Business Information Warehouse

Technology

SAP AG ● Neurottstraße 16 ● 69190 Walldorf ● Germany

Contents

Executive Summary **3**

Background 3

Key features of Business Information Warehouse: 3

Part of SAP's Business Framework 4

Business Information Warehouse Architecture Overview **5**

Data Warehouse Management and Administration **6**

Production Data Extractor 7

Staging Engine 7

Administrator's Workbench 8

Data Storage and Representation **9**

InfoCubes 10

Hierarchies 11

Meta Data Repository 11

Operational Data Store 12

Analysis and Reporting **12**

Slice and dice 13

Business Analysis Library 13

Report types and layout structure 14

Report Catalog Browser 15

Report presentation 17

Performance **17**

Openness **18**

Business Information Warehouse configurations **19**

Conclusion **19**

Introductory note

This White Paper gives an overview of the technological infrastructure of SAP's Business Information Warehouse (Business Information Warehouse). At the time of publication, Business Information Warehouse was undergoing trials in a number of pilot projects. Availability of Business Information Warehouse will be announced through the usual SAP channels.

The many pre-defined information models and business reports provided in the scope of supply of Business Information Warehouse are the subject of a separate paper, currently in preparation.

Executive Summary

Background

Today's decision-makers urgently need accurate information on production, sales and markets, finance and personnel. They need a complete and up-to-date picture of their business, and their business environment. However, that information is spread throughout the corporate IT structures and beyond, on a wide variety of platforms and in a huge range of applications. In other words, getting hold of vital facts and figures can be a complex and time-consuming task.

The recognized technical concept for meeting this challenge is known as data warehousing. A data warehouse is a separate application environment with a dedicated database drawing on diverse data sources and designed to support query and analysis.

The Business Information Warehouse (Business Information Warehouse) from SAP is a state-of-the-art, end-to-end data warehouse solution which gives knowledge workers rapid access to data from SAP systems, from other enterprise applications and from external sources, e.g. content providers such as Nielsen. Business Information Warehouse is SAP's new generation of data warehouse technology, raising the capabilities first shown in the Open Information Warehouse to a completely new level.

Based on proven R/3 client/server technology and exploiting SAP's acknowledged business expertise, Business Information Warehouse is a robust, functionality-rich solution which is easy to install, easy to use and easy to maintain.

Key features of Business Information Warehouse:

- Business Information Warehouse is a comprehensive, ready-to-go data warehouse solution. It includes a relational OLAP processor, automated data extraction and staging tools, a pre-configured meta data repository, a user-friendly front-end ("Business Explorer") with powerful reporting and analysis tools, and an Administrator's Workbench. Business Information Warehouse can therefore be implemented quickly, simply and cost-effectively.

- Data can be gathered automatically from R/3, from R/2 and, via open interfaces, from non-SAP applications and external sources.

- Business Information Warehouse includes an easy-to-use central administration tool. Modifications to Business Information Warehouse configuration can be made without difficulty in accordance with changes to the enterprise IT environment or emerging information needs.

- Pre-configured information models and reporting templates help users generate reports quickly and simply.
- Business Explorer displays reports in MS Excel (with Business Information Warehouse extensions), allowing users to exploit existing skills and familiar MS Excel features to create custom documents on the basis of Business Information Warehouse report views.
- Business Information Warehouse supports pre-defined standard reports and ad hoc analysis; both allow drill-down and multi-dimensional views.
- Users can define their own collections of favorite reports, enabling them to access relevant information at the "touch of a button".
- Business Information Warehouse is built for high performance. It resides on its own dedicated server. OLTP and reporting activities are therefore separated, and system performance is not compromised. Report caching mechanisms have also been carefully designed to maintain high performance.
- Business Information Warehouse embodies all the advantages of leading-edge R/3 Basis technology with its three-tier architecture: security, integrity, scalability, high availability and interoperability within a uniform environment.
- Business Information Warehouse includes a proven, consistent enterprise model for the application of business rules to data. This builds on the enterprise model of R/3 and greatly facilitates data warehouse implementation and management.
- Business APIs allow Business Information Warehouse to be integrated rapidly into the existing computing environment and enable specialized tools to be added, where required.

