Decision support systems

Overview

Decision support systems (DSS) are a subset of computer-based information systems (CBIS). The general term 'computer-based information systems' is a constellation of a variety of information systems such as office automation systems, transaction processing systems, management information systems and management support systems. Management support systems consist of DSS, expert systems and executive information systems. In the early 1970s, scholars in the CBIS area began to recognize the important roles information systems play in supporting managers in their semi-structured or unstructured decision-making activities. It was argued that information systems should exist only to support decisions, and that the focus of the information systems development efforts should be shifted away from structured operational control to unstructured critical decisions in organizations. Decisions are irreversible and have far-reaching consequences for the rest of organizational life. The importance of effective decision making can never be overemphasized. Decision making is, in effect, synonymous with management.

1 Evolution of decision support systems

Since the first electronic general-purpose computer was put into full operation in the early 1940s, data-processing techniques have been continuously advancing. It was in the late 1950s that many organizations began to utilize transaction processing systems (TPS) or electronic data processing (EDP) systems to automate routine clerical tasks such as payroll, inventory and billing. In the 1960s, we witnessed the emergence of management information systems (MIS) with the development of database management systems for collecting, organizing, storing and retrieving data (see MANAGEMENT INFORMATION SYSTEMS (MIS)). MIS were developed to extract valuable management information by aggregating and summarizing massive amounts of transaction data and allowing user-interactive managerial queries. The inclusion of simple modelling and statistical methods as a component of MIS permits computer systems to make routine (structured) decisions.

It was not until 1970 that scholars began to recognize the important roles computer-based information systems (CBIS) play in supporting managers in their semi-structured or unstructured decision-making activities (see DECISION MAKING AND IT/S). Since the 1970s, study of DSS has become an essential part of CBIS. In the 1980s, we witnessed another wave of information technologies, the artificial intelligence-based expert systems (ES), which are to replace and mimic human decision makers in making repetitive decisions in a narrow domain (see ARTIFICIAL INTELLIGENCE (AI); KNOWLEDGE-BASED SYSTEMS). During the mid 1980s, executive information systems (EIS) emerged as an important tool to serve the information needs of executives (see EXECUTIVE INFORMATION SYSTEMS). EIS provides timely and critical information which has been filtered and compressed for tracking and control purposes. The latest addition to CBIS is artificial neural networks (ANN). Neural network computing involves building intelligent systems to mimic human brain functions. ANN attempt to achieve knowledge processing based on the parallel processing method of human brains, pattern recognition based on experience, and fast retrieval of massive amounts of data (see NEURAL NETWORKS). Fuzzy logic, genetic algorithm, and intelligent agents are some of other intelligent techniques that can be used along with neural networks to improve the effectiveness of personal, group, and organizational decision making.

Table 1 summarizes an evolutionary pattern of CBIS and shows the focus of a CBIS from data and information to knowledge and wisdom. The critical information provided by EIS can be used to identify various symptoms of malfunctioning organizational activities in each functional department (see ORGANIZATIONAL INFORMATION AND KNOWLEDGE). These symptoms can be the basis of diagnosing managerial problems. Decision support systems (DSS) are human–computer decision-making systems to support managerial judgements, and intuitions to solve managerial problems by providing necessary information, generating, evaluating and suggesting decision alternatives. Most organizational problems need a combination of quantitative and qualitative data processing. EIS are to deal with those organizational problems that can be better solved by qualitative data processing. Other subsets of CBIS such as TPS and MIS provide data into DSS to be processed by DSS models and managerial judgements.

Table 1 Taxonomy of knowledge

		Technology analogy	*Management*	*Metaphor*
Data	EDP	Elements: H_2O, yeast bacteria, starch molecules	Muddling through	KNOW - NOTHING
Information	MIS	Ingredients: Flour, sugar, spices, fixed recipe for bread only (OR/MS) type	Efficiency (Measurement + search)	KNOW - HOW
Knowledge	DSS, ESS, AI	Choose among different recipes for bread	Effectiveness (decision making)	KNOW - WHAT
Wisdom	HSM, MSS	Why bread and not croissant	Explicability (judgment)	KNOW - WHY

Source: Zeleny (1987: 60)

Definition of decision support systems

Drawing on various definitions that have been suggested (Alter 1980; Bonczek *et al*. 1981; Keen and Scott-Morton 1978; Sprague and Carlson 1982) a DSS can be described as a computer-based interactive human–computer decision-making system that:

1 supports decision makers rather than replaces them;
2 utilizes data and models;
3 solves problems with varying degrees of structure: (a) non-structured (unstructured or ill-structured) (Bonczek *et al*. 1981); (b) semi-structured (Keen and Scott-Morton 1978); (c) semi-structured and unstructured (Sprague and Carlson 1982);
4 focuses on effectiveness rather than efficiency in decision processes (facilitating decision processes).

2 Architecture of decision support systems

As shown in Figure 1, a DSS consists of two major sub-systems – human decision makers and computer systems. Interpreting a DSS as only a computer hardware and software system is a common misconception. An unstructured (or semi-structured) decision by definition cannot be programmed because its precise nature and structure are elusive and complex (Simon 1960). The function of a human decision maker as a component of DSS is not to enter data to build a database, but to exercise judgment or intuition throughout the entire decision-making process (see DECISION MAKING AND IT/S).

Imagine a manager who has to make a five-year production planning decision. The first step of the decision-making process begins with the creation of a decision support model, using an integrated DSS program (DSS generator) such as Microsoft Excel, Lotus 1-2-3, Interactive Financial Planning Systems (IFPS)/Personal or Express/PC. The user interface sub-system (or dialogue generation and management systems) is the gateway to both database management systems (DBMS) and model-based management systems (MBMS). DBMS are a set of computer programs that create and manage the database, as well as control access to the data stored within it. the DBMS can be either an independent program or embedded within a DSS generator to allow users to create a database file that is to be used as an input to the DSS. MBMS is a set of computer programs embedded within a DSS generator that allows users to create, edit, update, and/or delete a model. Users create models and associated database files to make specific decisions. The created models and databases are stored in the model base and database in the direct access storage devices such as hard disks. From a user's viewpoint, the user interface sub-system is the only part of DSS components with which they have to deal. Therefore, providing an effective user interface must take several important issues into consideration, including choice of input and output devices, screen design, use of colours, data and information presentation format, use of different interface styles, etc. Today's decision support system generator provide the user with a wide variety of interface modes (styles): menu based interaction mode, command language style, questions and answers, form interaction, natural language processing based dialogue, and graphical user interface (GUI). GUIs use icons, buttons, pull-down menus, bars, and boxes extensively and have become the most widely implemented and versatile type. The interface system allows users access to:

(1) The data sub-system: (a) database (b) database management software; and (2) The model sub-system: (a) model base (b) model base management software.

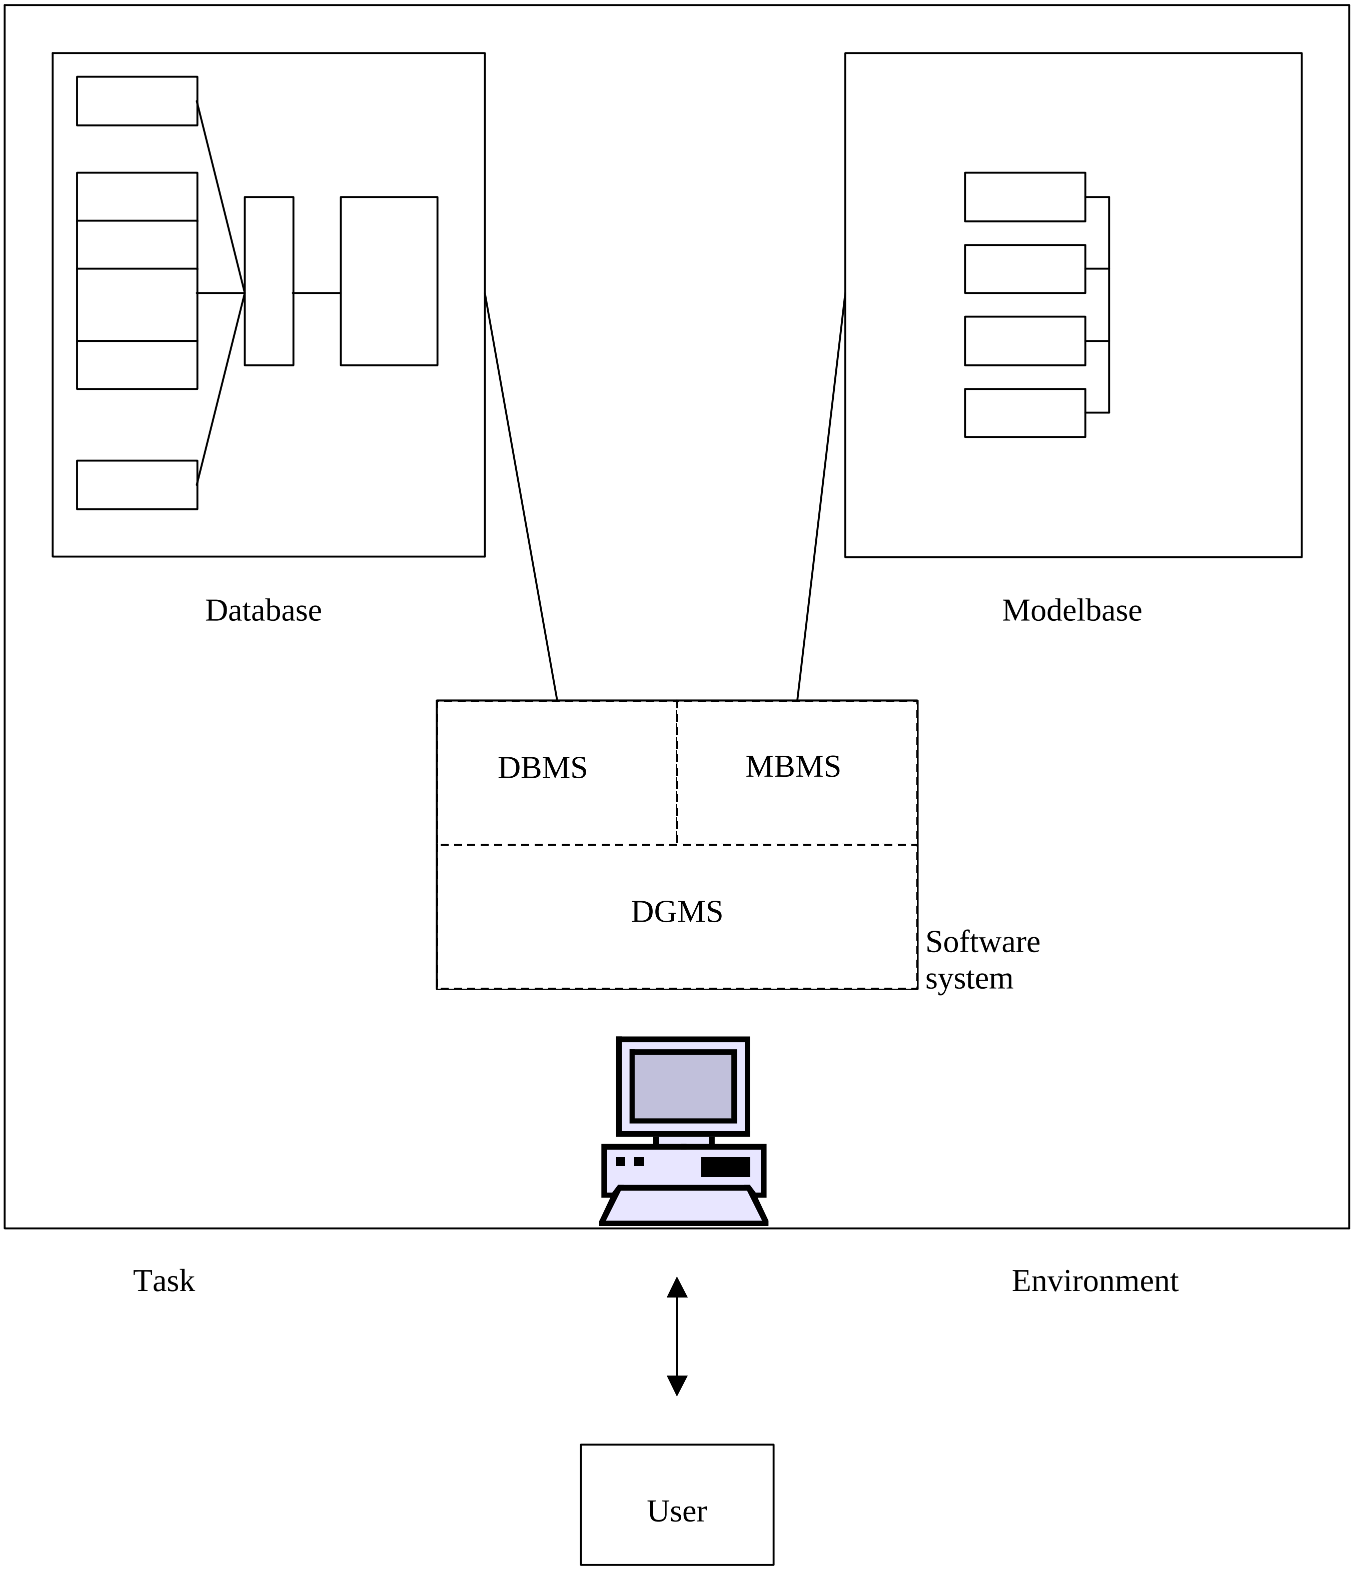

Figure 1 Components of decision support system
Source: Sprague and Carlson (1982: 29)

Decision-making processes and functions

DSS is distinguished from MIS in terms of focusing on effectiveness, rather than efficiency in decision processes (facilitating decision processes). An important performance objective of DSS is to support all phases of the decision-making process (Sprague and Carlson 1982). Simon's model of decision making describes human decision making as having three major steps: intelligence, design and choice. The term 'support' implies many different activities and tasks in each stage of the decision-making process (see DECISION MAKING AND IT/S).

In the intelligence stage, human decision makers play an important role in defining problems to be solved, based on the raw data obtained and information processed by transaction processing systems (TPS)/management information systems (MIS). Alter (1980: 73) suggests seven different types of DSS, based on the 'degree of action implication of DSS outputs' (that is, the degree to which the DSS's output could directly determine the decision). Among them, the following three DSS types are especially useful in the intelligence stage: (1) file drawer systems which allow online access only to particular data items; (2) data analysis systems which permit user(s) to retrieve, manipulate and display current and historical data; and (3) analysis information systems which manipulate the internal data from TPS and augment the internal data with external data using statistical packages and other small models to generate management information.

The majority of DSS in use today are developed to generate and evaluate decision alternatives via 'what-if' analysis and 'goal-seeking' analysis in the design and choice stages. *Accounting* models facilitate planning by calculating the consequences of planned actions on estimate-of-income statements, balance sheets and other financial statements. *Representational* models estimate the future consequences of actions on the basis of partially non-definitional models, including all simulation models. *Optimization* models generate the optimal solutions. *Suggestion* models lead to a specific suggested decision for a fairly structured task. Such systems perform mechanical calculations and leave little role for managerial judgement.

3 Decision support system sub-specialities

As Figure 2 shows, the study of DSS consists of the following three important groups of research areas:

1 Developing a specific DSS (labelled 'A' in Figure 2). Over the past three decades (1970-2000), about 500 specific functional DSS applications have been developed and published in English language journals (labelled 'B').
2 Developing DSS theory:
(a) developing theory on decision makers, data, model and interface (dialogue) (labelled 'F'–'I');
(b) developing theory on design, implementation and evaluation (labelled 'C', 'D' and 'E').
3 Study of contributing disciplines (labelled 'J').

The first group of research areas, labelled 'F'–'I', is based on the architecture of DSS heavily influenced by Sprague and Carlson while the second group of research areas, labelled as 'C'–'E', is influenced by the organizational perspectives of Keen and Scott-Morton (1978). The third group of research is DSS application development, labelled 'A' and 'B'.

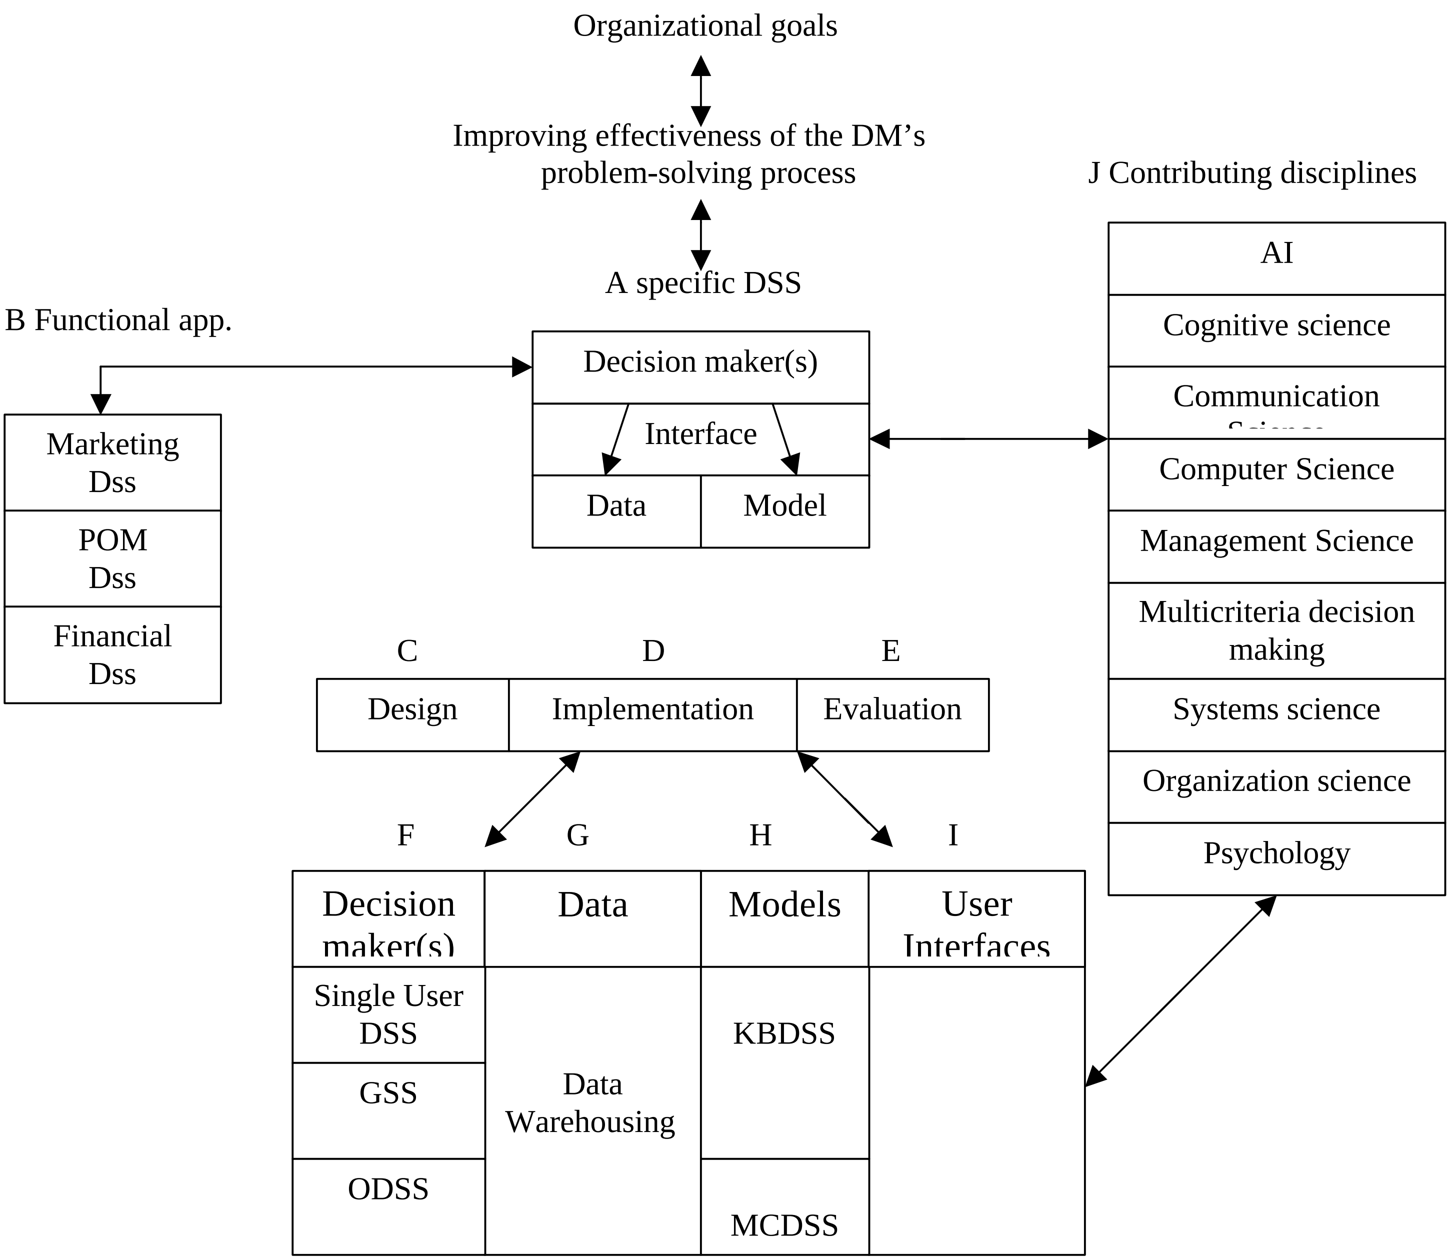

Figure 2 Theory applications and contributing disciplines of decision support system

Data/model management

Since model and data management in DSS are inseparable subjects, many DSS researchers continue to focus on both fields of data and model management. Data are facts which result from the observation of physical phenomena such as daily production quantity, daily sales quantity and inventory level of product A. A database is a collection of interrelated files. Database management systems are computer programs which are primarily concerned with managing a large amount of data in a physical storage such as hard disks and creating, updating and querying databases in an optimal way.

Data management in DSS is a necessary function primarily useful in the intelligence stage of the decision-making process, but not sufficient to support design and choice stages of decision-making processes. To adequately support these stages, DSS should be able to include the following activities: projection, deduction, analysis, creation of alternatives, comparison of alternatives, optimization and simulation (Sprague and Carlson

1982). In performing these essential tasks, DSS utilizes many types of management science/operations research (MS/OR) models. They include linear programming, integer programming, network models, goal programming, simulation and statistical models and spreadsheet modelling. (For a complete list of models used in DSS, see Eom *et al*. 1998.) All these models are stored in the model base. Model-based management systems are computer programs used as a part of a DSS generator to build models, restructure models and update models. In association with model management, multiple criteria decision making (MCDM) model embedded DSS and knowledge-based DSS have emerged recently as important DSS research sub-specialities (Eom and Min 1992; Eom 1996).

User interface sub-systems

The functions of the user interface (dialogue generation and management) sub-system is to:

1 allow the user to create, update, delete database files and decision models via database management systems and model-based management systems;
2 provide a variety of input and output formats. The formats include multi-dimensional colour graphics, tables and multiple windows in a screen;
3 provide different styles of dialogues (such as graphical user interfaces, menus, direct command languages, form interaction, natural language interaction, and questions and answers).

Research in user interface sub-systems has investigated several important issues in the designing, building, and implementing of a user interface. They include data /information display formats (for example, tabular versus graphics), cognitive and psychological factors, use of multimedia (multiple media combined in one application) and hypermedia (documents that contain several types of media linked by association), 3-dimemensional user interfaces, virtual reality and its impact on decision making, geographical information systems, and natural language processing.

Knowledge-based decision support systems

Another important emerging DSS sub-speciality is the study of knowledge-based decision support systems (KBDSS), which are hybrid systems of DSS and ES that help solve a broad range of organizational problems. In integrating DSS and ES, two basic approaches are discernible and labelled expert support systems (ESS) and intelligent support systems (ISS) (King 1993). The key differences between these two systems are as follows. ESS are to replace human expertise with machine expertise, while ISS are to amplify the memory and intelligence of humans and groups (King 1993). A broad range of real-world managerial problems can be better solved by using the analysis of both quantitative and qualitative data. Few would disagree with the notion that there are considerable benefits from integrating DSS and ES. The new integrated system (ESS or ISS) can support decision makers by harnessing the expertise of key organizational members. A bottleneck in the development of knowledge-based systems such as ESS is knowledge acquisition, which is a part of knowledge engineering – the process includes representation, validation, inferencing, explanation and maintenance.

Group DSS/Group support systems/Electronic meeting systems

Single user DSS and group DSS can be distinguished in many different ways in terms of purpose and components (hardware, software, people, procedures). First, group DSS and single user DSS have distinguishable purposes. DeSanctis and Gallupe (1985: 3) define a group DSS as 'an interactive computer-based system which facilitates solution of unstructured problems by a set of decision makers working together as a group'. A single user DSS can be simply defined by replacing 'a set of decision makers working together as a group' with 'a decision maker' (see TEAMS AND TEAMWORK).

Second, to support a set of decision makers working together as a group, group DSS have special technological requirements of hardware, software, people and procedures. Each member of the group usually has a personal computer which is linked to the personal computers of other group members and to one or more large public viewing screens, so that each member can see the inputs of other members or let other members see their work. Group DSS software also need special functional capabilities, in addition to the capabilities of single user DSS software, such as anonymous input of the user's ideas, listing group members' ideas, voting and ranking decision alternatives. The people component of group DSS should include a group facilitator, who leads the session by serving as the interface between the group and the computer systems.

Computer based information systems to support group activities have been conducted under the titles of group decision support systems (GDSS), computer-supported cooperative work (CSCW), group support systems

(GSS), collaboration support systems (CSS), and electronic meeting systems (EMS). GDSS have focused on decision making/ solving problems, while CSCW provide primarily a means to communicate more efficiently. However, these two types of systems, decision making focused systems and communication-focused systems, are becoming indistinguishable. There seems to be a consensus that GSS are a broad umbrella term referring to the collective of computer assisted technologies used to aid group efforts directed at identifying and addressing problems, opportunities, and the issues (Jessup and Valacich 1993).

Organizational decision support systems

An organizational decision support system is defined as 'a DSS that is used by individuals or groups at several work stations in more than one organizational unit who make varied (interrelated but autonomous) decisions using a common set of tools' (Carter *et al.* 1992: 19). According to the same source, an important goal of organizational DSS is to provide 'the glue that holds a large organization together and keeps its parts marching to the beat of the same drummer toward common goals'. The two key factors to achieving these outcomes are: (1) transmittal of consistent, timely information up and down the organizational hierarchy in forms that are appropriate to each decision maker; and (2) a set of decision-aiding models that use this information and that are appropriate for the decisions being made by each decision maker.

4 Sub-specialities based on organizational perspectives

Decision support system design

DSS design is the process of identifying the key decisions through decision analysis, specifying requirements of each DSS component to support key decisions identified through decision analysis.

DSS are designed and implemented to support organizational as well as individual decision making. Without a detailed understanding of decision-making behaviour in organizations, 'decision support is close to meaningless as a concept' (Keen and Scott-Morton 1978: 61). Organizational scientists classify organizational decision making in terms of several schools of thought: (1) the rational model which focuses on the selection of the most efficient alternatives, with the assumption of a rational, completely informed single decision maker; (2) the organizational process model which stresses the compartmentalization of the various units in any organization; (3) the satisficing model which reflects 'bounded rationality' to find an acceptable, good enough solution; and (4) other models.

Decision support system implementation

Use of some computer-based information systems such as TPS and MIS are, in most cases, mandatory. But decision support systems are voluntary systems. In regard to voluntary systems, DSS implementation research has been important for ascertaining the influence of success factors of DSS implementations. DSS implementation researchers are investigating the relationship between user-related factors and implementation success. User factors include cognitive style (the characteristic ways individuals process and utilize information to solve problems), personality (the cognitive structures maintained by individuals to facilitate adjustment to events and situations), demographics (age, sex and education), and user-situation variables (training, experiences and user involvement) (Alavi and Joachimsthaler 1992). future implementation research should be directed toward the development of causal models of user-related implementation factors. Furthermore, it is suggested that DSS researchers shift the research focus from user-related factors to the contextual variables. An important assumption on which the DSS implementation research is based is that DSS are voluntary systems. A recent survey of DSS suggests that an increasing number of DSS have become a strategic tool for organizational survival (Eom et al 1998). Thus, these systems are no longer voluntary ones. Future DSS implementation research must take this changing nature of DSS from voluntary systems to mandatory survival tools. Consequently, individual differences, cognitive styles, personality, demographics, and user-situational variables may become less critical success factors. Shifting the focus of implementation research from user-related factors to task-related, organizational, and external environmental factors may be necessary to reflect the changing decision environment in which organization must survive and prosper (Eom 2000).

Decision support system evaluation

Evaluation of DSS is concerned with analysing costs and benefits of DSS before and after DSS development and implementation. The unique nature of DSS evaluation is that although some DSS provide substantial cost saving and profit increases, measurements of benefits of DSS have been problematic as quantification of the positive impacts of improved decision process is difficult. Therefore, DSS evaluation research deals with the following methodologies: decision outputs, changes in the decision process, changes in managers' concepts of the decision situation, procedural changes, cost/benefit analysis, service measures and managers' assessment of the system's value (Keen and Scott Morton 1978).

5 Application development research

DSS application development is the fruit of DSS study. Theories developed from DSS research must be assimilated into the DSS development process. The next section briefly introduces the current status of DSS application development research in corporate functional management areas.

Applications of decision support systems

According to a survey (Eom *et al*. 1998), computer-based DSS are widely applied in both profit making (about 72 per cent) and non-profit organizations (about 28 per cent). In corporate functional management fields, production and operations management contain the largest number of application articles, followed by management information systems, marketing, finance, strategic management and multifunctional areas. Two functional fields are relatively minor fields for DSS application: international business and accounting/auditing. Table 2 lists some of the important application examples from the survey. Refer to http://cstl-hcb.semo.edu/eom/ORINSIHT.HTM for a more detailed classification of 271 DSS articles by application areas.

Table 2. DSS Applications in Corporate Functional Management

1. **Accounting/ Auditing**

Auditing health insurance claims
Estimating pencil manufacturing cost
Stochastic cost-volume-profit analysis

2 **Finance**

Asset-liability management
Cash management and debt planning
Capital budgeting
Evaluating financial risks
Financial analysis and diagnosis
Funding strategic product development
Locating banks
Managing portfolio
Planning in mergers and acquisitions
Selecting R&D project
Structuring optimal lease
Real estate appraisal and investment
Setting interest rates for money market deposit accounts
Small business financial planning

3 **Human Resources Management**

Manpower planning
Massive personnel assignment
Resolving labor management dispute
Tracking critical human resources information

.4 **International Business**

Allocating investment funds in MNCs
Analyzing international investment options
Planning global logistics
Planning global marketing/ production/ distribution

5 **Information Systems**

Data Communication

Evaluating LAN topologies
Designing a fiber optic WAN

DSS Generators

Application (domain)-independent system for supporting

- group decision making
- massive data retrieval and extraction
  - MCDM problems
  - consensus reaching processes
  - generating and exploring alternatives
  - decision conferencing
  - multicultural/ multilingual communication
- small groups decision making under uncertainties
- modeling tasks
- the Japanese style of group decision making

Systems Analysis, Design, Development & Evaluation

Designing online retail banking systems
Evaluating MIS Effectiveness
Joint application development
Optimizing MIS project portfolio mix
Systems analysis and planning
Strategic planning of system development

Information Resources Management

Planning information systems security
Supercomputer acquisition and capacity planning

6 **Marketing**

Allocating retail space in retail outlets
Competitive pricing and market share
Designing freight networks and integrated distribution systems
Distribution planning
Logistics planning and vehicle scheduling
Managing hazardous material shipments
Media-planning for advertising agencies
Measuring direct product profitability in retail merchandising
Selecting telemarketing sites

7 **Production and Operations Management**

A. Planning for Demand

Designing sampling procedures for accurate estimation of electrical demand
Forecasting rotable aircraft parts demand

B. Master Scheduling

Production planning
Production planning and control

C. Operations Scheduling and Control

C1. Manufacturing Industry

Designing electronics test facilities
Managing manufacturing logistics and dispatch
Operations scheduling
Operations control
Process planning
Planning for offshore drilling
Project scheduling and termination

C2. Service Industry

Airline arrival slot allocation system
Integrated management of fleet
Scheduling courier vehicle, flight, and day-off assignment for airlines cockpit crew
Train dispatching

D. Operations Design

Evaluating personal and machine productivity
Gasoline blending and planning refinery operations
Managing quality-control process
Product design
Selecting machines in integrated process planning
TQM consultant

E. Capacity Planning

Automating factory
Capacity planning
FMS configuration

FMS Scheduling and control
Just-in-time production
Line-balancing
Personnel assignment within FMS production cell
Setup reduction investment
Selecting robot

F. Inventory Planning

Inventory planning and control
Material requirement planning

G. Resource Management

Designing materials management processes
Procurement in business volume discount environments
Purchasing and material management of large projects
Selecting suppliers

8 **Strategic Management**

External environment and industry analysis
Strategic analysis of
- mergers and acquisitions
- multi-level (corporate, division, department) and mulitifuntional corporate planning
- product/ market position

Grand strategy selection (managing a portfolio of new product development research projects, and terminating projects)
Strategy control and evaluation
Decision conferencing for Strategic planning
Integrated strategic planning process support
Managing organizational crisis

9 **Multifunctional Management**

Multi-refinery, multi-period capital investments
Planning for expanding refining capacity
Budgeting and manpower planning
Strategic production and distribution planning
Manpower and vehicle scheduling
Integrated multifunctional systems for chemical production
Supporting reciprocally interdependent decisions

6 The future of decision support systems

A host of new tools and technologies are adding new capabilities to DSS/ESS and will reshape DSS developments in organizations. They include hardware and mathematical software developments, artificial intelligence techniques, the data warehouse/multidimensional databases (MDDB), data mining, online analytical processing (OLAP), enterprise resource planning (ERP) systems, intelligent agents, telecommunication technologies such as World Wide Web technologies, the Internet, and corporate intranets.

Single user decision support systems

Ever-increasing computing power makes it possible to solve a large-scale mathematical optimization model in a fraction of a second. The size of the problem solvable by commercial software is virtually unlimited, only dependent upon the size of random access memory of computers and the user's patience. Moreover, several solvers are built into the spreadsheet programs such as Microsoft Excel and Borland's Quattro-Pro, along with the capabilities of linking to databases and graphical user interfaces.

With the increasing trend of national and global communication networking, single user DSS will increasingly become a part of organization-wide distributed decision-making (DDM) systems. The DDM system consists of several single user DSS that work together and independently to make a sequential decision such as joint production/marketing decisions (Rathwell and Burns 1985). DDM systems work as a mechanism for integrating a number of separate DSSs that coexist in an organization, facilitating group cooperation between several DSSs in a distributed environment, and meeting the specific needs of group planning and group decision making.

Notable developments that will significantly affect the future development of DSS are the data warehouse, data mining and intelligent agents. The data warehouse is a subject-oriented, integrated, time-variant, and non-volatile (read only) collection of a relational/multidimensional database (MDDB) optimized for decision support, which is separated from operational databases. MDDB organizes data as an n-dimensional cube so that users deal with multidimensional data views such as product, region, sales, time, etc. with a faster query response time. Data mining, also known as Knowledge Data Discovery, refers to discovering hidden patterns/trends/classes/insights/relationships from data, and it attempts to automatically extract knowledge from the in large databases, either in the data warehouse or elsewhere (e.g., spreadsheets, transaction processing system files, etc.) (see DATA WAREHOUSING and DATA MINING).

Intelligent agents (known also as intelligent interfaces, or adaptive interfaces) research is an emerging interdisciplinary research area involving researchers from such fields as expert systems, DSS, cognitive science, psychology, databases, etc. Intelligent agents research has contributed to the emergence of a new generation of active and intelligent DSS and EIS. The active DSS will be equipped with the tools (stimulus agents) that will act as experts, servants, or mentors to decide when and how to provide advice and criticism to the user, while the user formulates and inquires about its problems under the continuous stimulus of electronic agents. The essence of active decision support activities includes monitoring decision making processes and stimulating creative ideas through carrying out insightful conversations with decision makers.

Knowledge-based decision support systems (Intelligent DSS)

An increasing number of systems are incorporating domain knowledge, modelling, and analysis systems to provide users the capability of intelligent assistance. Knowledge base modules are being used to formulate problems and decision models, and analyse and interpret the results. Some systems are adding knowledge-based modules to replace human judgments. Managerial judgements have been used to ascertain (assess) future uncertainty and to select assumptions on which decision models can be based. Some decisions are both knowledge and data intensive. Consequently, a large amount of data usually requires considerable efforts for their interpretation and use.

The knowledge-based DSS include a knowledge management component which stores and manages a new class of emerging AI tools such as machine learning and case-based reasoning and learning. These tools can obtain knowledge from prior data, decisions and examples (cases), and contribute to the creation of DSS to support repetitive, complex real-time decision making. Machine learning refers to computational methods/tools of a computer system to learn from experience (past solutions), data and observations, and consequently alter its behaviour, triggered by a modification to the stored knowledge. Artificial neural networks and genetic algorithms are the most notable approaches to machine learning.

The role of knowledge-based DSS should be to allow experts to broaden and expand their expertise, not to narrow it down. Zeleny suggests the important future direction of knowledge-based DSS development in this way (Zeleny 1987: 65):

> Trends toward narrow specialization of experts did take place in the society. The task of supportive expert systems should be to counteract such expertise-destructive trends, not to amplify them further. … Observe that so-called ‘knowledge explosion’ (also ‘information explosion’ or ‘information society’) are misnomers limited to the structural (‘surface’) knowledge only. …Even today we can aim for a ‘renaissance man’, especially with the help of expert systems constructed to move in the direction opposite of specialization, towards reintegration of knowledge, overcoming specialization, negating the ‘experts’ themselves.

The World Wide Web and Group/Organizational/Global DSS

The World Wide Web is increasingly being used as the client-server platform of many business organizations due to its network and platform-independence and very low software/installation/maintenance costs. More and more groupware will be inextricably tied to Internet technology. Especially, the World Wide Web is becoming an infrastructure for the next generation of decision support systems and groupware applications. Many groupware products, such as Lotus Development's Domino and Microsoft's Exchange, are integrating more Internet protocols into them. Microsoft's next version of Office suite is expected to completely remove the boundaries between the World Wide Web and groupware. Many companies are applying groupware technology to increase business-to-business collaborations (e.g. collaborations among the company, its customers, and its suppliers, a.k.a. superworkgroup software) over intranets and extranets (see COMPUTER-SUPPORTED COOPERATIVE NETWORK). Another development in the information systems area is the growing importance of enterprise resources planning (ERP) systems. ERP systems are a new generation of information systems packages that integrate information and information-based processes within and across functional areas in an organization. ERP has focused primarily on processing of transaction data resulting in the creation of the extensive, organizational databases of an organization that may consist of individual business units across the globe. The extensive databases created by the ERP system provide the platform for decision support, data warehousing, data mining, and executive support systems. integrated solutions provided by the ERP system are attributable to the use of the common database.

As we enter the age of the global village where geographical and temporal boundaries are shrinking rapidly, global DSS are emerging as the new frontiers in management information systems area. Over the next decade, DSS will focus on teams, work groups, and distributed, decentralized organizational structures (King 1993). Consequently, many organizations will increasingly design and implement group/organizational/global DSS. Global management support systems (MSS) will emerge as a key element in management decision making and as an essential weapon against global competitors. Supporting global business activities is becoming a most important and extremely complex task. To effectively cope with multinational managerial problems such as multiple currency management, foreign exchange risk management, global tax management and global consolidated reporting, global DSS are not enough. It is essential to develop an integrated global MSS which integrates EIS, artificial neural networks, ES with knowledge base captured from numerous experts in the same subject area as well as from a variety of specialists in international financial management, international accounting, international tax areas, and so forth.

SEAN B. EOM
SOUTHEAST MISSOURI STATE UNIVERSITY

Note

I am grateful to Professor R.H. Sprague of the University of Hawaii and Professor M. Zeleny of Fordham University for permission to use their works.

Further reading

(References cited in the text marked *)

* Alavi, M. and Joachimsthaler, E.A. (1992) 'Revisiting DSS implementation research: a meta-analysis of the literature and suggestions for researchers', *MIS Quarterly* 16 (1): 95–116. (A rigorous and quantitative review of the empirical DSS implementation literature.)

* Alter, S.L. (1980) *Decision Support Systems: Current Practice and Continuing Challenges*, Reading, MA: Addison-Wesley. (This book provides case studies and detailed information about usage and implementation patterns of DSS.)

* Bonczek, R.H., Holsapple, C.W. and Whinston, A.B. (1981) *Foundations of Decision Support Systems*, New York: Academic Press. (A classic DSS textbook which suggests an architecture of a language system, a knowledge system and a problem processing system.)

* Carter, G.M., Murray, M.P., Walker, R.G. and Walker, W.E. (1992) *Building Organizational Decision Support Systems*, Boston, MA: Academic Press. (This book provides step-by-step guidance to those who are interested in building organizational DSS.)

*Dennis, A.R., George, J.F., Jessup, L.M., Nunamaker, J.F., Jr., and Vogel, D.R. (1988) 'Information technology to support electronic meetings', MIS Quarterly 12(4): 591-624. (Electronic meeting systems are proposed to support group meetings, which may be distributed geographically and temporally.)

* DeSanctis, G. and Gallupe, B. (1985) 'Group decision support systems: a new frontier', *Data Base* 16 (2): 3–10. (An introduction to hardware, software, people and procedures of group DSS systems.)

DeSanctis, G. and Gallupe, B. (1987) 'A foundation for the study of group decision support systems', *Management Science* 33 (5): 589–609. (This influential group DSS article provides foundational concepts for further group DSS research guidance and future directions.)

* Eom, S.B., Sang, M., Lee, S.M., Kim, E. and Somarajan, C. (1998) 'A survey of decision support system applications (May 1988–1994)', *Journal of the Operational Research Society* 49 (2): 109–20 . (A survey on 271 specific DSS applications.)

* Eom, S.B. and Min, H. (1992) 'The changing role of multiple criteria in decision support systems', *Human Systems Management* 11 (3): 137–144. (An empirical study reporting that multiple criteria DSS have become 'mainstream' DSS research and have positioned themselves at the core of DSS study.)

* Eom, S.B. (1996) 'Mapping the intellectual structure of research in decision support systems through author cocitation analysis (1971–1993), *Decision Support Systems* 16 (4): 275–96. (This article reports influential authors, frequently cited DSS articles, and major DSS research sub-specialities and DSS reference disciplines.)

*Eom, S.B. (1997) 'Assessing the current state of intellectual relationships between the decision support system area and academic disciplines', in K. Kumar and J.I. DeGross (eds) *Proceedings of the Eighteenth International Conference on Information Systems*, Atlanta, GA: 167–82. (This article identifies the roots of decision support systems research and assesses contributions to DSS areas from reference disciplines.)

*Eom, S.B. (2000) ' Decision support systems implementation research: review of the current state and future directions', in: Proceedings of the Ninth International Conference on Information Systems Development, Christiansand, Norway (forthcoming) (The current state of DSS implementation research is reviewed based on author cocitation analysis of a bibliometrical data compiled over the period of 1991-1997 and suggest future implementation research directions.)

Gorry, G.A. and Scott-Morton, M.S. (1971) 'A framework for management information systems', *Sloan Management Review* 12 (3): 55–70. (This article emphasizes that information systems should exist only to support decisions of varying structuredness.)

*Jessup, L.M. and Valacich, J.S. (eds.) *Group Support Systems: New Perspectives*, New York, NY: Macmillan. (A collection of 17 group support system articles including introduction and overview, research issues, issues in the design, development, and management, and future directions).

* Keen, P.G.W. and Scott-Morton, M.S. (1978) *Decision Support Systems: An Organizational Perspective*, Reading, MA: Addison-Wesley. (The most frequently cited classic textbook presenting the broad conceptual overview of managerial decision making and the role of DSS.)

* King, D. (1993) 'Intelligent support systems: art, augmentation, and agents', in R.H. Sprague, Jr and H.J. Watson (eds), *Decision Support Systems: Putting Theory into Practice*, 3rd edn, Englewood Cliffs, NJ: Prentice Hall. (An account of integrated intelligent system development trends and the future of DSS and expert systems.)

* Kumar, K. and van Hillegersberg, J. (2000) 'ERP: Experiences and Evolution', *Communications of the ACM* 43 (4): 23–26. (This article discusses the evolution of ERP system in terms of technology and functionality during the 1980s and 1990s and concludes that the databases created by ERP will provide the platform for decision support.)
* Rathwell, M.A. and Burns, A. (1985) 'Information systems support for group planning and decision making activities', *MIS Quarterly* 9 (3): 254–71. (This study presents the concept and case studies of distributed decision-making systems.)
* Simon, H.A. (1960) *The New Science of Management Decision*, New York: Harper & Row. (This book provides foundational concepts of decision making and lays the groundwork for further development of DSS.)
Sprague, R.H., Jr (1980) 'A framework for the development of decision support systems', *MIS Quarterly* 4 (4): 1–26. (The most frequently cited article in the DSS literature.)
* Sprague, R.H., Jr and Carlson, E.D. (1982) *Building Effective Decision Support Systems*, Englewood Cliffs, NJ: Prentice Hall. (A classic DSS textbook with data-dialogue-model framework.)
* Zeleny, M. (1982) *Multiple Criteria Decision Making*, New York: McGraw-Hill. (One of the two most frequently cited multi-criteria decision making books.)
* Zeleny, M. (1987) 'Management support systems: towards integrated knowledge management', *Human Systems Management* 7 (1): 59–70. (This article focuses on the future direction of MSS development towards re-integration of expert knowledge.)
* Zeleny M. (1989) 'In this issue', *Human Systems Management*, 8 (2): 99–100. (An editorial introduction of a survey paper on multiple criteria decision making model-based DSS.)

See also: ARTIFICIAL INTELLIGENCE; BUSINESS INFORMATION; DECISION MAKING AND IT/S; DECISION MAKING, MULTIPLE-CRITERIA; ENTERPRISE RESOURCES PLANNING (ERP), EXECUTIVE INFORMATION SYSTEMS; INFORMATION TECHNOLOGY; KNOWLEDGE-BASED SYSTEMS; MANAGEMENT INFORMATION SYSTEMS (MIS); NEURAL NETWORKS; ORGANIZATIONAL LEARNING

Sean B. Eom, "Decision Support Systems," ***International Encyclopedia of Business and Management***, 2nd Edition, Edited by Malcolm Warner, International Thomson Business Publishing Co., London, London, England, 2001, forthcoming.

Chapter 1

Supporting Business Decision-Making

Good information is essential for fact-based decision-making.

Introduction

Beginning in the late 1970s, many vendors, practitioners and academics promoted computer-based Decision Support Systems (DSS). They created high expectations for DSS and much optimism about the prospects for improving decision-making. Despite the hyperbole, the success rate of DSS applications has been less than anticipated. Although the computing industry has transformed how business transactions and data are processed, managers have often been disappointed by attempts to use computers and information technology to support decision-making (cf., Drucker, 1998). Recently, because of technological developments, managers have become more enthusiastic about implementing innovative decision support projects. This is a positive development, but both managers and MIS practitioners need to discuss and review their expectations about Decision Support Systems before they begin important projects.

Anecdotes and research demonstrate that computer-based Decision Support Systems can provide managers with analytical capabilities and information that improves decision-making. In pursuing this goal, many different types of computerized Decision Support Systems have been built to help decision teams and individual decision makers. Some systems provide structured information directly to managers. Other systems help managers and staff specialists analyze situations using various types of models. Some DSS store knowledge and make it available to managers. Some systems support decision-making by small and large groups. Companies even develop DSS to support the decision-making of their customers and suppliers.

This book and chapter discuss how computers and information technology can support and improve business and managerial decision-making. The chapter begins with a short history of Decision Support and Management Information Systems; then the focus turns to examining the Decision Support Systems concept. Based on that examination, a revised framework for categorizing DSS is discussed. Finally, the new framework is linked to the traditional components of a Decision Support System.

A Brief History of Decision Support Systems

Prior to the mid-1960s, it was not cost effective to build large-scale information systems. The first Management Information Systems (MIS) were developed at about that time in large companies. MIS focused on providing managers with structured, periodic reports. Much of the information was from accounting and transaction systems.

In the late 1960s, a new type of information system became practical – model-oriented DSS or management decision systems. Two DSS pioneers, Peter Keen and Charles Stabell (1978), claim the concept of decision support evolved from “the theoretical studies of organizational decisionmaking done at the Carnegie Institute of Technology during the late 1950s and early '60s and the technical work on interactive computer systems, mainly carried out at the Massachusetts Institute of Technology in the 1960s.” Table 1.1 summarizes major developments in the evolution of Decision Support Systems concepts.

Evolution of DSS Concepts

1960s	1970s	1980s	1990s
MIS and structured reports	BrandAid	Key books	Data Warehouses
interactive systems research	MDS	GDSS	OLAP
theory development	RDBMS	EIS	Data mining
		Expert Systems	

Table 1.1. Evolution of DSS Concepts.

In 1971, Michael S. Scott Morton’s book **Management Decision Systems: Computer-Based Support for Decision Making** was published. In 1968-69 Scott Morton studied how computers and analytical models could help managers make a key decision. He conducted an experiment in which managers actually used a Management Decision System (MDS). Marketing and production managers used an MDS to coordinate production planning for laundry equipment. Scott Morton's research was a pioneering implementation, definition and research test of a model-based decision support system.

T.P. Gerrity, Jr. focused on Decision Support Systems design issues in his 1971 **Sloan Management Review** article titled "The Design of Man-Machine Decision Systems: An Application to Portfolio Management". His system was designed to support investment managers in their daily administration of a clients' stock portfolio. DSS for portfolio management have become very sophisticated since Gerrity began his research.

In 1974, Gordon Davis, a Professor at the University of Minnesota, published his influential text **Management Information Systems: Conceptual Foundations, Structure, and Development**. He asserted the MIS concept was “a substantial extension of the concepts of managerial accounting taking into consideration the ideas and techniques of management science and the behavioral theories of management and decision making (p. 8)."

Davis defined a Management Information System as "an integrated, man/machine system for providing information to support the operations, management, and decision-making functions in an organization. The systems utilize computer hardware and software, manual procedures, management and decision models, and a database (p. 5)."

Davis's Chapter 12 titled "Information System Support for Decision Making", and Chapter 13 titled "Information System Support for Planning and Control" created the setting for the development of a broad foundation for Decision Support Systems research and practice. MIS was in many ways beginning to converge with DSS concepts.

By 1975, J. D. C. Little was expanding the frontiers of computer-supported modeling. Little's DSS called Brandaid was designed to support product, promotion, pricing and advertising decisions. Little, in his **Management Science** article titled "Models and Managers: The Concept of a Decision Calculus" identified criteria for designing models to support management decision--making. His criteria included: robustness, ease of control, simplicity, and completeness of relevant detail.

Peter G. W. Keen and Michael Scott Morton's DSS textbook titled **Decision Support Systems: An Organizational Perspective** was published in 1978. Their text provided a comprehensive behavioral orientation to DSS analysis, design, implementation, evaluation and development.

In 1980, Steven Alter published his doctoral dissertation results in a book titled **Decision Support Systems: Current Practice and Continuing Challenge**. Alter's research expanded the framework for our thinking about management DSS. His case studies provided a firm descriptive foundation for identifying Decision Support Systems.

Bonczek, Holsapple, and Whinston’s (1981) book, **Foundations of Decision Support Systems**, created a theoretical framework for understanding the issues associated with designing Decision Support Systems. They identified four essential "aspects" or components common to all DSS: 1. A language system (LS) - all messages the DSS can accept; 2. A presentation system (PS) - all messages the DSS can emit; 3. A knowledge system (KS) -- all knowledge the DSS has stored and retained; and 4. A problem-processing system (PPS) -- the "software engine" that tries to recognize and solve problems during use of the DSS.

The book **Building Effective Decision Support Systems** by Ralph Sprague and Eric Carlson (1982) was an important milestone. It provided a practical, understandable overview of how organizations could and should build DSS. Although the book created some unrealistic expectations, the problem was more the limits of the existing technologies for building DSS than the limits of the concepts Sprague and Carlson presented.

In the mid-1980s, academic researchers developed software to support group and organizational decision-making (cf., DeSanctis and Gallupe, 1987). For the next 10 years, many research studies examined the impacts and consequences of GDSS.

Executive Information Systems (EIS) evolved from the single user Model-Driven Decision Support systems and improved relational database products. The first EIS used pre-defined information screens and were maintained by analysts for senior executives. Beginning in about 1990, data warehousing and On-Line Analytical Processing (OLAP) began broadening the realm of EIS and defined a broader category of Data-Driven DSS (cf., Dhar and Stein, 1997).

A detailed history on the origins of OLAP products by Nigel Pendse (1999) is available on the Web at URL http://www.olapreport.com/origins.htm. Pendse traces OLAP to APL, Express and Comshare's System W. He claims the first explicit Executive Information System product was Pilot Software's Command Center.

Today, a number of academic disciplines provide the substantive foundations for Decision Support Systems development and research. Database researchers have contributed tools and research on managing data and documents. Management Science and Operations Research have developed mathematical models for use in Model-Driven DSS and provided evidence on the advantages of modeling in problem solving. Cognitive Science, especially Behavioral Decision-Making research, has provided descriptive and empirical information that has assisted in DSS design and has generated hypotheses for DSS research. Some other important fields related to DSS include artificial intelligence, human-computer interaction, software engineering, and telecommunications.

A Conceptual Perspective

In the late 1970s, a number of companies developed interactive information systems that used data and models to help managers analyze semi-structured problems. These systems were called Decision Support Systems. DSS can be designed to support decision-makers at any level in an organization. They can support operations, financial management and strategic decision-making. Many of the more interesting DSS are targeted for middle and senior managers. DSS are also often designed for specific types of organizations like hospitals, banks or insurance companies. These systems are sometimes referred to as vertical market or industry-specific DSS.

DSS are both off-the-shelf and custom designed systems. DSS may support a small group of managers using a single personal computer or a large group of managers in a networked client-server environment. These latter systems are often called Enterprise-Wide DSS.

Characteristics of DSS

Although the term Decision Support System has many connotations, based on Steven Alter's (1980) pioneering research we can identify the following three major characteristics:

1. DSS are designed specifically to facilitate decision processes,

2. DSS should support rather than automate decision making, and

3. DSS should be able to respond quickly to the changing needs of decision makers.

Clyde Holsapple and Andrew Whinston, in their book **Decision Support Systems: A Knowledge-Based Approach** (1996), identified five characteristics one should expect to observe in a DSS (see pages 144-145). Their list is very general and somewhat abstract, but it provides an even broader perspective on the DSS concept. The Holsapple and Whinston characteristics are:

1. A DSS includes a body of knowledge that describes some aspects of the decision-maker's world, that specifies how to accomplish various tasks, that indicates what conclusions are valid in various circumstances, and so forth.

2. A DSS has an ability to acquire and maintain descriptive knowledge (i.e., record keeping) and other kinds of knowledge as well (i.e., procedure keeping, rule keeping, etc.).

3. A DSS has an ability to present knowledge on an ad hoc basis in various customized ways as well as in standardized reports.

4. A DSS has an ability to select any desired subset of stored knowledge for either presentation or deriving new knowledge in the course of problem recognition and/or problem solving.

5. A DSS can interact directly with a decision maker or a participant in a decision in such a way that the user has a flexible choice and sequence of knowledge-management activities.

Sprague and Carlson (1982) and others define Decision Support Systems broadly as interactive computer based systems that help decision-makers use data and models to solve ill-structured, unstructured or semi-structured problems. Bonczek, Holsapple and Whinston (1981) argued the "system must possess an interactive query facility, with a query language that ... is ... easy to learn and use (p. 19)". Various types of DSS help decision-makers use and manipulate very large databases; some help managers apply checklists and rules; others make extensive use of mathematical models.

Case studies from the past 25 years have demonstrated that it is possible to support management activities in many ways. Some DSS help managers by expediting access to information that would otherwise be unavailable or difficult to obtain; others contain explicit models that provide structure for particular decisions. Some systems are primarily tools for individuals working more or less alone on decision tasks; others serve primarily to support communication among people whose work must be coordinated.

Many terms are used for specific types of DSS including business intelligence, collaborative systems, data mining, data warehousing, knowledge management and on-line analytical processing. Software vendors use these more specialized terms for both descriptive and marketing purposes. What term we use for a system or software package is a secondary concern. Our primary concern is finding software and systems that meet a manager's decision support needs and provide appropriate management information.

Management Information

Managers and their support staffs need to consider what information and analyses are actually needed to support management and business activities. Some managers need both detailed transaction data and summarized data. Most managers only want summaries of transactions. Managers usually want lots of charts and graphs; a few only want tables of numbers. Many managers want information provided routinely or periodically and some want information available on-line and on demand. Managers want financial analyses and some managers want primarily "soft", non-financial or qualitative information.

In general, an Information System can provide business transaction information and it can help managers understand many business operations and performance issues. For example, a computerized system can help managers understand the status of operations, monitor business results, review customer preference data and investigate competitor actions. In all of these situations, management information and analyses should have a number of characteristics. Information must be both timely and current. These characteristics mean the information is up-to-date and available when managers want it. Also, information must be accurate, relevant and complete. Finally, managers want information presented in a format that assists them in making decisions. In general, management information should be summarized and concise and any support system should have an option for managers to obtain more detailed information.

Decision Support Systems need to provide current, timely information that is accurate, relevant and complete. A specific DSS must present information in an appropriate format that is easy to understand and manipulate. The information presented by a DSS may result from analysis of transaction data or it may be the result of a decision model or it may have been gathered from external sources. DSS can present internal and external facts, informed opinions and forecasts to managers. **Managers want the right information, at the right time, in the right format, and at the right cost.**

DSS versus MIS

How does a Decision Support System differ from a Management Information System? Let's begin drawing distinctions between these two terms by first examining the concepts Management Information System (MIS) and Information System (IS). Many authors have used the term MIS to describe a broad, general category of information systems. Also, MIS and IS are used interchangeably to describe a functional department in companies and organizations responsible for managing information systems and technology. A number of computing jobs are grouped together under the heading of MIS or IS professionals. Finally, the term Management Information Systems or MIS is used to identify an academic major and an area of scholarly inquiry in universities.

In the 1970s, an MIS generated periodic management reports. Today, managers use Data-Driven Decision Support Systems to meet their management reporting needs. When the term Management Information System is defined narrowly it refers to a management reporting system that provides periodic, structured paper-based reports. In contrast, Data-

Driven DSS are intended to be interactive, real-time systems that are responsive to unplanned as well as planned information requests and reporting needs. Model-Driven DSS are usually focused on modeling a specific decision or a set of related decisions (cf., Power, 1997).

Decision Support Systems should be defined as a broad category of analytical management information systems. DSS provide managers more control of their data, access to analytical tools, and capabilities for consulting and interacting with a distributed group of staff. An enterprise-wide Decision Support System is linked to a large data warehouse and serves many managers within one company. Also, a DSS is defined as an interactive system in a networked environment that helps a targeted group of managers make decisions. The primary focus in the following discussion is on various types of Decision Support Systems. The term MIS will be used sparingly and usually it will broadly refer to any information system that assists or supports managers in their various tasks.

Decision Support versus Transaction Processing Systems

Development of Decision Support Systems is one of the rapidly changing frontiers in the application of computers in organizations. One reason we study DSS is to understand how they differ from other systems. We have successfully implemented computer-based Transaction Processing Systems (TPS), but knowledge of building these operational systems is not adequate to create effective Decision Support Systems. So if DSS are to be successfully designed, developed and implemented, then both managers and many MIS professionals need a more sophisticated technical and philosophical understanding of Decision Support Systems.

Technology is creating new decision support capabilities, but much learning and discussion needs to occur to successfully exploit the technological possibilities. Decision Support Systems differ in many ways from operational Transaction Processing or Online Transaction Processing Systems. For example, a popular system that has been widely implemented is called Enterprise Resource Planning (ERP). ERP is **NOT** a Decision Support System even though the term suggests that decision-making and planning will be improved. In general, Enterprise Resource Planning is an integrated Transaction Processing System that facilitates the flow of information between all of the functional areas of a business. Recently, DSS have been built to help managers analyze the data from ERP systems.

This section discusses how Decision Support and Transaction Processing Systems differ. Let's begin by briefly reviewing the concept of a system.

What is a system?

The term system is used in many technology-related concepts including Decision Support System and Transaction Processing System -- both are computing or information systems. Managers and MIS specialists use the concept of a system frequently and yet it is hard for most of us to define and understand the concept.

Let's begin exploring this key term by defining a system as an interrelated set of components including people, activities, technology and procedures that are designed or intended to achieve a predefined purpose. A system receives input from its environment and the various subsystems or components of the system interact to produce outputs. Systems are defined in terms of their components. System components are surrounded by an imaginary boundary that separates a specific system from its environment. A system designer identifies both inputs from the environment as well as the outputs from the system. Systems also have feedback mechanisms to provide a means of controlling the operation of the system. Feedback is an output from a system that later reenters the system as an input.

Let's examine a simple conceptual specification of a system. The initial input into the system is a bank customer requesting a loan. The customer makes a request to a bank officer. The bank officer collects information from the customer and enters that information into a computerized form. A loan approval model is built into a computerized decision aid. Some people identify the computerized model as the actual decision support system. The banker uses the result from the computerized loan approval model to finalize the decision to approve or deny the loan. In some cases the loan information will need to be shared with a loan committee possibly using a group support system. The actual decision is then communicated to the customer either face-to-face or by a formal letter that may be generated by a computerized decision aid. Feedback comes from the customer.

This decision process and the overall conceptual system may include various Decision Support Systems. The bank's Transaction Processing System would be updated when the loan was made and the funds distributed. The loan is the primary transaction. Making the loan is the decision process. DSS can support making loans or a DSS can help analyze lending activity at the bank or predict lending activity and interest rates.

In a Decision Support System, the primary focus is often on the computerized components of the system. This is a narrow perspective for defining the components of a system; it is often helpful to define the DSS boundary to include a broader decision process that may involve people performing non-computerized tasks as well as more routine data gathering tasks. The users of the computerized tools are also part of the broader system. Finally, note that the actual communication or transmission of decisions may not occur using computerized systems. This step in a decision process needs to be considered in the design of the DSS and it should be included within the boundary of the system.

We need to define Decision Support Systems on both a conceptual level and a concrete, technical level. Both managers and DSS designers need to understand what they are trying to accomplish. The specific purpose of a proposed Decision Support System and its components need to be defined early in the design and development process.

Major differences

A major difference between Transaction Processing Systems and DSS is the general purpose of each type of system. Transaction Processing Systems are designed to expedite and automate transaction processing, record keeping, and simple business reporting of transactions. Decision Support Systems are intended to assist in decision-making and decision implementation. Transaction processing is however related to the design of DSS because transaction databases often provide data for decision-oriented reporting systems and data warehouses.

Transaction Processing Systems usually provide standard reports on a periodic basis and support the operations of a company. DSS are used on demand when they are needed to support decision-making. A manager typically initiates each instance of Decision Support System use, either by using the DSS herself or by asking a staff intermediary to use a DSS. Some managers and especially clerical employees use Transaction Processing Systems to support operations. DSS are designed for use by line managers and support staff. TPS record current information and maintain a database of transaction information. Some DSS use historical internal and external data for analysis. Other DSS focus on modeling current and future scenarios and incorporate historical data, forecasts, and assumptions. TPS emphasize data integrity and consistency; and although these qualities in a system are important, a DSS places it primary emphasis on flexibility and supporting ad hoc queries and analyses.

One can draw many distinctions between Transaction Processing Systems and DSS, but analysts and managers need to stay focused on the phrase "decision support" in the term Decision Support System. Decision Support Systems are intended to improve and speed-up the processes by which people make and communicate decisions. Thus the emphasis in building a DSS is on increasing individual and organizational decision making effectiveness rather than on increasing efficiency in processing operating data.

Examples of DSS Applications

Hundreds of DSS applications are described in professional journals like **Interfaces** and in Information Systems trade publications like **Information Week** (http://www.informationweek.com). Many DSS case studies are also available on the World-Wide Web. This section discusses various Decision Support Systems examples and a number of taxonomies of DSS.

One of the long-standing conclusions from reading DSS case studies is that what managers, vendors and consultants call DSS can "take on many different forms and can be used in many different ways (Alter, 1980, p. 71)." DSS vary in many ways. They differ in terms of who uses a specific system, that is some DSS are used by actual decision makers and some are used by intermediaries like marketing analysts or financial analysts. Some DSS focus on data, some on models and some on communications. DSS also differ in scope, some DSS are intended for one "primary" user and used “stand-alone” for analysis and others are intended for many users in an organization.

A few examples show the wide variety of DSS applications. Major airlines have DSS used by analysts for many tasks including pricing and route selection. Many companies have DSS that aid in corporate planning and forecasting. Specialists often use these DSS that focus on financial and simulation models. Investment evaluation and support systems are increasingly common. Frito-Lay has a DSS that aids in pricing, advertising, and promotion. Route salesmen use hand-held computers to support decision making activities. Many manufacturing companies use Manufacturing Resources Planning (MRP) software. This specific operational level DSS supports master production scheduling, purchasing, and materials requirements planning. More recent MRP systems support "what-if" analysis and simulation capabilities. Monsanto, FedEx and most transportation companies use DSS for scheduling trucks, airplanes and ships. The Coast Guard uses a DSS for procurement decisions. Companies like Wal-Mart have large data warehouses and use data mining software. Business Intelligence and Knowledge Management Systems are increasingly common. On the World-Wide Web one can find DSS that help track and manage stock portfolios, choose stocks, plan trips, and suggest gifts. DSS support distributed decision activities using groupware and a corporate intranet.

Alter’s Taxonomy

In 1980, Steven Alter (pps. 73-93) proposed a taxonomy of DSS. The next few paragraphs summarize his taxonomy and discuss some of the key issues for each type of DSS. Alter's taxonomy is based on the degree to which DSS output can directly determine the decision. The taxonomy is related to a spectrum of generic operations that can be performed by Decision Support Systems. These generic operations extend along a single dimension, ranging from extremely data-oriented to extremely model-oriented. DSS may involve retrieving a single item of information, providing a mechanism for ad hoc data analysis, providing pre-specified aggregations of data in the form of reports or "screens". DSS may also include estimating the consequences of proposed decisions and proposing decisions.

Alter's idea was that a Decision Support System could be categorized in terms of the generic operations it performs, independent of type of problem, functional area or decision perspective. Alter conducted a field study of 56 DSS that he categorized into seven distinct types of DSS. His seven types include:

- **File drawer systems** that provide access to data items. Examples include real-time equipment monitoring, inventory reorder and monitoring systems. Simple query and reporting tools that access OLTP fall into this category.

- **Data analysis systems** that support the manipulation of data by computerized tools tailored to a specific task and setting or by more general tools and operators. Examples include budget analysis and variance monitoring, and analysis of investment opportunities. Most data warehouse applications would be categorized as data analysis systems.

- **Analysis information systems** that provide access to a series of decision-oriented databases and small models. Examples include sales forecasting based on a marketing database, competitor analyses, product planning and analysis. OLAP systems fall into this category.

- **Accounting and financial models** that calculate the consequences of possible actions. Examples include estimating profitability of a new product; analysis of operational plans using a goal-seeking capability, break-even analysis, and generating estimates of income statements and balance sheets. These types of models should be used with “What if?” or sensitivity analysis.

- **Representational models** that estimate the consequences of actions on the basis of simulation models that include relationships that are causal as well as accounting definitions. Examples include a market response model, risk analysis models, and equipment and production simulations.

- **Optimization models** that provide guidelines for action by generating an optimal solution consistent with a series of constraints. Examples include scheduling systems, resource allocation, and material usage optimization.

- **Suggestion models** that perform the logical processing leading to a specific suggested decision for a fairly structured or well-understood task. Examples include insurance renewal rate calculation, an optimal bond-bidding model, a log cutting DSS, and credit scoring.

An understandable typology like Steven Alter's helps reduce the confusion for managers who are investigating and discussing Decision Support Systems. The taxonomy also helps users and developers communicate their experiences with DSS.

Other Taxonomies or Frameworks

Holsapple and Whinston (1996) identify 5 specialized types of DSS (see pp. 178-195). First they identify an evolving group of systems they call Text-Oriented DSS. This type of DSS supports a decision-maker by electronically keeping track of textually represented knowledge that could impact decisions. This type of system supports document creation, revision, viewing, searching and hypertext links. Holsapple and Whinston also discuss Database-Oriented DSS, Spreadsheet-Oriented DSS, Solver-Oriented DSS, and Rule-Oriented DSS. A solver is a general algorithm that can be customized to solve a specific instance of a more general class of problems. These last four types of DSS match up well with Alter’s categories.

Donovan and Madnick (1977) classified DSS as institutional or ad hoc DSS. Institutional DSS support decisions that are recurring. An ad hoc DSS supports problems that are not anticipated and that are not expected to reoccur. Hackathorn and Keen (1981) identified DSS in three distinct yet interrelated categories: Personal DSS, Group DSS and Organizational DSS. Many DSS are designed for a particular problem in a particular company, but some DSS are generic or ready-made DSS (cf., Turban and Aronson, 1998). Golden, Hevner and Power (1986) identified decision insight systems as a particular category of Model-Oriented DSS that uses decision analysis tools to help decision-makers structure decision situations and gain insight about possible solutions.

An Expanded DSS Framework

The terms frameworks, taxonomies, conceptual models and typologies are often used interchangeably. Taxonomies classify objects and typologies show how mutually exclusive types of things are related. The general desire is to create a set of labels that help people organize and categorize information. In this section we want to categorize the large number of computerized systems that support decision-making. Sprague and Watson (1996) argue typologies, frameworks or conceptual models are “often crucial to the understanding of a new or complex subject." A good framework shows the parts of a topic and how the parts interrelate.

A new, broader typology or framework than Alter's is needed today because DSS are much more common and more diverse than when he conducted his research and proposed his framework. Alter's typology is still relevant for categorizing some types of DSS, but not for all DSS. To keep the number of categories in a new framework manageable, one can and should simplify Alter's typology into three types of Decision Support Systems: Data-Driven, Model-Driven and Knowledge-Driven DSS. We can also categorize DSS in terms of internal and external users, specificity or function and technology. The following expanded DSS framework is probably not comprehensive and parsimonious, but it helps categorize the most common DSS currently in use. Some DSS are hybrid systems driven by more than one major DSS component. The framework focuses on one major dimension with 5 categories and 3 secondary dimensions.

Data-Driven DSS

Let's call the first category of Decision Support Systems **Data-Driven DSS.** These systems include file drawer and management reporting systems, data warehousing and analysis systems, Executive Information Systems (EIS) and Geographic Information Systems (GIS). Business Intelligence Systems are also examples of Data-Driven DSS. Data-Driven DSS emphasize access to and manipulation of large databases of structured data and especially a time-series of internal company data and some times external data. Simple file systems accessed by query and retrieval tools provide the most elementary level of functionality. Data warehouse systems that allow the manipulation of data by computerized tools tailored to a specific task and setting or by more general tools and operators provide additional functionality. Data-Driven DSS with Online Analytical Processing (OLAP) provide the highest level of functionality and decision support that is linked to analysis of large collections of historical data (cf., Dhar and Stein, 1997). Professor Paul Gray argues that in approximately 1993, "the data warehouse and the EIS people found one another, with the data warehouses obtaining their needed application and the EIS people receiving a new breath of life from expanding beyond the pretty screen."

Model-Driven DSS

A second category, **Model-Driven DSS,** includes systems that use accounting and financial models, representational models, and optimization models. Model-Driven DSS emphasize access to and manipulation of a model. Simple statistical and analytical tools provide the most elementary level of functionality. Some OLAP systems that allow complex analysis of data may be classified as hybrid DSS systems providing modeling, data retrieval and data summarization functionality. Model-Driven DSS use data and parameters provided by decision-makers to aid them in analyzing a situation, but they are not usually data intensive. Very large databases are usually not needed for Model-Driven DSS.

Knowledge-Driven DSS

The terminology for the this category of DSS is still evolving. Currently, the best term seems to be **Knowledge-Driven DSS.** Sometimes it seems equally appropriate to use Alter's term Suggestion DSS or the narrower term Management Expert System.

Knowledge-Driven DSS can suggest or recommend actions to managers. These DSS are person-computer systems with specialized problem-solving expertise. The "expertise" consists of knowledge about a particular domain, understanding of problems within that domain, and "skill" at solving some of these problems. A related concept is Data Mining. It refers to a class of analytical applications that search for hidden patterns in a database. Data mining is the process of sifting through large amounts of data to produce data content relationships. Tools used for building these systems are also called Intelligent Decision Support methods (cf., Dhar and Stein, 1997). Data Mining tools can be used to create hybrid Data-Driven and Knowledge-Driven DSS.

Document-Driven DSS

A new type of DSS, a **Document-Driven DSS** or Knowledge Management System, is evolving to help managers retrieve and manage unstructured documents and Web pages. A Document-Driven DSS integrates a variety of storage and processing technologies to provide complete document retrieval and analysis. The Web provides access to large document databases including databases of hypertext documents, images, sounds and video. Examples of documents that would be accessed by a Document-Based DSS are policies and procedures, product specifications, catalogs, and corporate historical documents, including minutes of meetings, corporate records, and important correspondence. A search engine is a powerful decision-aiding tool associated with a Document-Driven DSS (cf., Fedorowicz, 1993, pp. 125-136).

Communications-Driven and Group DSS

Group Decision Support Systems (GDSS) came first**,** but now a broader category of **Communications-Driven DSS** or groupware can be identified. This type of DSS includes communication, collaboration and decision support technologies that do not fit within those DSS types identified by Steven Alter. Therefore, Communications-Driven DSS need to be identified as a specific category of DSS. We will call these systems Communications-Driven DSS even though many people are more familiar with the term GDSS. A GDSS is a hybrid DSS that emphasizes both the use of communications and decision models. A Group Decision Support System is an interactive computer-based system intended to facilitate the solution of problems by decision-makers working together as a group. Groupware supports electronic communication, scheduling, document sharing, and other group productivity and decision support enhancing activities. We have a number of technologies and capabilities in this category in the framework -- GDSS Decision Rooms, two-way interactive video, White Boards, Bulletin Boards, and Email.

Inter-Organizational or Intra-Organizational DSS

A relatively new category of DSS made possible by new technologies and the rapid growth of the public Internet is **Inter-Organizational DSS.** These DSS serve a company's customers or suppliers. The public Internet is creating communication links for many types of inter-organizational systems, including DSS. An Inter-Organizational DSS provides stakeholders with access to a company’s intranet and authority or privileges to use specific DSS capabilities. Companies can make a Data-Driven DSS available to suppliers or a Model-Driven DSS available to customers to design a product or choose a product. Most DSS are **Intra-Organizational DSS** that are designed for use by individuals in a company as "stand-alone DSS" or for use by a group of managers in a company as a Group or Enterprise-Wide DSS. The prefix “intra” means the DSS is used within a specific organization and “inter” means the DSS is used more widely.

Function-Specific or General Purpose DSS

Many DSS are designed to support specific business functions or types of businesses and industries. We can call such DSS function-specific or industry-specific DSS. A **Function-Specific DSS** like a budgeting system may be purchased from a vendor or customized in-house using a more general-purpose development package. Vendor developed or “off-the-shelf” DSS support functional areas of a business like marketing or finance; some DSS products are designed to support decision tasks in a specific industry like a crew scheduling DSS for an airline. A task-specific DSS has an important purpose in solving a routine or recurring decision task. Function or task-specific DSS can be further classified and understood in terms of the dominant DSS component, that is as a Model-Driven, Data-Driven or Suggestion DSS. A function or task-specific DSS holds and derives knowledge relevant for a decision about some function that an organization performs (e.g., a marketing function or a production function). This type of DSS is categorized by purpose; Function-Specific DSS help a person or group accomplish a specific decision task. General-purpose DSS software helps support broad tasks like project management, decision analysis, or business planning.

Web-Based DSS

Finally, all of the above types of DSS can be implemented using Web technologies and we can call these systems Web-Based DSS. A **Web-Based DSS** is a computerized system that delivers decision support information or decision support tools to a manager or business analyst using a "thin-client" Web browser like Netscape Navigator or Internet Explorer. The computer server that is hosting the DSS application is linked to the user's computer by a network with the TCP/IP protocol. In many companies, a Web-Based DSS is synonymous with an intranet or Enterprise-Wide DSS. A company intranet is supporting a large group of managers using Web browsers in a networked environment. Managers often have Web access to a data warehouse as part of a DSS architecture. Today Web technologies are the primary tools used to create Inter-Organizational DSS that support the decision-making of customers and suppliers.

Web or Internet technologies are the leading edge for building DSS, but some Intra-Organizational DSS will continue to be built using traditional programming languages or fourth generation languages or application development tools using “thick-client” or mainframe enabling technologies.

Column one of Table 1.2 list five broad categories of Decision Support Systems that differ in terms of the DSS technology component, including Communications-Driven DSS, Data-Driven DSS, Document-Driven DSS, Knowledge-Driven DSS and Model-Driven DSS. Subsequent chapters explain these categories in more detail and identify development and implementation issues. The new DSS framework also categorizes Decision Support Systems by user groups – intra-organizational and inter-organizational. The new category called Inter-Organizational DSS helps us focus on the broadening of the DSS user group to include external stakeholders.

Dominant DSS Component	User Groups: Internal, External	Purpose: General, Specific	Enabling Technology
Communications **Communications-Driven DSS**	Internal teams, now expanding	Conduct a meeting Bulletin Board Help users collaborate	Web or Client/Server
Database **Data-Driven DSS**	Managers, staff, now suppliers	Query a Data Warehouse	Main Frame, Client/Server, Web
Document base **Document-Driven DSS**	Specialists and user group is expanding	Search Web pages Find documents	Web
Knowledge base **Knowledge-Driven DSS**	Internal users, now customers	Management Advice Choose products	Client/Server, Web
Models **Model-Driven DSS**	Managers and staff, now customers	Crew Scheduling Decision Analysis	Stand-alone PC

Table 1.2. A New DSS Framework.

From a different perspective, Decision Support Systems can be categorized by the purpose of the DSS. Many DSS have a narrow, focused, specific purpose rather than a general purpose. Finally, DSS can be categorized by the basic enabling technology. The Web is an important new development arena for DSS so it is crucial to examine and understand Web-Based DSS. We can use dominant DSS component, user group, purpose and enabling technology to categorize a specific system. For example, we may want to build a Model-Driven, Inter-Organizational, Product Design, Web-Based DSS.

Building Decision Support Systems

Traditionally, academics and practitioners have discussed building Decision Support Systems in terms of four major components – 1) the user interface, 2) the database, 3) the models and analytical tools, and 4) the DSS architecture and network (cf., Sprague and Carlson, 1982). This traditional list of components remains useful because it identifies similarities and differences between categories or types of DSS and it can help managers and analysts build new DSS. The DSS framework is based on the different emphases placed on DSS components when systems are actually constructed (see Figure 1.1).

Data-Driven, Document-Driven and Knowledge-Driven DSS need specialized database components. A Model-Driven DSS may use a simple flat-file database with fewer than 1,000 records, but the model component is very important. Experience and some empirical evidence indicate that design and implementation issues vary for Data-Driven, Document-Driven, Model-Driven and Knowledge-Driven DSS. Multi-participant systems like Group and Inter-Organizational DSS also create complex implementation issues. For instance, when implementing a Data-Driven DSS a designer should be especially concerned about the user's interest in applying the DSS in unanticipated or novel situations.

In creating an accounting or financial DSS simulation model, a developer should attempt to verify that the initial input estimates for the model are thoughtful and reasonable. In developing a representational or optimization model, the analyst should be concerned about possible misunderstandings of what the model means and how it can or cannot be used (cf., Alter, 1980, p. 92). Networking issues create challenges for many types of DSS, but especially for Communications-Driven systems with many participants, so-called multi-participant systems. Today architecture and networking issues are increasingly important in building DSS.

DSS should be built or implemented using an appropriate process. Many small, specialized Model-Driven DSS are built quickly. Large, Enterprise-Wide DSS are built using sophisticated tools and systematic and structured systems analysis and development approaches. Communications-Driven and Group DSS are often purchased as off-the-shelf software. Creating Enterprise-Wide DSS environments remains an iterative and evolutionary task. An Enterprise-Wide DSS grows and inevitably becomes a major part of the overall information systems infrastructure of an organization. Despite the significant differences created by the specific task and scope of a DSS, all DSS have similar technical components and share a common purpose, supporting decision-making.

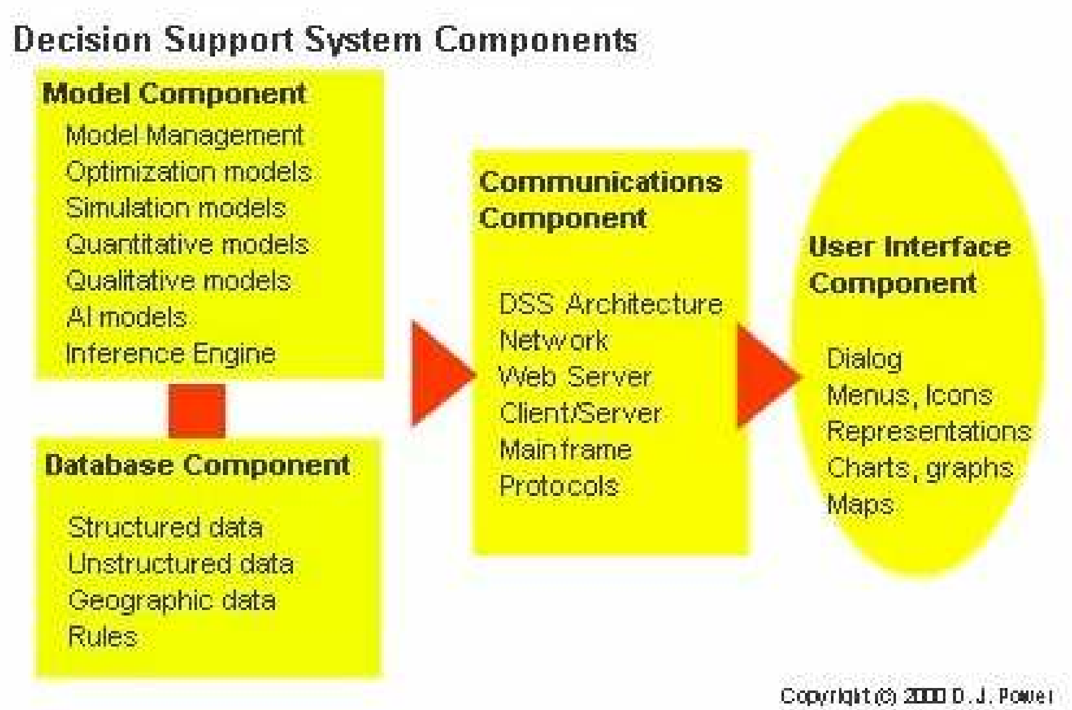

Figure 1.1. Traditional DSS components.

A Data-Driven DSS database is often a collection of current and historical structured data from a number of sources that have been organized for easy access and analysis. We are expanding the data component to include unstructured documents in Document-Driven DSS and "knowledge" in the form of rules in Knowledge-Driven DSS. Large databases of structured data in Enterprise-Wide DSS are often called data warehouses or data marts. DSS usually use data that has been extracted from all relevant internal and external databases. Managing information often means managing a database. Supporting management decision-making means that computerized tools are used to make sense of the structured data or documents in a database.

Mathematical and analytical models are the major component of a Model-Driven DSS. DSS models should be used and manipulated directly by managers and staff specialists. Each Model-Driven DSS has a specific set of purposes and hence different models are needed and used. Choosing appropriate models is a key design issue. Also, the software used for creating specific models needs to manage needed data and the user interface. In Model-Driven DSS the values of key variables or parameters are changed, often repeatedly, to reflect potential changes in supply, production, the economy, sales, the marketplace, costs, and/or other environmental and internal factors. Information from the models is then analyzed and evaluated by the decision-maker. Suggestion DSS use special models for processing rules or identifying relationships in data.

The DSS architecture and networking design component refers to how hardware is organized, how software and data are distributed in the system, and how components of the system are integrated and connected. A major issue today is whether DSS should be available using a Web browser on a company intranet and also available on the Global Internet. Managers and MIS staff both need to develop an understanding of the technical issues and the security issues related to DSS architectures, networks and the Internet. Networking is the key driver of Communications-Driven DSS.

Managers and DSS analysts both need to emphasize the user interface component. In many ways the user interface is the most important component. The tools for building the user interface are sometimes termed DSS generators, query and reporting tools, and front-end development packages. Much of the design and development effort should focus on building the user interface. We need to remember that the screens and displays in the user interface heavily influence how a manager perceives a DSS. What we see is the DSS!!

Conclusions and Commentary

The rapid growth of the World-Wide Web has created enormous opportunities for making more organizational information available to decision-makers. Web architectures permit Information Systems professionals to centralize and control information and yet easily distribute it in a timely manner to managers who need it. Also, the internal Internet or Intranet is providing many opportunities for delivering information from data warehouses, models and other tools to the desktop. The Web DSS permit and encourage further analysis and collaboration. The technologies and software associated with Decision Support Systems continues to change rapidly and development tools are overlapping for some applications. In general, managers and IS staff need to recognize that the overall technological and social context of DSS and business management is changing.

The new managers who are and will be using company Intranets and the Internet are more technologically sophisticated than the managers of the past. They will have high expectations for DSS, but in many ways they will be much better customers of computerized decision support. The DSS design and development environment is changing as rapidly as the software tools and in as positive a direction. The Web technologies will facilitate improved DSS tools at manager's desktops.

General Managers need broad knowledge of the managerial and technical issues associated with the various categories of Decision Support Systems. MIS professionals need this same general knowledge and they need specific skills in analysis, design and development of DSS.

In 1974, Gordon Davis wrote "The application of computer technology and MIS concepts has produced some spectacular successes and also some rather expensive failures." Both successes and failures will still occur. Failures occur in leading edge application areas and for what turn out to be overly ambitious projects. A shortage of MIS professionals is also slowing development in some areas and increasing failures of innovative systems. All of us need to recognize that resistance to change and insufficient user involvement contributes to DSS project failure in some situations. Managers need to resolve political issues associated with building new Decision Support Systems and providing greater access to management information. For example, senior managers need to address questions like: How should data be shared and how much data should be shared? Should all managers be required to use a DSS and support systems like email?

Managers and MIS practitioners need to consider at least five major issues associated with building and using Decision Support Systems. First, we must determine what business and decision processes should be computerized? And in some situations we need to ask what part of the process should be supported? In many companies this issue needs to be re-examined for current Decision Support Systems. Chapters 2 and 3 address this issue. Second, we must ask what data should be captured in processes and how should it be stored and integrated? Continuing to rely on existing decision processes may limit the information that can be provided to decision-makers. Chapter 4 discusses DSS design and development issues.

Third, we need to ask how data should be processed and presented to support decision-making? Chapter 5 emphasizes user interface design issues. Fourth, and perhaps the major issue is whether current Decision Support Systems are creating results that are "decision-impelling"? (based on Davis, 1974, p. 6). Chapters 7 to 11 review the possibilities for building innovative DSS.

Finally, we need to ask what information technology should be used for building DSS? Chapter 6 reviews DSS architecture and networking issues. Managers need some technical familiarity and sophistication to evaluate the wide-ranging set of technologies that are available for DSS applications. Understanding the various categories of Decision Support Systems that can be built begins the task of rationally answering the above questions. Subsequent chapters provide more elaboration and some details.

Decision Support Systems are not a panacea for improving business decisions. Most people acknowledge that managers need “good” information to manage effectively, but a DSS is not always the solution for providing “good” information. A DSS is limited by the data that can be obtained, the cost of obtaining, processing, and storing the information, the cost of retrieval and distribution, the value of the information to the user, and the capability of managers to accept and act on the information. Our capabilities to support decision-making have increased, but we still have very real technical, social, interpersonal and political problems that must be overcome when we build DSS. Chapter 12 addresses these issues and the evaluation of proposed DSS projects.

Audit Questions

1. Does your firm actively manage decision-relevant information?
2. Has your firm implemented any computerized systems to support decision-making?
3. Are you using any Decision Support Systems? If so, from what category?

Questions for Review

1. What is the MIS concept? How is it related to DSS?
2. What are the major characteristics of a DSS?
3. How does a Transaction Processing System differ from a Decision Support System?
4. What are the categories of DSS included in the proposed DSS framework?
5. What components are common to the design and implementation of computerized Decision Support Systems?
6. What were the two main streams of research that led to the evolution and development of Decision Support Systems?

Questions for Further Thought

1. Do managers need the support provided by DSS?
2. Is it realistic to use technology to support decision-making?
3. Do managers want to use decision support tools?
4. What experiences have you had using Decision Support Systems? What was a good experience? What is an example of a bad experience?

Internet Exercises

1. Find an example of a Decision Support System at a Web site. Use the DSS Framework and classify the DSS.
2. Search for the term DSS using 2 Web search engines.
3. Visit the Web sites of Information Week (www.informationweek.com), Internet Week (www.internetwk.com) and CIO (www.cio.com) and search for articles on key terms from this chapter like DSS, and MIS.

