

- business intelligence
- data warehousing
- data mining
- OLAP

Enabling CRM

Building the enterprise
customer data store

Agenda

Terminology & Concepts

- Operational Data Store (ODS)
- Enterprise Customer Data Store (CDS)

CDS Architecture

- Content
- Storage
- Processing

Case Study

- Opportunities Sought
- Key Considerations

Project Mgmt

- Highlights

- business intelligence
- data warehousing
- data mining
- OLAP

Terminology & Concepts

Operational Data Store (ODS)

- A concept that underlies the CDS
- Characteristics:
 - Subject Oriented
 - Customer, Vendor, Account, Product, etc.
 - Integrated
 - Data is cleansed, standardized and placed into a consistent data model
 - Volatile
 - UPDATES occur regularly, whereas data warehouses are refreshed via INSERTs to firmly preserve history
 - Current Valued
 - No INSERTs means no significant retention of historical data (don't take this literally...history can extend to an accounting cycle, for instance)
 - Detailed

Operational Data Store (cont.)

Examples

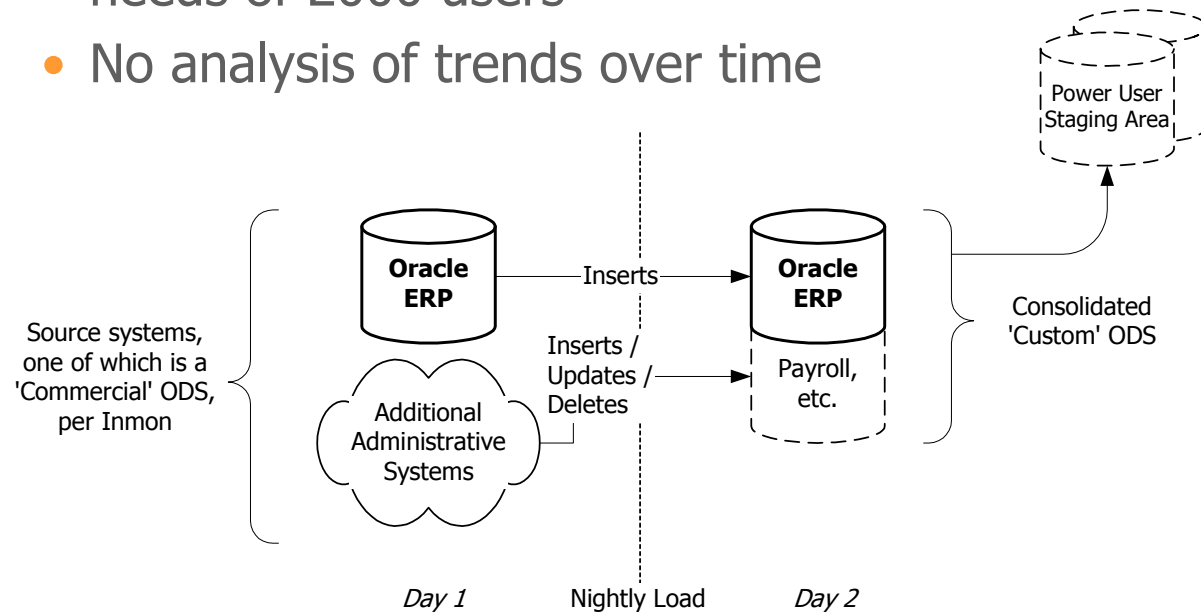
- Goldman Sachs
 - Disparate, global real estate investment details collected daily, then standardized and transmitted to HQ
- WalMart
 - Real-time transaction 'flash' sales report. Highly granular and can be aggregated as seen fit

Operational Data Store (cont.)

Examples (cont.)

Harvard

- General Ledger and other internal administrative data is off-loaded to ODS for the daily reporting needs of 2000 users
- No analysis of trends over time



Operational Data Store (cont.)