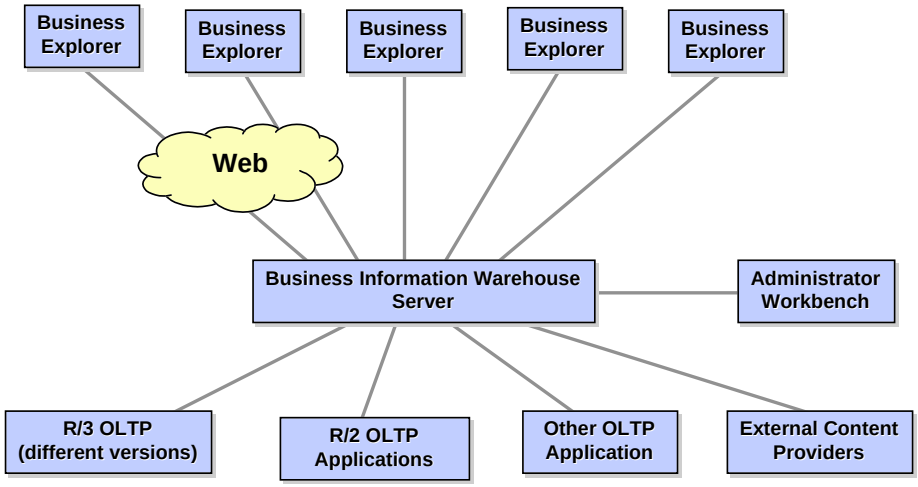

Fig. 1: Business Information Warehouse

Part of SAP's Business Framework

Business Information Warehouse is a separate R/3 system with its own release cycle. It is a Business Component of the Business Framework. Business Framework is SAP's strategic product architecture, designed to facilitate the seamless and rapid integration of new business functionality and IT technology into existing environments. Business Framework provides an open architecture, allowing its basic elements, the Business Components, to operate through standardized Business Application Programming Interfaces (BAPIs). Users benefit from faster implementation, non-disruptive maintenance and lower cost of ownership.

Business Information Warehouse Architecture Overview

Business Information Warehouse is a complete data warehouse solution, comprising the components shown in the figure below. The top layer of the figure depicts the Business Explorer, comprising Business Information Warehouse's client components for the end-user. The middle layer represents the Business Information Warehouse Server and the bottom layer the OLTP systems from which source data are extracted.

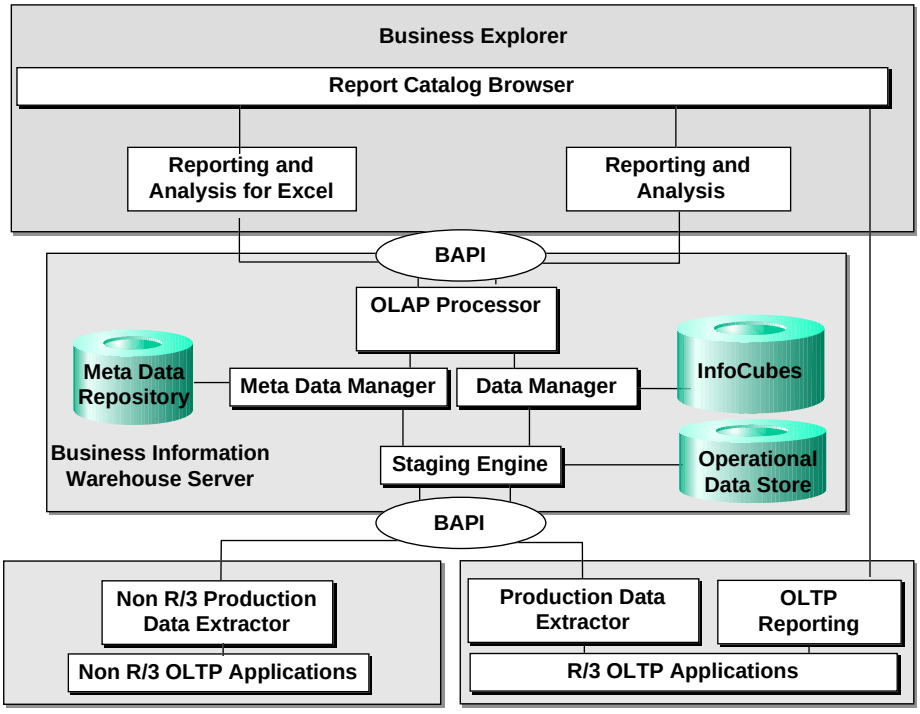

Fig. 2: Business Information Warehouse Architecture

Data warehouse management and administration (automated extraction, mapping, loading and maintenance) are performed by:

- The Production Data Extractor for R/3 OLTP applications
- The Production Data Extractor for non-R/3 OLTP applications
- The Staging Engine
- The Administrator's Workbench (not shown)

Data storage and representation are provided by:

- The InfoCubes in conjunction with the data manager
- The Meta Data Repository in conjunction with the meta data manager
- The optional Operational Data Store

Presentation of the data to the end-user for **analysis and reporting** is supported by:

- The OLAP Processor, the core business intelligence which allows exploitation of data for analysis/reporting
- The Business Explorer (the reporting environment front-end) including the Report Catalog Browser for effective organization of personal favorite reports, and the report and analysis components for standard reports and ad hoc analysis.
- The OLTP reporting component, for direct access to existing reporting functions within R/3 applications

These three central areas are described in more detail below.

Data Warehouse Management and Administration

Within any data warehouse, data must first be extracted from diverse sources and then mapped to the specific data structures required for analysis and reporting.

Most data warehouse solutions are designed to extract and map data at the level of individual database fields without recognizing the business relevance of that data. Data and their business relevance therefore have to be modeled and coded individually for each user organization, as do extraction and mapping. If changes are made to the underlying business processes, then this can have an impact on the entire flow of data and may result in significant maintenance overhead.

Business Information Warehouse, by contrast, recognizes the business role played by data and applies this knowledge to data extraction, to mapping and to the structures used to store the data in the warehouse. This eliminates much of the work usually associated with implementation and change management, and guarantees the consistency of data from end to end. This inherent "knowledge" of business rules and how they apply to data is based on Business Information Warehouse's meta data model which, in turn, is derived from the enterprise model of R/3.

