increases in data volumes, data usage and the need for scalability in data warehouse solutions:

1. Massively distributed sensor networks will generate huge quantities of data with opportunities for advanced analytics that are unfathomable by today's standards.
2. Decisioning services provided to front-line knowledge workers help transform the strategic vision of an organization to operational reality. The next step will be to deliver such capabilities to suppliers, distributors, customers and government agencies.
3. In yesterday's world, sophisticated analytics were often performed in separate data marts using specialized file systems. Advances in Teradata relational database management systems and third-party application technologies allow *in-database processing* to deliver significantly better performance for high-end analytics. Another example of this trend is the evolution of traditional extract, transform and load (ETL) tools to the extract, load and transform (ELT) approach, wherein transformations take place inside the scalable relational database rather than on an external server. Data mining, traditionally performed inside proprietary file systems, has also moved into the data warehouse. Implementations by KXEN and SAS demonstrate the power of in-database processing even for the most sophisticated analytics.
4. In tomorrow's world, structured content will constitute less than 20% of the volume in a data warehouse. New data types such as biometrics, images/video, sound/voice, geospatial, text and XML documents will dominate the storage and analytic resources in advanced data warehouses. This will require new tools for analysis and extensions to traditional relational structures for storing and processing complex data types.”<sup>30</sup>

### **The Accelerate Program: Family Power**

In the January 14, 2009 report of Current Economic Conditions by Federal Reserve District report (The Beige Book), the news was not positive.<sup>31</sup> All businesses feel the pinch in a recession, including Teradata's current and potential customers. At the same time, many companies that don't deploy business intelligence and data warehousing solutions experience major difficulties, not the least of which is trying to track down where they can cut costs. Some analysts reinforced the notion that recessions are the right time for business intelligence to prove itself. Alastair Taylor, a managing consultant at PA

Consulting, stated: "Significant business change often triggers big-bang, enterprise-wide BI, and the big change in the economic climate is forcing many customers to rethink their business models."<sup>32</sup> She added, "That climate is encouraging enterprises to explore potential efficiencies across the organization to enhance cross- and up-selling and identify areas for cost-cutting. Businesses want to find ways of assessing 'the relative return on investment today, tomorrow and in the future as the basis of where they put their capital'."

On November 19, 2008, Teradata's Accelerate Program was announced (See Exhibit Six). The Accelerate Program leverages Teradata's historical innovations centered on speed-up and scale-up to introduce comprehensive Teradata platform, LDM and professional services bundles that provide a wide variety of customers with everything needed to deploy an entry-level data warehouse with low risk at a competitive price point. The Accelerate Program hinges on the company's purpose-built platform family. Another main element of the Accelerate Program relates to the new positioning of Teradata products and service bundles squarely into the generic and timely *initiative-to-results* value stream where "time to value" has become the Program's mantra. In addition, investments in future product and service innovations were already well underway in areas like sensor technology, pervasive BI, non-traditional data types and in-database processing. By the end of January 2009, Darryl McDonald, CMO, was reflecting on the break-neck pace of new product and service introductions as he sat down to draft his monthly blog. Last month's blog had highlighted Gartner's most recent ranking of Teradata as clearly in its "Leadership Quadrant." He recalled a report he'd read shortly after the October Partner's conference by Ventana Research titled, "Teradata: Confident and Competitive." That title pretty much summed up 2008. He also knew that the "time to value" rallying cry behind the Accelerate Program was a solid market differentiator, and it was resonating with customers. In addition, he was aware that competitors had been touting that their BI customers were becoming members of what they referred to as the "Petabyte Club," but he knew that Teradata had an impressive list of customers that had reached that milestone quite awhile ago. Finally, he contemplated the Accelerate Program and how it offered customers specific solutions – would that change any of his marketing strategy? For now, he was facing the question of he was going to fit all of this into appropriate blog-speak. He sat down and penned the following:

***Product & Service Innovation Alignment Yields "Time to Value" Solutions***

Product/service bundles are not new. In fact, companies often assemble products and services into bundles just to be able to market a comparative cost advantage over competitors' offerings. Take my cable company, I counted three or more bundled offers in less than a two month span. The latest, "Save \$ by getting your fixed phone, internet and HDTV from your cable company, and we even provide a free cable modem and set-top box." Or, my cell phone provider, "Get the cell phone of your dreams at \$xx with only a two-year contract." Each different bundle promises this or that cost advantage over competitors' offerings. But bundling for the sole purpose of comparative cost advantage isn't the whole equation. When data warehouses and business intelligence capabilities are at stake, bundling should have one sole objective: each and every bundle has to be a *complete solution* that delivers both right-time and long-term customer value. Our customers must meet their objectives in a fitting time frame wherein a new initiative can manifest a stream of desired and valuable results. Anything else falls short.

The Teradata purpose-built platform family provides value-driven solutions so that in the Accelerate Program, each package is imbued with necessary hardware, software and professional services needed to meet customers' needs – all in a fitting time frame. These solutions haven't been slapped together for splashy headlines; they've been on our drawing board and have been perfected over a long period of time. As proof, consider that with the Teradata approach, customers can follow a seamless migration strategy as their data warehousing needs evolve. After all, the Teradata platform family and professional services solutions are the focal point for our employees, our customers and Teradata's business partners to innovate *in, on, under, around and with*.

Teradata began managing innovation priorities to align product and service evolutions in a coordinated way - even in the earliest days when our founders really did meet in a garage. Our evolution of intellectual property laid the foundation, and the alignment of product and service innovations with frequent prioritization and re-prioritization keeps us on the right track. Nowadays, the innovation areas we invest in include those like logical data model services, mixed workload capabilities, scalability, solution methodologies, etc., and they are all integrated - a venerable mix of product and service focuses - to meet customer and market requirements. At Teradata, it's not about the big release; it's how we think about changing the market again.

While CMO McDonald was working on his blog, a contemplative Scott Gnau, Vice President and General Manager of Teradata Research and Development, waited patiently in the CEO's reception area. Prior to the meeting, Scott held the final meeting with the 'Concept Team.' That team is basically Teradata's clearing house for vetting, prioritizing and building business cases for the company's innovation investments. The team operates using a well defined process model that engages the CTO, senior engineers, production specialists, marketing specialists, and it gathers input from many other teams – including customers. The Concept Team makes recommendation decisions in cycles, but not on a regular calendar basis. In fact, cycles are typically tied to major product release dates.