Brief Examples of DSS Implementations

Advanced Scout

IBM has prototyped software to help National Basketball Association (NBA) coaches and league officials organize and interpret the data collected at every game. Using software called Advanced Scout to prepare for a game, a coach can quickly review countless stats: shots attempted, shots blocked, assists made, personal fouls. But Advanced Scout can also detect patterns in these statistics that a coach may not have known about. Advanced Scout software provides an easy and meaningful way to process information. "It helps coaches easily mine through and analyze a lot of data and no computer training or data analysis background is required," says Dr. Inderpal Bhandari, computer scientist at IBM's T.J. Watson Research Center. Patterns found through analysis are linked to the video of the game. Coaches can look at just those clips that make up an interesting pattern (check http://www.research.ibm.com/scout/works.html).

BCA DSS (Base Closure and Analysis DSS)

An application called the Base Closure and Analysis DSS provided the U.S. Air Force with a robust methodology and common framework for analyzing the impact of various base closure scenarios. The software used a multi-layer, hierarchical filtering process to evaluate the relative impact of closing each base. Bases that posed minimum strategic, operational, social, and economic impact were placed at the top of the closure recommendation list. At any step, base closing committee members could review DSS-developed impact analyses to assist in determining which bases should be included in the next level of analysis. Using the DSS, the committee members could perform analyses using eight main criteria and 212 sub-criteria on which all bases were evaluated. These criteria, specified by DOD, focused on elements that impact operational effectiveness, including such items as alternate airfield availability, weather data, and facility infrastructure capacity (from URL http://www.strategy.com/success/msi_saf1.htm).

FedEx Business Intelligence System

Federal Express, based in Memphis, Tenn., rolled out Business Intelligence capabilities to a global base of 700 end-users. FedEx created a central, integrated data warehouse hub, which provides Web-based, real-time access to financial and logistical information necessary for planning and decision-making. The solution, from Pinnacle Solutions Inc., was deployed on a group of Dell PowerEdge servers running Windows NT Server 4.0. Data is stored in an Oracle database, and analytical queries are run against a separate server running Hyperion Essbase, an online analytical processing (OLAP) engine. Most access is from browsers over the corporate intranet, along with some standard client/server deployments using Excel spreadsheets.

ShopKo

In 1997, ShopKo developed a "Merchandise Data Warehouse." ShopKo stores carry 200,000 stock units of data. This results in massive amounts of data. Sales statistics on every stock unit in every store is collected daily and stored in a data warehouse. This central data repository is used in analysis, querying, and decision-making. The main strategy in developing the DSS tool was to allow ShopKo associates to query a common business repository for identification and analysis of business opportunities and exceptions. With this strategy, ShopKo stores are able carry the right merchandise at the right time in the right place while remaining current with changing demands due to seasons, trends, etc. Some of the important goals of this project were: improvement of sales analysis, understanding of inventory levels, determining market trends, and improvement of advertisement effectiveness. ShopKo extended its DSS capabilities to its store units by using a Web-based DSS.

Y2K GroupSystems Online

During the week leading up to and immediately following January 1, 2000, approximately 150 people participated in the crisis management activities 24 hours a day, 7 days a week using GroupSystems OnLine. Representatives from the Office of the Secretary of Defense, C3I, JCS Staff, Federal Emergency Management Agency (FEMA), the State Department and Legislative Affairs among others participated in crisis management sessions over a secure Intranet within the Pentagon known as the SIPRNET. Although major crises did not materialize during the course of the two-week period, some non-crisis events did occur that required internal action and decision making on the part of the Pentagon. GroupSystems OnLine was used to communicate information and it was used to provide input, discuss solutions and create reports of recommended action (cf., http://www.groupsystems.com).

Questions for discussion of the case examples:

1. Use the DSS framework to categorize the examples. What type of DSS was implemented?

2. Does each DSS seem useful? Would you use the system?

References

Alter, S.L. *Decision Support Systems: Current Practice and Continuing Challenge*. Reading, MA: Addison-Wesley, 1980.

Bonczek, R. H., C.W. Holsapple, and A.B. Whinston. *Foundations of Decision Support Systems.* New York: Academic Press, 1981.

Davis, G. *Management Information Systems: Conceptual Foundations, Structure, and Development.* New York: McGraw-Hill, Inc., 1974.

DeSanctis, G. and R. B. Gallupe. "A Foundation for the Study of Group Decision Support Systems." *Management Science*, 33, no. 5, May 1987.

Dhar, V. and R. Stein. *Intelligent Decision Support Methods: The Science of Knowledge*. Upper Saddle River, NJ: Prentice-Hall, 1997.

Donovan, J.J. and S.E. Madnick. “Institutional and Ad Hoc DSS and Their Effective Use”, *Data Base*, 8, no. 3, 1977.

Drucker, Peter. “The Next Information Revolution”. *Forbes*, August 24, 1998. URL http://www.forbes.com/.

Fedorowicz, J. "A Technology Infrastructure for Document-Based Decision Support Systems", in Sprague, R. and H. J. Watson, *Decision Support Systems: Putting Theory into Practice* (Third Edition), Prentice-Hall, 1993, pp. 125-136.

Gerrity, T. P., Jr. "The Design of Man-Machine Decision Systems". *Sloan Management Review*, vol. 12, no. 2, pp. 59-75, Winter 1971.

Golden, B., Hevner, A., and Power, D.J. Decision Insight Systems: A Critical Evaluation, *Computers and Operations Research*, 1986, 13(2/3), 287-300.

Hackathorn, R. D. and P.G.W. Keen. “Organizational Strategies for Personal Computing in Decision Support Systems”, *MIS Quarterly*, September 1981, pp.XX.

Holsapple, C.W. and A. B. Whinston. *Decision Support Systems: A Knowledge-based Approach*, Minneapolis, MN: West Publishing Co., 1996.

Inmon, W.H. "What is a Data Warehouse?", *PRISM*, vol. 1, No. 1, 1995, URL http://www.cait.wustl.edu/cait/papers/prism/vol1_no1/.

Keen, P. G. W. and M. S. Scott Morton. *Decision Support Systems: An Organizational Perspective.* Reading, MA: Addison-Wesley, Inc., 1978.

Little, J. D. C. "Models and Managers: The Concept of a Decision Calculus", *Management Science*, vol. 16, no. 8, pp. B466-485, April 1970.

Pendse, N. “The Origins of Today’s OLAP Products”, *The OLAP Report*, last updated Novemebr 4, 1999 at URL http://www.olapreport.com/origins.htm.

Power, D. J. "What is a DSS?", *DSStar*, vol. 1, no. 3, October 21, 1997, http://dssresources.com/papers/whatisadss.

Scott Morton, Michael S. *Management Decision Systems: Computer-based Support for Decision Making*. Boston, MA: Division of Research, Graduate School of Business Administration, Harvard University, 1971.

Sprague, R.H. and E.D. Carlson. *Building Effective Decision Support Systems.* Englewood Cliffs, NJ: Prentice-Hall, 1982.

Sprague, R. H. and Hugh J. Watson (editors). *Decision Support for Management.* Englewood Cliffs, N.J.: Prentice-Hall, Inc.: 1996.

Turban, E. *Decision Support and Expert Systems: Management Support Systems*. (Fourth Edition) Englewood Cliffs, NJ: Prentice Hall, Inc, 1995.

This chapter is a working draft. I would appreciate your comments and feedback. This chapter may be used for non-commercial or academic purposes during Year 2000. Last revised by D. J. Power, Sunday, September 24, 2000, email power@dssresources.com.

Chapter 1
Introduction to Decision Support Systems

School of Computer Studies

Mindanao State University – Iligan Institute of Technology

REFERENCES:

- Decision Support and Data Warehouse Systems by Efrem C. Mallach
- Computing Essentials by: O'Leary

Evolution of DSS

Early Computers	Was used to automate repetitive calculations
WW II	Under the pressure of this era, advances in computing was produced
1950's	E- computers where used in the commercial realm
Automatic data processing	Electronic Data processing (EDP) – punch card processing
Data processing	Manual data processing became a historical relic

Evolution of DSS

Transaction Processing	Describes the repetitive processing of common business events
Information Reporting System	17 – inch thick reports some are relevant some are not
Management Information System (MIS)	**Idea** to store all of a firms data. E.g: customers, ordes, inventory, production. . .
(1970's) DSS	Systems where made for a more specific types of decision
(1980's) Data Warehouse	From data to data warehouse

Information Systems (IS)

- is a combination of people, hardware, software, communication devices, network and data resources that processes (can be storing, retrieving, transforming information) data and information for a specific purpose.

Information Flow

- Based on needs of management
  - Top level management
  - Middle management
  - Supervisors
- Based on needs of organization

Information Flow: Supervisors

* Primarily vertical
* Provide status Reports

Information Flow: Middle Managers

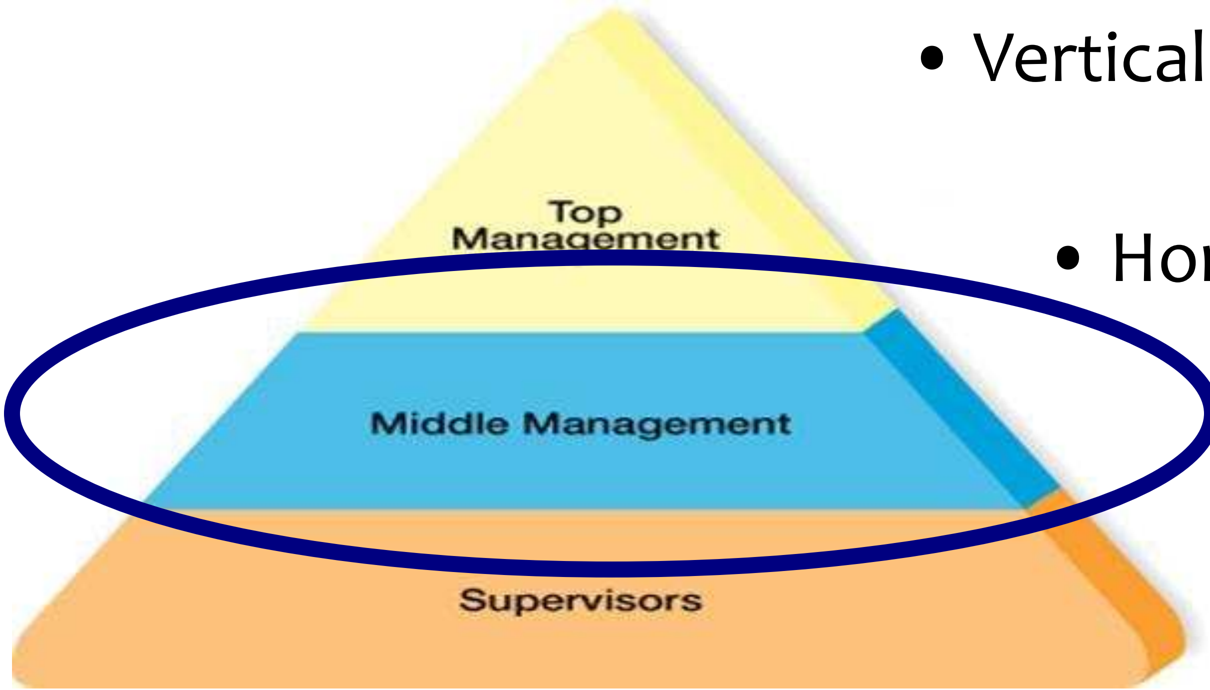

- Vertical
- Horizontal

Information Flow: Top Managers

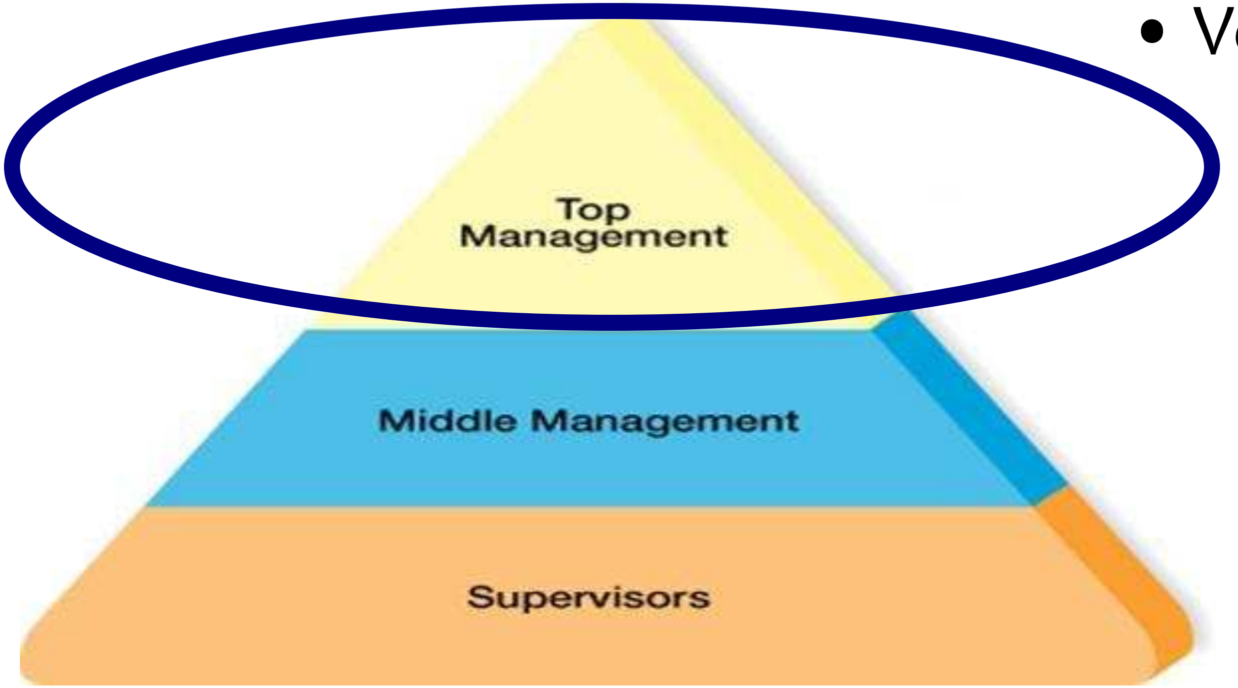

- Vertical
- External
- Horizontal

Information Systems (IS)

- **TPS** – Transaction Processing System
- **MIS** – Management Information System
- **DSS** – Decision Support System
- **OIS** – Office information System
- **EIS** - Executive Information System
- **PIS** – Personal Information System
- **WIS** – Workgroup Information System
- **ES** – Expert System
- **SIS** – Strategic Information System

Information Systems (IS)

- the mentioned IS’s are not distinct categories
- They overlap, interact, and supplement each other

Transaction Processing System (TPS)

- often form the foundation of a strategic information systems
- Keep track of routine operations and records
- An essential TPS is accounting
- Accounting TPS activities
  - Sales orders - Accounts payable
  - Accounts receivable - Payroll
  - Inventory control - General ledger

Management Information System (MIS)

- often form the foundation of a strategic information systems
- Produce standardized reports
  - Periodic
  - Exception
  - Demand
- Use databases created by TPS
- Integrate databases across departments

Decision Support System (DSS)

- Flexible tool for dealing with unanticipated questions
  - Analyzes data
  - Reports are not in fixed format
- Generally easy to use
- Analytical capabilities come from decision models

DSS Parts

- User
  - Makes decisions
- System software
  - Operating system
- Data
  - Internal from the organization
  - External from the organization
- Decision models
  - Strategic, tactical, operational

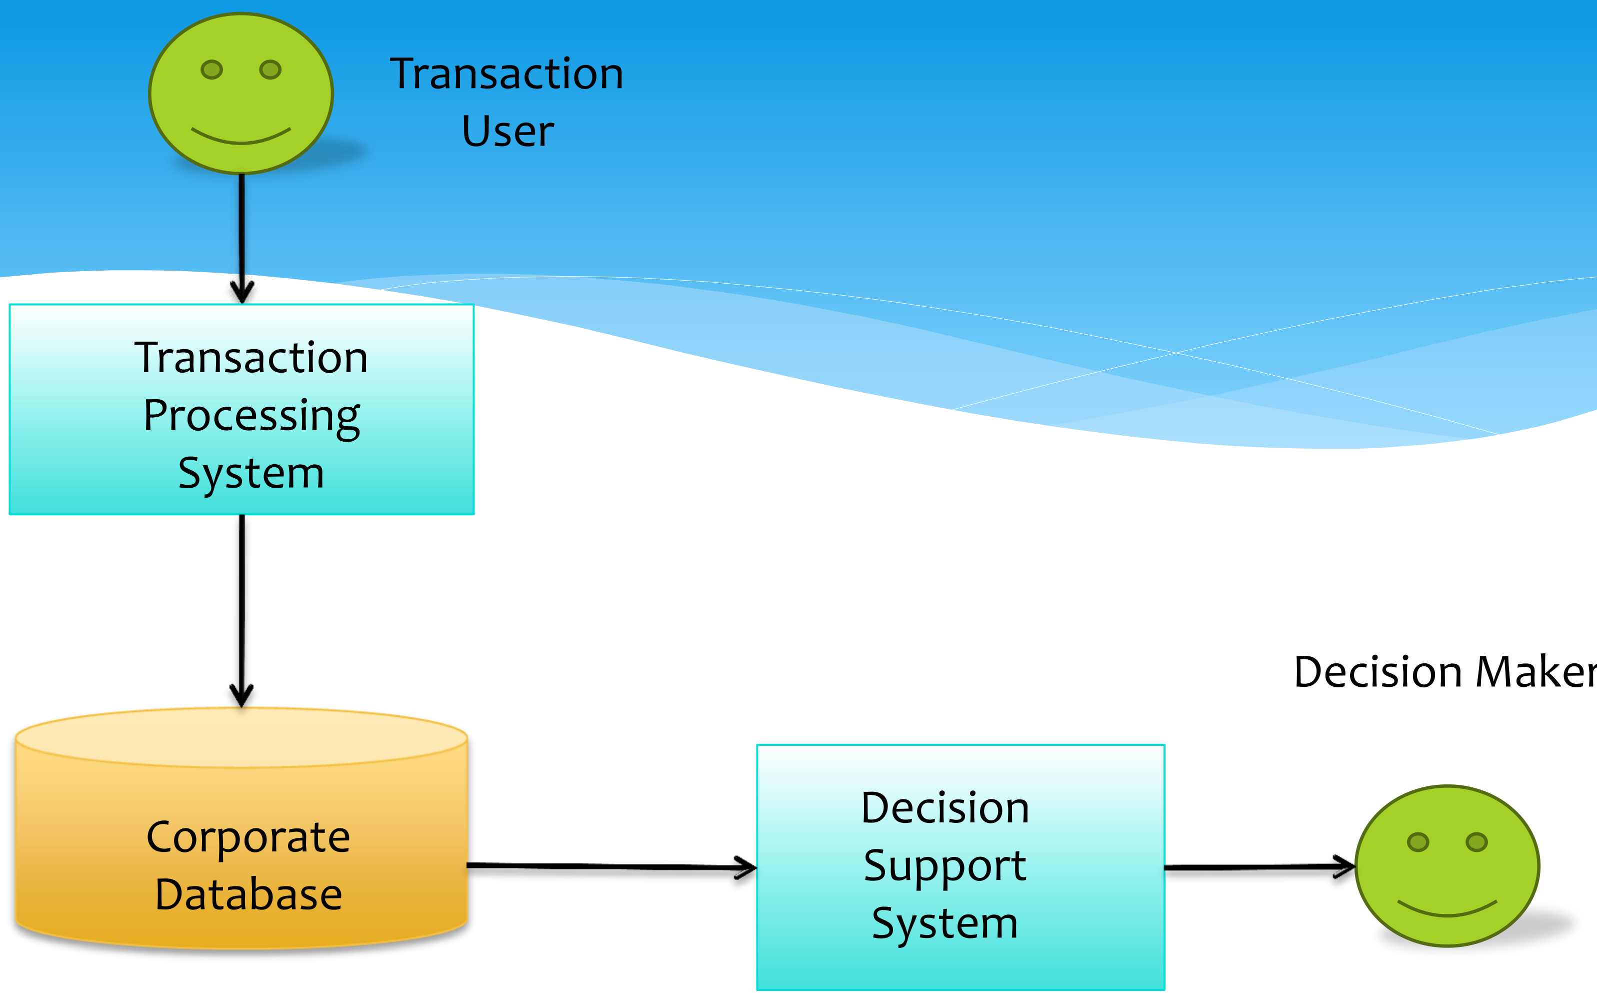

Flow Diagram of Transaction Data Feeding

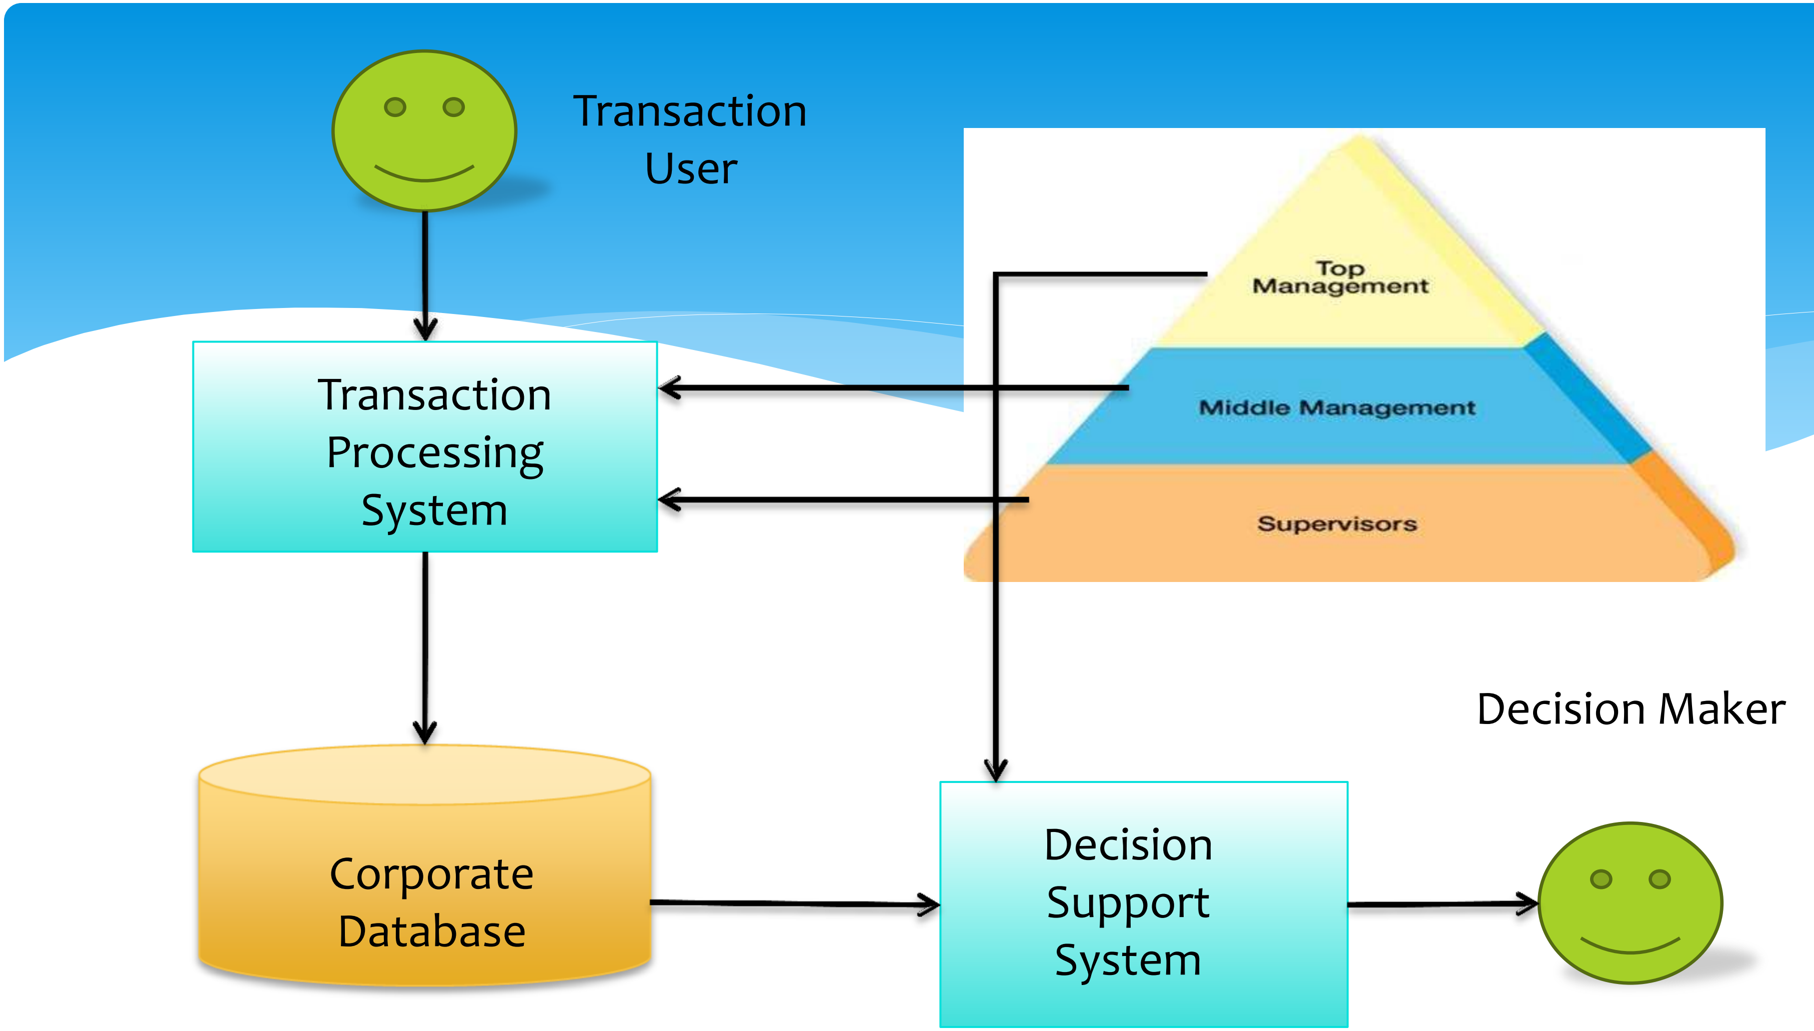

Flow Diagram of Transaction Data Feeding

Decision Support System (DSS)

- **Is therefore used to improve the effectiveness***(adequacy to accomplish a purpose; producing intended or expected result)* **rather than just the efficiency** *(performing or functioning in the best possible manner with the least waste of time and effort)* **of a decision making.**

100 %

Decision Support Content

Decision Support System

Executive Information System

Expert System

Information Reporting System

Workgroup Information System

Personal Information System

Office Information System

Transaction Processing System

0 %

Approximate Decision Support Content of Different Types of Information Systems

7 Characteristics of DSS :
4 are held in common and 3 are optional

All DSS. . .

- Are information systems
- Used by knowledge workers
- Used in making decisions
- Support, do not replace, people

And Some DSS. . .

- Used in semi structured decisions
- Incorporate models
- Incorporate a database

DSS Benefits

1. Improving personal efficiency
2. Improving problem solving
3. Facilitating communication
4. Promoting learning or training
5. Increasing organizational control

5 Categories of DSS

1. Communication-Driven DSS
2. Data-Driven DSS
3. Document-Driven DSS
4. Knowledge-Driven DSS
5. Model-Driven DSS

Communication-Driven DSS

Most communications-driven DSSs are targeted at internal teams, including partners. Its purpose are to help conduct a meeting, or for users to collaborate. The most common technology used to deploy the DSS is a web or client server.

Examples:

chats and instant messaging software, online collaboration and net-meeting systems.


Skype™ - sylvia_cooper1
Sylvia Cooper
James Cooper
New
Conversations
Contacts


Yahoo! Messenger
Benster
happy Friday everyone!
Friends
Alleycat212@hotmail.com
Dan Everest
hiking this weekend?
Jessica Green
I'm mobile
3 Viewers

Data-Driven DSS

Most data-driven DSSs are targeted at managers, staff and also product/service suppliers. It is used to query a database or data warehouse to seek specific answers for specific purposes. It is deployed via a main frame system, client/server link, or via the web.

Examples:

computer-based databases that have a query system to check (including the incorporation of data to add value to existing databases.

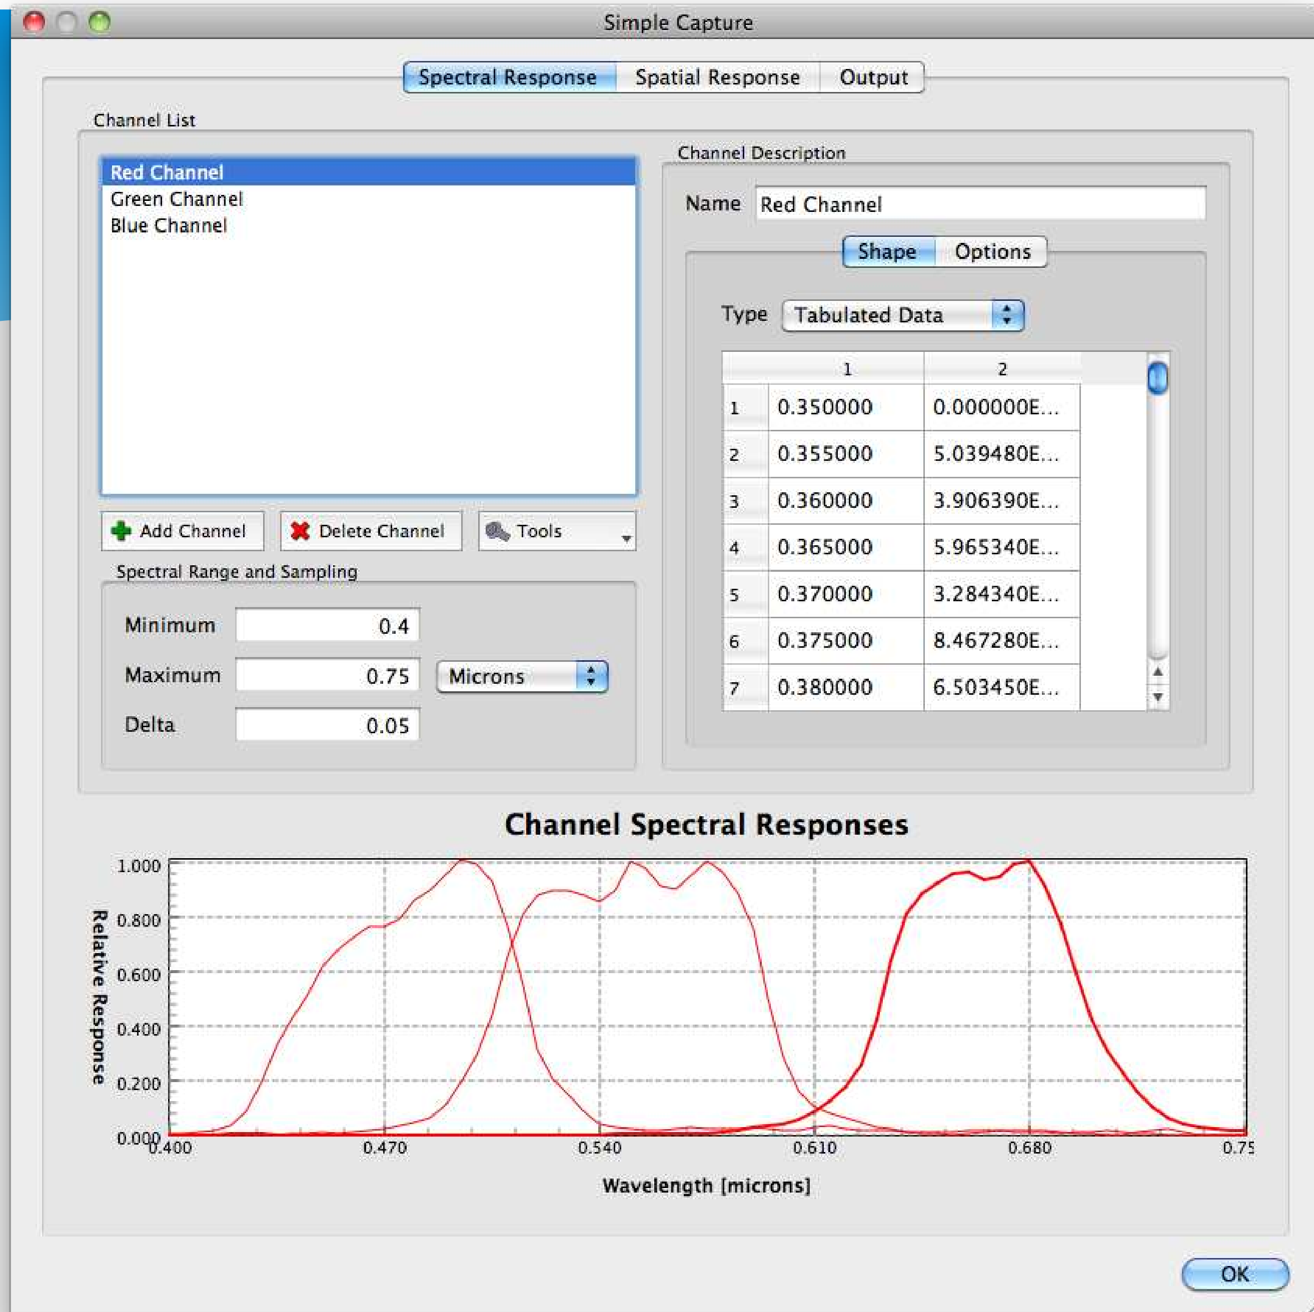
Simple Capture
Spectral Response
Spatial Response
Output
Channel List
Red Channel
Green Channel
Blue Channel
Add Channel
Delete Channel
Tools
Spectral Range and Sampling
Minimum 0.4
Maximum 0.75
Microns
Delta 0.05
Channel Description
Name Red Channel
Shape
Options
Type Tabulated Data
1 2
1 0.350000 0.000000E...
2 0.355000 5.039480E...
3 0.360000 3.906390E...
4 0.365000 5.965340E...
5 0.370000 3.284340E...
6 0.375000 8.467280E...
7 0.380000 6.503450E...
Channel Spectral Responses
Relative Response
1.000
0.800
0.600
0.400
0.200
0.000
0.400
0.470
0.540
0.610
0.680
Wavelength [microns]
OK

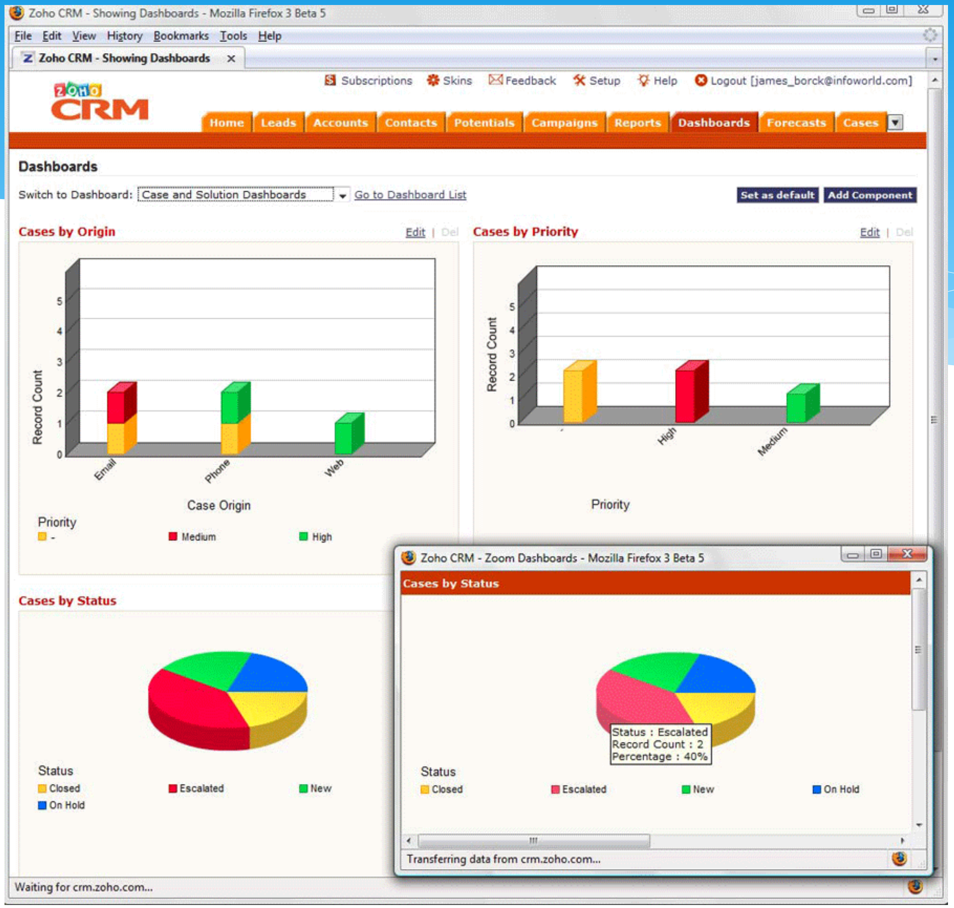

Zoho CRM - Showing Dashboards - Mozilla Firefox 3 Beta 5
File Edit View History Bookmarks Tools Help
Zoho CRM - Showing Dashboards
Subscriptions Skins Feedback Setup Help Logout [james_borck@infoworld.com]
Home Leads Accounts Contacts Potentials Campaigns Reports Dashboards Forecasts Cases
Dashboards
Switch to Dashboard: Case and Solution Dashboards Go to Dashboard List
Set as default Add Component
Cases by Origin
Edit | Del
Record Count
Email Phone Web
Case Origin
Priority
- Medium High
Cases by Priority
Edit | Del
Record Count
High Medium
Priority
Cases by Status
Status
Closed Escalated New On Hold
Zoho CRM - Zoom Dashboards - Mozilla Firefox 3 Beta 5
Cases by Status
Status : Escalated
Record Count : 2
Percentage : 40%
Status
Closed Escalated New On Hold
Transferring data from crm.zoho.com...
Waiting for crm.zoho.com...

Document-Driven DSS

Document-driven DSSs are more common, targeted at a broad base of user groups. The purpose of such a DSS is to search web pages and find documents on a specific set of keywords or search terms. The usual technology used to set up such DSSs are via the web or a client/server system.


Media Finder
Media Finder
Search
Downloads

Knowledge-Driven DSS

Knowledge-driven DSSs or 'knowledgebase' are they are known, are a catch-all category covering a broad range of systems covering users within the organization setting it up, but may also include others interacting with the organization - for example, consumers of a business.

Model-Driven DSS

Model-driven DSSs are complex systems that help analyze decisions or choose between different options. These are used by managers and staff members of a business, or people who interact with the organization, for a number of purposes depending on how the model is set up - scheduling, decision analyses etc.

END

Thank you for listening
God bless!

What is “Groupware?”

- **Groupware** is a term that refers to **technology** designed to help people collaborate and includes a wide range of applications. ...
  Collaborative **management** tools: Tools for **managing** group activities, e.g. project **management** systems, workflow systems,
  information **management** systems, etc.

What is "Groupware?"

- Tools (hardware, software, processes) that support person-to-person collaboration
- This can include e-mail, bulletin boards, conferencing systems, decision support systems, video and workflow systems, etc…
- Some common groupware acronyms:
  - Group Support Systems (GSS)
  - Group Decision Support Systems (GDSS)
  - Electronic Meeting Systems (EMS)
  - Bulletin Board Systems (BBS)
  - Group Collaboration Systems (GCS)
  - Computer-Supported Cooperative Work (CSCW) systems

Three types of interaction

All collaboration technology always implies:

- *1. Human - computer interaction*
- *2. Human - database interaction: (information seeking)*
    - Internets
    - Intra-net
    - Group networks
- *3. Mediated interpersonal interaction (communication)*

A Simple Classification of Groupware (adapted from Johansen, 1991)

Location	**Time**: Same	**Time**: Different
Same	**GDSS; Support for FtF Meetings,**	**Email, Bulletin board, Comp. Conf**
Different	**Teleconferencing, Instant Msg Chat, Whiteboard, Video**	**Email, Bulletin board, Comp. Conf Web based CS**

A GDSS Example

Video Conferencing

Groupware system

Five Basic group processes

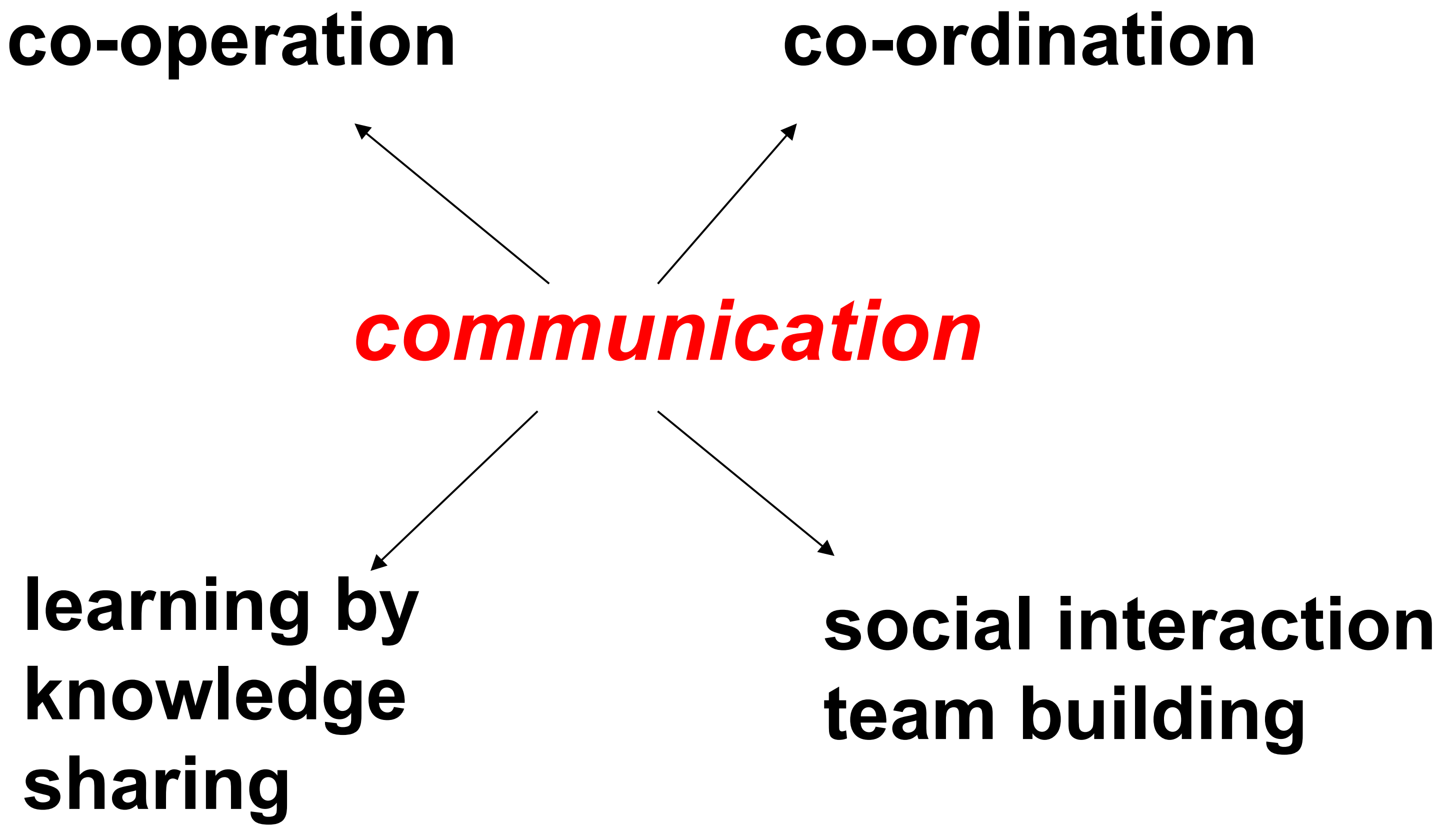

Dynamic Group Interaction model

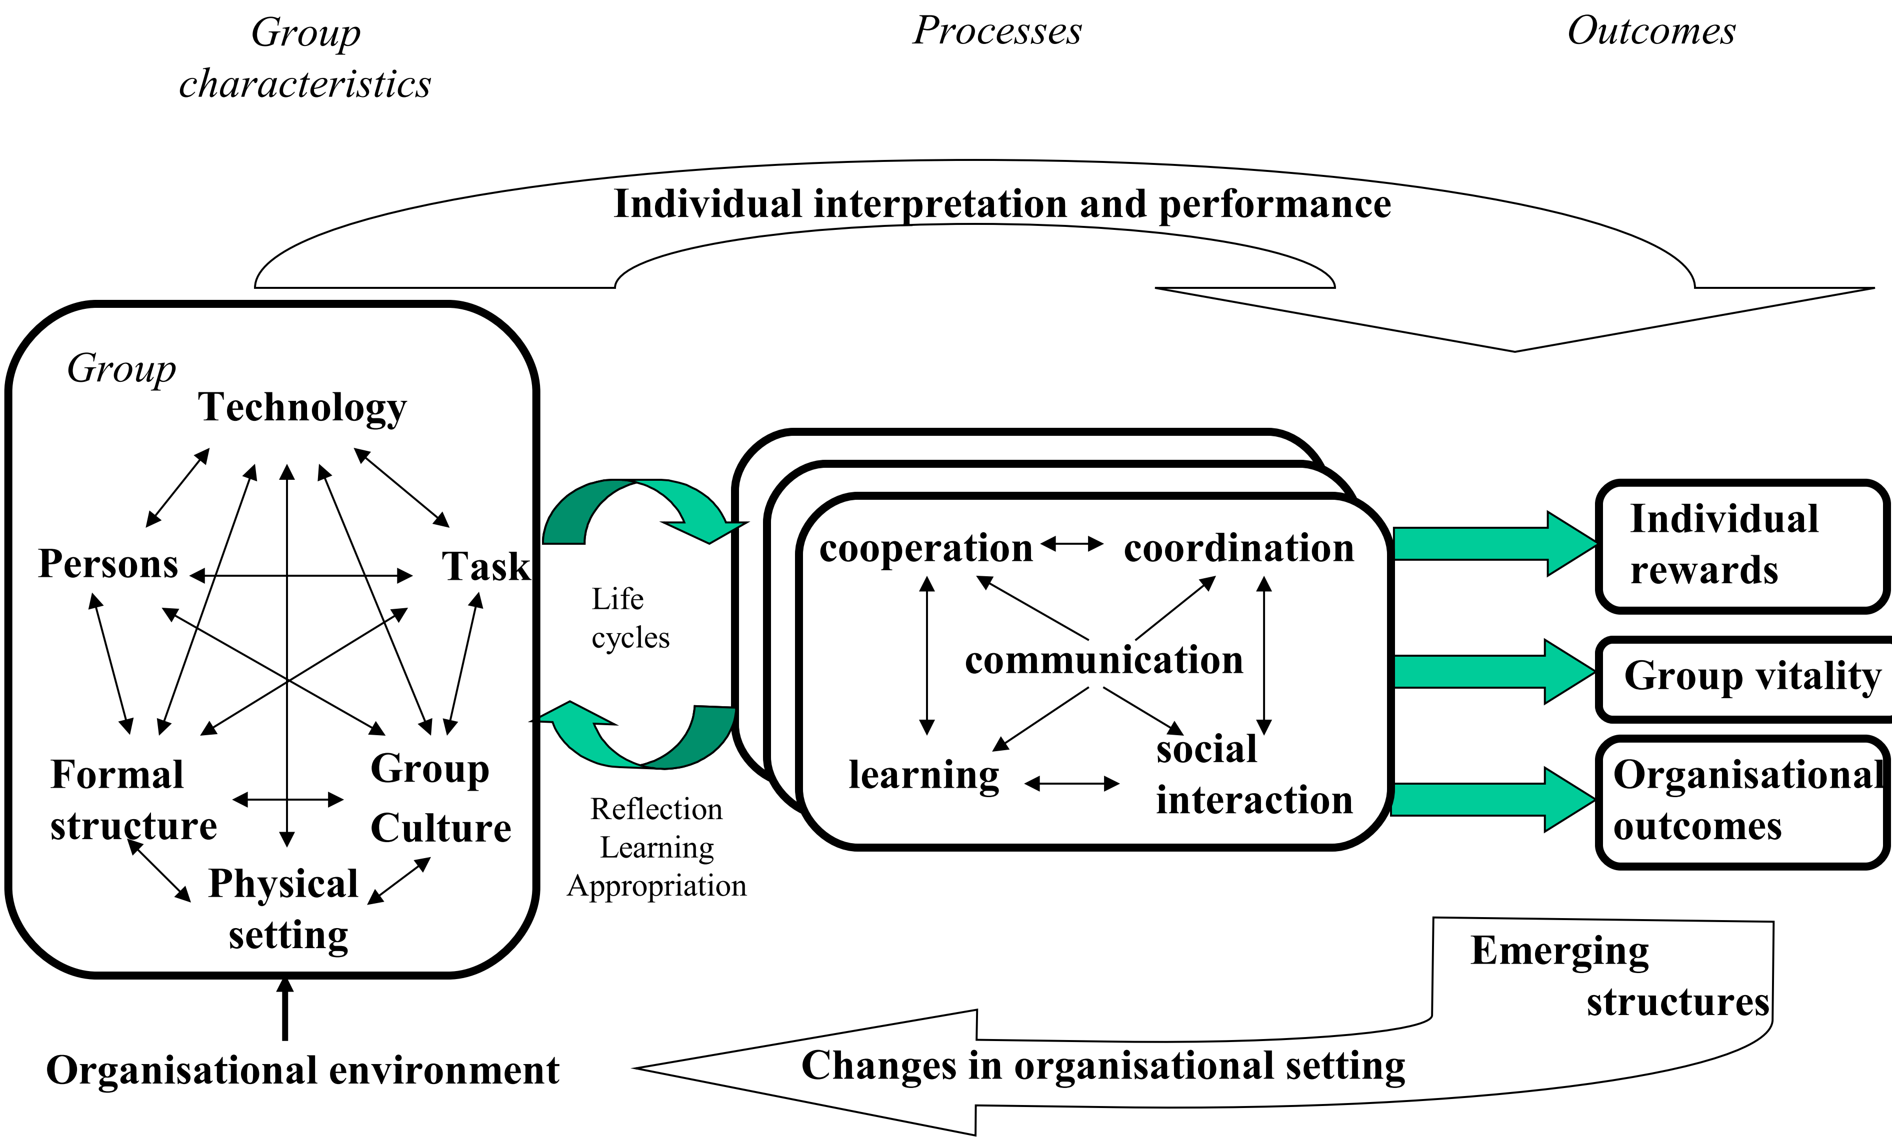

Basic Principles

- The effectiveness of a group can be expressed in terms of three types of outcomes, i.e. (quality and quantity of the)**products**, **individual ‘rewards’** and **vitality of the social relations**.
- Effectiveness depends on the quality of the individual preformance and six group processes, which have to match
- The quality of the group processes depends on the support of six conditions, and on the interaction with the environment.
- The six aspects of the context-of-use have to fit to each other.
- Groups develop and tools become adopted and adapted to, through interaction processes and feedback.

SUPPORT – MATCH – ADAPTATION

To evaluate the role of groupware tools.

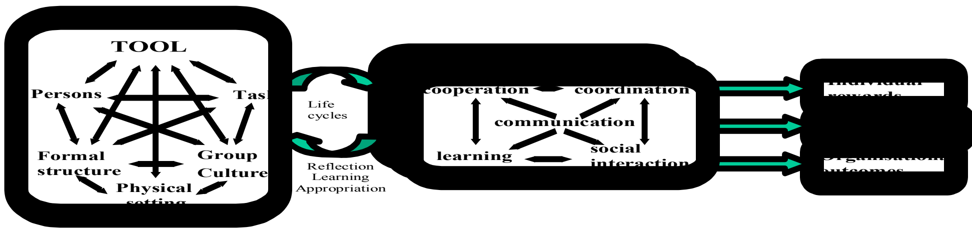

Evaluate

- Technical properties
- Degree of fit to task users, group, setting, other characteristics
- Degree of support for processes
- Effects on outcomes
- Possibilities to adapt

Evaluation issues in more detail

1. Describe the tool characteristics: *reliability, portability, maintainability, network performance, costs, infrastructural quality, security/privacy* and evaluate whether this is adequate (ISO-9126)
2. Describe the functionalities
3. Analyse the task and evaluate whether the functionalities fit the task
4. Analyse the users and evaluate whether the tool fits the users (*usability*)
5. Analyse the group (*structure, culture, setting*) and evaluate whether the tool fits the group
6. Evaluate whether the tool supports (or at least does not hinder) the group processes: *communication, co-operation, co-ordination, learning, social interaction*.
7. Evaluate whether the tool contributes to (or at least does not hinder) individual, group, organisational *outcomes*.
8. Evaluate to which extend the tools can be adapted to learning and new uses

Evaluation principles

- - Evaluation should be integrated in the design process from the very beginning.
- - The design and the evaluation process should be iterative and stakeholder centred; critical success factors of stakeholders should be formulated.
- - Evaluation can take place at three periods in the design life cycle
  - *a. Concept evaluation*
  - *b. Prototype evaluation*
  - *c. Operational assessment*

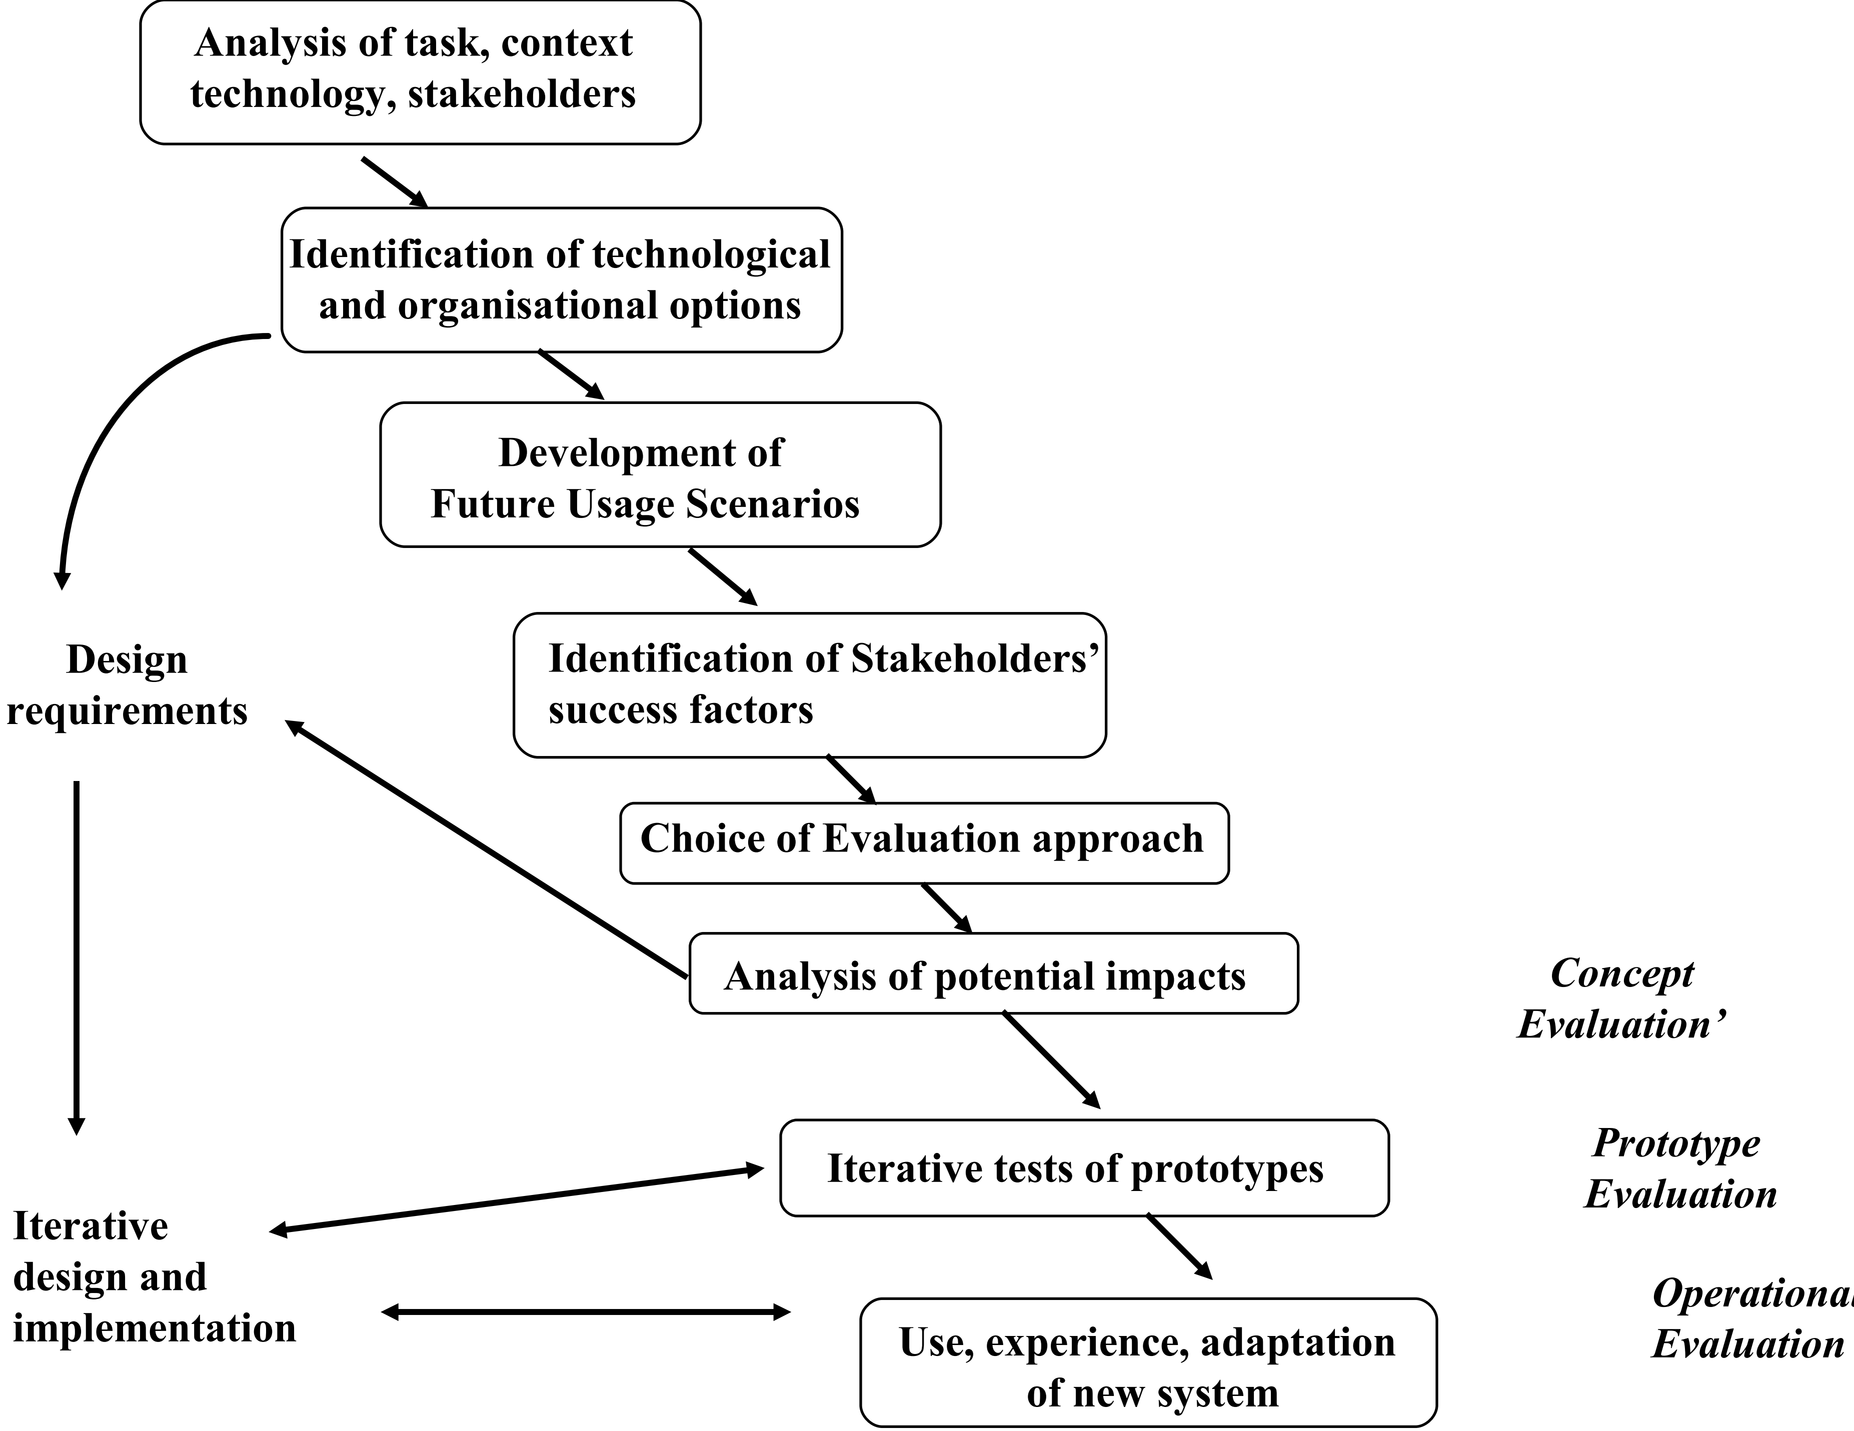

Analysis of task, context technology, stakeholders
Identification of technological and organisational options
Development of Future Usage Scenarios
Identification of Stakeholders' success factors
Choice of Evaluation approach
Analysis of potential impacts
Iterative tests of prototypes
Use, experience, adaptation of new system
Design requirements
Iterative design and implementation
Concept Evaluation'
Prototype Evaluation
Operational Evaluation

Lessons learned (1)

1. Groupware is part of a social system. Design not for a tool as such but for a new socio-technical setting.
2. Design for several levels of interaction, i.e. for user friendly human computer interaction, adequate interpersonal communication, group co-operation and organisational functioning.
3. Design in a participative way, i.e. users and possibly other stakeholders should be part of the design process from the beginning.
4. Analyse carefully the situation of the users. Success of collaboration technology depends on the use and the users, not on the technology. Introduction should match their skills and abilities, and also their attitudes, otherwise resistance is inevitable.
5. Analyse carefully the context, since success of collaboration technology depends on the fit to that context. The more a new setting deviates from the existing one the more time, energy and other resources should be mobilised to make it a success.

Lessons learned (2)

6. Introduce the new system carefully. Apply proper project management, find a champion, try a pilot, inform people intensively
7. Train and support end-users extensively
8. Measure success conditions and success criteria before, during and after the development process. Only in this way you can learn for future developments.
9. Plan for a long process of introduction, incorporation, evaluation and adaptation. Groupware is not a quick fix.
10. Despite careful preparations groupware is appropriated and adapted in unforeseen ways. Keep options open for new ways of working with the groupware, because this may result in creative and innovative processes.

Evaluation methods

1. *Inspection methods*
- Heuristic Evaluation:

2. *Performance analysis*
- Human Reliability Analysis

3. *Behaviour analysis*
- Diagnostic Recorder for Usability Measurement (DRUM)

4. *Effort and satisfaction*
- NASA-Task Load IndeX (NASA-TLX)
- Measuring the Usability of Multi-Media Systems (MUMMS)
- MultiMedia Communication Questionnaire (MMCQ)

Evaluation Methods (2)

5. *Task aspects and relations*
- Extended Delft Measurement Kit

6. N*etwork performance*

7. *System usage and interaction registration*
- Automatic registration of the use of the system
- Coding schemes for communication content

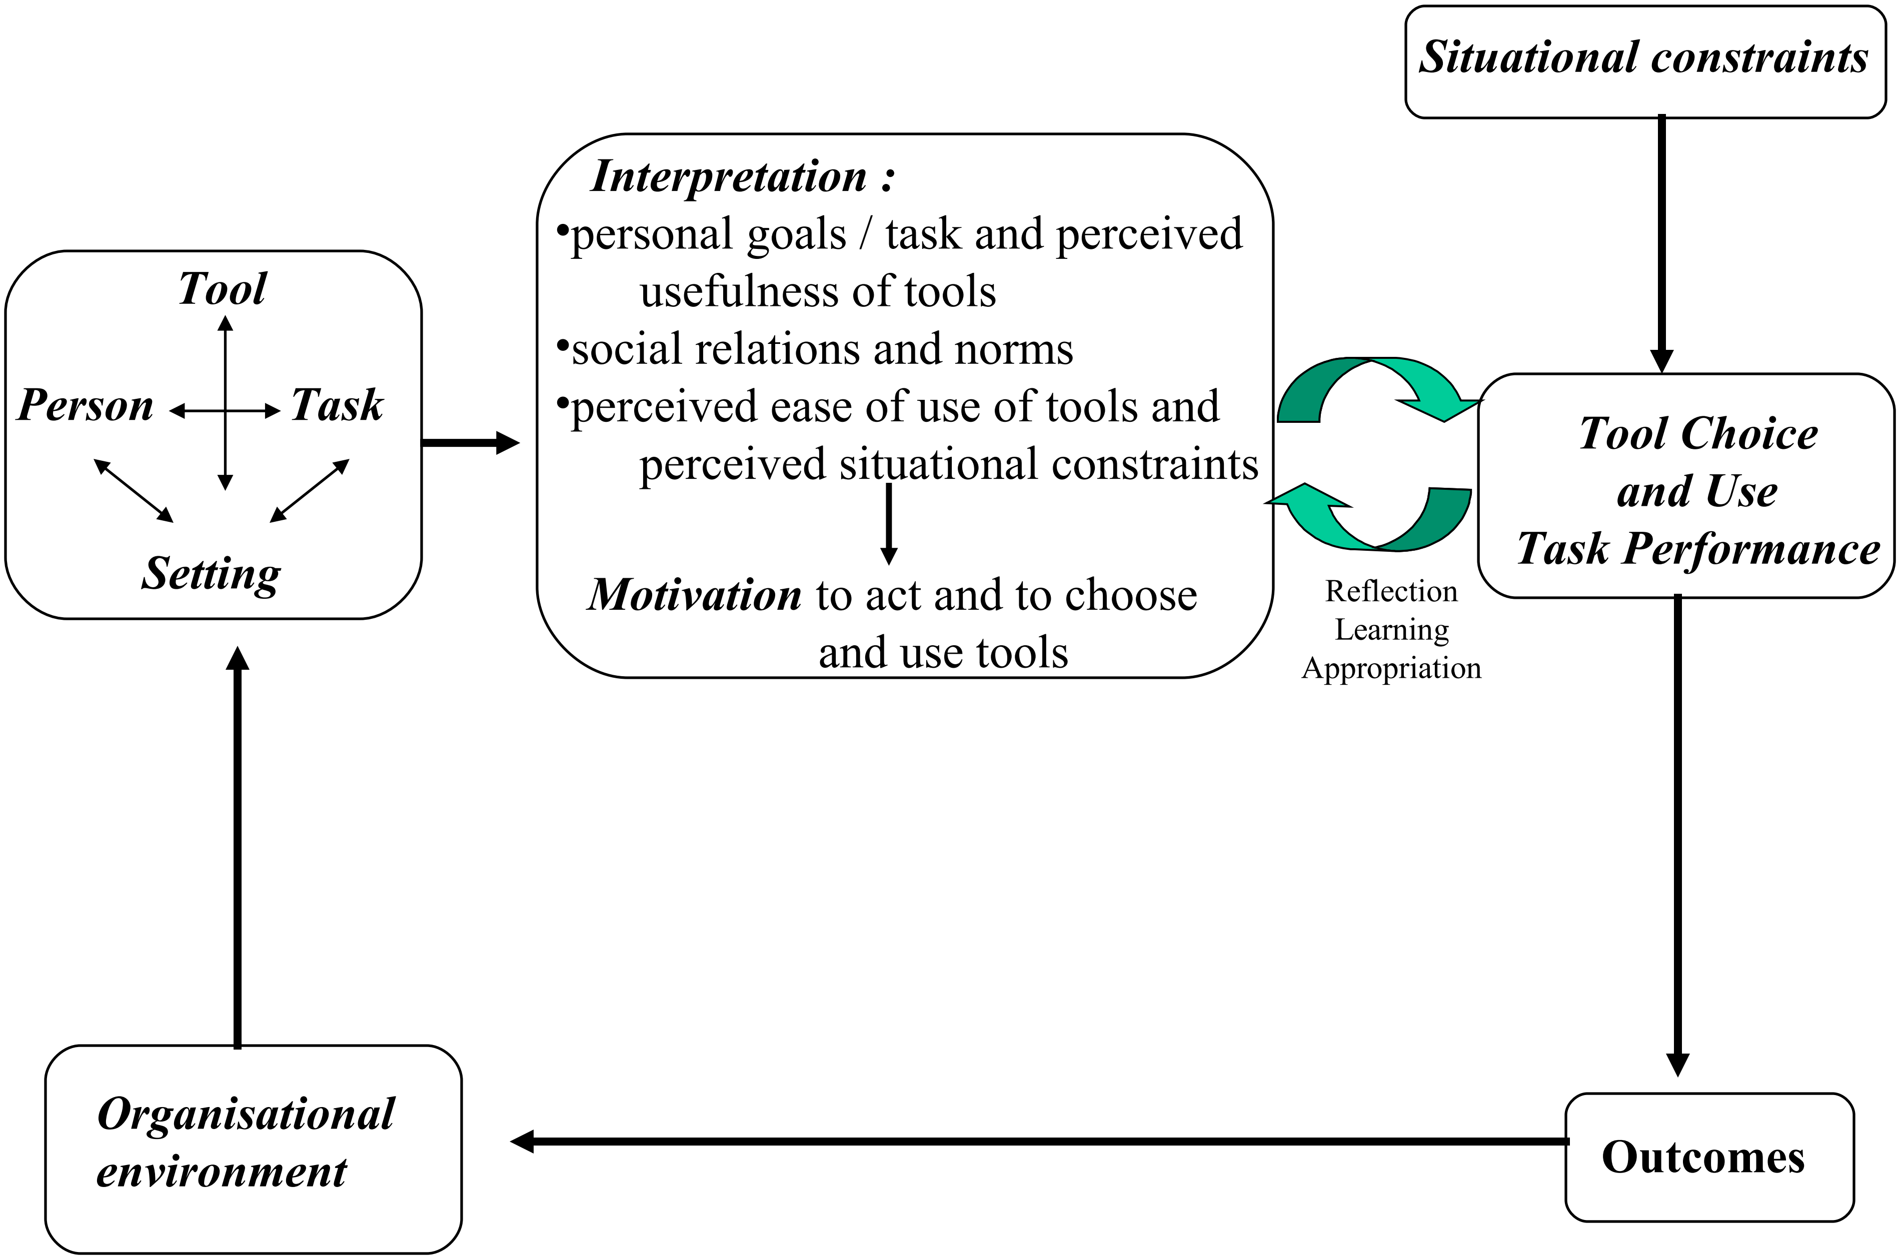
Situational constraints
Interpretation :
•personal goals / task and perceived usefulness of tools
•social relations and norms
•perceived ease of use of tools and perceived situational constraints
Motivation to act and to choose and use tools
Tool
Person
Task
Setting
Tool Choice and Use Task Performance
Reflection
Learning
Appropriation
Organisational environment
Outcomes

Diversity Hypothesis

Fully Dispersed | Three Subgroups | Two Subgroups

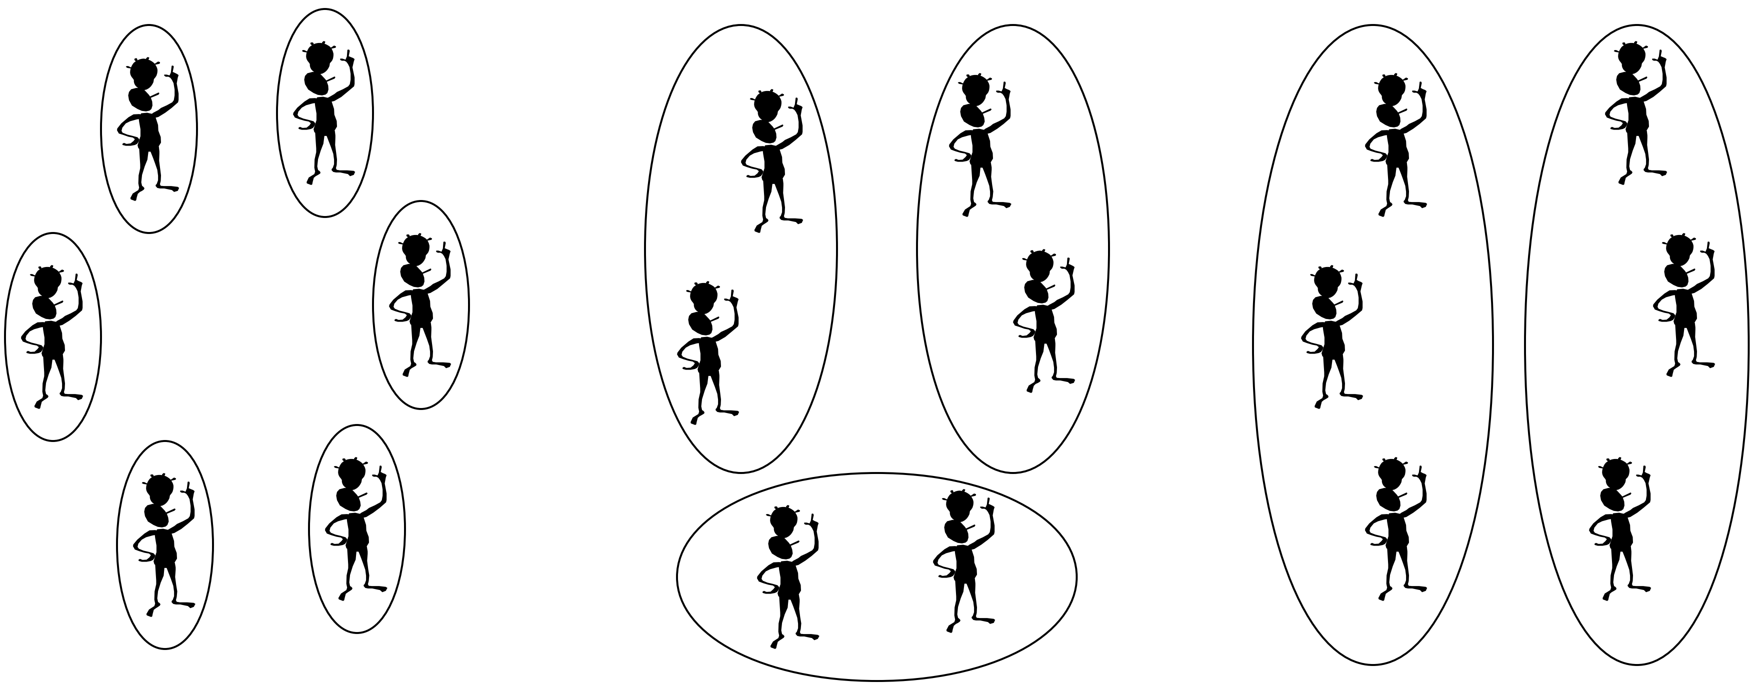

Most Conflict
Least Trust

Least Conflict
Most Trust

DECISION SUPPORT SYSTEMS

Neeraj Kumar Jain

Institute of Technology & Science
Feb 2020

INTRODUCTION

- Decision makers are faced with increasingly stressful environments – highly competitive, fast-paced, near real-time, overloaded with information, data distributed throughout the enterprise, and multinational in scope.
- The combination of the Internet enabling speed and access, and the maturation of artificial intelligence techniques, has led to sophisticated aids to support decision making under these risky and uncertain conditions.
- These aids have the potential to improve decision making by suggesting solutions that are better than those made by the human alone.
- They are increasingly available in diverse fields from medical diagnosis to traffic control to engineering applications.

DECISION SUPPORT SYSTEM

- A Decision Support System (DSS) is an interactive computer-based system or subsystem intended to help decision makers use communications technologies, data, documents, knowledge and/or models to identify and solve problems, complete decision process tasks, and make decisions.
- Decision Support System is a general term for any computer application that enhances a person or group's ability to make decisions.
- Also, Decision Support Systems refers to an academic field of research that involves designing and studying Decision Support Systems in their context of use.

SPREADSHEET-BASED DECISION SUPPORT SYSTEMS

- A DSS is made up of a model (or models), a source of data, and a user interface.
- When a model is implemented in Excel, it is possible to use Visual Basic for Applications (VBA) to make the system more efficient by automating interactive tasks that users would otherwise have to repeat routinely.
- VBA can also make the system more powerful by extending the functionality of a spreadsheet model and by customizing its use.

DECISION MAKING CHARACTERISTICS

- Decision is made based on the information available.
- At each part of the assessment, there may have to be iterative development to take account improvement in data that take place as the project proceeds.
- A project will not go ahead unless there is adequate funding.

MANAGEMENT

- ***Management is decision making***
- **The *manager* is a decision maker**
- **Organizations are filled with decision makers at different level.**
- **Management is considered as art: a talent acquired over years by trial-and-error.**
- **However decision making today is becoming more complicated:**
  - **Technology / Information/Computers : increasing→ More alternative to choose**
  - **Structural Complexity / Competition : increasing→ larger cost of error**
  - **International markets / Consumerism : increasing→ more uncertainty about future**
  - **Changes, Fluctuations : increasing→ need for quick decision**

Management Problems

- Most management problems for which decisions are sought can be represented by three standard elements – objectives, decision variables, and constraints.
- **Objective**
  - Maximize profit
  - Provide earliest entry into market
  - Minimize employee discomfort/turnover
- **Decision variables**
  - Determine what price to use
  - Determine length of time tests should be run on a new product/service
  - Determine the responsibilities to assign to each worker
- **Constraints**
  - Can't charge below cost
  - Test enough to meet minimum safety regulations
  - Ensure responsibilities are at most shared by two workers

TYPES OF PROBLEMS

- **Structured:** situations where the procedures to follow when a decision is needed can be specified in advance
  - Repetitive
  - Standard solution methods exist
  - Complete automation may be feasible
- Unstructured: decision situations where it is not possible to specify in advance most of the decision procedures to follow
  - One-time
  - No standard solutions
  - Rely on judgment
  - Automation is usually infeasible
- Semi-structured: decision procedures that can be pre specified, but not enough to lead to a definite recommended decision
  - Some elements and/or phases of decision making process have repetitive elements

DSS most useful for repetitive aspects of semi-structured problems

DSS IN SUMMARY

- A MANAGEMENT LEVEL COMPUTER SYSTEM Which:
  - COMBINES DATA,
  - MODELS,
  - USER - FRIENDLY SOFTWARE

FOR SEMISTRUCTURED & UNSTRUCTURED DECISION MAKING.

- It utilizes data, provides an easy-to-use interface, and allows for the decision maker's own insights.

WHY DSS?

- Increasing complexity of decisions
  - Technology
  - Information:
    - "Data, data everywhere, and not the time to think!"
  - Number and complexity of options
  - Pace of change
- Increasing availability of computerized support
  - Inexpensive high-powered computing
  - Better software
  - More efficient software development process
- Increasing usability of computers

Perceived Benefits

- decision quality
- improved communication
- cost reduction
- increased productivity
- time savings
- improved customer and employee satisfaction

A BRIEF HISTORY

- Academic Researchers from many disciplines has been studying DSS for approximately 40 years.
- According to Keen and Scott Morton (1978), the concept of decision support has evolved from two main areas of research: the theoretical studies of organizational decision making done at the Carnegie Institute of Technology during the late 1950s and early 1960s, and the technical work on interactive computer systems, mainly carried out at the Massachusetts Institute of Technology in the 1960s.
- It is considered that the concept of DSS became an area of research of its own in the middle of the 1970s, before gaining in intensity during the 1980s.

A BRIEF HISTORY

- In the middle and late 1980s, Executive Information Systems (EIS), group decision support systems (GDSS), and organizational decision support systems (ODSS) evolved from the single user and model-oriented DSS.
- Beginning in about 1990, data warehousing and on-line analytical processing (OLAP) began broadening the realm of DSS.
- As the turn of the millennium approached, new Web-based analytical applications were introduced.

HISTORY OF DSS

Goal: Use best parts of IS, OR/MS, AI & cognitive science to support more effective decision

APPROACHES TO THE DESIGN AND CONSTRUCTION OF DSS

- Studies on DSS development conducted during the last 15 years have identified more than 30 different approaches to the design and construction of decision support methods and systems.
- Interestingly enough, none of these approaches predominate and the various DSS development processes usually remain very distinct and project-specific.
- This situation can be interpreted as a sign that the field of DSS development should soon enter in its formalization stage.

A SUMMARY OF COMMERCIAL DSS SYSTEM

- A summary of commercial DSS system show seven types of DSS:
  - File Drawer Systems, that provide access to the data items.
  - Data Analysis systems, that support manipulation of data by computerized tools for a specific task.
  - Analysis Information systems, that provide access to a series of decision oriented databases and small models.
  - Accounting and financial models, that calculates the consequences of possible actions.
  - Representational model, that estimates the consequences of actions based on simulation models.
  - Optimization models, that provide guidelines for action by generating an optimal solution
  - Suggestion models, that perform the logical processing to a specific suggested decision for a task.

A MULTIDISCIPLINE STUDY

- It is clear that DSS belong to an environment with multidisciplinary foundations, including (but not exclusively):
  - Database research,
  - Artificial intelligence,
  - Human-computer interaction,
  - Simulation methods,
  - Software engineering, and
  - Telecommunications.

TAXONOMIES

- Using the mode of assistance as the criterion, *Power* (2002) differentiates five types for DSS:
  - communication-driven DSS,
  - data-driven DSS,
  - document-driven DSS,
  - knowledge-driven DSS, and
  - model-driven DSS.

MODEL-DRIVEN DSS

- A model-driven DSS emphasizes access to and manipulation of a statistical, financial, optimization, or simulation model. Model-driven DSS use data and parameters provided by users to assist decision makers in analyzing a situation; they are not necessarily data intensive. Dicodess is an example of an open source model-driven DSS generator (Gachet 2004).
- Other examples:
  - A spread-sheet with formulas in
  - A statistical forecasting model
  - An optimum routing model

DATA-DRIVEN (RETRIEVING) DSS

- A data-driven DSS or data-oriented DSS emphasizes access to and manipulation of a time series of internal company data and, sometimes, external data.
- Simple file systems accessed by query and retrieval tools provides the elementary level of functionality. Data warehouses provide additional functionality. OLAP provides highest level of functionality.
- Examples:
  - Accessing AMMIS data base for all maintenance Jan89-Jul94 for CH124
  - Accessing INTERPOL database for crimes by ....
  - Accessing border patrol database for all incidents in Sector ...

MODEL AND DATA-RETRIEVING DSS

- Examples:
  - Collect weather observations at all stations and forecast tomorrow's weather
  - Collect data on all civilian casualties to predict casualties over the next month

COMMUNICATION-DRIVEN DSS

- A communication-driven DSS use network and comminication technologies to faciliate collaboartion on decision making. It supports more than one person working on a shared task.
- examples include integrated tools like Microsoft's NetMeeting or Groove (Stanhope 2002), Vide conferencing.
- It is related to group decision support systems.

DOCUMENT-DRIVEN DSS

- A document-driven DSS uses storage and processing technologies to document retrieval and analysis. It manages, retrieves and manipulates unstructured information in a variety of electronic formats.
- Document database may include: Scanned documents, hypertext documents, images, sound and video.
- A search engine is a primary tool associated with document drivel DSS.

KNOWLEDGE-DRIVEN DSS

- A knowledge-driven DSS provides specialized problem solving expertise stored as facts, rules, procedures, or in similar structures. It suggest or recommend actions to managers.
- MYCIN: A rule based reasoning program which help physicians diagnose blood disease.

ARCHITECTURE

- Three fundamental components of DSS:
  - the database management system (DBMS),
  - the model management system (MBMS), and
  - the dialog generation and management system (DGMS).
- the Data Management Component stores information (which can be further subdivided into that derived from an organization's traditional data repositories, from external sources such as the Internet, or from the personal insights and experiences of individual users);
- the Model Management Component handles representations of events, facts, or situations (using various kinds of models, two examples being optimization models and goal-seeking models); and
- the User Interface Management Component is of course the component that allows a user to interact with the system.

A DETAILED ARCHITECTURE

- Even though different authors identify different components in a DSS, academics and practitioners have come up with a generalized architecture made of six distinct parts:
  - the data management system,
  - the model management system,
  - the knowledge engine,
  - The user interface,
  - the DSS architecture and network, and
  - the user(s)

TYPICAL ARCHITECTURE

- TPS: transaction processing system
- MODEL: representation of a problem
- OLAP: on-line analytical processing
- USER INTERFACE: how user enters problem & receives answers
- DSS DATABASE: current data from applications or groups
- DATA MINING: technology for finding relationships in large data bases for prediction

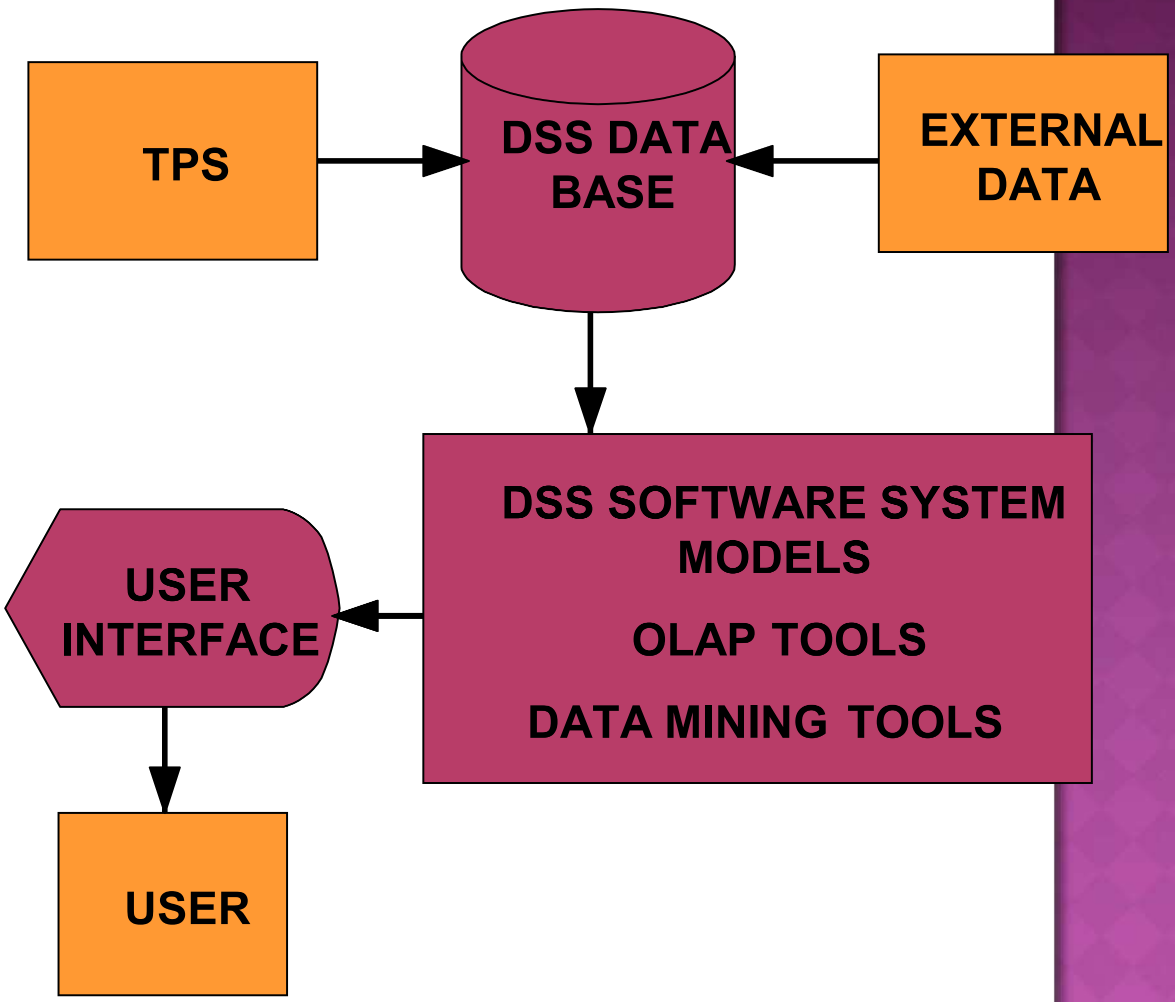

DSS MODEL BASE

- Model base
  - A software component that consists of models used in computational and analytical routines that mathematically express relations among variables
- Examples:
  - Linear programming models,
  - Multiple regression forecasting models
  - Capital budgeting present value models

APPLICATIONS

- There are theoretical possibilities of building such systems in any knowledge domain.
  - Clinical decision support system for medical diagnosis.
  - a bank loan officer verifying the credit of a loan applicant
  - an engineering firm that has bids on several projects and wants to know if they can be competitive with their costs.
  - DSS is extensively used in business and management. Executive dashboards and other business performance software allow faster decision making, identification of negative trends, and better allocation of business resources.
  - A growing area of DSS application, concepts, principles, and techniques is in agricultural production, marketing for sustainable development.
  - A specific example concerns the Canadian National Railway system, which tests its equipment on a regular basis using a decision support system.
  - A DSS can be designed to help make decisions on the stock market, or deciding which area or segment to market a product toward.