Business Information Warehouse features inherent "knowledge" of business rules and how they apply to data

Business Information Warehouse includes a number of components which allow data warehouse management tasks to be performed with a high degree of automation, aimed at fast implementation and low maintenance during ongoing data warehouse operation.

The three major components are the Production Data Extractors, the Staging Engine, and the Administrator's Workbench.

Production Data Extractor

The Production Data Extractor is a set of programs for the extraction of data from R/3 OLTP applications. These programs consolidate the transaction data and master data from OLTP applications that are required for analysis and reporting purposes. The extract data represent a consistent snapshot of OLTP data.

For the initial build, Business Information Warehouse transfers all OLTP data employed within the data warehouse. However, during ongoing operation only a subset of OLTP extract data (projection and selection) needs to be transferred to Business Information Warehouse. Special emphasis has been put on support for incremental data changes (delta updates), which significantly reduces the volume of data to be transferred. This is a major advantage over traditional data warehouse extractors, which operate on the plain OLTP tables and which have to transfer complete snapshots, not just changed data.

Business Information Warehouse is designed to keep data traffic to a minimum

Business Information Warehouse includes extract programs for all the major R/3 applications, including logistics, controlling, finance, and human resources.

Non-SAP data sources are connected using third-party or custom extraction tools and are regarded by Business Information Warehouse as source systems. However, these tools supply extract data directly to a Staging Engine BAPI on the Business Information Warehouse platform. The Staging Engine then performs the mapping and update rules and thus ensures the consistency of data used in reports. For non-SAP extract data, a rich library of format conversion functions is available, including year 2000 date conversion. In addition, a dedicated BAPI supports meta data exchange with third-party tools.

Support for non-SAP data sources via BAPIs

Staging Engine

Mapping and transformation are fully automated

The Staging Engine is employed to implement data mapping and transformation as a meta data driven process which is automated throughout. ALE, SAP's persistent messaging middleware, is used for data transport from OLTP platforms to the Business Information Warehouse server.

Initiated by its Scheduler, the Staging Engine requests an extract from an OLTP system. Typically, extract jobs include a selection criterion that determines the set of data to be extracted and transferred.

Production data then undergo a series of condensing, mapping and transformation steps before it arrives in the InfoCubes, the central data containers of Business Information Warehouse (cf. section InfoCubes). Starting from the OLTP side, the extract programs deliver data in an extract source structure. A source structure can be transaction data (characteristics and "key figures", i.e. the quantifiable values) or master data attributes, texts and hierarchies.

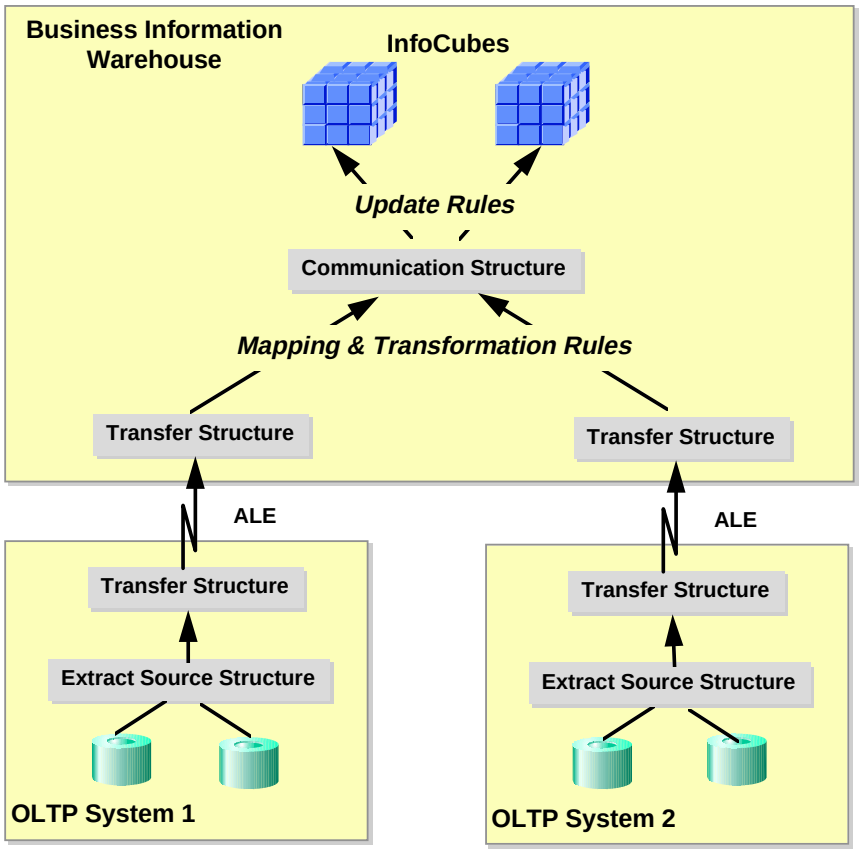

Fig. 3: Data Staging

Multiple transfer structures are mapped to a single communication structure used to homogenize data from disparate OLTP sources (including different releases of R/3) into a common format. Within the context of Business Information Warehouse, transfer structures and communication structures are known by the generic term "InfoSources". Finally, the InfoSource is mapped to one or more InfoCubes. At this stage, business logic can be

applied, ranging from aggregations to arbitrary calculations on key figures, enrichment with additional attributes and user-defined functions.