So today was the meeting he and the team had been preparing for. Today he would find out next year's R&D budget. He'd provided a projected budget that was more than what had been allocated than in 2008, but he knew Mike Koehler would have some tough questions. Scott's team had thought through a great many possibilities. Mike Koehler might ask, "If I decreased your projected budget by 5%, what innovation priorities would you eliminate?" In contrast, the question could be, "If I increased your projected budget by 10%, what new innovation priorities could you add?" Maybe he'd get both questions this year since there are solid reasons to invest in R&D in a volatile economy, but there are also other economic reasons that might require scaling back. Scott and his team had an analogy they often used to refer to each innovation in the budgeted portfolio. Some were 'little rocks' and some were 'big rocks,' where the size of the rock reflected both the significance of the innovation and the investment required to bring the innovation to fruition. Big rocks carried a bit more risk, and there were always some big rocks in the portfolio - this year was no exception. But there were little rocks as well. All of the rocks were important, but if he was faced with the 5% reduction, he knew there was one big rock that could be removed, or perhaps he could leave that and remove several little rocks. Either way, some important innovations were going to get left on the table. If there was an increase, he had to be ready to consider if he could pull them off given the resources at his disposal. What would it take to scale up if he added another big rock? Would he be better off adding several little rocks instead?

## Exhibit One: Third Quarter Financial Statements

### CONSOLIDATED STATEMENTS OF INCOME

For the Year Ended December 31	2008	2007	2006
<i>In millions, except per share amounts</i>			
<b>Revenue</b>			
Product revenue	\$ 849	\$ 884	\$ 807
Service revenue	913	818	740
<b>Total revenue</b>	<b>1,762</b>	<b>1,702</b>	<b>1,547</b>
<b>Operating expenses</b>			
Cost of products	302	312	289
Cost of services	511	474	429
Selling, general and administrative expenses	508	470	410
Research and development expenses	108	126	117
<b>Total operating expenses</b>	<b>1,429</b>	<b>1,382</b>	<b>1,245</b>
<b>Income from operations</b>	<b>333</b>	<b>320</b>	<b>302</b>
Other income, net	5	2	–
<b>Income before income taxes</b>	<b>338</b>	<b>322</b>	<b>302</b>
Income tax expense	88	122	110
<b>Net income</b>	<b>\$ 250</b>	<b>\$ 200</b>	<b>\$ 192</b>
<b>Net income per common share</b>			
Basic	\$ 1.40	\$ 1.11	\$ 1.06
Diluted	\$ 1.39	\$ 1.10	\$ 1.06
<b>Weighted average common shares outstanding</b>			
Basic	178.1	180.8	180.7
Diluted	179.8	181.3	180.7

The accompanying notes are an integral part of the consolidated financial statements.

Exhibit One: Third Quarter Financial Statements (continued)

CONSOLIDATED BALANCE SHEETS

At December 31	2008	2007
<i>In millions, except share amounts</i>		
<b>Assets</b>		
<b>Current Assets</b>		
Cash and cash equivalents	\$ 402	\$ 270
Short-term investments	40	–
Accounts receivable, net	451	507
Inventories	44	51
Other current assets	78	45
<b>Total current assets</b>	<b>1,015</b>	<b>873</b>
Property and equipment, net	88	94
Capitalized software, net	80	61
Goodwill	110	90
Deferred income taxes	109	140
Other assets	28	36
<b>Total assets</b>	<b>\$ 1,430</b>	<b>\$ 1,294</b>
<b>Liabilities and stockholders' equity</b>		
<b>Current liabilities</b>		
Accounts payable	\$ 99	\$ 120
Payroll and benefits liabilities	83	91
Deferred revenue	255	246
Other current liabilities	103	115
<b>Total current liabilities</b>	<b>540</b>	<b>572</b>
Pension and other postemployment plan liabilities	83	88
Other liabilities	30	3
<b>Total liabilities</b>	<b>653</b>	<b>663</b>
<b>Commitments and contingencies (Note 9)</b>		
<b>Stockholders' equity</b>		
Preferred stock: par value \$0.01 per share, 100.0 shares authorized, no shares issued and outstanding at December 31, 2008 and 2007, respectively	–	–
Common stock: par value \$0.01 per share, 500.0 shares authorized, 180.5 and 181.0 shares issued at December 31, 2008 and 2007, respectively	2	2
Paid-in capital	572	555
Treasury stock: 6.9 and no shares at December 31, 2008 and 2007, respectively	(137)	–
Retained earnings	329	79
Accumulated other comprehensive income (loss)	11	(5)
<b>Total stockholders' equity</b>	<b>777</b>	<b>631</b>
<b>Total liabilities and stockholders' equity</b>	<b>\$ 1,430</b>	<b>\$ 1,294</b>

The accompanying notes are an integral part of the consolidated financial statements.

## Exhibit One: Third Quarter Financial Statements (continued)

### CONSOLIDATED STATEMENTS OF CASH FLOWS

For the Year Ended December 31	2008	2007	2006
<i>In millions</i>			
<b>Operating activities</b>			
Net income	\$ 250	\$ 200	\$ 192
Adjustments to reconcile net income to net cash provided by operating activities:			
Depreciation and amortization	60	68	55
Stock-based compensation expense	21	17	9
Excess tax benefit from stock-based compensation	(1)	(1)	–
Deferred income taxes	38	80	(14)
Impairment on equity investment	3	–	–
Changes in assets and liabilities:			
Receivables	73	(128)	(29)
Inventories	7	(12)	(10)
Current payables and accrued expenses	(7)	75	1
Deferred revenue	13	52	15
Other assets and liabilities	(17)	36	–
<b>Net cash provided by operating activities</b>	<b>440</b>	<b>387</b>	<b>219</b>
<b>Investing activities</b>			
Purchases of short-term investments	(90)	–	–
Proceeds from sales and maturities of short-term investments	50	–	–
Expenditures for property and equipment	(19)	(50)	(20)
Additions to capitalized software	(52)	(45)	(46)
Purchased software licenses	(2)	(5)	(2)
Other investing activities and business acquisitions, net	(23)	(4)	(21)
<b>Net cash used in investing activities</b>	<b>(136)</b>	<b>(104)</b>	<b>(89)</b>
<b>Financing activities</b>			
Repurchases of Company common stock	(176)	–	–
Cash contributions from former parent	–	200	–
Transfer to former parent, net	–	(216)	(130)
Excess tax benefit from stock-based compensation	1	1	–
Other financing activities, net	8	1	–
<b>Net cash used in financing activities</b>	<b>(167)</b>	<b>(14)</b>	<b>(130)</b>
Effect of exchange rate changes on cash and cash equivalents	(5)	1	–
Increase in cash and cash equivalents	132	270	–
Cash and cash equivalents at beginning of year	270	–	–
<b>Cash and cash equivalents at end of year</b>	<b>\$ 402</b>	<b>\$ 270</b>	<b>\$ –</b>
<b>Supplemental data</b>			
Cash paid during the year for:			
Income taxes	\$ 33	\$ 1	\$ –
Interest	\$ 1	\$ –	\$ –

The accompanying notes are an integral part of the consolidated financial statements.