CHARACTERISTICS AND CAPABILITIES OF DSS

- The key DSS characteristics and capabilities are as follows:
  1. Support for decision makers in semistructured and unstructured problems.
  2. Support managers at all levels.
  3. Support individuals and groups.
  4. Support for interdependent or sequential decisions.
  5. Support intelligence, design, choice, and implementation.
  6. Support variety of decision processes and styles.
  7. DSS should be adaptable and flexible.
  8. DSS should be interactive ease of use.
  9. Effectiveness, but not efficiency.
  10. Complete control by decision-makers.
  11. Ease of development by end users.
  12. Support modeling and analysis.
  13. Data access.
  14. Standalone, integration and Web-based

DSS CHARACTERISTICS

(DSS In Action 1.5: Houston Minerals Case)

- **Initial risk analysis (management science)**
- **Model examination using experience, judgment, and intuition**
- **Initial model mathematically correct, but incomplete**
- **DSS provided very quick analysis**
- **DSS: flexible and responsive. Allows *managerial intuition* and *judgment***

INFORMATION SYSTEMS TO SUPPORT DECISIONS

	Management Information Systems	Decision Support Systems
Decision support provided	Provide information about the performance of the organization	Provide information and techniques to analyze specific problems
Information form and frequency	Periodic, exception, demand, and push reports and responses	Interactive inquiries and responses
Information format	Prespecified, fixed format	Ad hoc, flexible, and adaptable format
Information processing methodology	Information produced by extraction and manipulation of business data	Information produced by analytical modeling of business data

DEFINITIONS

- **DBMS** - System for storing and retrieving data and processing queries
- **Data warehouse** - Consolidated database, usually gathered from multiple primary sources, organized and optimized for reporting and analysis
- **MIS** - System to provide managers with summaries of decision-relevant information
- **Expert system** - computerized system that exhibits expert-like behavior in a given problem domain
- **Decision aid** - automated support to help users conform to some normative ideal of rational decision making
- **DSS** - provide automated support for any or all aspects of the decision making process
- **EIS** (Executive information system) - A kind of DSS specialized to the needs of top executives

MANAGEMENT INFORMATION SYSTEMS

- MIS
- Produces information products that support many of the day-to-day decision-making needs of managers and business professionals
- Prespecified reports, displays and responses
- Support more structured decisions

MIS REPORTING ALTERNATIVES

- Periodic Scheduled Reports
  - Prespecified format on a regular basis
- Exception Reports
  - Reports about exceptional conditions
  - May be produced regularly or when exception occurs
- Demand Reports and Responses
  - Information available when demanded
- Push Reporting
  - Information pushed to manager

INTRODUCTION

- IT tools help process information to create business intelligence according to:
  - OLTP
  - OLAP

INTRODUCTION

- ***Online transaction processing (OLTP)*** – the gathering of input information, processing that information, and updating existing information to reflect the gathered and processed information
  - Databases support OLTP
  - ***Operational database*** – databases that support OLTP

INTRODUCTION

- ***Online analytical processing (OLAP)*** – the manipulation of information to support decision making
  - Databases can support some OLAP
  - Data warehouses only support OLAP, not OLTP
  - Data warehouses are special forms of databases that support decision making

INTRODUCTION

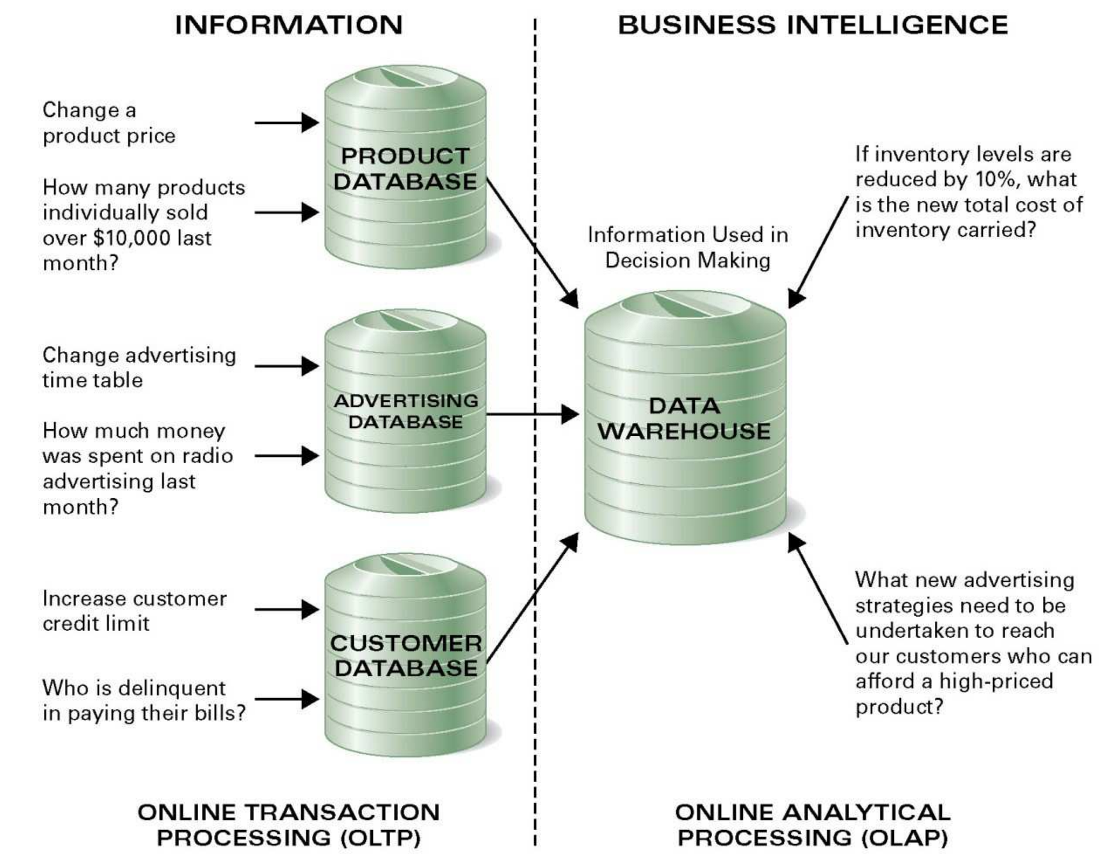

Figure 3.1

Building Business Intelligence

WHAT OLTP BASICALLY IS..?

Sr.No.	Data Warehouse (OLAP)	Operational Database(OLTP)
1	It involves historical processing of information.	It involves day-to-day processing.
2	OLAP systems are used by knowledge workers such as executives, managers, and analysts.	OLTP systems are used by clerks, DBAs, or database professionals.
3	It is used to analyze the business.	It is used to run the business.
4	It focuses on Information out.	It focuses on Data in.
5	It is based on Star Schema, Snowflake Schema, and Fact Constellation Schema.	It is based on Entity Relationship Model.
6	It focuses on Information out.	It is application oriented.
7	It contains historical data.	It contains current data.
8	It provides summarized and consolidated data.	It provides primitive and highly detailed data.
9	It provides summarized and multidimensional view of data.	It provides detailed and flat relational view of data.
10	The number of users is in hundreds.	The number of users is in thousands.

WHAT OLTP BASICALLY IS..?

11	The number of records accessed is in millions.	The number of records accessed is in tens.
12	The database size is from 100GB to 100 TB.	The database size is from 100 MB to 100 GB.
13	These are highly flexible.	It provides high performance.

WHAT OLTP BASICALLY IS..?

In a TPS, End users have online access to the system and to enterprise data, and directly initiate transactions.

Many users repeatedly process similar transactions, and require a fast response to each transaction such as order entry clerks, airline reservation clerks, or bank tellers.

They share an environment of programs and data.

WHAT OLTP BASICALLY IS..?

In a typical TPS:

· Many end users run the same or similar transactions, sharing the same databases and files.

· The system can schedule transactions on the basis of priority attributes.

· The transactions are invoked by online input and generate online output.

· The transactions are designed for a good user interface and fast response times.

WHAT OLTP BASICALLY IS..?

Transaction processing is an effective solution when end users want to:

· Process unscheduled single items in unpredictable volumes and sequence.
· Have immediate access to enterprise data that has been updated to reflect all previous transactions
· Change enterprise data immediately to reflect each transaction as it is processed.

EXECUTIVE INFORMATION SYSTEMS

- EIS
  - Combine many features of MIS and DSS
  - Provide top executives with immediate and easy access to information
  - About the factors that are critical to accomplishing an organization’s strategic objectives (Critical success factors)
  - So popular, expanded to managers, analysts and other knowledge workers

FEATURES OF AN EIS

- Information presented in forms tailored to the preferences of the executives using the system
  - Customizable graphical user interfaces
  - Exception reporting
  - Trend analysis
  - Drill down capability

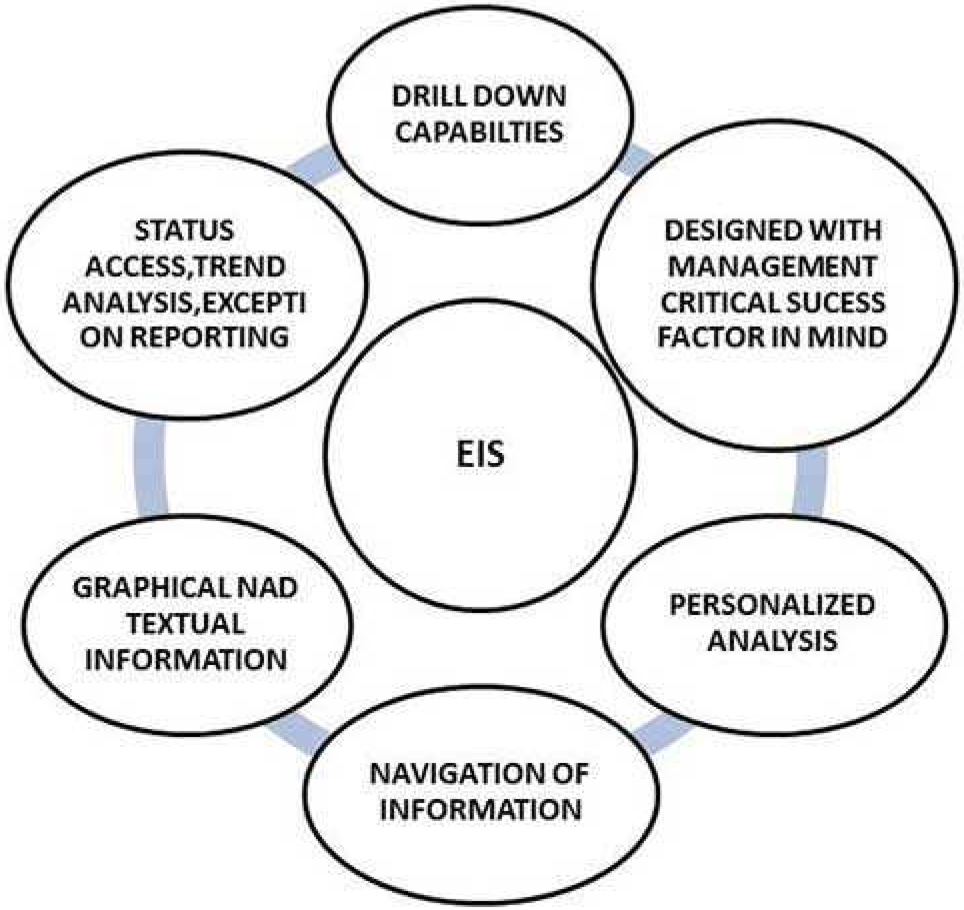

EIS IN DETAIL

- Executive support systems are intended to be used by the senior managers directly to provide support to non-programmed decisions in strategic management.
- These information are often external, unstructured and even uncertain. Exact scope and context of such information is often not known beforehand.
- This information is intelligence based –
- Market intelligence
- Investment intelligence
- Technology intelligence

EXAMPLES OF INTELLIGENT INFORMATION

Following are some examples of intelligent information, which is often the source of an ESS −

- External databases
- Technology reports like patent records etc.
- Technical reports from consultants
- Market reports
- Confidential information about competitors
- Speculative information like market conditions
- Government policies
- Financial reports and information

ADVANTAGES OF EIS

- Easy for upper level executive to use
- Ability to analyze trends
- Augmentation of managers' leadership capabilities
- Enhance personal thinking and decision-making
- Contribution to strategic control flexibility
- Enhance organizational competitiveness in the market place
- Instruments of change
- Increased executive time horizons.
- Better reporting system
- Improved mental model of business executive
- Help improve consensus building and communication
- Improve office automation
- Reduce time for finding information
- Early identification of company performance
- Detail examination of critical success factor
- Better understanding
- Time management
- Increased communication capacity and quality

DISADVANTAGES OF EIS

- Functions are limited
- Hard to quantify benefits
- Executive may encounter information overload
- System may become slow
- Difficult to keep current data
- May lead to less reliable and insecure data
- Excessive cost for small company

BUSINESS INTELLEGENCE SYSTEM

- The term 'Business Intelligence' has evolved from the decision support systems and gained strength with the technology and applications like data warehouses, Executive Information Systems and Online Analytical Processing (OLAP).
- Business Intelligence System is basically a system used for finding patterns from existing data from operations.

CHARACTERISTICS OF BIS

- It is created by procuring data and information for use in decision-making.
- It is a combination of skills, processes, technologies, applications and practices.
- It contains background data along with the reporting tools.
- It is a combination of a set of concepts and methods strengthened by fact-based support systems.
- It is an extension of Executive Support System or Executive Information System.
- It collects, integrates, stores, analyzes, and provides access to business information
- It is an environment in which business users get reliable, secure, consistent, comprehensible, easily manipulated and timely information.
- It provides business insights that lead to better, faster, more relevant decisions.

BENEFITS OF BIS

- Improved Management Processes.
- Planning, controlling, measuring and/or applying changes that results in increased revenues and reduced costs.
- Improved business operations.
- Fraud detection, order processing, purchasing that results in increased revenues and reduced costs.
- Intelligent prediction of future.

CAPABILITIES OF BIS

- Data Storage and Management –
  - Data ware house
  - Ad hoc analysis
  - Data quality
  - Data mining
- Information Delivery
  - Dashboard
  - Collaboration /search
  - Managed reporting
  - Visualization
  - Scorecard
- Query, Reporting and Analysis
  - Ad hoc Analysis
  - Production reporting
  - OLAP analysis

- THANKS

Information Systems

Presented by: Neeraj Kumar Jain
Institute Of Technology & Science
Mohan nagar

Outline

Definitions

Types of Information Systems

Information Systems Vs Information Technology

Expanding Roles of IS

Classification of IS

Enterprise Resource Planning

Information Systems Development

IS as Discipline

Information systems: Opportunities and Challenges

Conclusion

Definitions

Data

Raw facts such as an employee's name and number of hours worked in a week, inventory part numbers or sales orders.

Information

A collection of facts organized in such a way that they have additional value beyond the value of the facts themselves.

Data

$35,000 12 Units
$12,000 J. Jones
Western Region
$100,000 100 Units
35 Units

→ Data Processing →

Information

Salesperson: J. Jones
Sales Territory: Western Region
Current Sales: 147 Units = $147,000

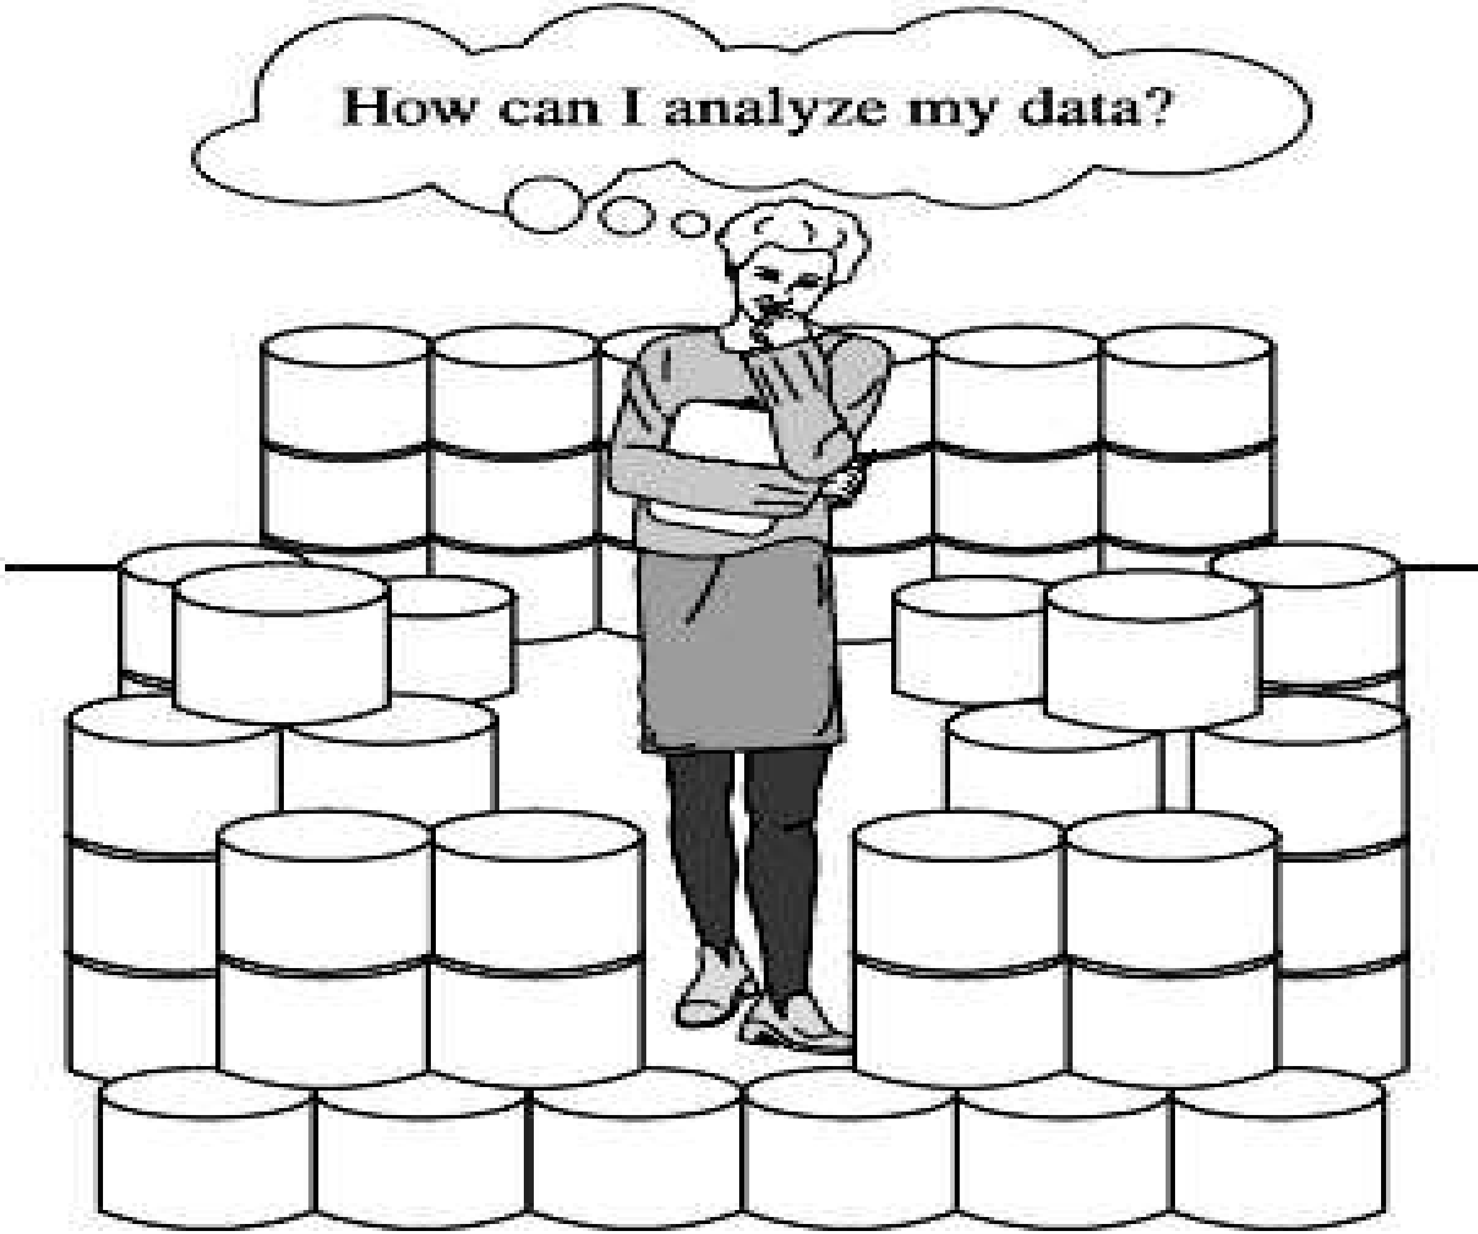

Figure 1.2 We are data rich, but information poor.

Figure 1.3 Data mining—searching for knowledge (interesting patterns) in your data.

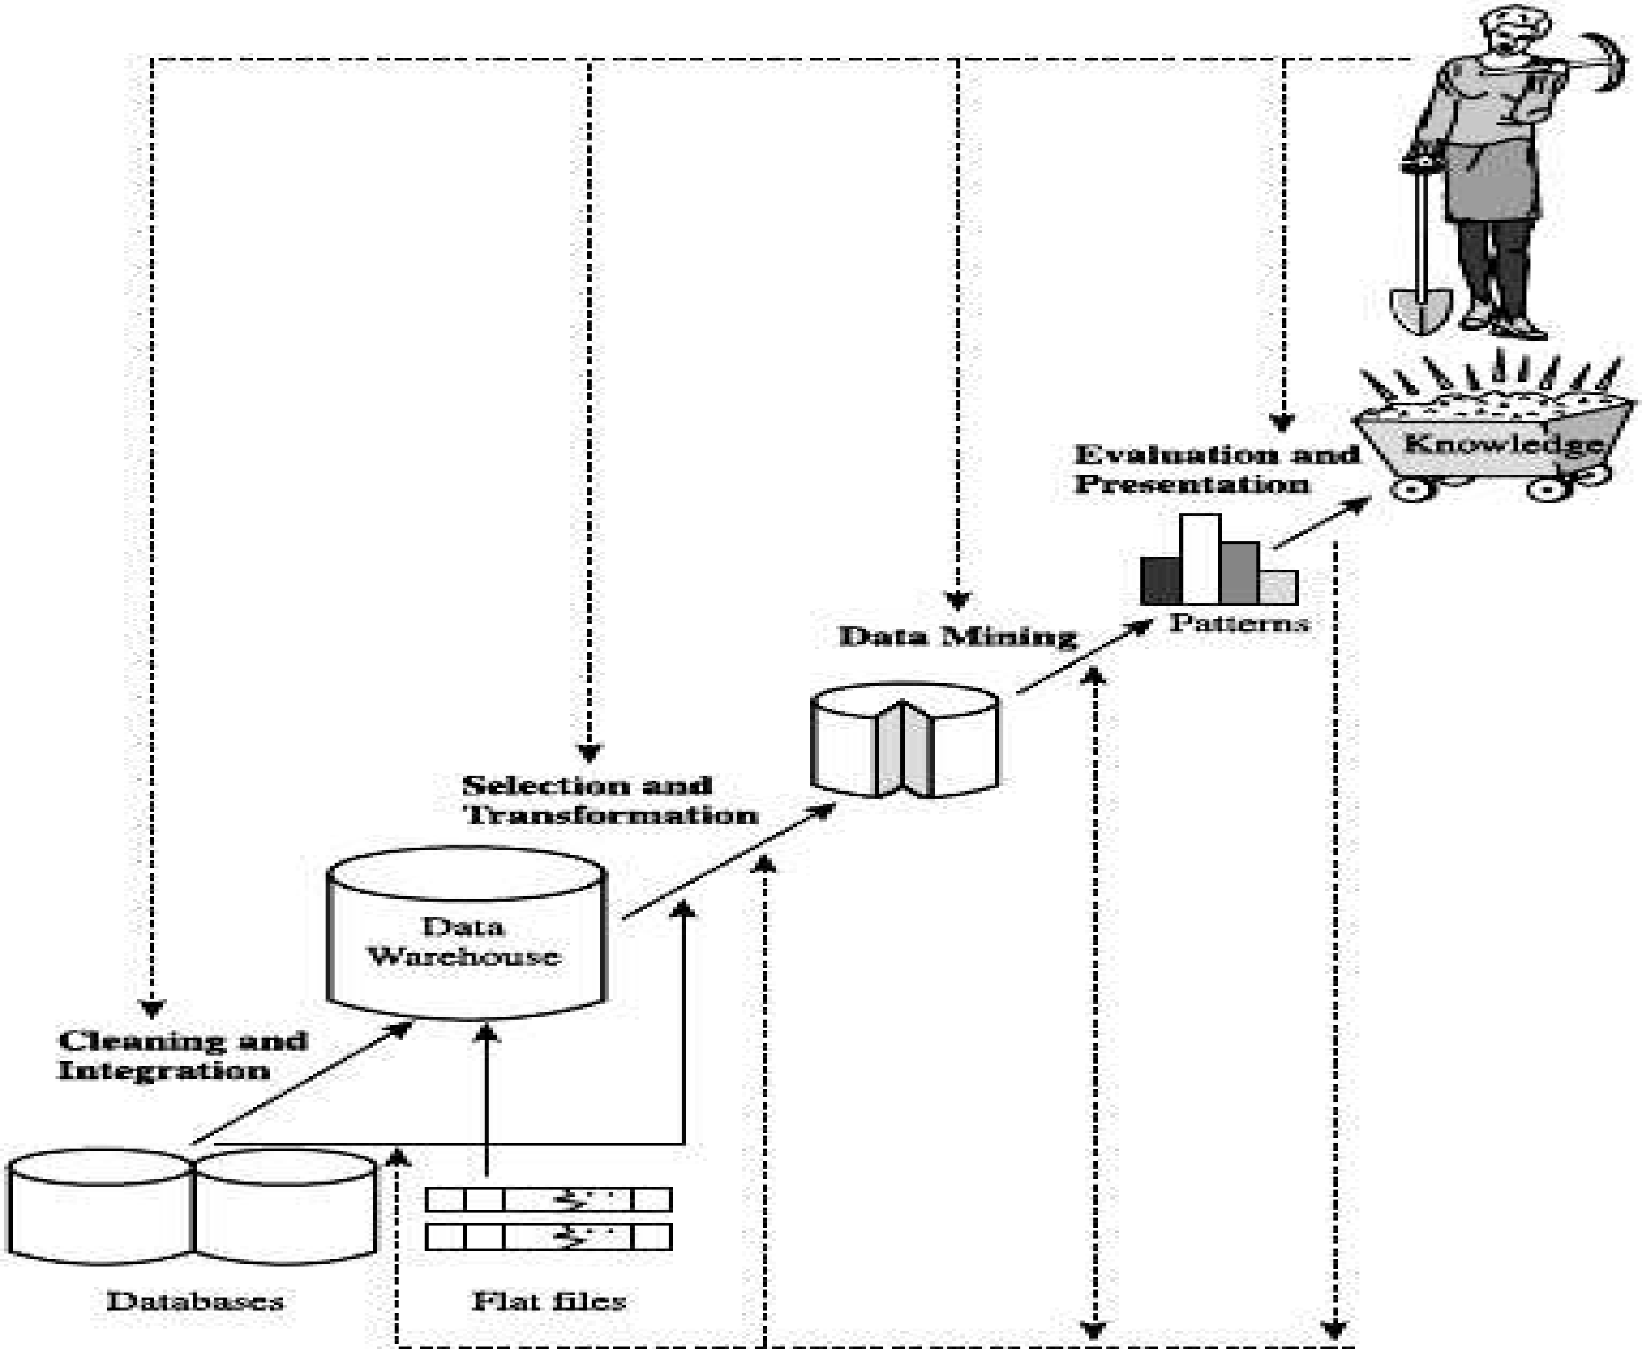

Figure 1.4 Data mining as a step in the process of knowledge discovery.

Definitions

Information Systems

An information system(IS) is typically considered to be a set of interrelated elements or components that collect(input), manipulate(processes), and disseminate (output) data and information and provide a feedback mechanism to meet an objective.

Open System

Close System

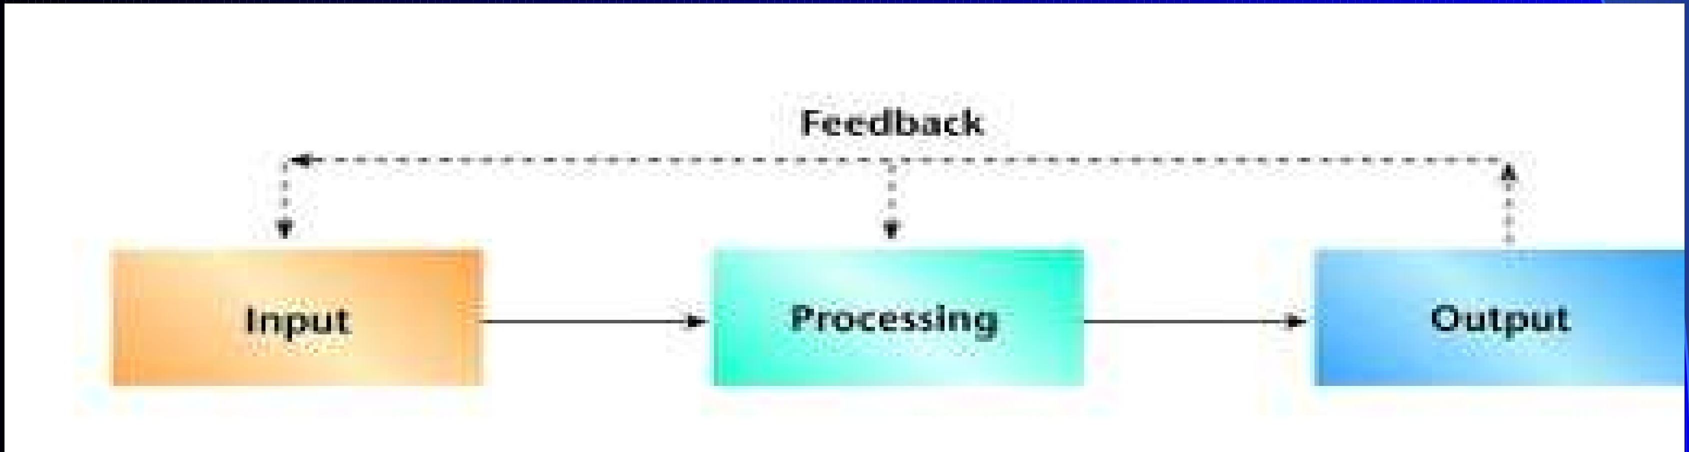

Types of Information Systems

1. Informal Information System
2. Formal Information System

Computer-based Information System

An Information System is an organized combination of people, hardware, software, communication networks and the data resources that collects, transforms and disseminates information in a organization.

IS Vs IT

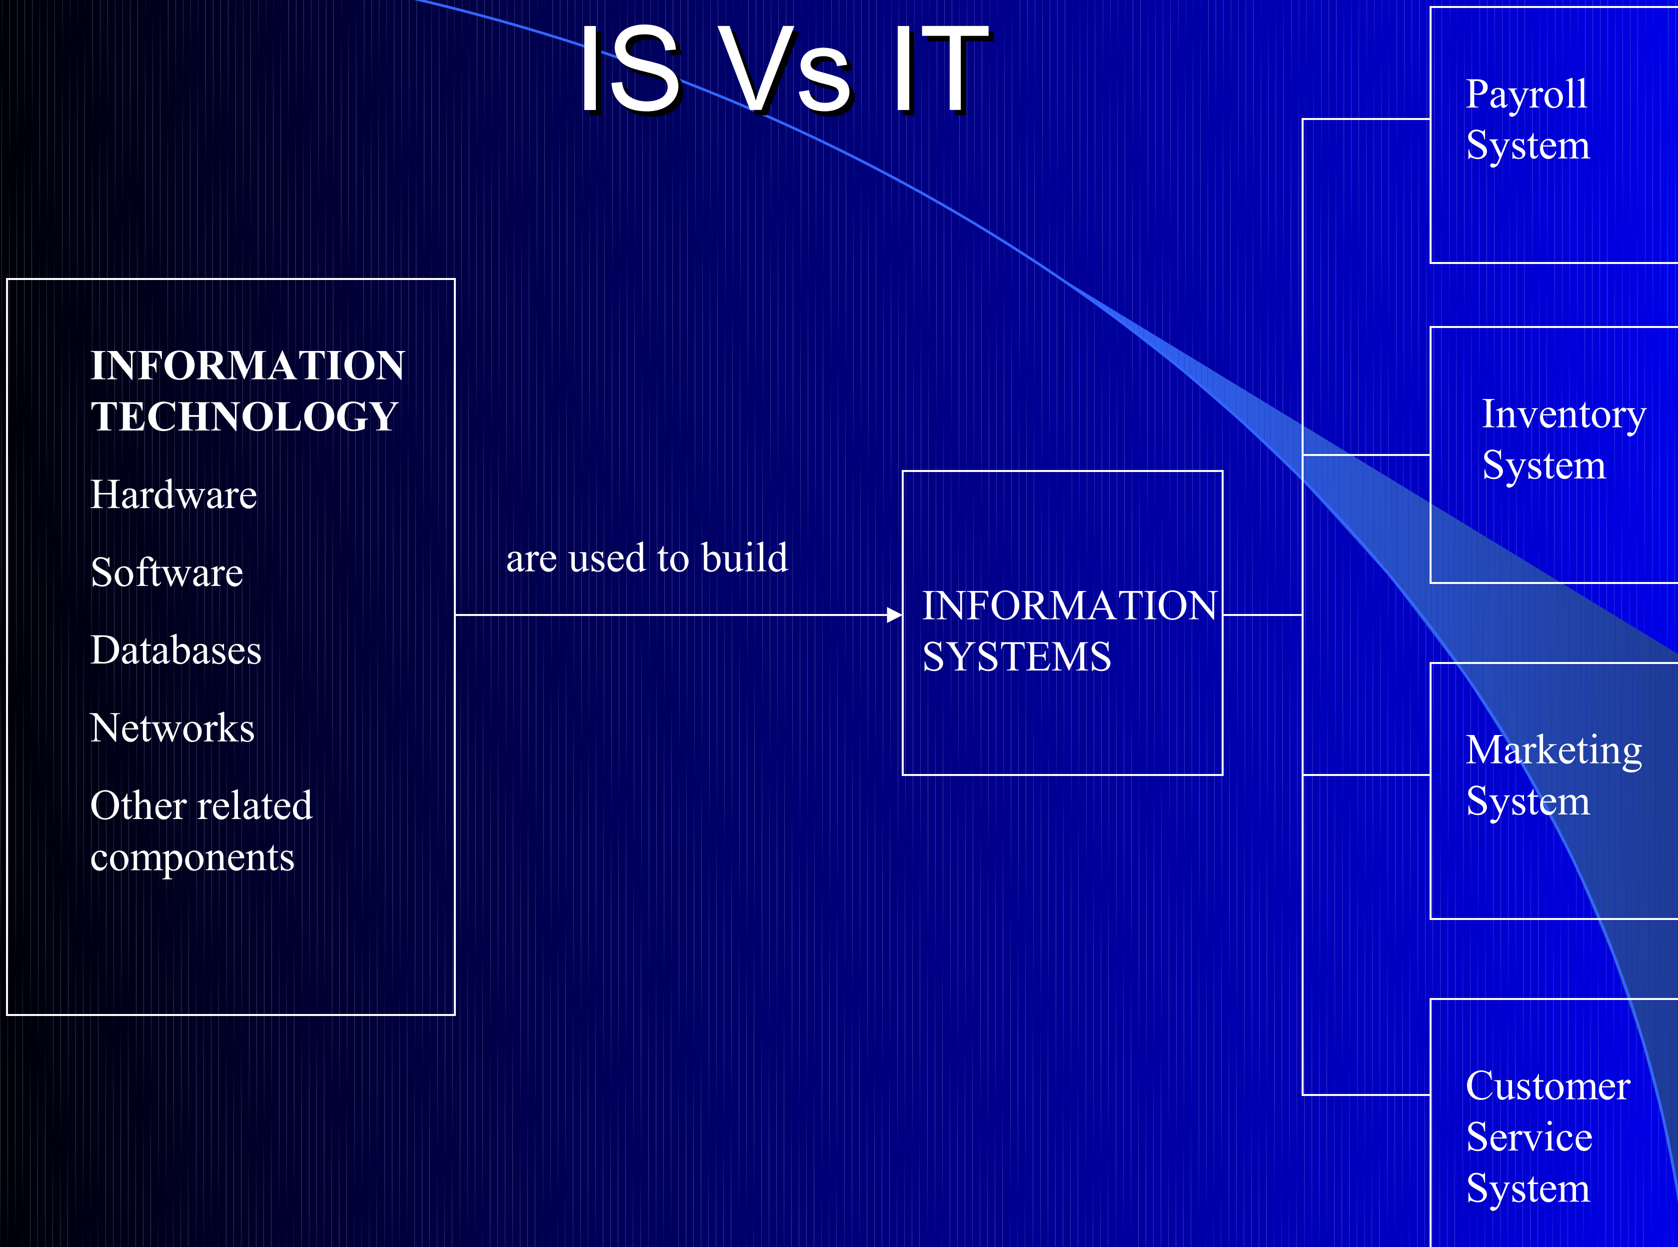
INFORMATION TECHNOLOGY
Hardware
Software
Databases
Networks
Other related components
are used to build
INFORMATION SYSTEMS
Payroll System
Inventory System
Marketing System
Customer Service System

Expanding Roles of IS

1. Data Processing: 1950s-1960s
2. Management Reporting: 1960s-1970s
3. Decision support: 1970s-1980s
4. Strategic and End User Support: 1980s-1990s
5. Global Internetworking: 1990s-2000s

Classification of IS

- Information Systems
  - Operations Support System
    - Transaction processing systems
    - Process control systems
    - Office automation systems
  - Management Support System
    - Management information systems
    - Decision support systems
    - Executive information systems

1. Operations support systems process data generated by business operations

Major categories are:

i) Transaction processing systems

ii) Process control systems

iii) Office automation systems

2. Management Support Systems provide information and support needed for effective decision making by managers

Major categories are

i) Management Information System

ii) Decision Support Systems

iii) Executive Information System

1. **Operations Support System**

i) Transaction processing systems

- Process business exchanges
- Maintain records about the exchanges
- Handle routine, yet critical, tasks
- Perform simple calculations

ii) **Process control systems** monitor and control industrial processes.

iii) **Office automation systems** automate office procedures and enhance office communications and productivity.

2. Management support systems provide information and support needed for effective decision making by managers

Major categories are:

i) **Management information systems**

- Routine information for routine decisions
- Operational efficiency
- Use transaction data as main input
- Databases integrate MIS in different functional areas

ii) Decision Support System

- Interactive support for non-routine decisions or problems
- End-users are more involved in creating a DSS than an MIS

iii) Executive information systems

provide critical information tailored to the information needs of executives

Other categories

a) Expert systems

b) End user computing systems

c) Business information systems

d) Strategic information systems

a) Expert Systems are knowledge-based systems that provides expert advice and act as expert consultants to the users

b) End user computing systems support the direct, hands on use of computers by end users for operational and managerial applications

c) Business information systems support the operational and managerial applications of the basic business functions of a firm

d) Strategic information systems provide a firm which strategic products, services, and capabilities for competitive advantage

Enterprise Resource Planning (ERP)

- Integrated programs that can manage a company’s entire set of business operations
- Often coordinate planning, inventory control, production and ordering

Information Systems Development

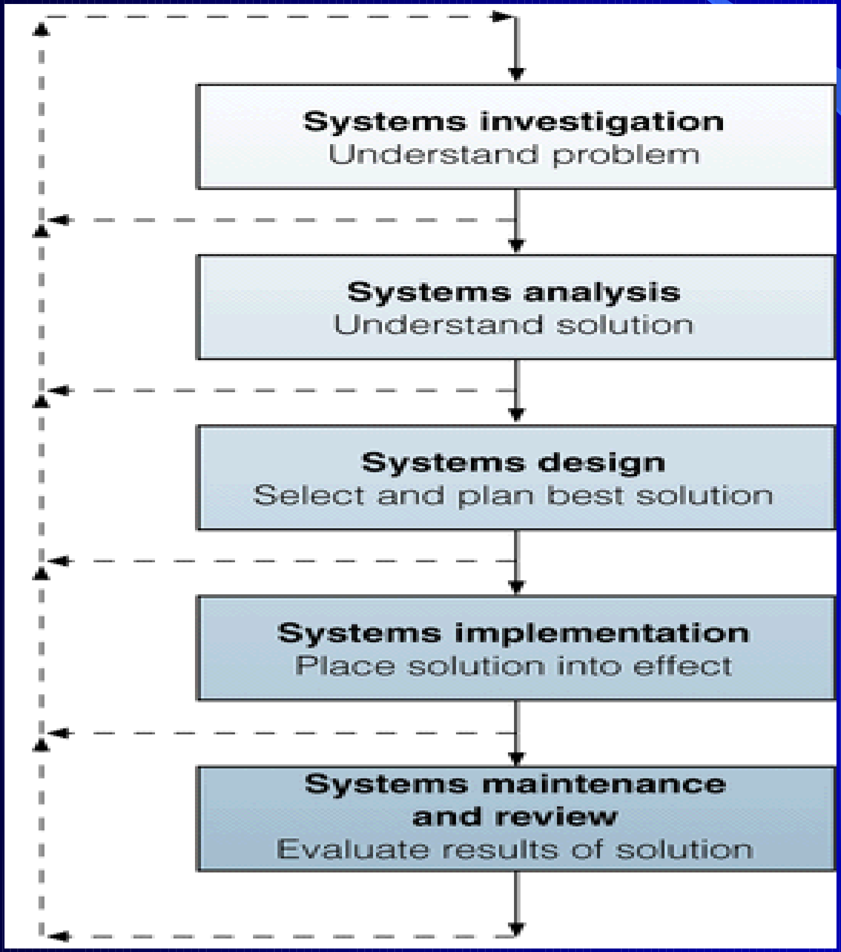

IS as Discipline

IS is an interdisciplinary field influenced by Computer Science, Political Science, Psychology, Operations Research, Linguistics, Sociology, and Organizational Theory.

Disciplines and Principles Used in the IS Field

Discipline	Key Principle
Computer science	Identify and understand the scientific principles of computing. Learn to develop more powerful and interactive computers.
Political science	Examine the political impact of information, both inside and outside the organization, and the effects of information systems on society.
Psychology	Understand and study how human beings reason and apply information to solve problems. Explore how to help human beings use information to become better decision makers.
Operations research	Develop tools, techniques, and methods to gather and process information efficiently and effectively. Understand the scientific and mathematical principles to get the maximum benefits from IS investments.
Linguistics	Understand the role and power of language in communications. Focus on developing more "humanistic" ways to communicate with computers.
Sociology	Identify the principles that govern society as they relate to information. Develop meaningful ethical guidelines for individuals, businesses, and society to work in an information-intensive world.
Organization theory and behavior	Assess how the influence of organizations and their culture influence individuals and society. Develop guidelines to help individuals succeed in different organizational settings.
Ergonomics	Determine ways to improve the interaction between human beings and machines. Explore ways to make computers a natural extension of human beings.

Challenges

1. Workforce downsizing
2. Information overload
3. Employee mistrust
4. Difficult to built
5. Security breaches

Opportunities

1. Enhanced global competitiveness
2. Capture market opportunities
3. Support corporate strategy
4. Enhance worker productivity
5. Improve quality of goods and services

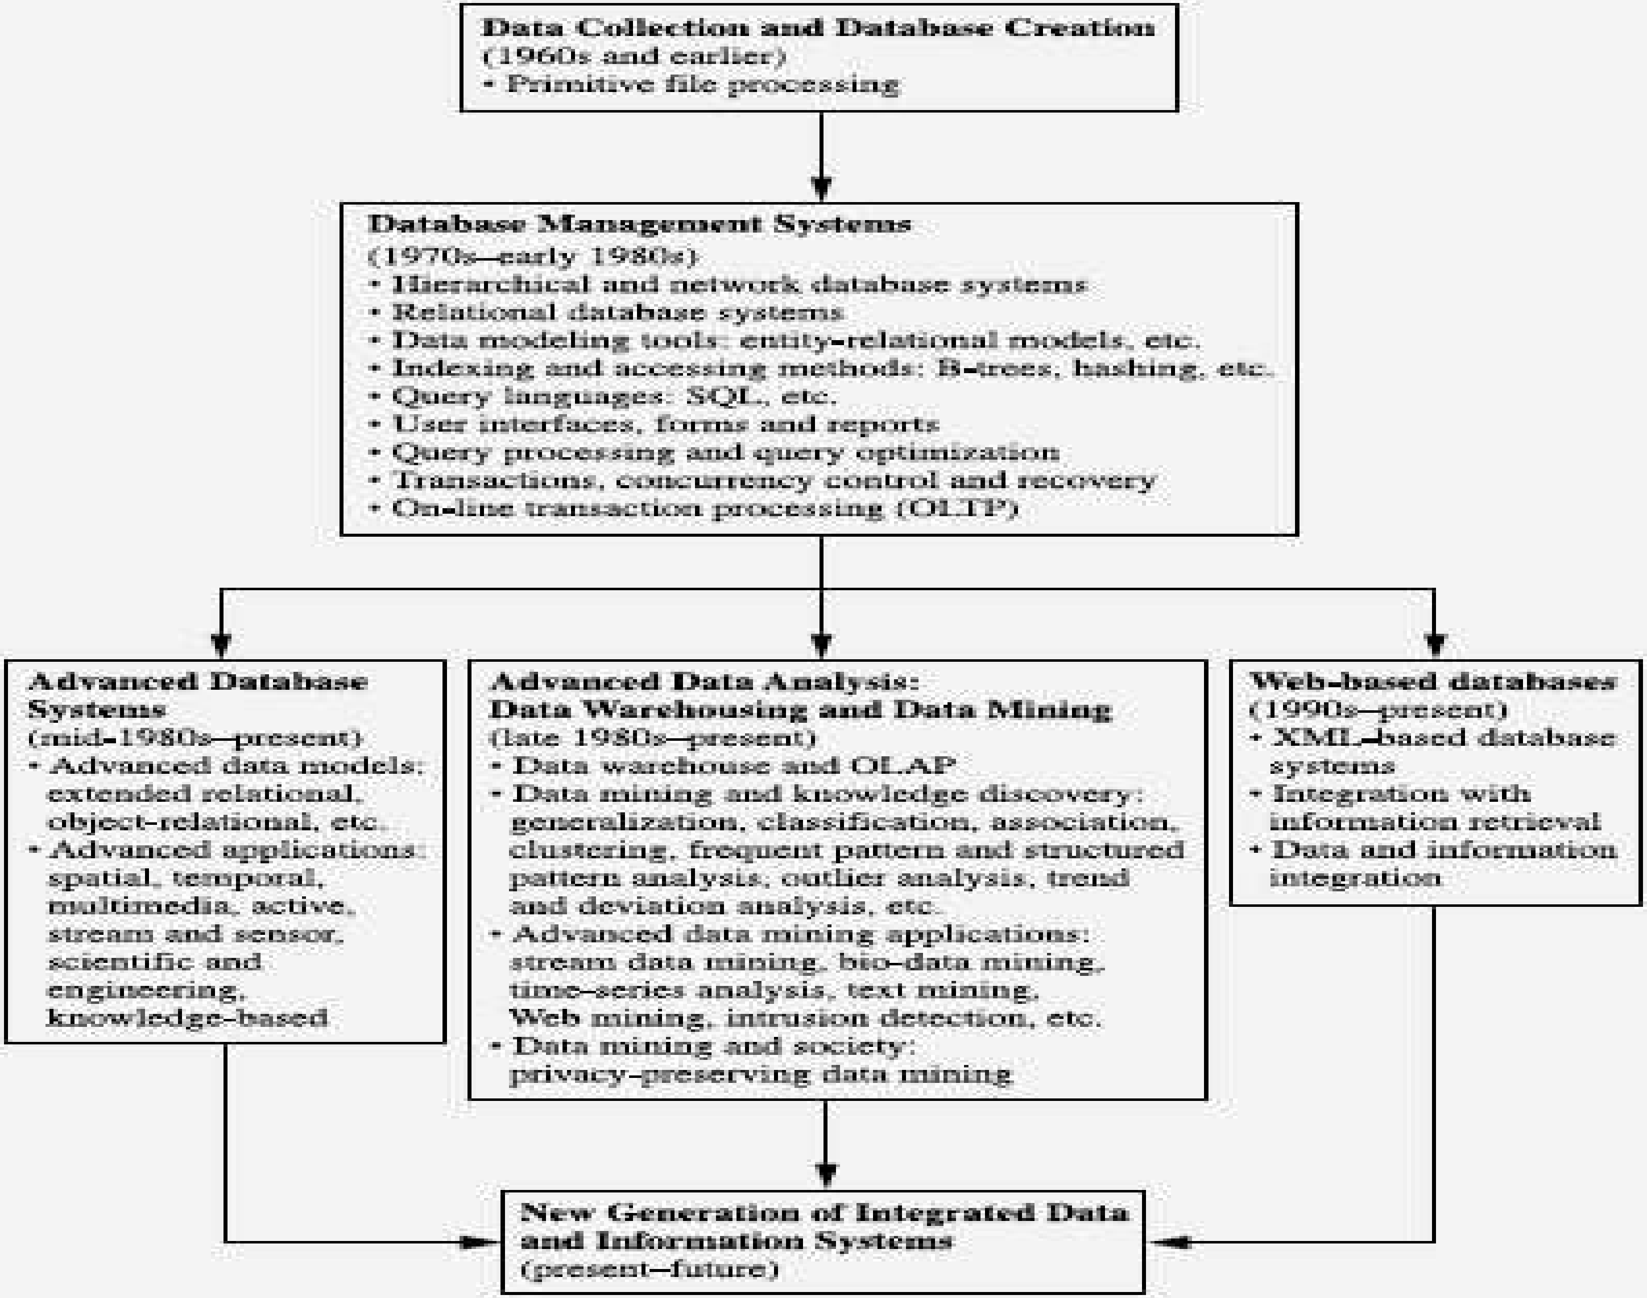

Figure 1.1 The evolution of database system technology.

Conclusion

Information Systems are indispensable to the business, industry, academia and any organization to meet the future challenges

Designing the data warehouse / data marts

Methodologies and Techniques

Data Mart

- A subset of a data warehouse that supports the requirements of a particular department or business function.

- Characteristics include:
  - Do not normally contain detailed operational data unlike data warehouses.
  - May contain certain levels of aggregation

Dependent Data Mart

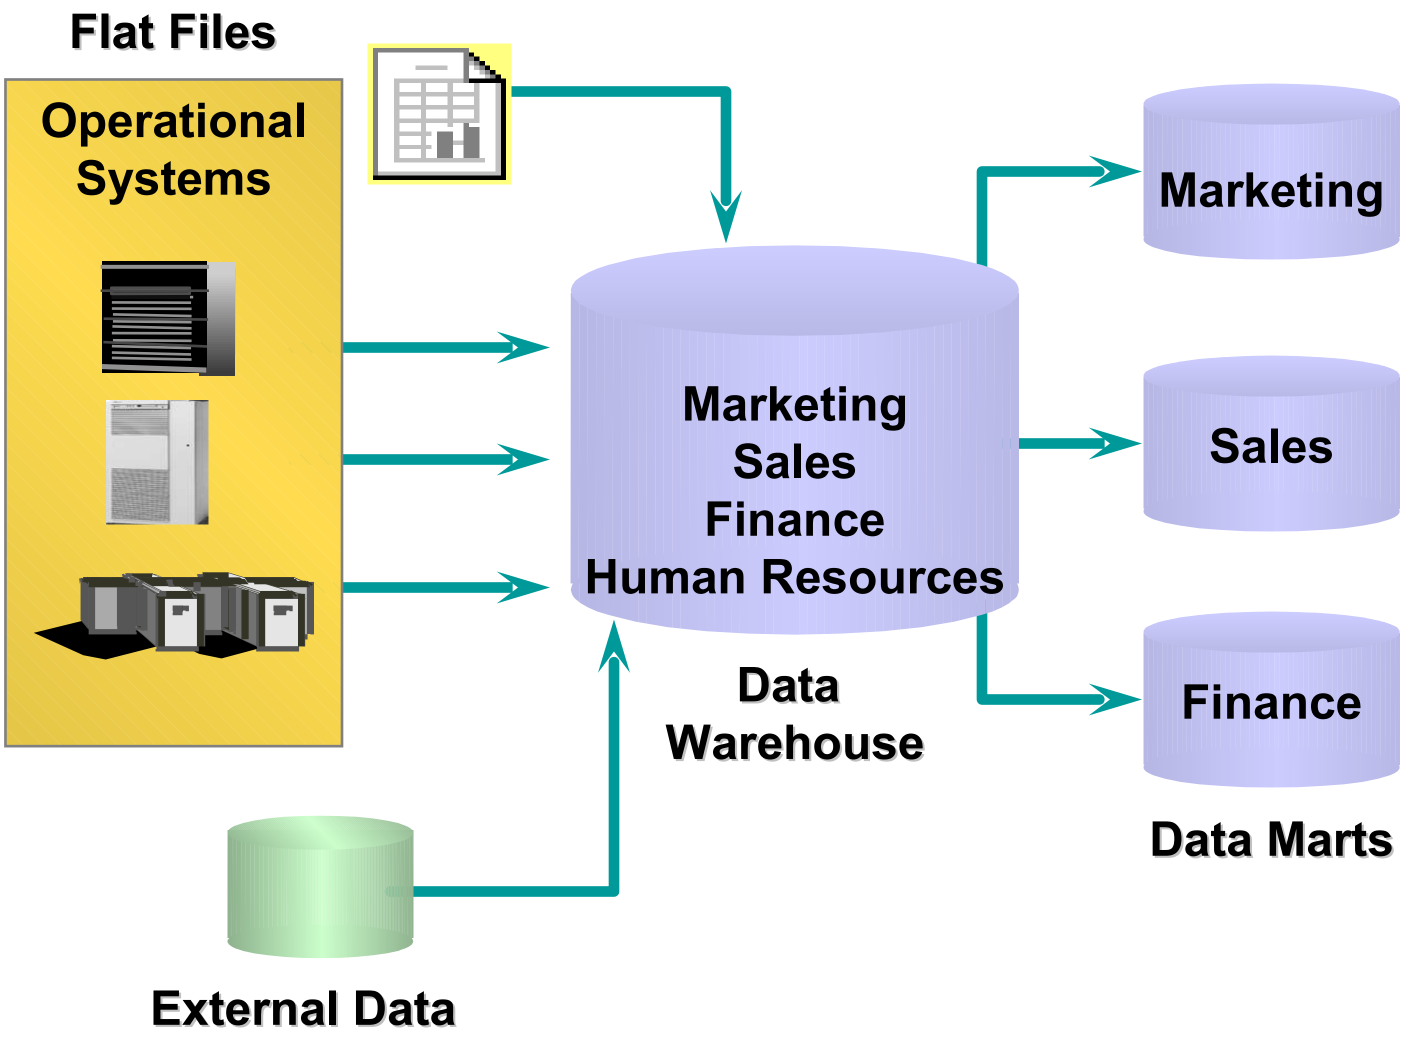

Independent Data Mart

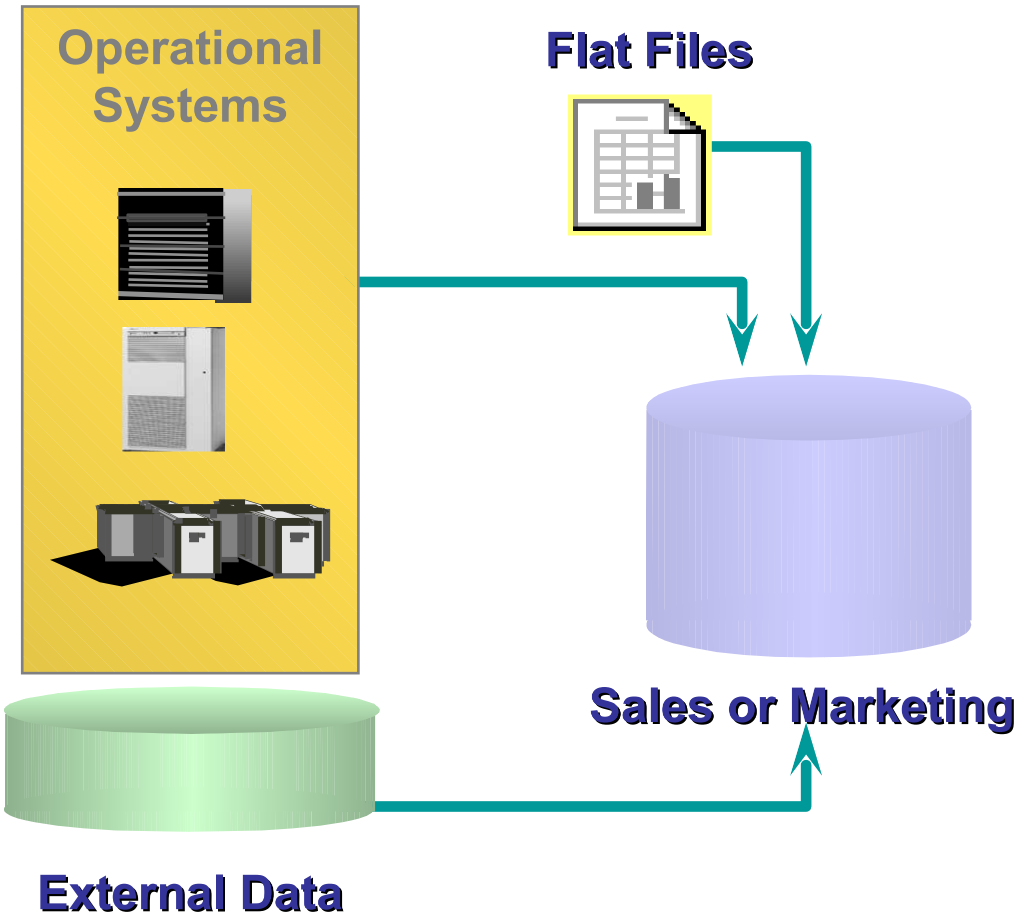

Reasons for Creating a Data Mart

- To give users more flexible access to the data they need to analyse most often.
- To provide data in a form that matches the <u>collective view</u> of a group of users
- To improve end-user response time.
- Potential users of a data mart are clearly defined and can be targeted for support

Reasons for Creating a Data Mart

- To provide appropriately structured data as dictated by the requirements of the end-user access tools.
- Building a data mart is simpler compared with establishing a corporate data warehouse.
- The cost of implementing data marts is far less than that required to establish a data warehouse.

Why Do We Need a Data Mart?

- Listed below are the reasons to create a data mart:
- To partition data in order to impose **access control strategies.**
- To speed up the queries by reducing the volume of data to be scanned.
- To segment data into different hardware platforms.
- To structure data in a form suitable for a user access tool.

Data Marts Issues

- **Data mart functionality**
- **Data mart size**
- **Data mart load performance**
- **Users access to data in multiple data marts**
- **Data mart Internet / Intranet access**
- **Data mart administration**
- **Data mart installation**

DATA WAREHOUSING AND DATA MINING

Neeraj Kr Jain

For BCA-VI

UNIT-II

Course Overview

The course:
what and how

0. Introduction

I. Data Warehousing

II. Decision Support and OLAP

III. Data Mining

IV. Looking Ahead

Demos and Labs

0. Introduction

Data Warehousing, OLAP and data mining: what and why (now)?

Relation to OLTP

A case study

demos, labs

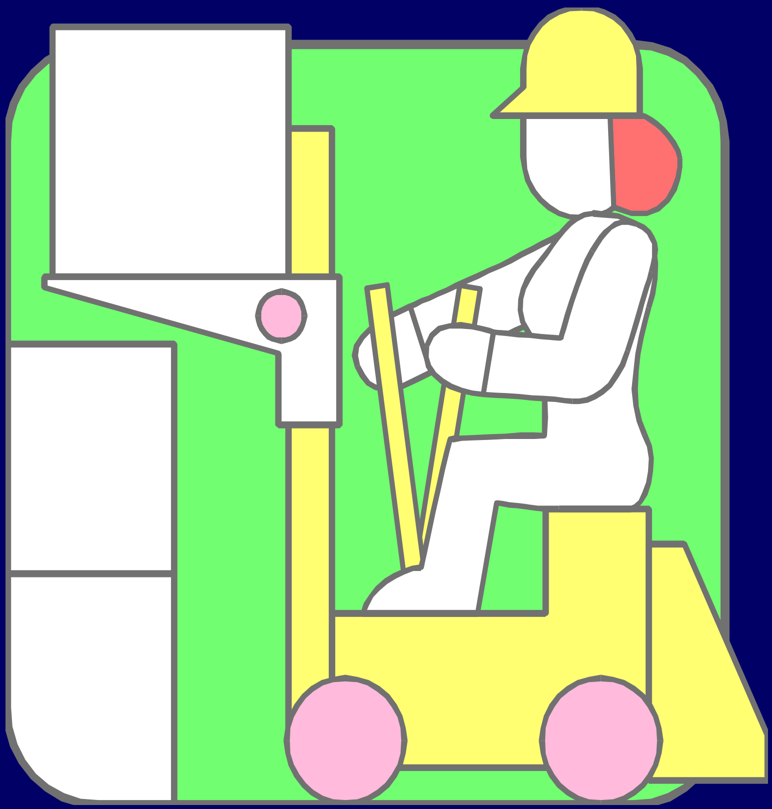

A producer wants to know….

Which are our lowest/highest margin customers ?

What is the most effective distribution channel?

Who are my customers and what products are they buying?

What product prom- -otions have the biggest impact on revenue?

Which customers are most likely to go to the competition ?

What impact will new products/services have on revenue and margins?

Data, Data everywhere yet ...

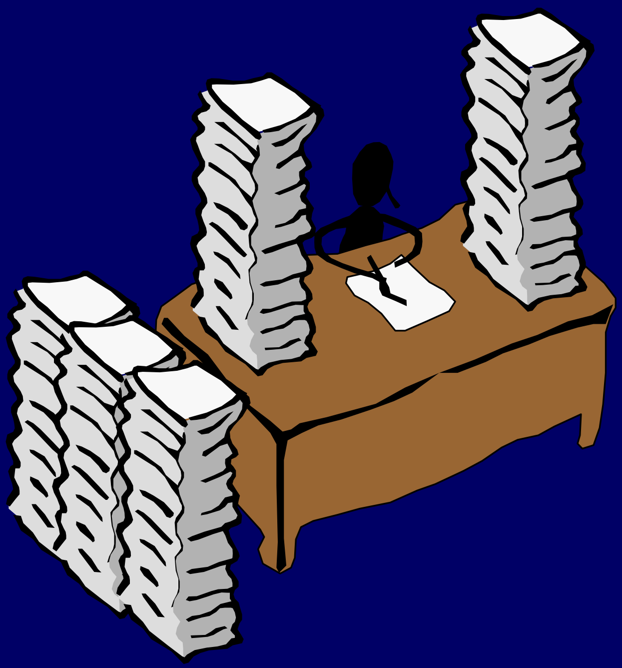

I can't find the data I need

- data is scattered over the network
- many versions, subtle differences

I can't get the data I need

- need an expert to get the data

I can't understand the data I found

- available data poorly documented

I can't use the data I found

- results are unexpected
- data needs to be transformed from one form to other

What is a Data Warehouse?

A single, complete and consistent store of data obtained from a variety of different sources made available to end users in a what they can understand and use in a business context.

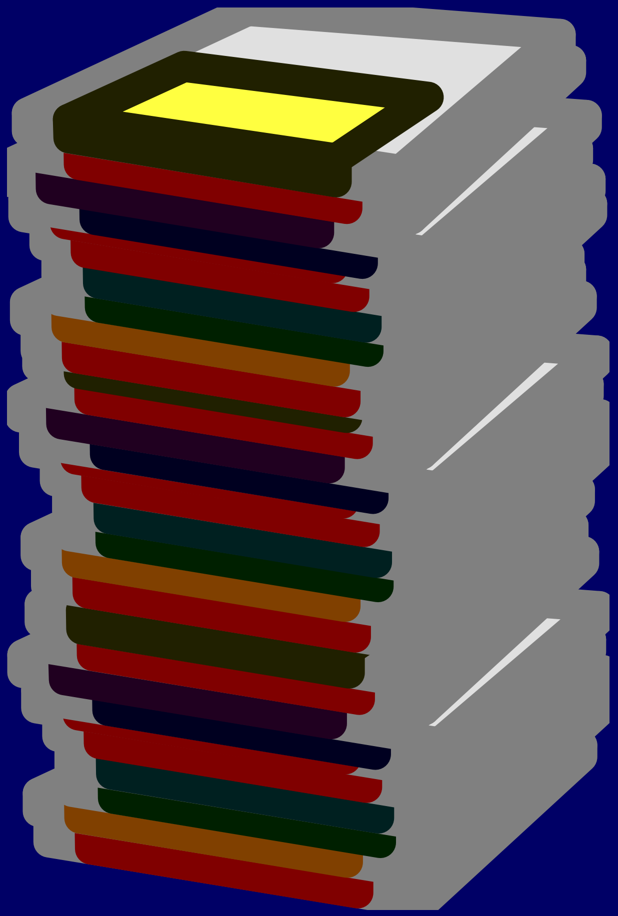

[Barry Devlin]

What are the users saying...

- Data should be integrated across the enterprise
- Summary data has a real value to the organization
- Historical data holds the key to understanding data over time
- What-if capabilities are required

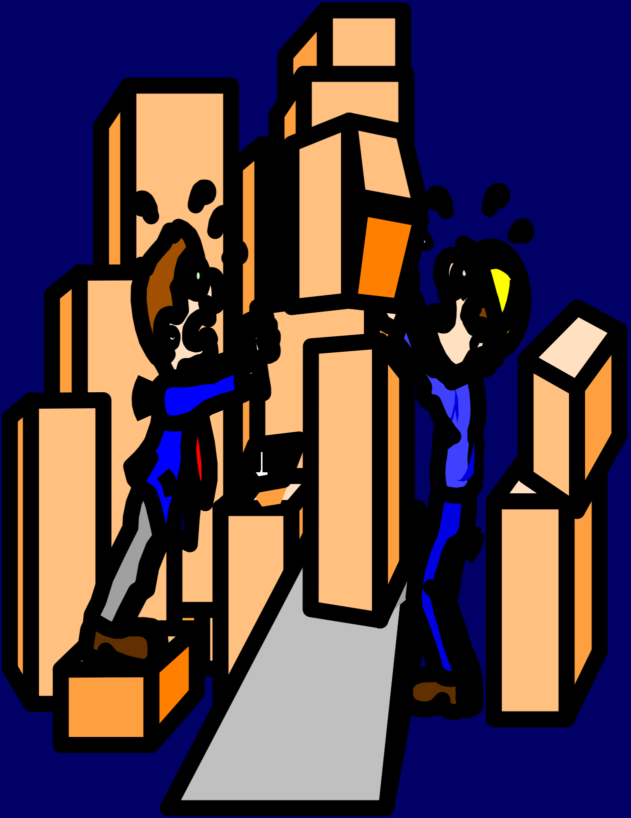

What is Data Warehousing?

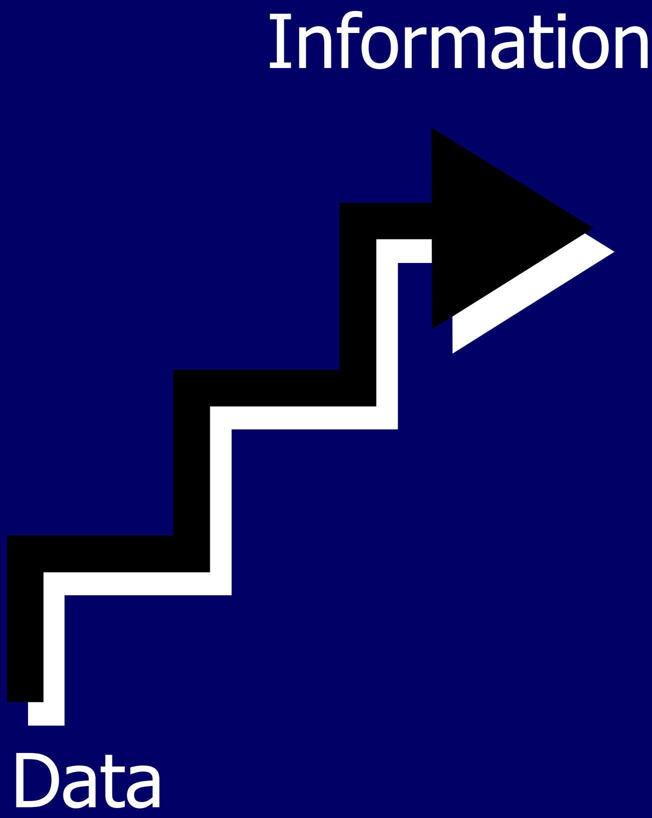

A process of transforming data into information and making it available to users in a timely enough manner to make a difference

[Forrester Research, April 1996]

Evolution

- 60’s: Batch reports
  - hard to find and analyze information
  - inflexible and expensive, reprogram every new request
- 70’s: Terminal-based DSS and EIS (executive information systems)
  - still inflexible, not integrated with desktop tools
- 80’s: Desktop data access and analysis tools
  - query tools, spreadsheets, GUIs
  - easier to use, but only access operational databases
- 90’s: Data warehousing with integrated OLAP engines and tools

Warehouses are Very Large Databases

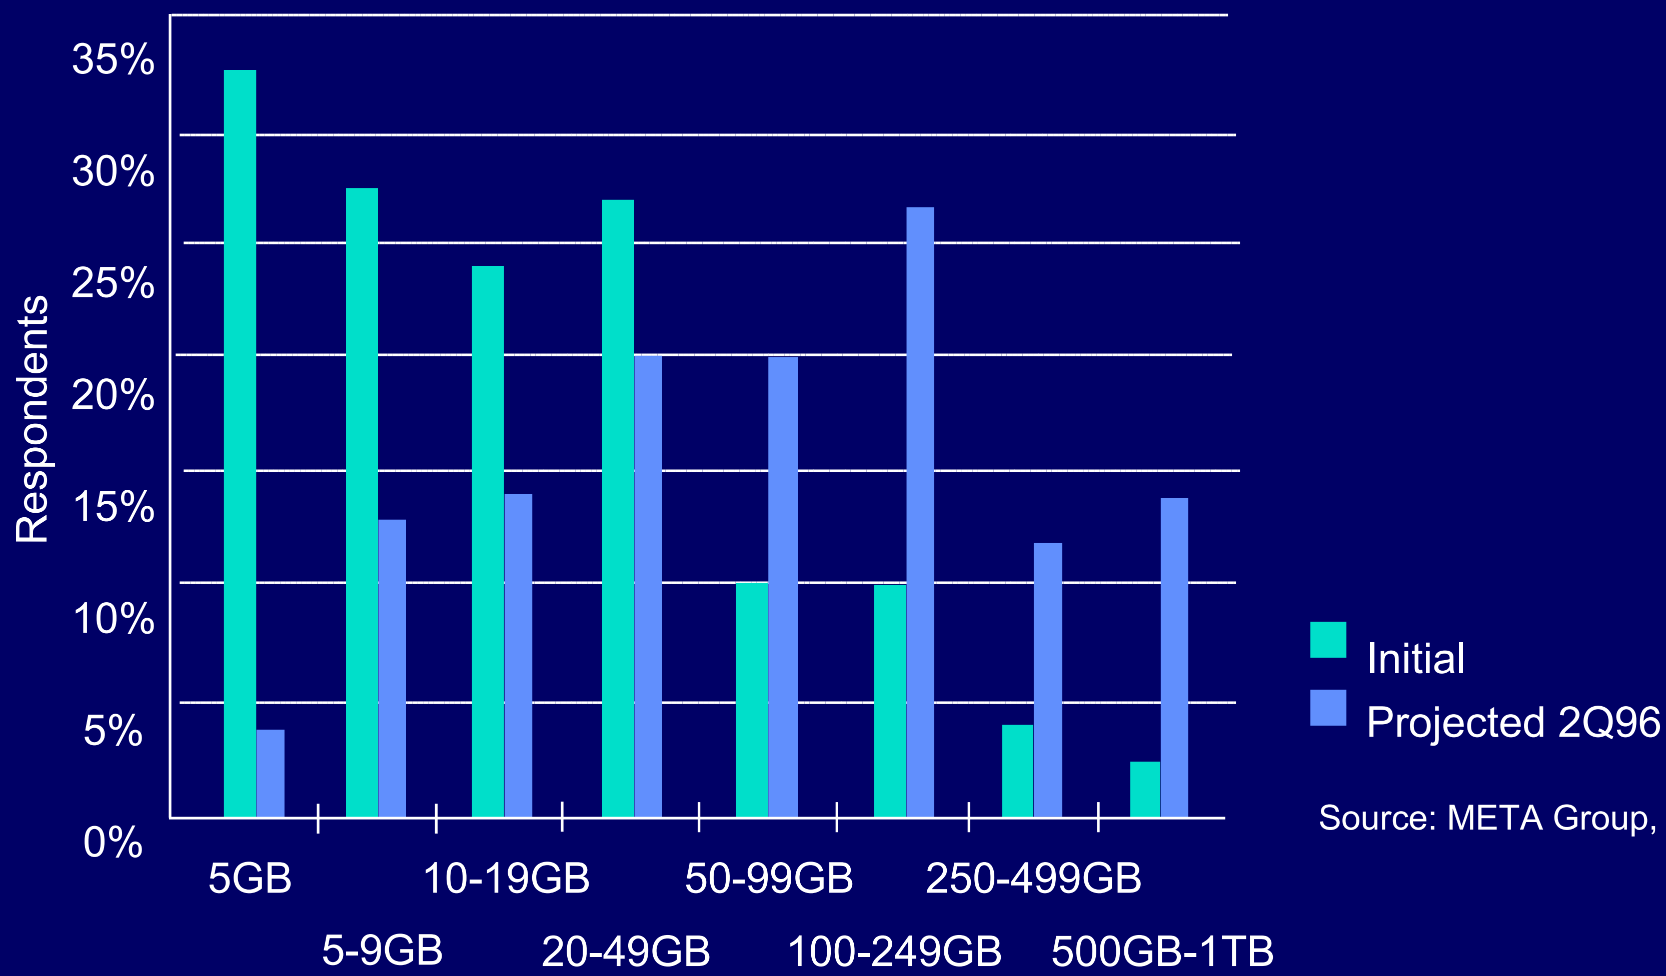

Source: META Group, Inc.

Very Large Data Bases

Terabytes -- 10^12 bytes: Walmart -- 24 Terabytes

Petabytes -- 10^15 bytes: Geographic Information Systems

Exabytes -- 10^18 bytes: National Medical Records

Zettabytes -- 10^21 bytes: Weather images

Zottabytes -- 10^24 bytes: Intelligence Agency Videos

Data Warehousing -- It is a process

- Technique for assembling and managing data from various sources for the purpose of answering business questions. Thus making decisions that were not previous possible
- A decision support database maintained separately from the organization's operational database

Data Warehouse

A data warehouse is a

subject-oriented

integrated

time-varying

non-volatile

collection of data that is used primarily in organizational decision making.

-- Bill Inmon, Building the Data Warehouse 1996

Data Warehouse Architecture

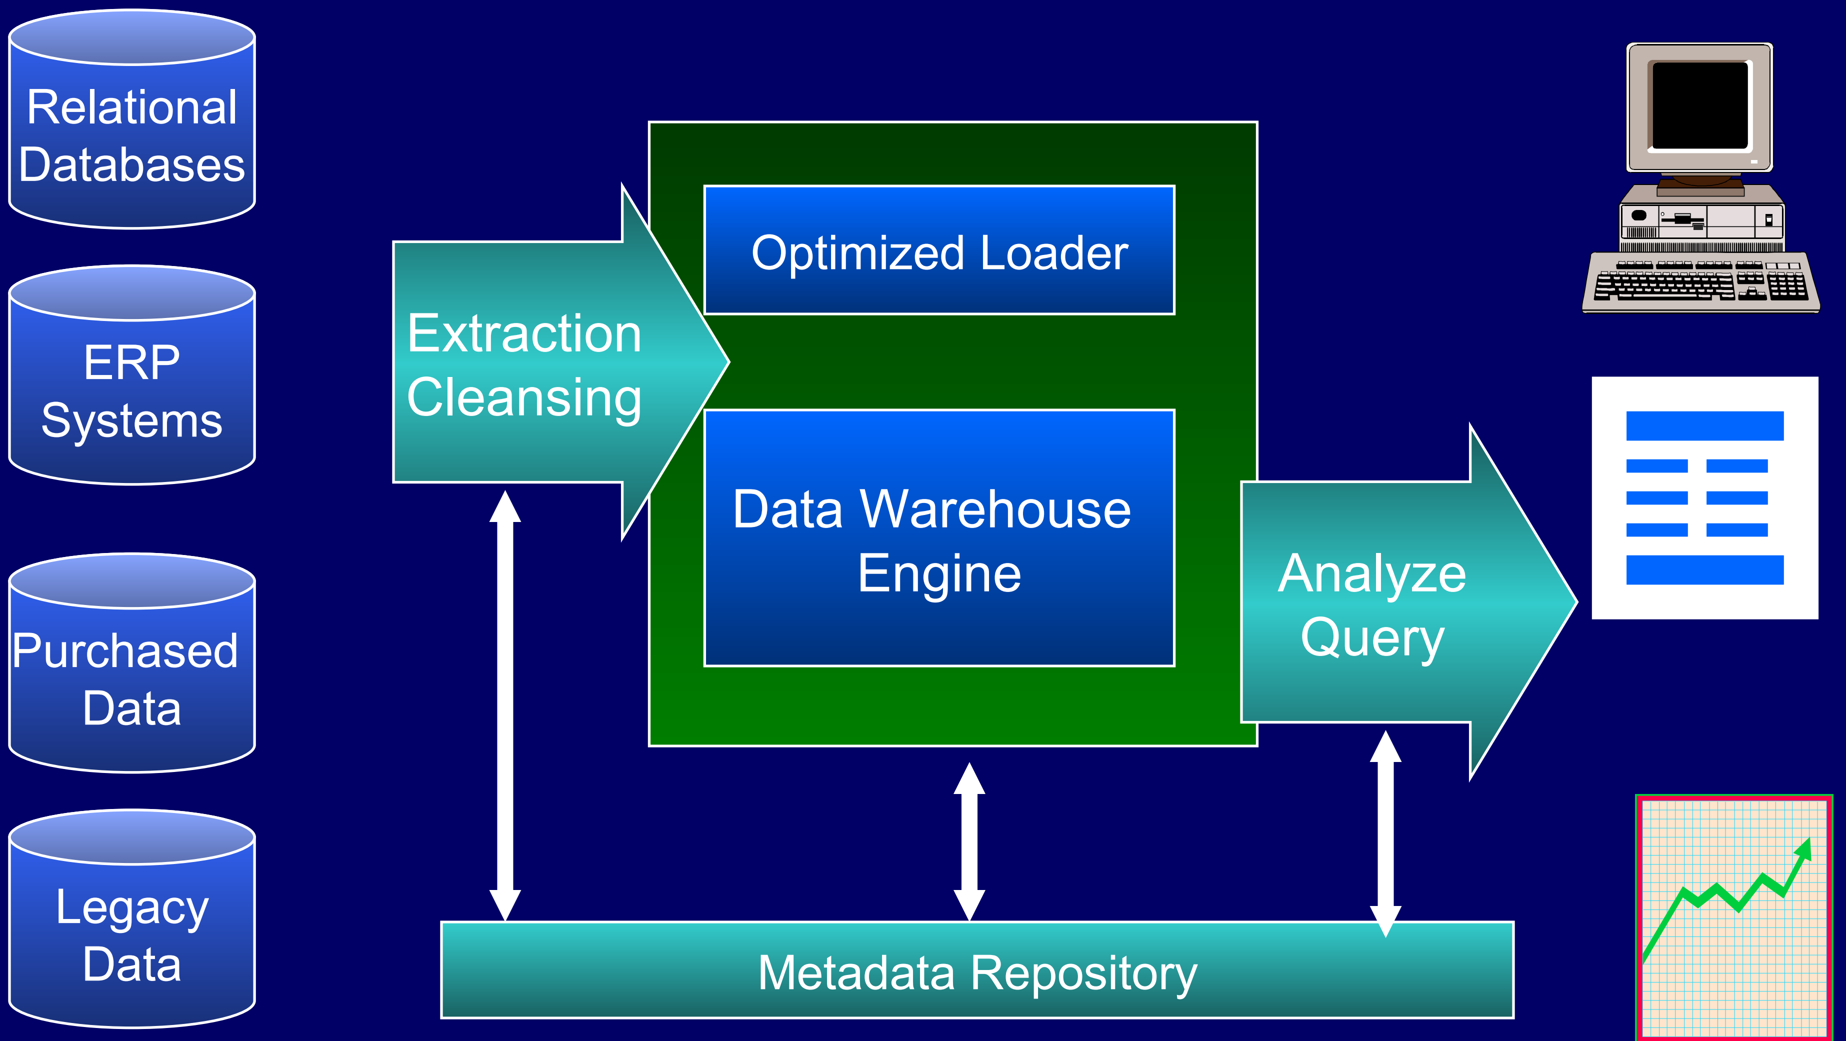

Data Warehouse for Decision Support & OLAP

Putting Information technology to help the knowledge worker make faster and better decisions

- Which of my customers are most likely to go to the competition?
- What product promotions have the biggest impact on revenue?
- How did the share price of software companies correlate with profits over last 10 years?

Decision Support

- Used to manage and control business
- Data is historical or point-in-time
- Optimized for inquiry rather than update
- Use of the system is loosely defined and can be ad-hoc
- Used by managers and end-users to understand the business and make judgements

Data Mining works with Warehouse Data

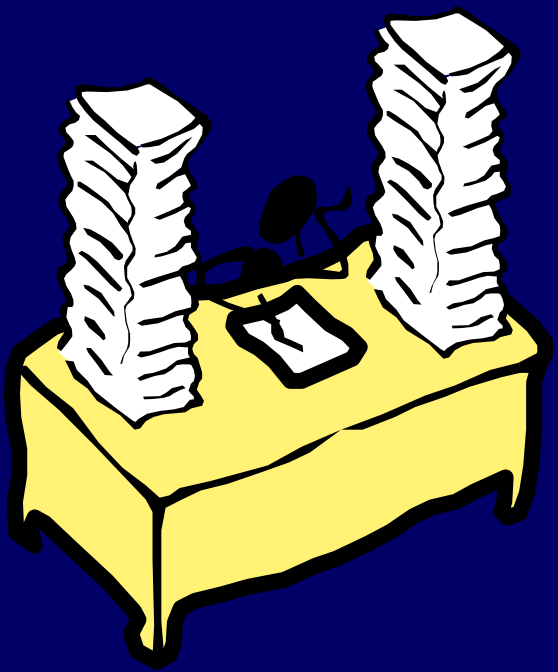

Data Warehousing provides the Enterprise with a memory

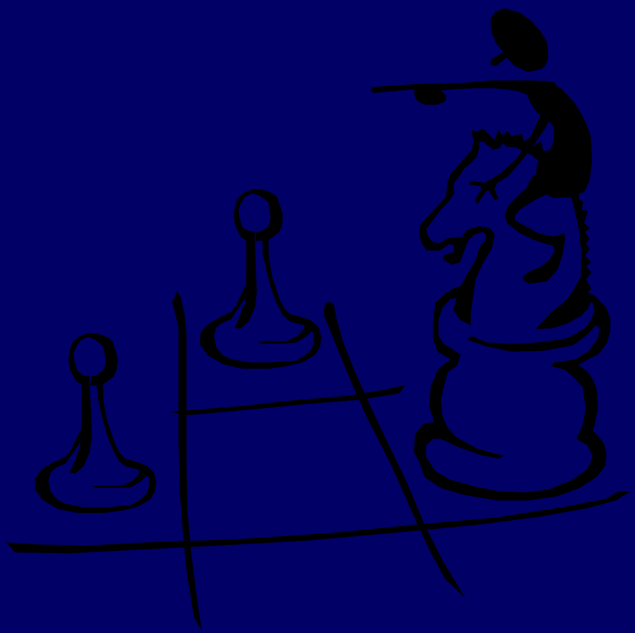

Data Mining provides the Enterprise with intelligence

We want to know ...

Given a database of 100,000 names, which persons are the least likely to default on their credit cards?

Which types of transactions are likely to be fraudulent given the demographics and transactional history of a particular customer?

If I raise the price of my product by Rs. 2, what is the effect on my ROI?

If I offer only 2,500 airline miles as an incentive to purchase rather than 5,000, how many lost responses will result?

If I emphasize ease-of-use of the product as opposed to its technical capabilities, what will be the net effect on my revenues?

Which of my customers are likely to be the most loyal?

Data Mining helps extract such information

Application Areas

Industry	**Application**
Finance	Credit Card Analysis
Insurance	Claims, Fraud Analysis
Telecommunication	Call record analysis
Transport	Logistics management
Consumer goods	promotion analysis
Data Service providers	Value added data
Utilities	Power usage analysis

Data Mining in Use

- The US Government uses Data Mining to track fraud
- A Supermarket becomes an information broker
- Basketball teams use it to track game strategy
- Cross Selling
- Warranty Claims Routing
- Holding on to Good Customers
- Weeding out Bad Customers

What makes data mining possible?

Advances in the following areas are making data mining deployable:

- data warehousing
- better and more data (i.e., operational, behavioral, and demographic)
- the emergence of easily deployed data mining tools and
- the advent of new data mining techniques.
  - -- Gartner Group

Why Separate Data Warehouse?

Performance

Op dbs designed & tuned for known txs & workloads.

Complex OLAP queries would degrade perf. for op txs.

Special data organization, access & implementation methods needed for multidimensional views & queries.

Function

Missing data: Decision support requires historical data, which op dbs do not typically maintain.

Data consolidation: Decision support requires consolidation (aggregation, summarization) of data from many heterogeneous sources: op dbs, external sources.

Data quality: Different sources typically use inconsistent data representations, codes, and formats which have to be reconciled.

What are Operational Systems?

- They are OLTP systems
- Run mission critical applications
- Need to work with stringent performance requirements for routine tasks
- Used to run a business!

RDBMS used for OLTP

- Database Systems have been used traditionally for OLTP
  - clerical data processing tasks
  - detailed, up to date data
  - structured repetitive tasks
  - read/update a few records
  - isolation, recovery and integrity are critical

Operational Systems

- Run the business in real time
- Based on up-to-the-second data
- Optimized to handle large numbers of simple read/write transactions
- Optimized for fast response to predefined transactions
- Used by people who deal with customers, products -- clerks, salespeople etc.
- They are increasingly used by customers

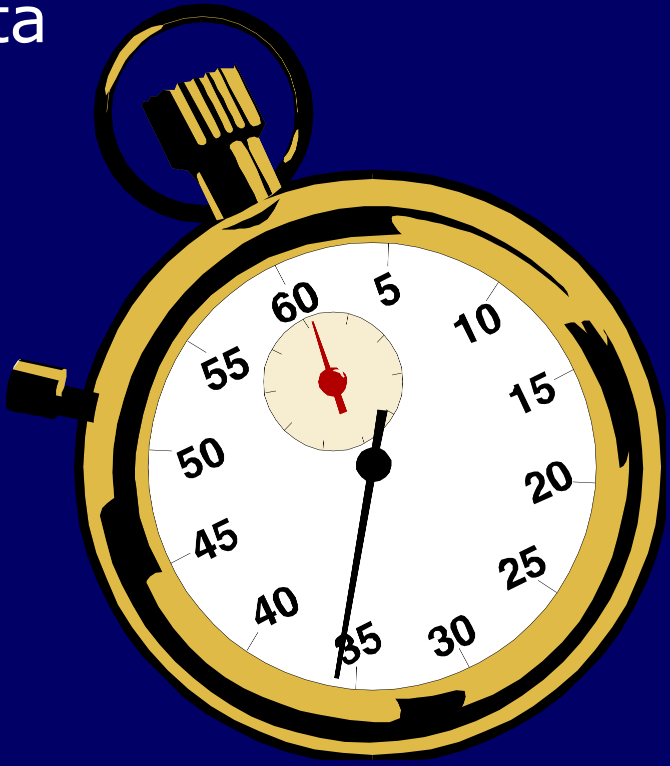

Examples of Operational Data

Data	Industry	Usage	Technology	Volumes
Customer File	All	Track Customer Details	Legacy application, flat files, main frames	Small-medium
Account Balance	Finance	Control account activities	Legacy applications, hierarchical databases, mainframe	Large
Point-of-Sale data	Retail	Generate bills, manage stock	ERP, Client/Server, relational databases	Very Large
Call Record	Telecomm-unications	Billing	Legacy application, hierarchical database, mainframe	Very Large
Production Record	Manufact-uring	Control Production	ERP, relational databases, AS/400	Medium

So, what's different?