Business Information Warehouse ensures the consistency of the load process. The transmission of the extract data is performed in an asynchronous fashion, giving the administrator the flexibility to run the three different tasks of building the extract, transferring it to Business Information Warehouse and updating the InfoCubes, at the same or independent times.

Administrator's Workbench

The Administrator's Workbench is a toolset with a user-friendly GUI that allows the Business Information Warehouse administrator to perform all data warehouse construction and maintenance tasks within a single, unified environment.

The component for meta data maintenance allows the administrator to specify and maintain InfoCube definitions as well as all technical meta data (such as connected source systems and their InfoSources with the corresponding mapping and transformation rules), and the business rules for InfoCube updates. SAP's pre-loaded definitions can be adapted to users' needs. It is also possible to add new InfoSources or InfoCubes specific to the user organization with equal ease.

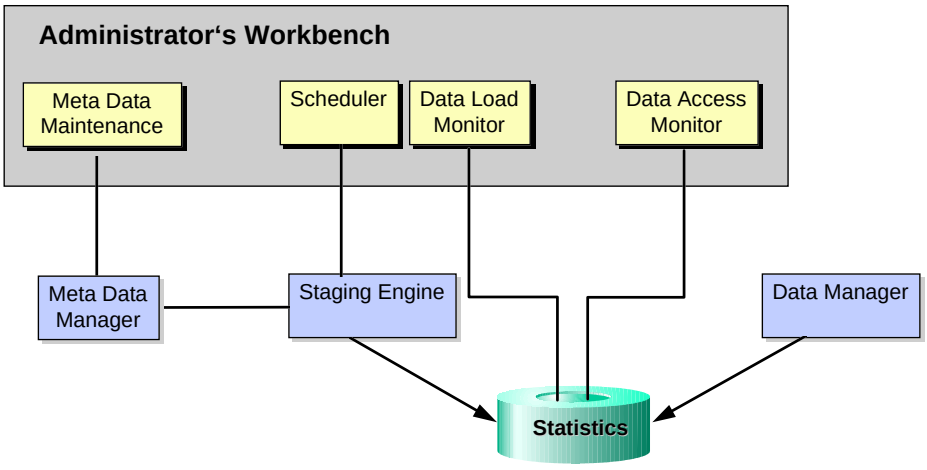

Fig. 4: Administrator's Workbench

The Business Information Warehouse Scheduler builds on the functionality of the R/3 Basis scheduling system

Once the source and target structures and mapping rules are set up, the administrator uses the Scheduler, the Data Load Monitor and the Data Access Monitor to control and supervise the ongoing operation of Business Information Warehouse. The Scheduler maintains extract and load jobs, which typically run at recurrent time intervals (e.g. once daily). The Business Information Warehouse Scheduler harnesses the proven power of the R/3 Basis scheduling system.

The Data Load Monitor supervises the load and staging processes, provides statistics on current and completed load jobs and notifies the administrator of exceptions. The Data Access Monitor provides statistics on Business

Information Warehouse usage, e.g. the frequency of query execution, usage of summary levels, etc.

The Administrator's Workbench is also used for tasks such as the definition of access rights and report catalog maintenance.

In summary, the Administrator's Workbench provides effective support not only for setting up Business Information Warehouse, but also for monitoring performance and usage, and for accommodating new needs.

Data Storage and Representation

Business Information Warehouse allows the end-user to look at data from many different perspectives, to transparently combine data from various sources and to drill-down from one level of detail to another. To make this possible, Business Information Warehouse stores data in multi-dimensional form.

InfoCubes

InfoCubes contain key figures and characteristics

InfoCubes are the central containers of data used in reports and analysis. They contain two types of data: key figures and characteristics. Key figures are quantifiable values, such as revenue. Characteristics are needed to compute and present key figures according to differing perspectives. Typical characteristics for revenue include the corresponding product or customer.

The aggregational behavior of key figures needs to follow certain business rules. Key figures such as revenue might be summed up by product and by time. Other key figures such as inventories are semi-additive, in other words a summation by plant with storage location makes sense, but summation by time does not. A third category of key figures is non-additive, such as ratios. Business Information Warehouse allows all these properties to be modeled, and the OLAP Processor guarantees that all business rules are satisfied and the computed views present valid results from a business point of view.

An InfoCube is a set of relational tables arranged according to the star schema: one large fact table at the center and several surrounding dimension tables. The fact table is used to store all key figures at the lowest level of granularity, while the dimension tables store the characteristics used for reporting and analyzing these key figures. Dimensions are considered to be mutually independent - only the fact table links the dimensions to the key figures. An example for the analysis of profitability by customer and product is shown below.

The effectiveness of the star schema and its derivatives has been proven in many data warehouse projects. It is a structure that supports efficient use of storage space and of CPU cycles, minimizing query response time.