Exhibit One: Third Quarter Financial Statements (continued)

CONSOLIDATED STATEMENTS OF CHANGES IN STOCKHOLDERS' EQUITY

	Common Stock		Treasury Stock		Paid-in Capital	Retained Earnings	Accumulated Other Comprehensive Income (Loss)	Parent Company Investment	Total	Comprehensive Income for the Year Ended
	Shares	Amount	Shares	Amount						
<i>In millions</i>										
<b>December 31, 2005</b>	-	-	-	-	-	-	\$ 15	\$ 502	\$ 517	
Employee stock compensation allocated from NCR								9	9	
Net income								192	192	\$ 192
Currency translation adjustments							3		3	3
Net transfers to parent								(130)	(130)	
<b>December 31, 2006</b>	-	-	-	-	-	-	\$ 18	\$ 573	\$ 591	\$ 195
Employee stock compensation allocated from NCR								11	11	
Net income						79		121	200	\$ 200
Net transfers to parent								(259)	(259)	
Contribution of net assets to Teradata Corporation and issuance of shares to parent (Note 1)	181	2			548		(30)	(446)	74	
Employee stock compensation plans					6				6	
Income tax benefit from stock compensation plans					1				1	
Currency translation adjustments							7		7	7
<b>December 31, 2007</b>	181	\$ 2	-	\$ -	\$555	\$ 79	\$ (5)	\$ -	\$ 631	\$ 207
Net income						250			250	\$ 250
Adjustments to net assets contributed from NCR (Note 1)					25		2		27	
Employee stock compensation plans	1				21				21	
Proceeds from employee stock purchase program and option exercises	1				8				8	
Repurchase of Company common stock, retired	(2)				(38)				(38)	
Income tax benefit from stock compensation plans					1				1	
Purchases of treasury stock, not retired			(7)	(137)					(137)	
Pension and postemployment benefit plans, net (Note 7)							10		10	10
Currency translation adjustments							4		4	4
<b>December 31, 2008</b>	181	\$ 2	(7)	\$ (137)	\$572	\$329	\$ 11	\$ -	\$ 777	\$ 264

The accompanying notes are an integral part of the consolidated financial statements.

## Exhibit Two: Teradata Partners

(Listing Courtesy of Teradata Corporation:  
Retrieved August 20, 2009 from <http://www.teradata.com/t/partners-list/>)

<a href="#">Ab Initio</a>	<a href="#">Aleri</a>	<a href="#">Attensity</a>
<a href="#">AIRCOM International</a>	<a href="#">Atanasoft</a>	<a href="#">BearingPoint</a>
<a href="#">Assetlink</a>	<a href="#">Bakbone Software</a>	<a href="#">Business Objects - An SAP Company</a>
<a href="#">Attunity</a>	<a href="#">BIS2</a>	<a href="#">Chordiant Software</a>
<a href="#">BEZ Systems, Inc.</a>	<a href="#">Capgemini</a>	<a href="#">Composite Software</a>
<a href="#">CA</a>	<a href="#">Cognos, An IBM Company</a>	<a href="#">DFA</a>
<a href="#">Clarabridge</a>	<a href="#">Deloitte</a>	<a href="#">EMC</a>
<a href="#">DataDirect Technologies</a>	<a href="#">Embarcadero Technologies</a>	<a href="#">Fox Technologies</a>
<a href="#">EDS</a>	<a href="#">FICO</a>	<a href="#">Hummingbird, a division of OpenText</a>
<a href="#">Epic Systems</a>	<a href="#">Guardium</a>	<a href="#">IBM Tivoli Storage Manager</a>
<a href="#">GoldenGate</a>	<a href="#">IBM IIS</a>	<a href="#">Information Builders</a>
<a href="#">IBM DataMirror</a>	<a href="#">Informatica</a>	<a href="#">LSI Corporation</a>
<a href="#">Infinian</a>	<a href="#">KXEN</a>	<a href="#">Novell</a>
<a href="#">Intelimedix</a>	<a href="#">MicroStrategy</a>	<a href="#">Portrait Software</a>
<a href="#">Microsoft</a>	<a href="#">Oracle</a>	<a href="#">SafeNet</a>
<a href="#">OptiMine Software</a>	<a href="#">Relational Solutions, Inc.</a>	<a href="#">Scorecard Systems</a>
<a href="#">Protegrity</a>	<a href="#">SAP</a>	<a href="#">SoftwareAG</a>
<a href="#">SAS</a>	<a href="#">Simba Technologies Inc.</a>	<a href="#">StrongMail Systems, Inc.</a>
<a href="#">Shiloh Technologies</a>	<a href="#">SPSS</a>	<a href="#">Symantec (Veritas) Software Corporation</a>
<a href="#">Speed-Trap.com, Ltd.</a>	<a href="#">SunGard</a>	<a href="#">Tableau</a>
<a href="#">Sun/StorageTek</a>	<a href="#">Syncsort Inc</a>	<a href="#">Trillium Software (Harte Hanks)</a>
<a href="#">Symark, Inc.</a>	<a href="#">TIBCO</a>	<a href="#">Vision Chain</a>
<a href="#">Talend</a>	<a href="#">Ventraq (TeleSciences)</a>	<a href="#">Webtrends</a>
<a href="#">Unica</a>	<a href="#">Ward Analytics</a>	
<a href="#">Voltage Security, Inc.</a>	<a href="#">WisdomForce</a>	
<a href="#">WhereScape Software Ltd</a>	<a href="#">Agilent</a>	
<a href="#">Accenture</a>	<a href="#">Aprimo</a>	