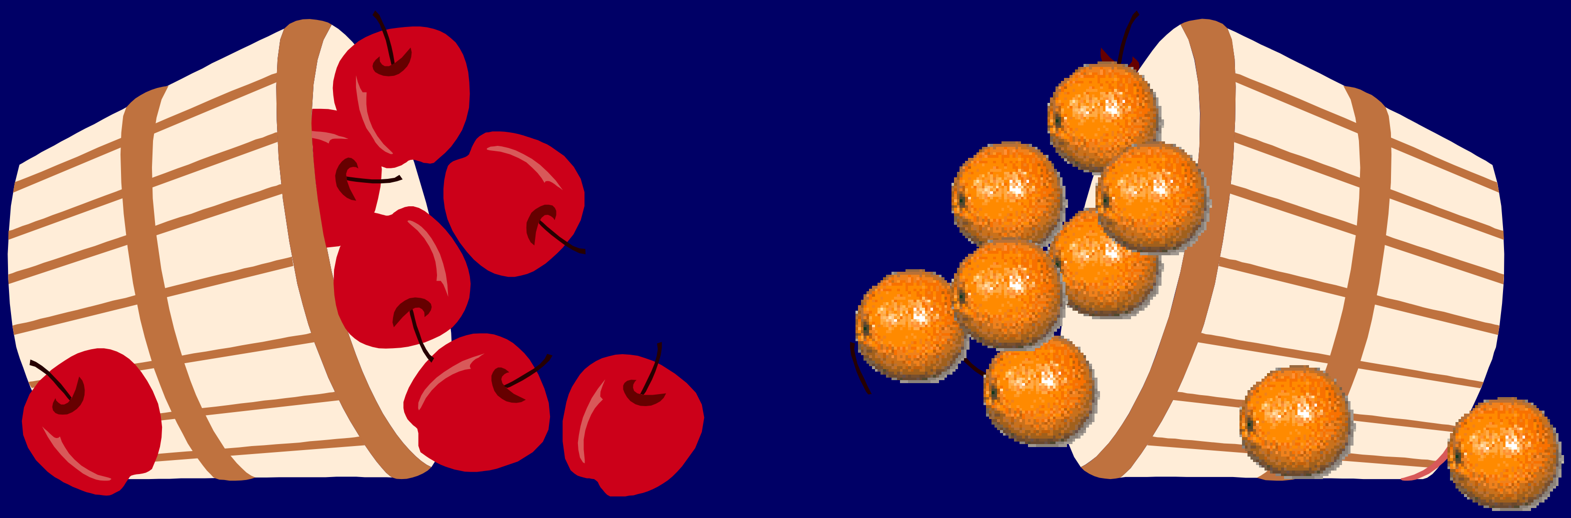

Application-Orientation vs. Subject-Orientation

Application-Orientation

Subject-Orientation

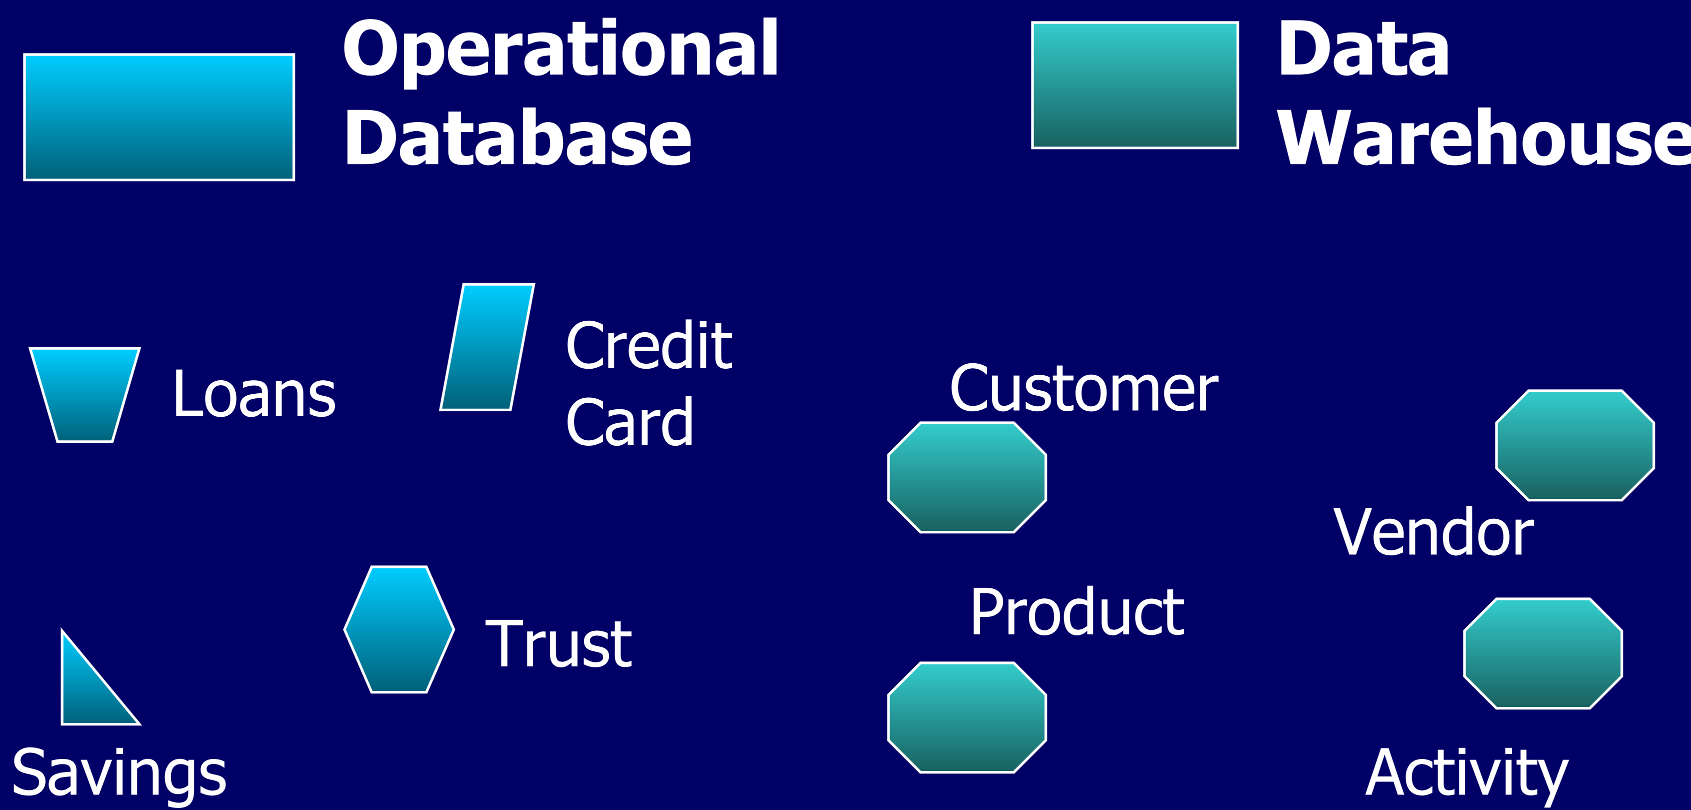

OLTP vs. Data Warehouse

OLTP systems are tuned for known transactions and workloads while workload is not known a priori in a data warehouse

Special data organization, access methods and implementation methods are needed to support data warehouse queries (typically multidimensional queries)

e.g., *average amount spent on phone calls between 9AM-5PM in Pune during the month of December*

OLTP vs Data Warehouse

OLTP

- Application Oriented
- Used to run business
- Detailed data
- Current up to date
- Isolated Data
- Repetitive access
- Clerical User

Warehouse (DSS)

- Subject Oriented
- Used to analyze business
- Summarized and refined
- Snapshot data
- Integrated Data
- Ad-hoc access
- Knowledge User (Manager)

OLTP vs Data Warehouse

OLTP	Data Warehouse
Performance Sensitive	Performance relaxed
Few Records accessed at a time (tens)	Large volumes accessed at a time(millions)
Read/Update Access	Mostly Read (Batch Update)
No data redundancy	Redundancy present
Database Size 100MB -100 GB	Database Size 100 GB - few terabytes

OLTP vs Data Warehouse

OLTP
- Transaction throughput is the performance metric
- Thousands of users
- Managed in entirety

Data Warehouse
- Query throughput is the performance metric
- Hundreds of users
- Managed by subsets

To summarize ...

OLTP Systems are used to *"run"* a business

The Data Warehouse helps to *"optimize"* the business

Data Warehouse Architecture

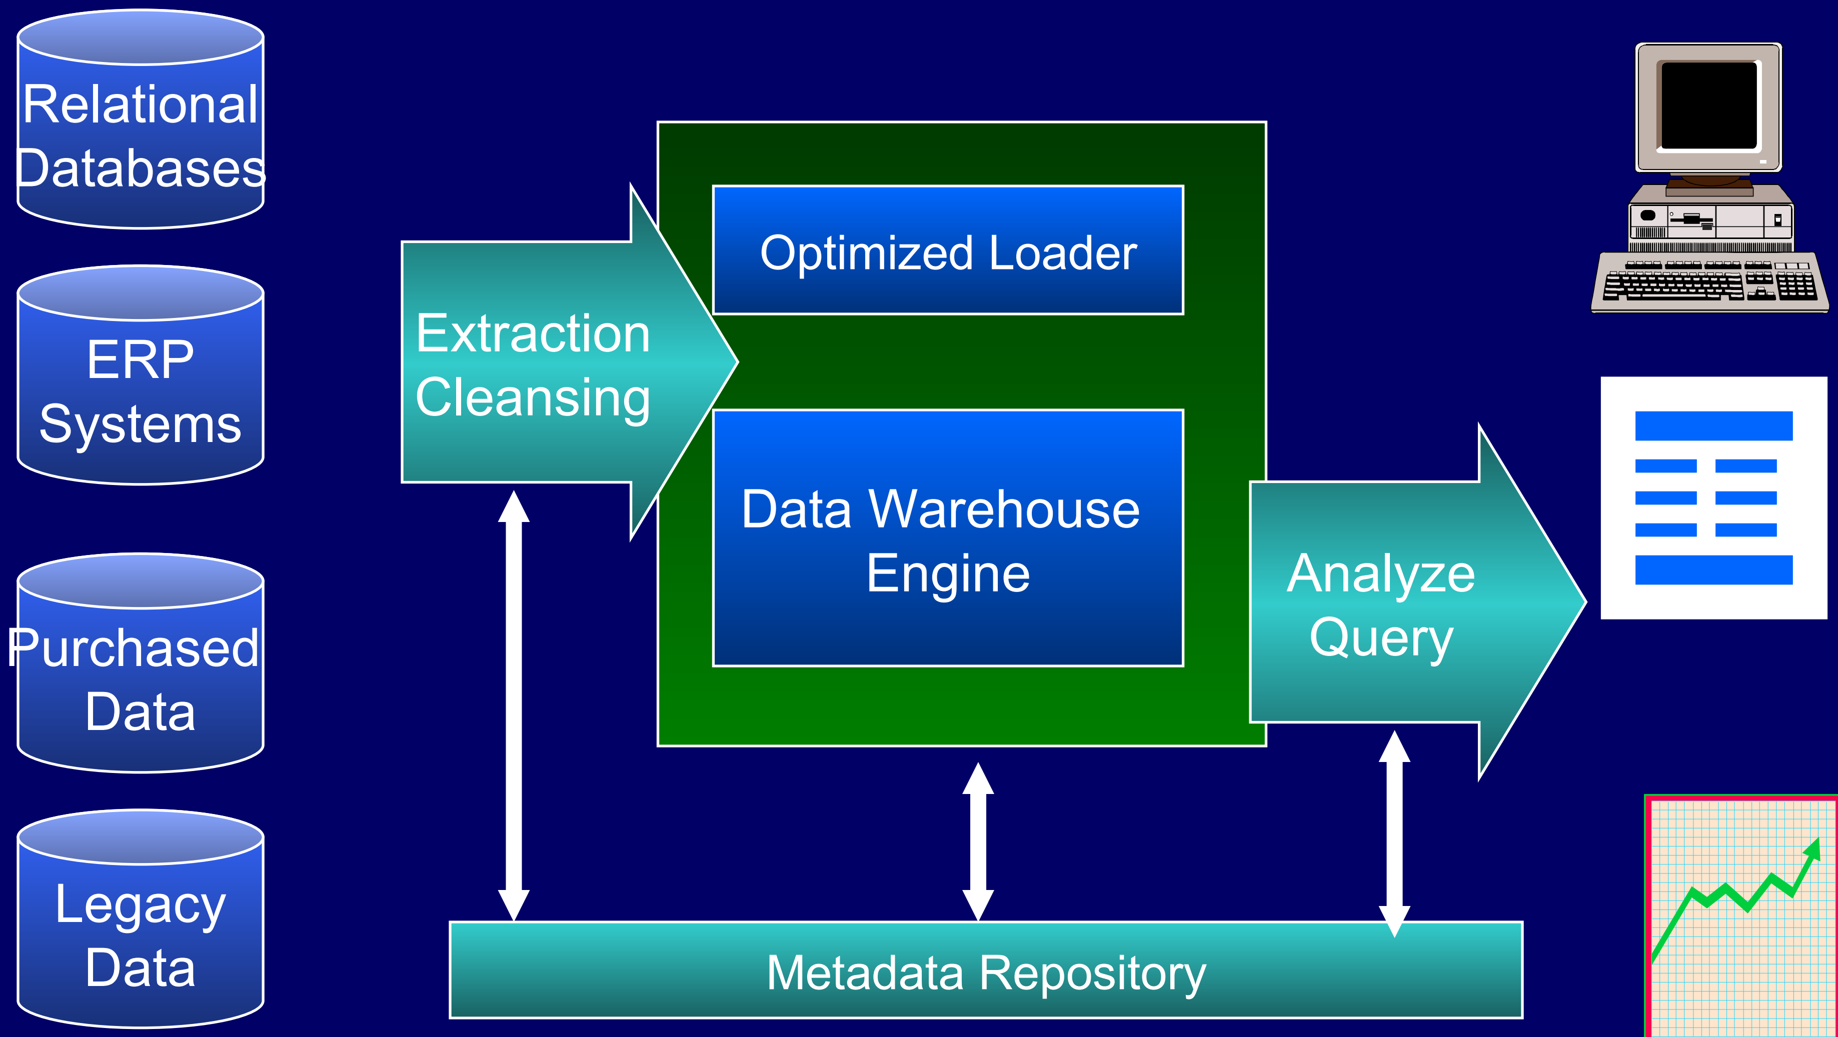

Components of the Warehouse

Data Extraction and Loading

The Warehouse

Analyze and Query -- OLAP Tools

Metadata

Data Mining tools

Loading the Warehouse

Cleaning the data before it is loaded

Source Data

- Typically host based, legacy applications
  - Customized applications, COBOL, 3GL, 4GL
- Point of Contact Devices
  - POS, ATM, Call switches
- External Sources
  - Nielsen's, Acxiom, CMIE, Vendors, Partners

Data Quality - The Reality

Tempting to think creating a data warehouse is simply extracting operational data and entering into a data warehouse

Nothing could be farther from the truth

Warehouse data comes from disparate questionable sources

Data Quality - The Reality

- Legacy systems no longer documented
- Outside sources with questionable quality procedures
- Production systems with no built in integrity checks and no integration
  - Operational systems are usually designed to solve a specific business problem and are rarely developed to a a corporate plan
    - "And get it done quickly, we do not have time to worry about corporate standards..."

Data Integration Across Sources

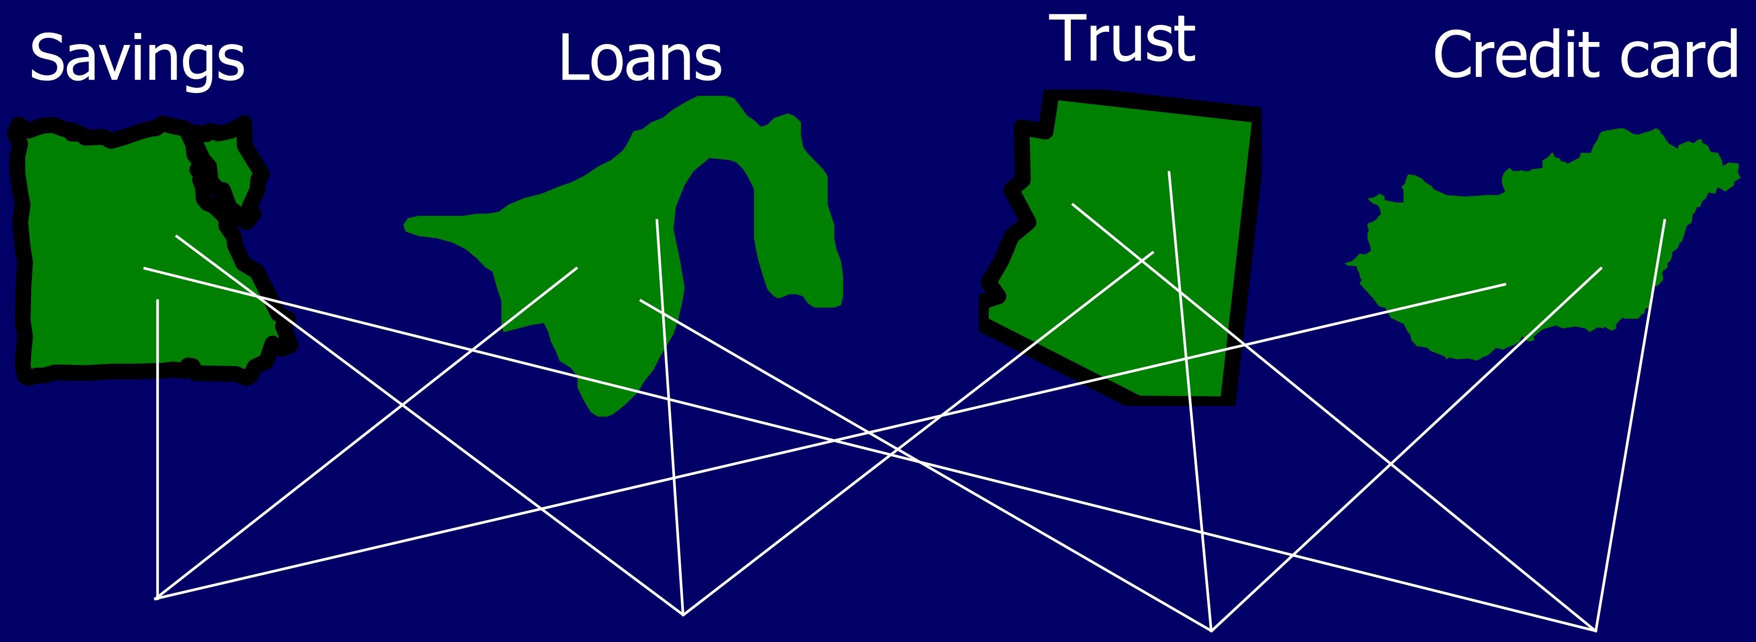

Same data
different name

Different data
Same name

Data found here
nowhere else

Different keys
same data

Data Transformation Example

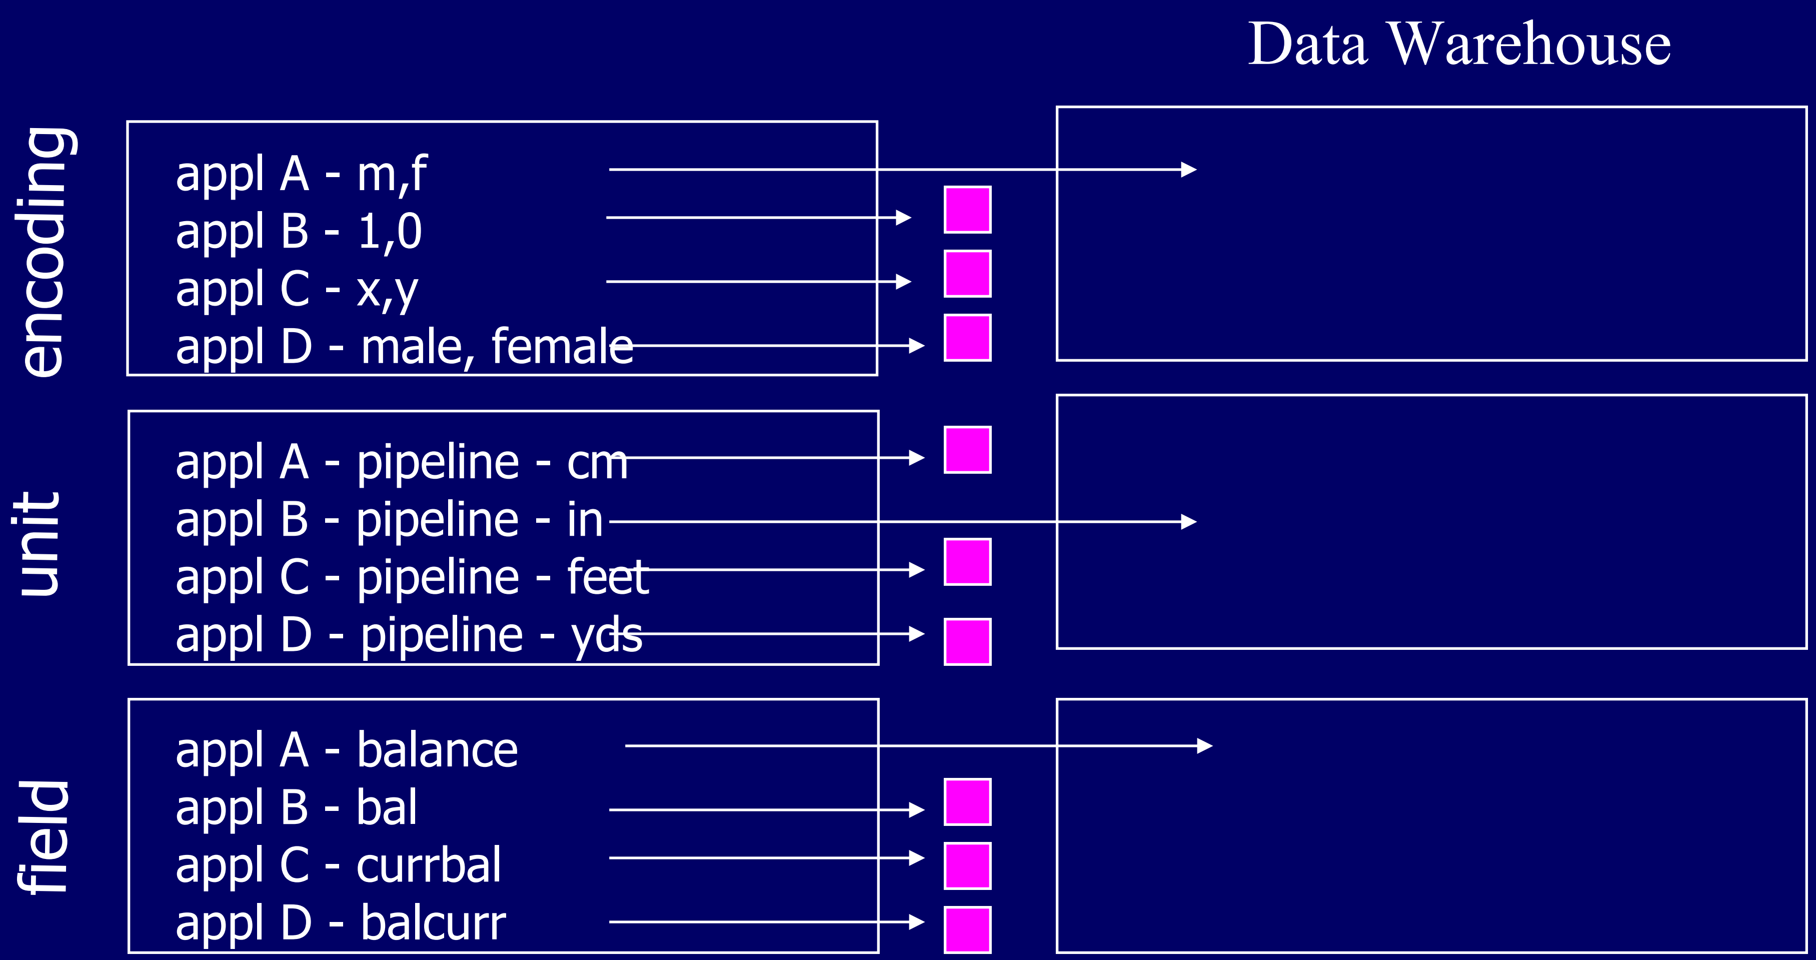

Data Integrity Problems

- Same person, different spellings
  - Agarwal, Agrawal, Aggarwal etc...
- Multiple ways to denote company name
  - Persistent Systems, PSPL, Persistent Pvt. LTD.
- Use of different names
  - mumbai, bombay
- Different account numbers generated by different applications for the same customer
- Required fields left blank
- Invalid product codes collected at point of sale
  - manual entry leads to mistakes
  - “in case of a problem use 9999999”

Data Transformation Terms

- Extracting
- Conditioning
- Scrubbing
- Merging
- Householding
- Enrichment
- Scoring
- Loading
- Validating
- Delta Updating

Data Transformation Terms

- Extracting
  - Capture of data from operational source in "as is" status
  - Sources for data generally in legacy mainframes in VSAM, IMS, IDMS, DB2; more data today in relational databases on Unix
- Conditioning
  - The conversion of data types from the source to the target data store (warehouse) -- always a relational database

Data Transformation Terms

Householding

- Identifying all members of a household (living at the same address)
- Ensures only one mail is sent to a household
- Can result in substantial savings: 1 lakh catalogues at Rs. 50 each costs Rs. 50 lakhs. A 2% savings would save Rs. 1 lakh.

Data Transformation Terms

- Enrichment
  - Bring data from external sources to augment/enrich operational data. Data sources include Dunn and Bradstreet, A. C. Nielsen, CMIE, IMRA etc...
- Scoring
  - computation of a probability of an event. e.g..., chance that a customer will defect to AT&T from MCI, chance that a customer is likely to buy a new product

Loads

- After extracting, scrubbing, cleaning, validating etc. need to load the data into the warehouse
- Issues
  - huge volumes of data to be loaded
  - small time window available when warehouse can be taken off line (usually nights)
  - when to build index and summary tables
  - allow system administrators to monitor, cancel, resume, change load rates
  - Recover gracefully -- restart after failure from where you were and without loss of data integrity

Load Techniques

- Use SQL to append or insert new data
  - record at a time interface
  - will lead to random disk I/O's
- Use batch load utility

Load Taxonomy

Incremental versus Full loads

Online versus Offline loads

Refresh

Propagate updates on source data to the warehouse

Issues:

- when to refresh
- how to refresh -- refresh techniques

When to Refresh?

- periodically (e.g., every night, every week) or after significant events
- on every update: not warranted unless warehouse data require current data (up to the minute stock quotes)
- refresh policy set by administrator based on user needs and traffic
- possibly different policies for different sources

Refresh Techniques

- Full Extract from base tables
  - read entire source table: too expensive
  - maybe the only choice for legacy systems

How To Detect Changes

- Create a snapshot log table to record ids of updated rows of source data and timestamp
- Detect changes by:
  - Defining after row triggers to update snapshot log when source table changes
  - Using regular transaction log to detect changes to source data

Data Extraction and Cleansing

- Extract data from existing operational and legacy data
- Issues:
  - Sources of data for the warehouse
  - Data quality at the sources
  - Merging different data sources
  - Data Transformation
  - How to propagate updates (on the sources) to the warehouse
  - Terabytes of data to be loaded

Scrubbing Data

- Sophisticated transformation tools.
- Used for cleaning the quality of data
- Clean data is vital for the success of the warehouse
- Example
  - Seshadri, Sheshadri, Sesadri, Seshadri S., Srinivasan Seshadri, etc. are the same person

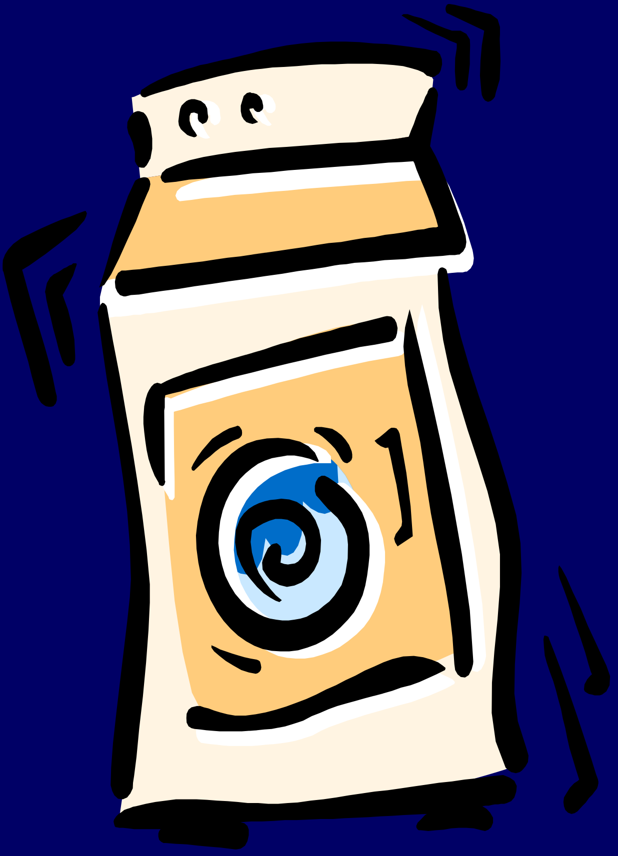

Scrubbing Tools

Apertus -- Enterprise/Integrator

Vality -- IPE

Postal Soft

Structuring/Modeling Issues

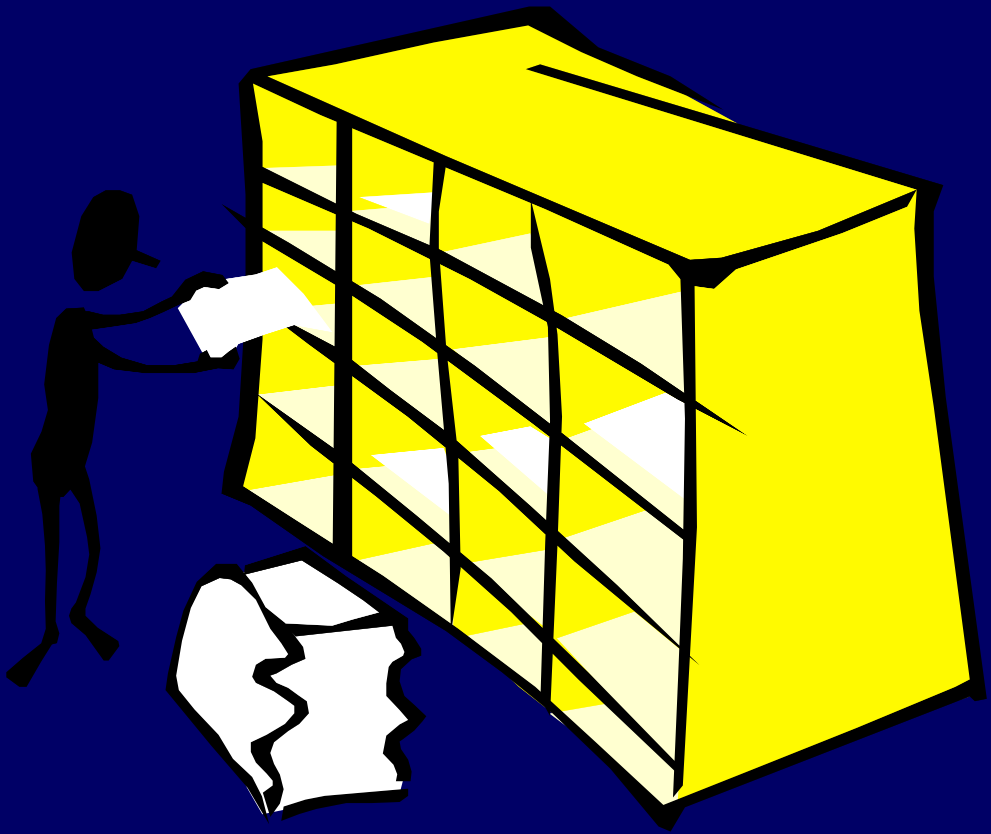

Data -- Heart of the Data Warehouse

- Heart of the data warehouse is the data itself!
- Single version of the truth
- Corporate memory
- Data is organized in a way that represents business -- subject orientation

Data Warehouse Structure

Subject Orientation -- customer, product, policy, account etc... A subject may be implemented as a set of related tables. E.g., customer may be five tables

Data Warehouse Structure

Time is part of key of each table

- base customer (1985-87)
  - custid, from date, to date, name, phone, dob
- base customer (1988-90)
  - custid, from date, to date, name, credit rating, employer
- customer activity (1986-89) -- monthly summary
- customer activity detail (1987-89)
  - custid, activity date, amount, clerk id, order no
- customer activity detail (1990-91)
  - custid, activity date, amount, line item no, order no

Data Granularity in Warehouse

- Summarized data stored
  - reduce storage costs
  - reduce cpu usage
  - increases performance since smaller number of records to be processed
  - design around traditional high level reporting needs
  - tradeoff with volume of data to be stored and detailed usage of data

Granularity in Warehouse

- Can not answer some questions with summarized data
  - Did Anand call Seshadri last month? Not possible to answer if total duration of calls by Anand over a month is only maintained and individual call details are not.
- Detailed data too voluminous

Granularity in Warehouse

- Tradeoff is to have dual level of granularity
  - Store summary data on disks
    - 95% of DSS processing done against this data
  - Store detail on tapes
    - 5% of DSS processing against this data

Vertical Partitioning

Acct. No	Name	Balance	Date Opened	Interest Rate	Address

Frequently accessed

Rarely accessed

Acct. No	Balance

Acct. No	Name	Date Opened	Interest Rate	Address

Smaller table and so less I/O

Derived Data

- Introduction of derived (calculated data) may often help
- Have seen this in the context of dual levels of granularity
- Can keep auxiliary views and indexes to speed up query processing

Schema Design

- Database organization
  - must look like business
  - must be recognizable by business user
  - approachable by business user
  - Must be *simple*
- Schema Types
  - Star Schema
  - Fact Constellation Schema
  - Snowflake schema

Dimension Tables

- Dimension tables
  - Define business in terms already familiar to users
  - Wide rows with lots of descriptive text
  - Small tables (about a million rows)
  - Joined to fact table by a foreign key
  - heavily indexed
  - typical dimensions
    - time periods, geographic region (markets, cities), products, customers, salesperson, etc.

Fact Table

- Central table
  - mostly raw numeric items
  - narrow rows, a few columns at most
  - large number of rows (millions to a billion)
  - Access via dimensions

Star Schema

A single fact table and for each dimension one dimension table

Does not capture hierarchies directly

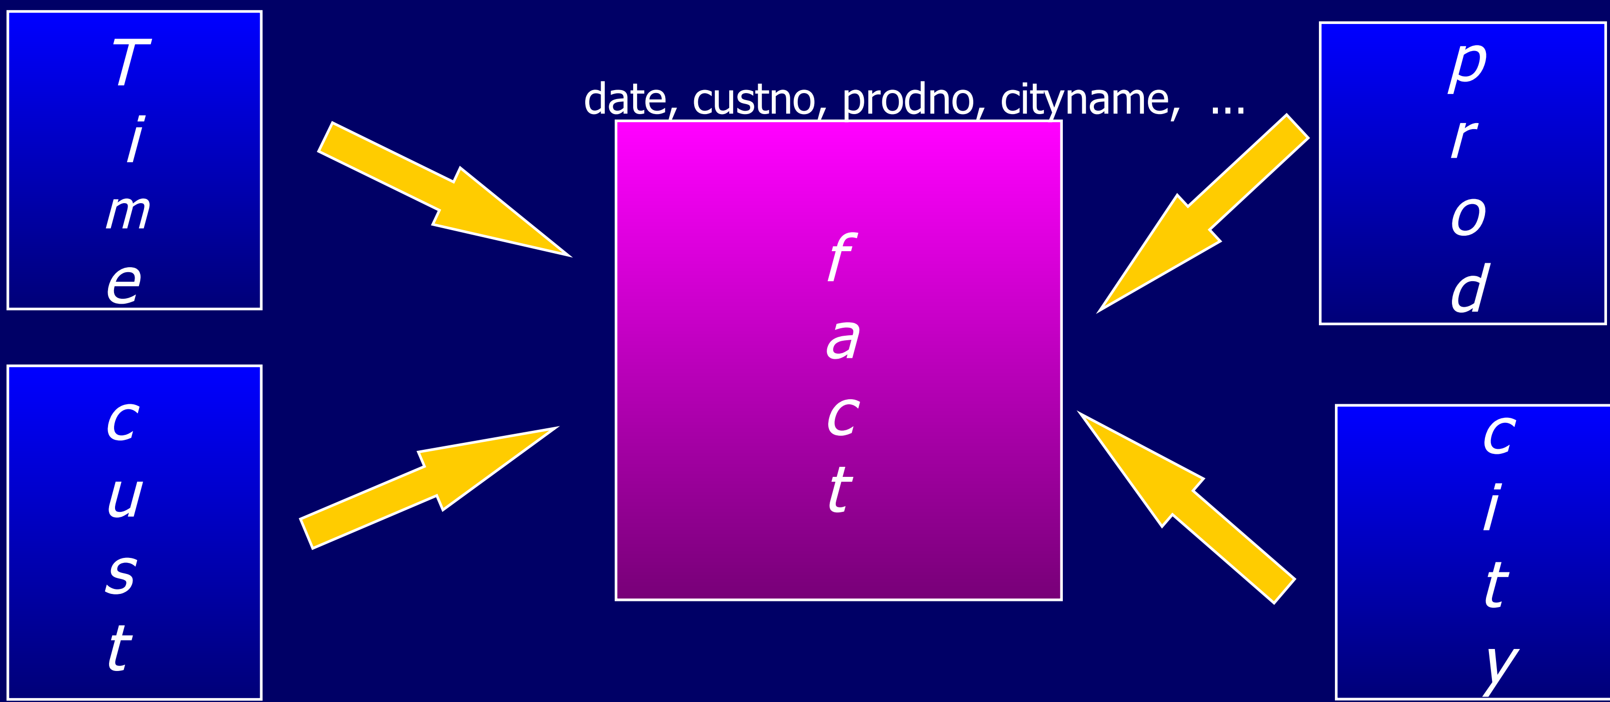

Snowflake schema

Represent dimensional hierarchy directly by normalizing tables.

Easy to maintain and saves storage

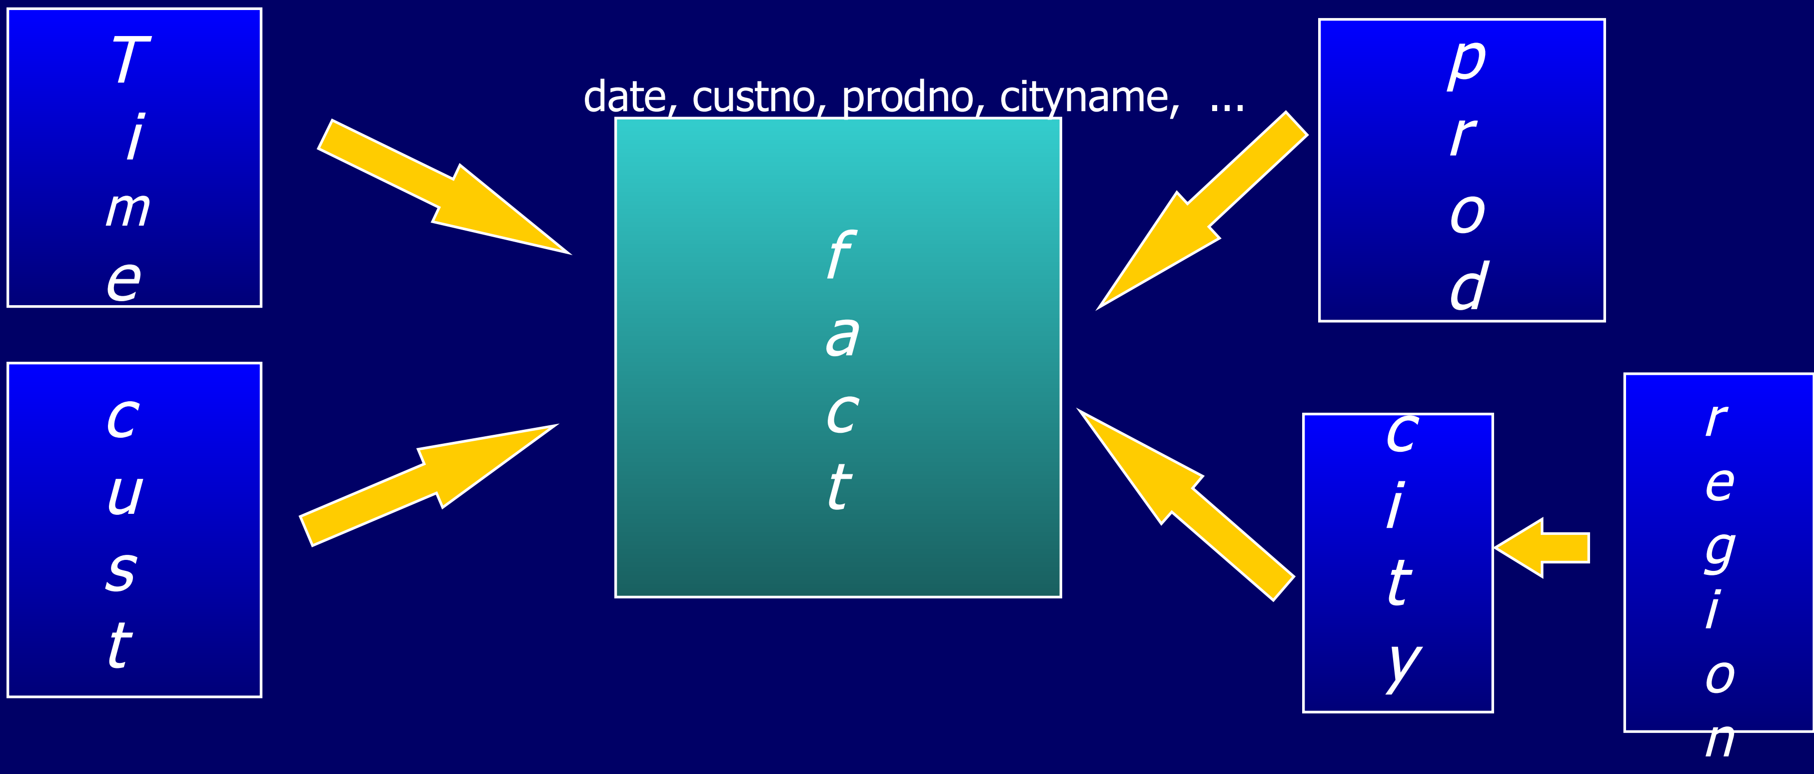

Fact Constellation

- Fact Constellation
  - Multiple fact tables that share many dimension tables
  - Booking and Checkout may share many dimension tables in the hotel industry

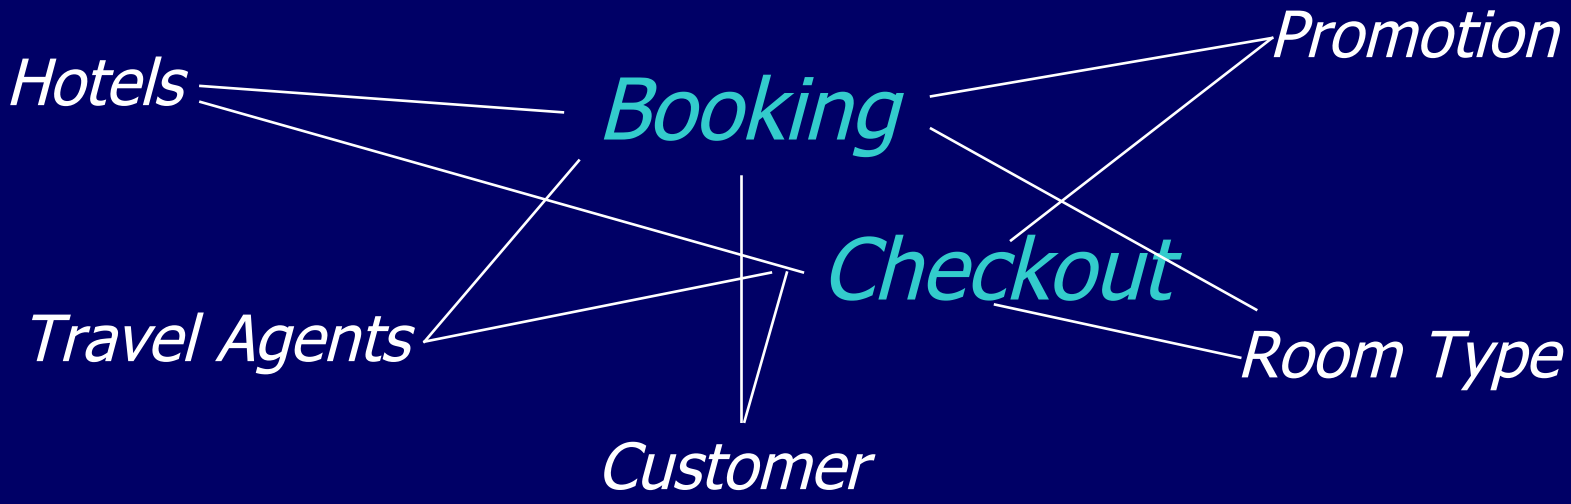

De-normalization

- Normalization in a data warehouse may lead to lots of small tables
- Can lead to excessive I/O’s since many tables have to be accessed
- De-normalization is the answer especially since updates are rare

Creating Arrays

Many times each occurrence of a sequence of data is in a different physical location

Beneficial to collect all occurrences together and store as an array in a single row

Makes sense only if there are a stable number of occurrences which are accessed together

In a data warehouse, such situations arise naturally due to time based orientation

- can create an array by month

Selective Redundancy

Description of an item can be stored redundantly with order table -- most often item description is also accessed with order table

Updates have to be careful

Partitioning

- Breaking data into several physical units that can be handled separately
- Not a question of *whether* to do it in data warehouses but *how* to do it
- Granularity and partitioning are key to effective implementation of a warehouse

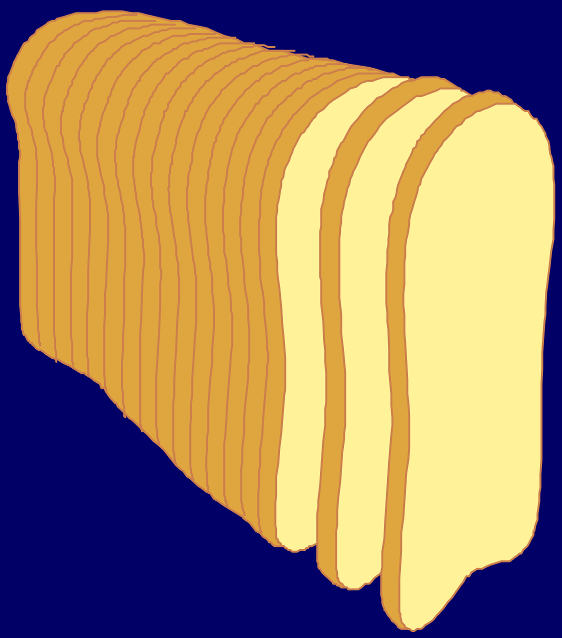

Why Partition?

- Flexibility in managing data
- Smaller physical units allow
  - easy restructuring
  - free indexing
  - sequential scans if needed
  - easy reorganization
  - easy recovery
  - easy monitoring

Criterion for Partitioning

- Typically partitioned by
  - date
  - line of business
  - geography
  - organizational unit
  - any combination of above

Where to Partition?

- Application level or DBMS level
- Makes sense to partition at application level
  - Allows different definition for each year
    - Important since warehouse spans many years and as business evolves definition changes
  - Allows data to be moved between processing complexes easily

Data Warehouse vs. Data Marts

What comes first

From the Data Warehouse to Data Marts

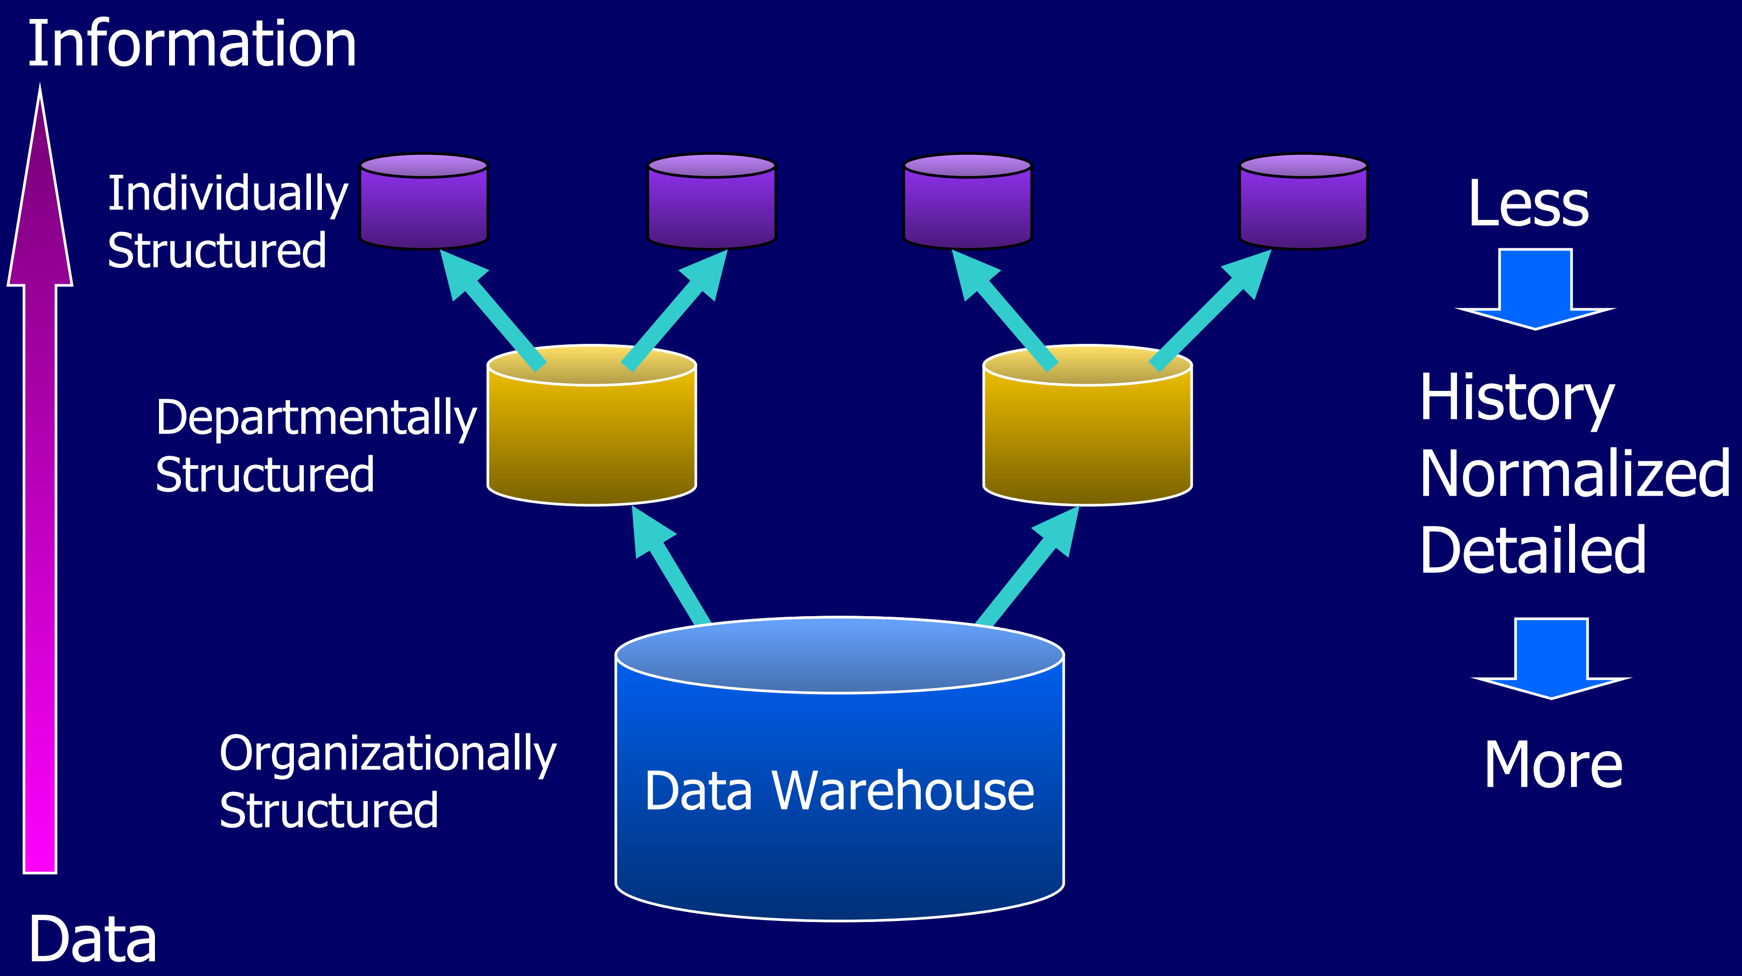

Data Warehouse and Data Marts

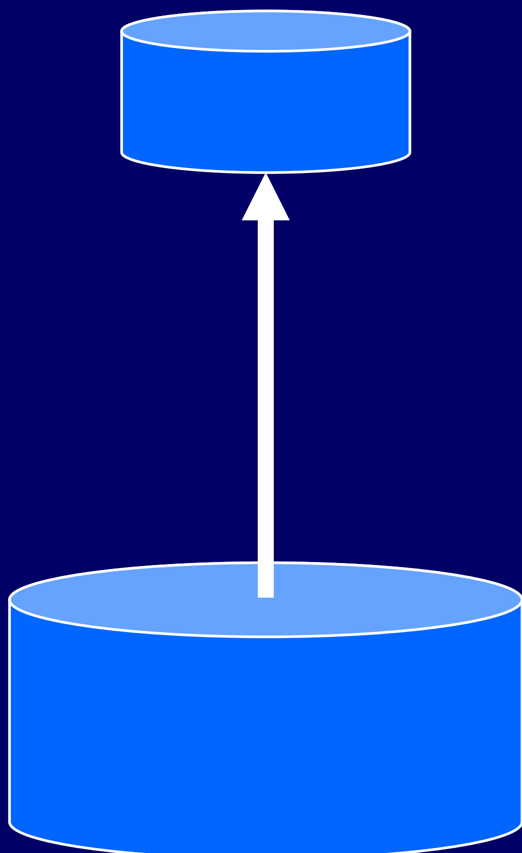

OLAP
Data Mart
Lightly summarized
Departmentally structured

Organizationally structured
Atomic
Detailed Data Warehouse Data

Characteristics of the Departmental Data Mart

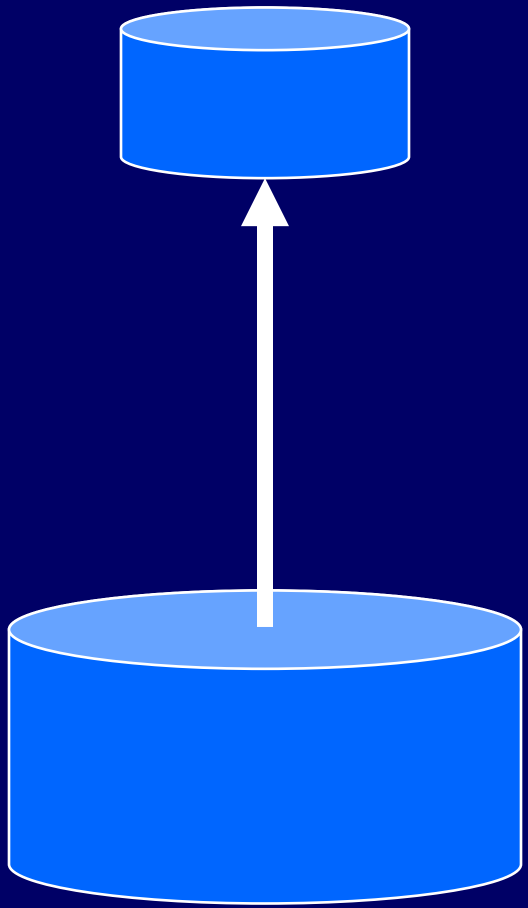

OLAP

Small

Flexible

Customized by Department

Source is departmentally structured data warehouse

Techniques for Creating Departmental Data Mart

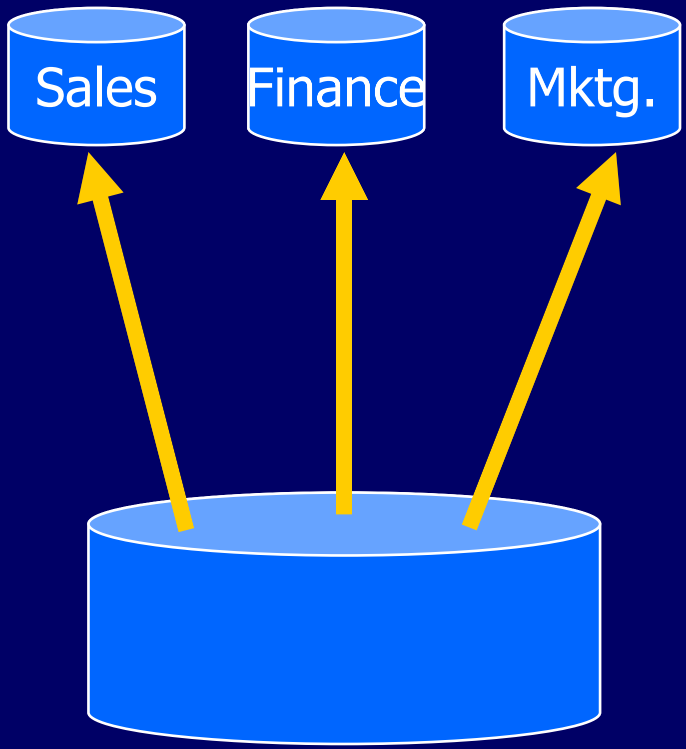

OLAP

Subset

Summarized

Superset

Indexed

Arrayed

Data Mart Centric

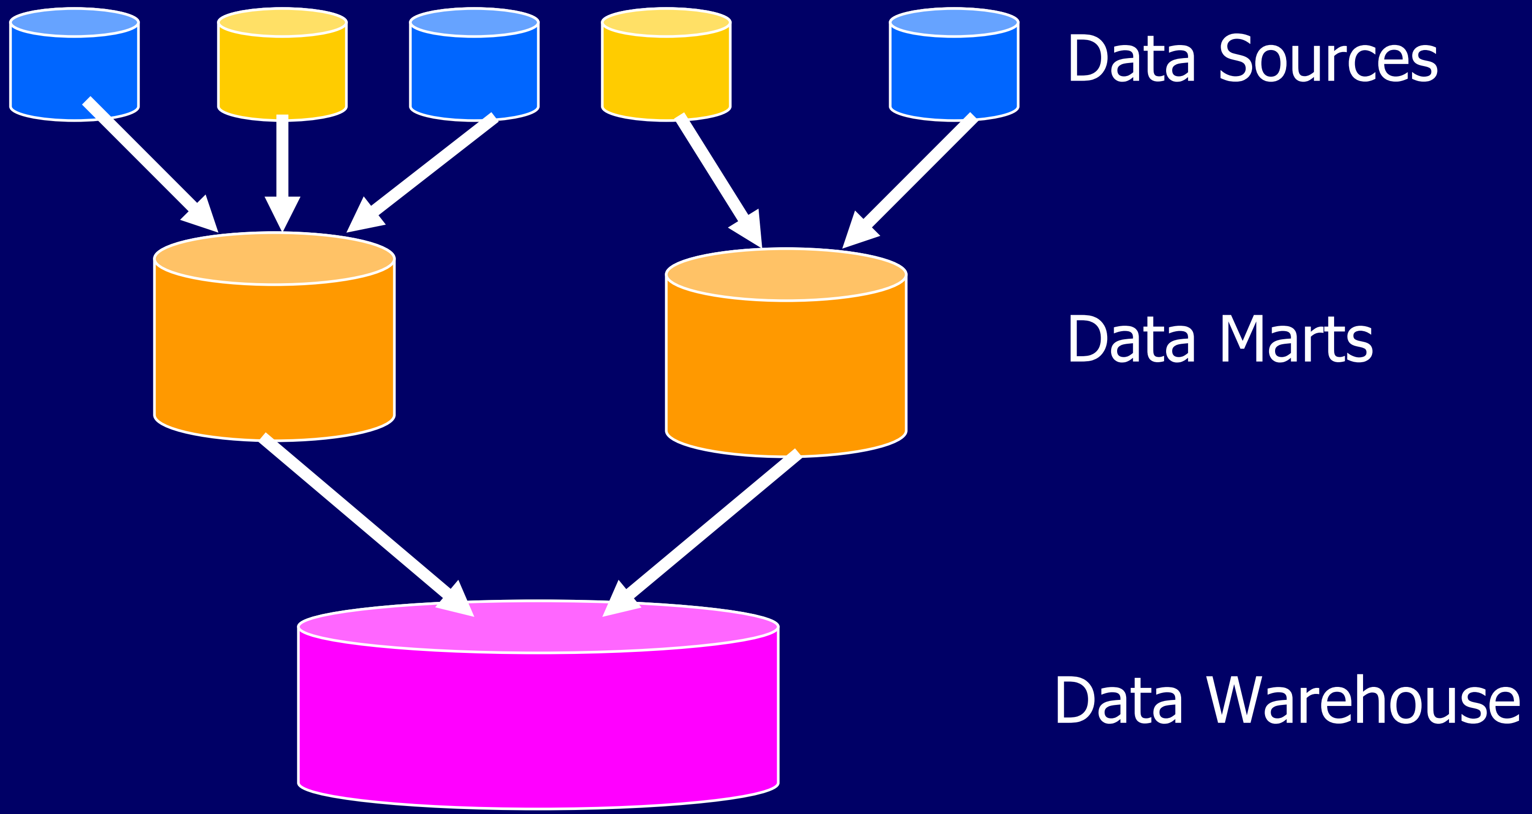

Problems with Data Mart Centric Solution

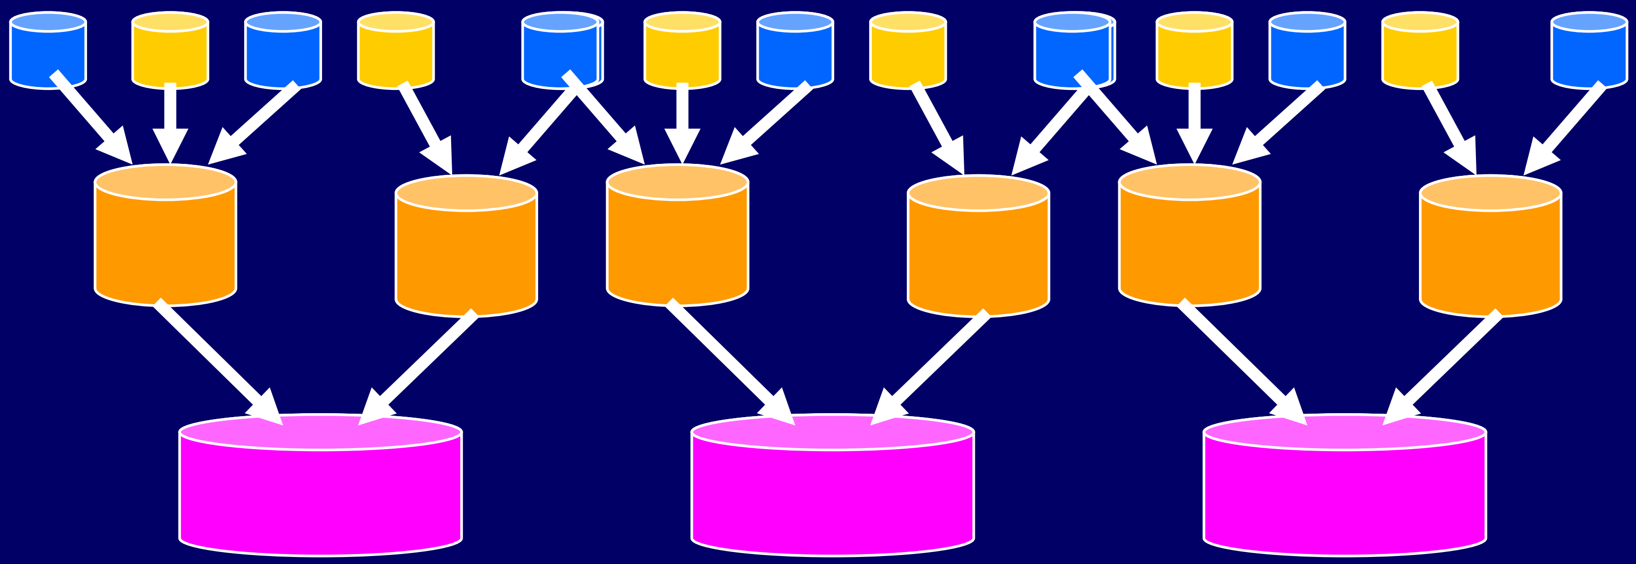

If you end up creating multiple warehouses, integrating them is a problem

True Warehouse

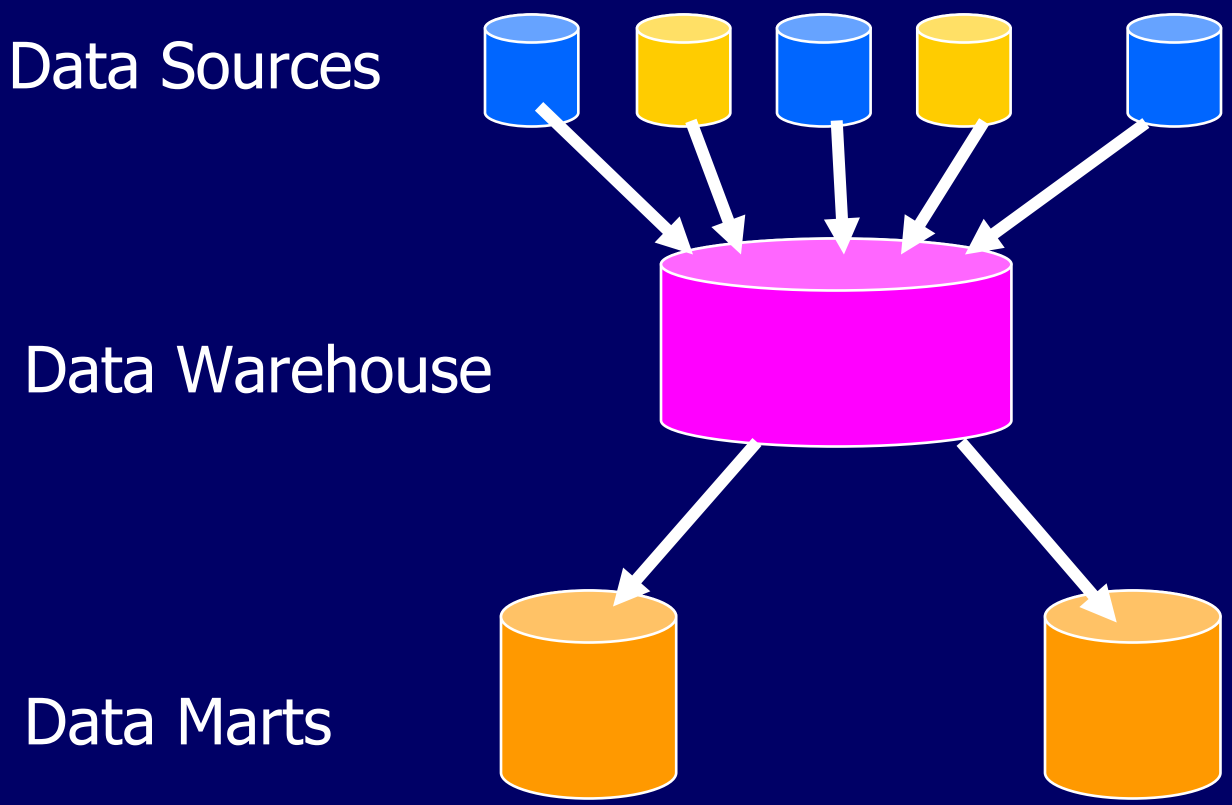

Query Processing

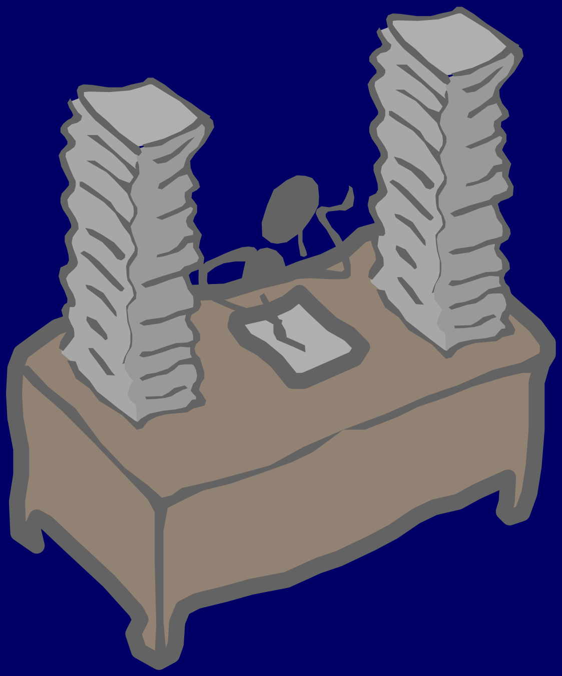

Indexing

Pre computed views/aggregates

SQL extensions

Indexing Techniques

- Exploiting indexes to reduce scanning of data is of crucial importance
- Bitmap Indexes
- Join Indexes
- Other Issues
  - Text indexing
  - Parallelizing and sequencing of index builds and incremental updates

Indexing Techniques

Bitmap index:

- A collection of bitmaps -- one for each distinct value of the column
- Each bitmap has N bits where N is the number of rows in the table
- A bit corresponding to a value v for a row r is set if and only if r has the value for the indexed attribute

BitMap Indexes

An alternative representation of RID-list

Specially advantageous for low-cardinality domains

Represent each row of a table by a bit and the table as a bit vector

There is a distinct bit vector Bv for each value v for the domain

Example: the attribute sex has values M and F. A table of 100 million people needs 2 lists of 100 million bits

Bitmap Index

gender	vote	
M	Y	
F	Y	
F	N	
M	N	
F	Y	
F	N	

Customer

gender (f)	vote (y)	result
0	1	0
1	1	1
1	0	0
0	0	0
1	1	1
1	0	0

Query : select * from customer where gender = 'F' and vote = 'Y'

Bit Map Index

Base Table

Cust	Region	Rating
C1	N	H
C2	S	M
C3	W	L
C4	W	H
C5	S	L
C6	W	L
C7	N	H

Region Index

Row ID	N	S	E	W
1	1	0	0	0
2	0	1	0	0
3	0	0	0	1
4	0	0	0	1
5	0	1	0	0
6	0	0	0	1
7	1	0	0	0

Rating Index

Row ID	H	M	L
1	1	0	0
2	0	1	0
3	0	0	0
4	0	0	0
5	0	1	0
6	0	0	0
7	1	0	0

Customers where **Region = W** ***And*** **Rating = M**

BitMap Indexes

Comparison, join and aggregation operations are reduced to bit arithmetic with dramatic improvement in processing time

Significant reduction in space and I/O (30:1)

Adapted for higher cardinality domains as well.

Compression (e.g., run-length encoding) exploited

Products that support bitmaps: Model 204, TargetIndex (Redbrick), IQ (Sybase), Oracle 7.3

Join Indexes

Pre-computed joins

A join index between a fact table and a dimension table correlates a dimension tuple with the fact tuples that have the same value on the common dimensional attribute

e.g., a join index on *city* dimension of *calls* fact table

correlates for each city the calls (in the *calls* table) from that city

Join Indexes

Join indexes can also span multiple dimension tables

- e.g., a join index on *city* and *time* dimension of *calls* fact table

Star Join Processing

Use join indexes to join dimension and fact table

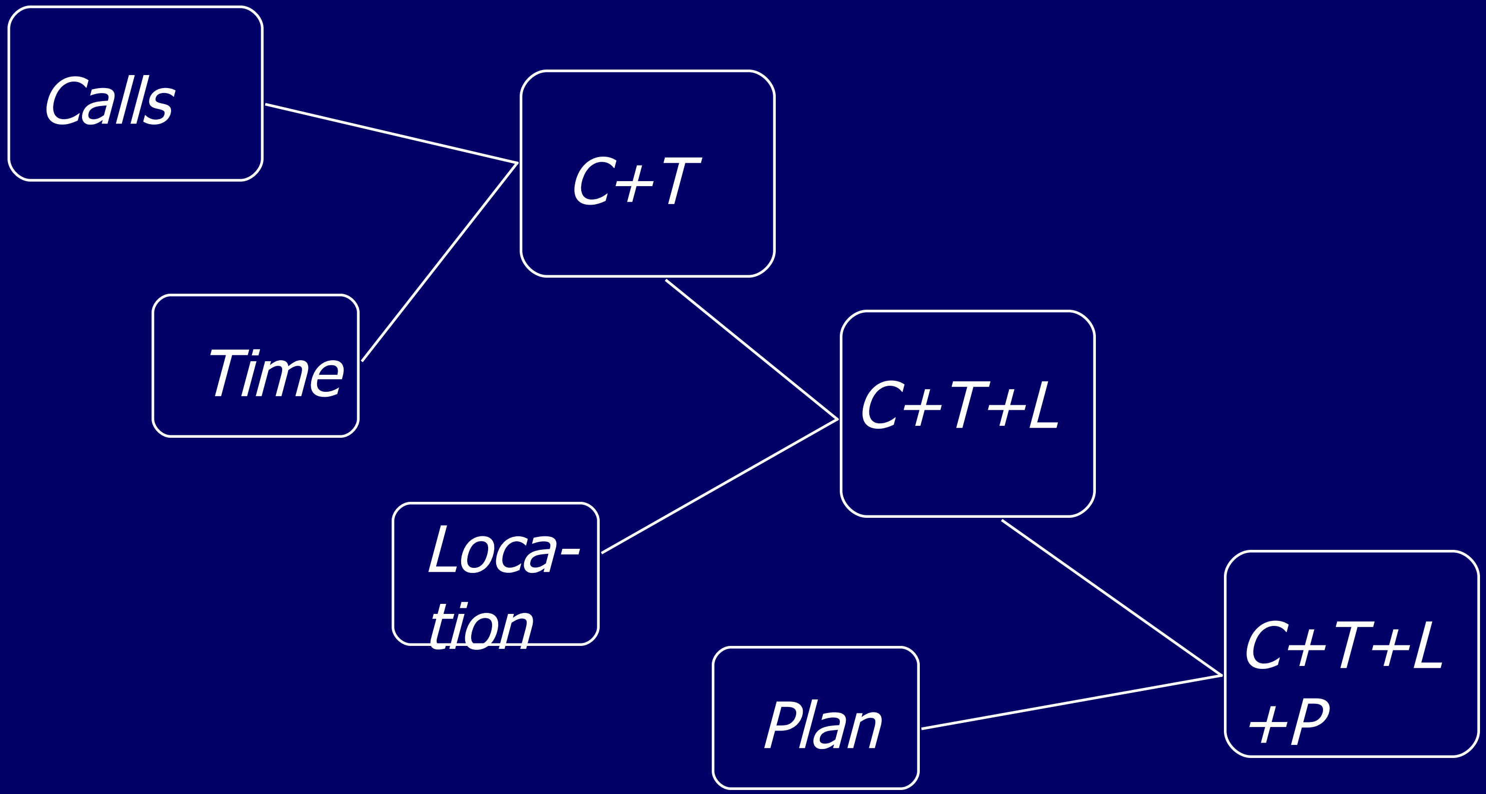

Optimized Star Join Processing

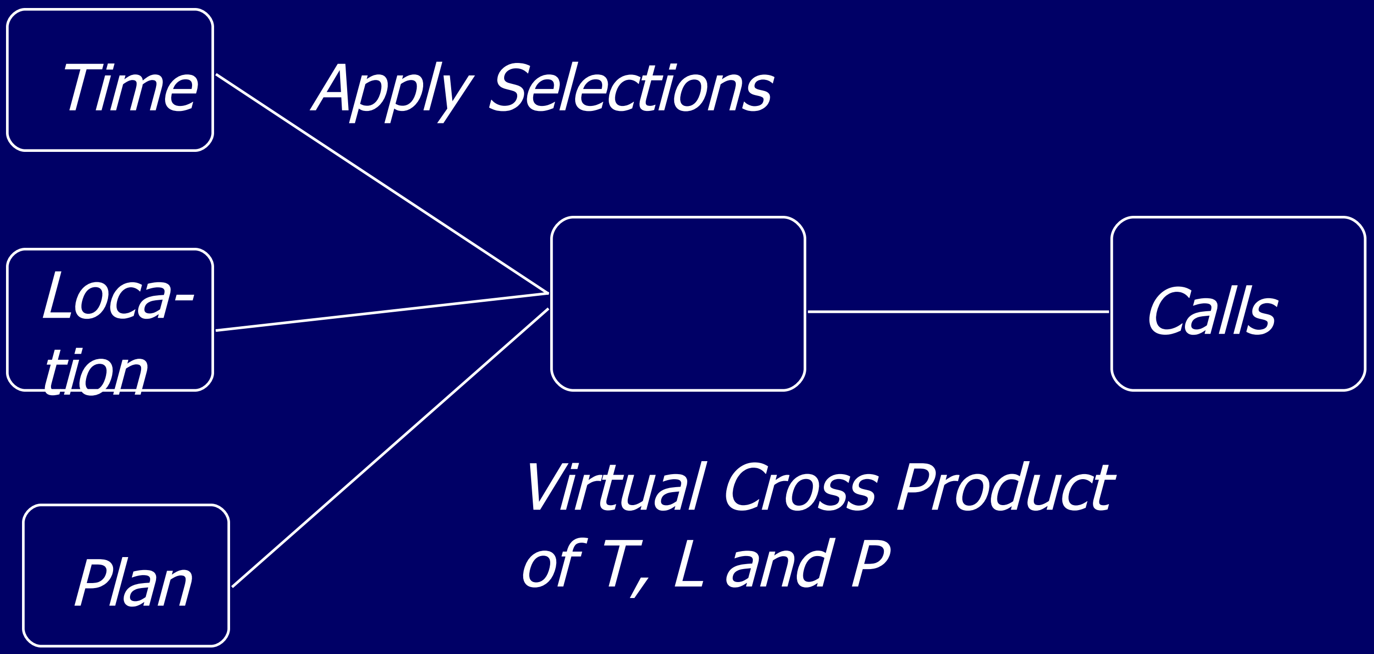

Bitmapped Join Processing

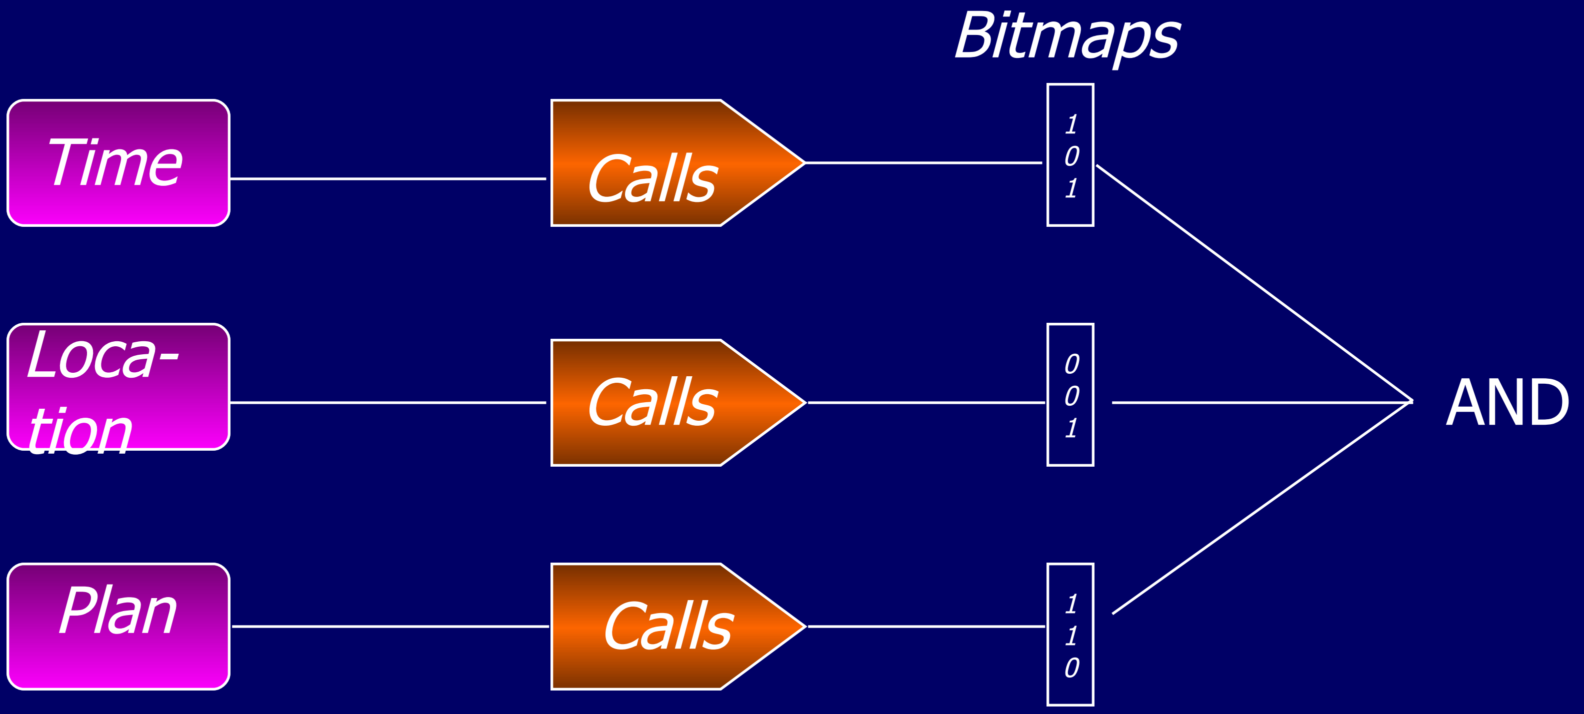

Intelligent Scan

- Piggyback multiple scans of a relation (Redbrick)
  - piggybacking also done if second scan starts a little while after the first scan

Parallel Query Processing

- Three forms of parallelism
  - Independent
  - Pipelined
  - Partitioned and “partition and replicate”
- Deterrents to parallelism
  - startup
  - communication

Parallel Query Processing

- Partitioned Data
  - Parallel scans
  - Yields I/O parallelism
- Parallel algorithms for relational operators
  - Joins, Aggregates, Sort
- Parallel Utilities
  - Load, Archive, Update, Parse, Checkpoint, Recovery
- Parallel Query Optimization

Pre-computed Aggregates

- Keep aggregated data for efficiency (pre-computed queries)
- Questions
  - Which aggregates to compute?
  - How to update aggregates?
  - How to use pre-computed aggregates in queries?

Pre-computed Aggregates

Aggregated table can be maintained by the

- warehouse server
- middle tier
- client applications

Pre-computed aggregates -- special case of materialized views -- same questions and issues remain

SQL Extensions

Extended family of aggregate functions

- rank (top 10 customers)
- percentile (top 30% of customers)
- median, mode
- Object Relational Systems allow addition of new aggregate functions

SQL Extensions

- Reporting features
  - running total, cumulative totals
- Cube operator
  - group by on all subsets of a set of attributes (month,city)
  - redundant scan and sorting of data can be avoided

Red Brick has Extended set of Aggregates

```
Select month, dollars, cume(dollars) as
run_dollars, weight, cume(weight) as
run_weights
from sales, market, product, period t
where year = 1993
and product like 'Columbian%'
and city like 'San Fr%'
order by t.perkey
```

RISQL (Red Brick Systems) Extensions

- Aggregates
  - CUME
  - MOVINGAVG
  - MOVINGSUM
  - RANK
  - TERTILE
  - RATIOTOREPORT
- Calculating Row Subtotals
  - BREAK BY
- Sophisticated Date Time Support
  - DATEDIFF
- Using SubQueries in calculations

Using SubQueries in Calculations

```
select product, dollars as jun97_sales,
        (select sum(s1.dollars)
        from market mi, product pi, period, ti, sales si
        where pi.product = product.product
        and   ti.year    = period.year
        and   mi.city    = market.city) as total97_sales,
        100 * dollars/
        (select sum(s1.dollars)
        from market mi, product pi, period, ti, sales si
        where pi.product = product.product
        and   ti.year    = period.year
        and   mi.city    = market.city) as percent_of_yr
from market, product, period, sales
where year = 1997
and   month = 'June' and city like 'Ahmed%'
order by product;
```

Course Overview

The course:
what and how

0. Introduction

I. Data Warehousing

II. Decision Support and OLAP

III. Data Mining

IV. Looking Ahead

Demos and Labs

II. On-Line Analytical Processing (OLAP)

Making Decision Support Possible

Limitations of SQL

"A Freshman in Business needs a Ph.D. in SQL"

-- Ralph Kimball

Typical OLAP Queries

Write a multi-table join to compare sales for each product line YTD this year vs. last year.

Repeat the above process to find the top 5 product contributors to margin.

Repeat the above process to find the sales of a product line to new vs. existing customers.

Repeat the above process to find the customers that have had negative sales growth.

What Is OLAP?

Online Analytical Processing - coined by EF Codd in 1994 paper contracted by Arbor Software*

Generally synonymous with earlier terms such as Decisions Support, Business Intelligence, Executive Information System

OLAP = Multidimensional Database

MOLAP: Multidimensional OLAP (Arbor Essbase, Oracle Express)

ROLAP: Relational OLAP (Informix MetaCube, Microstrategy DSS Agent)

*** Reference: http://www.arborsoft.com/essbase/wht_ppr/coddTOC.html**

The OLAP Market

- Rapid growth in the enterprise market
  - 1995: $700 Million
  - 1997: $2.1 Billion
- Significant consolidation activity among major DBMS vendors
  - 10/94: Sybase acquires ExpressWay
  - 7/95: Oracle acquires Express
  - 11/95: Informix acquires Metacube
  - 1/97: Arbor partners up with IBM
  - 10/96: Microsoft acquires Panorama
- Result: OLAP shifted from small vertical niche to mainstream DBMS category

Strengths of OLAP

It is a powerful visualization paradigm

It provides fast, interactive response times

It is good for analyzing time series

It can be useful to find some clusters and outliers

Many vendors offer OLAP tools

OLAP Is FASMI

Fast

Analysis

Shared

Multidimensional

Information

Nigel Pendse, Richard Creath - The OLAP Report

Multi-dimensional Data

"Hey…I sold $100M worth of goods"

Dimensions: Product, Region, Time
Hierarchical summarization paths

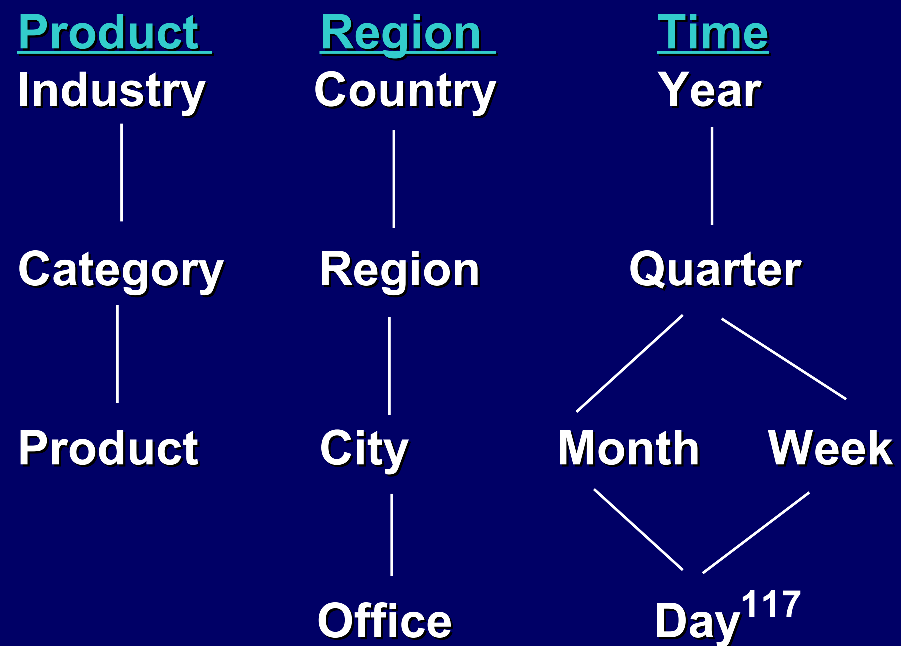

Data Cube Lattice

Cube lattice

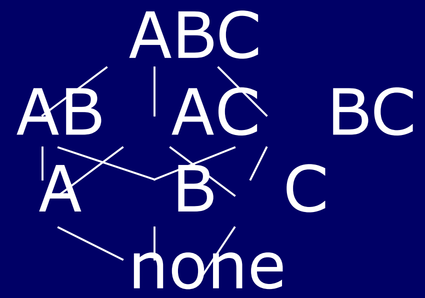

Can materialize some groupbys, compute others on demand

Question: which groupbys to materialze?