The star schema is an efficient way to store data

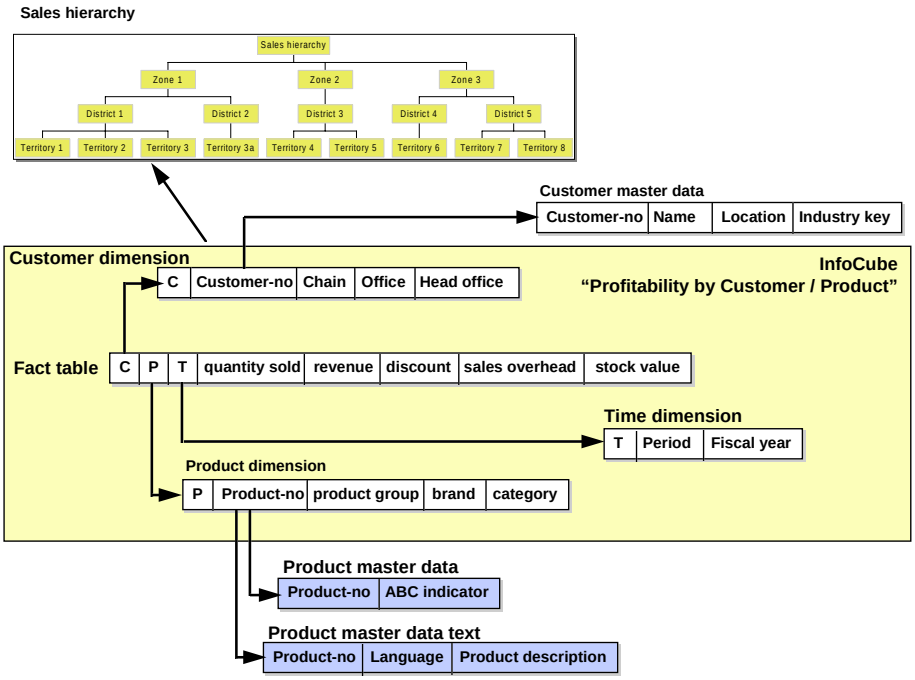

Fig. 5: InfoCube Example

In addition to this basic star schema, Business Information Warehouse allows master data and descriptive text to be stored in separate tables, shared between InfoCubes. Business Information Warehouse's master data structures are adapted from OLTP master data, and can also be extended to include other data. However, they contain only those attributes needed for presentation or as navigation criteria within reports.

Hierarchies

Hierarchical relationships for characteristics play an important role with regard to analytical and reporting activities. For example, cost centers are usually hierarchically structured; products are generally grouped into product family hierarchies, etc. Business Information Warehouse allows a hierarchy to be modeled over the value domain of one or more characteristics. These hierarchies are independent of InfoCubes and may be used as the basis for aggregation and drill-down criteria within reports, i.e. the ability of users to look at data at differing levels of detail.

Meta Data Repository

Meta data are the key to using information in a data warehouse

Meta data are central to the data warehouse concept. Put somewhat simplistically, meta data are ”information on information”, i.e. they are used to describe the source, history, and many other aspects of data. Meta data allow information stored in Business Information Warehouse to be used effectively for reporting and analysis, and ensure that all users have ”one version of the truth”.

Business Information Warehouse’s meta data repository contains both business-related and technical meta data. Business meta data include definitions, descriptions and rules used for reporting. Technical meta data describe structures and mapping rules for the data extraction and staging process. The repository maintains the various classes of meta data in the following catalogs:

The **InfoObject Catalog** describes all characteristics and key figures independently of their use in InfoCubes and InfoSources. This ensures the use of a uniform definition for all characteristics and key figures, guaranteeing a homogenous data model for data from highly diverse sources.

The InfoObject Catalog is also used to store definitions of derived key figures, i.e. figures calculated dynamically from the key figures in the database by a business analysis function at report execution time.

The **InfoCube Catalog** stores the InfoCube definitions, e.g. a description of the characteristics and key figures from which InfoCubes are composed.

The **Report Catalog** contains all report definitions and their descriptions. The Business Explorer front-end can view these definitions and descriptions, allowing the user to make informed report choices. The Report Catalog provides the OLAP Processor with instructions on which InfoCube data are to be selected, what business analysis functions are to be applied and how results are to be arranged for display by Business Explorer.

The **InfoSource Catalog** stores all InfoSource definitions, including information on mapping to InfoCubes.

A homogeneous data model across all sources

This way of storing and representing data ensures a consistent, homogeneous data model across all OLTP sources. Business Information Warehouse’s Staging Engine and OLAP processor are driven by these meta data. Any structural changes to, e.g. InfoSources or InfoCubes, are made and maintained at just one central location, ensuring that change management is easy, fast and cost-effective.

Operational Data Store

The optional Operational Data Store contains non-aggregated data, and serves several purposes. From an application point of view, reports may want to drill-down to a single document. A request of this type can then be satisfied by Business Information Warehouse without entering the OLTP system.

The Operational Data Store is organized as a set of flat tables, each assigned to a single InfoSource, which holds the extracted production data records. These records can be regarded as pseudo-documents, i.e. scrubbed OLTP documents composed of some fields from the original document that have been selected for transfer into Business Information Warehouse.

The second reason for the Operational Data Store is technical in nature. The Operational Data Store is used as an intermediate store for the staging process. Custom scrubbing and preparation tasks can be performed on a complete extract before data are mapped into InfoCubes. This supports, for example, complex aggregation methods during production data load, which operate on sets of records to compute the aggregate, possibly coming from different, heterogeneous source systems.