### Exhibit Three: Teradata History

1976 to 1979	The concept of Teradata grew out of research at the California Institute of Technology (Caltech) and from the discussions of Citibank's advanced technology group. Caltech showed "how;" the bank showed "why." The ideas came together in the winter and spring of 1979. Founders worked "out of the box" to design a revolutionary database management system for parallel processing with multiple microprocessors, specifically for decision support. The microprocessors, like horses, were harnessed to pull a large load rather than relying on one horse to do all the work.
1979	Teradata was incorporated on July 13, 1979, and started in a garage in Brentwood, Calif. The name Teradata symbolized the ability to manage terabytes (trillions of bytes) of data. The founders spent the next year raising financing. A seed round of \$150,000 closed in March 1980, which allowed design and patent work to begin in earnest.
1980	The first VC round, for \$2.5 million, closed in July 1980, which let the company hire the R&D team.
1983	The first beta system was shipped to Wells Fargo Bank just in time for Christmas.
1986	Fortune magazine named Teradata "Product of the Year." First profitable quarter (ending June).
1987	Teradata IPO in August.
1989	Teradata Corporation partnered with NCR Corporation to build the next generation of database computers.
1991	NCR was acquired by AT&T on September 19, 1991. In December of 1991, NCR announced its acquisition of Teradata.
1992	The first system over 1 terabyte (a trillion bytes) went live at Wal-Mart in January.
1994	Gartner named Teradata the "Leader in Commercial Parallel Processing."
1995	IDC consulting group named Teradata number one in MPP (massively parallel processing) in Computerworld magazine.
1996	A Teradata database was world's largest database with 11 terabytes (11 trillion bytes) of data. Gartner commented "...NCR's Teradata V2 has proven it can scale..." and The Data Warehouse Institute presented Teradata with its Best Practices Award in data warehousing.
1997	AT&T spun off NCR and, on January 1, 1997, NCR became an independent, publicly traded company once again.
	Teradata introduced its first event-driven marketing solution, "Relationship Optimizer," a customer relationship management (CRM) software application co-developed with a major banking customer.
	The database of one of Teradata's customers was the world's largest production database with 24 terabytes (24 trillion bytes) of user data. Teradata received The Data Warehouse Institute's Best Practices Award and DBMS Readers' Choice Award.
1998	Teradata was ported to Microsoft Windows NT.
1999	The database of one of Teradata's customers was the world's largest database in production with 130 terabytes (130 trillion bytes) of user data on 176 nodes.
	NCR consolidated its data warehousing business into a separate operating division.
2000	Teradata's first enterprise-class application for detailed customer profitability measurement, Value Analyzer (VA), was launched at Royal Bank of Canada, where it quickly became a documented success. Within nine months, VA was adopted by ten Teradata customers.
	NCR acquired Ceres Integrated Solutions, a provider of customer relationship management (CRM) software, reconfiguring the tool into a new product named Teradata CRM.
	Teradata acquired business partner Stirling Douglas Group of Canada, adding its Demand Chain Management (DCM) software to its growing family of enterprise-class analytical applications. DCM ensured efficient real-time inventory management via sophisticated forecasting capabilities leveraging detailed data from the Teradata Warehouse.
2001	Teradata more than doubled the lines of code (1.6 million to 3.8 million).

	Teradata introduced its Financial Management (FM) solution, an analytic architecture made up of hardware, software, professional consulting and support services.
2002	Teradata launched Teradata Warehouse 7.0. It was the first time in the history of data warehousing that any vendor extended decision making beyond corporate management to all functions across the organization, optimizing operational and strategic decisions.
2003	More than 120 industry-leading companies migrated from Oracle to Teradata after the launch of the Oracle-to-Teradata migration program.
	Teradata University Network was created to advance awareness of data warehousing in the academic community. Nearly 170 universities from 27 countries were represented in the network.
	Teradata Customer Relationship Manager (CRM) Version 5.0 was released.
2004	Teradata marked its 25th anniversary.
	Teradata and SAP, the world's leading provider of business software solutions, announced a technology partnership agreement to deliver analytic solutions to industries with high data volume requirements.
	Teradata and Siebel Systems, Inc. (now Oracle Business Intelligence), a leading provider of business applications software, announced a strategic partnership to immediately make available integrated and optimized products.
	According to InformationWeek (September 27, 2004), Wal-Mart's "single, centralized 423-terabyte Teradata system is Wal-Mart's competitive advantage...enabling the same data set for both buyers and suppliers."
	Teradata shipped the enhanced Teradata® Warehouse Miner 4.0 data-mining software and Teradata Profiler data-mining tool.
2005	Teradata launched Teradata Warehouse 8.1.
	Teradata acquired DecisionPoint Software, rebranding it Teradata Decision Experts.
	Teradata brought Linux operating system to enterprise-class data warehouses in addition to the existing choices of Windows® and its own UNIX operating system, UNIX MP-RAS.
2006	Gartner positioned Teradata in the Leaders Quadrant for multichannel campaign management.
	Teradata named one of the "Elite Global Information Technology Providers."
	Teradata completed SAP® integration certification.
	Teradata launched Enterprise Master Data Management (MDM) solution.
	Teradata Warehouse 8.2 expanded the ability of customers to deliver real-time intelligence with faster performance, a significant volume of concurrent queries, high system availability, event monitoring, easy integration into the enterprise, simplified system management and unparalleled support of short operational as well as long strategic workloads.
	Teradata shipped the NCR 5450 Server, which featured a boost of 18 percent in the price/performance ratio over previous models. It could coexist with six previous generations of servers and be operational in just hours.
	Teradata Enterprise Risk Management solution was announced for financial services to provide centralized infrastructure and analytic capabilities to identify best opportunities for profitable growth while meeting regulatory requirements with confidence.
	Microsoft Corporation and Teradata formed strategic alliance.
2007	NCR announced intention to separate into two independent companies to create two market-leading companies, NCR and Teradata, each focused on different businesses.
	Teradata named best global data warehouse-business intelligence appliance vendor by Intelligent Enterprise magazine.
	Teradata and DFA Capital Management Inc. announced partnership.