Question: what indices to create

Question: how to organize data (chunks, etc)

Visualizing Neighbors is simpler

	1	2	3	4	5	6	7	8
Apr								
May								
Jun								
Jul								
Aug								
Sep								
Oct								
Nov								
Dec								
Jan								
Feb								
Mar								

Month	Store	Sales
Apr	1	
Apr	2	
Apr	3	
Apr	4	
Apr	5	
Apr	6	
Apr	7	
Apr	8	
May	1	
May	2	
May	3	
May	4	
May	5	
May	6	
May	7	
May	8	
Jun	1	
Jun	2	

A Visual Operation: Pivot (Rotate)

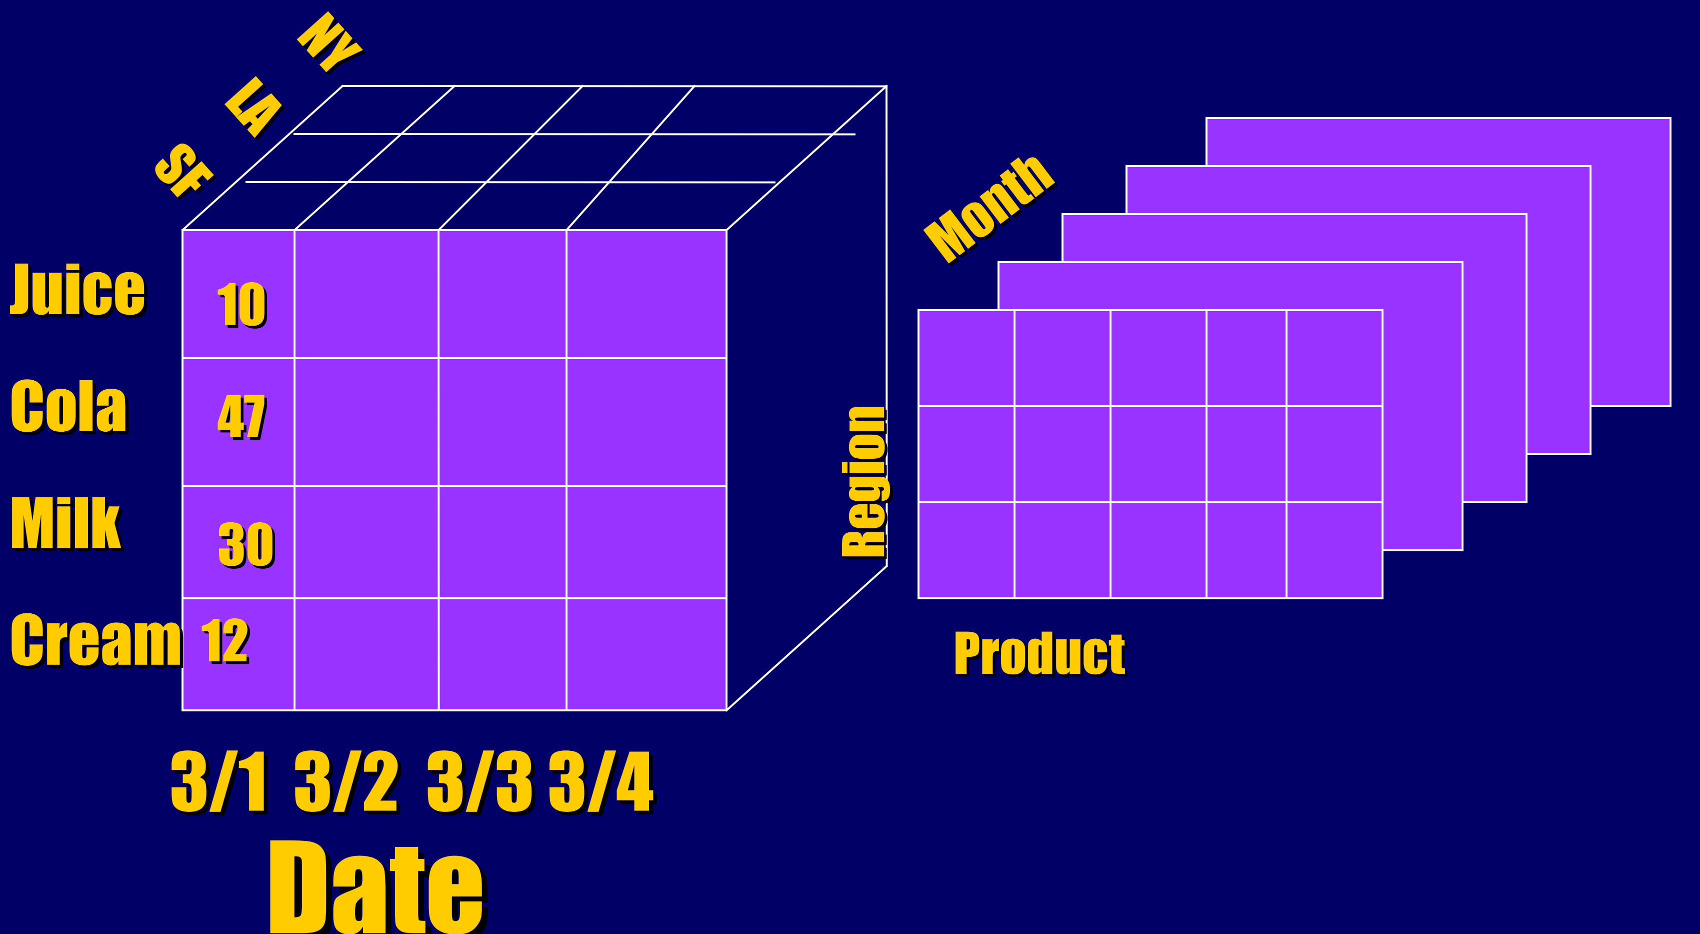

“Slicing and Dicing“

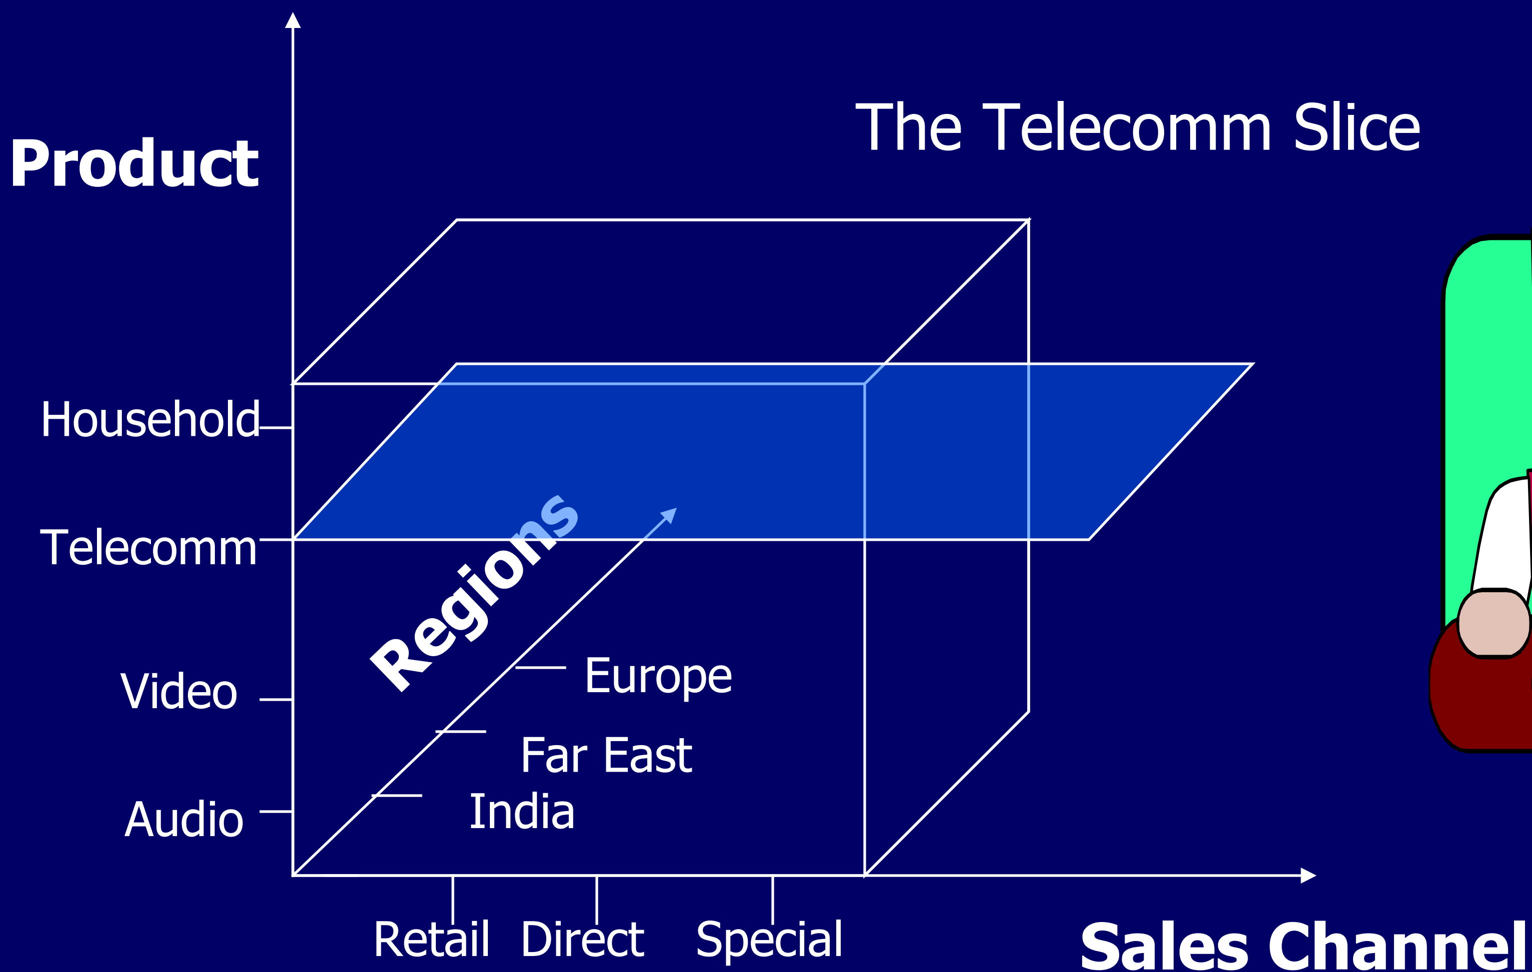

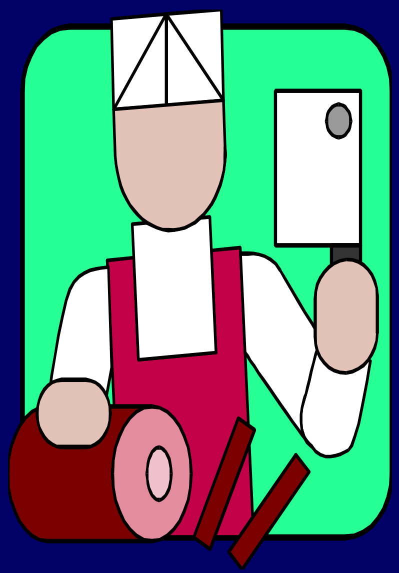

Roll-up and Drill Down

Higher Level of Aggregation

Roll Up

Sales Channel

Region

Country

State

Location Address

Sales Representative

Drill-Down

Low-level Details

Nature of OLAP Analysis

- Aggregation -- (total sales, percent-to-total)
- Comparison -- Budget vs. Expenses
- Ranking -- Top 10, quartile analysis
- Access to detailed and aggregate data
- Complex criteria specification
- Visualization

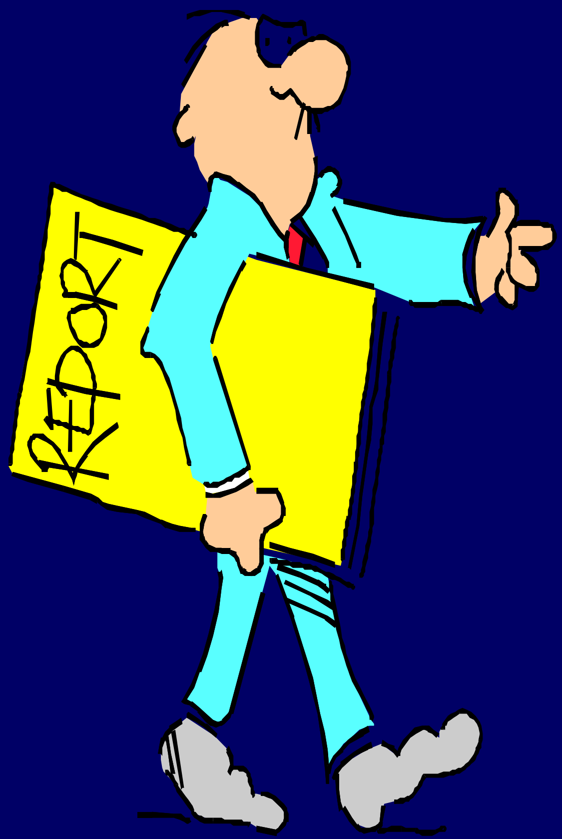

Organizationally Structured Data

Different Departments look at the same detailed data in different ways. Without the detailed, organizationally structured data as a foundation, there is no reconcilability of data

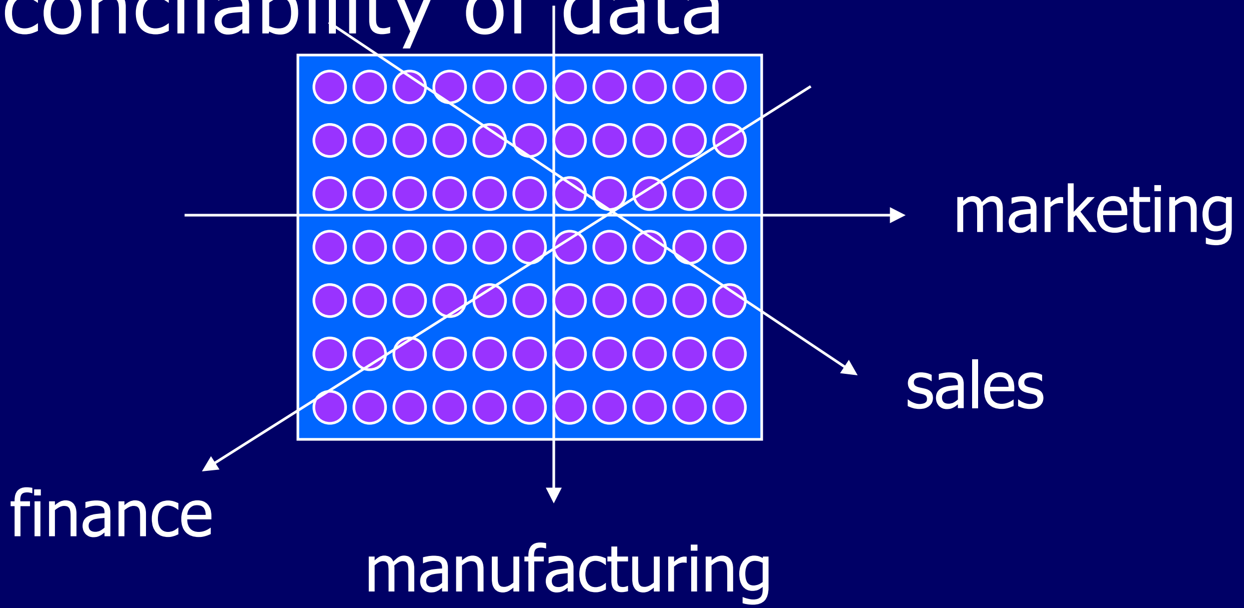

Multidimensional Spreadsheets

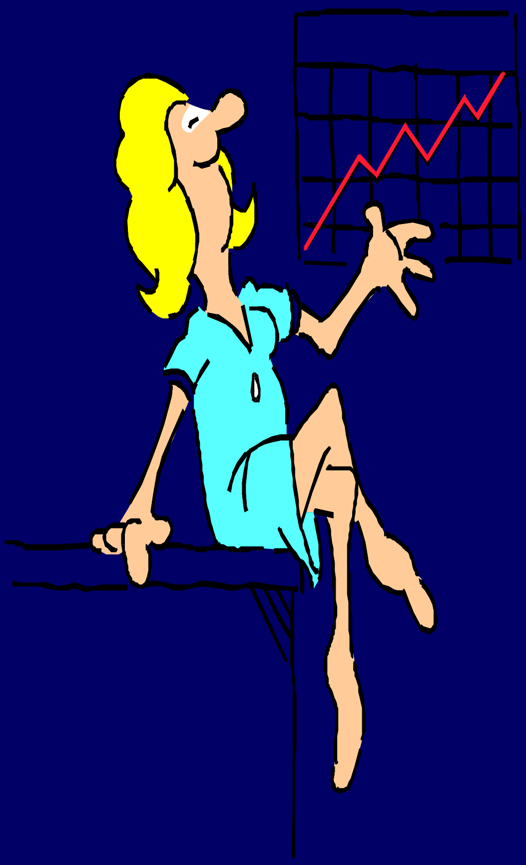

- Analysts need spreadsheets that support
  - pivot tables (cross-tabs)
  - drill-down and roll-up
  - slice and dice
  - sort
  - selections
  - derived attributes
- Popular in retail domain

OLAP - Data Cube

- Idea: analysts need to group data in many different ways
  - eg. Sales(region, product, prodtype, prodstyle, date, saleamount)
  - saleamount is a measure attribute, rest are dimension attributes
  - groupby every subset of the other attributes
    - materialize (precompute and store) groupbys to give online response
  - Also: hierarchies on attributes: date -> weekday,
    date -> month -> quarter -> year

SQL Extensions

- Front-end tools require
  - Extended Family of Aggregate Functions
    - rank, median, mode
  - Reporting Features
    - running totals, cumulative totals
  - Results of multiple group by
    - total sales by month and total sales by product
  - Data Cube

Relational OLAP: 3 Tier DSS

Data Warehouse | **ROLAP Engine** | **Decision Support Client**

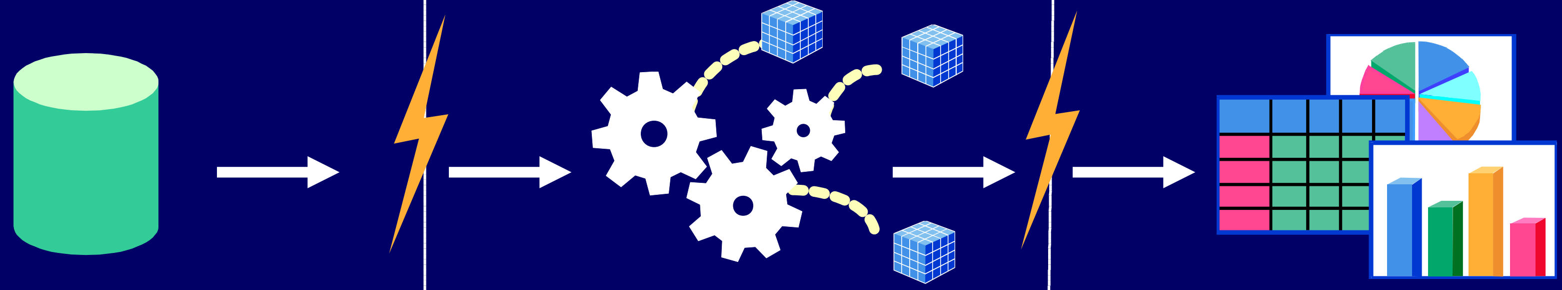

Database Layer	**Application Logic Layer**	**Presentation Layer**
Store atomic data in industry standard RDBMS.	Generate SQL execution plans in the ROLAP engine to obtain OLAP functionality.	Obtain multi-dimensional reports from the DSS Client.

MD-OLAP: 2 Tier DSS

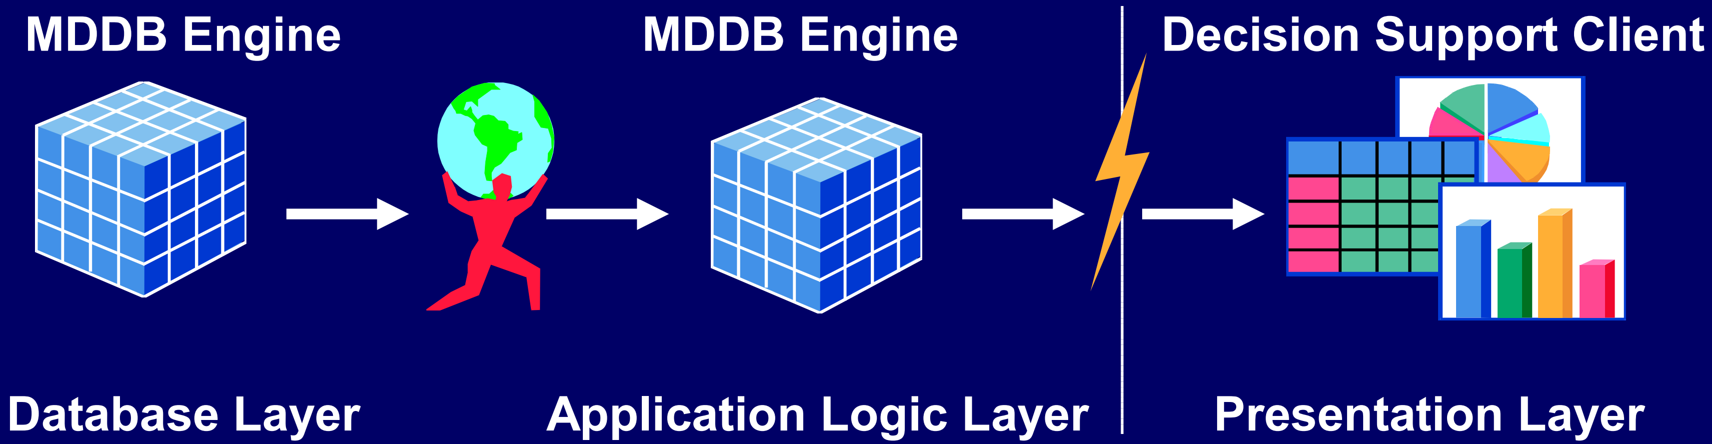

Store atomic data in a proprietary data structure (MDDB), pre-calculate as many outcomes as possible, obtain OLAP functionality via proprietary algorithms running against this data.

Obtain multi-dimensional reports from the DSS Client.

Typical OLAP Problems
Data Explosion

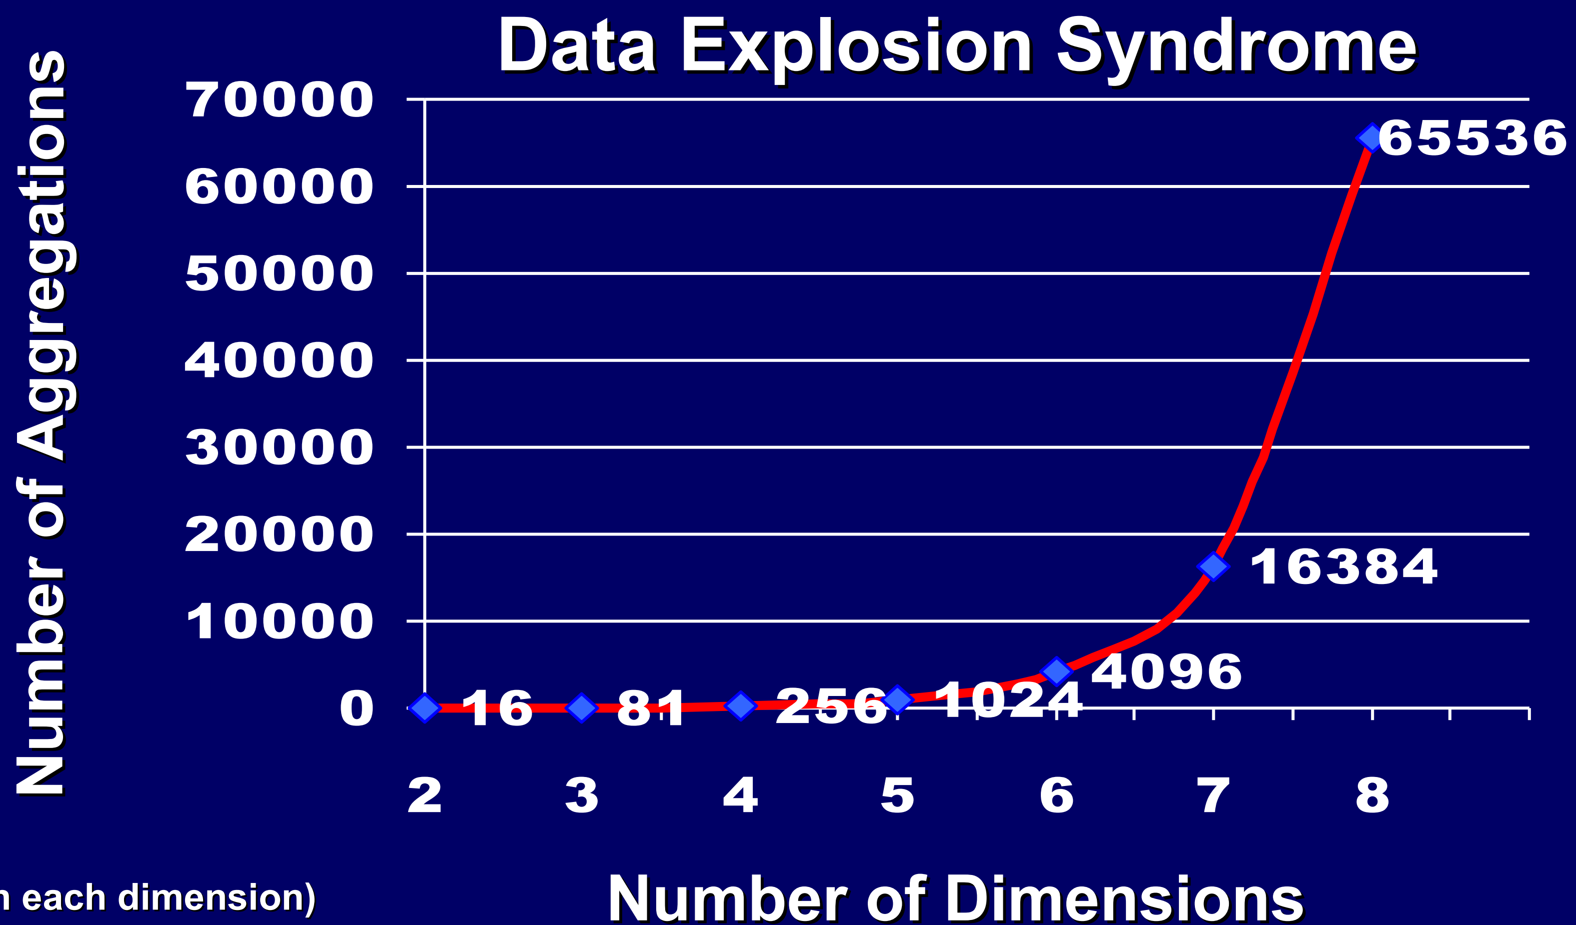

(4 levels in each dimension)

Metadata Repository

- Administrative metadata
  - source databases and their contents
  - gateway descriptions
  - warehouse schema, view & derived data definitions
  - dimensions, hierarchies
  - pre-defined queries and reports
  - data mart locations and contents
  - data partitions
  - data extraction, cleansing, transformation rules, defaults
  - data refresh and purging rules
  - user profiles, user groups
  - security: user authorization, access control

Metdata Repository .. 2

- Business data
  - business terms and definitions
  - ownership of data
  - charging policies
- operational metadata
  - data lineage: history of migrated data and sequence of transformations applied
  - currency of data: active, archived, purged
  - monitoring information: warehouse usage statistics, error reports, audit trails.

Recipe for a Successful Warehouse

For a Successful Warehouse

From Larry Greenfield, http://pwp.starnetinc.com/larryg/index.html

- From day one establish that warehousing is a joint user/builder project
- Establish that maintaining data quality will be an *ONGOING* joint user/builder responsibility
- Train the users one step at a time
- Consider doing a high level corporate data model in no more than three weeks

For a Successful Warehouse

Look closely at the data extracting, cleaning, and loading tools

Implement a user accessible automated directory to information stored in the warehouse

Determine a plan to test the integrity of the data in the warehouse

From the start get warehouse users in the habit of 'testing' complex queries

For a Successful Warehouse

- Coordinate system roll-out with network administration personnel
- When in a bind, ask others who have done the same thing for advice
- Be on the lookout for small, but strategic, projects
- Market and sell your data warehousing systems

Data Warehouse Pitfalls

- You are going to spend much time extracting, cleaning, and loading data
- Despite best efforts at project management, data warehousing project scope will increase
- You are going to find problems with systems feeding the data warehouse
- You will find the need to store data not being captured by any existing system
- You will need to validate data not being validated by transaction processing systems

Data Warehouse Pitfalls

Some transaction processing systems feeding the warehousing system will not contain detail

Many warehouse end users will be trained and never or seldom apply their training

After end users receive query and report tools, requests for IS written reports may increase

Your warehouse users will develop conflicting business rules

Large scale data warehousing can become an exercise in data homogenizing

Data Warehouse Pitfalls

- 'Overhead' can eat up great amounts of disk space
- The time it takes to load the warehouse will expand to the amount of the time in the available window... and then some
- Assigning security cannot be done with a transaction processing system mindset
- You are building a HIGH maintenance system
- You will fail if you concentrate on resource optimization to the neglect of project, data, and customer management issues and an understanding of what adds value to the customer

DW and OLAP Research Issues

- Data cleaning
  - focus on data inconsistencies, not schema differences
  - data mining techniques
- Physical Design
  - design of summary tables, partitions, indexes
  - tradeoffs in use of different indexes
- Query processing
  - selecting appropriate summary tables
  - dynamic optimization with feedback
  - acid test for query optimization: cost estimation, use of transformations, search strategies
  - partitioning query processing between OLAP server and backend server.

DW and OLAP Research Issues .. 2

Warehouse Management

- detecting runaway queries
- resource management
- incremental refresh techniques
- computing summary tables during load
- failure recovery during load and refresh
- process management: scheduling queries, load and refresh
- Query processing, caching
- use of workflow technology for process management

Products, References, Useful Links

Reporting Tools

Andyne Computing -- GQL
Brio -- BrioQuery
Business Objects -- Business Objects
Cognos -- Impromptu
Information Builders Inc. -- Focus for Windows
Oracle -- Discoverer2000
Platinum Technology -- SQL*Assist, ProReports
PowerSoft -- InfoMaker
SAS Institute -- SAS/Assist
Software AG -- Esperant
Sterling Software -- VISION:Data

OLAP and Executive Information Systems

- Andyne Computing -- Pablo
- Arbor Software -- Essbase
- Cognos -- PowerPlay
- Comshare -- Commander OLAP
- Holistic Systems -- Holos
- Information Advantage -- AXSYS, WebOLAP
- Informix -- Metacube
- Microstrategies --DSS/Agent
- Microsoft -- Plato
- Oracle -- Express
- Pilot -- LightShip
- Planning Sciences -- Gentium
- Platinum Technology -- ProdeaBeacon, Forest & Trees
- SAS Institute -- SAS/EIS, OLAP++
- Speedware -- Media

Other Warehouse Related Products

- Data extract, clean, transform, refresh
  - CA-Ingres replicator
  - Carleton Passport
  - Prism Warehouse Manager
  - SAS Access
  - Sybase Replication Server
  - Platinum Inforefiner, Infopump

Extraction and Transformation Tools

Carleton Corporation -- Passport

Evolutionary Technologies Inc. -- Extract

Informatica -- OpenBridge

Information Builders Inc. -- EDA Copy Manager

Platinum Technology -- InfoRefiner

Prism Solutions -- Prism Warehouse Manager

Red Brick Systems -- DecisionScape Formation

Scrubbing Tools

Apertus -- Enterprise/Integrator

Vality -- IPE

Postal Soft

Warehouse Products

Computer Associates -- CA-Ingres

Hewlett-Packard -- Allbase/SQL

Informix -- Informix, Informix XPS

Microsoft -- SQL Server

Oracle -- Oracle7, Oracle Parallel Server

Red Brick -- Red Brick Warehouse

SAS Institute -- SAS

Software AG -- ADABAS

Sybase -- SQL Server, IQ, MPP

Warehouse Server Products

- Oracle 8
- Informix
  - Online Dynamic Server
  - XPS --Extended Parallel Server
  - Universal Server for object relational applications
- Sybase
  - Adaptive Server 11.5
  - Sybase MPP
  - Sybase IQ

Warehouse Server Products

- Red Brick Warehouse
- Tandem Nonstop
- IBM
  - DB2 MVS
  - Universal Server
  - DB2 400
- Teradata

Other Warehouse Related Products

- Connectivity to Sources
  - Apertus
  - Information Builders EDA/SQL
  - Platimum Infohub
  - SAS Connect
  - IBM Data Joiner
  - Oracle Open Connect
  - Informix Express Gateway

Other Warehouse Related Products

- Query/Reporting Environments
  - Brio/Query
  - Cognos Impromptu
  - Informix Viewpoint
  - CA Visual Express
  - Business Objects
  - Platinum Forest and Trees

4GL's, GUI Builders, and PC Databases

Information Builders -- Focus

Lotus -- Approach

Microsoft -- Access, Visual Basic

MITI -- SQR/Workbench

PowerSoft -- PowerBuilder

SAS Institute -- SAS/AF

Data Mining Products

DataMind -- neurOagent

Information Discovery -- IDIS

SAS Institute -- SAS/Neuronets

Data Warehouse

W.H. Inmon, Building the Data Warehouse, Second Edition, John Wiley and Sons, 1996

W.H. Inmon, J. D. Welch, Katherine L. Glassey, Managing the Data Warehouse, John Wiley and Sons, 1997

Barry Devlin, Data Warehouse from Architecture to Implementation, Addison Wesley Longman, Inc 1997

Data Warehouse

W.H. Inmon, John A. Zachman, Jonathan G. Geiger, Data Stores Data Warehousing and the Zachman Framework, McGraw Hill Series on Data Warehousing and Data Management, 1997

Ralph Kimball, The Data Warehouse Toolkit, John Wiley and Sons, 1996

OLAP and DSS

Erik Thomsen, OLAP Solutions, John Wiley and Sons 1997

Microsoft TechEd Transparencies from Microsoft TechEd 98

Essbase Product Literature

Oracle Express Product Literature

Microsoft Plato Web Site

Microstrategy Web Site

Data Mining

Michael J.A. Berry and Gordon Linoff, Data Mining Techniques, John Wiley and Sons 1997

Peter Adriaans and Dolf Zantinge, Data Mining, Addison Wesley Longman Ltd. 1996

KDD Conferences

Other Tutorials

Donovan Schneider, Data Warehousing Tutorial, Tutorial at International Conference for Management of Data (SIGMOD 1996) and International Conference on Very Large Data Bases 97

Umeshwar Dayal and Surajit Chaudhuri, Data Warehousing Tutorial at International Conference on Very Large Data Bases 1996

Anand Deshpande and S. Seshadri, Tutorial on Datawarehousing and Data Mining, CSI-97

Useful URLs

Ralph Kimball's home page

http://www.rkimball.com

Larry Greenfield's Data Warehouse Information Center

http://pwp.starnetinc.com/larryg/

Data Warehousing Institute

http://www.dw-institute.com/

OLAP Council

http://www.olapcouncil.com/

Artificial Intelligence

UNIT-II
Of
Knowledge Management

Introduction

According to the father of Artificial Intelligence, John McCarthy, it is *“The science and engineering of making intelligent machines, especially intelligent computer programs”.*

Artificial Intelligence is a way

Introduction

AI is science & technology based on various disciplines as Computer Science, biology, psychology, linguistics, mathematics & engineering.

Major goal is to develop computer that can think, see, hear, walk, talk & feel.

Artificial intelligence is a science and technology based on disciplines such as Computer Science, Biology, Psychology, Linguistics, Mathematics, and Engineering. A major thrust of AI is in the development of computer functions associated with human intelligence, such as reasoning, learning, and problem solving.

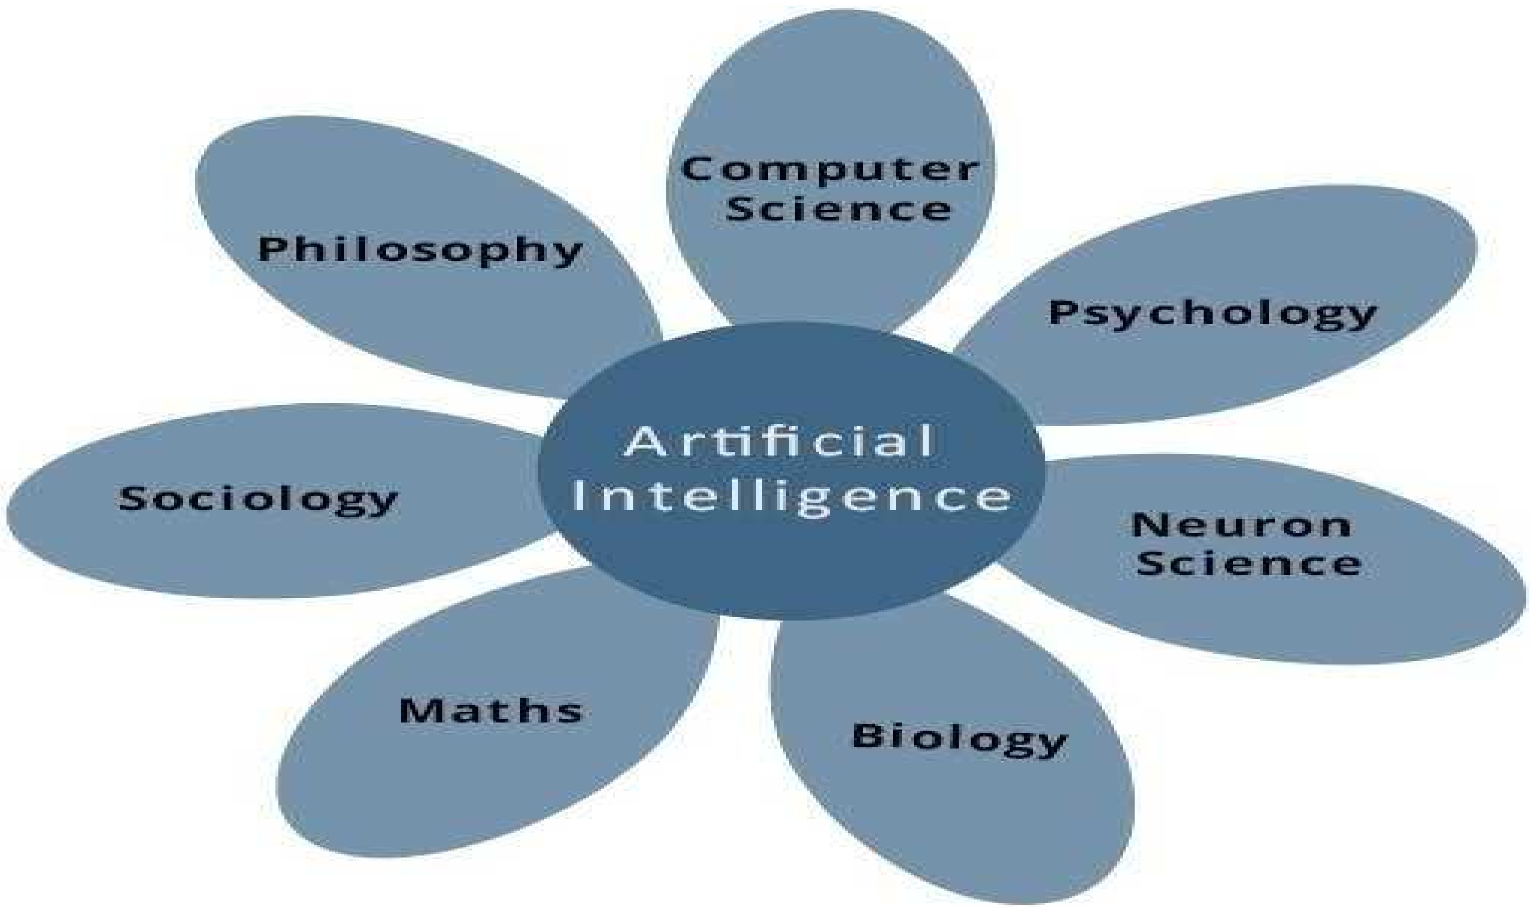

Programming Without and With AI

The programming without and with AI is different in following ways −

Programming Without AI	Programming With AI
A computer program without AI can answer the **specific** questions it is meant to solve.	A computer program with AI can answer the **generic** questions it is meant to solve.
Modification in the program leads to change in its structure.	AI programs can absorb new modifications by putting highly independent pieces of information together. Hence you can modify even a minute piece of information of program without affecting its structure.
Modification is not quick and easy. It may lead to affecting the program adversely.	Quick and Easy program modification.

Business & AI

Business & other organizations are striving to assist the human intelligence & productivity of knowledge workers with AI tools & techniques.

AI includes::

1. Natural Languages
2. Industrial Robots

Domains of AI

AI Applications can be grouped into three categories.

1. Cognitive Science
2. Robotics
3. Natural Interface

Cognitive Science

How the human brain works?

How human think & learn?

Result of such research is the basis to develop computer based applications of AI

Applications of Cognitive Science

1.Expert Systems

A CBIS uses its knowledge about a specific complex application area to act as an expert consultant to user.

System has knowledge base & software modules that performs inferences on knowledge to

Applications of Cognitive Science

2. Knowledge based Systems

AN IS which adds a knowledge base & some reasoning capabilities to the database & other components.

Applications of Cognitive Science

3. Adaptive Learning System

An IS that can modify its behavior based on information acquired as it operates.

Applications of Cognitive Science

4. Fuzzy Logic Systems

CBIS that can process data that are incomplete or partially correct & complete.

Can solve unstructured problem with incomplete knowledge by developing approx inferences and

Applications of Cognitive Science

5. Neural Networks

Can learn by processing sample problems & their solutions.

Capability to recognize patterns help to solve such problems on their own.

Applications of Cognitive Science

6. Genetic Algorithm

Uses mathematical techniques as randomizing to simulate processes that can generate increasing better solutions to the problems.

Applications of Cognitive Science 2

ROBOTICS

Basic disciplines of Robotics are AI, Engineering & mathematics.

This technology produces ROBOTS with computer intelligence & computer controlled humanlike physical capabilities

Applications of Cognitive Science 2

ROBOTICS applications include:

Visual perception
Tactility (Touch)
Dexterity (Skill in handling and manipulation)
Locomotion (Ability to move over any terrain)
Navigation (finding one’s way to a

Applications of Cognitive Science 3

Natural Interface:

Natural Languages & Speech recognition are major areas. Needs research & development in linguistics, psychology, computer science & other areas.

Efforts in this area include:

Natural Language, Multisensory

EXPERT SYSTEMS

A CBIS uses its knowledge about a specific complex application area to act as an expert consultant to user.

Inferences are performed on knowledge base & software modules to answer user queries.

Answers very specific questions

EXPERT SYSTEMS *Contd..*

Components of ES:

Knowledge Base

Software Resources

EXPERT SYSTEMS *Contd..*

Knowledge Base:

1. Facts about a specific subject area.
2. Heuristics that express the reasoning procedures of an expert on the subject.

EXPERT SYSTEMS *Contd..*

Software Resources:

1. Inference engine that processes the knowledge related to a specific problem.
2. User Interface to communicate with end user.

3. Explanation program for

EXPERT SYSTEMS *Contd..*

Hardware Resources:

1. Stand alone microcomputer system
2. Microcomputer workstations & terminals with full network
3. Special purpose computers

EXPERT SYSTEMS *Contd..*

People Resources:

1. Knowledge engineers

2. End Users

EXPERT SYSTEMS Applications

Using an ES means an interactive CB session in which:

1.Solution is explored as a wise consultant.

2.ES takes user questions, searches its KB for facts/rules/other knowledge.

3.Explain its reasoning process.

4.Gives expert advice to the user

Data Mining

BCA-VI

ITS Mohan Nagar

Data Mining

- New buzzword, old idea.
- Inferring new information from already collected data.
- Traditionally job of Data Analysts
- Computers have changed this.
  Far more efficient to comb through data using a machine than eyeballing statistical data.

- Data Mining is defined as extracting information from huge sets of data. In other words, we can say that data mining is the procedure of mining knowledge from data. The information or knowledge extracted so can be used for any of the following applications −
- Market Analysis
- Fraud Detection
- Customer Retention
- Production Control
- Science Exploration

Data Mining – Two Main Components

- Wikipedia definition: “Data mining is the entire process of applying computer-based methodology, including new techniques for knowledge discovery, from data.”
- **Knowledge Discovery**
  Concrete information gleaned from known data. Data you may not have known, but which is supported by recorded facts.
- **Knowledge Prediction**
  Uses known data to forecast future trends, events, etc. (ie: Stock market predictions)
- Wikipedia note: "some data mining systems such as neural networks are inherently geared towards prediction and pattern recognition, rather than knowledge discovery.“ These include applications in AI and Symbol analysis

Data Mining vs. Data Analysis

- In terms of software and the marketing thereof Data Mining != Data Analysis
- Data Mining implies software uses some intelligence over simple grouping and partitioning of data to infer new information.
- Data Analysis is more in line with standard statistical software (ie: web stats). These usually present information about subsets and relations within the recorded data set (ie: browser/search engine usage, average visit time, etc.)

Propositional vs. Relational Data

- Old data mining methods relied on Propositional Data, or data that was related to a single, central element, that could be represented in a vector format. (ie: the purchasing history of a single user. Amazon uses such vectors in its related item suggestions [a multidimensional dot product])
- Current, advanced data mining methods rely on Relational Data, or data that can be stored and modeled easily through use of relational databases. An example of this would be data used to represent interpersonal relations.
- Relational Data is more interesting than Propositional data to miners in the sense that an entity, and all the entities to which it is related, factor into the data inference process.

Key Component of Data Mining

- Whether Knowledge Discovery or Knowledge Prediction, data mining takes information that was once quite difficult to detect and presents it in an easily understandable format (ie: graphical or statistical)
- Data mining Techniques involve sophisticated algorithms, including Decision Tree Classifications, Association detection, and Clustering.
- Since Data mining is not on test, I will keep things superficial.

Uses of Data Mining

- **AI/Machine Learning**
  Combinatorial/Game Data Mining
  Good for analyzing winning strategies to games, and thus developing intelligent AI opponents. (ie: Chess)
- **Business Strategies**
  Market Basket Analysis
  Identify customer demographics, preferences, and purchasing patterns.
- **Risk Analysis**
  Product Defect Analysis
  Analyze product defect rates for given plants and predict possible complications (read: lawsuits) down the line.

Uses of Data Mining (Continued)

- **User Behavior Validation**
  Fraud Detection
  In the realm of cell phones
  Comparing phone activity to calling records.
  Can help detect calls made on cloned phones.

  Similarly, with credit cards, comparing purchases with historical purchases. Can detect activity with stolen cards.

Uses of Data Mining (Continued)

- **Health and Science**
  Protein Folding
  Predicting protein interactions and functionality within biological cells. Applications of this research include determining causes and possible cures for Alzheimers, Parkinson's, and some cancers (caused by protein "misfolds")

  Extra-Terrestrial Intelligence
  Scanning Satellite receptions for possible transmissions from other planets.
- For more information see Stanford’s Folding@home and SETI@home projects. Both involve participation in a widely distributed computer application.

Sources of Data for Mining

- Databases (most obvious)
- Text Documents
- Computer Simulations
- Social Networks

Privacy Concerns

- Mining of public and government databases is done, though people have, and continue to raise concerns.

- Wiki quote:
  "data mining gives information that would not be available otherwise. It must be properly interpreted to be useful. When the data collected involves individual people, there are many questions concerning privacy, legality, and ethics."

Prevalence of Data Mining

- Your data is already being mined, whether you like it or not.
- Many web services require that you allow access to your information [for data mining] in order to use the service.
- Google mines email data in Gmail accounts to present account owners with ads.
- Facebook requires users to allow access to info from non-Facebook pages. Facebook privacy policy:
  "We may use information about you that we collect from other sources, including but not limited to newspapers and Internet sources such as blogs, instant messaging services and other users of Facebook, to supplement your profile.
- This allows access to your blog RSS feed (rather innocuous), as well as information obtained through partner sites (worthy of concern).

Data Mining Techniques

There are several major ***data mining techniques*** have been developing and using in data mining projects recently including *association*, *classification*, *clustering*, *prediction, sequential patterns* and *decision tree*.

Data Mining Techniques Association

- Association is one of the best-known data mining technique. In association, a pattern is discovered based on a relationship between items in the same transaction. That's is the reason why association technique is also known as *relation technique*. The association technique is used in *market basket analysis* to identify a set of products that customers frequently purchase together.

- Retailers are using association technique to research customer's buying habits. Based on historical sale data, retailers might find out that customers always buy butter when they buy breads, and, therefore, they can put these items next to each other to save time for customer and increase sales.

Data Mining Techniques Classification

- Based on machine learning.
- Used to classify each item in a set of data into one of a predefined set of classes or groups.
- Makes use of mathematical techniques such as decision trees, linear programming, neural network and statistics.
- We develop the software that can learn how to classify the data items into groups.
- For example, we can apply classification in the application that "given all records of employees who left the company, predict who will probably leave the company in a future period." In this case, we divide the records of employees into two groups that named "leave" and "stay". And then we can ask our data mining software to classify the employees into separate groups.

Data Mining Techniques Clustering

- Clustering is a data mining technique that makes a meaningful or useful cluster of objects which have similar characteristics using the automatic technique.
- The clustering technique defines the classes and puts objects in each class, while in the classification techniques, objects are assigned into predefined classes.
- Book management in the library. Wide range of books on various topics available. The challenge is how to keep those books in a way that readers can take several books on a particular topic without hassle. By using the clustering technique, we can keep books that have some kinds of similarities in one cluster or one shelf and label it with a meaningful name. If readers want to grab books in that topic, they would only have to go to that shelf instead of looking for the entire library.

Data Mining Techniques Prediction

One of the data mining techniques that discovers the relationship between independent variables and relationship between dependent and independent variables.

For instance, the prediction analysis technique can be used in the sale to predict profit for the future if we consider the sale is an independent variable, profit could be a dependent variable. Then based on the historical sale and profit data, we can draw a fitted regression curve that is used for profit prediction.

Data Mining Techniques Sequential Patterns

- One of data mining technique that seeks to discover or identify similar patterns, regular events or trends in transaction data over a business period.

- In sales, with historical transaction data, businesses can identify a set of items that customers buy together different times in a year. Then businesses can use this information to recommend customers buy it with better deals based on their purchasing frequency in the past.

Data Mining Techniques Decision Trees

- One of the most common used data mining techniques because its model is easy to understand for users.
- In decision tree technique, the root of the decision tree is a simple question or condition that has multiple answers. Each answer then leads to a set of questions or conditions that help us determine the data so that we can make the final decision based on it. For example, We use the following decision tree to determine whether or not to play tennis:

Data Mining Techniques Decision Trees

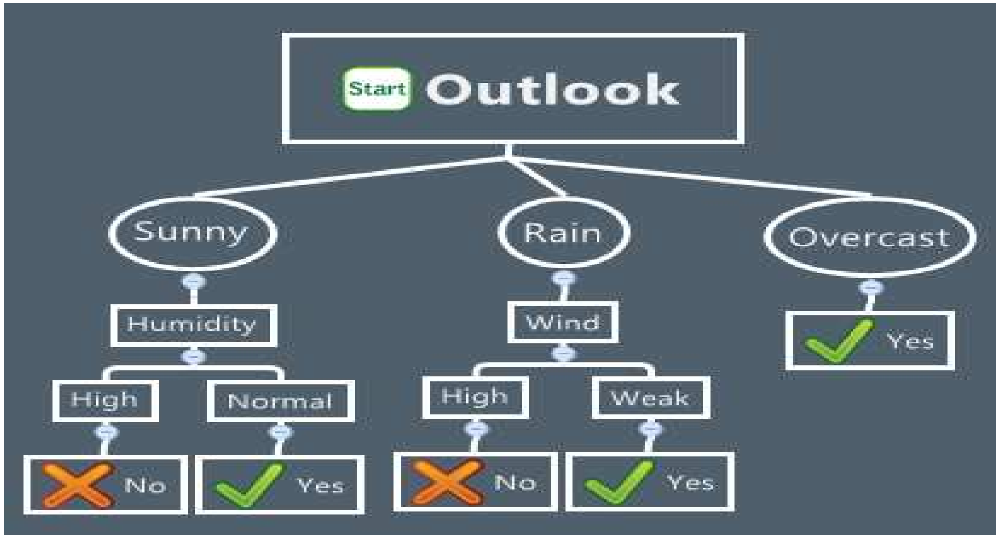

Starting at the root node, if the outlook is overcast then we should definitely play tennis. If it is rainy, we should only play tennis if the wind is the week. And if it is sunny then we should play tennis in case the humidity is normal.

Data Mining Techniques Link Analysis

Link analysis

In network theory, link analysis is a data-analysis technique used to evaluate relationships (connections) between nodes. Relationships may be identified among various types of nodes (objects), including organizations, people and transactions. Link analysis has been used for investigation of criminal activity (fraud detection, counterterrorism, and intelligence), computer security analysis, search engine optimization, market research, medical research, and art.

Data Mining Techniques Memory Based Reasoning MBR

- Memory-Based Reasoning(MBR) is
  - Identifying similar cases from experience
  - Applying the information from these cases to the problem at hand.
  - MBR finds neighbors similar to a new record and uses the neighbors for classification and prediction.
- It cares about the existence of two operations
  - Distance function ; assigns a distance between any two records
  - Combination function ; combines the results from the neighbors to arrive at an answer.
- Applications of MBR span many areas;
  - Fraud detection
  - Customer response prediction
  - Medical treatments
  - Classifying responses

Data Mining Techniques

We often combine two or more of those data mining techniques together to form an appropriate process that meets the business needs.

Data Mining works with Warehouse Data

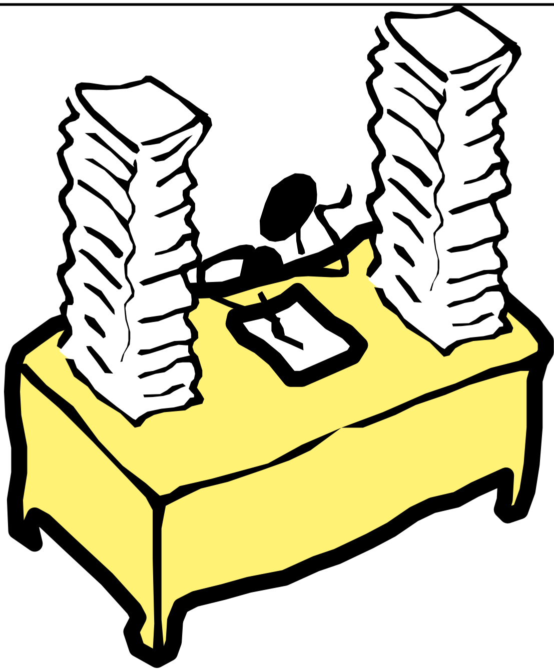

- Data Warehousing provides the Enterprise with a memory

Data Mining provides the Enterprise with intelligence

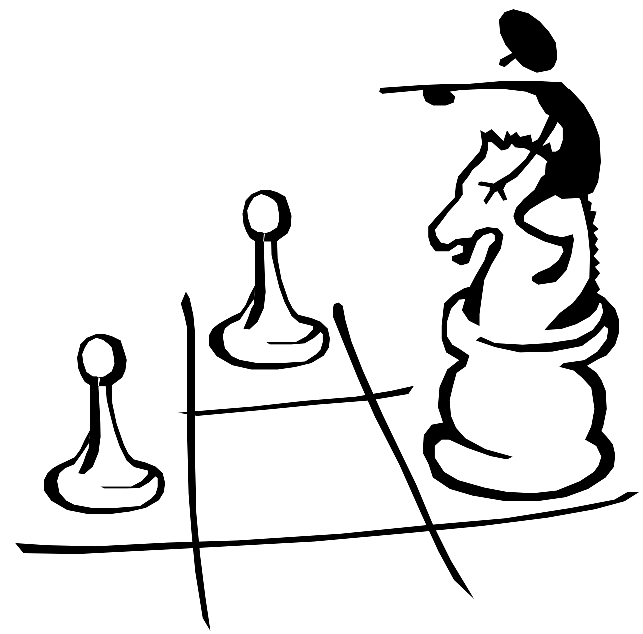

We want to know ...

- Given a database of 100,000 names, which persons are the least likely to default on their credit cards?
- Which types of transactions are likely to be fraudulent given the demographics and transactional history of a particular customer?
- If I raise the price of my product by Rs. 2, what is the effect on my ROI?
- If I offer only 2,500 airline miles as an incentive to purchase rather than 5,000, how many lost responses will result?
- If I emphasize ease-of-use of the product as opposed to its technical capabilities, what will be the net effect on my revenues?
- Which of my customers are likely to be the most loyal?

Data Mining helps extract such information

Data Mining in Use

- The US Government uses Data Mining to track fraud
- A Supermarket becomes an information broker
- Basketball teams use it to track game strategy
- Cross Selling
- Warranty Claims Routing
- Holding on to Good Customers
- Weeding out Bad Customers

What makes data mining possible?

- Advances in the following areas are making data mining deployable:
  - data warehousing
  - better and more data (i.e., operational, behavioral, and demographic)
  - the emergence of easily deployed data mining tools and
  - the advent of new data mining techniques.
    - -- Gartner Group

UNIT II

DATA WAREHOUSES

INTRODUCTION

- Organizations need business intelligence
- ***Business intelligence (BI)*** – knowledge about your customers, competitors, business partners, competitive environment, and internal operations to make effective, important, and strategic business decisions

INTRODUCTION

- IT tools help process information to create business intelligence according to:
  - OLTP
  - OLAP

INTRODUCTION

- ***Online transaction processing (OLTP)*** – the gathering of input information, processing that information, and updating existing information to reflect the gathered and processed information
  - Databases support OLTP
  - ***Operational database*** – databases that support OLTP

INTRODUCTION

- ***Online analytical processing (OLAP)*** – the manipulation of information to support decision making
  - Databases can support some OLAP
  - Data warehouses only support OLAP, not OLTP
  - Data warehouses are special forms of databases that support decision making

INTRODUCTION

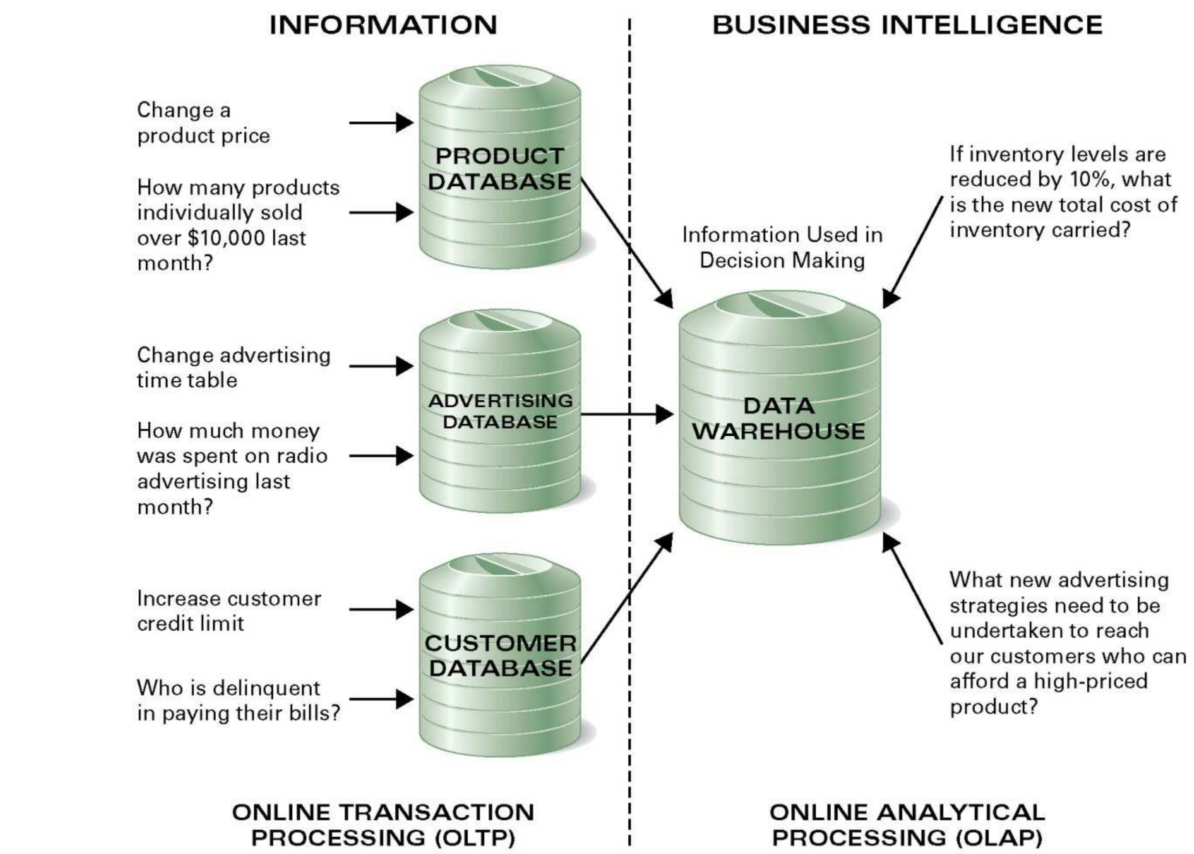

Figure 3.1

Building Business Intelligence

Understanding a Data Warehouse

- A data warehouse is a database, which is kept separate from the organization's operational database.
- There is no frequent updating done in a data warehouse.
- It possesses consolidated historical data, which helps the organization to analyze its business.
- A data warehouse helps executives to organize, understand, and use their data to take strategic decisions.
- Data warehouse systems help in the integration of diversity of application systems.
- A data warehouse system helps in consolidated historical data analysis.

Data Warehouse Features

The key features of a data warehouse are discussed below:

- **Subject Oriented** - A data warehouse is subject oriented because it provides information around a subject rather than the organization's ongoing operations. These subjects can be product, customers, suppliers, sales, revenue, etc. A data warehouse does not focus on the ongoing operations, rather it focuses on modelling and analysis of data for decision making.

- **Integrated** - A data warehouse is constructed by integrating data from heterogeneous sources such as relational databases, flat files, etc. This integration enhances the effective analysis of data.

- **Time Variant** - The data collected in a data warehouse is identified with a particular time period. The data in a data warehouse provides information from the historical point of view.

- **Non-volatile** - Non-volatile means the previous data is not erased when new data is added to it. A data warehouse is kept separate from the operational database and therefore frequent changes in operational database is not reflected in the data warehouse.

What Is a Data Warehouse?

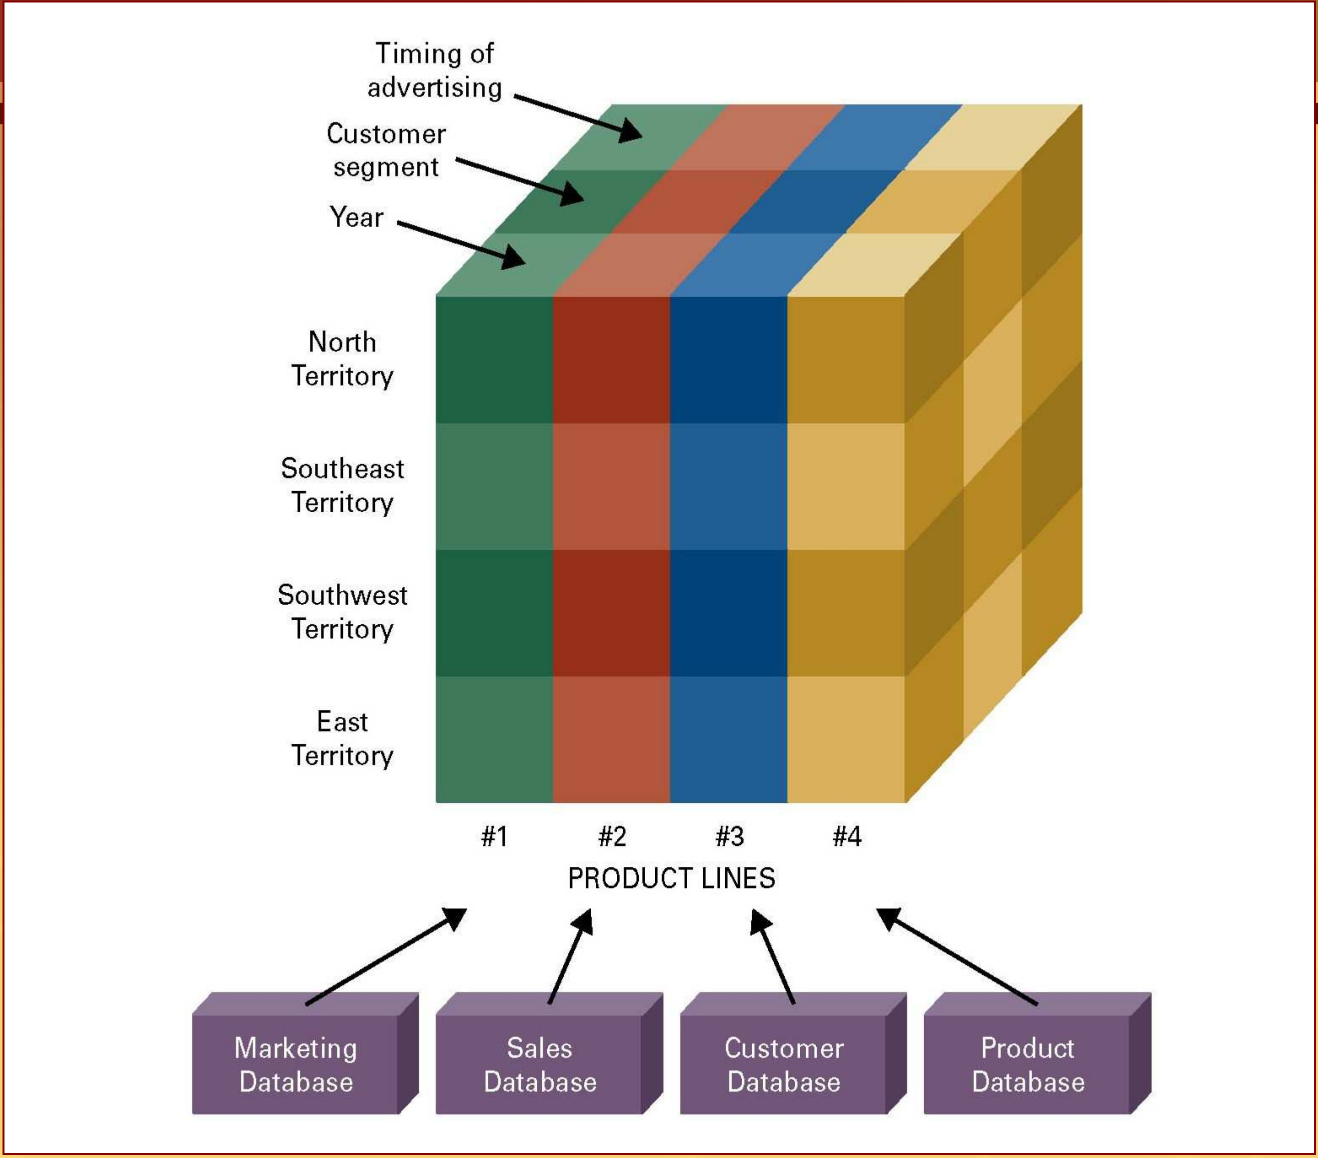

What Is a Data Warehouse?

- Multidimensional
- Rows and columns
- Also layers
- Many times called *hypercubes*
- What are the dimensions in Figure 3.8 on page 142?

What Are Data-Mining Tools?

- ***Data-mining tools*** – software tools that you use to query information in a data warehouse
  - Query-and-reporting tools
  - Intelligence agents
  - Multidimensional analysis tools
  - Statistical tools

What Are Data-Mining Tools?

Figure 3.9

The Data Miner's Tool Set

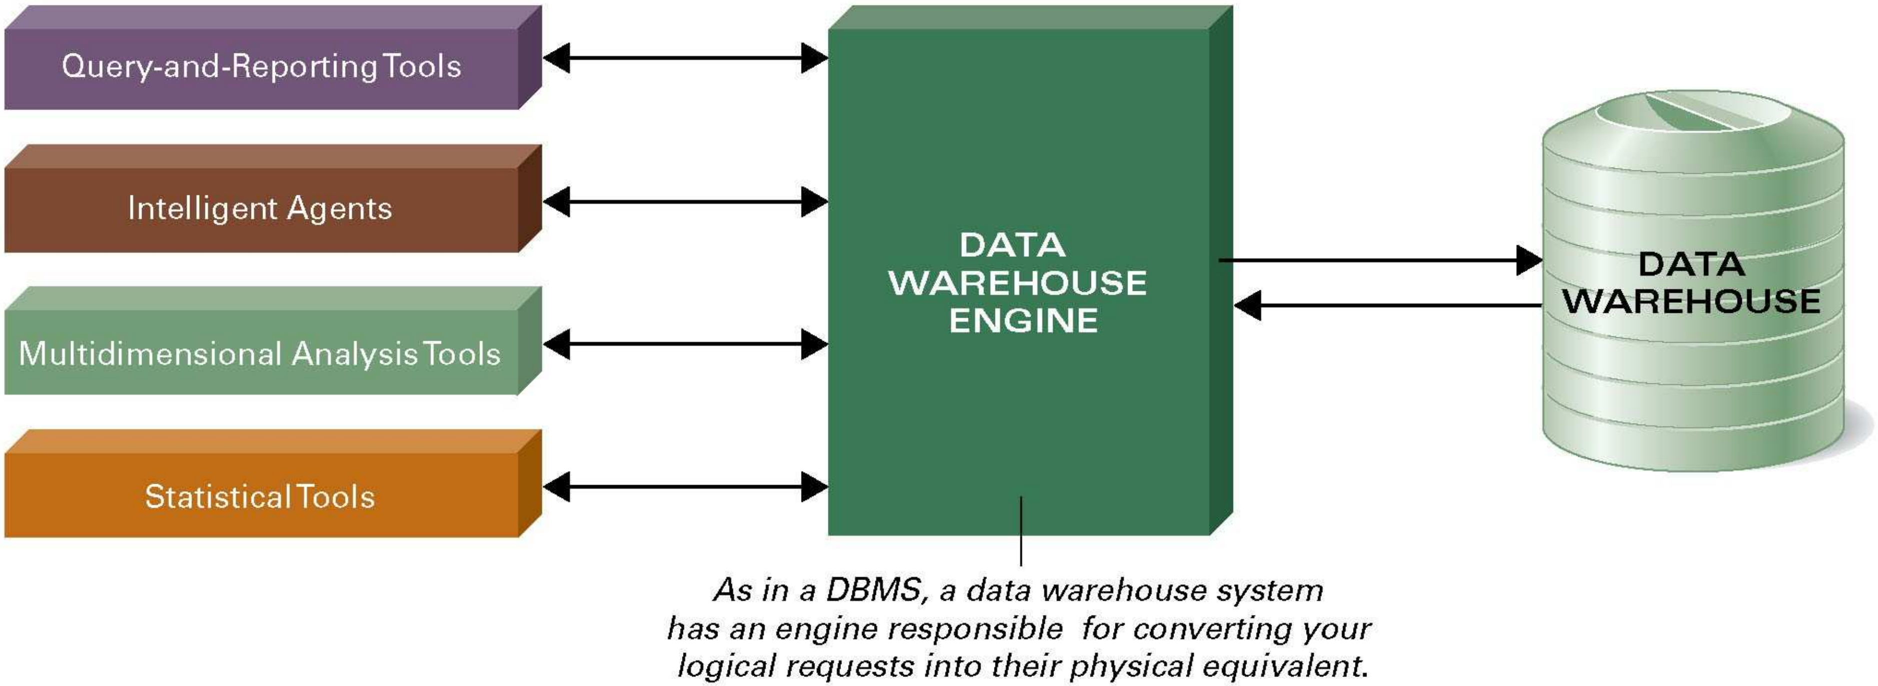

Query-And-Reporting Tools

- ***Query-and-reporting tools*** – similar to QBE tools, SQL, and report generators in the typical database environment

Intelligent Agents

- Use various artificial intelligence tools such as neural networks and fuzzy logic to form the basis for "information discovery" and building business intelligence
- Help you find hidden patterns in information
- Chapter 4 focuses more on these

Multidimensional Analysis Tools

- ***Multidimensional analysis (MDA) tools*** – slice-and-dice techniques that allow you to view multidimensional information from different perspectives
  - Bring new layers to the front
  - Reorganize rows and columns

Statistical Tools

- Help you apply various mathematical models to the information stored in a data warehouse to discover new information
  - Regression
  - Analysis of variance
  - And so on

Data Marts

- Data warehouses can support all of an organization's information
- Data marts have subsets of an organizationwide data warehouse
- ***Data mart*** – subset of a data warehouse in which only a focused portion of the data warehouse information is kept

Data Marts

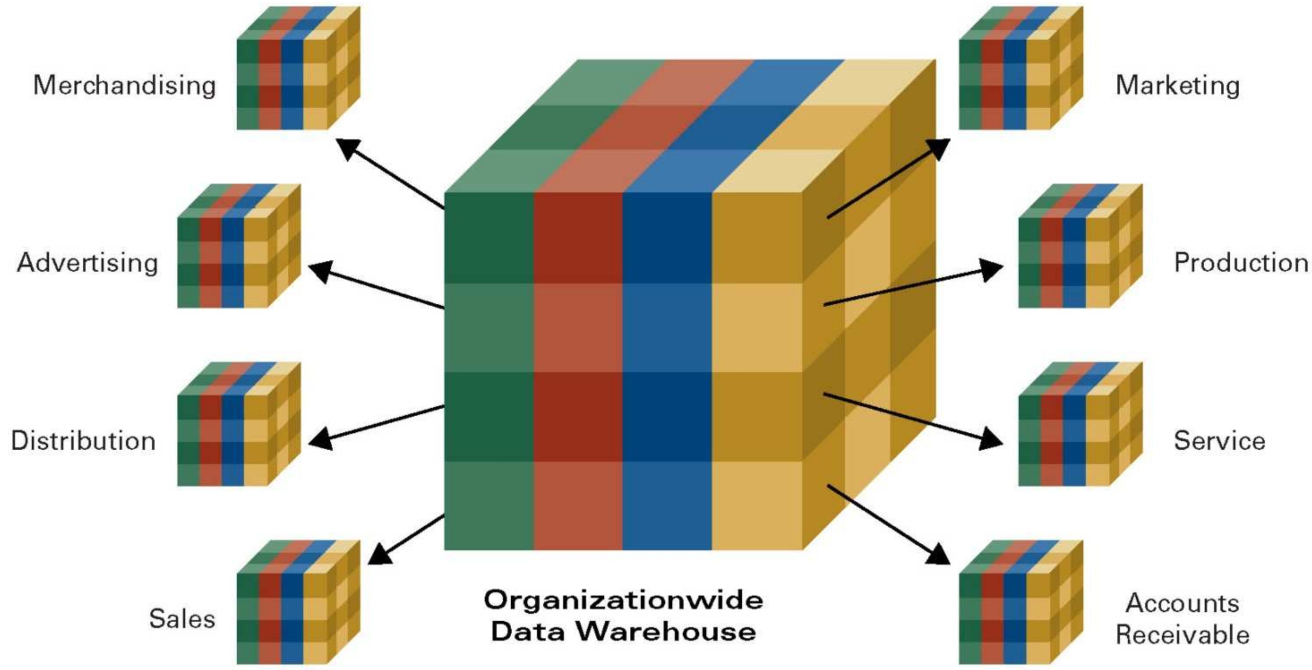

Figure 3.10

Data Marts Are Subsets of Data Warehouses

Data Mining as a Career Opportunity

- Knowledge of data mining can be a substantial career opportunity for you
  - Query and Analysis and Enterprise Analytic Tools (Business Objects)
  - Business Intelligence and Information Access tools (SAS)
  - Many in Cognos (the data warehouse leader)
  - PowerAnalyzer (Informatica)

Considerations in Using a Data Warehouse

- Do you need a data warehouse?
  - Perhaps database OLAP is sufficient
- Do all employees need the entire data warehouse?
  - If no, build smaller data marts
- How up-to-date must the information be?
- What data-mining tools do you need?

MANAGING THE INFORMATION RESOURCE

- Information is an organizational resource
- Just like people, capital, and equipment
- It must be managed effectively

MANAGING THE INFORMATION RESOURCE

- Who should oversee your organization’s information resource?
  - ***Chief information officer (CIO)*** – oversees an organization’s information resource
  - ***Data administration*** – plans for, oversees the development of, and monitors the information resource
  - ***Database administration*** – technical and operational aspects of managing information

MANAGING THE INFORMATION RESOURCE

- Is information ownership a consideration?
  - If you create information, you “own” it
  - You will also share it with others
  - Because you “own” it, you are responsible for its quality

MANAGING THE INFORMATION RESOURCE

- How "clean" must your information be?
  - Duplicate information (records) must be eliminated
  - Inaccurate information must be corrected
  - Information forms the basis of business intelligence
  - If your business intelligence is bad, you will make poor decisions

On-line Schematic

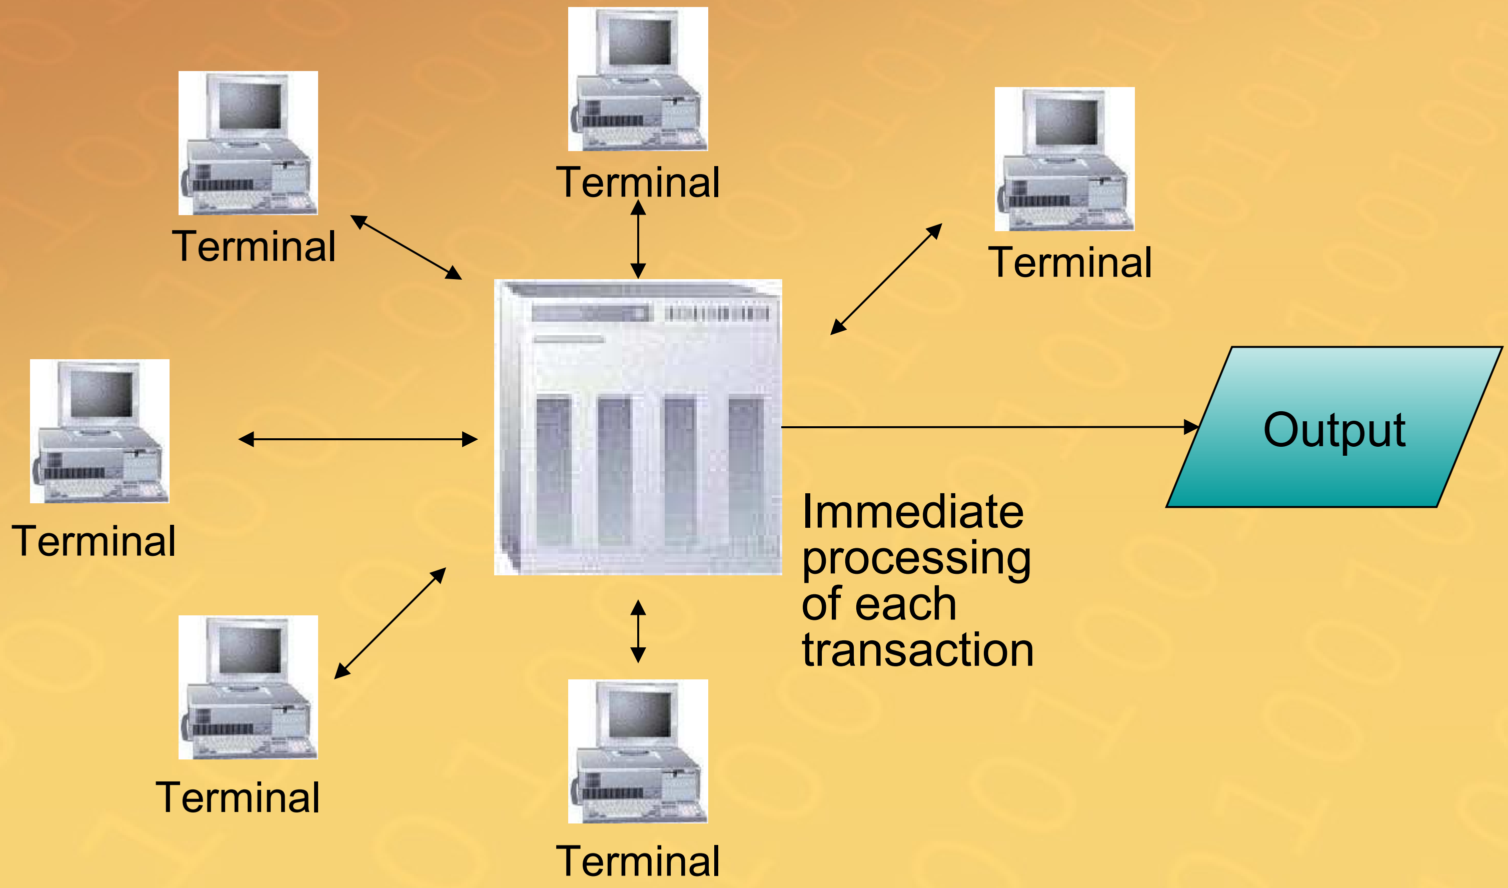

What OLTP basically is..?

Sr.No.	Data Warehouse (OLAP)	Operational Database(OLTP)
1	It involves historical processing of information.	It involves day-to-day processing.
2	OLAP systems are used by knowledge workers such as executives, managers, and analysts.	OLTP systems are used by clerks, DBAs, or database professionals.
3	It is used to analyze the business.	It is used to run the business.
4	It focuses on Information out.	It focuses on Data in.
5	It is based on Star Schema, Snowflake Schema, and Fact Constellation Schema.	It is based on Entity Relationship Model.
6	It focuses on Information out.	It is application oriented.
7	It contains historical data.	It contains current data.
8	It provides summarized and consolidated data.	It provides primitive and highly detailed data.
9	It provides summarized and multidimensional view of data.	It provides detailed and flat relational view of data.
10	The number of users is in hundreds.	The number of users is in thousands.

What OLTP basically is..?

11	The number of records accessed is in millions.	The number of records accessed is in tens.
12	The database size is from 100GB to 100 TB.	The database size is from 100 MB to 100 GB.
13	These are highly flexible.	It provides high performance.

What OLTP basically is..?

In a TPS, End users have online access to the system and to enterprise data, and directly initiate transactions.

Many users repeatedly process similar transactions, and require a fast response to each transaction such as order entry clerks, airline reservation clerks, or bank tellers.

They share an environment of programs and data.

What OLTP basically is..?

In a typical TPS:

· Many end users run the same or similar transactions, sharing the same databases and files.

· The system can schedule transactions on the basis of priority attributes.

· The transactions are invoked by online input and generate online output.

· The transactions are designed for a good user interface and fast response times.

What OLTP basically is..?

Transaction processing is an effective solution when end users want to:

· Process unscheduled single items in unpredictable volumes and sequence.
· Have immediate access to enterprise data that has been updated to reflect all previous transactions
· Change enterprise data immediately to reflect each transaction as it is processed.

What OLTP basically is..?

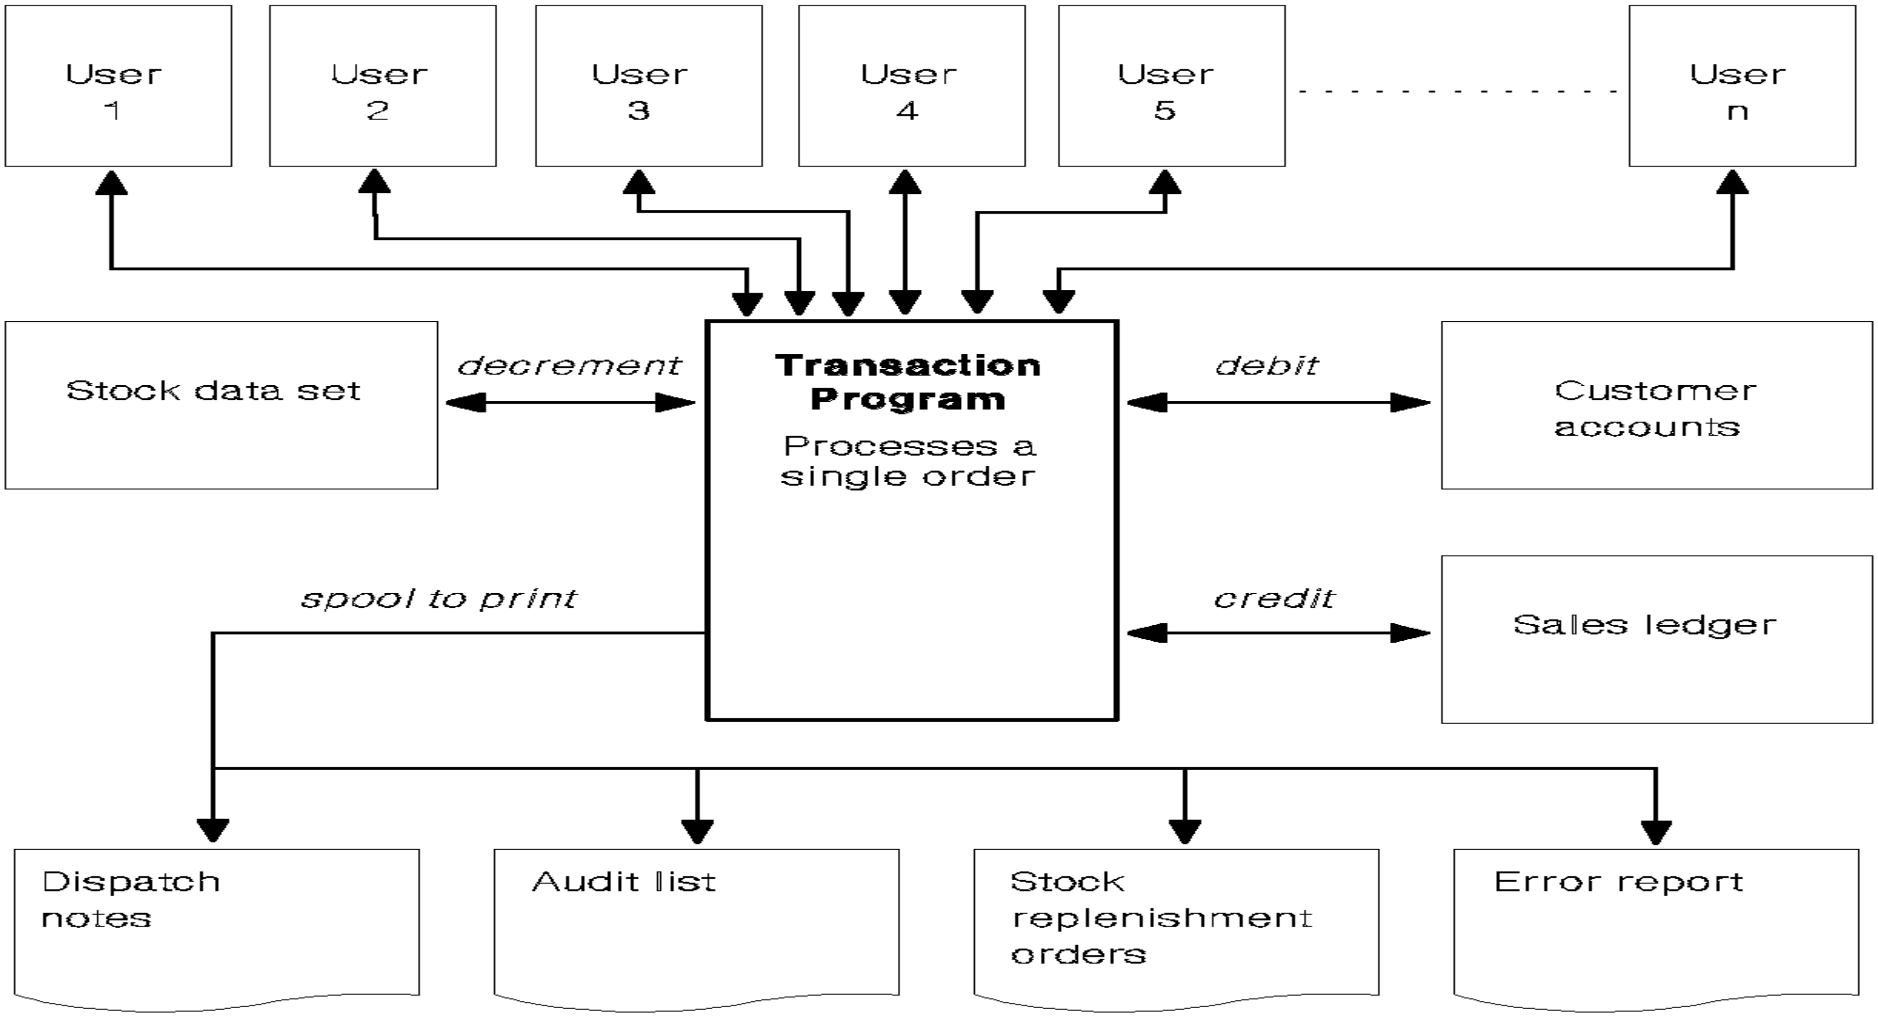

What OLTP basically is..?

