ORACLE
VALUE CHAIN PLANNING

ORACLE
ENGINEERED SYSTEMS

An Oracle White Paper
March 2013

Oracle In-Memory Consumption-Driven Planning

Overview and Business Value: Move Planning Closer to the Consumer

ORACLE

Disclaimer

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Table of Contents

Executive Overview

With the increasing availability of Point of Sale (POS) and other downstream data over the last decade, companies selling through retail channels have strived to use such data to drive their upstream supply chains. The availability of granular data, down to the item, store, and day level, offers the promise of transforming business processes in areas such as demand, replenishment, promotion, and new product planning. However, limitations in computing power have precluded taking full advantage of the huge amounts of data involved. Oracle's new product, Oracle In-Memory Consumption-Driven Planning (CDP), provides a solution. Based on an optimized version of Oracle's Demantra product line combined with and designed for the Oracle Engineered Systems platform, CDP offers an extreme step-change in performance and new functionality. Oracle Engineered Systems are the preferred platform for deploying Oracle Applications when performance and scalability are critical, because of the extensive performance optimizations in engineering of hardware and software to work together, and the innovations that are only available with the complete Oracle technology stack. Using the Oracle Exadata Database Machine, and optionally, the Oracle Exalogic Elastic Cloud or SPARC SuperCluster, CDP has been engineered from the ground up to provide unparalleled performance and enable business processes that would otherwise not be feasible. This paper outlines the capabilities and directions of CDP, applicable to multiple industries such as Consumer Packaged Goods and Durables, Consumer Electronics, Media and Entertainment, Automotive and Communications.

Introduction: The Demand-Driven Value Chain Vision

Over the last decade, many companies that sell through retail channels have been pursuing the vision of the demand-driven value chain, in order to drive their upstream supply chains based on actual end-customer demand – consumption, or sell-through to consumers – rather than shipments, or sell-in to Distribution Centers (DCs) and stores. This approach promises to eliminate the well-documented bullwhip effect, in which lack of visibility to downstream demand and inventory positions propagates upstream and leads to volatile order quantities as well as inventory build-ups and shortages. By monitoring sell-through rather than just sell-in, and minimizing the latency between actual consumer demand and its measurement, companies can gain a more accurate understanding of market conditions and adjust their production and distribution plans accordingly.

By deriving a sell-in forecast based on the sell-through forecast (taking into account channel inventory and lead times) companies can obtain a more accurate forecast of sell-in versus the traditional practice of forecasting sell-in directly. The lack of visibility resulting from observing only shipments can be eliminated and a direct relationship between consumer behavior and sell-in shipments can be established. With a more accurate forecast, companies can improve on-shelf availability, increase sales, increase turns, and decrease supply chain costs. Figure 1 below indicates companies' rationale for wanting to leverage downstream data, with improving forecast accuracy and higher customer service levels leading the list.

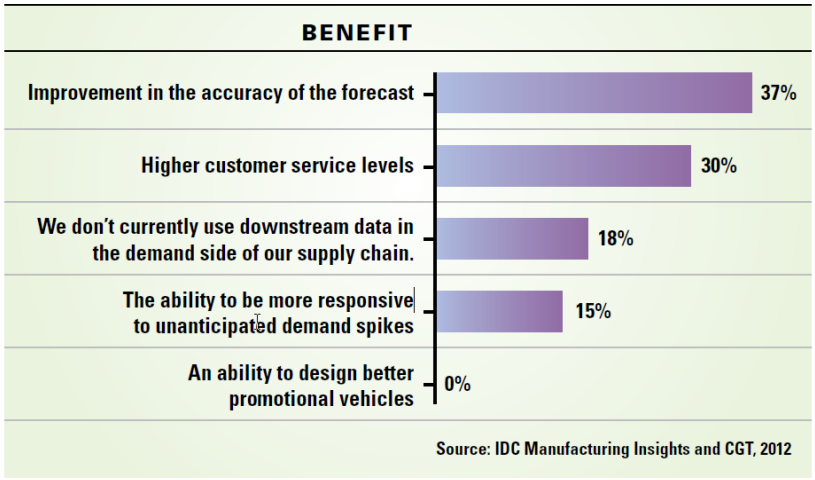

Figure 1: Benefits of Using Downstream Data in the Demand Side of Supply Chain (percent of respondents)[1]

[1] Source: 2012 Sales and Marketing Report, *Consumer Goods Technology*

Such a demand-driven approach is increasingly important in a business environment characterized by demand volatility, shorter product lifecycles, extended global supply chains, intense cost pressure, and (for manufacturers) larger and more demanding retailers. For obvious reasons, manufacturers of consumer electronics have been particularly interested in the demand-driven approach, and in North America they typically monitor POS data for retail accounts constituting over 80% of sales. Other companies that sell through retail channels are also interested in the demand-driven approach, including companies in the consumer goods, media & entertainment, and mobile telephony industries, as well as retailers and hybrid retailer-manufacturers. Figure 2 shows how usage of downstream data has increased significantly between 2011 and 2012, indicating the growing interest in become more demand-driven.