Classifications	Comment
Class 1	<ul style="list-style-type: none">■ Refresh cycle: Real-time■ Degree of transformation: Low due to compressed timeframes
Class 2	<ul style="list-style-type: none">■ Refresh cycle: 1/2 - 1 hour store and forward mechanism■ Degree of transformation: Medium
Class 3	<ul style="list-style-type: none">■ Refresh cycle: Daily. Traditional batch process■ Degree of transformation: High
Class 4	<ul style="list-style-type: none">■ Refresh cycle: Ad hoc, often involving preprocessed, value-added information from a data warehouse■ Degree of transformation: Highest

Enterprise Customer Data Store

■ Subject Area

- Customer

■ Goals

- Extend the lifetime, and lifetime value, of the customer relationship
 - If customers who provide most of the revenue stay on the books for longer periods of time, it has a positive effect on the company's long-term revenue position
 - Customers reward companies who treat them 'right' by continuing to buy over time
- Nurturing customer perceptions strengthens the company's brand
 - A strong brand and loyal customers helps attract new customers
- Optimize Return on Assets (ROA)

Enterprise Customer Data Store (cont.)

■ Examples:

- Acknowledge the relationship & the customer's importance
 - Nine West
 - Customers are asked for a name and address during checkout to sort new customers from existing ones. Transaction histories for existing customers--along with shoe size and width--appear on the POS terminal
 - Target
 - Providing a high quality guest experience is a key strategy for Target Corporation.
 - "Let's say we have a customer who has an excellent credit rating with us and has a Target card, and they have problems with a check at one of our Marshall Fields stores," explains Deb Bauman, Vice President of Development for Target Direct and Target Financial Services in Minneapolis. "Let's say they missed a payment on their Marshall Fields card but never missed a payment with their Target card. [With our ODS] we can bypass our ordinary rules and accept their payment at the Marshall Fields," Bauman says. "We wanted to improve the whole customer communication process."

Enterprise Customer Data Store (cont.)

■ Examples:

■ Optimize ROA

● Leader in Overnight Delivery

- Permit extended call times for high-profit customers placing inbound service calls
- Hinges on accurate calculation of profits, itself dependent on having comprehensive data from disparate systems (e.g., call detail records and sales & return information)

● Lucent

- Initiated ambitious 360-degree CRM program in 2000
- Single view of customer entailed integrating data from Siebel, Clarify, E.piphany, etc.
- Is attempting to ensure that warranty renewals are not only offered to customers (!), but offered at an optimal point in time

● Variable Pricing / Revenue Management?

Enterprise Customer Data Store (cont.)

■ Examples:

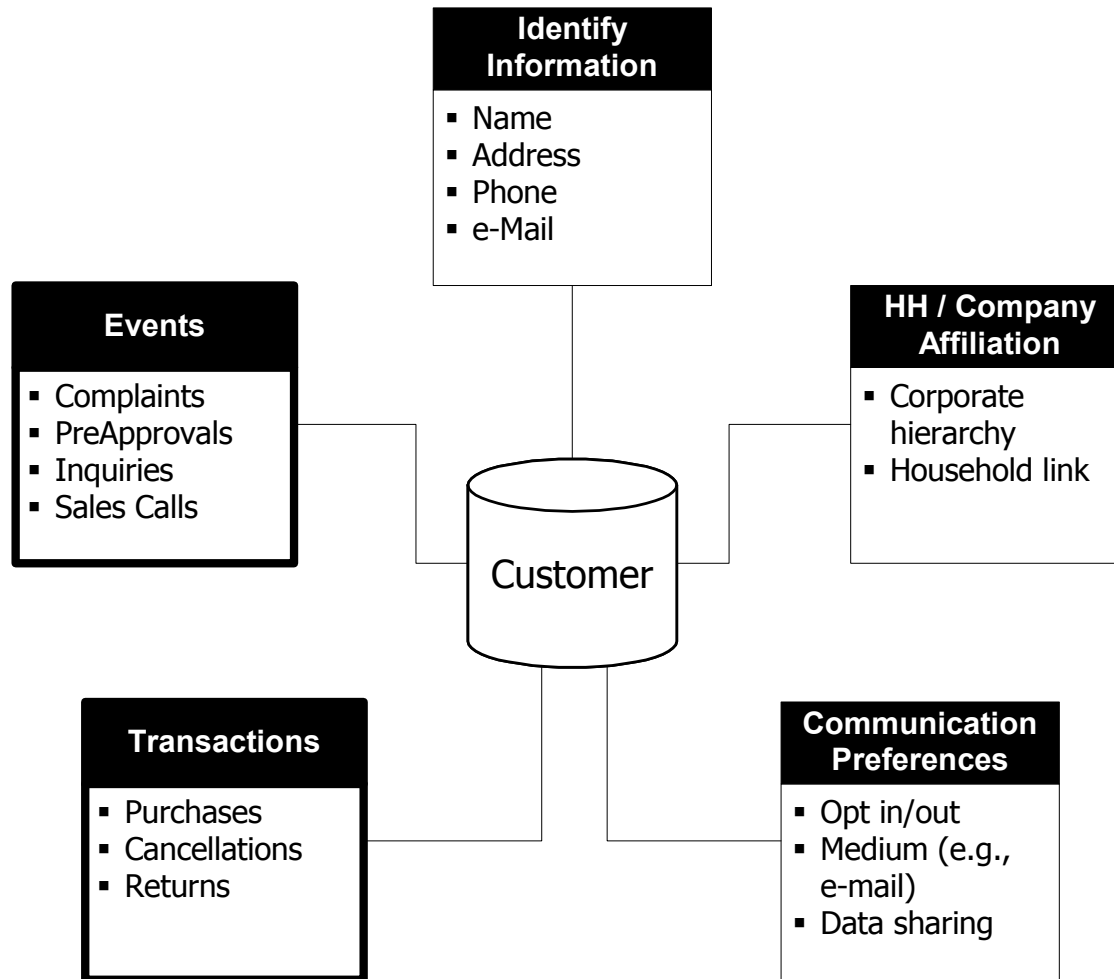
- Extending lifetime value of the customer (cross-channel) relationship
 - The Limited / Victoria's Secret
 - "Our analysis has found that customers buying Victoria's Secret products through all three channels -- retail stores, the Web and catalogs -- spend three to five times more than customers who buy through only one channel." William Lepler, The Limited's VP of CRM
 - Eddie Bauer
 - ""We're just starting to understand how to use channels more effectively," states Michael Boyd, Director of CRM for Eddie Bauer. "The ability to gather and crunch valuable data about customers from multiple channels, via the Web, catalog sales or store locations, is the force fueling electronic CRM"
 - Lands' End
 - "We were able to realize real cost savings very quickly," he says. "But in the long term, I think there's even more value in AbiliTec being able to integrate our various business channels." Dave Johnson, Vice President of Direct Marketing at Lands' End.