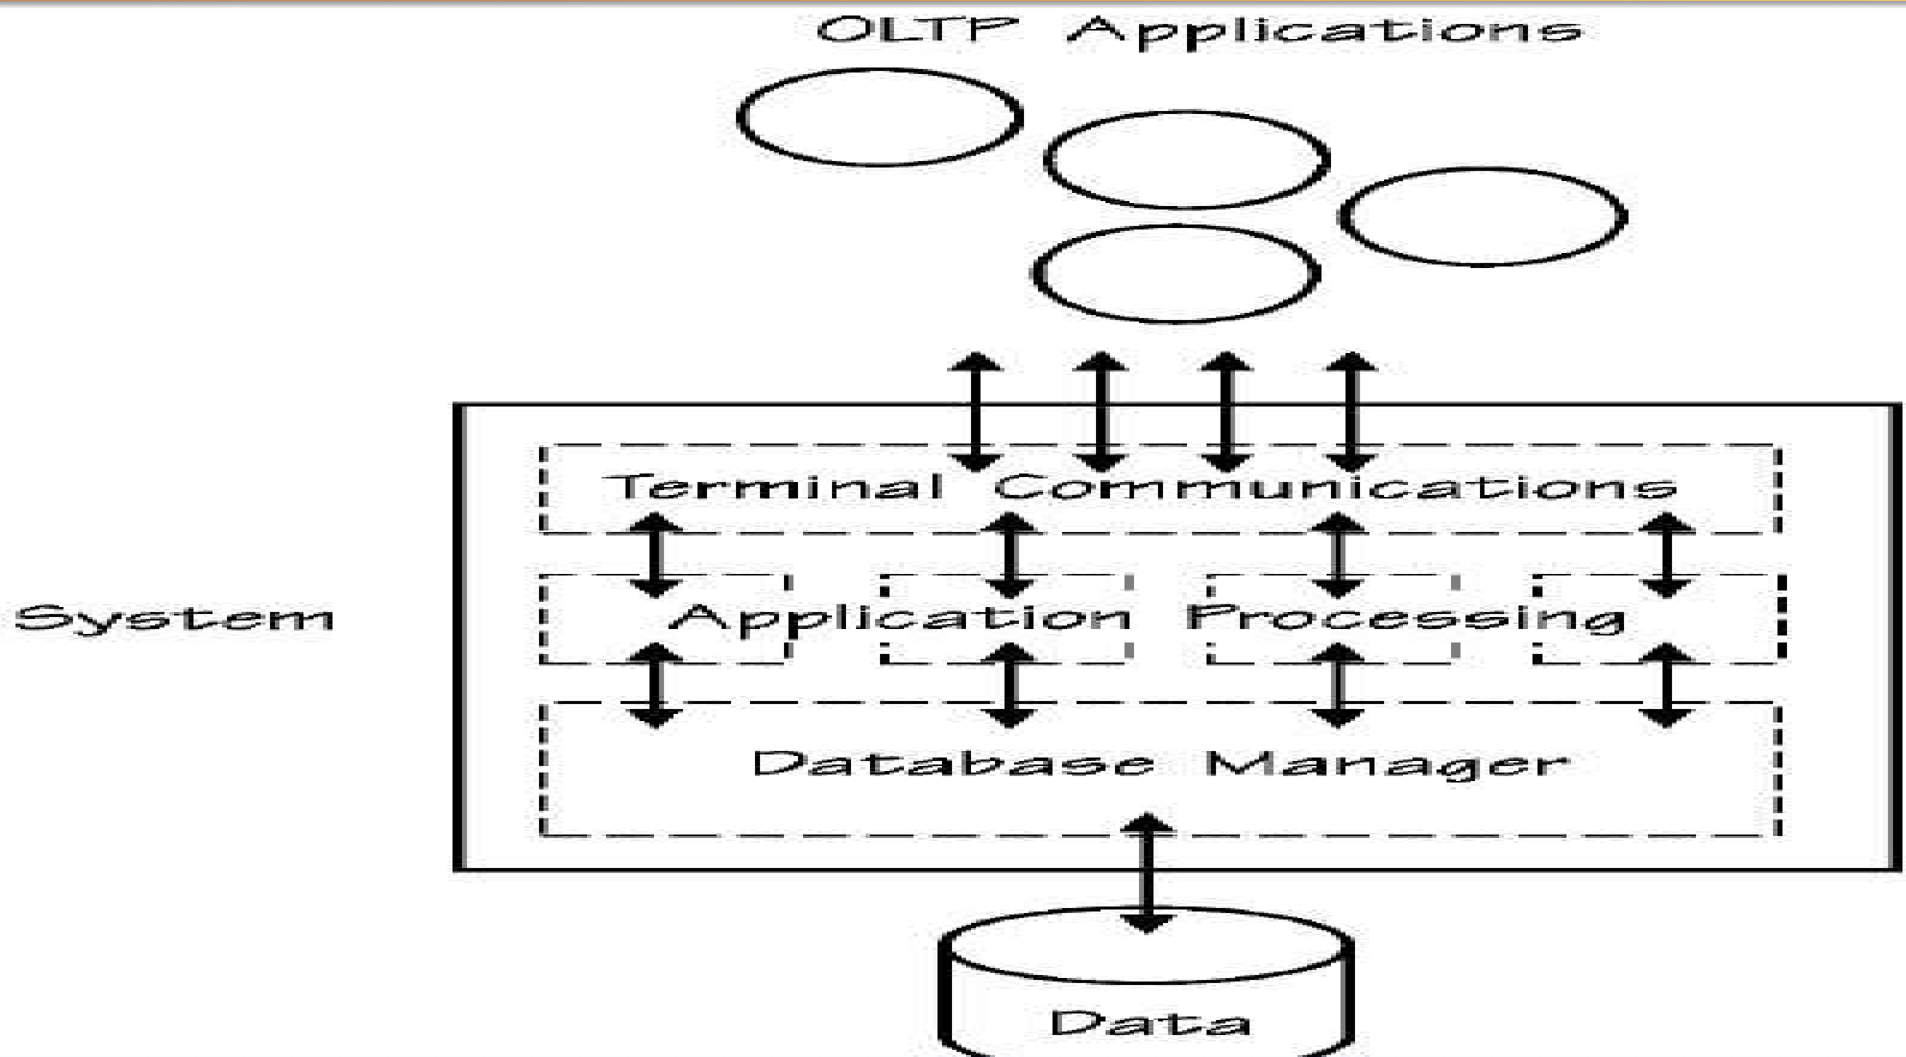

Warehouse Products

- Computer Associates -- CA-Ingres
- Hewlett-Packard -- Allbase/SQL
- Informix -- Informix, Informix XPS
- Microsoft -- SQL Server
- Oracle -- Oracle7, Oracle Parallel Server
- Red Brick -- Red Brick Warehouse
- SAS Institute -- SAS
- Software AG -- ADABAS
- Sybase -- SQL Server, IQ, MPP

Warehouse Server Products

- Oracle 8
- Informix
  - Online Dynamic Server
  - XPS --Extended Parallel Server
  - Universal Server for object relational applications
- Sybase
  - Adaptive Server 11.5
  - Sybase MPP
  - Sybase IQ

Warehouse Server Products

- Red Brick Warehouse
- Tandem Nonstop
- IBM
  - DB2 MVS
  - Universal Server
  - DB2 400
- Teradata

Data Mining Products

- DataMind -- neurOagent
- Information Discovery -- IDIS
- SAS Institute -- SAS/Neuronets

Recipe for a Successful Warehouse

For a Successful Warehouse

From Larry Greenfield, http://pwp.starnetinc.com/larryg/index.html

- From day one establish that warehousing is a joint user/builder project
- Establish that maintaining data quality will be an *ONGOING* joint user/builder responsibility
- Train the users one step at a time
- Consider doing a high level corporate data model in no more than three weeks

For a Successful Warehouse

- Look closely at the data extracting, cleaning, and loading tools
- Implement a user accessible automated directory to information stored in the warehouse
- Determine a plan to test the integrity of the data in the warehouse
- From the start get warehouse users in the habit of 'testing' complex queries

For a Successful Warehouse

- Coordinate system roll-out with network administration personnel
- When in a bind, ask others who have done the same thing for advice
- Be on the lookout for small, but strategic, projects
- Market and sell your data warehousing systems

Data Warehouse Pitfalls

- You are going to spend much time extracting, cleaning, and loading data
- Despite best efforts at project management, data warehousing project scope will increase
- You are going to find problems with systems feeding the data warehouse
- You will find the need to store data not being captured by any existing system
- You will need to validate data not being validated by transaction processing systems

Data Warehouse Pitfalls

- Some transaction processing systems feeding the warehousing system will not contain detail
- Many warehouse end users will be trained and never or seldom apply their training
- After end users receive query and report tools, requests for IS written reports may increase
- Your warehouse users will develop conflicting business rules
- Large scale data warehousing can become an exercise in data homogenizing

II. On-Line Analytical Processing (OLAP)

Making Decision Support Possible

Limitations of SQL

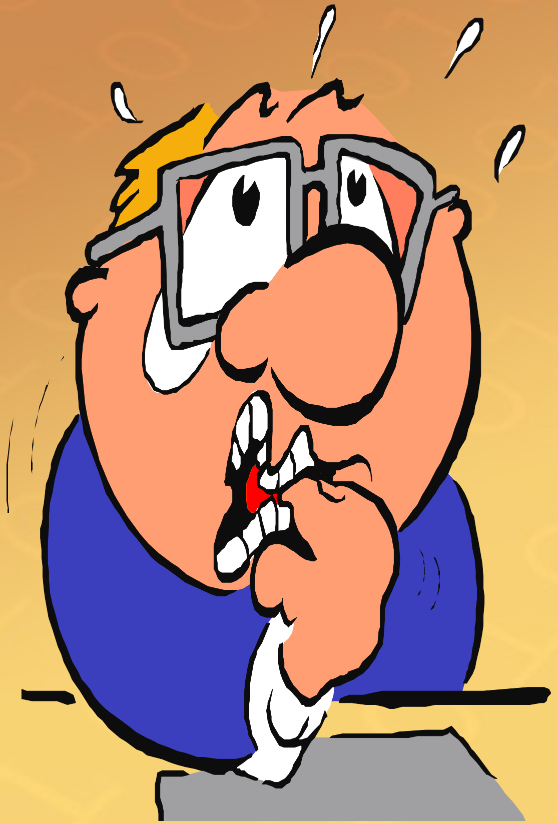

"A Freshman in Business needs a Ph.D. in SQL"

-- Ralph Kimball

Typical OLAP Queries

- Write a multi-table join to compare sales for each product line YTD this year vs. last year.
- Repeat the above process to find the top 5 product contributors to margin.
- Repeat the above process to find the sales of a product line to new vs. existing customers.
- Repeat the above process to find the customers that have had negative sales growth.

What Is OLAP?

- Online Analytical Processing - coined by EF Codd in 1994 paper contracted by Arbor Software*
- Generally synonymous with earlier terms such as Decisions Support, Business Intelligence, Executive Information System
- OLAP = Multidimensional Database
- MOLAP: Multidimensional OLAP (Arbor Essbase, Oracle Express)
- ROLAP: Relational OLAP (Informix MetaCube, Microstrategy DSS Agent)

*** Reference: http://www.arborsoft.com/essbase/wht_ppr/coddTOC.html**

Data Warehouse vs. Data Marts

What comes first

From the Data Warehouse to Data Marts

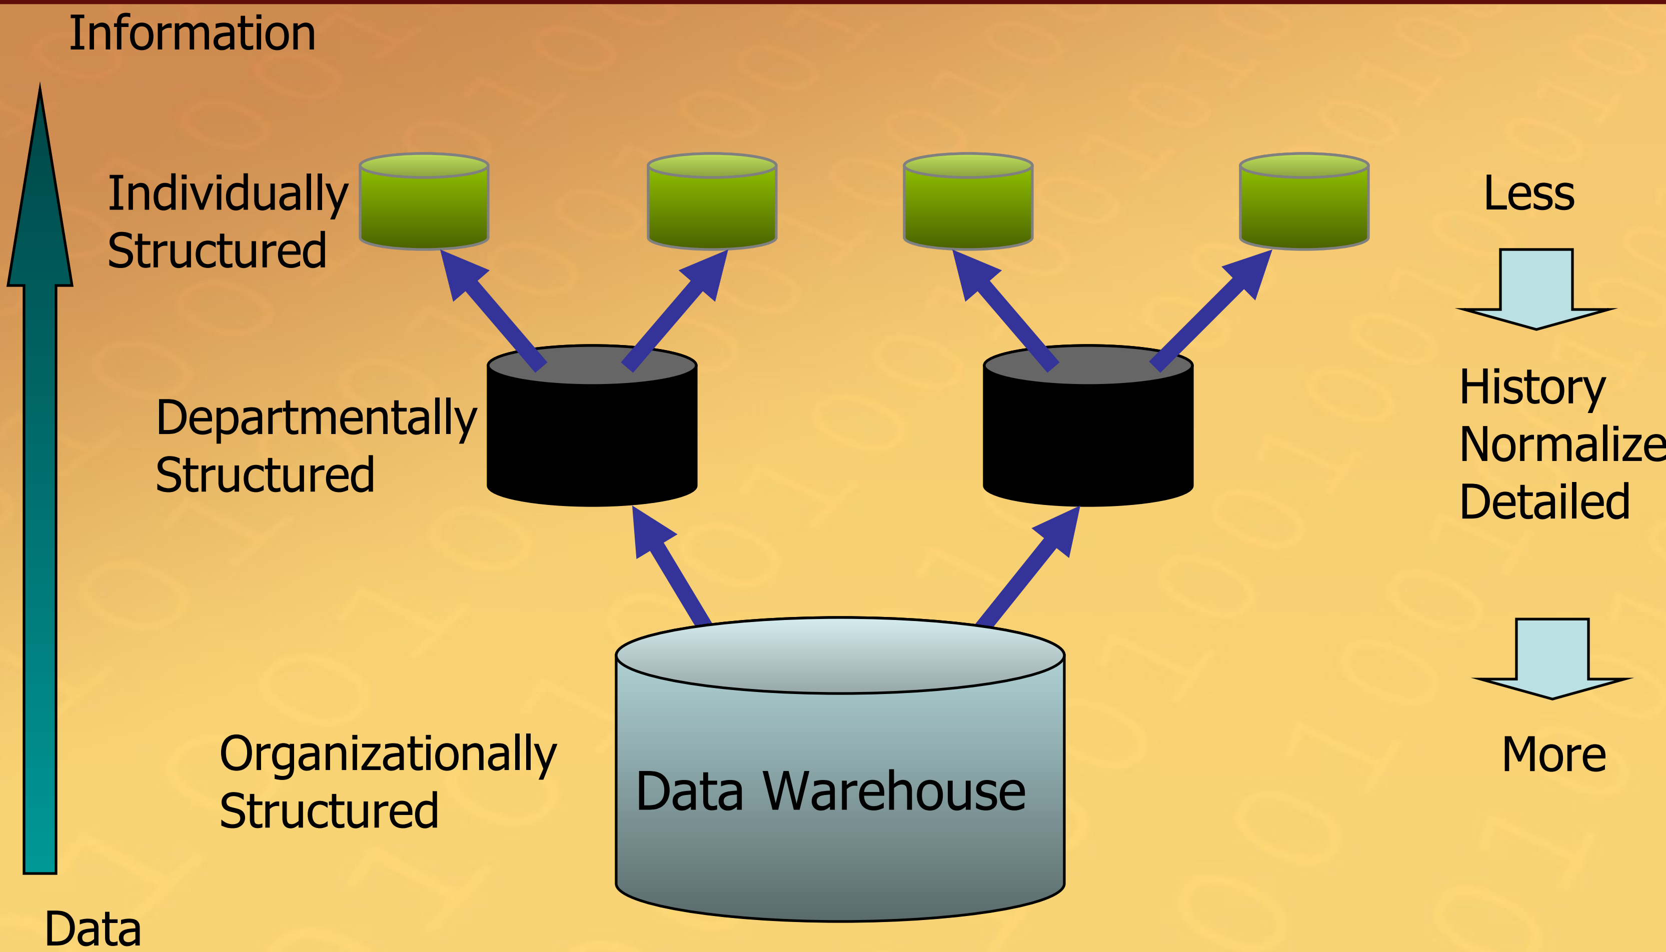

Data Warehouse and Data Marts

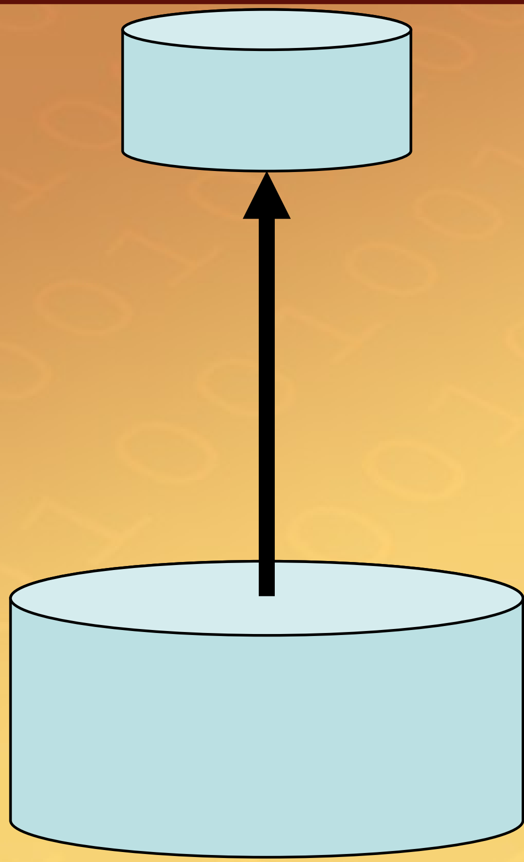

OLAP
Data Mart
Lightly summarized
Departmentally structured

Organizationally structured
Atomic
Detailed Data Warehouse Data

Characteristics of the Departmental Data Mart

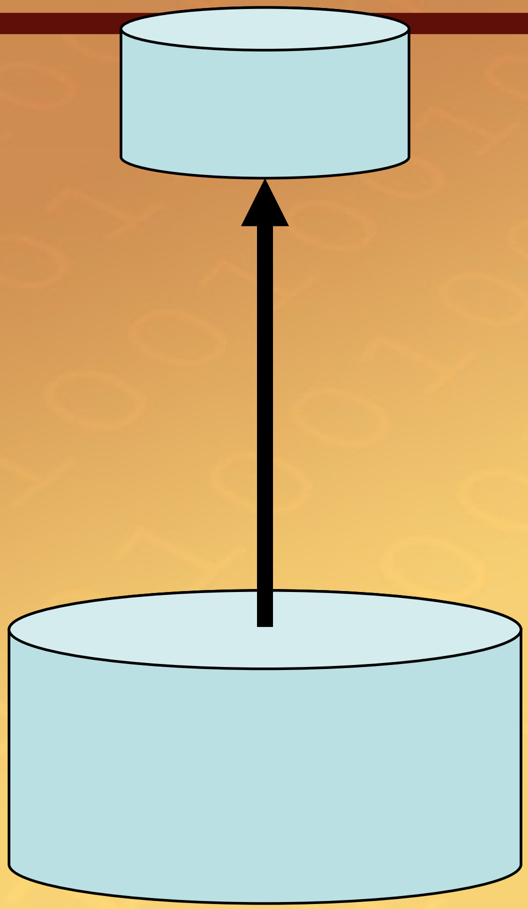

- OLAP
- Small
- Flexible
- Customized by Department
- Source is departmentally structured data warehouse

Techniques for Creating Departmental Data Mart

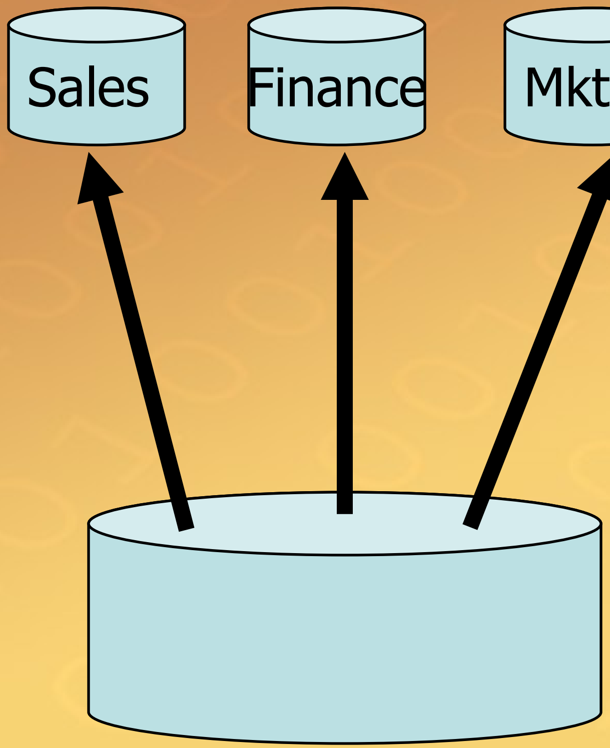

- OLAP
- Subset
- Summarized
- Superset
- Indexed
- Arrayed

Data Mart Centric

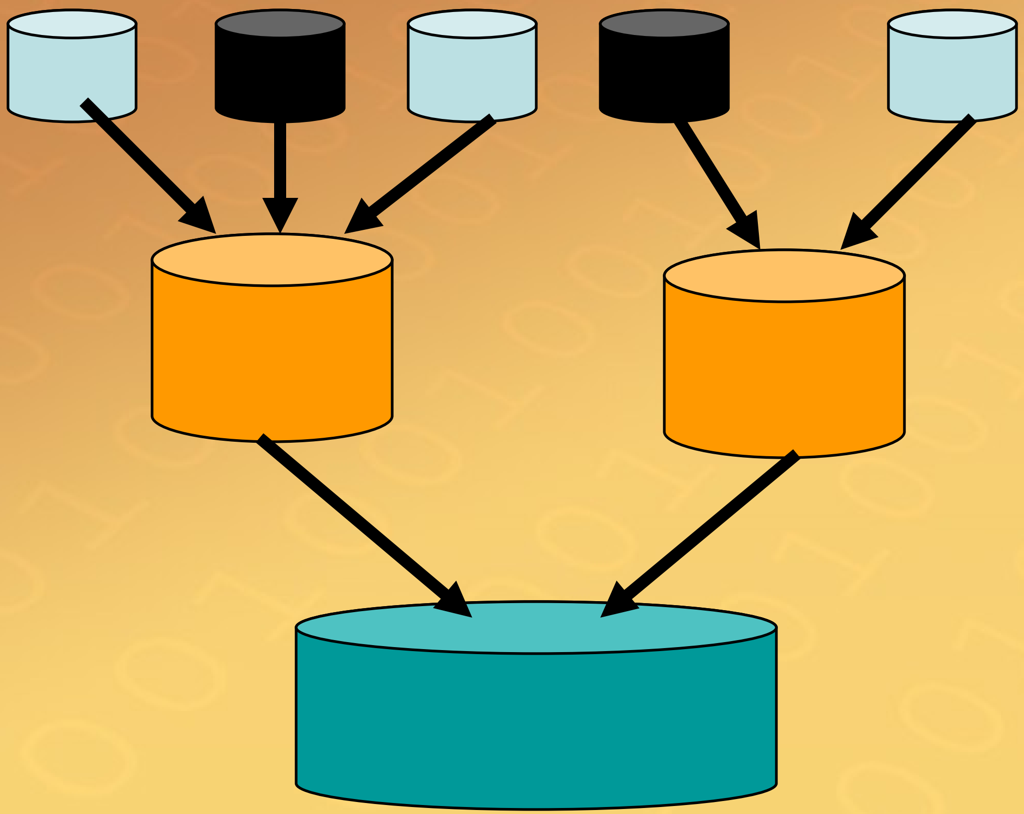

Data Sources

Data Marts

Data Warehouse

Problems with Data Mart Centric Solution

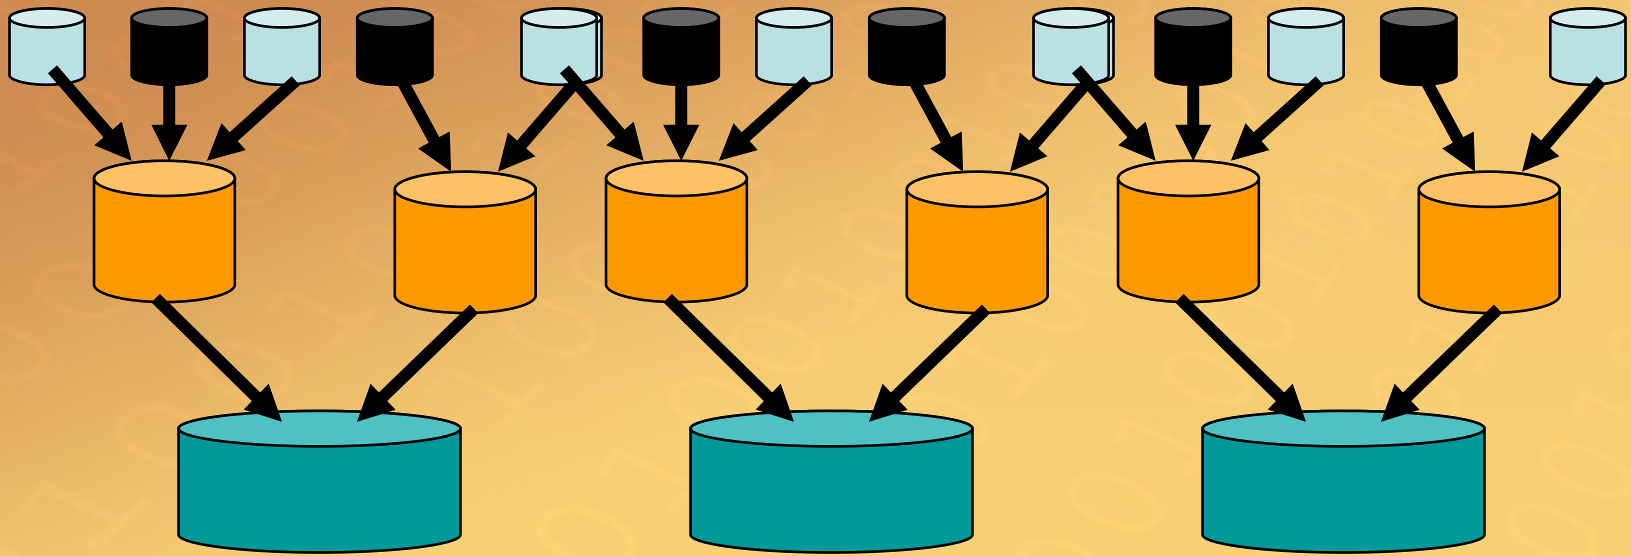

If you end up creating multiple warehouses, integrating them is a problem

True Warehouse

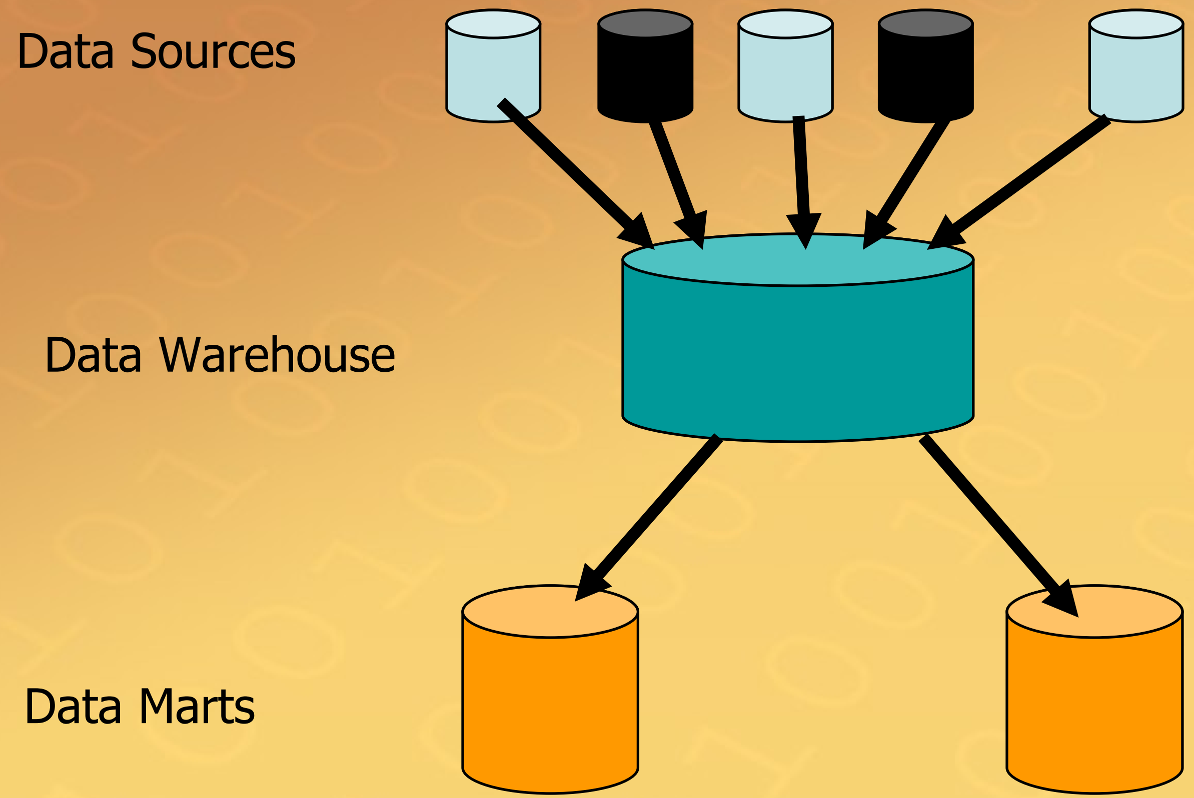

Examples of Operational Data

Data	Industry	Usage	Technology	Volumes
Customer File	All	Track Customer Details	Legacy application, flat files, main frames	Small-medium
Account Balance	Finance	Control account activities	Legacy applications, hierarchical databases, mainframe	Large
Point-of-Sale data	Retail	Generate bills, manage stock	ERP, Client/Server, relational databases	Very Large
Call Record	Telecomm-unications	Billing	Legacy application, hierarchical database, mainframe	Very Large
Production Record	Manufact-uring	Control Production	ERP, relational databases, AS/400	Medium

Data Warehouse Architecture

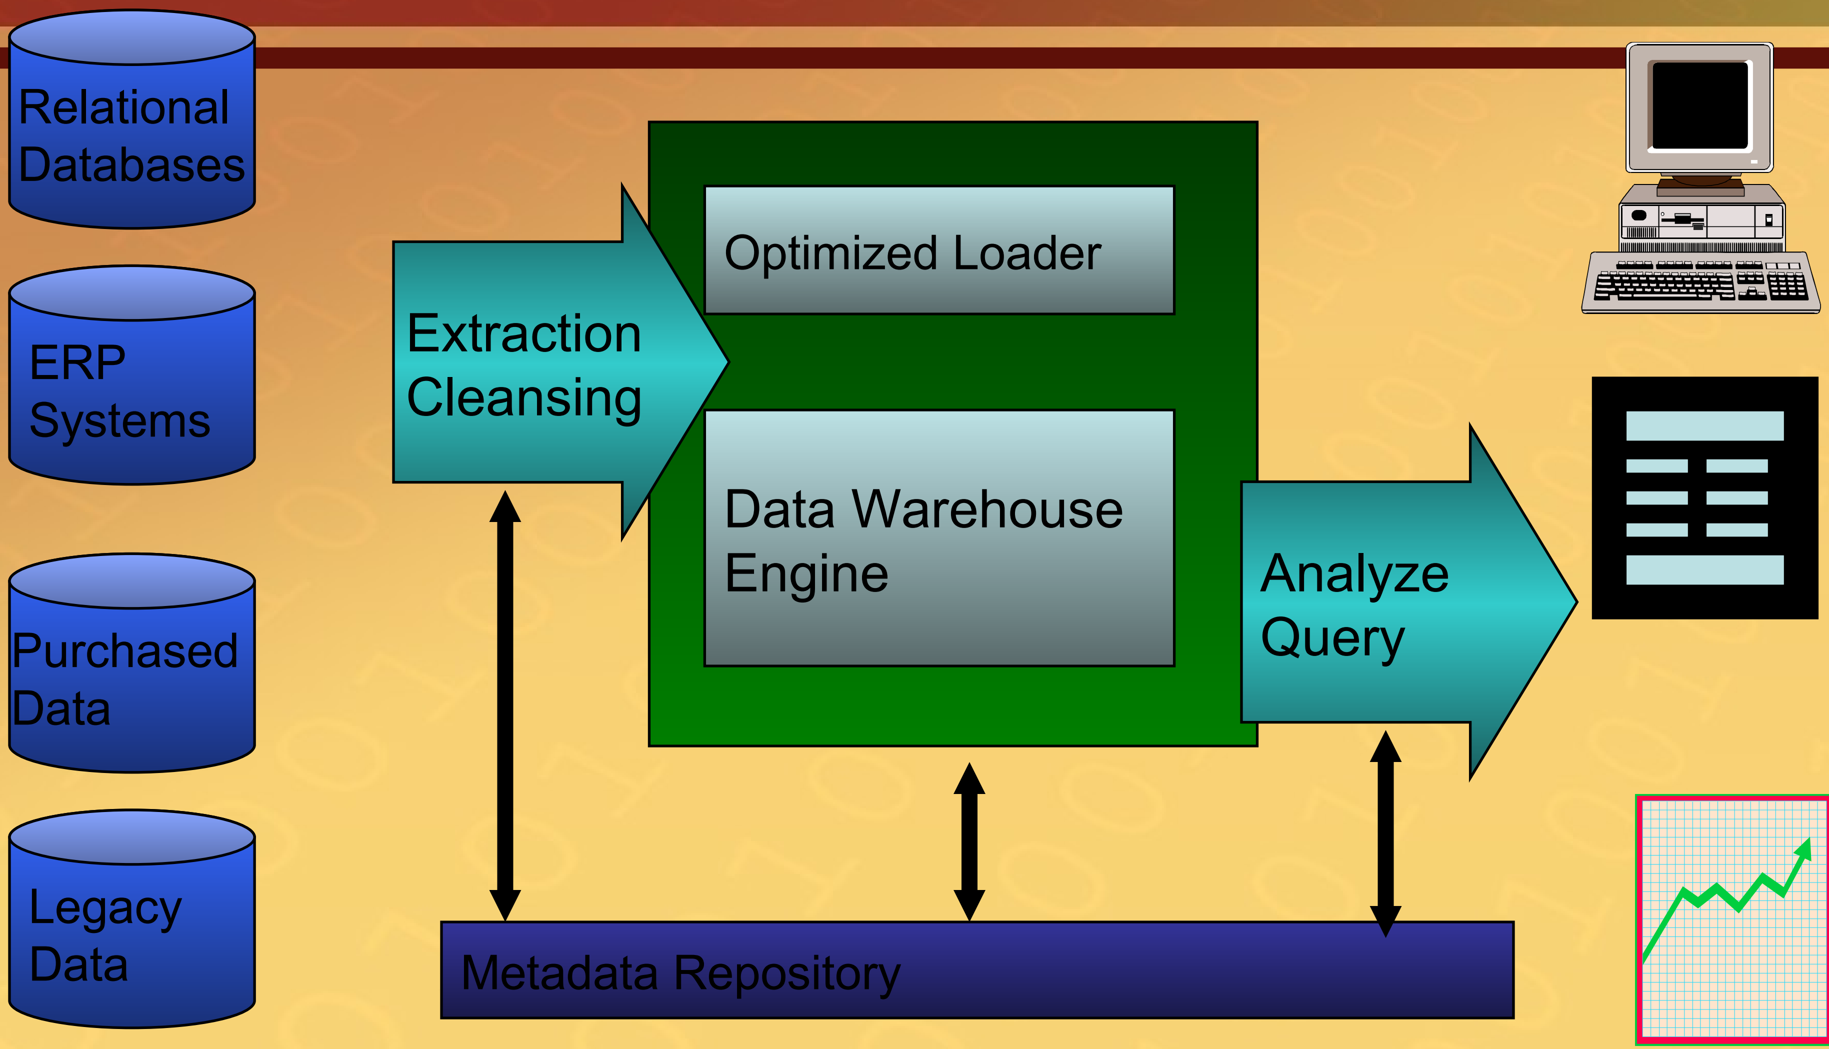

What is a Data Warehouse?

A single, complete and consistent store of data obtained from a variety of different sources made available to end users in a what they can understand and use in a business context.

[Barry Devlin]

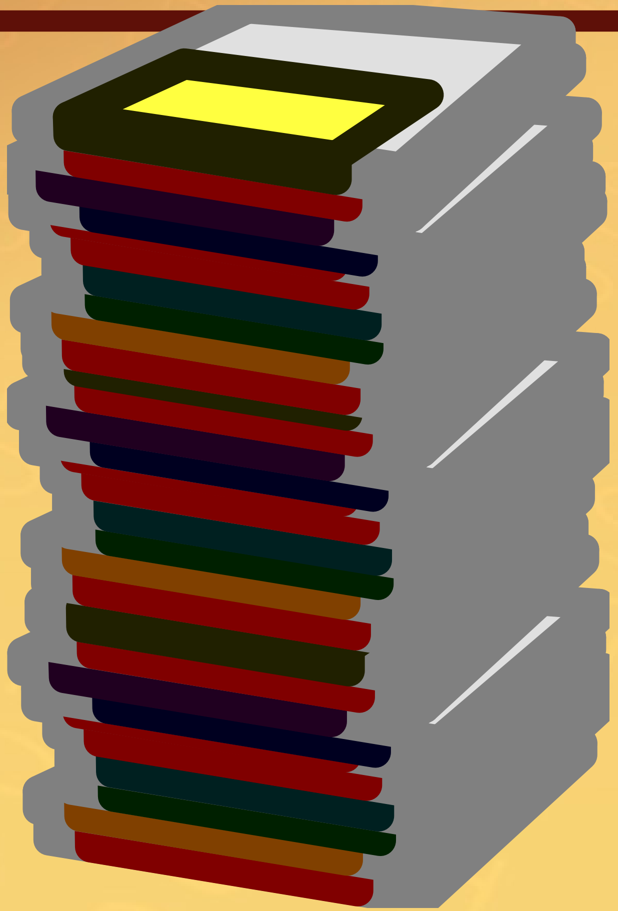

What are the users saying...

- Data should be integrated across the enterprise
- Summary data has a real value to the organization
- Historical data holds the key to understanding data over time
- What-if capabilities are required

What is Data Warehousing?

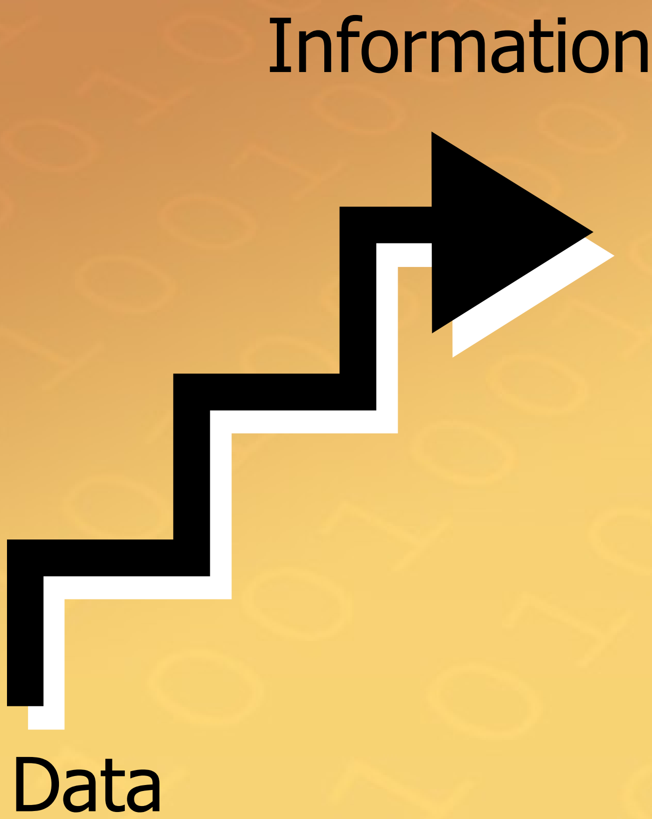

A process of transforming data into information and making it available to users in a timely enough manner to make a difference

[Forrester Research, April 1996]

Evolution

- 60’s: Batch reports
  - hard to find and analyze information
  - inflexible and expensive, reprogram every new request
- 70’s: Terminal-based DSS and EIS (executive information systems)
  - still inflexible, not integrated with desktop tools
- 80’s: Desktop data access and analysis tools
  - query tools, spreadsheets, GUIs
  - easier to use, but only access operational databases
- 90’s: Data warehousing with integrated OLAP engines and tools

Warehouses are Very Large Databases

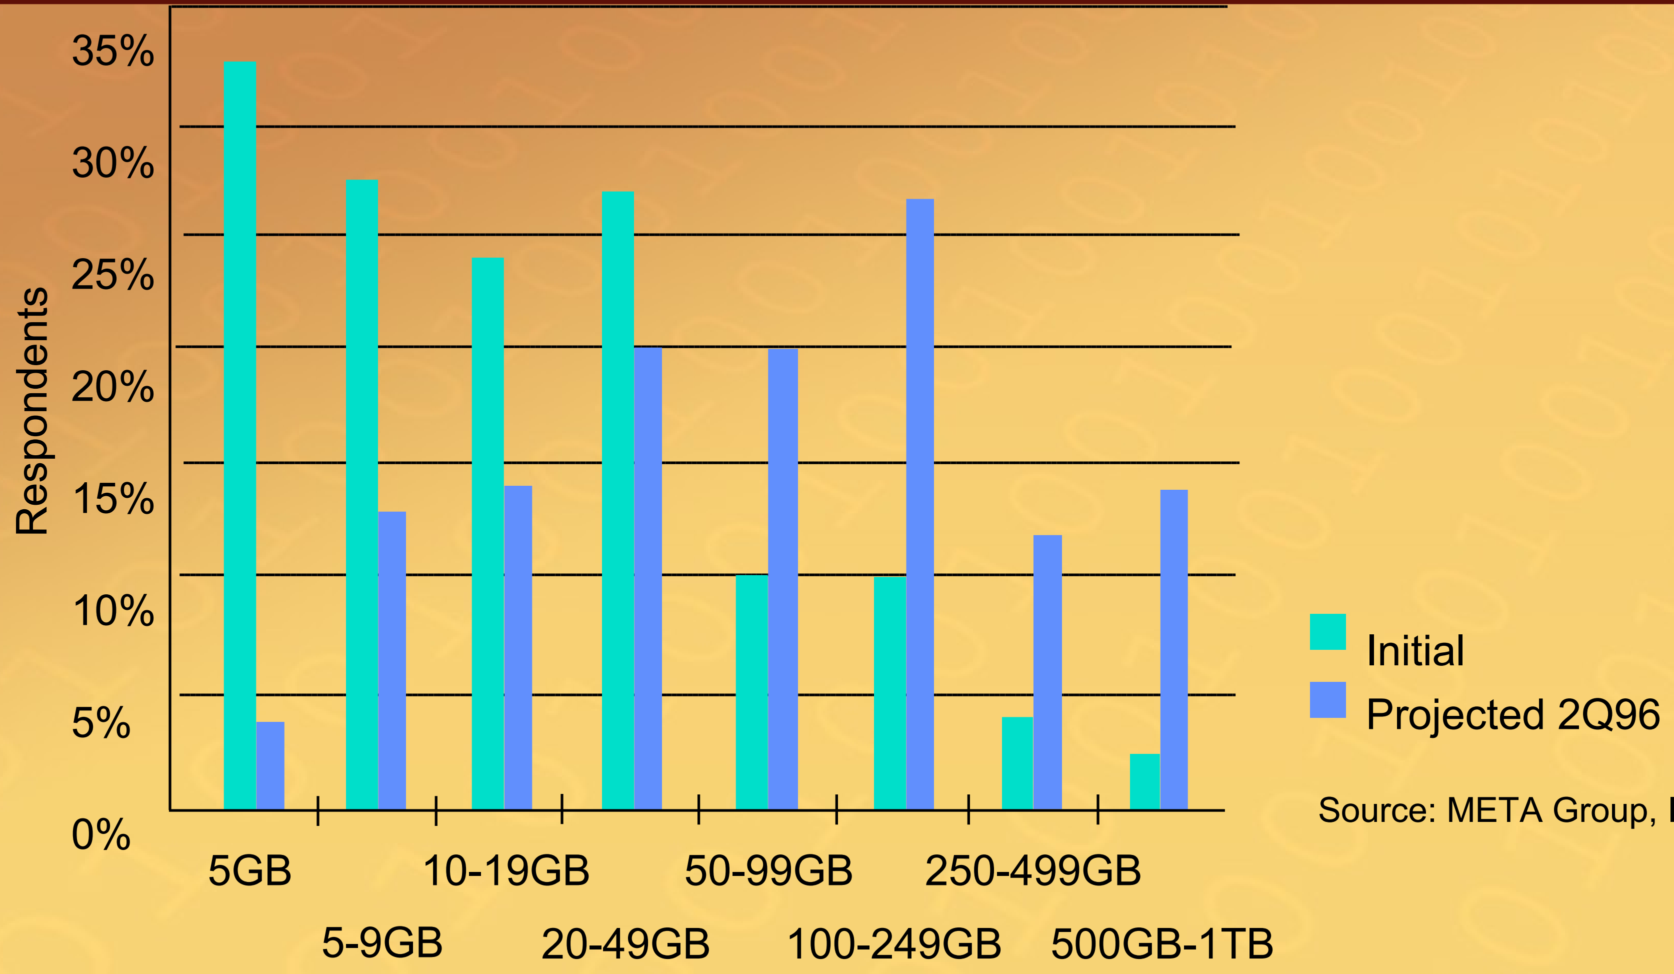

Source: META Group, Inc.

Very Large Data Bases

- Terabytes -- 10^12 bytes: Walmart -- 24 Terabytes
- Petabytes -- 10^15 bytes: Geographic Information Systems
- Exabytes -- 10^18 bytes: National Medical Records
- Zettabytes -- 10^21 bytes: Weather images
- Zottabytes -- 10^24 bytes: Intelligence Agency Videos

Data Warehousing -- It is a process

- Technique for assembling and managing data from various sources for the purpose of answering business questions. Thus making decisions that were not previous possible
- A decision support database maintained separately from the organization's operational database

Data Warehouse

- A data warehouse is a
  - subject-oriented
  - integrated
  - time-varying
  - non-volatile

  collection of data that is used primarily in organizational decision making.

-- Bill Inmon, Building the Data Warehouse 1996

A producer wants to know….

Which are our lowest/highest margin customers ?

Who are my customers and what products are they buying?

What is the most effective distribution channel?

What product prom- -otions have the biggest impact on revenue?

Which customers are most likely to go to the competition ?

What impact will new products/services have on revenue and margins?

Data, Data everywhere yet ...

- I can’t find the data I need
  - data is scattered over the network
  - many versions, subtle differences

I can’t get the data I need

need an expert to get the data

I can’t understand the data I found

available data poorly documented

I can’t use the data I found

results are unexpected

data needs to be transformed from one form to other

Business Analysis Framework

The business analyst get the information from the data warehouses to measure the performance and make critical adjustments in order to win over other business holders in the market. Having a data warehouse offers the following advantages:

Since a data warehouse can gather information quickly and efficiently, it can enhance business productivity.

A data warehouse provides us a consistent view of customers and items, hence, it helps us manage customer relationship.

A data warehouse also helps in bringing down the costs by tracking trends, patterns over a long period in a consistent and reliable manner.

Business Analysis Framework

To design an effective and efficient data warehouse, we need to understand and analyze the business needs and construct a **business analysis framework**. Each person has different views regarding the design of a data warehouse. These views are as follows:

The top-down view - This view allows the selection of relevant information needed for a data warehouse.

The data source view - This view presents the information being captured, stored, and managed by the operational system.

Business Analysis Framework

The data warehouse view - This view includes the fact tables and dimension tables. It represents the information stored inside the data warehouse.

The business query view - It is the view of the data from the viewpoint of the end-user.

Three-Tier Data Warehouse Architecture

Generally a data warehouses adopts a three-tier architecture. Following are the three tiers of the data warehouse architecture.

Bottom Tier - The bottom tier of the architecture is the data warehouse database server. It is the relational database system. We use the back end tools and utilities to feed data into the bottom tier. These back end tools and utilities perform the Extract, Clean, Load, and refresh functions.

Three-Tier Data Warehouse Architecture

Middle Tier - In the middle tier, we have the OLAP Server that can be implemented in either of the following ways.

- By Relational OLAP (ROLAP), which is an extended relational database management system. The ROLAP maps the operations on multidimensional data to standard relational operations.
- By Multidimensional OLAP (MOLAP) model, which directly implements the multidimensional data and operations.

Three-Tier Data Warehouse Architecture

Top-Tier - This tier is the front-end client layer. This layer holds the query tools and reporting tools, analysis tools and data mining tools.

Three-Tier Data Warehouse Architecture

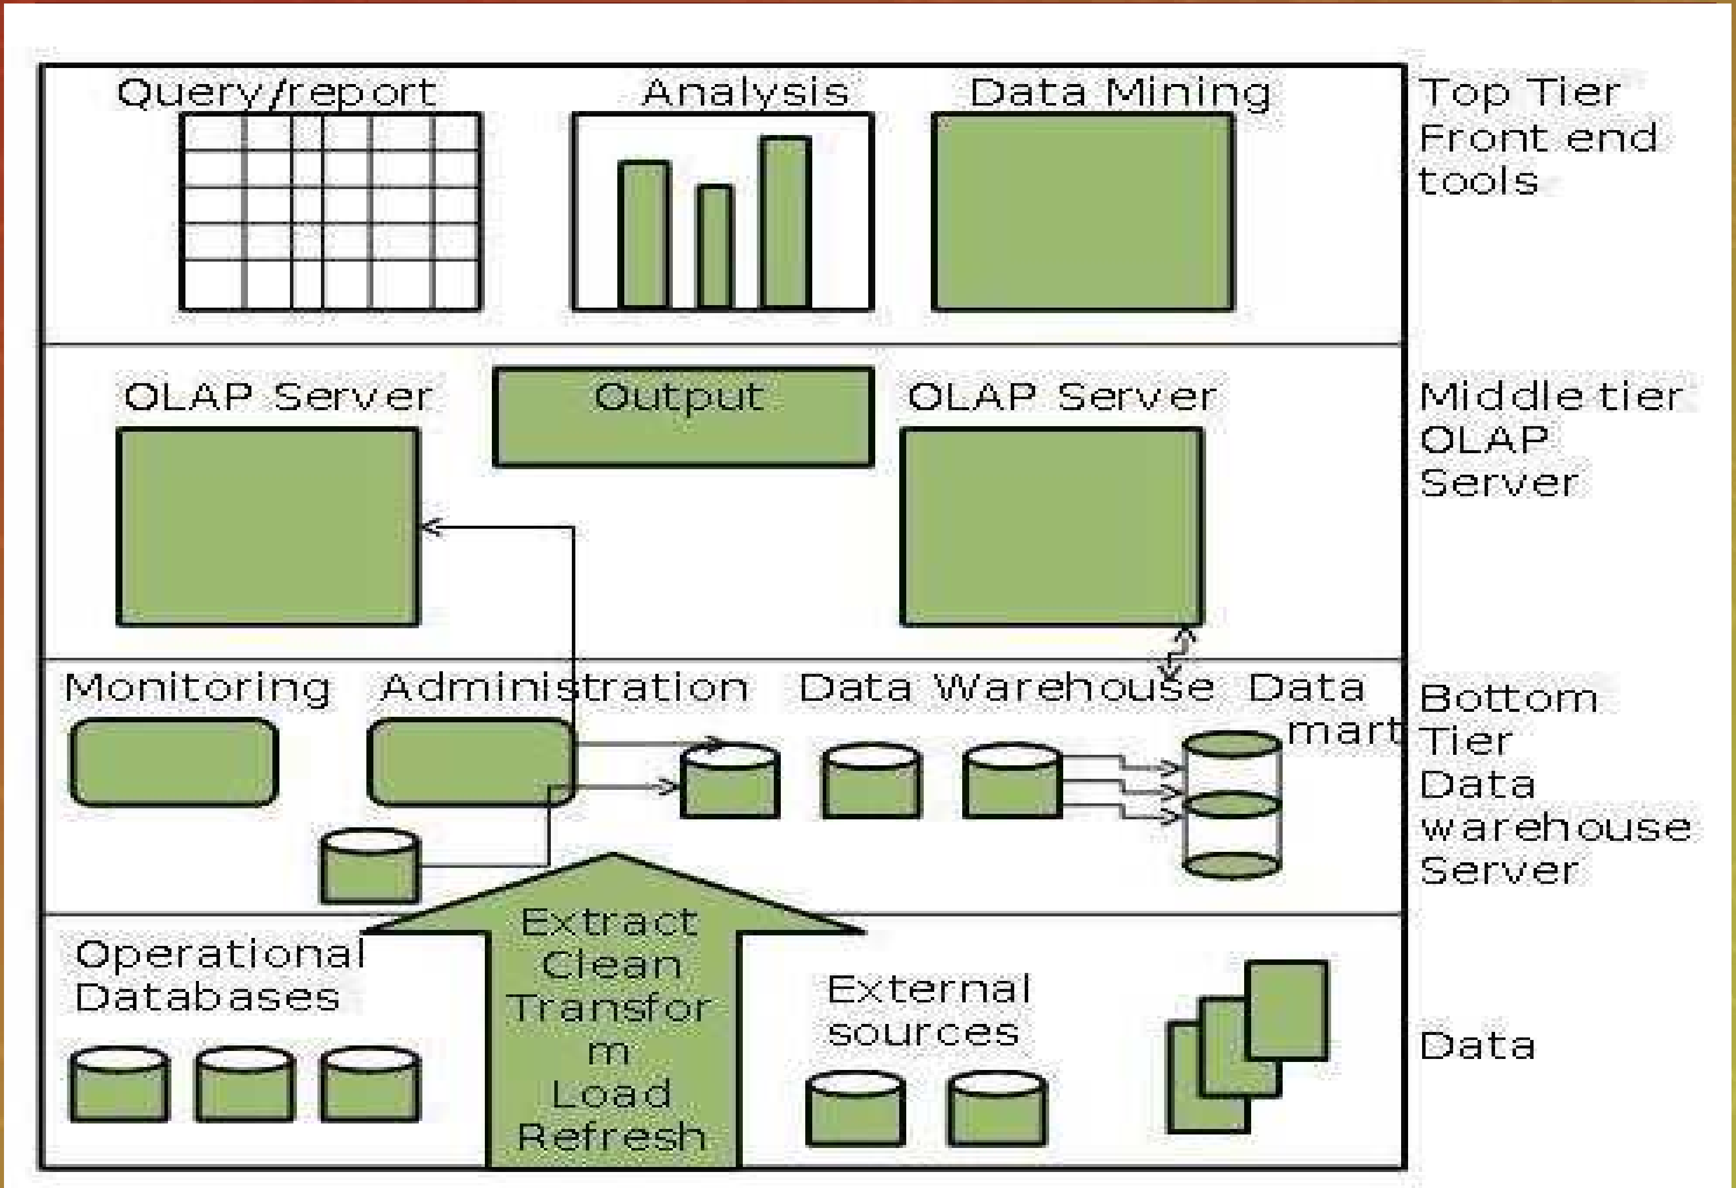

Data Warehouse Models

From the perspective of data warehouse architecture, we have the following data warehouse models:

Virtual Warehouse

Data mart

Enterprise Warehouse

Data Warehouse Models

Virtual Warehouse

The view over an operational data warehouse is known as a virtual warehouse. It is easy to build a virtual warehouse. Building a virtual warehouse requires excess capacity on operational database servers.

Data Warehouse Models

Data Mart

Data mart contains a subset of organization-wide data. This subset of data is valuable to specific groups of an organization.

In other words, we can claim that data marts contain data specific to a particular group. For example, the marketing data mart may contain data related to items, customers, and sales. Data marts are confined to subjects.

Data Warehouse Models

Points to remember about data marts:

Window-based or Unix/Linux-based servers are used to implement data marts. They are implemented on low-cost servers.

The implementation data mart cycles is measured in short periods of time, i.e., in weeks rather than months or years.

The life cycle of a data mart may be complex in long run, if its planning and design are not organization-wide.

Data marts are small in size.

Data marts are customized by department.

The source of a data mart is departmentally structured data warehouse.

Data mart are flexible.

Data Warehouse Models

Enterprise Warehouse

An enterprise warehouse collects all the information and the subjects spanning an entire organization.

It provides us enterprise-wide data integration.

The data is integrated from operational systems and external information providers.

This information can vary from a few gigabytes to hundreds of gigabytes, terabytes or beyond.

ABC of Knowledge Management

*Freely extracted from the NHS National Library for Health at http://www.library.nhs.uk/knowledgemanagement/
by [Géraud Servin]*

*Creator: NHS National Library for Health: Knowledge Management Specialist Library
Contributor: [Caroline De Brún]
Publication Date: July 2005*

Table of Contents

1 WHAT IS KNOWLEDGE MANAGEMENT?

Knowledge management is based on the idea that an organisation's most valuable resource is the knowledge of its people. Therefore, the extent to which an organisation performs well, will depend, among other things, on how effectively its people can create new knowledge, share knowledge around the organisation, and use that knowledge to best effect.

If you have read any of the huge array of knowledge management books and articles that are currently available, you are possibly feeling slightly bewildered. Perhaps you are wondering whether knowledge management is just the latest fad and hoping that if you ignore it, it will eventually go away. Let's be honest – knowledge management is surrounded by a great deal of hype. But if you can put the hype to one side, you will find that many of the tools, techniques and processes of knowledge management actually make a great deal of common sense, are already part of what you do, and can greatly help you in your job.

1.1 What is knowledge management?

Many of us simply do not think in terms of managing knowledge, but we all do it. Each of us is a personal store of knowledge with training, experiences, and informal networks of friends and colleagues, whom we seek out when we want to solve a problem or explore an opportunity. Essentially, we get things done and succeed by knowing an answer or knowing someone who does.

Fundamentally, knowledge management is about applying the collective knowledge of the entire workforce to achieve specific organisational goals. The aim of knowledge management is not necessarily to manage all knowledge, just the knowledge that is most important to the organisation. It is about ensuring that people have the knowledge they need, where they need it, when they need it – the right knowledge, in the right place, at the right time.

Knowledge management is unfortunately a misleading term – knowledge resides in people's heads and managing it is not really possible or desirable. What we can do, and what the ideas behind knowledge management are all about, is to establish an environment in which people are encouraged to create, learn, share, and use knowledge together for the benefit of the organisation, the people who work in it, and the organisation's customers (or in the case of the NHS, patients).

1.2 What is knowledge?

Academics have debated the meaning of "knowledge" since the word was invented, but let's not get into that here. A dictionary definition is "the facts, feelings or experiences known by a person or group of people" (Collins English Dictionary). Knowledge is derived from information but it is richer and more meaningful than information. It includes familiarity, awareness and understanding gained through experience or study, and results from making comparisons, identifying consequences, and making connections. Some experts include wisdom and insight in their definitions of knowledge. In organisational terms, knowledge is generally thought of as being "know how", or "applied action". The last point is an important one. Today's organisations contain a vast amount of knowledge and the NHS is certainly no exception. However, in applying knowledge management principles and practices in our organisation, knowledge is not our end, but the means for further action. What we are trying to do is to use our knowledge to get better at doing what we do, i.e. health care and health care improvement.

1.3 Why do we need knowledge management?

Knowledge management is based on the idea that an organisation's most valuable resource is the knowledge of its people. This is not a new idea – organisations have been managing "human resources" for years. What is new is the focus on knowledge. This focus is being driven by the accelerated rate of change in today's organisations and in society as a whole. Knowledge management recognises that today nearly all jobs involve "knowledge work" and so all staff are "knowledge workers" to some degree or another – meaning that their job depends more on their knowledge than their manual skills. This means that creating, sharing and using knowledge are among the most important activities of nearly every person in every organisation.

It is easy to see the importance of knowledge in the health sector. As clinicians, managers and other practitioners, we all rely on what we know to do our jobs effectively. But....

Do we know everything we need to know or are there gaps in our knowledge? Of course there are. Medical advances are being made all the time so there is always new knowledge to be learned. Government policies are constantly evolving, as are management practices. The current modernisation programme requires us to let go of what we knew and to learn and apply new knowledge. Changing doctor-patient relationships are requiring us to revisit our whole approach to the provision of health care. And of course, every new patient that comes through our door brings a potential new learning opportunity.

Do we share what we know? The NHS is made up of over a million individuals in hundreds of organisations, each of which have their own knowledge. Is the knowledge of individuals available to the whole organisation? Is the knowledge or organisations available to the whole NHS? Not at present. How many times have we lost valuable knowledge and expertise when a staff member moves on? How many times have we "reinvented the wheel" when we could have learned from someone else's experience? How many times have patients suffered as a result of the "postcode lottery"?

Do we use what we know to best effect? Not always. In the NHS Plan, the NHS was described as "a 1940s infrastructure operating in the 21st century". Clearly our knowledge has not always been applied to best effect, and we have fallen behind the times. How many times have we had an idea about how a process or an activity could be improved, but felt we lacked the time or resources to do anything about it? How many times have we had an idea that might help our colleagues, but we keep quiet because our colleagues might not appreciate us "telling them how to do their job"? How many times have we implemented a new initiative, only to find we reverted back to the "old way" a few months later? Perhaps we have had insights about how our patients" needs could be better met, but there was no forum for us to share and explore those insights so we just forgot about it.

These are just a few examples.

Almost everything we do in the NHS is based on our knowledge. If we do not constantly update and renew our knowledge, share our knowledge, and then use that knowledge to do things differently and better, then our people, our organisations, our patients and the general public will ultimately suffer. We know this because it has already happened. As The NHS Plan (2000) affirms, in spite of our many achievements, the NHS has failed to keep pace with changes in our society. What can transform that, along with the current investment and modernisation programme, is harnessing the vast collective knowledge of the people working in the NHS, and using it to best effect. That is why we need knowledge management.

1.4 What does knowledge management involve?

Knowledge management is essentially about facilitating the processes by which knowledge is created, shared and used in organisations. It is not about setting up a new department or getting in a new computer system. It is about making small changes to the way everyone in the organisation works. There are many ways of looking at knowledge management and different organisations will take different approaches. Generally speaking, creating a knowledge environment usually requires changing organisational values and culture, changing people's behaviours and work patterns, and providing people with easy access to each other and to relevant information resources.

In terms of how that is done, the processes of knowledge management are many and varied. As knowledge management is a relatively new concept, organisations are still finding their way and so there is no single agreed way forward or best practice. This is a time of much trial and error. Similarly, to simply copy the practices of another organisation would probably not work because each organisation faces a different set of knowledge management problems and challenges. Knowledge management is essentially about people – how they create, share and use knowledge, and so no knowledge management tool will work if it is not applied in a manner that is sensitive to the ways people think and behave.

That being said, there are of course a whole raft of options in terms of tools and techniques, many of which are not new. Many of the processes that currently fall under the banner of knowledge management have been around for a long time, but as part of functions such as training, human resources, internal communications, information technology, librarianship, records management and marketing to name a few. And some of those processes can be very simple, such as:

- providing induction packs full of "know how" to new staff;
- conducting exit interviews when staff leave so that their knowledge is not lost to the organisation;
- creating databases of all publications produced by an organisation so that staff can access them from their desk;

- providing ongoing learning so that people can constantly update their knowledge;
- encouraging people with a common interest to network with each other;
- creating electronic filing systems that can be searched in a number of ways, making the information much easier to find;
- redesigning offices to be open plan so that staff and managers are more visible and talk to each other more;
- putting staff directories online so that people can easily find out who does what and where they are;
- creating intranets so that staff can access all kinds of organisational information and knowledge that might otherwise take a great deal of time and energy to find.

1.5 Some "textbook" definitions of knowledge management

Here are a few definitions:

- "Clinical knowledge management means enhancing the identification, dissemination, awareness and application of the results of research relevant to clinical practice in health and social care."
  Jeremy Wyatt
- "The creation and subsequent management of an environment, which encourages knowledge to be created, shared, learnt, enhanced, organised and utilized for the benefit of the organisation and its customers."
  Abell & Oxbrow, tfpl Ltd, 2001
- "Knowledge management is a process that emphasises generating, capturing and sharing information know how and integrating these into business practices and decision making for greater organisational benefit."
  Maggie Haines, NHS Acting Director of KM
- "The capabilities by which communities within an organisation capture the knowledge that is critical to them, constantly improve it, and make it available in the most effective manner to those people who need it, so that they can exploit it creatively to add value as a normal part of their work."
  BSI's A Guide to Good Practice in KM
- "Knowledge is power, which is why people who had it in the past often tried to make a secret of it. In post-capitalism, power comes from transmitting information to make it productive, not from hiding it!"
  Peter Drucker
- "Knowledge management involves efficiently connecting those who know with those who need to know and converting personal knowledge into organisational knowledge."
  Yankee Group
- "Knowledge management is not about data, but about getting the right information to the right people at the right time for them to impact the bottom line."
  IBM
- "The capability of an organization to create new knowledge, disseminate it throughout the organization and embody it in products, services and systems."
  Nonaka & Takeuchi, 1995
- "Knowledge management is a relatively young corporate discipline and a new approach to the identification, harnessing and exploitation of collective organisational information, talents, expertise and know-how."
  Office of thee-Envoy, 2002
- "Knowledge management is the explicit and systematic management of vital knowledge and its associated processes of creating, gathering, organizing, diffusion, use and exploitation. It requires turning personal knowledge into corporate knowledge that can be widely shared throughout an organization and appropriately applied."
  David J Skyrme, 1997

2 PRINCIPLES AND PROCESSES OF KNOWLEDGE MANAGEMENT

A "rough guide" to some of the main general approaches to knowledge management.

2.1 Right knowledge, right place, right time

Some people mistakenly assume that knowledge management is about capturing all the best practices and knowledge that people possess and storing it in a computer system in the hope that one day it will be useful. In fact this is a good example of what knowledge management is not about! Consider this: how often has information or knowledge been pushed at you when you don't need it – paper, emails, training, another irrelevant meeting? Then later, when you do need it, you vaguely remember seeing something relevant but can't find it. Some surveys suggest that professional workers spend ten per cent of their time looking for information they know is somewhere. And if what you want is in people's heads, and they're not always around, how can you access it when you need it? What if you don't even know whose head it's in, or if they'd be willing to share it with you?

In a nutshell, good knowledge management is all about getting the **right knowledge, in the right place, at the right time.**

The right knowledge is the knowledge that you need in order to be able to do your job to the best of your ability, whether that means diagnosing a patient, making a decision, booking a referral, answering a patient's question, administering a treatment, training a new colleague, interpreting a piece of research, using a computer system, managing a project, dealing with suppliers etc. Information and knowledge can usually be found in a whole variety of places – research papers, reports and manuals, databases etc. Often it will be in people's heads – yours and other people's. The right place, however, is the point of action or decision – the meeting, the patient helpline, the hospital bedside, behind the reception desk and so on. The right time is when you (the person or the team doing the work) need it.

2.2 Types of knowledge: explicit and tacit

Knowledge in organisations is often classified into two types: explicit and tacit.

1. **Explicit knowledge** is knowledge that can be captured and written down in documents or databases. Examples of explicit knowledge include instruction manuals, written procedures, best practices, lessons learned and research findings. Explicit knowledge can be categorised as either structured or unstructured. Documents, databases, and spreadsheets are examples of structured knowledge, because the data or information in them is organised in a particular way for future retrieval. In contrast, e-mails, images, training courses, and audio and video selections are examples of unstructured knowledge because the information they contain is not referenced for retrieval.
2. **Tacit knowledge** is the knowledge that people carry in their heads. It is much less concrete than explicit knowledge. It is more of an "unspoken understanding" about something, knowledge that is more difficult to write down in a document or a database. An example might be, knowing how to ride a bicycle – you know how to do it, you can do it again and again, but could you write down instructions for someone to learn to ride a bicycle? Tacit knowledge can be difficult to access, as it is often not known to others. In fact, most people are not aware of the knowledge they themselves possess or of its value to others. Tacit knowledge is considered more valuable because it provides context for people, places, ideas and experiences. It generally requires extensive personal contact and trust to share effectively.

2.3 Types of knowledge: old and new

Most knowledge management strategies generally have one (or sometimes both) of two thrusts. The first is to make better use of the knowledge that already exists within the organisation, and the second is to create new knowledge.

Making better use of the knowledge that already exists within an organisation ("old" knowledge) often begins with "knowing what you know". Very often leading managers comment: "if only we knew what we knew". Too frequently, people in one part of the organisation reinvent the wheel or fail to solve a problem because the knowledge they need is elsewhere in the organisation but not known or accessible to them. Hence the first knowledge management initiative of many companies is that of finding out what they know, and taking steps to make that knowledge accessible throughout the organisation. Specific approaches might include conducting a knowledge audit, mapping the organisation's knowledge resources and flows, making tacit knowledge more explicit and putting in place mechanisms to move it more rapidly to where it is needed.

Creating new knowledge can equally be approached in a number of ways such as through training, hiring external resources, bringing different people and their knowledge together to create fresh knowledge and insights, etc. It is also about innovation – making the transition from ideas to action more effective. Many managers mistakenly believe this is about R&D and creativity. In fact there is no shortage of creativity in organisations – not just in R&D but everywhere. The real challenge is not to lose these creative ideas and to allow them to flow where they can be used.

In reality, the distinction between "old" and "new" knowledge is not always that clear. Innovation will often draw on lessons from the past, particularly those that have been forgotten, or those that can be put together in new combinations to achieve new results. Similarly, the application of (old) knowledge almost always involves some adaptation, and so in the process of adaptation, new knowledge is created. At the end of the day, the quality of knowledge does not depend on whether it is "old" or "new" but rather whether it is relevant. Whether it is old or new hardly matters. The question is: does it work in practice?

2.4 Ways with knowledge: collecting and connecting

Knowledge management programmes tend to have both a "collecting" and a "connecting" dimension.

The collecting dimension involves linking people with information. It relates to the capturing and disseminating of explicit knowledge through information and communication technologies aimed at codifying, storing and retrieving content, which in principle is continuously updated through computer networks. Through such collections of content, what is learned is made readily accessible to future users.

Current examples in the NHS include various intranets, the National electronic Library for Health, the CLIP database, The Cochrane Library, and many more. This collecting dimension is often the main emphasis of many European and US knowledge programmes. However it has its limitations. Even where comprehensive collections of materials exist, effective use may still need knowledgeable and skilled interpretation and subsequent alignment with the local context to get effective results, just as reading a newspaper article on brain surgery does not qualify or enable a reader to conduct brain surgery. An organisation that focuses completely on collecting and makes little or no effort at connecting (see below) tends to end up with a repository of static documents.

The connecting dimension involves linking people with people – specifically people who need to know with those who do know, and so enhancing tacit knowledge flow through better human interaction, so that knowledge is diffused around the organisation and not just held in the heads of a few. Connecting is necessary because knowledge is embodied in people, and in the relationships within and between organisations. Information becomes knowledge as it is interpreted in the light of the individual's understandings of the particular context. Examples of connecting initiatives include skills directories and expert directories – searchable online staff directories that give much more detail about who does what and who knows what, collaborative working, communities of practice – networks of people with a common interest, and various "socialisation" activities designed to support knowledge flows. This connecting dimension tends to be the main emphasis in Japanese knowledge programmes. However an organisation that focuses entirely on connecting, with little or no attempt at collecting, can be very inefficient. Such organisations may waste time in "reinventing wheels".

Most knowledge management programmes aim at an integrated approach to managing knowledge, by combining the benefits of both approaches and achieving a balance between connecting individuals who need to know with those who do know, and collecting what is learned as a result of these connections and making that easily accessible to others. For example, if collected documents are linked to their authors and contain other interactive possibilities, they can become dynamic and hence much more useful.

2.5 Ways with knowledge: people, processes and technology

One popular and widely-used approach is to think of knowledge management in terms of three components, namely people, processes and technology:

- People: Getting an organisation's culture (including values and behaviours) "right" for knowledge management is typically the most important and yet often the most difficult challenge. Knowledge management is first and foremost a people issue. Does the culture of your organisation support ongoing learning and knowledge sharing? Are people motivated and rewarded for creating, sharing and using knowledge? Is there a culture of openness and mutual respect and support? Or is your organisation very hierarchical where "knowledge is power" and so people are reluctant to share? Are people under constant pressure to act, with no time for knowledge-seeking or reflection? Do they feel inspired to innovate and learn from mistakes, or is there a strong "blame and shame" culture?
- Processes: In order to improve knowledge sharing, organisations often need to make changes to the way their internal processes are structured, and sometimes even the organisational structure itself. For example, if an organisation is structured in such a way that different parts of it are competing for resources, then this will most likely be a barrier to knowledge sharing. Looking at the many aspects of "how things are done around here" in your organisation, which processes constitute either barriers to, or enablers of, knowledge management? How can these processes be adapted, or what new processes can be introduced, to support people in creating, sharing and using knowledge?
- Technology: A common misconception is that knowledge management is mainly about technology – getting an intranet, linking people by e-mail, compiling information databases etc. Technology is often a crucial enabler of knowledge management – it can help connect people with information, and people with each other, but it is not the solution. And it is vital that any technology used "fits" the organisation's people and processes – otherwise it will simply not be used.

These three components are often compared to the legs of a three-legged stool – if one is missing, then the stool will collapse. However, one leg is viewed as being more important than the others – people. An organisation's primary focus should be on developing a knowledge-friendly culture and knowledge-friendly behaviours among its people, which should be supported by the appropriate processes, and which may be enabled through technology.

3 GENERAL CONCEPTS

3.1 A brief history of knowledge management

Knowledge management as a conscious discipline would appear to be somewhere between five and fifteen years old. It evolved from the thinking of academics and pioneers such as Peter Drucker in the 1970s, Karl-Erik Sveiby in the late 1980s, and Nonaka and Takeuchi in the 1990s. During that time, economic, social and technological changes were transforming the way that companies worked. Globalisation emerged and brought new opportunities and increased competition. Companies responded by downsizing, merging, acquiring, reengineering and outsourcing. Many streamlined their workforce and boosted their productivity and their profits by using advances in computer and network technology. However their successes in doing so came with a price. Many lost company knowledge as they grew smaller. And many lost company knowledge as they grew bigger – they no longer "knew what they knew".

By the early 1990s a growing body of academics and consultants were talking about knowledge management as "the" new business practice, and it began to appear in more and more business journals and on conference agendas. By the mid-1990s, it became widely acknowledged that the competitive advantage of some of the world's leading companies was being carved out from those companies' knowledge assets such as competencies, customer relationships and innovations. Managing knowledge therefore suddenly became a mainstream business objective as other companies sought to follow the market leaders.

Many of these companies took the approach of implementing "knowledge management solutions", focusing almost entirely on knowledge management technologies. However they met with limited success, and so questions began to be asked about whether knowledge management wasn't simply another fad that looked great on paper, but in reality did not deliver. In fact for a while, it looked as if knowledge management was destined to be confined to the "management fad graveyard". However on closer inspection, companies realised that it wasn't the concept of knowledge management that was the problem as such, but rather the way that they had gone about approaching it. Reasons for their limited success included:

- The focus was on the technology rather than the business and its people.
- There was too much hype – with consultants and technology vendors cashing in on the latest management fad.
- Companies spent too much money (usually on "sexy" technologies) with little or no return on their investments.
- Most knowledge management literature was very conceptual and lacking in practical advice, which led to frustration at the inability to translate the theory into practice – "it all makes so much sense but why isn't it working?"
- Knowledge management was not tied into business processes and ways of working.
- It was seen as another laborious overhead activity or yet another new initiative.
- A lack of incentives – employees quite rightly asked the "what's in it for me?" question.
- There wasn't sufficient senior executive level buy in.

Fortunately companies are now recognising these early mistakes and are beginning to take a different approach to knowledge management – one in which the emphasis is more on people, behaviours and ways of working, than on technology. Of course there are still some sceptics who believe that knowledge management is just a fad. But according to a number of company surveys, it would seem that they are in a minority. A more popular view is that knowledge management may not remain as a distinct discipline, but rather will become embedded in the way organisations work. This can be compared to Total Quality Management which was the "in thing" in the 1980s; nobody talks about "TQM" any more, but many of its principles and practices are an integral part of how most organisations operate. It looks likely that this could also be the future for knowledge management.

3.2 The "knowledge economy"

"As we enter the 21st century we are moving into a new phase of economic and social development, which can usefully be referred to as a "knowledge economy", in which knowledge will be a key determining factor in organizational and economic success or failure. The most effective organizations in the knowledge economy will be those which recognize and best harness the crucial role that knowledge plays both inside and outside their organisation."
From: **Knowledge Enhanced Government: A strategy for the UK Office of the e-Envoy**, July 2002

The government's objective is to make the UK one of the world's leading knowledge economies.

3.3 Knowledge management in the public sector

In both the private and public sectors, more and more organisations are beginning to take responsibility for managing knowledge as a means to create value. But what does "value" mean in the context of the public sector? Public sector organisations are not usually seeking a competitive advantage, so why bother with knowledge management? If we go back to our definition of knowledge as "the capacity for effective action" (see the section What is KM?) then this probably better describes the expectations of government and public services. Every public service involves a wide range of relationships between policy makers, service providers, local authorities, the general public and various other interested parties such as voluntary and community sector organisations, the private sector etc. If we think about the many interactions within and between these groups, and their impact on policy and service provision, then we begin to see the scope for knowledge management in the public sector. How does one of these various parties share an experience and introduce one policy driven initiative with that of another for the benefit of all concerned? How can everyone involved have an awareness of the "bigger picture" as well as their own individual standpoints? How can all parties be better prepared to act?

In recent years there has been a number of government policies aimed at equipping the public sector to function more effectively in an information society. These have included:

- our Information Age (HMSO, 1998) – the de facto UK national information policy
- open for learning, open for business (National Grid for Learning, 1998) – establishing a commitment to a national grid for learning
- modernising government (HMSO, 1999) – committed government to modernising public services so that all would be capable of being delivered by computer by 2005
- e-government (Cabinet Office, 2000) – a strategic framework for public services in the information age

Building on this, subsequent developments have focused on making better use of the tacit knowledge within, and improving knowledge transfer across, the public sector.

The Office of the e-Envoy's UK Annual Report 2000 announced the development of a cross-government knowledge management system, focusing on the creation of a Knowledge Network – "a unified cross-government communications infrastructure to enable officials in all government departments and associated bodies … to communicate electronically with each other and share common, secure access to databases, discussion forums, web-based community sites and "knowledge pools"."

From there, a new programme of modernisation led by the Office of the e-Envoy known as Knowledge Enhanced Government (KEG) was launched. The KEG team is working with the major central government departments in ensuring that there are departmental teams and processes in place to support participation in KEG. The Department of Health is already a key player in these processes.

As part of KEG, the Office of the e-Envoy has recently considered the development of a knowledge management policy framework to provide a holistic view of knowledge management and recommendations for activity. Early proposals have suggested that this framework could be based around ten key areas of activity:

1. knowledge capture – policies and processes for identifying and capturing explicit and tacit knowledge.
2. knowledge transfer – policies and processes for transferring knowledge among and between its various sources and forms.
3. knowledge retention – policies and processes for retaining organisational knowledge, especially during periods of organisational change.
4. content management – policies and processes for efficiently managing the organisational knowledge base.
5. knowledge capital – policies and processes for measuring and developing the government's human and social capital.
6. enabling communities – policies and processes for promoting and supporting knowledge-based community working across and between departments.
7. supporting a knowledge culture – policies and processes to create the necessary cultural changes to embed the knowledge management ethos into working practices.
8. knowledge partnerships – policies and processes for promoting and supporting knowledge partnerships between central government and key partners such as local government, departmental agencies, non-departmental public bodies, voluntary and community organizations etc.
9. supporting key business activities – policies and processes to support key business activities in government such as project management, the legislative process, delivery monitoring etc.
10. knowledge benchmarking – policies and processes for benchmarking current knowledge management capabilities and practices against UK and international best practice, and for improving performance.

For more information about Knowledge Enhanced Government and related initiatives, see the Office of the e-Envoy website at http://archive.cabinetoffice.gov.uk/e-envoy/index-content.htm.

4 GETTING STARTED

With such a wide range of definitions, philosophies, methodologies, tools and techniques, approaching knowledge management can initially seem quite daunting. In starting out, many practitioners tend to offer the following types of advice:

4.1.1 Review your options

It is useful to gain a broad understanding of the variety of approaches to knowledge management. Not only are there many alternatives, but also some of them differ quite widely from others in their methods. Before selecting your approach or approaches, try to explore the many options open to you.

4.1.2 Don't get too hung up on "the best"

There is no single "right" way to approach knowledge management. Knowledge management methods are as varied as the organisations in which they are implemented. Every organisation is different and so its approach to knowledge management will need to reflect its own particular circumstances. There is no "one size fits all". The "best" approach will be one that works well for your organisation.

4.1.3 Keep it simple – avoid rocket science

There is still quite a lot of confusion about what knowledge management actually is and what it involves. Don't add to that confusion by blinding people with rocket science and textbook definitions. Get clear on what knowledge management means for your organisation. Then make the concepts of knowledge management real for others in your organisation. Use simple definitions and simple language to explore real problems and opportunities. Create a clear, tangible picture of the benefits of knowledge management as they relate to your organisation's specific goals and circumstances.

4.1.4 Learn while doing

Avoid the temptation to wait until you have "mastered" the theory of knowledge management before getting started on the practice. (The theory is constantly evolving, so the chances are you will never master it). One of the best ways to learn is "on the job". You can learn a great deal from what others have done, but you will only learn what does and doesn't work for your organisation when you actually get started and do something.

4.1.5 Celebrate what you're already doing

Start from where you are, with what you have. In most organisations there will already be examples of good knowledge management practice – except they won't usually be thought of as knowledge management. Look around your organisation for current activities that might already be related to knowledge management – not necessarily big projects or initiatives, but simple, day-to-day ways of doing things. Look for teams or groups that are currently sharing knowledge, and make connections with these people. Find out how it is benefiting those people and the organisation as a whole. Celebrate and build on these examples of good practice.

4.1.6 Look at your organisation's goals

Given that knowledge management is not an end in itself, but rather a means to achieving organisational goals, then this is a logical place to start. Look at both the long-term goals and short to medium-term objectives of your organisation: what are they? How might knowledge management help you to achieve them? Then look at what people – teams and individuals – do in your organisation. What are the services they provide? What activities and processes do they perform in order to provide those services? How might they be done better for the benefit of individual staff, the organisation a whole, and your patients? What knowledge do people need in order to do their jobs? What knowledge might they need in order to do them better? How can you acquire, create, use and share that knowledge to bring that about? In what ways are you already doing so? How might you do it better?

4.1.7 Look for needs, problems and pains

Another good place to start is with what some managers call "needs, problems and pains". These are the things that are not working well in your organisation: things that are getting in the way of people doing a good job, things that irritate people and make their lives difficult, things that hamper the quality of your service to patients. Talk to people and start to build up a list of some of the major needs, problems and pains in your organisation. From there, you can select one or several of these with which to start, and look at how you might resolve it using knowledge management principles and practices. A great advantage of this approach is that it can allow you to achieve "quick wins". These are problems that are generally fairly simple and quick to resolve, but their resolution has a big impact and the results are clearly visible. Quick wins can be very useful in demonstrating the potential benefits of knowledge management to both managers and staff – there is nothing like real results to win people over.

4.1.8 Start small

Attempting to launch an organisation-wide knowledge management programme without building the evidence first is unfortunately a common mistake, but one to be avoided. Some organisations prefer to "dip their toe in the water" with one or two initiatives before considering a formal knowledge management strategy; others choose not to create a formal strategy at all, choosing instead to take a more informal or incremental approach. Either way, whether you choose to create a formal knowledge management strategy or not, a large-scale, high-cost, "big bang" roll-out is not recommended. Knowledge management is more an iterative process of continuous development. Hence, it is far better to gradually introduce a series of practical, manageable changes. Then, as interest develops, you can look to expand your initiatives.

4.1.9 Don't take off without a pilot

When looking to implement any major new initiative, conducting a pilot is essential. A pilot involves "test driving" the initiative on a relatively small scale in order to learn what works and what doesn't, make any necessary changes accordingly, and gather clear, demonstrable evidence about the benefits, before rolling out the initiative on a larger scale. This means that when you come to roll it out, you have already made most of your mistakes, and you have something that has been proven to work well in practice. In terms of securing resources and support, this is a whole different proposition to having an idea in theory.

4.1.10 Remember the "big three": people, processes, technology

In implementing knowledge management tools and techniques, never forget the importance of creating the right kind of environment. Your organisation's people, processes and technology will at all times be acting as either enablers of, or barriers to, the effective use of your knowledge management tools. You need to identify the barriers and remove them, and build on the enablers. If you have already tried to implement something and it hasn't worked, this is where you need to look. If you are about to implement something, look before you leap.

4.1.11 The ultimate aim: institutionalisation

Granted, you are just starting out with knowledge management. This is the beginning of the road. However it is worth keeping one eye on the horizon further down that road. It is useful to bear in mind that success in knowledge management does not involve building up a big new department or a whole network of people with "knowledge" in their job title. You may need to do these things to some degree in the medium-term. However the ultimate aim is for knowledge management to be fully "institutionalised". Or in other words, so embedded in the way your organisation does things, so intrinsic in people's day-to-day ways of working, that nobody even talks about knowledge management any more – they just do it. So if you are a knowledge manager, you will know that you have fully succeeded when you have worked yourself out of a job!

4.2 KM toolbox – inventory of tools and techniques

The following "toolbox" presents some of the most common tools and techniques currently used in knowledge management programmes. The aim is to give an introduction, to present an overview of what is involved, and to provide some pointers to further resources.

1 After Action Reviews (AARs)
A tool pioneered by the US army and now widely used in a range of organisations to capture lessons learned both during and after an activity or project.

2 Communities of Practice
Widely regarded as "the killer KM application', communities of practice link people together to develop and share knowledge around specific themes, and are already being established in the NHS.

3 Conducting a knowledge audit
A systematic process to identify an organisation's knowledge needs, resources and flows, as a basis for understanding where and how better knowledge management can add value.

4 Developing a knowledge management strategy
Approaches to developing a formal knowledge management plan that is closely aligned with an organisation's overall strategy and goals.

5 Exit interviews
A tool used to capture the knowledge of departing employees.

6 Identifying and sharing best practices
Approaches to capturing best practices discovered in one part of the organisation and sharing them for the benefit of all.

7 Knowledge centres
Similar to libraries but with a broader remit to include connecting people with each other as well as with information in documents and databases.

8 Knowledge harvesting
A tool used to capture the knowledge of "experts" and make it available to others.

9 Peer assists
A tool developed at BP-Amoco used to learn from the experiences of others before embarking on an activity or project.

10 Social network analysis
Mapping relationships between people, groups and organisations to understand how these relationships either facilitate or impede knowledge flows.

11 Storytelling
Using the ancient art of storytelling to share knowledge in a more meaningful and interesting way.

12 White pages
A step-up from the usual staff directory, this is an online resource that allows people to find colleagues with specific knowledge and expertise.

4.3 After Action Reviews

4.3.1 What are after action reviews?

An after action review (AAR) is a discussion of a project or an activity that enables the individuals involved to learn for themselves what happens, why it happened, what went well, what needs improvement and what lessons can be learned from the experience. The spirit of an AAR is one of openness and learning – it is not about problem fixing or allocating blame. Lessons learned are not only tacitly shared on the spot by the individuals involved, but can be explicitly documented and shared with a wider audience.

After action reviews were originally developed and are extensively used by the US Army.

4.3.2 What are the benefits?

What makes after action reviews so powerful is that they can be applied across a wide spectrum of activities, from two individuals conducting a five minute AAR at the end of a short meeting to a day-long AAR held by a project team at the end of a large project. Activities suitable for AARs simply need to have a beginning and an end, an identifiable purpose and some basis on which performance can be assessed. Other than that, there are few limits.

Some examples of when to use an AAR are: when you have introduced a new set of procedures or ways of working; after a busy winter season in which capacity was stretched; following the introduction of a new computer system; after a major training activity; after a shift handover; following a piece of research or a clinical trial; after performing surgery; etc.

AARs are excellent for making tacit knowledge explicit during the life of a project or activity and thus allowing you to capture it. Learning can be captured before a team disbands, or before people forget what happened and move on to something else. Despite the name ("after action"), they do not have to be performed at the end of a project or activity. Rather, they can be performed after each identifiable event within a project or major activity, thus becoming a live learning process in which lessons learned can be immediately applied. In fact this is where AARs can add the greatest value.

AARs provide insights into exactly what contributes to the strengths and weaknesses of a project or activity, including the performance of each individual involved, of the project leader, the team as a whole, and the various processes involved.

AARs are also a useful tool for developing your employees, which they do by providing constructive, directly actionable feedback in a non-threatening way because they are not linked to employee assessment. Similarly, they gives people an opportunity to share their views and ideas and to be heard.

4.3.3 How do I go about it?

AARs can be grouped into three types: formal, informal and personal. Although the fundamental approach involved in each is essentially the same, there is some variation in how they are conducted.

Formal AARs tend to be conducted at the end of a major project or event (learning after doing). They require some preparation and planning, but are not difficult as they take the form of a simple meeting. This meeting may take place over a couple of hours or a couple of days, depending on the scale of the project. Steps and tips for successful formal AARs include:

1. **Call the meeting as soon as possible and invite the right people**
   AARs should be conducted as soon as possible after the event. The reasons are simple – memories are fresh, participants are available and where appropriate, learning can be applied immediately. As well as the project manager and the key members of the project, it may be useful to invite the project client or sponsor and also members of any project teams who are about to embark on a similar project. However, be aware that the presence of external people may inhibit some team members.

2 **Create the right climate**
The ideal climate for an AAR is one of trust, openness and commitment to learning. AARs are learning events, not critiques, and so should not be treated as performance evaluation. There are no hierarchies in AARs – everyone is regarded as an equal participant and junior members of the team should feel free to comment on the actions of senior members. Make it clear that the purpose of the meeting is to help future projects run more smoothly by identifying the learning points from this project.