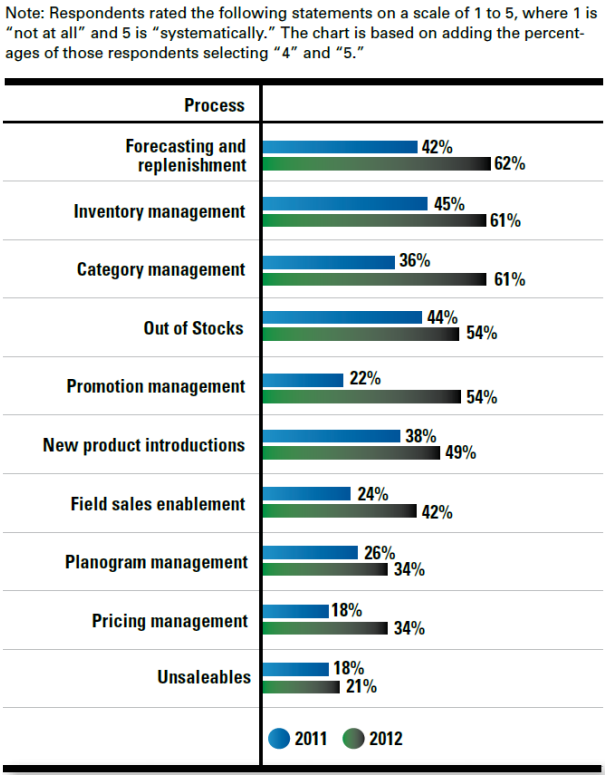

Figure 2: Responses to the Question, "How much are you using downstream data to power each process?" [2]

[2] Source: Retailer/Supplier Shared Data Study 2012, *Consumer Goods Technology / Retail Info Systems News*

The use of POS data by manufacturers has evolved over time, as illustrated in Figure 3. Initially, POS data was used primarily on an account by account basis, at the initiative of individual account teams, using separate decentralized systems. As companies increasingly recognized the value of downstream data for enterprise use, many companies started using aggregated data, typically at the DC and weekly levels, for enterprise planning purposes such as forecasting and replenishment. More granular use was precluded by limited data availability, poor data accuracy and the high cost of the computing power needed to manage such large volumes of data.

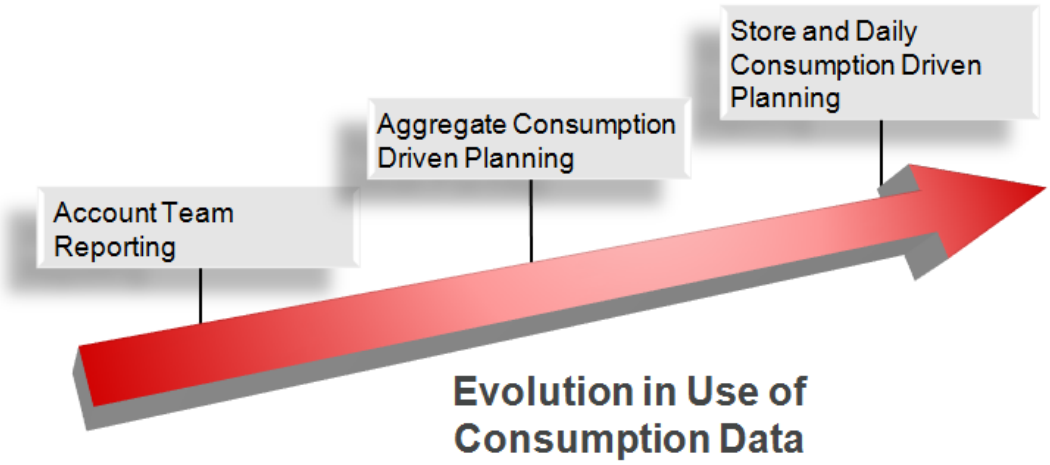

Figure 3: Evolution in Use of Consumption Data

With increasing availability and accuracy of store and daily data from retailers such as Walmart and Target and improvements in information technology, early adopters are starting to plan at a store and daily level. Their reasons are typically one of the following:

- They (manufacturer or retailer) are already doing store-level replenishment or shipping to retail DCs for cross-docking. In the past perhaps they were forecasting based on store-level shipments or orders, or all planning was done at an aggregate DC level. But with the availability of store-level POS data and the ability to create accurate store-level forecasts and replenishment plans, companies can create superior plans, optimally trading off service and inventory levels. In addition, the move from weekly- to daily-level planning can further increase forecast accuracy when reacting to promotions and holidays and optimize replenishment plans, particularly if there are multiple store deliveries per week.
- They are not doing store-level replenishment, but want to improve DC-level forecasts based on insights and information available only at the store level, including consumer behavior and store inventory. Demographic-based behavior is most meaningful at the store level and is required for

accurate modeling in a highly dynamic market. Store-level inventory is necessary to accurately forecast shipments. By deriving a DC-level forecast based on store-level forecasts and inventory plans, DC-level forecasts, assortments and inventories can be improved. The availability of daily demand and inventory levels greatly contributes to the ability to react to consumer behavior and ensure DC planning achieves better service levels.