	Agilent Technologies and Teradata established first-of-its-kind partnership to integrate network and customer data for telecommunications industry.
	Teradata® Warehouse Miner 5.1 launched, enabling data-mining technology from multiple vendors to be embedded and run within the Teradata Enterprise Data Warehouse, helping businesses to predict the future. The open analytic environment expanded data-mining technology options while leveraging the Teradata database for in-database mining.
	Teradata announced the availability of its Teradata 5500 Server, using approximately 75 percent less energy and coexisting with multiple generations of Teradata servers, thereby protecting the technology investment of customers.
	More than 850 universities from nearly 70 countries were represented in the Teradata University Network.
	Teradata completed its spinoff from NCR Corporation, and shares of its common stock began regular trading on the New York Stock Exchange (NYSE) under the symbol "TDC" on Monday, October 1, 2007.
	Teradata launched Teradata 12, an innovative, advanced database delivering traditional data warehousing for strategic planning, along with usable intelligence to frontline operations throughout the enterprise.
	Teradata announced a key strategic partnership with SAS involving deeper technical integration of their respective products and coordinated marketing, sales and services activities.
	Teradata announced the availability of Teradata Relationship Manager, its newest marketing automation software portfolio.
2008	Teradata introduced a new family of platforms addressing many customer needs, from entry-level to active enterprise-level, all with the proven power of the Teradata 12.0 database engine. The new, cost-effective platform family includes the Teradata 550 SMP (symmetric multiprocessing), a departmental data warehouse; the Teradata 2550, an entry-level data warehouse; and the Teradata 5550, an active data warehouse-class platform.

Exhibit Four: Introduction to Teradata Parallel Processing  
 (Diagrams courtesy of Teradata Corporation)

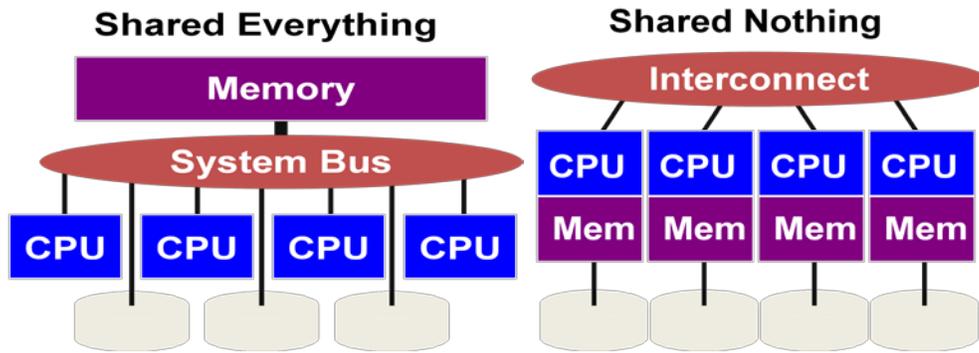


Exhibit 4a: Shared Memory Alternatives

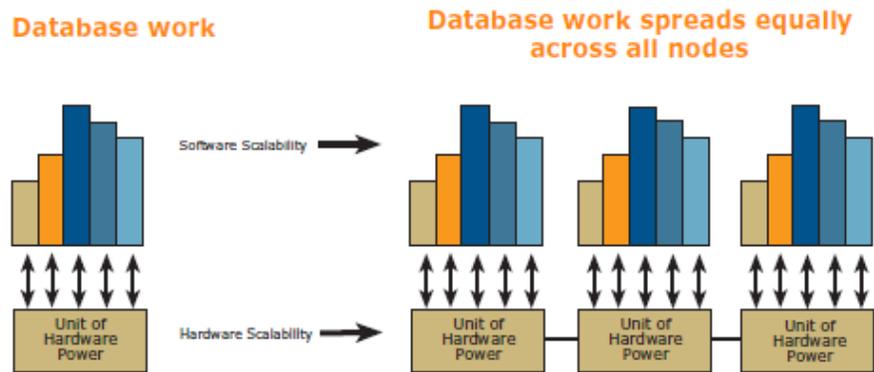


Exhibit 4b: Spreading the Workload

Exhibit Five: BI Value Streams  
 (Diagrams Courtesy of Teradata Corporation)