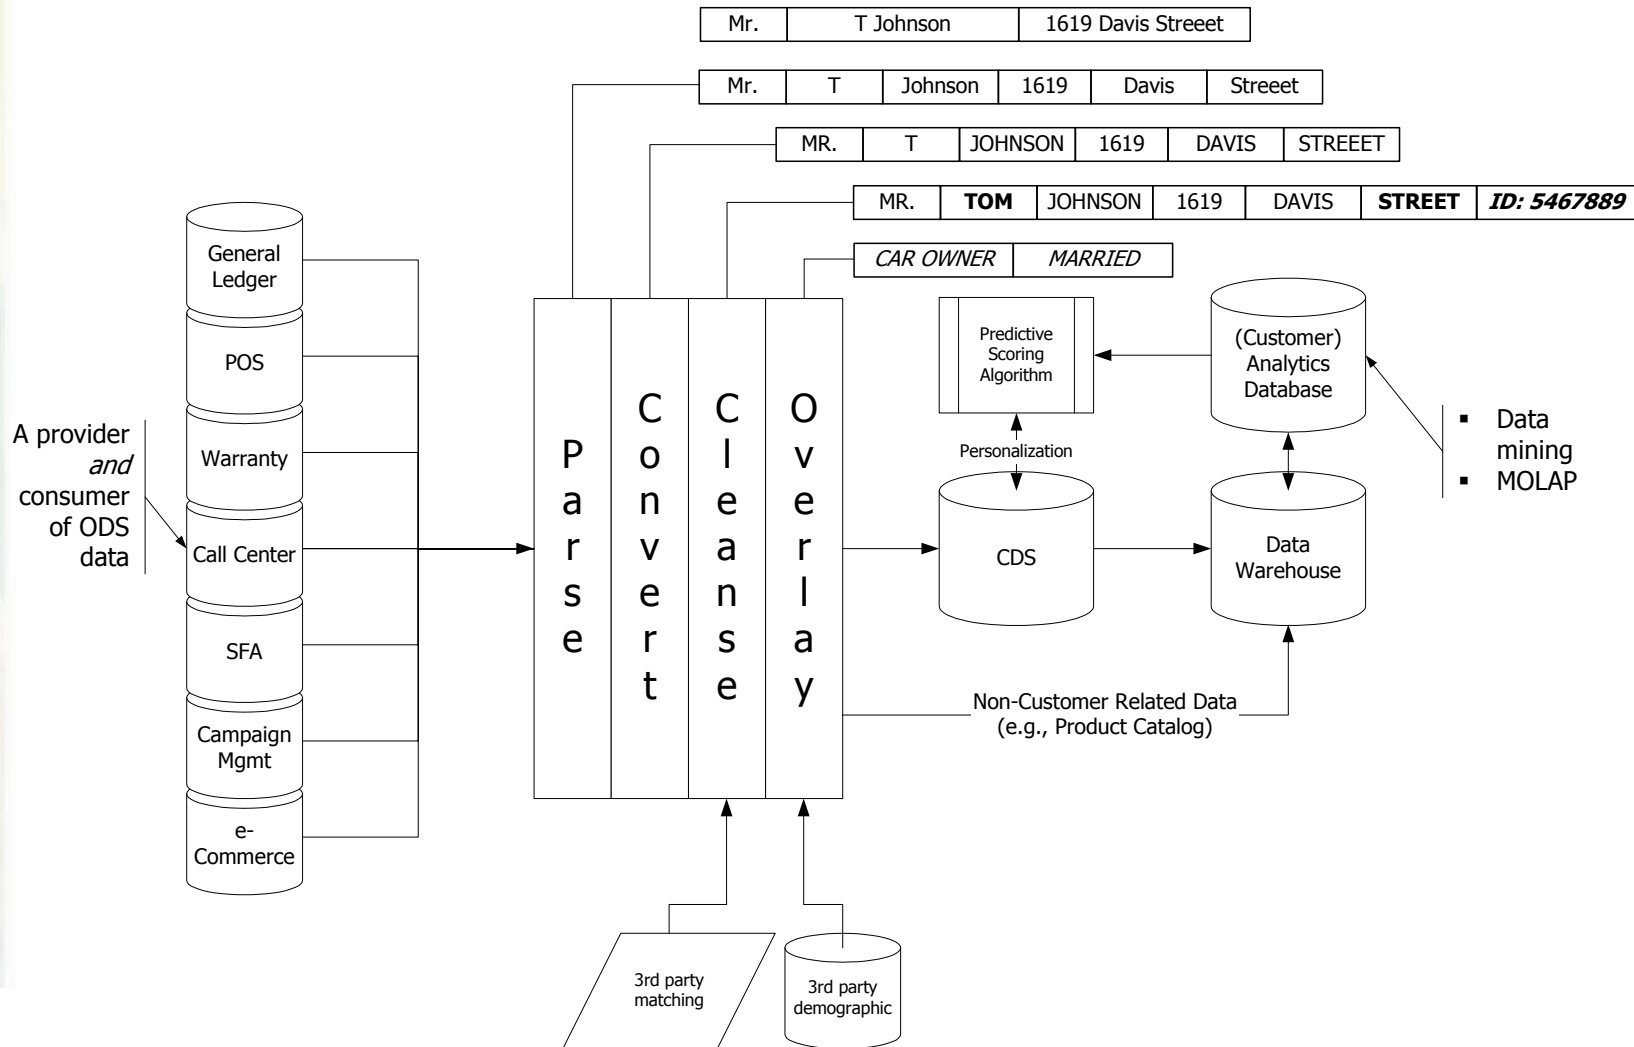
- business intelligence
- data warehousing
- data mining
- OLAP