3 **Appoint a facilitator**
Ideally an AAR should be facilitated. (Certainly a formal AAR should be facilitated but informal AARs and personal AARs need not be so). The main purposes of the facilitator are to help the team to learn by drawing out answers, insights and previously unspoken issues; to ensure that everyone has an opportunity to contribute; and to help create the right climate and ensure that blame is not brought in. The facilitator should be someone who was not closely involved in the project, so that they can remain objective.

4 **Revisit the objectives and deliverables of the project**
Ask "what did we set out to do?" and "what did we actually achieve?". You might like to revisit the original project plan at this stage. You might also decide to construct a flow chart of what happened, identifying tasks, deliverables and decision points. This can help you to see which parts of the project were particularly effective or ineffective.

5 **Ask "what went well?". Find out why, and share learning advice for the future**
It is always a good idea to start with the positive points. Here you are looking to build on best practice as well as learning from mistakes. For each point that is made about what went well, keep asking a "why?" question. This will allow you to get to the root of the reason. Then press participants for specific, repeatable advice that others could apply in similar situations.

6 **Ask "what could have gone better?". Find out what the problems were, and share learning advice for the future**
Notice that you are not simply asking "what went wrong?" but rather "what could have gone better?". This way you can learn not only from mistakes, but also from any aspects of the project that got in the way of delivering even more. Hence the focus is not on failure, but on improvement. Even if no mistakes are made as such there is almost always scope for improvement. Again, for each point that is made, keep asking a "why?" question to get to the root of the reason. Then again, press participants for specific, repeatable advice that others could apply in similar situations: what would we do differently next time?

7 **Ensure that everyone feels fully heard before leaving the meeting**
It is important that participants do not leave the meeting feeling that they have not been heard or that things have been left unsaid. A useful technique here is to ask them for a numerical rating of the project: "looking back, how satisfied are you with the project: marks out of ten?". People who have said the project was fine will often still score it an eight, which enables you to then ask "what would have made it a ten for you?".

8 **Recording the AAR**
It is important to have a clear and interesting account of the AAR and its learning points, both as a reminder to those involved and in order to effectively share that learning with others. You should aim to include things like: lessons and guidelines for the future; some background information about the project to help put these guidelines into a meaningful context; the names of the people involved for future reference; and any key documents such as project plans or reports. Bear in mind who will be using your account and ask yourself if you were to be the next project leader, would this account and the lessons in it be of benefit to you?

9 **Sharing the learning**
As well as distributing your account of the AAR to the project team, you need to consider who else could benefit from it. For example, you may be aware of another team that is about to embark on a similar project. You also need to make your learning more widely available so that people working on similar projects in the future might also benefit; your document therefore needs to be stored somewhere it can be easily found and accessed by those it could help. This may be in a library, or in some kind of knowledge database or on an intranet.

Informal AARs tend to be conducted after a much smaller event such as a meeting or a presentation (learning after doing), or a following a specific event during a wider project or activity (learning while doing). They require much less preparation and planning and can often be done on the spur of the moment, as the format is simple and quick – a "pencil and paper" or flip chart exercise. In an open and honest meeting, usually no longer than half an hour, each participant in the event answers four simple questions:

> What was supposed to happen?
> What actually happened?
> Why were there differences?
> What did we learn?

Personal AARs are a simple matter of personal reflection. For example, take a few minutes to reflect on something you did yesterday such as a patient consultation, dealing with a complaint or making a specific telephone call. Ask yourself the four AAR questions above. What does that tell you about what you could do differently tomorrow?

4.3.4 Are there any other points I should be aware of?

> It is worth repeating is that AARs are learning events, not critiques. It is therefore vital that they are not treated as performance evaluation. The quality of an AAR depends on the willingness of participants to be open; this is unlikely to happen if they fear they are going to be assessed or blamed.
> Studies on the learning process show that the less time that elapses between discussing a lesson and applying it at work, the more effective the application. This would suggest that AARs are most valuable when used to "learn while doing".

4.3.5 More information

Books

Chapter 10: Networking and communities of practice
In: Learning to fly: practical lessons from one of the world's leading knowledge companies
Collison C, Parcell G. Oxford: Capstone, 2001

Web-sites

Post-mortem to living practice: After Action Review

After Action Reviews: 3 Step Process

Powerpoint presentation on After Action Reviews

AAR Case Studies

Step-by-step guide to writing AARs

U.S. Army Corps of Engineers guide to After Action Reviews

David Gurteen's Introduction to After Action Reviews

4.4 Communities of Practice

4.4.1 What are communities of practice?

A community of practice (CoP) is a network of people who share a common interest in a specific area of knowledge or competence and are willing to work and learn together over a period of time to develop and share that knowledge.

Etienne Wenger is credited with coining the term "community of practice" and he defines them as "groups of people who share a concern, a set of problems, or a passion about a topic, and who deepen their knowledge and expertise by interacting on an ongoing basis." He also believes that learning is a social activity and that people learn best in groups.

Communities can vary quite widely in their characteristics. Some exist for years while others form around a specific purpose and disband once that purpose has been achieved. Members may be very similar e.g. consultant gynaecologists, or they may be multi-disciplinary, such as is often the case in communities that are formed around addressing a specific challenge. Some may be small and localised while others will be geographically dispersed "virtual communities" that communicate primarily by telephone, e-mail, online discussion groups and video conferencing, etc.

Communities of practice differ from the usual notion of a team or work groups in a number of fundamental ways:

> Voluntary membership
Whereas teams and work groups are formed by management, membership of a community of practice is voluntary;

> Specific focus
Teams and work groups are formed to focus on a specific objective or activity, while communities of practice are not necessarily; they may have some stated goals, but they are more general and fluid;

> No expectation of tangible results
Teams and work groups are required to deliver tangible results, whereas communities of practice are not necessarily;

> Existence defined by group members
Teams and work groups are disbanded or reorganised once they have achieved their goals, while communities of practice last as long as their members want them to last.

Communities of practice exist in some form in every organisation – whether they have been deliberately created and labelled as such or not. The challenge for knowledge managers is to support them in such a way that they make a positive contribution to creating and sharing organisational knowledge.

Communities of practice are already being established in the NHS, based around the National Library for Health (NLH).

4.4.2 What are the benefits?

In her book "The Complete Idiot's Guide to Knowledge Management" Melissie Clemmons Rumizen calls communities of practice "the killer knowledge management application". Communities of practice:

> provide a valuable vehicle for developing, sharing and managing specialist knowledge;

> avoid reinventing the wheel;

> cut across departmental boundaries and formal reporting lines;

> can be more flexible than traditional reporting units;

> generate new knowledge in response to problems and opportunities;

> provide early warning of potential opportunities and threats;

> can be vehicle for cultural change (creating a knowledge sharing culture);

> are largely self-organising.

As well as the organisational benefits, communities of practice also provide benefits for individual community members, including:

- having access to expert help to expand horizons, gain knowledge and seek help in addressing work challenges;
- members often feel more conscious of, and confident in, their own personal knowledge;
- provides a non-threatening forum to explore and test ideas or validate courses of action;
- can foster a greater sense of professional commitment and enhance members" professional reputation.

4.4.3 How do I go about it?

There is a wide range of approaches to creating and developing communities of practice, and a wide range of resources full of guidelines and pointers. As a starting point, these are a few key areas for consideration:

Getting started

Communities of practice are organic and self-organising. Ideally they should emerge naturally. Organisations that have tried to create communities "from the top down" have often failed. Communities can however, be "seeded". Any area or function of your organisation where knowledge is not evenly distributed is a potential target for a community of practice. However, the impetus for a new community usually comes from the recognition of a special need or problem. From there, next steps will revolve around:

1. Defining the scope
   What is the domain of knowledge? At the heart of every community is a domain of knowledge; that domain can be either based around a professional discipline or on some specific problems or opportunities;
2. Finding participants
   Who can make a major contribution to this community? Who are the subject experts, and possible co-ordinators, facilitators, and librarians and/or knowledge managers? Will membership be open or by invitation only?
3. Identifying common needs and interests
   What are the core issues within the domain of knowledge? What are members interested in and passionate about? How do they hope to benefit from membership of the community?
4. Clarifying the purpose and terms of reference
   What are the specific needs or problems that need to be addressed? What is the community setting out to achieve? How will the community benefit the organisation? What are its values and ways of working? How will it be structured, organised and resourced?

It can often help to launch a community with a meeting or workshop so that members can meet each other and begin to develop relationships, and also spend some time together exploring and agreeing their purpose, terms of reference and ways of working.

Developing and sustaining

Once the initial enthusiasm of the set-up phase has passed, communities can easily wane and fade away unless they are actively developed and sustained.

1. Maintaining members' interest and involvement
   The ongoing success of a community depends on members' continued interest and involvement. A good co-ordinator will be constantly seeking to maintain that using a variety of methods. For example:
   - ensuring that members of virtual communities meet face to face at least once a year to keep personal relationships alive;
   - allowing plenty of time for socialising at gatherings;
   - ensuring that the wider organisation supports members in taking time to participate; motivating and rewarding people for their contribution;

- ➔ introducing new and challenging perspectives in the subject area from time to time, either from within the community or from external experts.

2 Growing the community
In the life of any community, members will come and go, and there will usually be a need for ongoing recruitment – both to replace lost members and to "keep things fresh". Similarly, roles and responsibilities will often be rotated between members over time. The ongoing success of the community will be affected by how well new members are welcomed and integrated into it.

3 Developing the body of knowledge
At this stage the community will probably be taking a more proactive and formal role in assuming responsibility for the relevant body of knowledge, with typical activities including:

- ➔ creating knowledge maps;
- ➔ organising knowledge resources;
- ➔ identifying and seeking to fill knowledge gaps;
- ➔ Here, the roles of librarians and/or knowledge managers will be particularly important.

4 Moving the agenda forward and adding value
Communities thrive when they are supported and valued by the organisation. This is a "two-way street" so it is important that a community develops in alignment with overall organisational goals, rather than to its own agenda. This will increase the chances of ongoing support from the organisation, such as:

- ➔ providing resources;
- ➔ recognition and reward of community members and particularly co-ordinators;
- ➔ help in removing barriers to community membership;
- ➔ and involvement of communities in key management decisions and problem-solving

However, at the same time, care is needed not to "over-formalise" or "institutionalise" the community.

Closure

Communities can naturally fade away and this is not always a bad thing. Sometimes a natural ending is reached – a group of people or a practice reach a natural conclusion. Other times a community can break up and in its place, a number of "sub-communities" based around particular specialist subjects emerge. Either way, when a community fades it is important to celebrate its life and achievements, and to ensure that the relevant body of knowledge is captured and/or transferred.

4.4.4 Are there any points I should be aware of?

- The successful cultivation of communities of practice requires a fine balance between giving them enough support and direction to ensure their value, while at the same time not imposing too much structure and therefore risking losing the informal social relationships that underpin their effectiveness.
- Successful communities or practice require a simultaneous focus on two key areas:
  - ➔ developing the practice;
  - ➔ and, developing the community.

Developing the community involves a focus on the social structure – the sum of the social relationships built up within a community. Particular roles that are important and should be explicitly recognised are:

> **Leader (or coordinator)**
> recognised in the organisation at large as the spokesperson for this community; organises and co-ordinates the community's interactions and activities;

> **Facilitator(s)**
> Facilitates the interactions within the community, e.g. in face-to-face meetings, and steers the agenda of online interactions;

> **Librarian or knowledge manager**
> manages the explicit knowledge resources of the community.

You might consider providing training and support for these roles, for example in co-ordination and moderation techniques.

Developing the practice looks at the community's inputs and outputs – the resources that the community uses and develops. These resources consist not only of information and knowledge resources such as documents, databases, a web-site, etc. but also the processes and practices within the community. These include ways of developing and enhancing the knowledge base such as through peer group reviews of emerging best practice, and ways of communicating new knowledge developed within the community to the wider organisation. Many communities are becoming the focal point within their organisations for documenting best practice, identifying valuable external resources, writing case studies, and developing frameworks, techniques and tools for their particular knowledge domain.

4.4.5 Resources and references

Books

Chapter 8: Communities of practice – the killer application
In: The complete idiot's guide to knowledge management
Clemmons Rumizen, M.(2002) Madison, WI: CWL Publishing Enterprises

Chapter 10: Networking and communities of practice
In: Learning to fly: practical lessons from one of the world's leading knowledge companies.
Collison C, Parcell G.(2001) Oxford: Capstone

Cultivating communities of practice
Wenger E. (2002) Massachusetts: Harvard University Press

Articles

Knowledge is the enemy of disease
Brice, A and Gray, M. CILIP Update, 2003, March

Working together – communities of practice in family medicine
Endsley, S and Kirkegaard, M and Linares A. Family Practice Management, 2005, January

A survey of current research on online communities of practice
Johnson, CM. Internet and Higher Education, 2001, 4(1)

A sense of community: the role of CoPs in knowledge management
Lelic, S. Knowledge Management, 2001, 10 October

Communities of practice and organizational performance
Lesser, EL and Storck, J. IBM Systems Journal, 2001, 40(4)

Web-sites

Etienne Wenger and Communities of Practice

Tools

Getting to 7 – Cultivating Communities of Practice: the 7 Stages of Development

4.5 Conducting a knowledge audit

4.5.1 What is a knowledge audit?

The term "knowledge audit" is in some ways a bit of a misnomer, since the traditional concept of an audit is to check performance against a standard, as in financial auditing. A knowledge audit, however, is a more of a qualitative evaluation. It is essentially a sound investigation into an organisation's knowledge "health". A typical audit will look at:

> What are the organisation's knowledge needs?

> What knowledge assets or resources does it have and where are they?

> What gaps exist in its knowledge?

> How does knowledge flow around the organisation?

> What blockages are there to that flow e.g. to what extent do its people, processes and technology currently support or hamper the effective flow of knowledge?

The knowledge audit provides an evidence-based assessment of where the organisation needs to focus its knowledge management efforts. It can reveal the organisation's knowledge management needs, strengths, weaknesses, opportunities, threats and risks.

4.5.2 What are the benefits?

Benefits of a knowledge audit include:

> helping the organisation clearly identify what knowledge is needed to support overall organisational goals and individual and team activities.

> giving tangible evidence of the extent to which knowledge is being effectively managed and indicates where improvements are needed.

> providing an evidence-based account of the knowledge that exists in an organisation, and how that knowledge moves around in, and is used by, that organisation.

> presenting a map of what knowledge exists in the organisation, and where it exists, revealing both gaps and duplication.

> identifying pockets of knowledge that are not currently being used to good advantage and therefore offer untapped potential.

> providing a map of knowledge and communication flows and networks, revealing both examples of good practice and blockages and barriers to good practice.

> presenting an inventory of knowledge assets, allowing them to become more visible and therefore more measurable and accountable, and giving a clearer understanding of the contribution of knowledge to organisational performance.

> supplying vital information for the development of effective knowledge management programmes and initiatives that are directly relevant to the organisation's specific knowledge needs and current situation.

Some examples of situations in which a knowledge audit can be beneficial include:

> you are about to embark on creating a knowledge management strategy and so need to establish exactly "where you are now"

> people are having difficulty in finding the information and knowledge they need to make key decisions

> useful sources of information and knowledge are frequently stumbled across by accident

> there is duplication of information and knowledge gathering activities across different departments or teams, and hence duplication of costs

> questions are being raised about the value of knowledge management systems, initiatives or investments

> when findings from research and development are not making their way into practice quickly enough.

4.5.3 How do I go about it?

There are a wide variety of approaches to conducting a knowledge audit, with varying levels of coverage and detail. As a general rule, most knowledge audits will involve some or all of the following:

Identifying knowledge needs

The first step in most knowledge audits involves getting clear about precisely what knowledge the organisation and the people and teams within it need in order to meet their goals and objectives.

A knowledge audit provides a systematic way of finding this out to some level of detail. Common approaches taken to collating this information include questionnaire-based surveys, interviews and facilitated group discussions, or a combination of these.

In asking people about knowledge needs, it is important to provide a point of focus, as “knowledge” can be seen as being quite conceptual and therefore difficult to articulate. To get around this, and to ensure that you are concentrating on vital knowledge, invite people to think about their goals and objectives, and the core processes, activities and decisions that they perform in the course of their day-to-day work. You might ask them to also consider their main problems and challenges, and how might faster access to better knowledge help them in that regard.

It is always beneficial to begin a knowledge auditing process with identifying knowledge needs. This enables you to then use your understanding of these needs to guide the rest of the auditing process, and therefore be sure that you are focusing on the knowledge that is important to the organisation.

Drawing up a knowledge inventory

A knowledge inventory is a kind of stock-take to identify and locate knowledge assets or resources throughout the organisation. It involves counting and categorising the organisation’s explicit and tacit knowledge. In the case of explicit knowledge, this will include things like:

> what knowledge we have – numbers, types and categories of documents, databases, libraries, intranet websites, links and subscriptions to external resources etc.?

> where the knowledge is – locations in the organisation, and in its various systems?

> organisation and access – how are knowledge resources organised, how easy is it for people to find and access them?

> purpose, relevance and “quality” – why do these resources exist, how relevant and appropriate are they for that purpose, are they of good “quality” e.g. up-to-date, reliable, evidence-based etc.?

> usage – are they actually being used, by whom, how often, what for?

In the case of tacit knowledge, the inventory will focus on people and look at things like:

> Who we have – numbers and categories of people

> Where they are – locations in departments, teams and buildings

> What they do – job levels and types

> What they know – academic and professional qualifications, core knowledge and experience

> What they are learning – on the job training, learning and development.

The knowledge inventory gives you a snapshot of your knowledge assets or resources. By comparing your inventory with your earlier analysis of knowledge needs, you can begin to identify gaps in your organisation’s knowledge as well as areas of unnecessary duplication. This is also explored in greater detail in the next step.

Analysing knowledge flows

While an inventory of knowledge assets shows what knowledge resources your organisation has, an analysis of knowledge flows looks at how that knowledge moves around the organisation – from where it is to where it is needed. In other words, how do people find the knowledge they need, and how do they share the knowledge they have? Again, the knowledge flow analysis looks at both explicit and tacit knowledge, and at people, processes and systems:

The relative focus in this stage is on people: their attitudes towards, habits and behaviours concerning, and skills in, knowledge sharing and use. This will usually require a combination of questionnaire-based surveys followed up with individual interviews and facilitated group discussions.

In terms of processes, you will need to look at how people go about their daily work activities and how knowledge seeking, sharing and use are (or are not) part of those activities. In most organisations, there will be pockets of good knowledge management practice (though they may not be called knowledge management). You will also need to look at what policies and practices currently affect the flows and usage of information and knowledge, for example are there existing policies on things like information handling, records management, web publishing? Are their other wider policies and practices that, while not directly related to knowledge management, act as enablers or barriers to good knowledge practice?

On the systems side, some assessment is needed of key capabilities that will be used in any recommended actions or solutions. This includes the technical infrastructure: information technology systems, content management, accessibility and ease of use, and current actual levels of use. In short, to what extent do your systems effectively facilitate knowledge flows, and help to connect people with the information and other people they need.

An analysis of knowledge flows will allow you to further identify gaps in your organisation's knowledge and areas of duplication; it will also highlight examples of good practice that can be built on, as well as blockages and barriers to knowledge flows and effective use. It will show where you need to focus attention in your knowledge management initiatives in order to get knowledge moving from where it is to where it is needed.

Creating a knowledge map

A knowledge map is a visual representation of an organisation's knowledge. There are two common approaches to knowledge mapping:

1. The first simply maps knowledge resources and assets, showing what knowledge exists in the organisation and where it can be found
2. The second also includes knowledge flows, showing how that knowledge moves around the organisation from where it is to where it is needed.

Clearly the second approach provides the most complete picture for the knowledge auditor. However, the first is also useful, and in some organisations is made available to all staff to help people locate the knowledge they need.

4.5.4 Are there any other points I should be aware of?

- Be clear about your purpose. The knowledge audit is not a quick or simple process, and so the time and effort required needs to be justified by a clear purpose and a set of actions that will be taken as a result of what the audit reveals.
- When conducting a knowledge audit, bear in mind the widely-accepted statistic that around 80% of an organisation's knowledge is tacit, hence beware of focusing too much time and energy on explicit knowledge and not enough on tacit knowledge.
- The ease or difficulty that you have in gathering and collating the information you need as part of the audit process is itself a good indicator of the status of your current knowledge management capabilities.
- If you decide to commission a knowledge audit from external consultants, be aware that the quality and depth of work that comes under the general banner of "knowledge auditing" varies quite. Many vendors use the term "knowledge audit" to describe what is in fact an information audit – which will only look at explicit knowledge. Auditing tacit knowledge is probably where the greater challenge lies, and is hence the area in which expert help is likely to be most valuable.

4.5.5 More information

Know Map – the Knowledge Management, Auditing and Mapping Magazine, has a range of resources about knowledge auditing available on their site.

4.6 Developing a knowledge management strategy

4.6.1 What is a knowledge management strategy?

A knowledge management strategy is simply a plan that describes how an organisation will manage its knowledge better for the benefit of that organisation and its stakeholders. A good knowledge management strategy is closely aligned with the organisation's overall strategy and objectives.

4.6.2 What are the benefits?

A good, clear knowledge management strategy can help to:

> increase awareness and understanding of knowledge management in your organisation
> articulate the business case and identify potential benefits
> gain senior management commitment
> attract resources for implementation
> communicate good knowledge management practice
> give you a clear, communicable plan about where you are now, where you want to go, and how to plan to get there
> give you a basis against which to measure your progress

4.6.3 How do I go about it?

There are many ways to approach the development of a knowledge management strategy, as well as many ways of presenting the strategy document itself – there is no "one size fits all". Larger organisations will probably need a detailed, formal strategy document whereas for a smaller organisation something briefer and less formal might be more appropriate.

The strategy document

As a general guideline, a strategy of any kind tends to include answers to three key questions: where are we now, where do we want to be, and how do we get there?

A relatively brief and informal knowledge management strategy might be structured around these three questions and include things like:

1. Where are we now?
   An assessment of the current situation. How does current knowledge management practice (or lack of it) affect the organisation's ability to meet its goals? How does it affect the effectiveness of individuals and teams? To what extent do the organisation's culture, processes and systems currently act as enablers of, or barriers to, good knowledge management practice?

2. Where do we want to be?
   An outline of what knowledge management will do for the organisation. How will it help the organisation and the people in it to meet their objectives? What might "good knowledge management practice" look like for this organisation specifically? How will you know when you are there i.e. how will you measure the progress and value of your efforts?

3. How do we get there?
   Describing the specific actions that will be taken to get to where you want to be. An action plan covering the three key elements of people, processes and technology: what specific knowledge management tools and processes will you use; how will you motivate people and realign your organisational culture to a "knowledge friendly" one, and how will you develop the supporting technological infrastructure? Also needs to include details of resources required, deliverables, time-scales and responsibilities.

For the larger organisation requiring a more formal and detailed strategy, David Skyrme (http://www.skyrme.com/) suggests the following format:

- > **Executive summary** – no more than 1-2 pages.
- > **Background** – giving sufficient context about what initiated this strategy and where this document fits within the wider context.
- > **The case for knowledge management** – starting with your organisation's definition of knowledge management, then explaining the contribution that better knowledge management will make to your organisation, based on core organisational objectives.
- > **Current knowledge management situation** – highlight existing knowledge management activities and experience, outlining the benefits and explaining how these can be built upon; expose barriers to further progress.
- > **Stakeholders' challenges and knowledge needs** – summarise the key issues and knowledge needs of the organisation and relevant stakeholders (e.g. leaders, staff, patients, relevant NHS and government authorities etc.); include an assessment of the existing quality and accessibility of knowledge resources.
- > **Knowledge management vision and strategy overview** – it is often useful to encapsulate an inspiring vision and mission in one or two sentences each; this is followed by some key knowledge management objectives.
- > **Details of strategy** – outline the list of activities and projects to be implemented; it is useful to group these into specific themes or areas of action; typical themes might include: knowledge management tools and techniques; people and cultural aspects; knowledge management skills development; technology; leadership and governance (who will own and drive the strategy); communications (how will the strategy be promoted and rolled out); and measurement (how will performance and progress be measured).
- > **Action plan** – give details of deliverables, time-scales, resources and budgets required for all actions, and reiterating the benefits.
- > **Dependencies** – highlight critical dependencies such as the availability of key personnel, approval of budgets etc.; also spell out the impact of "doing nothing".
- > **Conclusions/Next Steps** – a simple outline of what needs to happen next to move the agenda forward and translate the strategy into action.
- > **Appendices** – typical appendices might include the findings of a knowledge audit, some background material on knowledge management such as definitions, summaries of any existing knowledge management projects or initiatives, etc.

Developing your strategy

In developing a knowledge management strategy, various practitioners offer a range of tips, some of which are outlined here:

1 **Start with your organisation's strategy and objectives**
The most important factor in guiding a knowledge management strategy is the organisation's overall strategy and goals. Given that the whole purpose of knowledge management is to help the organisation to achieve its goals, the knowledge management strategy should describe precisely that. In order to do that, you need to understand what your organisational goals are, and how you are currently performing against them. Talk to key people throughout your organisation about strategy and goals. Look at what various departments or functions are doing. Discuss plans for the future, and look at factors that influence reaching goals. Get a feel for how sub-optimal knowledge management might be currently limiting the organisation in achieving its goals, and how better knowledge management might help it to achieve them.

Look for gaps that could prevent the organisation from achieving its goals. As you talk to people, be on the look out for the issues that are really causing them problems – their "pains". As well as problems, look for opportunities – not only the chance to fix things, but also the chance to do something new or better. Needs, problems, pains and opportunities give you an opening to use knowledge to make a difference.

As well as being be an integral part of the wider organisational strategy, a knowledge management strategy should also be coherent with human resources and information technology strategies.

2 **Conduct a knowledge audit**
A knowledge audit is an investigation into an organisation's knowledge management "health". A typical audit will look at:

- ➔ What are the organisation's knowledge needs?
- ➔ What knowledge assets or resources does it have and where are they?
- ➔ What gaps exist in its knowledge?
- ➔ How does knowledge flow around the organisation?
- ➔ What blockages are there to that flow?
- ➔ To what extent do its people, processes and technology currently support or hamper the effective knowledge management?

The knowledge audit can reveal the organisation's knowledge management needs, strengths, weaknesses, opportunities, threats and risks. It provides an evidence-based assessment of where the organisation needs to focus its knowledge management efforts.

3 **Think about people, processes and technology**
When planning your approach to knowledge management, be sure to address each of the three key aspects of people, processes and technology. It is often said that any knowledge management strategy that does not incorporate all three is destined to fail.

4 **Think about capturing versus connecting**
A key decision in developing your strategy and in selecting knowledge management tools and techniques involves looking at the relative focus on explicit and tacit knowledge – in other words, do you want to focus on connecting people with information, or on connecting people with people? Of course this is not an "either/or" decision and most knowledge management strategies tend to involve a combination of the two; the optimal balance between them will depend on your organisational context.

5 **Balance a long-term vision with quick wins**
A good strategy will reflect a balance between "quick-wins" and building a sustainable knowledge management capability into the long-term. The advantage of quick wins is that they allow people to see immediate benefits, and therefore they are more likely to give their support.

As well as seeking a number of quick wins, try not to be over-ambitious in the short to medium-term. Avoid long lists of things to do. You cannot change an organisation culture and ingrained work habits overnight. Pick a few core activities where you can make a difference, and prioritise and focus on those. At the same time, do keep your long-term vision in view.

6 **What's in it for me?**
Gaining support and acceptance for your strategy and ultimately embedding knowledge management into the organisation is about winning "hearts and minds". Think constantly about addressing the "what's in it for me?" question that those whose contribution is needed will invariably ask (and even if they don't ask it in so many words, you can be fairly sure they are thinking it). Always anticipate that question from all of those involved – senior managers, budget-holders, middle managers, staff, patients, those departments and functions whose support you will need such as human resources and information technology. In answering the "what's in it for me?" question, consider the three key levels of "me": myself, my team/department/function, and my organisation as a whole.

7 **Build the evidence with pilots**
The vast majority of knowledge management practitioners who have learned from direct experience strongly recommend using a pilot project as a "test bed" before launching any new knowledge management initiatives. Pilots have a number of advantages: they allow you to test an approach with a small group of users to find what works and what doesn't, and to refine your approach and "get it right" before rolling out across the wider organisation. This means that when rolling out, you already have evidence to demonstrate that what you are advocating actually works in practice. Similarly, your learning and "mistakes" have taken place in a contained environment, so they will not have a negative impact on the organisation as a whole view of knowledge management. You are

therefore strongly advised to build pilot projects into your knowledge management strategy before seeking to launch any major new initiatives.

4.6.4 Are there any other points I should be aware of?

- It is important to define precisely what knowledge management means for your organisation. There is no single agreed definition "out there" and given that knowledge management as a concept essentially borrows from a range of other disciplines, there is a great deal of misunderstanding about what is actually involved. People from an information management background might have one viewpoint; those working in information technology will tend to have another, those in human resources another still, etc. A clear and common understanding of what it means in your organisation is therefore essential.
- Don't think you have to wait until you have a knowledge management strategy in place before you can "do" knowledge management. More often than not, knowledge management initiatives begin before there is a strategy. In fact many practitioners actively advocate it, believing that a strategy only becomes appropriate once knowledge management initiatives have "had their honeymoon period" and are ready to be formally organised and endorsed.
- A common mistake is a strategy that is too theoretical and "dry". Many knowledge management strategies read as if they have come straight from a textbook (and some probably have). Your strategy needs to be "real', written in the language of your organisation, and relevant to your organisation's situation. Similarly, be creative in making it interesting and bringing it alive.
- Again, don't forget the "what's in it for me?" question. Clearly demonstrate the benefits of knowledge management throughout your strategy. How will it reduce costs and time, improve performance, increase efficiency, reduce risk, etc.? Use real examples.

4.7 Exit interviews

4.7.1 What are exit interviews?

Traditionally, exit interviews are conducted with employees leaving an organisation. The purpose of the interview is to provide feedback on why employees are leaving, what they liked or didn't like about their employment and what areas of the organisation they feel need improvement. Exit interviews are one of the most widely used methods of gathering employee feedback, along with employee satisfaction surveys.

More recently, the concept of exit interviewing has been revisited and expanded as a knowledge management tool, as a way of capturing knowledge from leavers. Rather than simply capturing human resources information, the interview also aims to capture knowledge about what it takes to do the job.

4.7.2 What are the benefits of exit interviews?

> vital knowledge is not lost to the organisation when people leave

> the learning curve of new people joining the organisation is shortened

> they can be done relatively quickly and inexpensively

> they can result in the leaver having a more positive view of the organisation

Done correctly, exit interviews can be a win-win situation for both the organisation and the leaver. The organisation gets to retain a portion of the leaver's knowledge and make it available to others, while the leaver gets to articulate their unique contributions to the organisation and to "leave their mark".

4.7.3 How do I go about it?

Traditional exit interviews can be conducted in a variety of ways: face-to-face, over the telephone, using a written questionnaire, or via the Internet using an exit interview management system. In a knowledge-focused exit interview, a face-to-face interview is needed.

You will need to think carefully about the information you would like to gather before the interview and start your preparations early. While the traditional exit interview will tend to collect mainly human resources information, the primary focus of the knowledge-focused interview is on knowledge that would be helpful to the next person who will do the job or to others in the organisation doing similar jobs.

Start planning the handover and exit interview as soon as you know a person is leaving. Identify who in the organisation might benefit from that person's knowledge and what they will need to know. Then work out a plan to capture the leaver's knowledge during the time remaining before they leave. This should include both explicit knowledge (knowledge that is already documented such as in files and e-mails, and knowledge that can be easily documented), and tacit knowledge (knowledge that is less easy to capture and that needs to be explained or demonstrated).

In the case of explicit knowledge, make sure the leaver moves relevant files – both hard copy and electronic – into shared folders or a document library. Ask them to prune and organise these files and to create role and task folders or notes for their successor.

For tacit knowledge, you will need to interview the leaver face-to-face. Prepare for the interview by reviewing the key tasks the person does based on a job description or annual performance plan. You can then use that information as the basis for discussing how they go about those tasks, what knowledge and skills they need, any problems or pitfalls to be aware of etc. Find out about their network of contacts and sources of knowledge. If possible, create an overlap period between the leaver and their successor so that a "live" handover can be done.

When conducting exit interviews, think carefully about who will be the interviewer. Someone from the Human Resources Department conducts traditional exit interviews. However this need not be the case in the knowledge-focused interview. Often a peer or a relevant subject expert will be most appropriate. Over and above the obvious interpersonal and interviewing skills needed, you will need to consider issues of trust and honesty. For example, if an employee has had a difficult relationship with a manager or colleague, that person might not be best placed to conduct the interview. Whoever you select, make sure they are appropriately skilled and trained.

4.7.4 Are there any other points I should be aware of?

- Traditional exit interviews are usually only appropriate for employees who voluntarily resign or retire rather than those who are fired or made redundant. In the case of the knowledge-focused interview, much will depend on the extent to which the organisation has a culture that encourages knowledge sharing.
- Be clear about who will use the knowledge gathered and how it will be used, before you begin to gather it; the purpose of the interview is not to gather knowledge per se, but to gather useful knowledge that will actually be used.
- The less you capture knowledge on a regular basis, the more you need to capture it at exit. However you may decide that you could gain more value from capturing knowledge at more regular intervals. For example, The Post Office uses exit interviews as one part of a series of "cradle-to-grave" interviews to collect knowledge, using a method called 3E. The three Es are Entry, Expert and Exit. Entry interviews allow you to gather knowledge when employees first join the organisation when they have "new eyes" and a fresh perspective, and also to ask them what they would like to know to help them "get up to speed". Expert interviews are conducted as they develop skills and become experts in a particular role or field. For more information about this wider approach, see knowledge harvesting.

4.7.5 More information

Leverage exit interviews to collect key knowledge
by Pamela Holloway – Workforce Management.
(Registration required but this is free)

Tips and Techniques for Effective Exit Interviews
by Pamela Holloway.

Disappearing knowledge: are exit interviews the wit's end?
by David Skyrme – I3 Update, 2001, November, No 55

4.8 Identifying and sharing best practices

4.8.1 What is identifying and sharing best practices?

The sharing of practices is often one of the first things to be carried out in a knowledge management initiative. In most organisations it is already being done to some degree. This often begins with common practices such as instruction manuals or "how to" guidelines. The next step from there is to identify and share best practices.

A best practice is simply a process or a methodology that represents the most effective way of achieving a specific objective. Some people prefer to use the term "good practice" as in reality it is debatable whether there is a single "best" approach – and, of course, approaches are constantly evolving and being updated. So, another way of defining a best practice is one that has been proven to work well and produce good results, and is therefore recommended as a model.

Much of best practice knowledge is tacit – held in people's heads and not always easy to document. Therefore, most best practice programmes combine two key elements: explicit knowledge such as a best practices database (connecting people with information), and methods for sharing tacit knowledge such as communities of practice (connecting people with people). These two approaches are complementary. A database can provide enough information for a potential user of the best practice to find it and decide if it is worth pursuing further. However, the best way of sharing best practices is "on the job" and so communities and personal contact with others who have used the best practice is key.

4.8.2 What are the benefits?

The essence of identifying and sharing best practices is to learn from others and to re-use knowledge. Effective sharing of best practices can help organisations to:

> identify and replace poor practices
> raise the performance of poor performers closer to that of the best
> avoid reinventing the wheel
> minimize re-work caused by use of poor methods
> save costs through better productivity and efficiency
> improve services to patients

Best practice programmes are most appropriate in organisations where processes are quite well developed and where a certain amount of knowledge and experience has been accumulated. They are most useful where an organisation has several units or people performing similar tasks but who are widely dispersed and so do not tend to learn from each other through day-to-day contact.

4.8.3 How do I go about it?

In "Best Practices in Best Practices", David Skyrme recommends a 6-step approach to identifying and sharing best practices. This is summarised here. The overall approach is aimed at documenting the essential features of a best practice, giving pointers to relevant experts in that practice, deducing general guidelines, diffusing basic knowledge, and using subject matter experts to apply and adapt the practices in a new context.

The key steps are as follows:

1 **Identify users' requirements**
This step may sound obvious, but it is not uncommon for someone given the task of capturing best practices to start by designing a database, when clearly this is a case of putting the cart before the horse. Start by considering where you can really add value. Look at what areas of the organisation need attention because of poor performance or difficult challenges. Who can most benefit from better knowledge and understanding of best practices? How will they access and use them?

2 **Discover good practices**
There are various methods of identifying best practices. One approach is to look at who is producing excellent results and is therefore likely to be using good practices. Having discovered these people, you will then need to discern which parts of their overall approach or methods being used are relevant practices such as subject matter experts, internal auditors, consultants and peers. A range of alternative approaches for identifying best practices can be found within various knowledge management tools. These include communities of practice, after action reviews, knowledge harvesting and exit interviews. Don't necessarily limit your search to only include practices within your organisation; much can be learned from the practices of other organisations in your field, or even organisations in other industries.

3 **Document good practices**
Best practice descriptions are usually kept in a database in a standard format. A typical template might include the following sections:

- ➔ **Title**: short, descriptive title; this can be accompanied by a short abstract.
- ➔ **Profile**: several short sections outlining processes, function, author, keywords, etc.
- ➔ **Context**: where is this applicable? What problems does it solve?
- ➔ **Resources**: what resources and skills are needed to carry out the best practice?
- ➔ **Description**: what are the processes and steps involved?
- ➔ **Improvement measures**: are there performance measures associated with this practice?
- ➔ **Lessons learned**: what proves difficult? What would the originators of the practice do differently if they were to do it again?
- ➔ **Links to resources**: experts contact details, workbooks, video clips, articles, transcripts of review meetings, tools and techniques used.

The aim at this stage is not to describe the practice in great detail, but to give enough information to allow users of the database to decide whether it matches their needs and where they can find further information. A key consideration is how you can organize and classify the information in your database so that users can readily find what they need.

4 **Validate best practices**
A practice is only "good" or "best" if there is a demonstrable link between what is practised and the end result. In most organisations, and especially in areas where practices are constantly evolving, rigorous cause-and-effect analysis is impracticable. Hence a degree of subjective judgement is needed as to what constitutes "best". A common approach is to have a panel of reviewers compromising internal and external subject experts and peers, who evaluate a potential best practice against their knowledge of existing practice. It is equally important to ensure that you seek input and feedback from customers (i.e. the ultimate beneficiaries, such as patients) of the best practices.

In the context of the NHS, a further important consideration is that of evidence-based practice. When identifying and validating best practices, it is important to ensure that these are based on a combination of both on-the-job experience and sound research evidence.

5 **Disseminate and apply**
While a database of best practices is a useful starting point, most organisations find it essential to complement this with face-to-face knowledge sharing about those best practices. This is where the real value is added. Not only does it help the recipient dig beneath the explicit knowledge and gain more in depth insights, but it can also provide a two-benefit in that dialogue between the conveyor of best practice knowledge and the recipient can enrich the knowledge of both.

Common ways of sharing best practice knowledge include: communities of practice; improvement groups or quality circles in which teams within an organisation meet regularly to discuss ways of improving a process; visits to other departments or organisations with good performance; organised learning events such as share fairs or knowledge cafés, that bring people together to share specific knowledge and experience; job secondments or exchanges; etc.

6 **Develop a supporting infrastructure**
To successfully implement a best practice programme, you need to ensure you have the required infrastructure in place. This infrastructure is often developed as part of a wider knowledge management strategy. Typically, several generic aspects need attention:

- ➔ The people to facilitate and drive the process through its initial stages, until it becomes embedded in the organisation's ways of working (e.g. a best practices team, or a network of best practices co-ordinators).
- ➔ The technical infrastructure for document sharing and databases.
- ➔ The content management infrastructure to ensure that best practices are documented and classified electronically in a way that makes them easy to find.

4.8.4 Are there any other points I should be aware of?

- > Establishing a programme to identify and share best practice is not generally a "quick fix" solution for organisations that are relatively new to knowledge management. Setting up the required processes and infrastructure can be quite a big task, unless you already have some aspects of a knowledge management infrastructure in place.
- > As with an knowledge management initiative, don't forget the importance of motivation and culture. The ease with which good practices emerge and are shared depends on the culture of your organisation. If there is a "not invented here" culture, then good practices will be slow to emerge and spread, as each part of the organisation will defend its own way of doing things rather than learning from, and shearing with, others. Where people are generally encouraged to seek out knowledge and learning, best practices are more likely to emerge and spread.
- > Try not to get too prescriptive about best practices. Rather than putting in rigid rules that say "this is best practice and you should follow it", focus more on encouraging people to develop and share best practices voluntarily.
- > Do not make the mistake of focusing on capturing best practices for the sake of capturing them. Focus on how they can be used to add value. Who are the users? What are their issues? What kind of knowledge do they need to perform better? How might they best assimilate that knowledge?
- > You will need to actively promote your best practice resources. Otherwise you may end up with databases and people that are under-used and not fulfilling their potential.
- > Be sure to demonstrate the benefits and the evidence. Use case examples to show the benefits of sharing best practices, and as far as possible demonstrate how a best practice has contributed to better performance.
- > Remember that best practice is constantly evolving. Therefore feedback mechanisms must be built in so that the value of existing best practices is constantly assessed, and feedback used to create further improvements.
- > Resist the temptation to focus on explicit knowledge – it cannot be emphasised enough that databases of best practices are insufficient. Databases point to examples and people, but it is through people that deep knowledge is transferred.

4.9 Knowledge centres

4.9.1 What are knowledge centres?

In short, an enhanced version of a library. The "enhancement" lies in a wider focus on knowledge as well as on information: a knowledge centre typically provides a focus for collecting, organising and disseminating both knowledge and information. This does not necessarily mean that the knowledge centre will actually perform all of these activities itself. Rather, it will create a framework and provide leadership, co-ordination, guidance and expertise.

4.9.2 What are the benefits?

A knowledge centre can bring core knowledge management responsibilities and activities under a single umbrella rather than leaving it to dispersed individuals and teams. Economies of scale can therefore be achieved through:

> avoiding duplication of effort and resources;

> pooling expertise;

> achieving bulk purchasing discounts;

> reusing knowledge and information in a variety of contexts.

4.9.3 How do I go about it?

The services that a typical knowledge centre might provide include:

> Maintaining and developing knowledge repositories e.g. the organisation's intranet, key information databases and collections.

> Providing content management services such as cataloguing, indexing and developing taxonomies for electronic knowledge repositories.

> Coordinating the capturing of knowledge from projects and assignments and incorporating it into knowledge bases such as databases of best practices and/or case studies.

> Identifying and forming links with sources of important knowledge, both inside and outside the organisation.

> Providing pointers to people as well as to information – connecting people who need help with people who can provide it, identifying subject experts, maintaining a skills database, connecting people who share similar needs or are working on similar problems, etc.

> Providing a "one stop shop" for multiple knowledge and information needs.

> Providing pointers to resources and/or training in information and knowledge skills.

Good knowledge centres will put as much emphasis on connecting people with people – "know-who" – as they do on connecting people with information and document collections. They will be concerned with "active" not "archive" knowledge, so need to be fully up to speed with what is happening in the organisation including current priorities and work in progress – "who is doing what now".

Knowledge centres may also be created for very specific goals. For example in the 1990s, consulting firm Ernst & Young created three knowledge centres, each with a distinct remit:

1. the Centre for Business Innovation would create new knowledge through research,
2. the Centre for Business Technology would use existing knowledge to create predefined methods and automated tools, and
3. the Centre for Business Knowledge would gather and store the firm's internal and external knowledge and information resources.

The services of the latter included a library, a call centre for answering consultant requests, and a database of consultant skills. Managers of the centre also had responsibility for identifying and tracking subject matter experts, and for organising knowledge networks around each key domain of knowledge within the business. Another key task of the centre was to develop a knowledge architecture and taxonomy, in order to specify the categories and terms in which the firm needed to gather and store knowledge. Key areas of knowledge were represented by "Power Packs" – structured sets of online resources and materials including answers to frequently encountered issues.

4.9.4 Are there any other points I should be aware of?

Knowledge centres, while similar to libraries, are not the same. A knowledge centre is based on the idea that knowledge resides primarily in people rather than in documents or computer systems. Hence in a knowledge centre, there is a strong emphasis on connecting people with each other, as well as with information.

4.10 Knowledge harvesting

4.10.1 What is knowledge harvesting?

Knowledge harvesting is an approach that allows the tacit knowledge or know-how of experts and top performers in an organisation to be captured and documented. This know-how can then be made available to others in various ways such as through training programmes, manuals, best practices and knowledge management databases. Knowledge in organisations exists in two forms: explicit knowledge, which is easily captured and shared; and tacit knowledge, which is more experiential and intuitive, and so is less easy to articulate. Knowledge harvesting is about trying to make some of the tacit knowledge more explicit. Its aim is to help organisations make better and wider use of their existing knowledge by extracting it from the heads of a few key people and making it available to a much wider range of people.

4.10.2 What are the benefits?

The ultimate goal of knowledge harvesting is to capture an expert's decision-making processes with enough clarity that someone else could repeat the same processes and get the same results.

Knowledge harvesting can be effectively used in a range of situations such as:

- When an organisation wants to "know what it knows".
- When knowledge and information are needed for a specific, clearly defined, purpose.
- To capture the knowledge of employees who are leaving the organisation or department.
- To gather knowledge to support a process of change or improvement.
- To kick-start a knowledge management programme by quickly generating a body of expert knowledge about a subject and making it available across the organisation.
- As an ongoing practice, as part of a wider knowledge management strategy.

The benefits of knowledge harvesting include:

- The knowledge of a few key individuals is made readily available to others who need it.
- Individuals can access experts' knowledge when and where they need it, without being dependent on the availability of that expert.
- Vital knowledge is not lost to the organisation when people leave.
- The learning curve of new people joining the organisation is shortened.
- The tangible knowledge assets of the organisation can be increased.
- Productivity and efficiency can be improved, as people can use existing expertise rather than having to go through their own trial-and-error experiences.
- It can be done relatively quickly and inexpensively.

4.10.3 How do I go about it?

While there is no set formula for knowledge harvesting, there are some general guidelines that facilitate the process. These can be broken down into a number of steps.

Focus

Decide on what specific knowledge and expertise you want to capture, and be clear about what the benefits will be. It is neither possible nor desirable to capture everything that everyone knows.

You need to focus on the knowledge that is most important to the success of your organisation. Start by looking at your organisation's goals and objectives. What do you need to do better, or continue to do well, in order to achieve those objectives? How does knowledge support you in doing that?

Examples of key knowledge might be:

- knowledge about, or a relationship with, a particular type of patient or a supplier
- key operational processes
- a key system, technology or piece of equipment
- a specific illness, disease or treatment
- the organisational culture, the internal infrastructure, "how to get things done around here"

Understand your target audience

It is important to understand who will be using the knowledge that you are capturing before you start to capture it. This will help you ensure you capture the right knowledge at the right level, and make it available in the most appropriate ways. Consider who will be your target audience, how many of them there are, where they are located, what their needs are – what do they need to know about this specific subject; what is their current level of knowledge and experience of it; how will they apply the knowledge; what access to they have to various media such as an intranet; etc.?

Find your experts

Identify the experts – the people who have the knowledge and know-how you are seeking to capture. If you have a staff directory that includes details people's skills and knowledge then this is a good place to start. Otherwise you might look at key documents on a subject and see who authored them, or ask managers and staff working in the area. Bear in mind that experts are not necessarily the most senior people in the organisation. Once you have found your experts, you can then collate some relevant background information about them including job descriptions, roles and responsibilities, education and training, work experience etc.

Choose your harvesters

An effective harvester (interviewer) is crucial. Much of the success of knowledge harvesting relies on the ability of the interviewer to elicit the right knowledge from experts. Making tacit knowledge explicit can be difficult – people often don't "know what they know" and so helping people to talk about what they know, and then capturing that effectively, is a key skill. It is generally recommended that you use a trained harvester – whether you hire an external consultant, or develop and train someone in-house. In the latter case, consider people with strong communication, interpersonal and interviewing skills, such as recruiters, researchers, trainers, counsellors or nurses.

Harvest: interview your experts

The best way to capture tacit knowledge is using one-to-one, face-to-face interviews with your experts. The interviews will involve asking them to talk about what they do and to describe specific situations in which they have applied specific know-how. Interviews need to be well prepared in advance, including drafting a topic guide or a list of questions. Examples of questions might include:

- Describe a time when…?
- What's the first thing you do?
- How do you know to do that?
- How do you know when to do it?
- What do you do next? Why?
- What usually happens?
- What happens if something else is done?
- What would happen if…?
- Who else is involved?
- What are some common mistakes or misconceptions?
- What is the most important thing to remember when you're doing this?
- Describe how you currently help others learn how to do this?

- What are the main obstacles that prevent them from achieving the same results as you?
- What are examples of support materials, documents, procedures, manuals, research evidence, check-lists that are relevant?
- What would make this process easier to understand?
- What would make this process easier to achieve?
- Etc.

In order to effectively capture the responses, you will need either a tape recorder or a second person to transcribe the interview. Some practitioners recommend a process in which the harvester conducts initial interviews with experts, and then presents the results to a group representing the eventual users of that knowledge. Any gaps in what the users need to know, or in their understanding of what has been captured, can then be used to form the basis of a second round of expert interviewing. This process of cycling between experts and eventual users can be invaluable in ensuring a fit between what is needed and what is being captured.

Organise, package and share

Once the knowledge has been gathered it can then be edited, organised and presented (or "packaged") into a form that meets the needs of its users. This may be a check-list, a manual or a set of guidelines etc. that can then be made available either in hard copy or (ideally, assuming your users have easy access to a computer) in a knowledge database or on the organisation's intranet. In some cases, the information is loaded into interactive software to provide an online tool to help users through relevant decision-making processes. For example, such a system might provide a variety of multiple-choice questions that guide the user to define a problem and apply the relevant criteria to solve it.

Apply, evaluate and adapt

It is important to ensure that the knowledge you have captured is being accessed and applied and that users are getting value from it. You will also need to consider its value over time: knowledge harvesting can result in relatively static documents that will, at some point, become out-of-date and so they will need to be continually refreshed if they are to retain their value.

4.10.4 Are there any other points I should be aware of?

- Before embarking on a knowledge harvesting programme, you need to consider whether your organisation's culture is one that encourages knowledge sharing. Successful knowledge gathering and sharing is unlikely to happen if people feel they would be at a disadvantage by sharing their knowledge. For example, experts may feel that their status or job security depends on keeping their knowledge to themselves. For more information about organisational culture, see People.
- Before you begin, be sure that you are clear on how you intend to package and make available the knowledge you have harvested and that you have the resources to do so. Otherwise you could end up with a stock of potentially useful knowledge that is going to waste.
- Not all tacit knowledge can be made explicit. There will always be aspects of know-how and experience that remain tacit. For those aspects, you will need to apply other knowledge management tools. The challenge is therefore to determine how much of the tacit knowledge in your organisation can be harvested and made explicit, and how much is best approached in another way.
- Some knowledge management practitioners feel that it is a mistake to focus on capturing and documenting tacit knowledge. Their view is that there is greater value in connecting people with each other so that they can share their tacit knowledge through "live" discussion and collaboration, and so they favour knowledge management tools such as communities of practice, storytelling, white pages and expertise directories, etc. In practice, it is wise to look at a combination of approaches, and adapt them to the specific needs and circumstances of your organisation.

4.10.5 More information

KnowledgeHarvesting.org
The website of a US-based knowledge management consulting firm specialising in knowledge harvesting. The site has a useful range of training documents concerning knowledge harvesting.

4.11 Peer assists

4.11.1 What are peer assists?

A peer assist is simply a process where a team of people who are working on a project or activity call a meeting or workshop to seek knowledge and insights from people in other teams. While seeking help from peers is certainly not new, the formal use of this process as a knowledge management tool and the coining of the term "peer assist" were pioneered by British Petroleum (BP).

4.11.2 What are the benefits?

> Peer assists are part of a process of what BP calls "learning before doing", in other words gathering knowledge before embarking on a project or piece of work, or when facing a specific problem or challenge within a piece of work. The benefits of peer assists are therefore quickly realised: learning is directly focused on a specific task or problem, and so it can be applied immediately.

> A peer assist allows the team involved to gain input and insights from people outside the team, and to identify possible new lines of enquiry or approach – in short, reusing existing knowledge and experience rather than having to reinvent the wheel. Peer assists also have wider benefits: they promote sharing of learning between teams, and develop strong networks among people.

> Peer assists are relatively simple and inexpensive to do: they do not require any special resources or any new, unfamiliar processes.

> It is worth using a peer assist when a team is facing a challenge, where the knowledge and experience of others will really help, and when the potential benefits outweigh the costs of travel.

4.11.3 How do I go about it?

There is no single right way to hold a peer assist. The following is a method that has worked well for BP.

1. Clarify your purpose
   Peer assists work well when the purpose is clear and you communicate that purpose to participants. Define the specific problem you are seeking help with, and be sure that your aim in calling a peer assist is to learn something (rather than seeking endorsement for a decision you have already made).

2. Has the problem already been solved?
   Do some research to find out who else has already solved or tackled a similar problem. Also, share your peer assist plans with others, as there may be other teams who are currently tackling a similar problem who could also benefit from participating in the peer assist.

3. Get a facilitator
   You will need a facilitator from outside the team, to make sure the meeting participants reach their desired outcome. The facilitator also may or may not record the event: be sure to agree on that before the meeting.

4. Timing is important
   Ensure that you plan a date for the peer assist that is early enough in your project to make use of the input you receive and to do something different on the basis of what you have learned. A frequent mistake is to hold the meeting too close to the decision date to make a real impact. Consider that you might get a different response to the one you expect: will you have time to do anything about it?

   The length of a peer assist depends on the complexity of the problem and tends to be somewhere between half a day and two days long.

5 **Select the participants**
Once you are clear on your purpose, select participants who have the diversity of knowledge, skills and experiences needed for the peer assist. Six to eight people is a good number. Look "across" the organisation rather than "up" it – hierarchies can hamper the free exchange of knowledge whereas peers tend to be much more open with each other and can challenge without feeling threatened. Avoid the temptation to select "the usual suspects": if the same experts are selected for peer assists again and again, you may be limiting the number of fresh ideas and perspectives available to you. Similarly, seek to select people who will challenge your ways of thinking and working and perhaps offer a different angle, rather than looking for people who will validate your current approach. You might consider inviting people from outside your organisation.

6 **Get clear about the deliverables**
Get clear on what you hope to achieve during the peer assist and then plan the time to achieve that. The deliverables should comprise options and insights rather than providing an answer. It is up to the person or team who called the peer assist to then make the relevant decisions, based on what is learned. Provide the participants with any briefing materials in advance so that they have adequate time to prepare.

7 **Allow time for socialising**
Allow time in your agenda for the teams to get to know one another; this might be a dinner the night before or time for coffee at the start of the day. It is important to build rapport so that the group can work openly together.

8 **Define the purpose and set the ground rules**
At the start of the meeting, ensure that everyone is clear about the purpose of the peer assist and their roles within it. The role of the host team is to listen in order to understand and learn. The role of the visiting team is to share knowledge and experience to help resolve the challenge without adding to the workload. Agree that where there are areas of contention, you will focus on the activity rather than the individual people involved.

9 **Start by sharing information and context**
Divide the meeting time roughly into four equal parts. During the first quarter, the host team will present the context, history and their future plans regarding the problem or challenge in question. Keep this part short and sharp – you only want to say enough to get the visiting team started in the right direction. Remember that the purpose of the peer assist is to learn rather than tell.

When communicating the problem or challenge about which you are seeking input, be prepared for it to be redefined as part of the peer assist process. It may be that the problem you have identified is in fact the symptom of a further problem and the peer assist will help you identify the root cause.

10 **Encourage the visitors to ask questions and give feedback**
In the second quarter, the visitors consider what they have heard, and then begin by discussing what they have heard that has surprised them, and what they expected to hear but haven't. The host team should take a back seat at this stage and simply listen; in some cases they may even opt to leave the room. The visitors then consider what else they need to know to address the problem and where might they find that knowledge. It may be that they want to make some telephone calls and talk to some other people, or request some data or reports. Remember, they are not seeking to solve the problem but to offer some options and insights based on their own knowledge and experience.

11 **Analyse what you have heard**
The third quarter of the meeting is for the visiting team to then analyse and reflect on what they have learned and to examine options. Again, the home team remains largely in the back seat; it might be appropriate to involve one or two of them, provided that they continue to listen and learn rather than closing off options or seeking to draw conclusions too early.

12 **Present the feedback and agree actions**
In the fourth and final quarter of the meeting, the visiting team presents their feedback to the host team and answers any questions. The presentation will be along the lines of "what we have learned, what options we see, and what has worked elsewhere". As with all feedback, this should start with the positive – what has been done well, and then what options there are to do things differently. When presenting what has worked elsewhere, presenters should simply tell the story rather than prescribing "you should…"

In closing, the person who called the peer assist should acknowledge the contribution of the visiting team, and also commit to when he or she will get back with an action list of what the team are going to do differently.

Finally, invite the visiting team to reflect on what they have learned and what they will take away and apply. Learning is never one-way.

4.11.4 Are there any other points I should be aware of?

- In the context of the NHS, an important consideration is that of evidence-based practice. When conducting peer assists, you will need to ensure that lessons learned are based on a combination of both on-the-job experience and sound research evidence.
- You might wish to carry out an After Action Review following your peer assist to look at whether the process went according to plan, what was different and why, and what can you learn from that for the next time.
- While the peer assist process is designed to provide input for a specific purpose or project, consider who else might benefit from the lessons learned. Always look out for opportunities to share and re-use knowledge and learning.

4.11.5 Resources and references

As the peer assist process was pioneered by BP, the above information was taken exclusively from the following two key sources. These sources easily provide enough information to get started.

Collison, Chris and Parcell Geoff. (2001) **Learning to fly: practical lessons from one of the world's leading** knowledge companies. Oxford: Capstone. Chapter 6: Learning from your peers. (This book as a whole is well worth a read – refreshingly low on theory and jargon, and high on sound, practical advice based on proven results).

Collison, Chris Collison Parcell, Geoff. **Learning before doing: BP's peer assist process**. Knowledge Management Magazine, 2001, Volume 4, Issue 10 An article that draws on the information in the chapter 6 of the above book.

4.12 Social Network Analysis

4.12.1 What is social network analysis?

"Social network analysis is the mapping and measuring of relationships and flows between people, groups, organisations, computers or other information/knowledge processing entities." (Valdis Krebs, 2002).

In the context of knowledge management, social network analysis (SNA) enables relationships between people to be mapped in order to identity knowledge flows: who do people seek information and knowledge from? Who do they share their information and knowledge with? In contrast to an organisation chart which shows formal relationships – who works where and who reports to whom, a social network analysis chart shows informal relationships – who knows who and who shares information and knowledge with who. It therefore allows managers to visualise and understand the many relationships that can either facilitate or impede knowledge creation and sharing. Because these relationships are normally invisible, SNA is sometimes referred to as an "organisational x-ray" – showing the real networks that operate underneath the surface organisational structure.

4.12.2 What are the benefits?

Once social relationships and knowledge flows can be seen, they can be evaluated and measured. The results of social network analyses can be used at the level of individuals, departments or organisations to:

> identify teams and individuals playing central roles – thought leaders, key knowledge brokers, experts, etc.;
> identify isolated teams or individuals;
> detect information bottlenecks;
> spot opportunities for knowledge flow improvements;
> accelerate the flow of knowledge and information across functional and organisational boundaries;
> improve the effectiveness of formal communication channels;
> target opportunities where increased knowledge flow will have the most impact;
> raise awareness of the importance of informal networks.

4.12.3 How do I go about it?

The process of social network analysis typically involves the use of questionnaires and/or interviews to gather information about the relationships between a defined group or network of people. The responses gathered are then mapped using a software tool specifically designed for the purpose (see Resources and references below for examples). This data gathering and analysis process provides a baseline against which you can then plan and prioritise the appropriate changes and interventions to improve the social connections and knowledge flows within the group or network.

Key stages of the process will typically include:

> Identifying the network of people to be analysed (e.g. team, work group, department).
> Gathering background information – interviewing managers and key staff to understand the specific needs and problems.
> Clarifying objectives, defining the scope of the analysis and agreeing the level of reporting required.
> Formulating hypotheses and questions.
> Developing the survey methodology and designing the questionnaire.
> Surveying the individuals in the network to identify the relationships and knowledge flows between them.
> Use a software mapping tool to visually map out the network.

- Reviewing the map and the problems and opportunities highlighted using interviews and/or workshops.
- Designing and implementing actions to bring about desired changes.
- Mapping the network again after a suitable period of time.

4.12.4 Are there any other points I should be aware of?

In order for SNA maps to be meaningful, it is important to know what information you need to gather in order to build a relevant picture of your group or network. Good survey design and questionnaire design are therefore key considerations.

Questions will be typically based on factors such as:

- Who knows who and how well?
- How well do people know each other's knowledge and skills?
- Who or what gives people information about xyz?
- What resources do people use to find information/feedback/ideas/advice about xyz?
- What resources do people use to share information about xyz?

4.12.5 More information

Hanneman, Robert. Introduction to social network methods.
An online book available here

International Network for Social Network Analysis

4.13 Storytelling

4.13.1 What is storytelling?

"Back in the mists of time when only the monks and the monarchy could write, there three ways in which we learnt: first by having a go at it. When that didn't work, or you wanted to improve – by watching someone who knew how to do it. Then at the end of the day when the sun had set and it got too dark to see what anyone was doing – by listening to that "someone" tell you about the time when they..." (Weaver – Grazing Animals Project, 2003)

Storytelling is quite simply the use of stories in organisations as a communication tool to share knowledge. Traditionally, organisational communications have had a tendency to be somewhat dry and lacking in inspiration. Storytelling uses a range of techniques to engage, involve and inspire people, using language that is more authentic (everyday language as opposed to "textbook buzzword speak") and a narrative form that people find interesting and fun.

Storytelling has of course existed for thousands of years as a means of exchanging information and generating understanding. Similarly, it has always existed in organisations – otherwise known as "the grapevine". However, as a deliberate tool for sharing knowledge it is quite recent but growing very rapidly, to the extent that it is becoming a favoured technique among an increasing number of management consultants.

4.13.2 What are the benefits?

When used effectively, storytelling offers numerous advantages over more traditional organisational communication techniques:

> Stories communicate ideas holistically, conveying a rich yet clear message, and so they are an excellent way of communicating complicated ideas and concepts in an easy-to-understand form. Stories therefore allow people to convey tacit knowledge that might otherwise be difficult to articulate; in addition, because stories are told with feeling, they can allow people to communicate more than they realise they know.

> Storytelling provides the context in which knowledge arises as well as the knowledge itself, and hence can increase the likelihood of accurate and meaningful knowledge transfer.

> Stories are an excellent vehicle for learning, as true learning requires interest, which abstract principles and impersonal procedures rarely provide.

> Stories are memorable – their messages tend to "stick" and they get passed on.

> Stories can provide a "living, breathing" example of how to do something and why it works rather than telling people what to do, hence people are more open to their lessons.

> Stories therefore often lead to direct action – they can help to close the "knowing-doing gap" (the difference between knowing how to do something and actually doing it).

> Storytelling can help to make organisational communication more "human" – not only do they use natural day-to-day language, but they also elicit an emotional response as well as thoughts and actions.

> Stories can nurture a sense of community and help to build relationships.

> People enjoy sharing stories – stories enliven and entertain.

4.13.3 What can stories be used for?

Stories can be used for all manner of purposes in an organisation. Different purposes will tend to require different kinds of stories. Steve Denning (http://www.SteveDenning.com) outlines 8 possible purposes for using storytelling in organisations:

1. **Storytelling to ignite organisational change**
   Experience has shown that storytelling can be highly effective as a change agent, even in change-resistant organisations. Telling an appropriate story can stimulate people to think actively about the implications of change and to projecting themselves into visions of the future, enabling them to better understand what it will be like to be doing things in a different way, rather than being given vague, abstract concepts about it.

2. **Storytelling for communications**
   In contrast to the conventional approach which views communications as the sending of a message from a communicator to a recipient, storytelling is based on a more interactive view of communication. Because the listener imaginatively recreates the story in his or her own mind, the story is not perceived as coming from outside, but rather as something that is part of the listener's own identity. The idea becomes the listener's own.

3. **Storytelling to capture tacit knowledge**
   Tacit knowledge can be a multi-layered and multi-dimensional thing and as such it is often difficult to articulate (for example, have you ever tried to explain to someone who can't swim how to swim, without actually showing them?). Stories can provide a way of allowing people to express and share tacit knowledge in rich and meaningful ways, rather then being forced to articulate it in more "structured" ways that can detract from its value.

4. **Storytelling to embody and transfer knowledge**
   Similarly, a simple story can communicate a complex multi-dimensioned idea, not simply by transmitting information as a message, but by actively involving the listeners in co-creating that idea. Furthermore, as a story is told and retold, it changes, and so the knowledge embodied in it is constantly being developed and built upon.

5. **Use of stories for innovation**
   The use of storytelling in innovation and knowledge creation can encourage people to move away from linear thinking towards a more multi-dimensional view, to see new connections between things, and also to marry scientific logic with a more creative or intuitive approach.

6. **Storytelling to build community**
   There is something about stories that brings people together and fosters a sense of community. Storytelling is non-hierarchical, it unlocks feelings and emotions as well as thought processes, and hence it helps to build relationships and trust.

7. **Storytelling to enhance technology**
   People often find it difficult to communicate about technology. Users can have trouble articulating their needs and expectations, while experts can have difficulty "talking in plain English". Wherever there is a gap in language and understanding, storytelling can provide a bridge, by communicating the real essence of what each party is trying to get across.

8. **Storytelling for individual growth**
   Storytelling is a skill, and one that draws on a number of other key skills, mostly relating to interpersonal communication. The development of these skills is an important component of most knowledge management programmes.

4.13.4 What makes a "good" story?

Larry Prusak (see Storytelling: passport to the 21st century http://www.creatingthe21stcentury.org/) defines 4 attributes of a good story:

1. **Endurance**
   Good stories endure. They may change a little – or even a lot, but the key lessons remain the same. They also need to be succinct enough for people to remember.
2. **Salience**
   Good stories are relevant to their audience, they have a point, and they have emotional impact.
3. **Sense-making**
   Good stories explain something, make sense of something. Perhaps they show you how to behave in particular situation, how to resolve a problem, or why something happened the way it did. They have a prescriptive normative value: do x and y will occur.
4. **Comfort level**
   To be effective, stories must make sense within the context of the listener's experience – they need to ring true.

Other tips (from Steve Denning – http://www.SteveDenning.com) include:

> **Fact versus fiction**
Storytelling can be counter-productive when the story told is not true. A story can be factually accurate while being authentically untrue and many corporate communications take this form, particularly those that are told more as a public relations exercise than as a means to promote genuine learning.

> **Oral versus written stories**
In the written word there is a distance between the speaker and the spoken, and so in an organisational context, it can lack some authenticity. Practitioners have found that oral storytelling has a greater impact than putting stories into booklets or videos or online. This doesn't mean that written stories can't achieve good effects, but that they work in different kinds of ways.

> **The "happy ending"**
Steve Denning reports having had no success in telling a story along the lines of: "Let me tell you about an organisation that didn't implement knowledge management and it went bankrupt." In other words, focus on the positive.

> **The "hero"**
A story needs to be told from the perspective of a single protagonist, someone who everyone in the organisation can instantly understand, empathise with, resonate with their dilemma, and understand what they were going through.

> **The "plot"**
A story needs to have a certain strangeness or incongruity – something that is remarkable and therefore grabs attention. ("That's remarkable that you could get an answer to a question like that in such a short time frame"). But it is nevertheless plausible (email exists, the web exists).

> **A beginning, middle, and an end**
A story needs to embody whatever it is you are seeking to get across as fully as possible. Don't leave loose ends.

> **Timing**
A story should be as recent as possible – older stories can work, but the fresher the better. "This happened last week" conveys a sense of urgency.

4.13.5 Are there any other points I should be aware of?

- Storytelling is not a panacea – it doesn't always work. Storytelling can only be as good as the underlying idea being conveyed. If that idea is unsound, storytelling may well reveal its inadequacy.
- Even when the underlying idea is good, there are times when storytelling is inappropriate or ineffective. For example: routine situations in which nothing new, unexpected or different happened; or situations that require objectivity in reporting.
- Storytelling does not replace analytical thinking. It supplements it by helping to give it context and meaning. Abstract analysis is often easier to understand when seen through the lens of a well-chosen story.
- Try to avoid telling a story for the first time at a high profile, high-risk occasion. Test the story in advance on a variety of similar audiences, so that you know exactly the effect that the story will have.
- When using the knowledge contained in the stories of others to support your own decisions, consider how you will balance that anecdotal knowledge with evidence-based knowledge: how will you assess and integrate the knowledge from stories?
- We are all storytellers and spend much of our lives telling stories whether we realise it or not. However we can all get better at storytelling, particularly at using stories to achieve specific effects. Understanding how and why storytelling works and learning what kinds of stories work in different situations, and what kinds of effects different kinds of stories have, can enable us to be more adept storytellers in an organisational context.

4.13.6 More information

SteveDenning.com
http://www.stevedenning.com
Steve Denning is widely regarded as one of the main "gurus" of storytelling. His website has a collection of materials on knowledge sharing and storytelling, and also includes a facility where you can also e-mail Steve direct with questions and comments.

Storytelling: passport to the 21st century
http://www.creatingthe21stcentury.org/
A website in which four leading thinkers on knowledge management explore storytelling. An extremely content-rich site, almost like an online book.

4.14 White Pages

4.14.1 What are "white pages"?

An organisational "white pages" is a tool to help people to find others in their organisation that have the knowledge and expertise they need for a particular task or project. It is like a staff directory, but rather than simply listing people's names, job titles, departments and contact details, it includes details about their knowledge, skills, experience and interests.

"White pages" are electronic rather than paper-based, so that users can search it in a variety of ways, just like they might perform a search on the Internet.

"White pages" are often also known as experts' directories, expertise directories, skills directories or capabilities catalogues.

4.14.2 What are the benefits?

A "white pages" directory is particularly beneficial in organisations that are over a certain size or that are spread around in different locations, and so people don't have the opportunity to get to know each other well. Specific benefits include:

> "White pages" are technologically quite simple to create

> They can be extremely effective in helping organisations to "know what they know"

> They allow people to find the tacit knowledge they need, by easily finding the people who have it

> They can underpin all of the organisations various initiatives to connect people with people, and to learn from others

> A "white pages" is not necessarily aimed at those embarking on a major project or piece of work; often the greatest value comes from a multitude of simple ten-minute conversations in which people ask each other for a quick word of advice or a steer in the right direction.

By way of an example, can you find an asthma expert who has considerable experience in a specific treatment, has successfully used that treatment with children under five, and is currently in or around the Birmingham area, all in under a minute? A good "white pages" could enable you to do that (assuming of course that such a person exists!).

4.14.3 How do I go about it?

Be clear about your aims

First, be clear about your aims. Using a "white pages" to find people is a means to an end, not an end in itself. How do you intend for people to use it? For what purposes do you envisage them using the system to find people? How will they approach and use the system? It is vital to be clear on this before you begin designing any system. Talk to people in your organisation to find out about their needs and views. Talk to people in other organisations who have already implemented a "white pages" to find out what you can learn from their experiences.

Ownership and onus

Opinions vary about whether to make individuals' inclusion a "white pages" compulsory or voluntary, and similarly whether to create and manage entries centrally or provide a template for individuals to create and update their own entries. Organisations such as BP-Amoco and Texaco who have implemented successful "white pages" strongly favour the voluntary approach in which individuals create their own entries if they so choose. Their experience would seem to show that ownership needs to be with the people contributing to, and using, the system.

This has a number of advantages. First, it creates a sense of personal responsibility for the system which in turn fosters support; second, it allows people to present their entries in a way that reflects how they want to be known rather than how the organisation sees them; and hence third, it helps to create a "living" system that reflects real personalities and therefore encourages personal relationships.

Balance formal with informal information

While the purpose of a "white pages" is to help people find others with relevant knowledge and expertise, the chances of them actually acting on that information and calling that person will be greatly increased if they feel they "know" them. This sense of "knowing" or familiarity can be created to some extent by including some personal information and a photograph in people's entries. Allow people to be creative in how they present themselves. For example, at BP people are encouraged up upload photographs of themselves at home or at play – perhaps with their children or enjoying their favourite sport – rather than using a more sterile passport-style photograph.

What to include

Common fields found in a "white pages" include:

- Name
- Job title
- Department or team
- A brief job description and/or description of what is currently being worked on and what has been worked on in the past
- Relevant professional qualifications
- An uploaded CV
- Areas of knowledge and expertise (selected from a pre-defined list of subjects/terms; people might also rank their knowledge, for example from "extensive" to "working knowledge" to "basic")
- Main areas of interest
- Key contacts – both internal and external
- Membership of communities of practice or other knowledge networks
- Personal profile
- Photograph
- Contact information

Organising entries for ease of loading and retrieval

In order to encourage people to create entries, you will need to make it easy for them. Most organisations use a simple template into which users enter their information. In creating a template, think not only about ease of entry, but also about how users will search the system to retrieve information. You will need a common language or taxonomy to describe information in the essential fields, in particular those relating to knowledge, expertise, areas of work and interests. You might like to create fixed terms and options for these fields that users can select from a menu or a selection of tick-boxes. This could also be supplemented with a box for users to enter free text, perhaps with some suggested terms alongside it to guide their use of language.

In contrast, personal information can of course be relatively unstructured – leave scope for more creativity and free expression here!

Keeping it current

A "white pages" must be maintained and kept up-to-date. People are constantly moving locations, changing jobs, and adding to their knowledge and skills. If your "white pages" is linked with your human resources system, then job details and contact information can be automatically updated. Alternatively if individuals have sole responsibility for their own entries, then you might build a reminder process into your system, whereby an e-mail is sent automatically to remind users who haven't updated their entries since a certain time period, such as three to six months. Similarly, be sure to build information about the "white pages" into processes for new joiners and leavers, so that new joiners know about the system and are encouraged add their entry, and leavers remember to either delete their entry or delegate it to someone else to "own" (assuming they are happy for people to still contact them after they have left).

Encouraging use

You will need to actively market your "white pages". Don't assume that if you create it, people will automatically use it. Your marketing efforts will need to encourage both participation and use; the two are inextricably linked as you need a certain amount of submissions for people to see the "white pages" as being worth using. Possible ideas might include posters, presence at events such as learning fairs, nominating champions to promote the "white pages" in various parts of the organisation, or competitions that give prizes to the first departments in which everyone is uploaded, or for those with the best success stories of how using the "white pages" has helped them in their job. Be sure to focus on the benefits in your marketing efforts – people will want to know "what's in it for me?".

4.14.4 Are there any other points I should be aware of?

> A "white pages" need not just include individuals – for example you might like to include formal communities of practice, project teams, etc.

> Similarly, a "white pages" need not just cover internal people; you can also have a similar system, or a section, for suppliers of various types (e.g. IT outsourcing, consultancy services, recruitment agencies, etc.), and for other organisations with which you work or collaborate, both within and outside the NHS.

> You can add further value to your "white pages" by linking it with other knowledge management tools, such as those available on an intranet. For example you might have collaborative working tools or best practice databases that list relevant contacts – these contact listings can be linked directly into the "white pages" – and vice versa.

> Be careful when using the term "expert" – it can be quite a "political" one and may create hierarchies; if some people are considered as experts, this might make others feel that their knowledge is less valuable so it may discourage their contribution.

> Be aware of issues relating to data protection – check with your legal department to ensure that your "white pages" will comply with relevant requirements, and to create a policy on its correct use.

5 DEVELOPING THE KM ENVIRONMENT

Whatever knowledge management tools and techniques you use, they are unlikely to work in isolation – they need to be supported by the right kind of environment. The three key elements of that environment are outlined here, namely:

- People
- Processes
- Technology

Your organisation's people, processes and technology will at all times be acting as either enablers of, or barriers to, effective knowledge management. You will need to identify the barriers and remove them. You will probably also need to build on existing enablers and create additional ones. This is often where the greatest knowledge management challenges lie.

5.1 People

5.1.1 Introduction

Of the three components of knowledge management – people, processes and technology – the most important is undoubtedly people. Why? Because creating, sharing and using knowledge is something that is done by people. Processes and technology can help to enable and facilitate knowledge management, but at the end of the day it is people who either do it or don't do it. A number of organisations have learned this through bitter experience. Of those companies that led the way in the early days of knowledge management, many focused primarily on processes and technology – to their cost. Having made significant investments in the latest systems, they then found that people simply did not use them and so the systems ended up being confined to what became known as "the ". Since then, organisations have learned that it is people who "make or break" knowledge management initiatives.

5.1.2 Why people don't want to share knowledge – or do they?

There is a traditional view that knowledge sharing is not a natural act and that people need to be coerced or cajoled into it. In fact why not take a few moments right now to think about some of the values, attitudes and behaviours in your organisation that constitute barriers to seeking, sharing and using knowledge? For example:

- "Knowledge is power"
- "I don't have time"
- "I've got too much real work to do"
- "That's not my job"
- "You're just using other people's ideas and taking the credit"
- "I want to do things my way"
- "This is how it's always been done"
- "I'd like to help, but my manager won't like it if I waste time doing things for another team"
- "That's not how we do things around here"
- "I don't trust them"
- "Are you telling me how to do my job?"
- "I'm already suffering from information overload"
- "We're not allowed to make mistakes, let alone admit to them, share them or learn from them"
- "Don't bother others by asking them for help, work it out for yourself"
- "You should already know all the answers"
- "It's just another management fad; if I ignore it, it'll eventually go away"

- "What's in it for me?"
- "No"

These are just a few! However it may surprise you to learn that there is also a view that knowledge sharing is in fact a very natural act and that we are already doing it all the time. If you take a few moments to watch people both at work and at play, you can see the evidence daily: in corridors, by the coffee machine, on the phone, by e-mail, at the pub, etc. – people are freely sharing knowledge all the time. Similarly, knowledge management consultants have reported that in their experience of working with a range of organisations, people want to share. They want to make a valuable contribution to their organisations, they like to see their knowledge being used, they want to help their colleagues, and they want to learn from others who they trust and respect.

So why does the "people" aspect of knowledge management tend to be such a challenge for most organisations?

Because our organisational cultures get in the way, they give rise to, and reinforce, behaviours that inhibit knowledge sharing. Most of us in the Western world have been trained to believe in individual effort and competition, and this from an earlier age than you might realise – remember at school how knowledge sharing was called cheating? Since then, our working environments have largely perpetuated this way of thinking. We compete for jobs, salaries, promotions, recognition, status, power, budgets and resources, always believing that if someone else has something then there's less of it left for us. Put simply, we have been trained not to share.

Awareness of this is the first step to overcoming it. It is important to understand that we all carry this kind programming and we all need to take responsibility for unlearning it and rethinking our old philosophies. Contrary to popular belief, experience is increasingly showing that people are generally willing to share, but they need a supportive, encouraging and safe environment in which to do so. Sadly, most organisational cultures have some way to go before they can claim to provide such an environment.

5.1.3 The two big makers or breakers: culture and behaviour

Essentially there are two key aspects of "people" that you will need to address when introducing knowledge management into an organisation: organisational culture and individual behaviour. The two are inextricably linked.

5.1.4 Organisational culture

Effective knowledge management requires a "knowledge sharing" culture to be successful. What exactly is organisational culture? The short answer is that culture is "the way we do things around here". A more complete answer is that an organisational culture is a set of values, beliefs, assumptions and attitudes that are deeply held by the people in an organisation. They influence the decisions people make and they ways in which they behave. In organisations that recognise only individual achievement, people are rewarded for their personal knowledge and have no incentive to share it. In a knowledge sharing culture, people can be rewarded for individual achievements, but are also recognised and rewarded for their knowledge sharing and contributions to team efforts. Key characteristics of a knowledge sharing culture include the following:

- top leadership sees knowledge as a strategic asset and provides incentives and support for knowledge management processes;
- the organisation focuses on the development and exploitation of its knowledge assets;
- tools and processes for managing knowledge are clearly defined;
- knowledge creation, sharing and use are a natural and recognised part of the organisation's processes, not separate from normal work processes;
- groups within the organisation cooperate instead of compete with each other;
- knowledge is made accessible to everyone who can contribute to it or use it;
- rewards and performance evaluations specifically recognise contributions to, and use of, the organisation's knowledge base; communication channels and a common technology infrastructure enable and enhance knowledge management activities.

Organisational cultures run deep: the older and the bigger the organisation, they deeper they will tend to run. Which brings us to the question: to what extent can we change organisational culture? There is some debate about this, but the common view is that culture can be changed, but usually not without a great deal of time and effort. Think about the last time you tried to change somebody's mind about just one thing: multiply that by the number of people in your organisation then add to it the cohesive power of shared beliefs, and you begin to get an idea of the task at hand. In other words, taking on the entire organisational culture at on go is simply not feasible. The good news is that there is another approach: individual behaviour.

5.1.5 Individual behaviour

If knowledge management is new to an organisation, it requires changes in individual behaviour. Individuals must be encouraged to incorporate knowledge management activities into their daily routines. This includes activities relating to seeking out knowledge when they have questions or problems, finding and using existing knowledge rather than reinvesting the wheel, sharing their own knowledge, learning from others' experience and helping others to learn from theirs.

While people's behaviours are largely a function of the organisational culture, they are easier to see and to identify as "makers or breakers" – enablers or barriers – to knowledge sharing. This is best approached from the context of your current objectives, issues and the day-to-day work of your employees. By changing the way people behave and by showing them new ways of working that can make their jobs easier and more successful, you can not only change their behaviour, but also affect the underlying cultural assumptions that drive people's behaviour in the first place. In other words, people learn best by doing, rather than being told.

Of course for individual behaviours to change in a sustained way, there needs to be a conducive organisational culture, which brings us back to the earlier point that the two are inextricably linked.

5.1.6 How do we make the changes?

Assuming that people will generally share knowledge if the barriers and disincentives to doing so are removed, then you can seek to bring about lasting changes in both individual behaviours and organisational culture by:

- focusing on changing individual behaviours first
- understanding the barriers to knowledge sharing and seeking to eliminate them
- introducing policies and practices that enable and encourage knowledge sharing
- understanding your organisational culture and working within it rather than against it while gradually working to change it

Here are some approaches and issues to consider:

5.1.7 Culture – work with it while you work towards changing it

If the people in your organisation hold a fundamental belief that asking for help is a sign of weakness, then immediately launching a peer assist programme might not be the best way forward. If people prefer to seek information from other people, then loading endless documents into knowledge databases is unlikely to work. And if people feel that they are not allowed to make mistakes and that to admit to mistakes might be dangerous, then you may need to wait until this has started to shift before introducing after action reviews. In other words, if you pit yourself against the organisational culture, you are fairly likely to lose. Far better to work within it, at least initially, and then seek to change it from the inside.

For example, you might have something you feel is good practice that you want to share, but people in your organisation have a "not invented here" attitude and your good ideas have been ignored in the past. Instead of trying to sell your idea, ask for help to improve your practice. You may well find that not only do you receive plenty of input to help you improve it further, but others are suddenly more interested in finding out more with a view to applying it – because they have contributed to its development.

5.1.8 Lead by example

Actions speak louder than words. Nobody likes to be told to change their behaviour by someone who is clearly not exhibiting that behaviour themselves – and rightly so. Good leadership is key. Even if leaders are supporters of knowledge management, they still might need some coaching. Knowledge seeking and sharing behaviours may well be as new to your leaders as to everyone else. They need to be shown the way, and then be seen to be leading the way. For example, do leaders openly and actively share knowledge about what they are doing, where the organisation is going, what their plans for the future are, how things are financially? Do they gather knowledge from a range of people throughout the organisation as part of their decision-making processes? Do they seek honest input and feedback from both staff and patients? Do they listen, and where appropriate, do they act on it?

As well as leaders, middle managers are also very important in knowledge management. For most people in an organisation, the person who most affects their day-to-day work is their line manager or supervisor. These managers are often evaluated on how their individual section or department performs which means that the focus of their attention may not be on the bigger picture. Like senior managers, middle managers will most likely need some coaching to change their behaviours.

Finally, don't forget that you, as a knowledge management change agent, will need to lead by example too.

5.1.9 Align rewards and recognition

As with any change, whenever people are asked to do something differently, they need a good reason: what's in it for me? If people believe they will benefit from sharing knowledge, either directly or indirectly, they are more likely to share.

When looking at reward and recognition, remember that different people are motivated by different things; some by money, others by status, some by knowledge, others by freedom etc. A good reward system will recognise this.

In seeking to create and sustain a knowledge sharing culture, you will need to address your organisation's formal rewards that are embedded in your human resources policies and practices including salaries, bonuses, promotions etc. Most organisations' formal reward systems still reward individual effort and knowledge. To create a culture that supports knowledge creation, sharing and re-use, you will need to recognise and reward those behaviours. However some practitioners recommend focusing on informal rewards and recognition in the initial stages; they suggest that seeking to change formal reward structures very early on in the process could be damaging as you might lose the support of people who feel threatened. While the first behavioural changes are taking place, people need a safe space to learn and readjust without being assessed or penalised.

Informal rewards and incentives need not be financial, nor need they be complicated. A number of studies have shown that one of the most effective incentives is simple recognition. For example, you might decide to personalise knowledge: "John Broadbent's Guide to Winter Capacity Planning" or "Camden NHS Trust's Booking Process". This simple approach can increase the credibility of the knowledge, thereby increasing its likely use, and also make those who created and shared it feel valued and credited.

Be sure that you reward only valuable knowledge – knowledge that is actually used. Organisations who have offered incentives to staff to submit documents to a database or other knowledge system have often ended up with systems full of worthless information that nobody uses. Similarly, do not just focus on rewarding people who share their knowledge. This is only part of the knowledge equation. At the end of the day, you are seeking to encourage people to use and reuse knowledge, so reward the user for reusing and building on existing knowledge rather than wasting time and energy reinventing the wheel.

5.1.10 Make knowledge work part of everyone's job

Of all the reasons people have for not sharing knowledge, not having time is one of the most common. People are too busy with "real work". Knowledge work needs to be recognized as "real work" – an integral part of everyone's job. People need time to seek out knowledge, to reflect, to share what they know, to change the way they do things based on knowledge and learning received. They need to know that these activities are regarded not only as acceptable, but important, by the organisation. By making knowledge sharing a formal part of people's responsibilities, using it in job descriptions, and incorporating it into performance appraisal processes, you can clearly demonstrate the importance of knowledge work and begin to lay the foundations for a real knowledge culture.

5.1.11 Develop relationships

People share things better with people they know and trust. If people don't share personal relationships or mutual trust, they are unlikely to share knowledge of high value. Similarly, whether or not people seek out and use the knowledge of others depends if they know and trust the source of the knowledge. People also generally prefer to learn from their peers than from managers telling them what to do. And studies show that people will more often than not prefer to contact someone they know for information before searching a knowledge database.

While you cannot shift an organisational culture to one of openness and mutual trust overnight, you can make significant progress by helping and encouraging individuals and teams to form new and better relationships. As organisations get bigger, people get busier, and technology creates increasingly "virtual" ways of communicating with each other (across the internet, by e-mail, via video conferencing etc.), opportunities for developing relationships can seem few and far between – unless you make it a deliberate strategy.

While early practitioners of knowledge management focused on technology, the current view is that the greatest value can be realised by building relationships and connecting people with people, using tools such as communities of practice, peer assists, learning events, coaching and mentoring, and others.

5.1.12 Educate people about what is involved and skill them to do it

Given that most of us have not been educated or trained to share, and therefore, we simply don't know how to carry out one of the core activities of knowledge management. In many cases people simply don't realise what they know, or they don't realise the value of what they know. Even if they do, they may not know with whom to share or how to share what they know.

As with any other behavioural change, you need to show people clearly what is expected of them and what is involved, and then give them the skills to do it. You need to show people what creating, sharing and using knowledge looks like – both in general terms, and specifically within your organisation. You may also need to show them what exactly you mean by "knowledge". Knowledge can seem very conceptual, at least to begin with; it is not always obvious to people what it is they need to know, what they currently know, and how that might be useful to others.

In short, you need to train people in knowledge management skills. Educate them about what knowledge is valuable, how to create it, find it, evaluate it, share it, use it, adapt it, reuse it etc. Ensure that essential communication skills are also looked at. For example, knowledge sharing works better if people develop active listening skills. Active listening is where people spend time understanding what the other person really means, instead of focusing on what their own response will be and queuing up to speak. Another important skill in knowledge sharing is giving and receiving feedback – both positive and "negative". Challenging another person's beliefs or approaches in a way that causes neither offence nor defence is not always easy, nor is receiving such a challenge. Similarly, many people feel equally intimidated about both giving and receiving compliments and praise.

You will also need to ensure that people have enough information about the context in which they are working; for example, to effectively seek and share knowledge, people need some understanding of organisational strategies and goals, of the interrelationships between different functions and teams in the organisation, of what knowledge is most valuable to the organisation, and how it can be used to best effect.

Start early by building aspects of knowledge work into your organisation's induction programmes. (Note: if you don't have a formal induction programme, are there informal processes in place to ensure that new recruits get the knowledge they need to learn what they need to know?). Look at integrating aspects of knowledge work into other general training programmes: people learn and apply learning most effectively when knowledge work is seen as an integral part of their day-to-day job rather than a separate add-on.

5.1.13 Demonstrate the value

It is important that people understand the benefits of knowledge sharing on a number of levels: benefits to the organisation, benefits to patients, and benefits to them personally. The more you can clearly demonstrate these benefits, the more people are likely to be open to change. Be ready to answer the inevitable "why should I, what's in it for me?" question.