Exhibit 5a: Insight to Action Value Stream

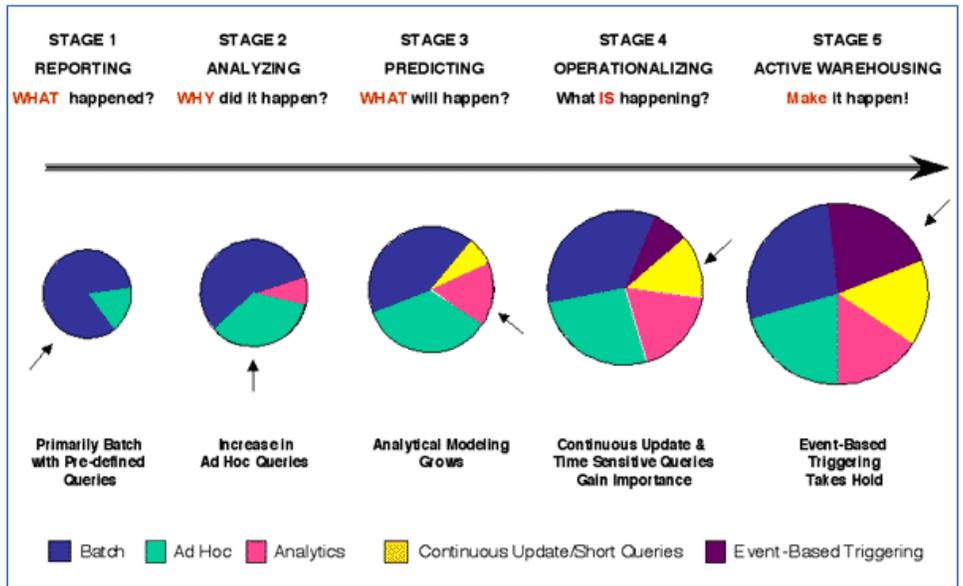


Exhibit 5b: BI Value Stream Evolution

**Exhibit Six: Teradata Accelerate**  
(Tables and Diagrams Courtesy of Teradata Corporation)

	TARGET CUSTOMER PROFILE	TIME	STARTING PRICE
<b>FOUNDATION PACKAGES</b>			
<b>Teradata Accelerate for Do-It-Yourself</b>	Companies looking to deploy a proven analytical platform that are skilled in defining and building data warehouses. Professional Services assistance is included to jump-start the initial data load and build sample queries.	~ 10 days	\$350,000
<b>Teradata Accelerate for Load and Go</b>	Organizations that want to load data into a data warehouse to get quick access to analytics for business users. This production environment will support ongoing business analysis.	~ 30 days	\$440,000
<b>Teradata Accelerate for Application Re-Host</b>	Companies that want to re-host data and applications from existing data marts or operational reporting systems on a more powerful, proven analytical platform.	~ 120 days	\$550,000
<b>Teradata Accelerate for Data Warehousing</b>	First-time data warehouse users that want an affordable, low-risk data warehouse in a rapid timeframe.	~ 120 days	\$830,000
<b>APPLICATION PACKAGES</b>			
<b>Teradata Accelerate for Finance</b>	Finance organizations that seek a more flexible reporting environment through accelerated data sourcing. Pre-built, finance-oriented data models are included to quickly transform enterprise resource planning data into meaningful information.	~ 70-80 days	\$700,000
<b>Teradata Accelerate for Marketing</b>	Marketing departments that want to segment customers, personalize customer communication and define collateral. The departments will be able to execute campaigns and perform analytics on customers and campaigns to understand program effectiveness.	~ 70 days	\$670,000
<b>INDUSTRY PACKAGES</b>			
<b>Teradata Accelerate for Demand Signal Repository</b>	Customer goods manufacturers wanting to deploy an analytic environment to better understand customer buying behavior.	~ 60 days	\$400,000
<b>Teradata Accelerate for Gaming</b>	Gaming companies that want to deploy a business intelligence (BI) solution.	~ 90 days	\$600,000
<b>Teradata Accelerate for Retail</b>	Retail organizations that want to enhance their data environment with intelligent decision-making applications in the areas of sales and inventory, assortment planning, market-basket analysis and store-performance management.	~ 90 days	\$640,000

Figure 6a: Accelerate Program Bundles

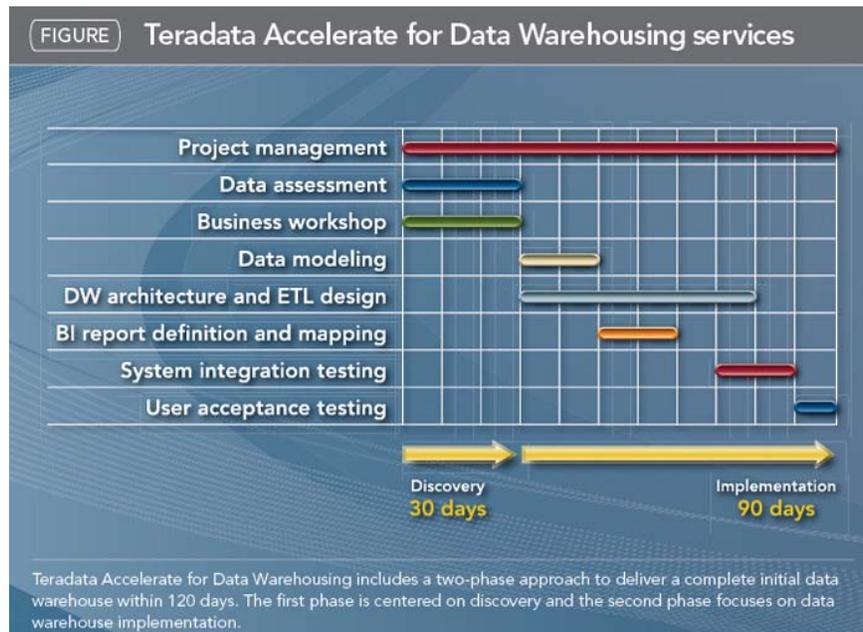


Figure 6b: Example Accelerate Program Services & Timeline

- ✓ *Teradata Database 12.0 — the market-leading data warehousing database with built-in functionality to provide for easy deployment and manageability*
- ✓ *One member of the Teradata Purpose-Built Platform Family — Teradata Active Enterprise Data Warehouse, Teradata Extreme Data Appliance, Teradata Data Warehouse Appliance or Teradata Data Mart Appliance*
- ✓ *Teradata software — tools and utilities to make it easy to rapidly load and query data*
- ✓ *Teradata Professional Services — consulting and support services to quickly get up and running*
- ✓ *Teradata training — supports users as they learn to get the most value out of their data warehouse investment*

Figure 6c: Example Components Bundled as Part of a Teradata Accelerate Package

## Exhibit Seven: Interview with the Chief Technology Officer and the Chief Development Officer

*“Independence accelerates innovation”*

### A discussion with Teradata visionaries Stephen Brobst and Scott Gnau.

By Anne Wainscott-Sargent

On Oct. 1, Teradata marked its first anniversary as an independent company. It's been a year of major technology and partnership milestones. Recently, *Teradata Magazine* sat down with Stephen Brobst, chief technology officer, and Scott Gnau, chief development officer, to explore key innovation highlights from the past year and what Teradata customers can expect in the upcoming year.

#### **Q: What have been the key technology innovation milestones since Teradata was spun off from NCR?**

**Brobst:** Since the spinoff, we delivered a major release of the Teradata Database—Teradata 12.0 at the end of last year, which had significant performance gains from optimizer enhancements. And Teradata 13.0 was announced this year with significant new features, especially in the area of storage virtualization. I consider this accomplishment a demonstration of the acceleration of our innovation. In addition, we now have a multi-product family that ranges from departmental data marts using SMP [shared-memory parallelism] technology all the way up to active enterprise data warehousing appliances at the high end of the business intelligence [BI] spectrum. Previously, we covered this range of requirements with a single product line using a single form factor. Now, we have many form factors to deliver at a variety of price-performance and functionality points.

**Gnau:** We certainly have developed two of our most feature-rich releases ever, and we've done it back to back. This has required a lot of innovation in areas that are not traditionally sexy areas for software developers, like the process we use for code integration, code and unit test, and all the back-office operations that have allowed our developers to be more productive and get more innovative content into the release.

With Teradata 12.0, we reinvented ourselves in terms of query optimization; we took something that is widely regarded as the leader in the industry and made it better. When I say "better," I mean the combination of cool technology targeted at a sweet spot for our customers in terms of query optimization. These innovations help not only sophisticated users of Teradata but also dramatically improve integration with tools and applications and allow Teradata analytics to scale even better for the masses.