Growing Data Volumes Present Challenges

A small number of early adopters have been doing store and daily-level planning for some time, using Oracle software and that of other vendors. Yet the massive increase in data involved (7 times going from weekly to daily, perhaps up to 100 times going from DC to store-level), particularly for a large number of stock-keeping units (SKUs) and stores, presents significant challenges. Figure 4 illustrates how data volumes might scale, in terms of transactions (Tx, meaning number of item-location-time periods) going from traditional DC-weekly systems to store-daily systems.

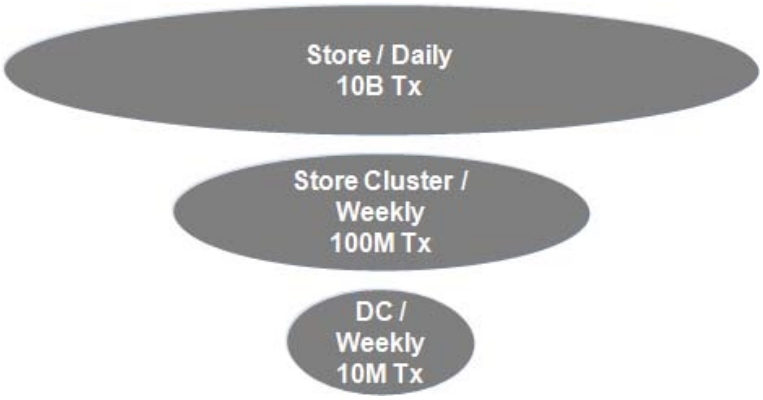

Figure 4: Data Volumes Grow with Increasing Data Granularity

Making such systems workable when moving from 1X to 100+X data scale is a daunting task, which for most companies has been a barrier to achieve value. It requires a high level of technology expertise and investment to scale to this level with traditional hardware and software architectures. The only way to get there has involved extensive customizations and the use of multiple application and database systems. In addition, the large scale of data will have substantial impact on the ability to have such a solution serve as an integrated system to also support a demand planning or sales and operations planning (S&OP) process. Application functional configurations that may be routine at a smaller scale may require more careful consideration of performance implications and may not be feasible. Standard data and batch processes will be lengthy, and it may be difficult to complete them during off-hours (and perhaps impossible for global businesses operating around the clock). Ongoing system maintenance is likely to require more effort and have a large impact on IT budgets. For these reasons most companies that have embarked on this journey end up taking a costly custom application route which limits the success, value and quality that can be achieved by professionally engineered applications.

Oracle Engineered Systems Provide a New Approach

A new approach is needed to enable planning at a more granular level. Oracle's new product, Oracle In-Memory Consumption-Driven Planning, leverages the latest innovations in hardware and software technology from Oracle to provide a solution with unparalleled performance and scalability. CDP uses Oracle's industry-leading Demantra Demand Management product as a starting point, enhanced to run on a combination of Oracle Exadata or SPARC SuperCluster and optionally Oracle Exalogic. CDP also includes new out-of-the-box scalable functionality to enable consumption-driven planning and replenishment business processes. The result is a product that can scale to handle the largest planning challenges at a daily- and store-level.

Oracle Engineered Systems provide the extreme step-change in scalability and cost-performance to move granular consumption-driven planning from custom early adopters to the mainstream market. Exadata provides massive data processing throughput, with enough flash memory to store and manage an entire planning data set of a size not feasible with traditional disk-based systems. Exadata also provides unique storage servers integrated with the database that optimizes large dataset processing. Oracle Exalogic with Exabus and Exadata can support hundreds of concurrent Demantra forecast engines running in parallel to solve the largest forecasting problems in minutes instead of days.

By offering a pre-configured bundle of hardware and software, Oracle CDP, Oracle Exadata and Oracle Exalogic reduce implementation time, cost, and risk. Oracle Exadata and Exalogic can be purchased in 1/8, 1/4, 1/2, and full rack configurations depending on data set size and performance requirements to provide the most economical solution. A single purchase of Oracle Exadata can host several Value Chain Planning Solutions together, providing unparalleled integration speeds and reduced maintenance efforts.

Following is more detail on Oracle Engineered Systems and the Exadata and Exalogic hardware.

Introduction to Engineered Systems

Oracle Engineered Systems combine best-of-breed hardware and software components with game-changing technical innovations. Designed, engineered, and tested to work best together, Oracle Engineered Systems can power the cloud or streamline data center operations to make traditional deployments even more efficient. The components of Oracle Engineered Systems are preassembled for targeted functionality and then—as a complete system—optimized for extreme performance. By taking the guesswork out of these highly available, purpose-built solutions, Oracle delivers a solution that is integrated across every layer of the Oracle technology stack—a simplicity that translates into less risk and lower costs for your business. Only Oracle can innovate and optimize at every layer of the stack to simplify data center operations, drive down costs, and accelerate business innovation.

Oracle Exadata Database Machine

Oracle Exadata Database Machine is Oracle's database platform delivering extreme performance for database applications including Online Transaction Processing, Data Warehousing, Reporting, Batch Processing, or Consolidation of mixed database workloads. Oracle Exadata is a pre-configured, pre-tuned, and pre-tested integrated system of servers, networking and storage all optimized around the

Oracle Database. Because Oracle Exadata is an integrated system, it offers superior price-performance, availability and supportability. Exadata frees users from the need to build, test and maintain systems and allows them to focus on higher value business problems.