Analysis and Reporting

Reports are not static, they support extensive data exploration

The user accesses information held in InfoCubes through reports. For each report, there are numerous different views that display information according to slice and dice operations (cf. section Slice and dice performed by the user. In other words, a report is not a static snapshot, but a gateway through which the user can explore the underlying data.

Business Information Warehouse provides all end-users with a uniform reporting environment, known as Business Explorer. Business Explorer can be used for all types of reporting activities, from ad-hoc analysis to highly standardized reporting.

The OLAP Processor builds the reports on top of the InfoCubes and provides the methods needed to perform slice and dice operations.

Reports can be presented in MS Excel, allowing the user to utilize Excel's formatting and layout features.

For user-friendly organization of reports and fast access to favorites, Business Explorer includes the Report Catalog Browser.

Slice and dice

Reports are not "frozen in time"; users can explore data and trigger report views which show information from a new perspective.

Reports can be used to look at data from many different viewpoints

The user can, for example, move from a general view to one which is more detailed (known as "slicing" or "drill-down"), or from a very detailed level to one which is more aggregated. He can also decide to view data from a different perspective by introducing a completely different analysis criterion ("dicing" or changing view).

One example, illustrated below, would be a manager examining sales figures. He can look at sales for a particular continental region, for instance Europe, then drill-down to sales for individual countries within Europe. He may now

decide that this is the level of detail at which he wishes to remain, but that he now wants to change view to see sales listed according to product groups.

It is also possible for him to jump from one type of report to another, transferring the "current focus" to the new report, i.e. he could switch to a report on the order pipeline, but would retain the focus on Europe. It is even possible to call up reporting functions from within the underlying OLTP R/3 application.

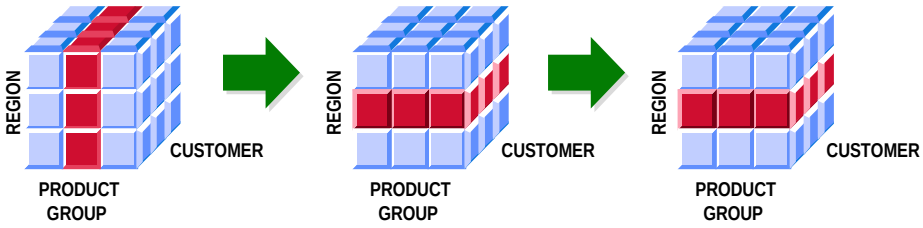

Fig. 6: Slice and Dice

Slice and dice operations are possible because data are stored in InfoCubes in multidimensional form; i.e. InfoCubes contain key figures plus characteristics. The OLAP Processor uses these characteristics to present the various views of key figures and of dynamically computed figures (derived key figures) in accordance with user needs.

Business Analysis Library

The Business Analysis Library is a collection of powerful analysis functions. These functions are used to convert the data contained in InfoCubes into business information of relevance to the end-user's needs. Together with slice and dice capabilities, the analysis functions of the Business Analysis Library are the "building blocks" on which the rich reporting and analysis features of Business Information Warehouse are based:

- **Aggregate functions** include:
  sum, count, count distinct, min, max, first, last, average, average over period sum/count.
  Many types of data need to be viewed at an "aggregate" level, rather than at their greatest level of detail, i.e. sales of an individual product need to be viewed at regional or national level.

- **Comparison functions** include:
  difference, ratio, percent, percent difference, share, correlation.
  Comparison functions are essential for tasks such as the breakdown of the shares of sales by region, or for displaying the difference between budgeted costs and actual costs.

- **Sequence functions** include:
  sort, cumulative sum, tertiles, quartiles, top/last n, top/last n%, classification, dual classification, ABC analysis.
  Sequence functions are needed for sorting by a particular characteristic, e.g. costs by region, and also for answering complex questions such as "Which customers account for 80% of turnover?"

- **Exception conditions** include thresholds on:
  absolute values, top / last n, top / last n%, trends, etc.
  Exception conditions are used, for instance, to identify a sudden drop in sales or to highlight other unusual deviations from expected values.

- **Financial functions** such as currency conversion, business period conversion, and internal business volume elimination.
  Financial functions support complex business structures. In international companies, for example, currencies need to be converted before monetary amounts can be compared or consolidated. With Business Information Warehouse, figures can be entered in local currency and conversion made according to the applicable rates.

Report types and layout structure

Reports are created by first selecting characteristics and key figures of the underlying InfoCubes and derived key figures from the meta data repository. Derived key figures are shared between reports and thus ensure a consistent interpretation across reports employed for different areas of application. The second step in creating a report is to define a layout structure that determines where key figures and characteristics are to be displayed, i.e. in which columns and which rows.

Business Information Warehouse supports 3 types of reports:

- reports with a default layout structure,
- reports with one-dimensional layout templates, and
- reports with two-dimensional layout templates.

Reports with a default layout structure are often used for exploratory tasks. They are easy and fast to define and have a simple layout structure: the characteristics' values are displayed as row headers and selected key figures as columns or vice versa.