A big breakthrough in innovation with Teradata 13.0 is storage virtualization and the enhancement of our multi-temperature offer. We also are setting the stage for more heterogeneous storage subsystems—and being able to tailor service level and performance with market requirements. What's innovative is how we have integrated this capability into the product and how it will be automatic, easy to manage and will provide not only better performance but also enhanced flexibility for customers. These are hallmarks of Teradata technology.

Adding Linux as an operating system to our lineup provides great benefit to our customers in terms of overall performance improvement and robustness of feature and function. This is a whole new paradigm for us—leveraging a quality, open-source operating system—to benefit not only our customers but also the larger Linux user community.

#### **Q: What is Teradata's position as an appliance market player?**

**Brobst:** Speaking historically, we were the first company to offer an appliance—Teradata has always delivered a purpose-built solution. We've now broadened the appliance spectrum with additional family members. The Teradata Data Warehouse Appliance 2550 is a new version of our entry-level data warehouse/data mart appliance, and it has even better price performance than the 2500. Not only did we introduce the family, but we also introduced new technology into the family at a much faster rate than any of our competitors.

**Gnau:** The Teradata Data Warehouse Appliance 2550 includes key technology to optimize for this space to make Teradata a premier player at the entry level, while also providing our Teradata Data Warehouse Appliance customers with options to grow into a full-scale EDW [enterprise data warehouse] over time. It is a great entry point based on core technology but also provides a roadmap into the best-integrated, real-time platform in the industry.

**Q: What are the key new performance features in the new release?**

**Brobst:** Essentially, we increased I/O performance and workload management capabilities as well as what we call "clicking" capability for extra-high availability. The analytical capability of data warehouses is becoming more and more mission-critical within organizations. It's no longer acceptable to have "optional availability" from the data warehouse. People expect the data warehouse to be open for business at all times—more and more, data warehousing is a 24x7 endeavor.

**Q: What other innovations do you want to highlight?**

**Brobst:** In addition to these new releases, one of the biggest innovations of the past year is our joint partnership with SAS for in-database processing of advanced analytics. Historically, the advanced analytics had to be performed using a proprietary file system with a lot of data movement and duplication from a data warehouse to disparate file systems. Now, the advanced analytics are executing directly inside the database. That means a much quicker turnaround time on advanced analytics and a much lower cost, because you don't have to have shadow copies of all the data files.

We're working with SAS to build extensions to the Teradata Database to encapsulate these advanced analytics inside the parallel execution engine. This has not previously been possible in relational database technology. Teradata is the only relational database that has these capabilities with SAS.

**Q: Open-source technology is definitely coming into its own within the industry. Can you tell us more about Teradata's new Viewpoint portal?**

**Brobst:** Teradata Viewpoint is a portal framework that leverages open-source technology to allow customers and third-party integrators to build their own portals. Teradata Viewpoint has portals to measure performance, monitor systems health and keep track of system events such as when the next backup is scheduled.

We are in the process of setting up a "developer's zone" that will allow software engineers to share software written within the Teradata Viewpoint framework. Our goal is to leverage the open-source framework to accelerate the delivery of functionality through the contribution of our many smart customers, partners and professional services consultants.

**Gnau:** The developer's zone will launch in 2009. It is an extension of our whole passion around the accessibility of our solution. In areas where it makes sense, we want to provide toolkits and allow innovation to grow from anywhere it can, so partners and customers can take advantage of the framework we built. In the spirit of the whole open-source community approach to innovation, we're looking to provide a place where they can share ideas, learn from each other and reuse assets. Ultimately, we think that will turn into an advantage for Teradata, making it easier to integrate and making the surrounding tool ecosystem much better and much more effective than if we just tried to go it alone.

**Q: Do you feel Teradata is a more interesting place for programmers these days with all this diversity happening in the product family?**

**Gnau:** Absolutely. It's not just the new platform family, but it's all the extensions to the functionality and accessibility of the database, making our programmers think more holistically about very complicated systems. That creates an optimal environment for intellectually capable people. And this is the key to really deploying valuable

solutions, not just reporting what happened but allowing our customers to drive change in their own industries by knowing the details and predicting disruption.

What's most exciting for me is seeing ideas being formulated and working in that creative process where people are just flat out inventing and thinking of things that haven't been thought of before. It becomes even more exciting when this innovation turns into a product that we bring to market, like our new appliances this year. All contain significant amounts of innovation, they target a new kind of audience and they actually contain software content and inventions that were created here.

**Q: What can customers look forward to in the area of innovation at Teradata?**

**Brobst:** Going forward, our customers should expect to see more additions to the Teradata platform family. We'll use a variety of form factors to deliver different price-performance points to meet the analytics requirements within an enterprise. We will also be expanding our delivery of advanced data types and enhancing the virtualization of Teradata's data warehousing capability. A lot of cool stuff is coming in the workload management arena as well.

**Gnau:** We continue to be motivated by a healthy paranoia, and we are supported by some tremendous breakthroughs in the generic computing environment, like multi-core growth within the microprocessors, like solid state disk devices and like blindingly fast bandwidth in interconnects. If you take a motivated group of really smart people who have a little paranoia and you give them tools like that from the industry, they're going to be able to think of things that even a year or two years ago would have been impossible. Certainly, new products, offers and partnerships like Teradata Viewpoint and the SAS integration demonstrate how we are expanding our scope in ways that would have been considered non-traditional for Teradata just a year or two ago. **T**

Anne Wainscott-Sargent is a writer/editor and PR consultant based in Atlanta, specializing in the health, government, data warehousing and wireless/satellite communications markets.