Oracle Exadata uses a scale out architecture for database servers and storage. This architecture maintains an optimal storage hierarchy from memory to flash to disk. Smart Scan query offload has been added to the storage cells to offload database processing. Oracle Exadata implements Smart Flash Cache as part of the storage hierarchy. Oracle Exadata software determines how and when to use the Flash storage for reads and writes as well as how best to incorporate Flash into the database as part of a coordinated data caching strategy. A high-bandwidth low-latency InfiniBand network running specialized database networking protocols connects all the components inside an Oracle Exadata Database Machine. In addition to a high performance architecture and design, Oracle Exadata offers the industry's best data compression to provide a dramatic reduction in storage needs.

The following Oracle Exadata features are used by CDP:

- **Exadata Smart Flash Cache** uses Flash memory to dramatically reduce the time to read and write database and log records. The intelligence in Smart Flash Cache transparently moves active database blocks from disk to flash in real time, thus ensuring that "hot" data is in Flash memory when the next access occurs. Blocks that should not be in Flash are similarly recognized, maximizing the amount of space in Flash for active data.
- **Exadata Smart Scan** speeds up data-intensive queries by leveraging the processing power of Exadata Storage Servers to scan and filter out results. By moving queries to storage instead of moving the data to the database servers, long-running reports and queries often complete 10x faster than on conventional systems.
- The use of **InfiniBand** as the networking fabric within Exadata ensures the lowest latency for messages and the highest bandwidth for data transfers. High-speed transactions as well as data-intensive queries and reports reap the benefits.
- **Exadata Scale-Out Storage** enables the full performance of Exadata to be realized against large and growing databases, without fear of bottlenecks. As the database size grows and storage capacity is added to Exadata, storage performance and networking bandwidth scale in equal proportion.
- **IORM** (I/O Resource Manager) allocates I/O bandwidth across different applications and databases, based on a prioritized allocation plan, to ensure that the most important applications get the performance they need when they need it.

Oracle Exalogic

Oracle Exalogic is an Engineered System on which enterprises deploy Oracle business applications, Oracle Fusion Middleware or third-party software products. Oracle Exalogic comes pre-built with compute nodes, memory, flash storage and centralized storage, all connected using InfiniBand in a high redundancy architecture delivering five-nines availability, with fault tolerance and zero-down-time maintenance.

Oracle Exalogic dramatically improves performance of Oracle Applications, Fusion Middleware and third party applications without requiring code changes. It reduces costs across the application lifecycle, from initial set-up to on-going maintenance, as compared to conventional hardware platforms. Oracle has made unique optimizations and enhancements in Oracle Exalogic firmware, Oracle Exalogic software, Oracle Fusion Middleware, and Oracle Applications. These include on-chip network virtualization based on near zero latency InfiniBand fabric, high-performance Remote Direct Memory Access, workload management in Oracle WebLogic server and optimizations in Oracle Coherence and Oracle Traffic Director. Oracle Exalogic includes support for a highly optimized version of the Oracle VM, which significantly outperforms comparable virtualization solutions and is an ideal consolidation platform for Oracle Applications. There are templates available to simplify the installation, deployment and configuration of Oracle Applications on Exalogic.

Exalogic is optional for CDP. If used, CDP takes advantage of the following Exalogic features:

- **Oracle WebLogic Optimizations on Exalogic:** Applications using WebLogic benefit from a number of optimizations for thread efficiency, faster interprocess communication and higher message throughput. An optimized work scheduler for Exalogic balances the number of threads per core available on Exalogic systems, providing better application processing efficiency. WebLogic Server has changed to use shared byte buffers instead of array copies when passing data, improving application interprocess communication performance and a 66% reduction in number of objects created. This reduces heap usage and results in fewer expensive garbage collections for Applications. WebLogic also optimizes socket calls to reduce lock contention on Exalogic, allowing fewer threads to process a larger number of message requests.
- **Oracle jRockit Optimizations on Exalogic**: On Exalogic, the JVM makes more efficient network I/O calls using collections of chunked data resulting in higher throughput for Java applications. The JVM optimizes object management with fewer copies resulting in reduced garbage collection and less heap size resulting in better Java application performance.
- **Exalogic Exabus:** Applications running on Exalogic utilize Exabus, the underlying InfiniBand fabric, which provides low latency and high throughput eliminating I/O bottlenecks in every application layer. Applications components are typically deployed in more than one server and Exabus provides low latency for I/O across nodes on same Exalogic rack. Access to ZFS storage device over Exabus greatly reduces latency for log file writes and other file access operations. For applications running on Exalogic and accessing database tier on Exadata, Exabus delivers faster I/O, reduces CPU usage on both the mid-tier and DB-tier and provides higher connection pooling efficiency.
- **Oracle VM for Exalogic:** Exalogic Oracle VM can be used to sub-divide a physical compute node into multiple virtual machines to increase application deployment efficiency while maintaining application performance. Oracle VM has been engineered for tight integration with Exalogic Exabus I/O backplane using a technique called Single Root I/O Virtualization (SR-IOV) ensuring Oracle VM significantly outperforms comparable hypervisors from other leading vendors. The benefit of this approach is unmatched application performance. In an Exalogic configuration, the impact of virtualization on application throughput and latency is negligible.