CDS Architecture

What Goes in the CDS?



How is the Data Processed?



More Sample Transformations

Transformation	Example
Parse data into piece parts	'Intelligent' promotion code: <ol style="list-style-type: none">102 (1st mailing of 2002)HG (Product indicator; Hard Goods)345267 (Prospect ID)DB (Source code; Dun & Bradstreet)
Convert	1/0/Blank --> M/F/Null
Summarize	Return details stored in full, but POS purchase orders lightly summarized to transaction level (1 circular saw + 1 circular saw is stored in CDS as simply qty = 2 circular saws)

How is the Data Stored?

- Both dimensional star schemas and normalized ER models can be used
 - Transactional nature of CDS may call for normalized--or at least snowflake--approach, however, to permit selective and rapid UPDATES against a single row, rather than a range of rows

- business intelligence
- data warehousing
- data mining
- OLAP

Case Study:

Target / Compaq Zero Latency Enterprise



Class 1 CDS Problem Overview

■ Mandate

- Consolidate information across 24 'guest' databases to give customers a consistently positive experience across all touchpoints
- Enforce a single corporate interaction policy...a standard set of business rules

■ Initial Workload

- Twenty internal clients spanning all functional areas
- A three month marathon gathering business requirements and developing a conceptual data model

Program Objectives

Category	Comment
Marketing	<ul style="list-style-type: none">■ Marketing desires more targeted communications to retain existing guests, and to acquire new guests whose profiles fit their brand segments■ Target also wants to increase the profitability of each of their guests
Brand Management	<ul style="list-style-type: none">■ Brand Management requires Target to deliver a consistent image across operating companies and channels■ “Brand collision” can be avoided when they know what interactions that guest has had with them■ Target does not want a campaign to down-sell a guest from an already profitable relationship with another operating company, and they don’t want to over-saturate their best guests

Program Objectives (cont.)

Category	Comment
e-Commerce	<ul style="list-style-type: none">■ Requires a guest database to provide tailored services and generate near real-time product recommendations■ Target's guests like the convenience of shopping online and they seek a unified, cross-channel view of their activity available 24 hours a day. Both in-store and Internet experiences are memorable & brand-defining to their guests
Guest Services	<ul style="list-style-type: none">■ Guest Service improves when team members have real-time access to a customer's profile, lifetime value and interaction history

Program Objectives (cont.)

Category	Comment
Merchandise Planning	<ul style="list-style-type: none">■ Merchandising analysts will have access to market basket-level detail required to gain in-depth knowledge of personal preferences and behaviors. This information is needed to take their inventory management, ad effectiveness, and pricing strategies to the next level
Risk Management	<ul style="list-style-type: none">■ Target's goal is to accurately identify their worst customers while minimizing the impact of loss prevention programs on their best guests



Key Considerations

- Real-time messaging
 - How to detect and propagate changes from legacy systems to CDS?
 - Very different from batch processes most warehouse designers are used to
- Linking customer records
 - How to rapidly join records across systems that use different IDs and store variations of essentially the same customer?
- The CDS as the system of record
 - Once we have a working system, should users update it as they would the legacy system?