A number of studies have shown that by far the most effective incentive for producing lasting change is when the process of sharing knowledge is rewarded, supported, and celebrated, by the organisation. If this is not the case, then any artificial rewards and incentives will have little effect. In other words, knowledge management should provide intrinsic rewards to the people who use it. For example, does a particular knowledge system or process enable its users do their job more easily, more efficiently or more effectively? Does it help them provide a better service to their patients? Do people receive greater recognition from peers as key contributors and experts? Is their work faster, more accurate, more rewarding?

This is the bottom line: if knowledge management helps people to do their work, and the organisation's culture supports it, then people will most likely adopt it; if it doesn't, then they probably won't. And probably rightly so!

There are many ways to demonstrate and reinforce value, even in the early stages when knowledge management is new and the benefits have not yet been fully realised. Again, you can build a knowledge element into training programmes, using case studies and action-based learning to demonstrate the value of good knowledge sharing practices. Storytelling can also be a very effective tool here, as can creating knowledge "champions" or "heroes".

5.1.14 Create champions and heroes

A useful approach to showing people the benefits of knowledge sharing and to encourage them to change their behaviour is to create "knowledge champions" and/or subject "experts" dotted around the organisation. Every organisation has people who are naturally "knowledge savvy" – that is just their way of doing things: they love to learn and to share what they know with others. Similarly, every organisation has its "early adopters" – those who are first to change their behaviour and adopt new ways of working. Find these people and celebrate them as "heroes'; publicly recognise and reward their behaviour; encourage them to tell stories about what they did and what were the benefits. Create semi-formal roles that recognise this behaviour as a role model and allow these people to spend some time sharing their approach with, and inspiring, others.

5.1.15 Make it easy

Finally, remember that barriers to knowledge sharing are not just related to culture and behaviour. There are also barriers that relate to organisational structures and processes, and to technology. If you want people to change their ways of working, then you need to make it as easy and painless as possible for them. You need to identify, and as far as possible eliminate, these other barriers. Otherwise, even with the best will in the world, seeking and sharing knowledge may simply be more effort that it is worth for people. If technology is slow and unreliable, if different people use different systems and therefore cannot communicate or share documents easily, if structures promote hierarchies and internal competition rather than peer relationships and co-operation, if processes are highly task-oriented rather than people-oriented, then people will find knowledge sharing a challenge.

Similarly, when seeking to eliminate barriers and to introduce knowledge management tools, be sure to do so in a way that is integrated with people's day-to-day working practices. A common mistake in knowledge management is to introduce technology and processes and then sit back and wait for people use them – "if we build it, they will come". If systems and processes are created in a way, which is not integrated with how people actually work, then they will not be used. It is critical to include users in the development of knowledge tools so that this all too common and costly mistake can be avoided.

5.1.16 Cultural change is not just a knowledge management issue

A final word on seeking to change organisational cultures and individual behaviours: this is neither quick nor easy, but for effective knowledge management, it is not optional.

However, you should not expect knowledge management to carry the full weight of cultural change. Culture is critical to knowledge management, but it is equally critical in either enabling or disabling most other organisational processes. Cultural change is too big a task for knowledge management to take on alone. A better approach is to combine initiatives and present a common vision and focus that integrates knowledge management with overall organisational learning and performance improvement. This may well be all the more important in the content of the NHS, where people are already becoming "initiative-weary". That being said, there is currently a tremendous opportunity to align knowledge management with the context of the massive transformation currently under way in the form of the ten-year modernisation programme, set out in the NHS Plan. Delivery of the Plan will require change and transformation on a vast scale and on a number of levels throughout the NHS. Knowledge management is a natural partner to such transformation, as it will require major cultural change, new ways of working, and a strong focus on knowledge and learning. Strike while the iron is hot!

5.1.17 More information

Lelic Simon. "Your say": creating a knowledge sharing culture. Inside Knowledge, 2001, 4(5)

5.2 KM Processes

5.2.1 Introduction

Knowledge management processes are the activities or initiatives you put in place to enable and facilitate the creation, sharing and use of knowledge for the benefit of your organisation. Processes also refer to your organisation's general infrastructure and processes (or "ways of doing things"), and the extent to which these act as enablers of, or barriers to, good knowledge management practice. Hence the "process" component involves looking at:

> Organisational processes and infrastructure – and whether they currently help or hinder knowledge management;

> Knowledge management processes and infrastructure – the "process" element of what needs to be put in place to make knowledge management happen (as well as people/culture and technology).

5.2.2 Organisational processes and infrastructure

Every organisation has a structure and processes, and these operate on a number of levels. The buildings, in which you work, and their geographical location, provide a physical structure. The way the organisation is divided into departments and functions provides another form of structure. How people are organised into hierarchies and the relationships between them provides another. The way that resources are allocated – finances, technology, equipment, etc. – provides yet another. Each of these types and layers of structure will have an impact on how knowledge is created, shared and used in an organisation. For example:

> Does everyone in your organisation have ready access to a computer? Do they know how to use it?

> Is everyone located in the same building or are they dispersed across different buildings or even different towns or regions?

> Within each building, how is the space organised? Are people shut off from each other in offices with closed doors or is the space more open? Are managers located in the same areas as their teams, are they visible and accessible, or are they hidden away in a private area? Are there areas where people can simply "be" together – such as a café, or chairs and tables near a coffee machine, or informal "breakout rooms" – in addition to formal office and meeting space?

> What is the nature of the relationship between various departments and functions? Is it competitive or collaborative? How is this sustained, for example do departments have to compete for resources? Or is there a higher "status" attached to some departments over others?

> Is your organisation very hierarchical with lots of layers of management and staff, and long chains of command? Or is it a flatter, more functional structure? Do people's job titles reflect that hierarchy and imply status, or do they simply describe what a person does?

> How do people go about their jobs? Are there set processes and procedures in place to do particular jobs that people must follow? Or is there scope for creativity and initiative? Do these processes include knowledge components? Do people have time to seek and share knowledge and to reflect on it as they go about their work, or are they always under pressure to get the job done and produce results?

Often, the best way to find out whether and how an organisation's infrastructure and processes are helping or hindering people is to ask them. But before you do, be aware of the impact of both infrastructure and culture on people's willingness to tell the truth – does your organisation make it safe for them to speak their mind openly?

5.2.3 Knowledge management processes

In bringing knowledge management into your organisation, you will need to select and implement a number of processes that will help your organisation to be better at creating, finding, acquiring, organising, sharing and using the knowledge it needs to meet its goals. There are many such processes, including for example:

> Conducting knowledge audits to identify knowledge needs, knowledge resources and knowledge flows

> Creating knowledge strategies to guide the overall approach

- Connecting people with people to share tacit knowledge using approaches such as communities of practice or learning events
- Connecting people with information to share explicit knowledge using approaches such as best practices databases, and using content management processes to ensure that explicit knowledge is current, relevant and easily accessible
- Creating opportunities for people to generate new knowledge, for example through collaborative working and learning
- Introducing processes to help people seek and use the knowledge of others such as peer assists
- Teaching people to share knowledge in ways that inspire people by using storytelling techniques
- Encouraging people to prioritise learning as part of their day-to-day work, by learning before, during and after the tasks and projects they have performed

You can find more details of each of these in the KM toolbox. Some knowledge management processes are fairly new to organisations but many are not – they are simply being considered from a new perspective, that of focusing on knowledge. There is no "perfect" process nor is there a "one size that fits all". Your choice of processes will depend on the nature of your organisation.

5.2.4 Knowledge management infrastructure

A knowledge management infrastructure includes the knowledge management processes (as above) put in place to ensure good knowledge management practice, and also the organisational infrastructure that is created to enable these processes – the essential management and staff roles and responsibilities that need to be put in place to support the new processes and initiatives. In other words, the people who will take the lead in driving it all forward and bringing about the necessary changes. This infrastructure may have a number of levels, depending on the size and structure of your organisation. For example:

- Ownership and a "home"
  Where will knowledge management "live" within your organisation? Who "owns' it? Who is accountable for results? Knowledge management can reside in a range of places in organisations such as within information management, information technology, human resources, training, corporate universities, research and development, support services, or as a separate function reporting directly to the board. When making this decision, think not only about practicalities, but also about what message you are conveying about knowledge management by the "home" you are giving it, and also what impact that "home" is likely to have on the direction and development of your knowledge management efforts. For example, if knowledge management is part of IT (information technology), might issues of people and organisational culture take a back seat to technology? Or, if knowledge management is part of research and development, might there be too much focus on creating and finding new knowledge and not enough on reusing the knowledge you already have? There are no "right" answers here, but an awareness of these kinds of issues is key.
- Knowledge managers and the core team
  Your knowledge management efforts will need a core team of managers and co-ordinators to lead the way and drive the changes – to secure budgets and resources, provide direction, oversee and co-ordinate efforts, give encouragement and assistance, and monitor and evaluate progress and value. Again, the nature of your core team will depend largely on the size and structure of your organisation. A large organisation may need a Chief Knowledge Officer (or equivalent) supported by a network of Knowledge Managers and perhaps also Knowledge Co-ordinators, while a smaller organisation may simply need a single Knowledge Manager.
- Steering committees and senior supporters
  Management buy-in and support, especially at senior level, is vital to any knowledge management programme. Similarly, the more support you have from the various different functions and departments across the organisation, the better, as this can greatly speed the adoption of knowledge management. Having a steering committee with representatives from various functions can also help you to create better solutions: you get input from a range of perspectives and types of expertise, and can also clearly see the "big picture" across the organisation, allowing you to better prioritise resources and approaches.
- Knowledge brokers and champions
  In addition to your core team, there will also be people throughout the organisation whose job it is

(or part of whose job it is) to gather and share knowledge on a day-to-day basis. These people will already exist in most organisations even if there have previously been no deliberate knowledge management efforts. Such people include researchers, information workers, librarians, writers/editors/publishers, website producers, help desk advisers, internal communications people, team secretaries and administrators, etc.; they might be part of a central service such as a library or publications department, or they may be spread throughout the organisation in various departments and functions. Either way, you will need to identify these people and bring them on board, given that they are already acting as "brokers" or "champions" of knowledge and knowledge working.

> Support from outside
Your core team might want to attend some courses or conferences, do some research and reading, and make some contacts in other organisations, to "get up to speed" in knowledge management. Even if your core team comprises people with considerable experience in the field, knowledge management is a rapidly-evolving discipline and so there are always new developments of which they will need to keep abreast. Attendance at events, contacts with knowledge managers in other organisations, journal subscriptions and joining professional membership bodies can all be useful. You may also wish to bring in more specific external support in the form of knowledge management consultants for any projects for which you feel you do not have the required expertise in-house.

Obviously in smaller organisations this infrastructure will be much simpler, although the same principles will still need to be applied, albeit in less formal ways.

Similarly, whatever the size of your organisation, this infrastructure is likely to "thin out" and simplify in the long term, as knowledge management becomes integrated into the organisation and knowledge management practices become part of "the way we do things around here". As that begins to happen, dedicated knowledge management roles and functions will probably disappear, but this will take some time – it will not happen overnight.

In the meantime, you will need to create, resource and maintain this knowledge management infrastructure to drive and support your knowledge management processes.

5.3 KM Technology

5.3.1 IT and knowledge management

In the early days of knowledge management, there was a strong focus on information technology (IT). As knowledge management became the latest buzzword, technology vendors were quick to spot an opportunity to sell "knowledge management solutions" and many of the companies that led the way in knowledge management were quick to buy – to their cost. Having made significant investments in the latest systems, they then found that people simply did not use them and so the systems ended up being confined to what became known as "the knowledge management graveyard". These companies learned the hard way that knowledge management is about people, processes and technology – in that order of priority.

That being said, technology is an important enabler of many, if not most, knowledge management initiatives. Technology can support and enable knowledge management in two main ways:

1. It can provide the means for people to organise, store and access explicit knowledge and information, such as in electronic libraries or best practices databases.
2. It can help to connect people with people so that they can share tacit knowledge, such as through white pages, groupware or video conferencing.

Much of the early focus on technology was driven by an over-focus on explicit knowledge – on "getting things down" and into high-level databases. However, given the current view that up to 80% of an organisation's knowledge is always going to be in people's heads, there is a growing interest in technologies that support communication and collaboration between people.

Technology adds value when it reduces the cost, time and effort needed for people to share knowledge and information. However if it is not closely aligned with organisational needs and with people's ways of working, or if it results in information overload and so people can no longer make sense of it all, then even with the best technology in the world, you will end up right back at square one: people still cannot easily find the knowledge an information they need. The importance of this cannot be overemphasised.

The reality is that technology can only fulfil some of our needs. And how well it fulfils them depends critically on managing the knowledge behind them – content management, assigning knowledge roles etc. There are many tools that can help enable individuals and organisations to be more effective at accessing and sharing their knowledge. How well we exploit these opportunities depends more on good knowledge management than on finding the "best" piece of technology. In other words, technology by itself does not create shared knowledge: it needs to be supported by, and integrated with, relevant people and processes. Tom Davenport, a prominent author on knowledge management, is often quoted as offering the following rule of thumb: your investment in technology in terms of both cost and effort should stay under one third of the total knowledge management effort – otherwise you are going wrong somewhere.

So, what kinds of technology are we talking about? The following is a brief and simply overview, aimed at giving the non-technical manager an overall idea of some of the knowledge-enabling technologies currently available.

5.3.2 Groupware

Groupware is a term for software specifically designed for groups of people, not just individuals. As the name suggests, groupware allows groups of people to share information and to coordinate their activities over a computer network. Examples of popular proprietary groupware packages are Lotus Notes, Novell GroupWise and Microsoft Exchange. Groupware packages are diverse in the functions they offer. Most include a shared database where team members can work on common documents and hold electronic discussions. Some include group schedulers, calendars and/or e-mail. Others focus on real-time meeting support. Combined, these pieces allow team members to work on a single document, discuss ideas online, maintain records, and prioritise and schedule teamwork and meetings. A true groupware package should include several of these functions, not just one.

In recent years, intranets have emerged as cheaper and more open alternatives to proprietary groupware products so many companies are giving up proprietary groupware in favour of intranets.

5.3.3 Intranets

An intranet is simply a private Internet. Internet-type services are installed onto an organisation's internal computer network which enables it to then provide web pages and related services such as e-mail, discussion boards, access to shared documents and databases, and collaboration tools such as shared calendars and project management tools. An intranet can convey information in many forms, not just web pages but documents, tables, spreadsheets and images; it can host applications and databases. Above all, it provides connectivity that allows people to collaborate, wherever they are located.

Intranets can support knowledge sharing in a variety of ways, including:

> Ease of access and use: the use of World Wide Web browsers provides a low cost and user-friendly interface to information and applications.

> Universal access to information: information can be kept on any "server" on the network, and can be accessed from anywhere within the intranet.

> Connecting people with people: intranets can simplify interaction between people in different locations through applications such as email and discussion boards.

> Informal networks: publishing information and making contact can be quick and informal on an intranet.

> Providing a "one stop knowledge shop": an intranet can provide a single access point for internal information and knowledge, as well as providing "gateways" to the Internet for access to external information resources.

5.3.4 Connecting people with people: collaborative tools

Collaborative tools are simply electronic tools that support communication and collaboration – people working together.

Essentially they take the form of networked computer software that lets different people coordinate their work activities.

There are a number of key considerations and characteristics to bear in mind when looking at collaborative tools. These include:

> **Time**
> Is the collaboration taking place simultaneously (e.g. video conferencing) or at different times (e.g. email)?

> **Place**
> Is the collaboration taking place in the same location or at different locations?

> **Information richness**
> How much and what types of information can be conveyed? For example video conferencing conveys body language and tone of voice, whereas e-mail focuses almost exclusively on the written word and given this lack of context, can be prone to misinterpretation.

> **Social presence**
> How well can the tool help people to connect with each other and form relationships? For example an email has low social presence while a face-to-face meeting has a high social presence.

> **Technology**
> What technology is needed? How comfortable are people with technology? How easy to use is the tool? How much training will people need?

In short, no one tool is ideal for all situations.

Collaborative tools can provide a number of benefits, such as:

- allowing people to work together in teams, over a network, irrespective of location or time
- enabling the sharing of tacit knowledge between a wide range of people
- the ability to access the knowledge of experts wherever they are located
- savings on meeting costs – travel and subsistence, meeting rooms, etc.

The various tools can be provided as part of a groupware package, over an intranet, or in some cases as standalone tools. Common collaborative tools include the following:

> **Email**
> A simple electronic version of written mail, and undoubtedly the most widely used collaborative tool. Messages are sent via an electronic network and attachments can be added such a copies of documents and presentations. Email can be used between individuals, or to broadcast messages to a wider audience.

> **Discussion boards**
> Discussion boards (also known as message boards, bulletin boards or chat rooms) give people the ability to post and reply to messages in a common area. Sometimes a leader or facilitator will moderate these boards. Their purpose is to provide an "informal meeting place" a bit like a café. People can ask for advice and share information around topics of interest. Discussion boards are often used to support communication within communities of practice.

> **Video conferencing**
> Video conferences can either be done using specialized video facilities, or from people's desktops using computer software. Video conferencing works well for situations that require a degree of trust and relationship building, for discussing issues and exploring ideas, and in situations where you don't need a detailed permanent record to be generated automatically. Consideration needs to be given to the quality of the video link, as many of the benefits can be lost through poor quality. Also, be aware that not everyone likes, or feels comfortable with, video conferencing An alternative is audio (telephone) conferencing, which tends to work best when participants already know each other.

> **Project support tools**
> There are a number of tools that enable work groups and project teams to share documents and exchange messages across different locations in "real time". For example, when a group is working on a shared document, there needs to be a tool to make the document centrally available, allow people to make changes, synchronise the changes made, and ensure that the most up-to-date version is clearly available. Similarly, remote project teams can take advantage of "electronic whiteboards" to brainstorm together, generate lists of options, draw or map concepts visually to aid understanding, display and analyse data together etc.

> **Work flow tools**
> Work flow tools are developed to model typical processes that take place in organisations. They enable people to work together on shared tasks, with some of the core process knowledge embedded in the design of the work flow software application. An example would be a purchasing or transaction process, starting with the creation of an order and ending with the supply of goods. Where several people and a sequence of documents and processes are involved, automation can help speed up the process and also provide information about what stage the process is at, at any given time.

> **E-learning tools**
> E-learning is a rapidly growing field and uses information technology to deliver learning and training to people electronically at their desktop. There is a wide variety of tools and technologies available to support e-learning, many of which include facilities for learners in different locations to work together on assignments, case studies and projects.

> **Virtual working tools**
> At the highly sophisticated end of the spectrum, technologies are emerging that allow the knowledge and expertise of a person in one location to be directly applied in another location in real time. Such technologies allow knowledge to be not only shared, but applied, remotely. For example, in 2001 a pioneering surgical procedure was tested in which two surgeons in New York operated on a patient in France, using joysticks and voice commands to direct three robotic arms in the operating room. This was the first instance of remote surgery on a human. Similar technologies have already been used quite extensively in fields such as engineering.

5.3.5 Connecting people with information: managing content

Whether you use an intranet or some other form of groupware to network and share documents, applications and collaborative tools across your organisation, you will need processes in place to ensure that users can easily and quickly find the information they need. You need to consider content management. There are three critical aspects of managing content:

> **Collecting the content**
> Including issues such as: where will the content come from; who will collate it; how will they find and evaluate sources to ensure that quality and reliability of content; how will they ensure it meets users' needs both now and in the future, as needs change; how will they weed out out-of-date content; how will you ensure that your content complies with issues such as copyright, legal liability, data protection, and information risk and security?

> **Organising the content**
> How will the content be organised so that people can easily find what they need, when they need it? How will content be classified and indexed, and what terms and language will you use? Will you use taxonomies? A thesaurus?

> **Retrieving and using the content**
> How will people find and access the information they need? What combination of navigation tools will you offer them – menus, maps, search engines? What balance will you strike between "pushing" information to users (e.g. through alerting services) or waiting for users to "pull" information out for themselves (e.g. using search engines)?

Be aware that while there are a number of content management systems and software packages available, an important element is people to manage the content; this function is often best performed by people with a background in librarianship and information management. Some examples of tools and processes used in content management are:

> **Taxonomies**
> A taxonomy is a hierarchical structure for organising a body of knowledge; it gives a framework for understanding and classifying that knowledge – how to group it and how the various groups relate to each other. In content management, the purpose of taxonomy is to organise information so that users can more easily navigate their way through it. Taxonomies can be generated either manually or automatically using a software programme.

> **Thesauri**
> A thesaurus is a list of the various terms and language that are used to describe a body of knowledge, and which specifies the relationship between the terms: antonyms and synonyms,

broader terms and narrower terms, etc. In content management, the aim of a thesaurus is to enable content to be indexed in a variety of ways so that different users who tend to use different terms can still find it.

> Search engines
A search engine is a piece of software that carries out searches for information across multiple sources. Search engines vary widely in their level of sophistication. Some simply allow users to search for documents that contain a specific word or phrase, which can leave users having to sift through great deal of irrelevant information. More advanced search engines allow users to construct more specific searches, enabling them to narrow their search and reduce the amount of irrelevant material retrieved.

> Portals
A portal is a website or a web page that provides your main point of entry into an intranet or the Internet, and which gathers and integrates information from various sources into a single location. Portals are essentially "personalised gateways" – a kind of one-stop-shop for information that is personalised, either to an organisation's needs or to individual people's needs.

The purpose is to avoid information overload by providing at each person's desktop access to the specific information and tools they need to do their job, while filtering out those they don't need. Think about your computer desktop for example: you will probably have arranged your applications and files in a way that suits you, making the ones that you use most the easiest to find. A portal does the same thing, except through an intranet or the Internet. On the Internet, for example, you may have set up a personalised portal on AOL, Yahoo or Freeserve.

Portals are relatively new in organisations, largely because an effective portal is technically difficult to create, and so the technology is still evolving.

5.3.6 Knowledge creation technologies

As well as technologies designed to enable the sharing of knowledge, there is also an increasing number of tools aimed at supporting the creation of knowledge – helping to generate information and knowledge from data. A few examples are briefly mentioned here for general interest only:

> Data mining
Tools that analyse data in very large databases and look for trends and patterns that can be used to improve organisational processes.

> Information visualization
Computer-supported interactive visual representations of abstract data to help improve understanding.

> Decision trees
Provide a structure in which alternative decisions and the implications of taking those decisions can be displayed and evaluated.

> Root cause analysis
A method or series of actions taken to find out why a particular failure or problem exists, and correcting those causes.

5.3.7 More information

Four articles by David Skyrme Associates available online:

Intranets: sharing organizational knowledge
Insights, No. 25

Getting to grips with groupware
Insights, No. 7

Is content king?
I3 Update No. 59

Portals: panacea or pig?
I3 Update No. 44

6 MEASURING THE EFFECTS OF KNOWLEDGE MANAGEMENT

6.1 Why measure?

Measurement is undoubtedly the least developed aspect of knowledge management, which is not surprising given the difficulties in defining it let alone measuring it. In fact some practitioners feel that measurement is premature at this stage and that trying to measure knowledge before you fully understand how knowledge is created, shared and used is likely to lead you to focus on the wrong things. Elaborate measurement systems, they say, cannot currently be justified because we simply do not yet know enough about the dynamics and impact of knowledge.

That being said, in practice, few organisations have the luxury of being allocated resources to implement something without being required to demonstrate its value. Without measurable success, enthusiasm and support for knowledge management is unlikely to continue. And without measurable success, you are unlikely to be able to what works and what doesn't and therefore make an informed judgement regarding what to continue doing, and what to adjust.

6.2 What to measure? Common measurement approaches

There are a number of approaches that are increasingly being used to measure the value of, and progress in, knowledge and knowledge management in organisations. Some of the more common approaches are outlined here for the purposes of providing a general overview.

6.2.1 Measuring the impact of knowledge management on the organisation's performance

Given that the whole point of knowledge management is to improve the performance of your organisation and to help it to achieve its objectives, the best and most logical approach is tie-in measurement of knowledge management with your organisation's overall performance measurement systems. This can be done either at an organisational level, or for individual projects and processes.

However, one limitation of this approach is that if knowledge management practices are made an integral part of work, you cannot be sure of the relative contribution of those knowledge management practices to the success of a project or process, versus other factors. In view of this, O'Dell and Grayson, in Chapter 12 of their book "If only we knew what we knew: the transfer of internal knowledge and best practice (1998)" recommend a two-pronged approach that seeks to measure both outcomes and activities.

Measuring outcomes focuses on the extent to which a project or a process achieves its stated objectives. The success of the project or process serves as a proxy measure for the success of the knowledge management practices embedded in it. In other words, knowledge management is seen as an integral tool for improving a project or process, rather than as a separate thing. For example, outcomes might be measured in terms of the reduced cost of a process, improved efficiency, the reduction in time taken to do it, the improved quality of delivery, etc.

Measuring activities then shifts the focus onto the specific knowledge management practices that were applied in the project or process. What were the specific knowledge management activities behind this practice and what was their effect? In measuring activities, you are looking specifically at things like how often users are accessing, contributing to, or using the knowledge resources and practices you have set up. Some of these measures will be quantitative ("hard") measures such as the number and frequency of hits or submissions to an intranet site per employee. However these measures only give part of the picture – they do not tell you why people are doing what they are doing. Hence to complete the picture, you will also need qualitative ("soft") measures by asking people about the attitudes and behaviours behind their activities.

6.2.2 The balanced scorecard

An increasingly popular approach to measuring an organisation's performance, and one that is being widely adopted in knowledge management, is the balanced scorecard. The advantage of this approach in knowledge management terms is that it directly links learning to process performance, which in turn is linked with overall organisational performance. Developed by Kaplan and Norton, the balanced scorecard focuses on linking an organisation's strategy and objectives to measures from four key perspectives: financial, customers, internal processes, and learning and growth. In contrast to traditional accounting measures, the balanced scorecard shifts the focus from purely financial measures to include three key measures of intangible success factors. These roughly equate to the three components of intellectual capital – namely human capital (learning), structural capital (processes), and customer capital. The four perspectives can be framed as follows:

1 Financial
How do we look to our "shareholders" (or governing bodies)?

2 Customer
How do our patients see us? Are we meeting their needs and expectations?

3 Internal processes
What do we need to do well in order to succeed? What are the critical processes that have the greatest impact on our patients and our financial objectives?

4 Learning and growth
How can we develop our ability to learn and grow in order to meet our objectives in the above three areas?

This knowledge management, which is about learning and growth, is measured as an integral and yet distinct part of overall organisational performance.

The balanced scorecard approach can be applied to individual initiatives as well as to a whole organisation.

6.2.3 Return on investment (ROI)

Most initiatives that require resources will be expected to show a return in investment – what benefits did we get to justify the costs involved – and knowledge management in usually no exception. The problem is that both the costs and the benefits of knowledge management can be notoriously difficult to pin down. While the costs associated with an investment in information technology can be relatively straightforward to identify, other costs can be less so, such as for projects that involve an amalgam of resources from across the organisation, or those inherent in challenging an organisation's culture. On the benefits side, how do you measure things like increased knowledge sharing, faster learning or better decision-making?

A number of approaches have been developed for showing financial returns on knowledge assets. Such approaches tend to be rather complex, and therefore are probably more appropriate to organisations that are reasonably advanced in their knowledge management efforts, rather than just starting out.

6.2.4 The knowledge management life cycle

Some organisations measure the progress of their knowledge management activities in terms of their maturity – how far "down the line" they are in implementing knowledge management practices and ways of working. The American Productivity and Quality Center has developed a framework known as Road Map to Knowledge Management Results: Stages of Implementation. The aim is to provide organisations with a map to guide them from getting started right through to "institutionalising" knowledge management – embedding it in the organisation and making it an integral part of the way an organisation works. The map has five stages:

1 Get started

2 Develop a strategy

3 Design and launch a knowledge management initiative

4 Expand and support

5 Institutionalise knowledge management

There are measures associated with each stage.

6.2.5 Employee surveys

Given the importance of people in knowledge management, employee surveys can be a useful additional to your measurement toolbox. Surveys can be used to assess aspects of organisational culture and the extent to which people's opinions, attitudes and behaviours are, or are not, changing. Obviously such surveys measure people's subjective perceptions and these may or may not reflect reality, but in many ways that can be their very benefit, as people's perceptions will determine their behaviours with respect to knowledge management. In order to be effective, it is vital that any such surveys are carried out by people with the required expertise, whether that is through in-house capabilities or by hiring external consultants.

6.2.6 Measuring the value of knowledge assets

As well as measuring the progress and value of knowledge management initiatives, organisations are also developing ways to measure the value of their knowledge assets. The traditional balance sheet is increasingly being regarded as an incomplete measure of an organisation's worth, as it does not place a value on intangible assets such as knowledge or intellectual capital. As already mentioned, intellectual capital is commonly regarded as having three components: human capital (the knowledge and skills of people), structural capital (the knowledge inherent in an organisation's processes and systems), and customer capital (customer relationships). There are a number of key models for measuring the value of intellectual capital. Among the best-known are:

- The Skandia Navigator and its associated Value Creation Model
  Developed by Swedish financial services company Skandia, this approach uses the metaphor of a house whose roof represents an organisation's financial assets and whose foundations represent innovation and renewal. The model includes a long list of measures, which are organised into five categories, namely: financial, customer, process, renewal and development, and human.
- Sveiby's Intangible Assets Monitor
  Developed by knowledge management pioneer Karl Erik Sveiby, the monitor categorises intangible assets into human competence, internal structure and external structure, with further subdivisions into indicators of efficiency and utilisation, stability, and growth and renewal.
- Intellectual Capital Services' IC- Index
  Originally developed in Scandinavia and Australia by Johan and Göran Roos, the index identifies four categories of intellectual capital: relationship, human, infrastructure and innovation; it then looks at the relative importance of each, and also at the impact of changes in intellectual capital.
- Philip M'Pherson's Inclusive Value Methodology (IVM)
  A model in which users create hierarchies of intangibles to which they assign value ratings according to priorities, then a computer model determines the overall value rating and tests for areas of risk.

6.3 How to measure?

Melissie Clemmons Rumizen outlines the following steps in developing measures, in Chapters 19-22 of her book "The complete idiot's guide to knowledge management (2002)":

6.3.1 Revisit your goals

Your starting point for measuring any knowledge management initiative will be the original goals of that initiative: what is it that you set out to achieve? Developing measures will often lead you to get clearer about how you define your goals in the first place; if your goals are not concrete and clear enough, then measuring your success or progress against them will be difficult. Hence ensure that your goals define clearly what constitutes success in measurable terms.

6.3.2 Know the audience for your measures

In defining success, you will often find that different people have different ideas about what constitute success. Managers who approve the allocation of resources will want to know about the returns on their investment. Users of the knowledge management initiative will want to know how it has benefited them and whether their participation has been worthwhile. Other beneficiaries of the initiative, such as patients, will want to know how they have gained.

6.3.3 Define the measures

Define what exactly you are going to measure, and what measurement approach or approaches you intend to take. Ensure that your measures are:

> Valid
  They actually measure what they are intended to measure rather than something else,

> Reliable
  They give consistent results

> Actionable
  They give information that can be acted upon if necessary.

6.3.4 Decide what data will be collected and how it will be collected

This is a process of "putting the meat on the bones" – spelling out the details: what data will be collected, who will collect it, how, when, where, etc.?

6.3.5 Analysing and communicating the measures

When analysing and presenting the results, be sure to refer back to your original goals and your audience. Aim to present results in a way that answers their questions in a meaningful way, rather than simply presenting facts and figures.

6.3.6 Review your combination of measures

Monitor and evaluate how your measures are working. Developing measures is a process of trial and error – don't necessarily expect to get it right first time. Similarly, remember that as objectives and situations change over time, so will your measures need to.

Additional pointers emphasised by other practitioners include:

> Measuring for the sake of measuring is a waste of time – be sure that you are measuring for a specific purpose or purposes.

> Be sure that some kind of action or decision will be taken as a result of your measures

> Don't try to measure everything; instead, focus on what is important. Trying to measure too much not only requires a great deal of work, it also tends to dilute the important issues.

> If your organisation already has a measurement system in place, then you can use those measures.

> If your knowledge management initiatives work, then you might assume that this will show up in your organisation's other performance measures. Of course there is no guarantee that existing measures are good ones so you might like to look into them, but there are two major advantages to "piggy-backing" on existing measures: first, they are already accepted practice in the organisation, and second, they are most likely measuring things that are important to the organisation.

6.3.7 More information

American Productivity and Quality Center (APQC)
Road Map to Knowledge Management Results: Stages of Implementation

The Balanced Scorecard Collaborative

7 KNOWLEDGE MANAGEMENT GLOSSARY OF TERMS

The glossary (go directly to http://libraries.nelh.nhs.uk/knowledgemanagement/default.asp?page=GLOSSARY) is arranged in alphabetical order. Where terms are also known by other names, or where two or more terms are closely related, you can click on links from one to the other.

If there is a term that is not included and you think it should be, if you feel that a better explanation of a term is needed, or if you want to suggest an alternative explanation, please, please contact us.

Information Management and Mining

Knowledge Management

BCA-VI

March, 2020

Information Management and Mining

- Data Vs. Information: Is there a difference?
  - Information is a human interpretation of knowledge
- Databases and querying
- Data mining

What is Data Mining?

- **Many Definitions!**
  - Non-trivial extraction of implicit, previously unknown and potentially useful information from data
  - Exploration & analysis, by automatic or semi-automatic means, of large quantities of data in order to discover meaningful patterns

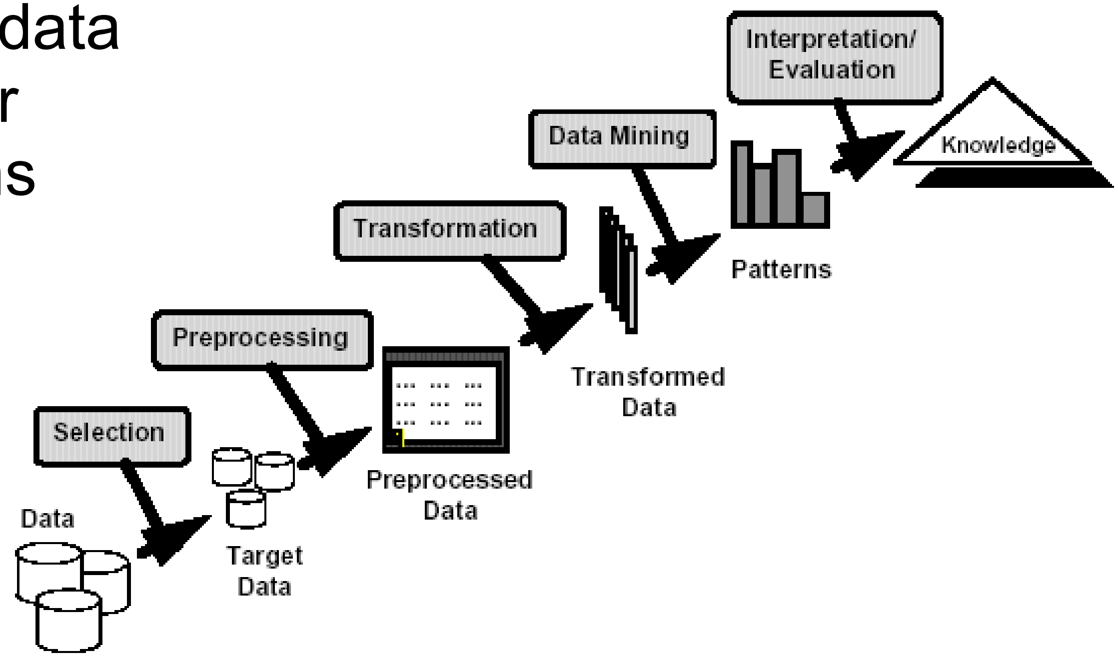

What is (not) Data Mining?

What is not Data Mining?

– Look up phone number in phone directory

– Query a Web search engine for information about "Amazon"

What is Data Mining?

– Certain names are more prevalent in certain US locations (O'Brien, O'Rurke, O'Reilly… in Boston area)

– Group together similar documents returned by search engine according to their context (e.g. Amazon rainforest, Amazon.com,)

Origins of Data Mining

- Draws ideas from machine learning/AI, pattern recognition, statistics, and database systems
- Traditional Techniques may be unsuitable:
  - Enormity of data
  - High dimensionality of data
  - Heterogeneous, distributed nature of data

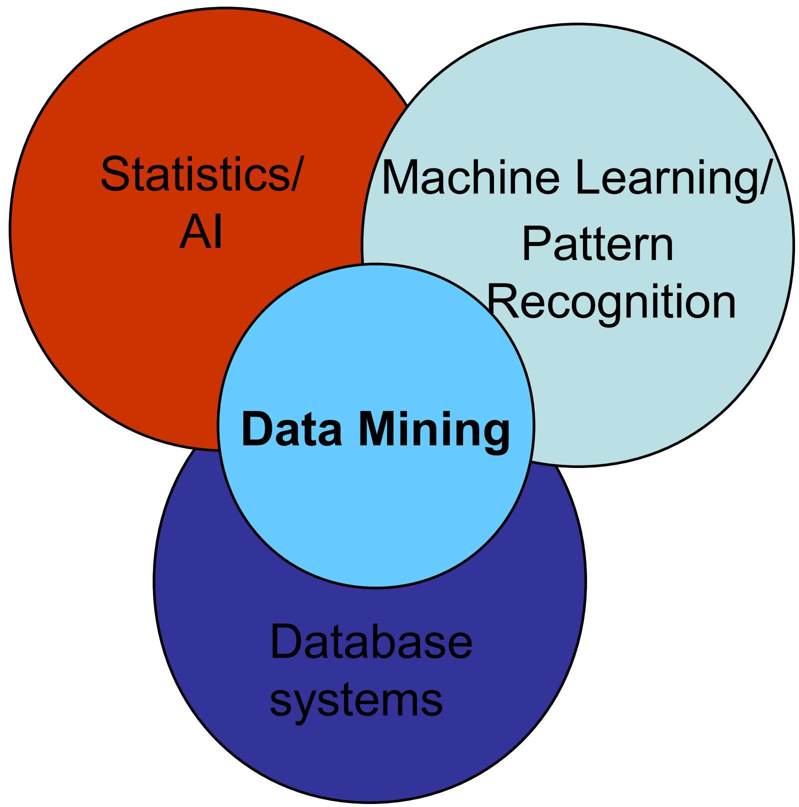

Data Mining Tasks

- Description Methods
  - Find human-interpretable patterns that describe the data.
- Prediction Methods
  - Use some variables to predict unknown or future values of other variables.

Data Mining Tasks...

- Classification [Predictive]
- Clustering [Descriptive]
- Association Rule Discovery [Descriptive]
- Sequential Pattern Discovery [Descriptive]
- Regression [Predictive]
- Deviation Detection [Predictive]

Classification Example

categorical categorical continuous class

Tid	Refund	Marital Status	Taxable Income	Cheat
1	Yes	Single	125K	No
2	No	Married	100K	No
3	No	Single	70K	No
4	Yes	Married	120K	No
5	No	Divorced	95K	Yes
6	No	Married	60K	No
7	Yes	Divorced	220K	No
8	No	Single	85K	Yes
9	No	Married	75K	No
10	No	Single	90K	Yes

Refund	Marital Status	Taxable Income	Cheat
No	Single	75K	?
Yes	Married	50K	?
No	Married	150K	?
Yes	Divorced	90K	?
No	Single	40K	?
No	Married	80K	?

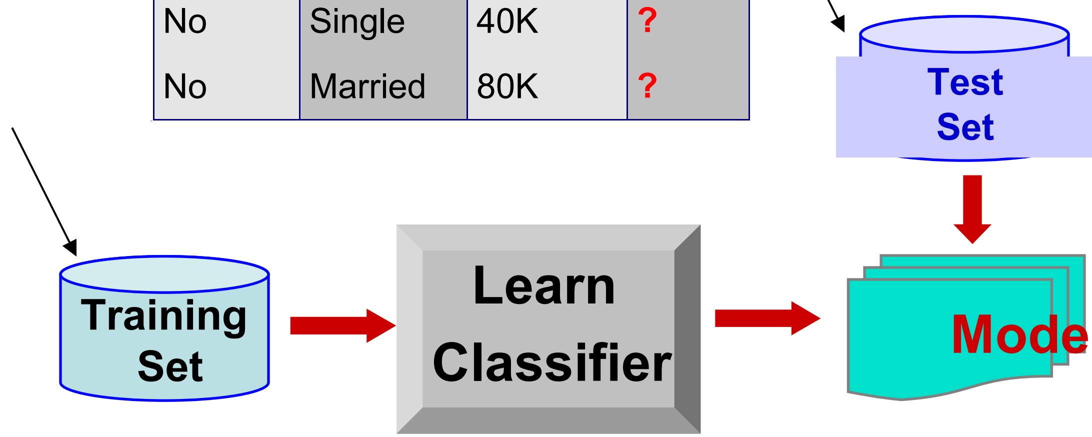

Classification: Application

- Fraud Detection
  - Goal: Predict fraudulent cases in credit card transactions.
  - Approach:
    - Use credit card transactions and the information on its account-holder as attributes.
      - When does a customer buy, what does he buy, how often he pays on time, etc
    - Label past transactions as fraud or fair transactions. This forms the class attribute.
    - Learn a model for the class of the transactions.
    - Use this model to detect fraud by observing credit card transactions on an account.

Clustering

- Given a set of data points, each having a set of attributes, and a similarity measure among them, find clusters such that
  - Data points in one cluster are more similar to one another.
  - Data points in separate clusters are less similar to one another.

Illustrating Clustering

Intracluster distances are minimized

Intercluster distances are maximized

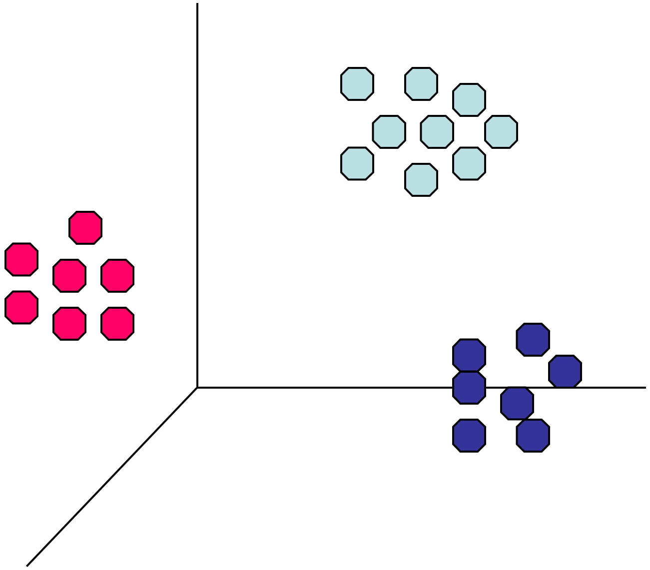

Clustering: Application

- Document Clustering:
  - Goal: To find groups of documents that are similar to each other based on the important terms appearing in them.
  - Approach: To identify frequently occurring terms in each document. Form a similarity measure based on the frequencies of different terms. Use it to cluster.

Illustrating Document Clustering

- Clustering Points: 3204 Articles of Los Angeles Times.
- Similarity Measure: How many words are common in these documents (after some word filtering).

Category	***Total Articles***	***Correctly Placed***
Financial	555	364
Foreign	341	260
National	273	36
Metro	943	746
Sports	738	573
Entertainment	354	278

Association Rule Discovery

- Given a set of records each of which contain some number of items from a given collection;
  - Produce dependency rules which will predict occurrence of an item based on occurrences of other items.

TID	*Items*
1	**Bread, Coke, Milk**
2	**Beer, Bread**
3	**Beer, Coke, Diaper, Milk**
4	**Beer, Bread, Diaper, Milk**
5	**Coke, Diaper, Milk**

Rules Discovered:

{Milk} --> {Coke}
{Diaper, Milk} --> {Beer}

Association Rule Discovery: Application

- Supermarket shelf management.
  - Goal: To identify items that are bought together by sufficiently many customers.
  - Approach: Process the point-of-sale data collected with barcode scanners to find dependencies among items.
  - A classic rule --
    - If a customer buys diaper and milk, then he is very likely to buy beer.
    - So, don’t be surprised if you find six-packs stacked next to diapers!

Regression

- Predict a value of a given continuous valued variable based on the values of other variables, assuming a linear or nonlinear model of dependency.
- Greatly studied in statistics, neural network fields.
- Examples:
  - Predicting sales amounts of new product based on advetising expenditure.
  - Predicting wind velocities as a function of temperature, humidity, air pressure, etc.
  - Time series prediction of stock market indices.

Deviation/Anomaly Detection

- Detect significant deviations from normal behavior
- Applications:
  - Credit Card Fraud Detection

  - Network Intrusion Detection

Typical network traffic at University level may reach over 100 million connections per day

Why Mine Data? Commercial Viewpoint

- Lots of data is being collected and warehoused
  – Web data, e-commerce
  – purchases at department/ grocery stores
  – Bank/Credit Card transactions

- Computers have become cheaper and more powerful
  – New computational frameworks (Map-Reduce from Google)
- Competitive Pressure is Strong
  – Provide better, customized services for an *edge*

Why Mine Data? Scientific Viewpoint

- Data collected and stored at enormous speeds (GB/hour)
  - remote sensors on a satellite
  - telescopes scanning the skies
  - microarrays generating gene expression data
  - scientific simulations generating terabytes of data

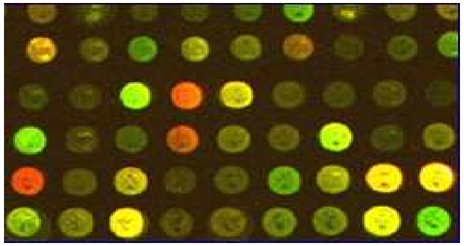

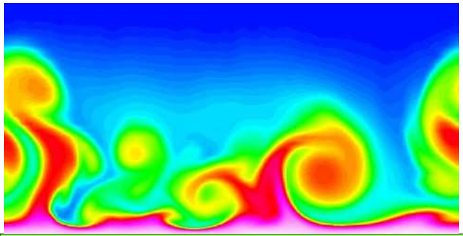

- Traditional techniques infeasible for raw data
- Data mining may help scientists
  - in classifying and segmenting data
  - in Hypothesis Formation

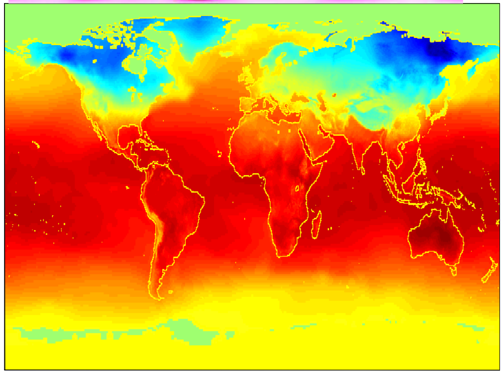

Challenges in Data Mining

- Scalability
- Dimensionality
- Complex and Heterogeneous Data
- Data Quality
- Data Ownership and Distribution
- Privacy Preservation
- Streaming Data

Data Mining Techniques Decision Trees

- One of the most common used data mining techniques because its model is easy to understand for users.

- In decision tree technique, the root of the decision tree is a simple question or condition that has multiple answers. Each answer then leads to a set of questions or conditions that help us determine the data so that we can make the final decision based on it. For example, We use the following decision tree to determine whether or not to play tennis:

Data Mining Techniques Decision Trees

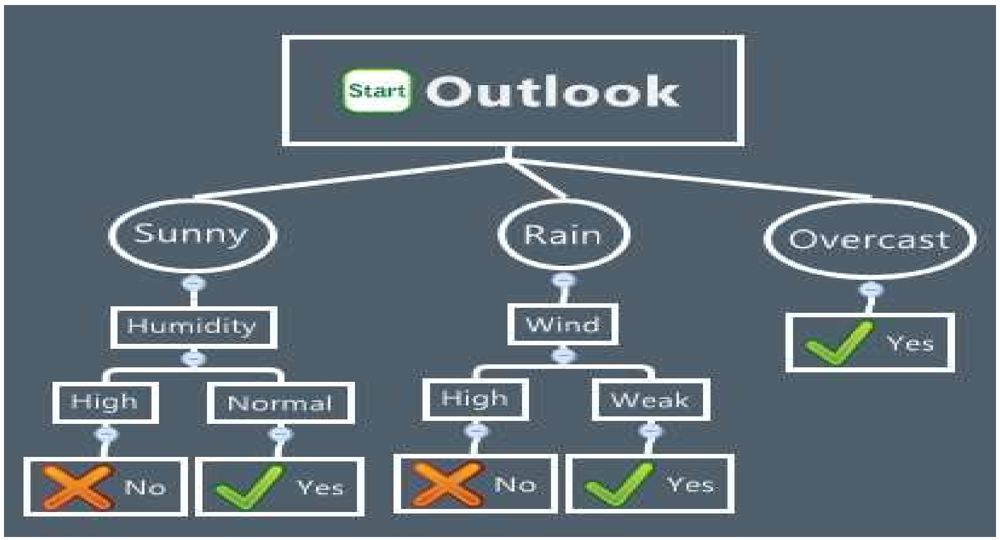

Starting at the root node, if the outlook is overcast then we should definitely play tennis. If it is rainy, we should only play tennis if the wind is the week. And if it is sunny then we should play tennis in case the humidity is normal.

Data Mining Techniques Link Analysis

Link analysis

In network theory, link analysis is a data-analysis technique used to evaluate relationships (connections) between nodes. Relationships may be identified among various types of nodes (objects), including organizations, people and transactions. Link analysis has been used for investigation of criminal activity (fraud detection, counterterrorism, and intelligence), computer security analysis, search engine optimization, market research, medical research, and art.

Data Mining Techniques Memory Based Reasoning MBR

- Memory-Based Reasoning(MBR) is
  - Identifying similar cases from experience
  - Applying the information from these cases to the problem at hand.
  - MBR finds neighbors similar to a new record and uses the neighbors for classification and prediction.

- It cares about the existence of two operations
  - Distance function ; assigns a distance between any two records
  - Combination function ; combines the results from the neighbors to arrive at an answer.

- Applications of MBR span many areas;
  - Fraud detection
  - Customer response prediction
  - Medical treatments
  - Classifying responses

Data Mining Techniques

We often combine two or more of those data mining techniques together to form an appropriate process that meets the business needs.

Data Mining works with Warehouse Data

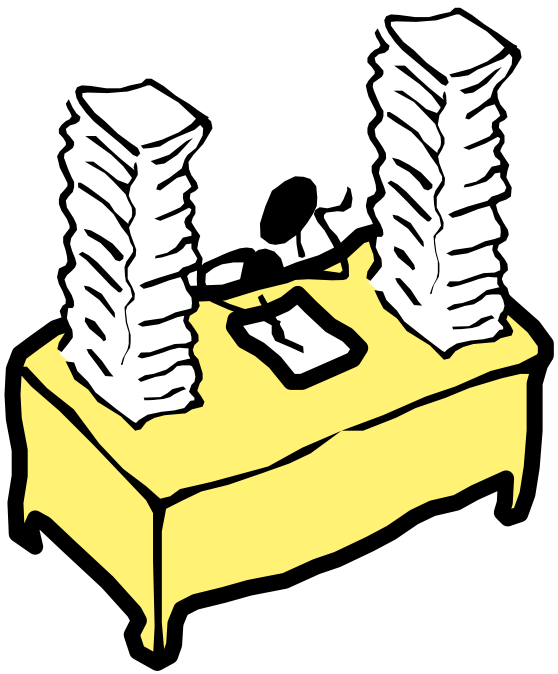

- Data Warehousing provides the Enterprise with a memory

Data Mining provides the Enterprise with intelligence

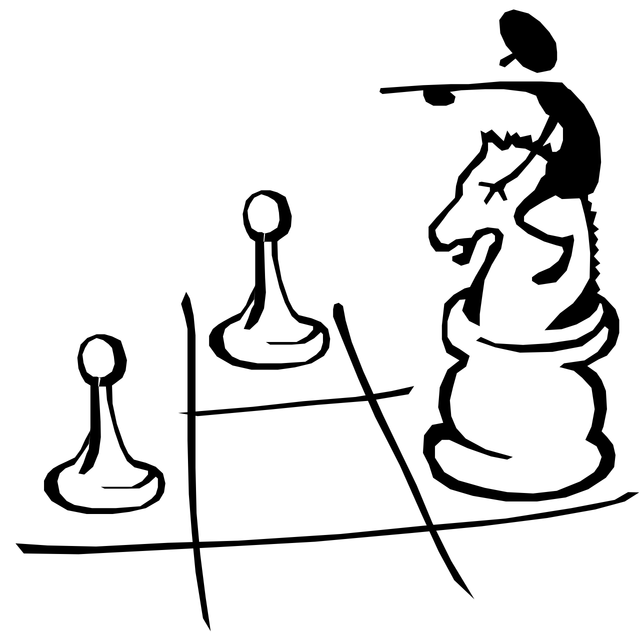

We want to know ...

- Given a database of 100,000 names, which persons are the least likely to default on their credit cards?
- Which types of transactions are likely to be fraudulent given the demographics and transactional history of a particular customer?
- If I raise the price of my product by Rs. 2, what is the effect on my ROI?
- If I offer only 2,500 airline miles as an incentive to purchase rather than 5,000, how many lost responses will result?
- If I emphasize ease-of-use of the product as opposed to its technical capabilities, what will be the net effect on my revenues?
- Which of my customers are likely to be the most loyal?

Data Mining helps extract such information

Data Mining in Use

- The US Government uses Data Mining to track fraud
- A Supermarket becomes an information broker
- Basketball teams use it to track game strategy
- Cross Selling
- Warranty Claims Routing
- Holding on to Good Customers
- Weeding out Bad Customers

What makes data mining possible?

- Advances in the following areas are making data mining deployable:
  - data warehousing
  - better and more data (i.e., operational, behavioral, and demographic)
  - the emergence of easily deployed data mining tools and
  - the advent of new data mining techniques.
      - -- Gartner Group

Genetic Algorithms

Genetic Algorithm (GA) is a search-based optimization technique based on the principles of **Genetics and Natural Selection**. It is frequently used to find optimal or near-optimal solutions to difficult problems which otherwise would take a lifetime to solve. It is frequently used to solve optimization problems, in research, and in machine learning.

Genetic Algorithms

Introduction to Optimization
Optimization is the process of **making something better**. In any process, we have a set of inputs and a set of outputs as shown in the following figure.

Genetic Algorithms

Optimization refers to finding the values of inputs in such a way that we get the "best" output values. The definition of "best" varies from problem to problem, but in mathematical terms, it refers to maximizing or minimizing one or more objective functions, by varying the input parameters.

The set of all possible solutions or values which the inputs can take make up the search space. In this search space, lies a point or a set of points which gives the optimal solution. The aim of optimization is to find that point or set of points in the search space.

Knowledge Management

WHAT IS KNOWLEDGE MANAGEMENT

- Knowledge management, or KM, is the process through which organizations generate value from their intellectual property and knowledge-based assets

- KM involves the creation, dissemination, and utilization of knowledge

- It is also viewed as the intersection between People, Processes and Technology

WORKING DEFINITION

Knowledge Management is the explicit and systematic management of vital knowledge – and its associated processes of creation, organisation, diffusion, use and exploitation.

2 KEY THRUSTS

Sharing existing knowledge
"Knowing what you know"

Knowledge for Innovation
"Creating <u>and</u> Converting"

KNOWLEDGE MAP

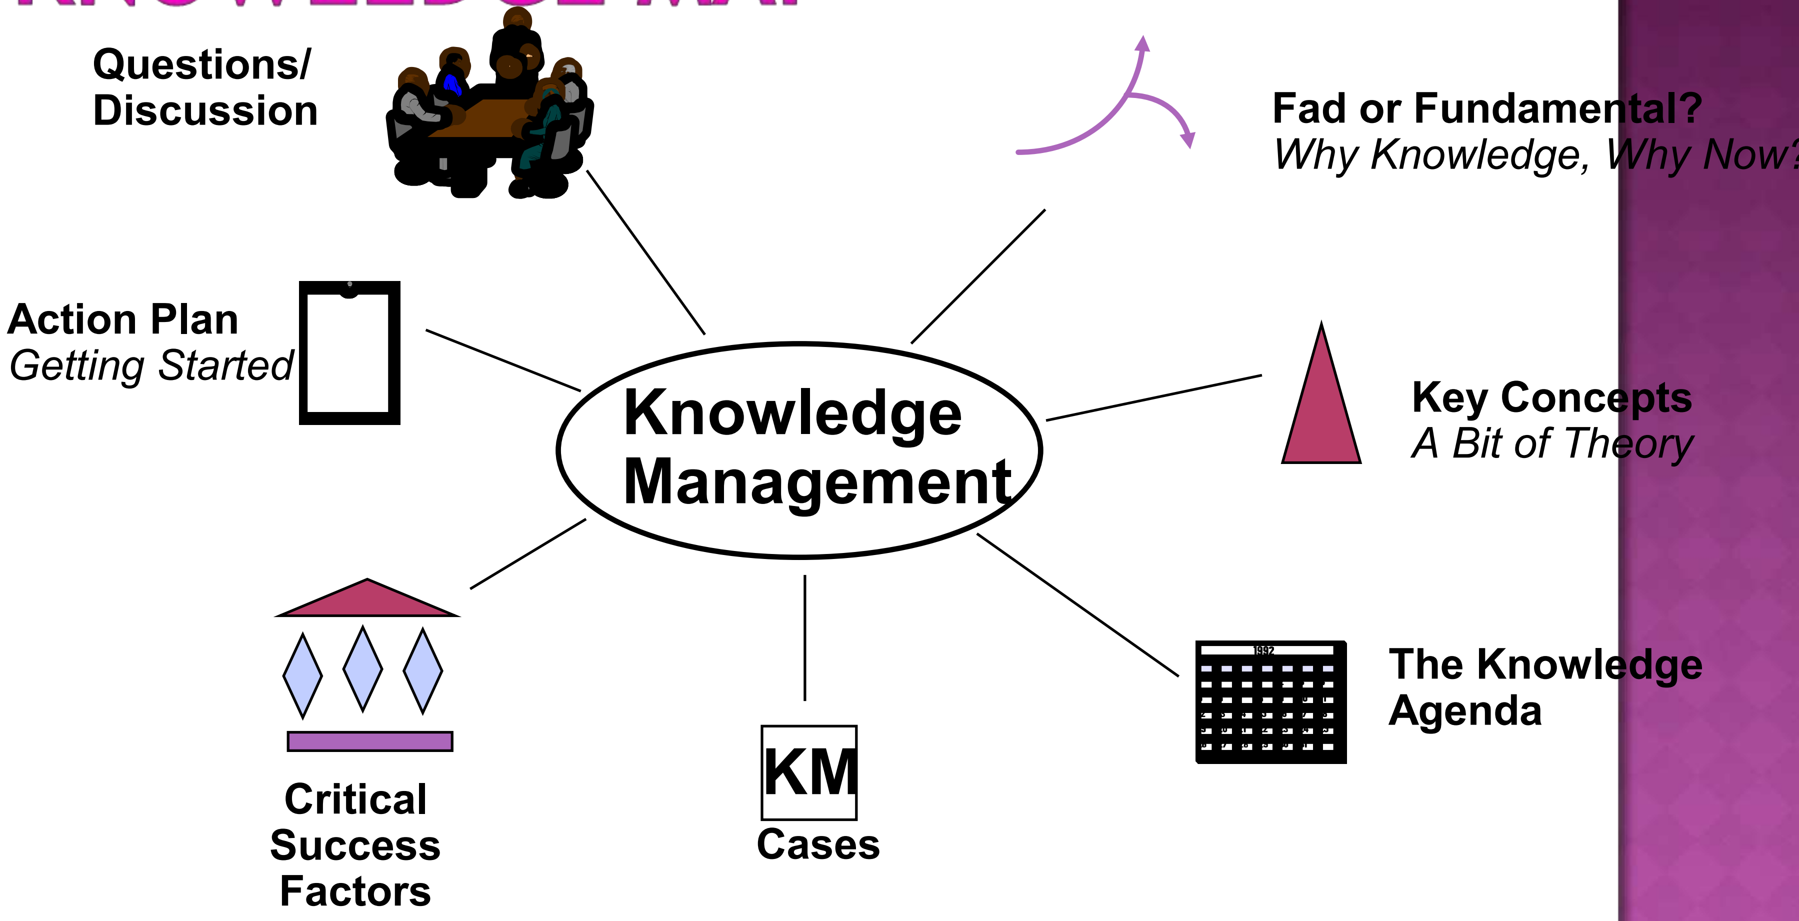

FAD OR FUNDAMENTAL?

Innovation in Products, Services and Processes

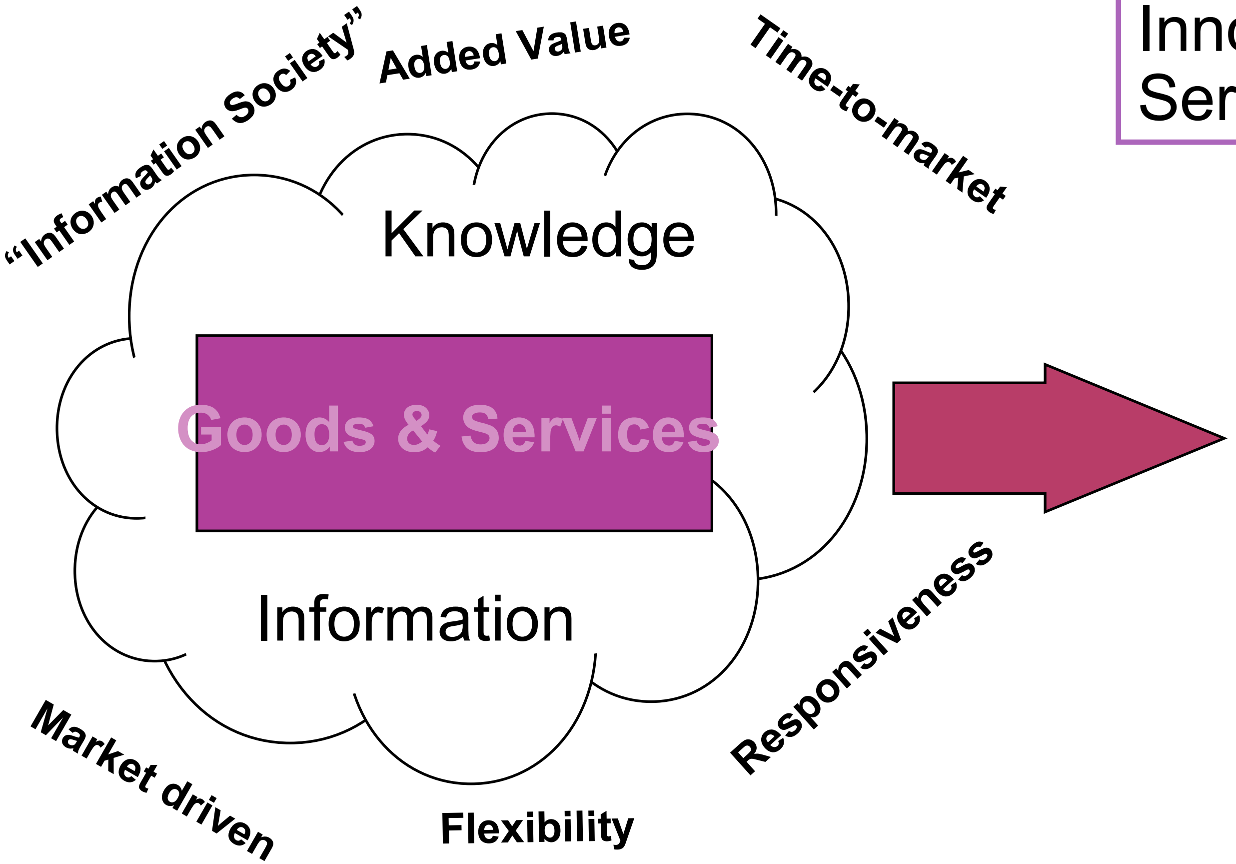

- Global Customers
- Changing Needs
- Time-to-market
- 'Smart' Products
- Customization
- Service
- Quality
- Intangibles

ROOTS OF KNOWLEDGE MANAGMENT

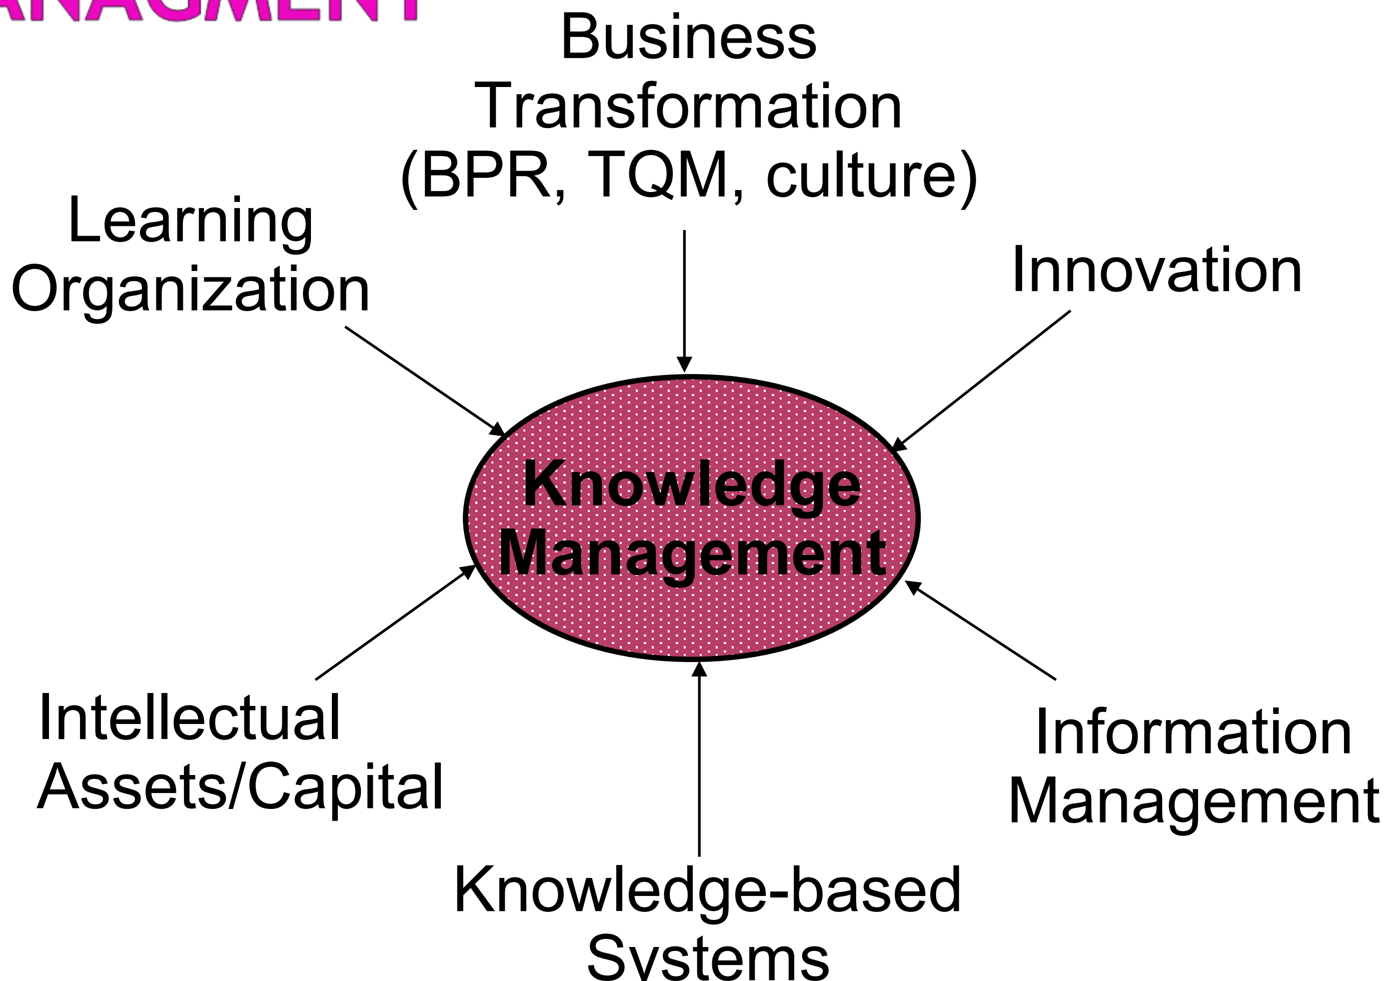

KNOWLEDGE IS DIFFERENT (1)

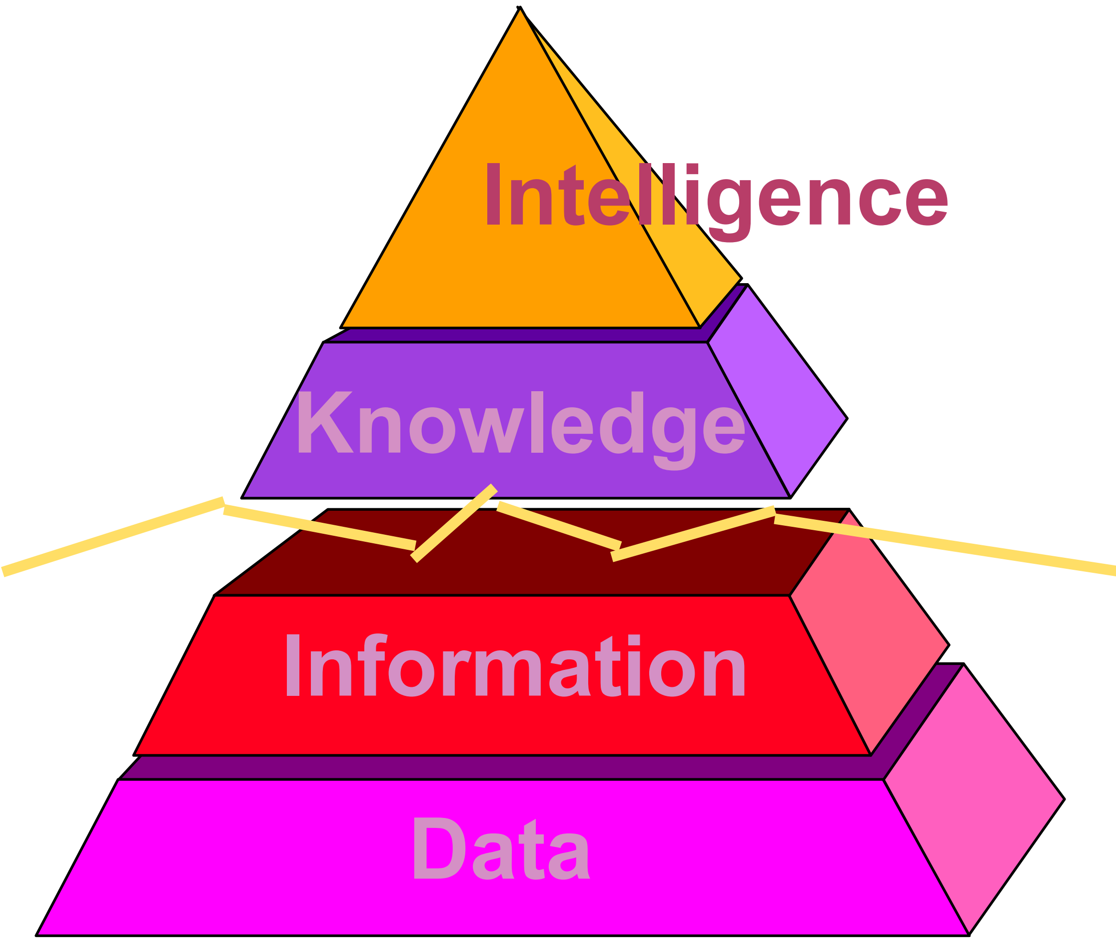

Human, judgemental

Contextual, tacit
Transfer needs learnin

Codifiable, explicit
Easily transferable

KNOWLEDGE MANAGEMENT COMPONENTS

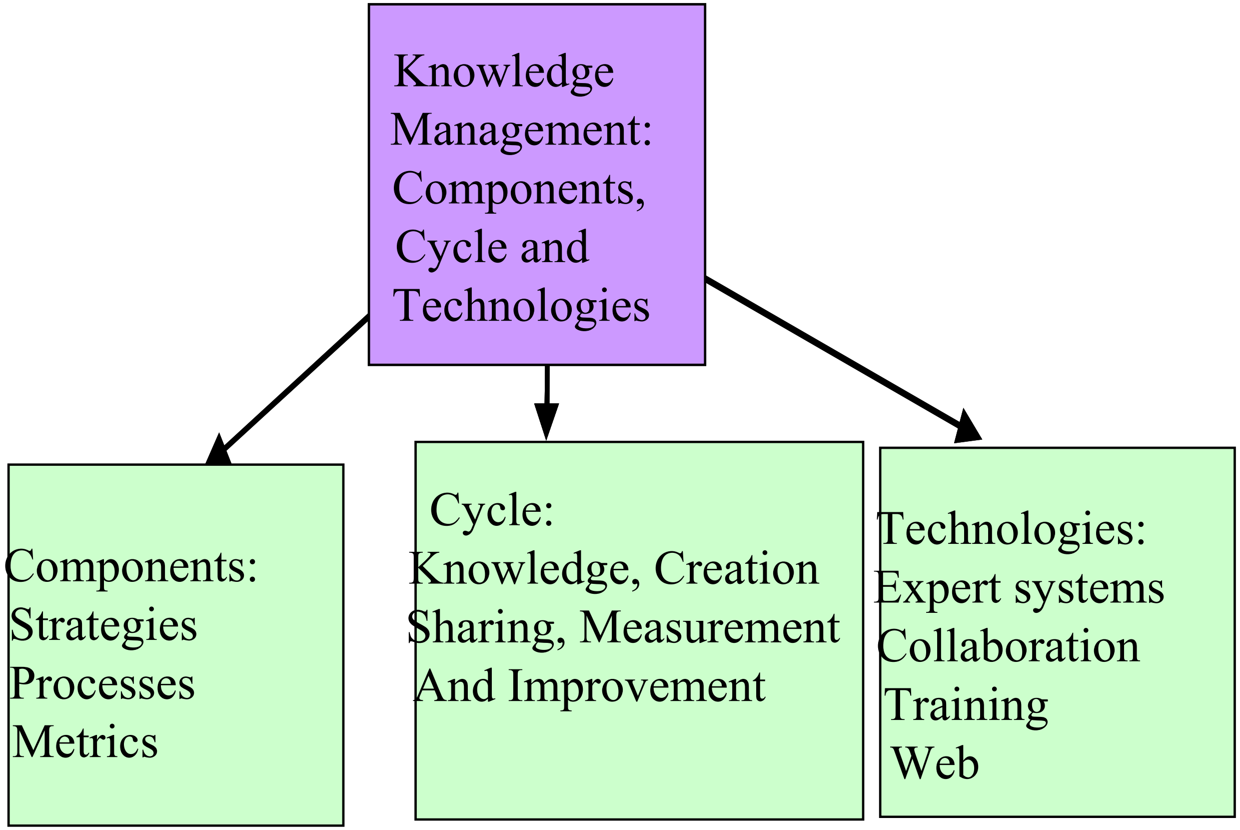

KNOWLEDGE MANAGEMENT CYCLE

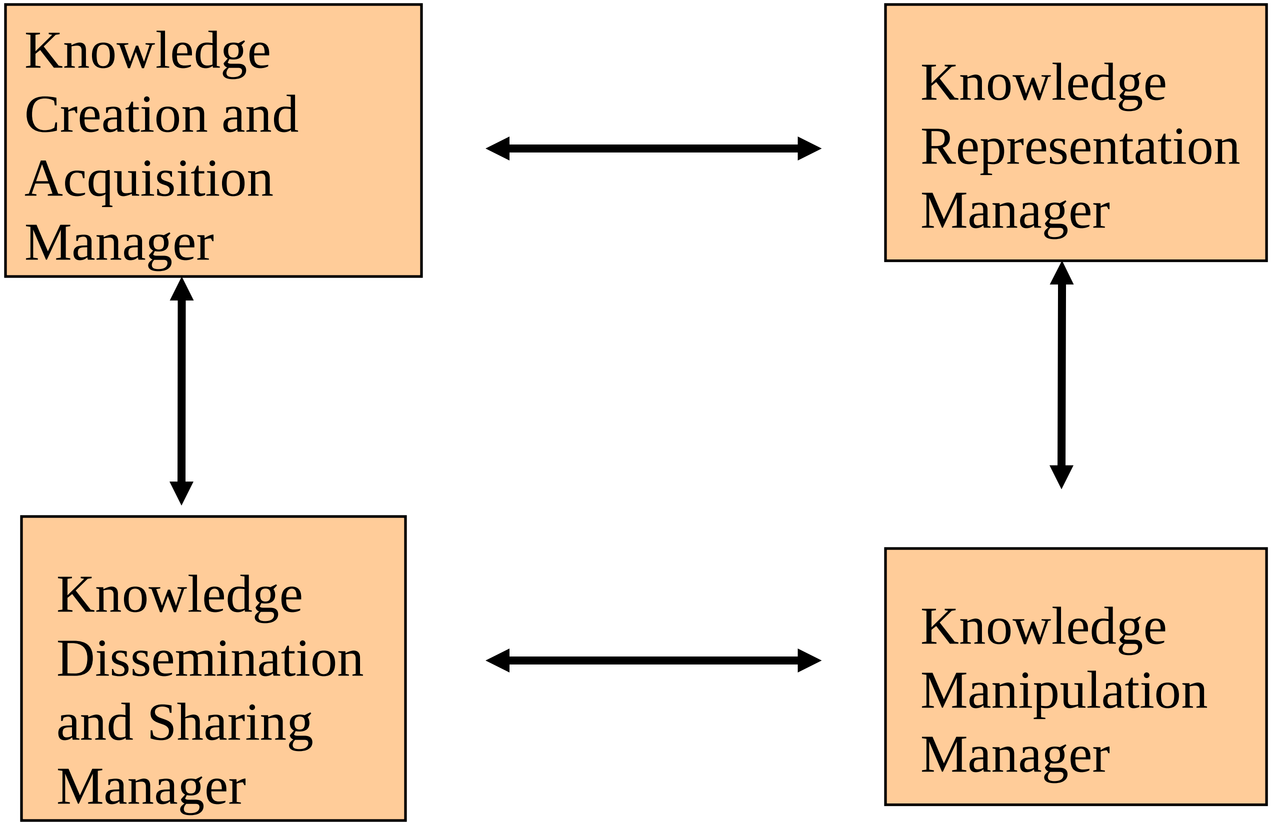

KNOWLEDGE MANAGEMENT ARCHITECTURE

- Knowledge management requires several components:
  - Access to both internal and external information sources,
  - Repositories that contain explicit knowledge,
  - Processes to acquire, refine, store, retrieve, disseminate and present knowledge,
  - Organizational incentives and management roles to support these activities,
  - People who facilitate, curate, and disseminate knowledge within the organization.
  - Information technology to provide automation support for many of the above activities,

KM ARCHITECTURE

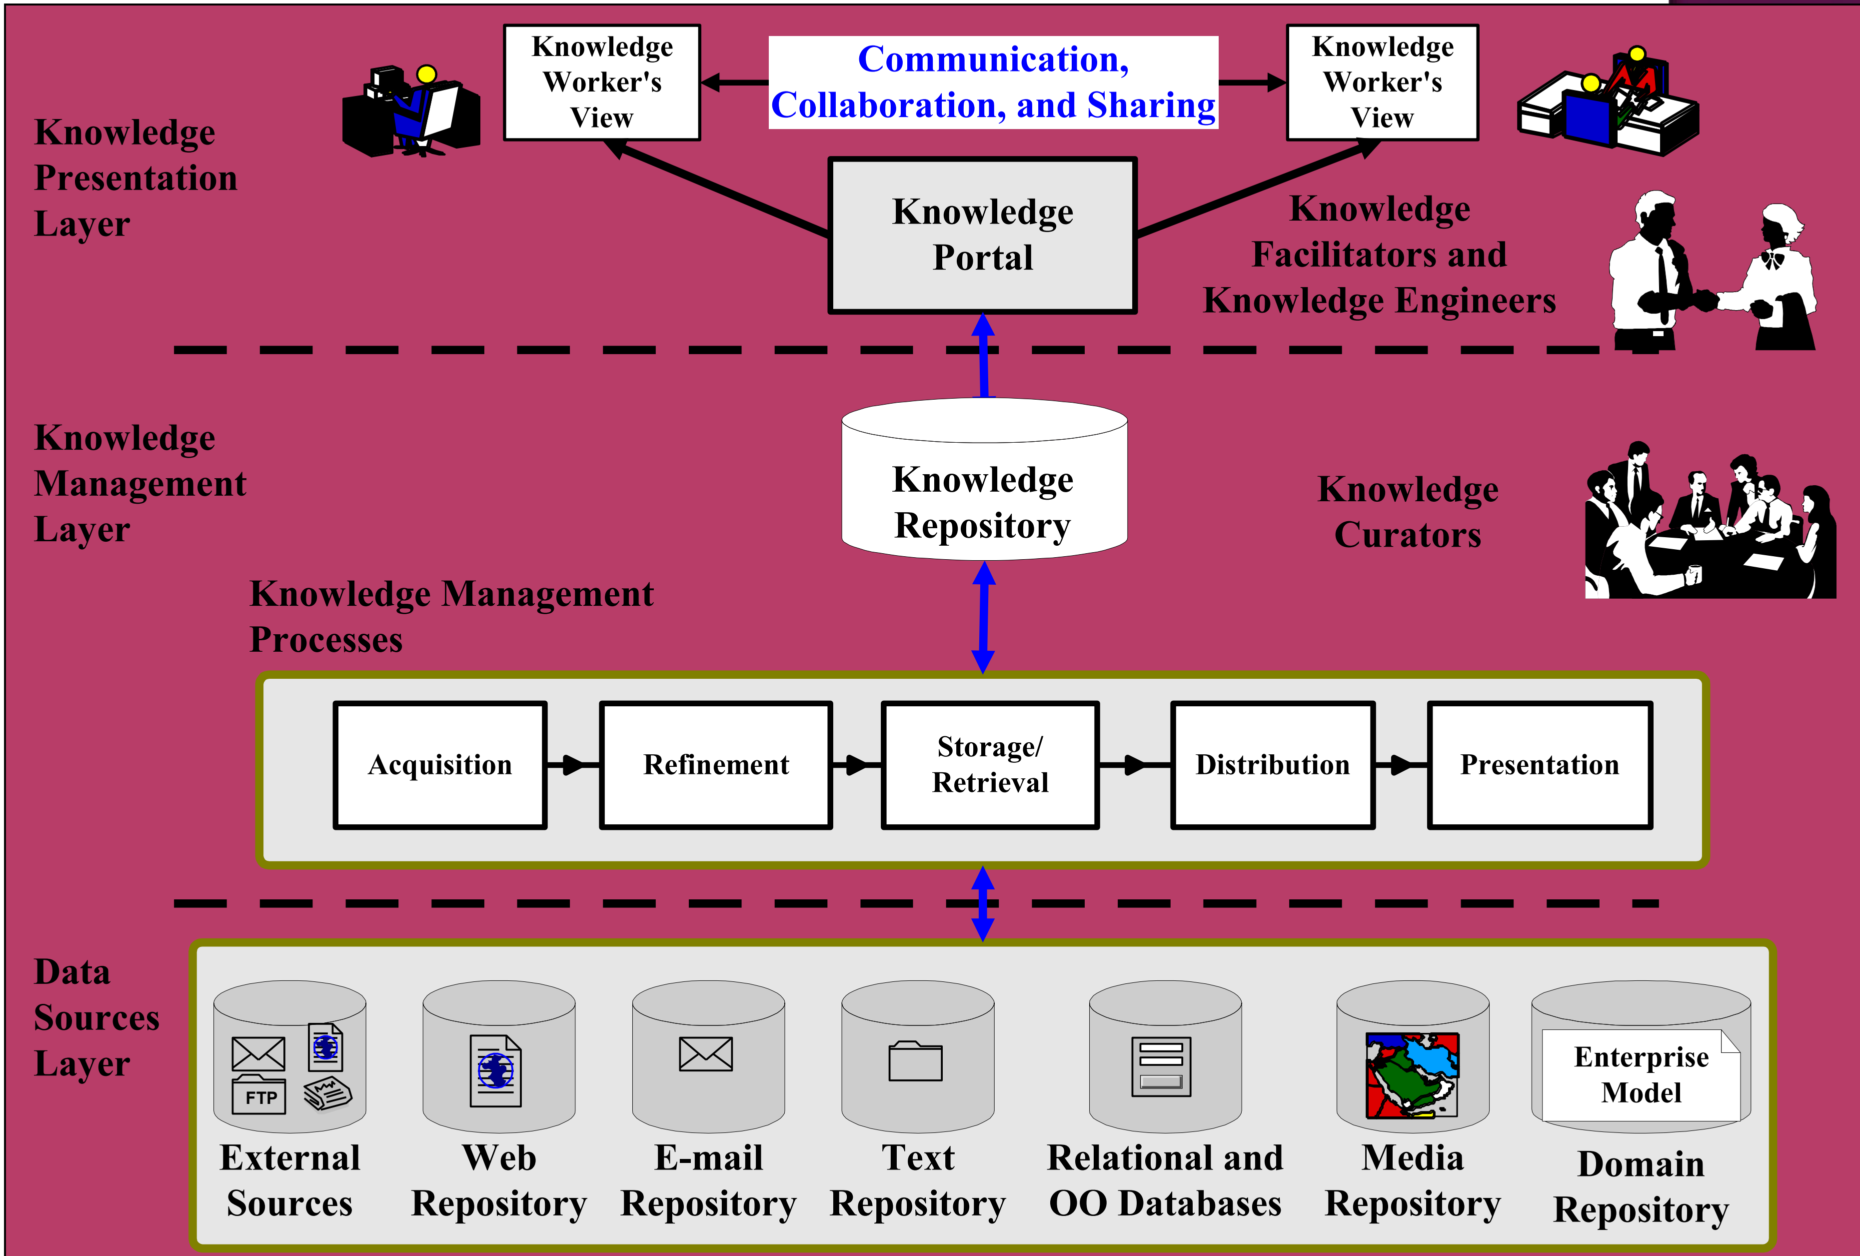

KNOWLEDGE MANAGEMENT PROCESS MODEL.

Acquisition

- Expertise
- Domain Model
- Business Rules
- Ownership; Federation Agreements, Data Sources
- External Sources and Formats.
- Wrappers
- Politics of data

Refinement

- Data Cleansing
- Indexing
- Metadata Tagging
- Concept Formulation
- Information Integration
- Ontology & Taxonomy
- Knowledge Curation.

Storage/ Retrieval

- Storage and indexing of Knowledge
- Concept-based Retrieval
- Retrieval by Author, Content, Threads, etc.
- Knowledge Security.

Distribution

- Intranet & Internet
- Knowledge Portals
- XML
- Active Subscriptions
- Discussion Groups.
- Digital Rights Management

Presentation

- User Profiles for dynamic tailoring links.
- Knowledge creation, update annotation, and storage in Knowledge Repository.
- Collaboration Environments

SKM: STRATEGIES, PROCESSES, METRICS, TECHNIQUES

- Security Strategies:
  - Policies and procedures for sharing data
  - Protecting intellectual property
  - Should be tightly integrated with business strategy
- Security processes
  - Secure workflow
  - Processes for contracting, purchasing, order management, etc.
- Metrics
  - What is impact of security on number of documents published and other metrics gathered
- Techniques
  - Access control, Trust management

WHAT IS KNOWLEDGE MANAGEMENT

- Gartner group: KM is a discipline that promotes an integrated approach to identifying and sharing all of an enterprise's information assets, including databases, documents, policies and procedures as well as unarticulated expertise and experience resident in individual workers
- Peter Sange: Knowledge is the capacity for effective action, this distinguishes knowledge from data and information; KM is just another term in the ongoing continuum of business management evolution

Some Excerpts from Knowledge Management

- Knowledge is Created by Human Beings
- Human needs and motivation lead us naturally to create knowledge
- Everybody is a knowledge worker
- People choose to share their knowledge
- Knowledge management is not about technology
- Knowledge is born in chaotic processes that take time

KM: Strategy, Process and Metrics

- Strategy
  - Motivation for KM and how to structure a KM program
- Process
  - Use of KM to make existing practice more effective
- Metrics
  - Measure the impact of KM on an organization

TRANSFORMING KNOWLEDGE

- Tacit to Tacit
  - E-meetings, Synchronous collaboration (chat)
- Explicit to Tacit
  - Visualization, Browsable video/audio of representations
- Tacit to Explicit
  - Answering questions, Annotations
- Explicit to Explicit
  - Text search, Document categorization

DATA MINING FOR KNOWLEDGE MANAGEMENT

- Data Mining is a key technology for knowledge management
- Mine the data to determine the competitor strategy to improve business; also to enhance one's own strategy
- Targeted marketing to customers to improve sales
- Determine strategies for employee retention and benefits
- In summary data mining is key to better business intelligence and business intelligence is key to effective knowledge management

KNOWLEDGE RESOURCE EXCHANGE

- The challenge is create value through alliances and collaborations
- The partner resource exchange model:
  - Two partner resource exchange in which the resources each partner contributes can be measured against the dimensions of tacitness, specificity, and complexity.
  - A variable is defined that reflects the degree to which a given partner contributes tacit, specific and complex knowledge resources to an alliance
- The concepts can be extended to include multiple partners and multiple dimensions