SPARC SuperCluster

Oracle's SPARC SuperCluster is the world's most efficient multi-purpose engineered system, delivering extreme efficiency, cost savings, and performance for consolidating mission critical applications and rapidly deploying cloud services. Oracle's SPARC SuperCluster represents a complete, pre-engineered, and pre-tested high-performance enterprise infrastructure solution that is faster and easier to deploy than a collection of individual database and application servers. The system combines innovative Oracle technology—the computing power of Oracle's SPARC servers, the performance and scalability of Oracle Solaris, the Sun ZFS Storage Appliance, the optimized database performance of Oracle Database accelerated by Oracle Exadata Storage Servers, and a high-bandwidth, low-latency InfiniBand network fabric—into a scalable, engineered system that is optimized and tuned for consolidating mission-critical enterprise applications

Oracle's SPARC SuperCluster provides both the capacity for growth, as well as the fine-grained server virtualization needed to isolate individual application components. With multiple layers of enterprise application infrastructure consolidated onto a high-performance, highly available SPARC SuperCluster system, deployment speed, application performance, and availability can all be optimized. Designed as a pre-configured, pre-tested, and ready-to-deploy SPARC SuperCluster engineered system, the solution provides a complete and optimized infrastructure solution for applications, built around robust compute, networking, storage, virtualization, and management resources. The result is a system that is orders of magnitude easier to manage, and up to five times faster to deploy than alternatives, all while occupying considerably less real estate requiring less power. Furthermore, the SPARC SuperCluster system provides full built-in redundancy resulting in a highly reliable infrastructure without single point of failure. An issue with one component will not impact other components of the system offering true isolation. Customers can consolidate multiple environments with minimum disruption, without fear of performance degradation, and the ability to achieve required service levels.

CDP Functional Overview

Leveraging the increased performance made possible by Oracle Engineered Systems, Oracle In-Memory Consumption-Driven Planning enables comprehensive, time-phased planning capability at the store and day level. Following is an overview of this functionality.

A Complete Forecasting and Replenishment Solution

CDP offers a complete solution for forecasting and replenishment planning. Oracle Demand Signal Repository (or alternative solution) can be used as a source of harmonized downstream consumption data, and then CDP manages the forecast, calculates safety stock levels, and comes up with a time-phased replenishment plan. The complete cycle is illustrated in Figure 5.

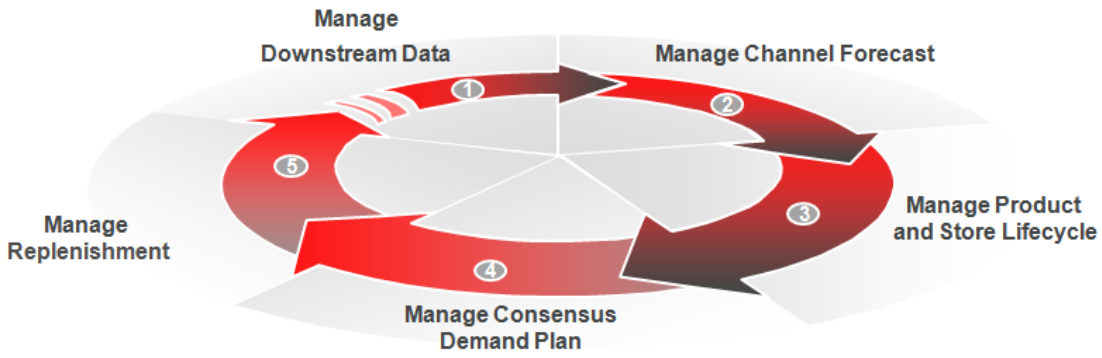

Figure 5: Complete Process Enabled by Oracle Demand Signal Repository (optional) and Oracle In-Memory Consumption-Driven Planning

The five steps are detailed below:

1. **Manage Downstream Data**

Oracle Demand Signal Repository captures, cleanses, validates, and harmonizes POS and other types of downstream data from multiple sources and stores it in the Oracle Database. (Oracle Demand Signal Repository is optional, and customers may use an alternative product for capturing POS data.) Hundreds of out-of-the-box metrics, and dozens of out-of-the-box dashboards and reports, based on Oracle Business Intelligence Enterprise Edition, can be used to analyze the data. The data is then brought into Oracle CDP via a prepackaged integration.

2. **Manage Channel Forecast**

CDP forecasts sell-through demand based on downstream POS data. It then calculates a sell-in forecast by determining what needs to be shipped to maintain target inventory levels, which are in turn based on time-phased weeks of supply targets and the sell-through forecast.

3. **Manage Product and Store Lifecycle**

CDP supports forecasting of new products and phasing-out of old products, with the ability to control the introduction timing and life stages of a product either globally or at varying granularity, including the store level. New products can be based on existing products and can inherit their data and store association. Similar functionality exists for creating new stores and phasing-out stores, as well as allowing for store remodels and resets.