Real-Time Messaging

- Uses a publish/subscribe pattern to 'push' changed information out to appropriate applications (CDS) as changes occur
 - Customer transaction COMMITed to database
 - Database 'adapter' detects change and places record in message queue managed by message broker, who ensures that record is successfully sent ASAP
 - Certain data transformation may take place at this point, such as converting EDI to XML
 - Compaq ZLE uses Mercator
- Target's transaction rates: 100 / second, but not all will involve customer record cleansing, discussed next

Linking Customer Records

- ETL/EAI tools cannot perform this specialized task. ZLE relies on Trillium and Acxiom's AbiliTec.
- Let's review Acxiom's AbiliTec service:
 - Establish TCP/IP or leased line to Acxiom Data Network. Next, via XML, transmit:
 - Consumer/Business Name (XML compliant...no ampersand, for instance)
 - Address (parsed, preferably)
 - ZIP
 - SSN (optional)

Linking Customer Records (cont.)

- Specify whether you wish a maintained or derived match to take place
 - *Derived* matches rely exclusively on pattern algorithms and USPS data
 - *Maintained* matches use pattern algorithms and USPS data, but primarily depend on a database of name/address variations (a slowly changing dimension!)
 - Before move and marriage:
J.Doe, San Francisco, CA
 - After move and marriage:
Jane Smith, Dallas, TX
- System returns a 'link'. That is, Acxiom's unique identifier for that individual or business

ZLE Sample Architecture Components

Vendor	Category
Trillium	<ul style="list-style-type: none">■ Cleansing and deduplication
Acxiom	<ul style="list-style-type: none">■ Customer demographics
Mercator	<ul style="list-style-type: none">■ Message broker
HNC	<ul style="list-style-type: none">■ Blaze Rules advisor
SAS	<ul style="list-style-type: none">■ Data mining
Microstrategy	<ul style="list-style-type: none">■ Multi-dimensional ROLAP database
Siebel	<ul style="list-style-type: none">■ Call center
Compaq	<ul style="list-style-type: none">■ Tandem NonStop™ Himalaya™ servers, providing clustered, fault tolerant 24x7 uptime■ NonStop™ SQL database

The CDS as the System of Record

- Avoid this temptation, though technology to make this viable is progressing
 - “It is patently a mistake to try to shift auditability and adjustments to the ODS”
(Source: Building the Operational Data Store; 2nd Edition; B. Inmon)
 - Reverse scrubbing
 - Can integration server/transformation engine/ETL process work in reverse to rebuild the old values we parsed and transformed? Not likely!
 - Do we really want to maintain data entry edit rules *both* in CICS and in our new server-side (e.g., EJB) apps?

- business intelligence
- data warehousing
- data mining
- OLAP

Project Management



Project Management Highlights

These may sound obvious, but very few projects have the discipline or skill sets needed to quickly cycle through these steps...

- Define/prioritize *measurable* business drivers
- Define the business processes that are affected
- Define the information required by, and generated from each of the business processes
- Generate CRM requirements; Compare to existing capabilities
- Define strategies for addressing requirements; Communicate requirements to vendors
- Document the initiative



- business intelligence
- data warehousing
- data mining
- OLAP

Appendix

Victoria's Secret

- Victoria's Secret's database has 21 million customer names, but The Limited selected about 125,000 highly loyal lingerie customers to be tested for their cross-channel experiment.
- Female customers were chosen due to their spending patterns. Some were store-only customers, some were Web-only and some were catalog-only. The Limited segmented these customers into five cells, each with the goal of introducing a customer to a new channel. One cell included catalog shoppers whom The Limited wanted to direct to the Web, while another included Web shoppers whom The Limited wanted to direct to the catalog. Each cell contained 25,000 customers.
- Each customer received a similar mail piece but with a different message. For example, if the customer was a store shopper, she was invited to visit the Web site or order a catalog. If she was a catalog shopper, she was invited to visit a nearby Victoria's Secret retail store to sample the fragrance.
- 400% ROI was claimed, merely as a result of moving catalog-only clients to the retail store. These catalog-only shoppers also were selected based on having retail shops near them.