*Teradata Magazine-December 2008*

## End Notes

---

- <sup>1</sup> "Teradata Announces 2008 Fourth Quarter and Full-Year Results," Reuters News, February 12, 2009, <http://www.reuters.com/article/pressRelease/idUS155581+12-Feb-2009+BW20090212> , Accessed April 21, 2009.
- <sup>2</sup> The history of Teradata can be found at <http://www.teradata.com/t/page/42649/> or contact Mike O'Sullivan, Ph: 937-242-4786, [mike.osullivan@teradata.com](mailto:mike.osullivan@teradata.com).
- <sup>3</sup> Teradata: Confident and Competitive, <http://www.ventanaresearch.com/blog/commentblog.aspx?id=2841>, Accessed 10/14/08, Ventana Research.
- <sup>4</sup> Hagerty, J., AMR Research, Quoted by Daniel, D., "Business Intelligence Mergers and Acquisitions: Has BI Grown Up?" *CIO.com*, November 14, 2007.
- <sup>5</sup> Cusumano, Michael A. and Annabelle Gawer, "The Elements of Platform Leadership," *MIT Sloan Management Review*, Vol. 43, No. 2, Spring 2002, 51-58.
- <sup>6</sup> "Data Warehouse Appliance," *Wikipedia*, [http://en.wikipedia.org/wiki/Data\\_warehouse\\_appliance](http://en.wikipedia.org/wiki/Data_warehouse_appliance), Accessed 1/20/2009.
- <sup>7</sup> Morgan, T. P., "SUSE Linux 10 is the First 64-Bit Platform for Teradata Data Warehouses," *The Linux Beacon*, September 19, 2006, <http://www.itjungle.com/tlb/tlb091906-story06.html> , Accessed 4/21/2009.
- <sup>8</sup> "Teradata Delivers Linux for Enterprise-Class Data Warehouses," <http://www.teradata.com/t/page/156964/index.html> , 9/18/2006, Accessed 1/20/2009.
- <sup>9</sup> Interview with Scott Gnau, Vice President and General Manager, Research and Development, Teradata, 07/27/2008.
- <sup>10</sup> Feinberg, D and M.S. Beyer, *Magic Quadrant for Data Warehouse Database Management Systems*, Gartner Core Research Note G00138797, 9/12/2006.
- <sup>11</sup> Gary, Loren, "Ambidextrous Innovation," *Harvard Management Update*, April 2003, Article Reprint Number U0304B.
- <sup>12</sup> Meyer, M. H. and Seliger, R. "Product Platforms in Software Development," *Sloan Management Review*, Fall 1998, V. 40, no. 1.
- <sup>13</sup> *Ibid*, p. 62.
- <sup>14</sup> "Grosch's Law," *Wikipedia*, [http://en.wikipedia.org/wiki/Grosch's\\_law](http://en.wikipedia.org/wiki/Grosch's_law) , Accessed 1/20/2009.
- <sup>15</sup> Neches, P., D. Hartke, R. Baran, D. Woodcock, A. Papachristidis, "Dynamically Partitionable Parallel Processors," *United States Patent Number 4,925,311*, May 15, 1990, <http://www.google.com/patents?id=KwnkAAAAEBAJ> , Accessed 1/20/2009.
- <sup>16</sup> Fisher, Lawrence M. "An Old-Timer Bets on a New Line," *New York Times*, September 18, 1990. <http://query.nytimes.com/gst/fullpage.html?res=9C0CEFD1030F93BA2575AC0A966958260> , Accessed 1/20/2009.
- <sup>17</sup> DeWitt, David J. and Jim Gray, "Parallel Database Systems: The Future of High Performance Database Systems," *Communications of the ACM*, Vol. 35, No. 6, 1992, p. 85-98.
- <sup>18</sup> Interview with Jim Dietz, Product Marketing Manager for Platforms, Teradata, November 11, 2008.
- <sup>19</sup> Ebel, D., O. Naik, R. Fowler and D. Patriarca, "System and Method for Planning and Implementing a Data Warehouse Solution," *United States Patent Number 7,092,968*, August 15, 2006, <http://www.google.com/patents?id=XrF6AAAAEBAJ> , Accessed 1/20/2009.
- <sup>20</sup> Reuters, "Active Enterprise Intelligence Gains Momentum Among Leading Companies," September 26, 2008, <http://www.reuters.com/article/pressRelease/idUS217037+26-Sep-2008+BW20080926> , Accessed 1/20/2009.
- <sup>21</sup> Walter, T. "Teradata 13.0: A Virtual Performance," *Teradata Magazine Online*, <http://www.teradata.com/tdmo/v08n04/Tech2Tech/AskTheExperts/Teradata13.aspx> Accessed 1/23/2009.
- <sup>22</sup> Hargett, K., S. Brobst, P. Lessing, W. Matera, M. Crosby and L. Young, "Method for Developing Data Warehouse Logical Data Models Using Shared Subject Areas," US 7,359, 906 , April 15, 2008,

---

<http://www.google.com/patents?id=9QupAAAAEBAJ&dq=developing+data+warehouse+logical+data+models> , Accessed 1/20/2009.

<sup>23</sup> Interview with Barb Vostmyer, Senior Manager for Industry Solutions, Teradata, 12/3/2008.

<sup>24</sup> Ibid.

<sup>25</sup> Miller, T., B. Tute and A. Rollins, “Analytical Logical Data Model,” United States Patent Number 6,553,366, April 22, 2003,

<http://www.google.com/patents?id=9xgPAAAAEBAJ&dq=analytical+logical+data+model> , Accessed 1/20/2009.

<sup>26</sup> Martin, James, *The Great Transformation*, (Amacom: New York), 1995.

<sup>27</sup> Whittle, R. and C. Myrick, *Enterprise Business Architecture: The Link Between Strategy and Results*, (Auerbach: Boca Raton, FL), 2005.

<sup>28</sup> Winter, R. “EBay Turns to Analytics as a Service,” InformationWeek.com, October 11, 2008,

[http://www.informationweek.com/news/business\\_intelligence/warehouses/showArticle.jhtml?articleID=210800736](http://www.informationweek.com/news/business_intelligence/warehouses/showArticle.jhtml?articleID=210800736) Accessed 1/23/2009.

<sup>29</sup> Imhoff, C. “Teradata Partners Users Conference Very Enlightening and Entertaining,” *Blog Archive*, September 21, 2005, [http://www.b-eye-network.com/blogs/imhoff/archives/2005/09/teradata\\_partne.php](http://www.b-eye-network.com/blogs/imhoff/archives/2005/09/teradata_partne.php), Accessed 1/20/2009.

<sup>30</sup> Interview with Stephen Brobst, Chief Technology Officer, Teradata, September 9, 2008.

<sup>31</sup> *The Beige Book*, The Federal Reserve, January 19, 2009,

<http://www.federalreserve.gov/FOMC/BeigeBook/2009/20090114/default.htm> , Accessed 1/20/2009.

<sup>32</sup> Everett, C., Business Intelligence Tools Up for the Recession, ZDNet UK, 12/1/2008

<http://www.zdnetasia.com/news/software/0.39044164.62048860.00.htm> , Accessed 1/20/2009.