4. **Manage Consensus Demand Plan**

Once a statistical forecast is arrived at it can be seamlessly incorporated into Oracle Demantra Demand Management to ensure a more accurate demand forecast to drive the supply planning process. Consensus planning functionality using Oracle Demantra Real-Time S&OP (optional) enables different organizations, such as Sales, Marketing, Operations, and Finance to come together in a common environment and generate a single number enterprise forecast.

5. **Manage Replenishment**

CDP computes target safety stock, with a choice of 4 different approaches, which can be assigned at product and location levels:

- Statistical calculation based on expected forecast error and desired service level
- Days of supply coverage
- Lead time coverage
- Fixed percent

Using safety stock, sell-in forecast, minimum and maximum inventory constraints, in-transit, and on order quantities, CDP comes up with a recommended time-phased single-tier replenishment plan to feed to enterprise execution systems. Customers desiring a multi-echelon inventory and replenishment plans can alternatively feed CDP forecasts into Oracle Advanced Supply Chain Planning.

Managing Multi-Channel Demand

CDP's flexible hierarchies and configurability enable consumption-driven planning for heterogeneous distribution channels, in which companies are distributing products through multiple channels with different levels of hierarchy and different degrees of data availability, as illustrated in Figure 6. In this diagram, we see a company selling through its own stores, through third party retailers (shipping to both retail DCs and direct to store), and selling though distributors (both with and without sell-through data available). Each of the five scenarios shown in the diagram requires a different set of sell-through to sell-in calculations, represented by the pink arced-arrows. CDP is designed to accommodate such multi-channel calculations.

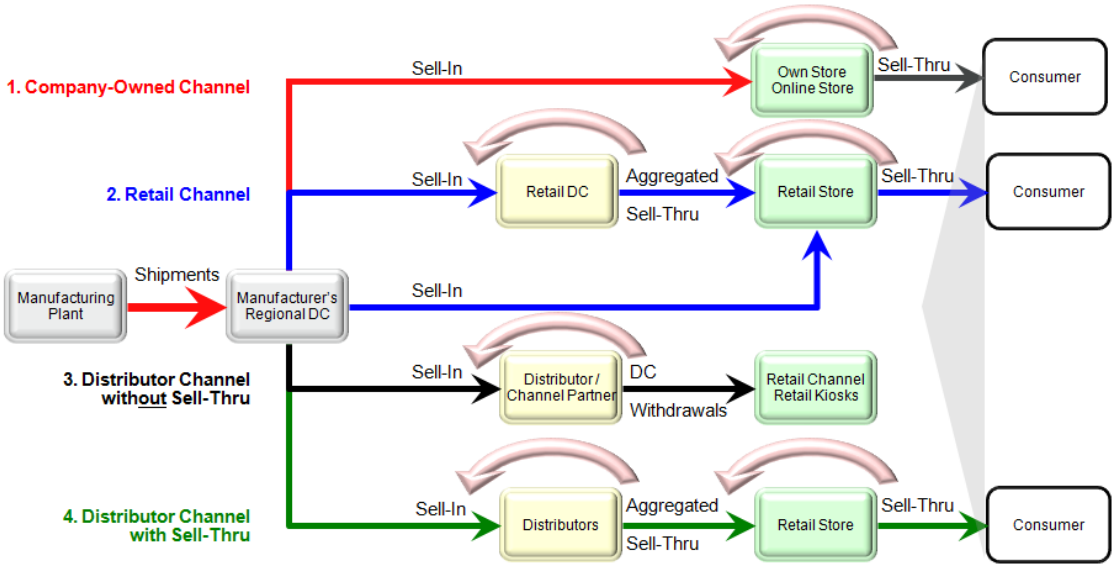

Figure 6: Multi-Channel Demand Management

Granular and Aggregate Planning in a Single System

One of the most important enhancements in CDP is the ability to maintain data at different levels. For instance, if a customer wants to plan 30% of its accounts at the store level using consumption data, but continues with traditional shipment-based DC-level planning with the other 70%, CDP can maintain information at both store and DC levels, so that the performance impact associated with doing daily store-level planning is isolated from the other accounts. In addition to benefiting performance, this feature makes it that much easier to do granular consumption-driven planning and traditional DC-level demand management together in a single system, all rolling up to an enterprise-level plan, eliminating the need to have separate systems and data integrations to consolidate plans.

Business Benefits

Oracle In-Memory Consumption-Driven Planning offers a step change in planning performance and scalability and enables new capabilities not previously possible. The following five business benefits are the result:

Benefit #1: Reduced Planning Cycle Time and Data Latency

The improved performance and scalability translates into tangible business benefits in several ways:

1. **Reduce batch processing time**: Batch processing is critical for planning applications. Whether it is an intra-day, daily or weekly batch process, there is always a critical time constraint in a planning system. For example, the forecast for Asia needs to complete within a certain time window before the Europe plan starts (24x7 business operations). The number of hours available in the batch window is shrinking as businesses take on a larger global footprint but the amount of data processed continues to grow as more data is leveraged to drive various aspects of the business. Oracle Engineered Systems are designed to scale the business without the worry of running out of time to complete the planning cycles. Taking advantage of the new performance characteristics of set-based operations available on Exadata, data processing operations can be achieved in minutes instead of hours. This enables customers to reduce their batch cycle times and at the same time expand their data set to more granular data, or execute multiple planning runs within the same time window. Using the massive scale-out capabilities of the Oracle Exalogic systems, CDP can coordinate hundreds of cores of computing power for planning operations.
2. **Move to more frequent planning cycles**: The highly improved performance from Oracle Engineered Systems and subsequent reduced batch processing time when applied to memory and I/O intensive processes in CDP enable companies to change their planning operations. Planning organizations can now consider moving from weekly to daily planning, from daily to intra-day (shift, store replenishment frequency), enabling them to examine changing conditions in the value chain network faster and prevent or address problems much faster.
3. **Increase the planning scope**: Planning scope can be defined as the product of planning horizon (how many time periods does a customer plan in the future) and planning and data granularity (product versus category level, store versus warehouse). Batch processing times continue to grow as companies expand their solution to new areas of the business or decide to improve accuracy by

examining more granular data.

For example, in demand management granular data can provide more visibility into your customer demand by loading more details regarding customer or geographic attributes. It can also mean moving to more detailed information in the inter-day or intra-day buying patterns. Also, the total number of SKULs (SKUs X distribution sites X sales channels) directly drives the processing time of forecast operations. The larger the SKUL count the more independent time series that need to be modeled. Large systems can have millions of SKULs being managed. Oracle Engineered Systems allow the planning engines to utilize the scale-out compute power of Exalogic Elastic Compute Cloud by distributing the forecast across all the cores of Exalogic Cloud, rather than reducing the number of SKULs modeled in the system. In discrete manufacturing industries, companies which model both end-item forecasts as well as product-option forecasts greatly increase the number of managed time series in the system, and generate improved forecast accuracy of their overall product portfolio including the optional components selected by their customers.

4. **Enable new business processes:** The increase in scalability also enables new business processes that previously were not feasible. CDP supports high-volume intra-day store level forecasting and replenishment processes, as often required in single echelon replenishment for distributors and retail operations. Traditionally, these type of systems are separated from the core unit-forecasting processes to ensure that both processes can leverage the available capacity in memory and I/O to complete within the desired times and do not adversely impact each other. Typically, separate database and middle tier instances would be configured and tuned to handle the loads. With Oracle Engineered Systems, companies can now operate a single model and bring all data together for enterprise-wide analysis and decision making. In these unified environments the need for data integration and synchronization which occurs in separated systems is eliminated with each report accessing and presenting all relevant pieces of information.

To summarize, the improvements in granularity provide critical visibility to the nature of the business and give insight to planning engines and end-users. Oracle CDP allows companies to move into these new planning directions while at the same time taking the risk out of larger scale planning implementations that in the past have been fraught with high complexity, high cost, and long implementation times. Oracle Engineered Systems deliver value to the business by allowing them to expand into these new areas, gain new insights and optimize their value chains while staying in control of the entire process.

Benefit #2: Increased Availability and Scalability

As companies grow, either organically or via acquisitions, the planning processes need to scale up to handle the increased scale. Companies also benefit from using a single hardware platform to support a larger volume of transactions and end users. Companies continue to transition from multiple disconnected planning systems to more consolidated corporate planning solutions. Whether this is an upgrade of older planning systems or a rollout of a new planning process across multiple regions or business units, such initiatives drive cost saving in infrastructure management, alignment of business processes, corporate visibility to the overall supply chain and synergies across business units to

optimize the global supply chain. CDP is designed to handle the scalability required for such planning systems. Other business requirements which can greatly benefit from being able to scale include:

- For companies in high growth industries (i.e. mobile or e-commerce) or rapid market expansion (i.e., emerging markets) growth in the volume of transactions associated with growth in a company's business can mean the need to plan for larger planning footprints.
- Growth in the volume of transactions can also result from a specific business event such as an acquisition, restructuring, or merger of company divisions.
- Large systems become more complex and the simplicity provided by Oracle Engineered Systems enables a company to scale to meet their needs without the growth in complexity of managed systems or having to replace hardware as new size thresholds are reached.

Oracle CDP has been enhanced in several areas to take advantage of Oracle Engineered Systems in the following ways:

- Demand Management and Predictive Trade Planning Forecasting processes have been optimized with Oracle Exadata using optimized data processing approaches. Set-based SQL methods have been applied throughout the data processing architecture from data loading to planning operations, batch processing and integrations.
- The Forecasting Engine Cluster has been aligned with the Exalogic cluster and Oracle Virtual Machine (OVM) deployment to enable maximal advantage of Exalogic scale-out capabilities.
- The Forecasting Engine Cluster has been extended to enable new algorithmic processes to leverage the Demantra cluster framework with Exadata. New replenishment processes can now be executed across the cluster and scaled-out on Exalogic.
- User actions such as a Forecast simulation or Batch Logic Engine runs can be rapidly executed across the cluster using the new enhancements to the VCP Planning Cluster.