Templates can even be used to standardize calculation methods across the enterprise

More sophisticated layout capabilities are supported by means of templates. As the name suggests, these templates can be re-used and modified. It is possible to use existing templates or to create new ones. Templates can be an extremely useful way not only of creating sophisticated reports with very little effort but also for standardizing the way certain indicators or parameters, such as contribution margin, are calculated throughout an enterprise.

A one-dimensional template allows specification of characteristics as column headers below which the key figures are arranged. The row headers are the same as for ad-hoc reports: characteristics that determine the actual slice and dice level. A typical use of a one-dimensional template would be to display budgeted costs and actual costs - and the difference between the two.

A two-dimensional template allows the user to specify arbitrary characteristics for the row headers and column headers, even individual calculations for each cell. A template of this kind is typically used for well-defined company standard reports such as the balance sheet.

Within all three reports, support is provided for all navigation and analysis features, including slice and dice, the computation of derived key figures and the application of business analysis functions.

Report Catalog Browser

Given such powerful reporting features and the large number of different reports that can be defined by both the individual user and his organization as a whole, one question remains: how does the user find the information he needs quickly and easily?

Business Information Warehouse's Business Explorer helps the user organize information according to his personal requirements by means of the Report Catalog Browser component.

Clusters of favorite reports

Studies have shown that most users need only a limited set of reports on a regular basis. The Report Catalog Browser allows the user to create and maintain clusters of favorite reports, each cluster grouped around a heading of the user's choice.

Favorite reports can be displayed with the click of a mouse

Reports can be defined as favorites or rearranged within and between clusters very simply, by drag and drop. Favorite reports are displayed using a new presentation style optimized for clarity and fast activation. When the user touches a report title, a structural preview and a textual abstract of the report are displayed, helping the user decide whether the report provides the information he requires. If the report has identified an exception this is flagged to the user by means of special highlighting.

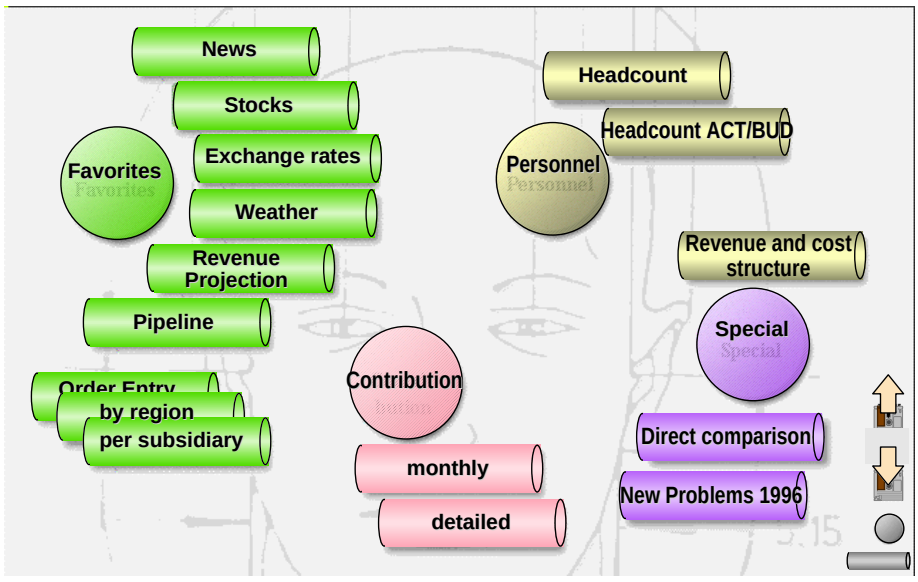

Fig. 7: Favorite Reports

Catalogs of standard reports

In addition to the clusters of personal favorites, the Report Catalog Browser supports report catalogs in the form of nested folders. Catalogs are available for pre-defined standard reports used by the individual user, by groups of users and by the company as a whole. The latter is derived from a SAP-defined report catalog and tailored to the specific user organization's needs. The end-user can select and execute any report from the catalogs for which he holds authorization. Catalogs, catalog reports and access rights are specified using the Administrator's Workbench.

SAP-supplied standard reports

SAP has created standard reports of immediate relevance to decision-makers

Building on its proven business expertise, SAP has created predefined standard reports tailored to the information needs of specific types of user, such as the financial controller or the sales manager. These standard reports are described in more detail in a separate paper.

Security

Access to sensitive data can be controlled at several different levels

Access to reports that contain sensitive data needs to be restricted to authorized persons only. Business Information Warehouse provides extensive access control on different levels. Access rights may be defined for a report as a whole, for certain key figures (e.g. salary in a personnel InfoCube), and even for specific values of a characteristic (e.g. a certain cost center). Access rights can be specified or modified by using the Administrator's Workbench.

Report presentation

Business Explorer offers a variety of options for displaying report views. For pre-defined standard reports, the main method is to display results in MS Excel, exploiting Excel's powerful layout capabilities. Business Explorer's user interface elements are embedded in Excel and comply with the Microsoft Office standard. As Excel is a widely used application, many end-users can make use of their existing skills and will not have to become accustomed to an unfamiliar environment.

Presentation of results in MS Excel for the creation of sophisticated documents for further distribution

Presentation in Excel also allows users to take a report view and manipulate it in the form of a document: i.e. they can add their own personal calculations or notes, charts and graphics, combine different reports in one workbook, or distribute it via e-mail.