Benefit #3: Increased User Satisfaction via Improved Response

End user system response times are critical in planning applications. Planners work with large data sets and execute large data transactions. This is not similar to traditional OLTP processes (with noted exceptions like "month-end" processing) that support a very large numbers of users executing smaller transactions. Instead, planning systems have the opposite profile - a smaller number of users executing a very large number of transactions for analysis or summary reporting. In planning systems the number of transactions can be massive, requiring the processing of millions of rows for a user's operations.

CDP benefits from the performance provided by Oracle Engineered Systems. There are several key factors to keep in mind that are driven by improved user response times:

1. **Reduction in operating costs** – Improvement in employee productivity as a result of reduced planning cycle time and reduced number of plans can improve operational efficiency and lead to reductions in operational costs.

2. **Improvement in top line revenue** – Improved response times allow planners to review and audit more forecasts, assimilate more feedback on the forecast from product and sales managers in real time and evaluate more forecast model scenarios. This leads to improved forecast accuracy, which drives improved inventories, higher order fulfillment and improved sales.
3. **Greater user acceptance** – With improved system performance, user adoption and daily use of planning systems is increased. Users are more likely to shift from decentralized spreadsheets to a central environment that provides them and the rest of the business real-time visibility and accountability of planning decisions.

For example, a simple forecast change for a product group in an account can translate into an update of 100K rows for each time period, even though the user is looking at a single row. With the power of Oracle Exadata, CDP worksheets can load and refresh up to 10 times faster. User updates can be 10 times faster. Forecast simulations can be 10 times faster. This enables planners to rapidly evaluate multiple different business scenarios, collaborate with peers and improve forecast accuracy.

Benefit #4: Lower Total Cost of Ownership, Faster Time to Value

In more traditional implementations Oracle Applications are developed to support a myriad of different hardware options and possible configurations. The Oracle software is installed and configured on customer-selected hardware systems. Oracle standard applications provide many features for tailoring the setup to take advantage of the specific hardware characteristics that the customer is using. The time taken to install, tune and tailor the solutions for your hardware can be time-consuming and costly, requiring expertise on the application, expertise on the hardware, and the knowledge of the proper alignment between these. This time delay itself has a cost to the business, as well as the effort and its cost. Add to that the risk of not achieving the optimal configuration and also the risk of revising hardware in cases where the hardware components (compute, transform, storage) were unbalanced for the planning profiles targeted. Oracle Engineered Systems provide a highly valuable alternative.

Oracle Engineered Systems reduce the time-to-value and implementation cost for the business. Oracle CDP drives shorter deployment and configuration times when using Oracle Engineered Systems as they are available as Oracle Virtual Machine Templates (OVM Templates). These OVMs are ready to be deployed on Oracle Engineered Systems as preconfigured certified application instances ready to run and contain a full computing configuration pre-installed including Oracle Enterprise Linux, Oracle WebLogic Server, CDP, combined with settings, configuration and tuning already done. These OVMs can also be deployed automatically to an Oracle Exalogic machine using Oracle Server that manages the OVM instances on an Exalogic Machine.

For consumption data specifically, the data storage can be multi-Terabyte schemas for large systems. Oracle's Advanced Compression and Oracle Exadata's Hybrid Columnar Compression solution can provide significant storage cost reduction for consumption data driven application environments.

CDP has been designed, built, scaled, and tested from the ground-up as a hardware and software package specifically and exclusively for Oracle Engineered Systems. This allows us to provide highly accelerated implementations due to targeted platform development which takes advantage of unique capabilities of the Oracle technology stack throughout the product design, including development,

testing, and performance tuning. When customers deploy the system, the software and hardware are deployed in unison and operate as a designed cohesive solution, reducing overall cost of deployment and ownership.

Benefit #5: New Levels of Supply Chain Performance

By offering previously unattainable levels of performance, CDP enables forecasting and replenishment planning at the store and daily level. The benefits are improved on-shelf availability, increased sales, reduced inventories, and lower supply chain costs. Organizations can manage demand and replenishment across multiple channels, at different levels of granularity, eliminating the need for separate systems. With the resulting visibility to demand and inventories, companies can make better supply chain decisions and improve company performance.

Conclusion

Oracle In-Memory Consumption-Driven Planning offers the breakthrough in computing performance needed for addressing the most challenging demand and replenishment planning applications. This is the result of bringing together hardware and software engineered together to provide extreme data throughput performance while maintaining production reliability for large data transactions that planning systems thrive on. Oracle Exadata provides massive data processing throughput, while Exalogic optionally offers scale-out of planning runs. The result is reduced planning cycle time and data latency; increased application availability and transaction scalability; increased user satisfaction via improved response time; and lower total cost of ownership and faster time to value. At the same time CDP takes the risk out of large scale planning implementations that in the past have been fraught with high complexity, high costs, and slow return on investment. With this type of performance, running forecasting and replenishment plans at the store and daily level is possible. Oracle In-Memory Consumption-Driven Planning is thus helping to bring manufacturers closer to achieving the vision of the demand-driven value chain.

Oracle In-Memory Consumption-Driven Planning
March 2013
Authors: Michael Liebson
Contributing Authors: Bart Feldman, Vikash Goyal, Nadav Zivelin

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com

Oracle is committed to developing practices and products that help protect the environment

Hardware and Software, Engineered to Work Together