For more ad-hoc and exploratory tasks (which typically have simpler layout requirements) Business Information Warehouse also provides an additional presentation option which is independent of Excel. However, all three types of report layouts can be used in conjunction with either presentation option.

Business Information Warehouse supports reporting in intranet and Internet environments

In Internet and intranet environments the Business Explorer runs under a Web Browser for report catalog browsing and report execution.

In conclusion, Business Explorer provides an easy-to-understand environment for exploring and exploiting data. It helps the user to access the information he needs with a minimum of effort and time.

Performance

Performance is of vital importance within any data warehouse if users are to get full benefit from available information. Business Information Warehouse's architecture includes several features which ensure high performance.

- **Dedicated server:**
  Business Information Warehouse resides on it own dedicated server. Like the R/3 OLTP system there is a number of configuration options to provide highly scalable performance as the number of users grows. Business Information Warehouse makes use of the flexibility of R/3 three-tier client/server technology. All Business Information Warehouse server components may reside on the same hardware system, or the database can be placed on a dedicated database server, and multiple application servers can be added as the number of users grows.

- **InfoCube Summary Levels:**
  InfoCubes are designed for high query performance. They can easily grow to some ten gigabytes in size. Naturally, a user does not want to browse through this amount of data. Instead higher aggregated levels are required. For common aggregations, the Data Manager maintains materialized aggregates (summary levels) which dramatically speed up queries. Summary levels are transparent to the OLAP Processor and therefore to the end-user. The administrator can add or remove summary levels according to the usage pattern of the user community.

- **Persistent Report Cache:**
  Once the OLAP Processor has queried data from InfoCubes it maintains a report cache in memory which includes additionally derived key figures. A cached report, i.e. the aggregated data of a report at a given point in time, can be stored persistently on disk. This option is very useful for standard reports on data which only change very gradually and which are accessed frequently. When a report of this type is executed, it reloads the report data from the persistent cache rather than querying the database.

- **Batch reports:**
  Long running reports can be scheduled as a batch process. The result is stored in the persistent cache, ready for execution e.g. after a monthly data load for a consolidated company report.

- **Excel Preview:**
  Report views presented in Excel can be stored as documents and presented as such. The OLAP Processor is only activated when the data need to be refreshed or a new view needs to be computed.

Openness

Based on the Business Framework, Business Information Warehouse is an end-to-end yet open system. Several interfaces are provided to allow for custom extensions:

- A BAPI on top of the OLAP Processor for presentation of and navigation through a report. This interface might be used by a specialized GUI builder, e.g. for display of revenue figures on a geographical map.
- Access to plain Business Information Warehouse data and meta data. All Business Information Warehouse data are stored in transparent tables. However, not all business rules handled by the OLAP Processor are available on this level. For instance, a data-mining tool can access an InfoCube's data, load data into its specialized data store and then perform exploratory analysis on that data.
- The load interface BAPI of the Staging Engine. As described in section Staging Engine, this interface is intended for loading non-SAP data to Business Information Warehouse, e.g. through custom extract programs, or through a third-party data warehouse management tool.

In addition to these interfaces for other tools, there is a number of user exits in the Staging Engine for custom transformation and validation programs during production data load. It is therefore possible for the user to apply his own special-purpose transformation rules, such as the conversion of heterogeneous product codes used in diverse OLTP sources into a uniform product coding system.

In other words, although ideally suited to an SAP environment, Business Information Warehouse is not a proprietary solution. It can be combined with additional tools and non-SAP data sources.

Business Information Warehouse configurations

Data marts can be set up for specialized information needs

Business Information Warehouse supports a variety of configurations. The standard approach is to create a central, enterprise-level Business Information Warehouse (a data warehouse). It also possible to create multiple, independent Business Information Warehouses for the specialized needs of individual business units or departments (data marts). Staged extract and consolidation from multiple Business Information Warehouses to a central Business Information Warehouse is also foreseen. The Business Information Warehouse architecture allows consolidation on several different levels. Business Explorer provides a single point of end-user entry for one or more Business Information Warehouse servers.

Conclusion

Business Information Warehouse is a key Business Component of the SAP Business Framework. It allows user-friendly yet sophisticated business reporting and decision support. It is a complete, high-quality solution based on leading client/server technology and SAP's acknowledged business expertise.

In particular, Business Information Warehouse is the following:

- A fast solution: as a Business Component of SAP's Business Framework, Business Information Warehouse can be implemented at high speed and low cost.
- An open solution: Business Information Warehouse is not limited to R/3 or R/2. Through open interfaces, Business Information Warehouse can be combined with other non-SAP data sources and non-SAP tools.
- A functionality-rich solution: Business Information Warehouse offers a wide range of powerful reporting and analysis features for effective exploration and exploitation of data.
- An adaptable solution: if changes are made to business processes or to the IT environment, Business Information Warehouse is designed to allow adaptation during ongoing system operation. In addition, the Administrator's Workbench provides effective support for easy Business Information Warehouse maintenance and quick implementation of changes.
- A robust, business-driven solution: based on SAP's proven expertise in R/3 client/server technology and in real-world business processes, Business Information Warehouse is built to answer the specific information needs of decision-makers in all industries, reliably and effectively.