

INFORMATION WAREHOUSE IN THE INSURANCE INDUSTRY

Document Number GG24-4341-00

August 1994

International Technical Support Organization
San Jose

Take Note!

Before using this information and the product it supports, be sure to read the general information under "Special Notices" on page xiii.

First Edition (August 1994)

This edition applies to the following products:

- DataPropagator Relational Version 1 Release 1, Program Number 5622-244
- DataHub/2 Version 1 Release 1, Program Number 5667-134
- DataGuide/2 Version 1 Release 1, Program Numbers 5622-487 and 5622-488
- FlowMark Version 1 Release 2, Program Number 5621-290
- DataPropagator NonRelational Version 1 Release 2, Program Number 5696-705
- DataRefresher Version 3 Release 1, Program Number 5696-703
- Visualizer Query Version 1 Release 0, Program Number 5871-BBB
- S/390 Parallel Query Server

Order publications through your IBM representative or the IBM branch office serving your locality. Publications are not stocked at the address given below.

An ITSO Technical Bulletin Evaluation Form for readers' feedback appears facing Chapter 1. If the form has been removed, comments may be addressed to:

IBM Corporation, International Technical Support Organization
Dept. 471, Building 070B
5600 Cottle Road
San Jose, California 95193-0001

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright International Business Machines Corporation 1994. All rights reserved.**

Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Abstract

This publication is one of three publications that relate Information Warehouse architecture and products to industry applications and requirements. These three publications are:

- *Information Warehouse in the Finance Industry*
- *Information Warehouse in the Insurance Industry*
- *Information Warehouse in the Retail Industry.*

The publications describe the *Information Warehouse Architecture I* and emphasize the following products:

- DataPropagator Relational
- DataHub
- DataGuide
- FlowMark
- DataPropagator NonRelational
- DataRefresher
- Visualizer.

These products provide a variety of functions defined in *Information Warehouse Architecture I*.

This publication is intended for business analysts acting as consultants to an Information Warehouse implementation project and technical professionals who are designing Information Warehouse solutions in the Insurance industry. A knowledge of the Information Warehouse framework is assumed.

DS

(100 pages)

Contents

PART 1. INTRODUCTION	1
Chapter 1. Industry Library Introduction	3
1.1 Library at a Glance	3
1.2 Terminology	4
1.3 Introduction to Solution Threads	4
PART 2. THE BUSINESS VIEW	9
Chapter 2. Insurance Industry Perspective	11
2.1 Insurance Industry Trends	13
2.1.1 Marketing	13
2.1.2 Distribution	13
2.1.3 Customer Service	14
2.1.4 Organizational Vitality	14
2.1.5 Competitive Forces	15
2.2 Information Technology in the Insurance Industry	15
2.2.1 Information Technology Potential	16
2.3 Industry Challenges	16
2.4 Key Systems	17
2.5 Key Requirements	17
Chapter 3. Insurance Industry Business Requirements	19
3.1 Insurance Industry Examples	20
3.1.1 Product-to-Market	20
3.1.2 The Health Industry	20
3.2 Technology Requirements	21
3.2.1 Data Replication Requirement	21
3.2.2 Workflow Management Requirement	22
3.2.3 Workflow Management: Product-to-Market	23
3.3 Informational Application Requirement	24
PART 3. THE TECHNOLOGY VIEW	27
Chapter 4. Insurance Application Architecture	29
4.1 IAA Modeling Terminology and Concepts	31
4.1.1 Modeling	31
4.2 IAA Framework	32
4.2.1 Models	32
4.2.2 Dimensions	33
4.3 Data Model Commonality	34
4.4 IAA Data Model	36
Chapter 5. Information Warehouse Framework	41
5.1 Value of the Information Warehouse Framework	42

5.2 Why Data Replication	43
5.2.1 Operational Systems	43
5.2.2 Database Technology	44
5.2.3 Cost of Data Access	44
5.2.4 Historical Data	45
5.2.5 Ownership	45
5.2.6 Point-in-Time Data	45
5.2.7 Reconciliation	45
5.3 The Information Warehouse Architecture	45
5.4 Using the Information Warehouse Architecture	47
5.5 Access Enablers	48
5.5.1 Embedded SQL	50
5.5.2 SQL Call Level Interface	50
5.5.3 Distributed Relational Database Architecture	51
5.6 The Insurance Industry	51
5.7 Workflow Management Requirement	51
5.7.1 Workflow Management Function	52
5.7.2 Workflow Management Interface	52
5.8 Decision Support: Informational Applications	52
5.8.1 Informational Application Function	53
5.8.2 Informational Application Interfaces	53
Chapter 6. Insurance Solution Thread Overview	55
6.1 The Meta-data Requirement	57
6.2 Reconciliation	59
6.3 Solution Components	59
Chapter 7. Organization Asset Data	61
7.1 Business View	63
7.2 Mapping the Organization Asset Data	63
7.3 Managing the Organization Asset Data	64
Chapter 8. Data Replication Tools	67
8.1 Tool Selection	68
8.2 DataRefresher	69
8.2.1 Operating Environment	69
8.3 Data Sources and Targets	70
8.3.1 Data Sources	71
8.3.2 Target Systems	72
Chapter 9. Workflow Management	73
9.1 Workflow Management Technology	74
9.1.1 Workflow Management Users	74
9.1.2 Terminology	74
9.1.3 Information Warehouse Framework Integration	76
9.1.4 Workflow Management Coalition	76
9.2 FlowMark	77
9.2.1 Definition and Documentation of Processes	77
9.2.2 Testing Processes	78
9.2.3 Running Processes	78
9.3 Insurance Solution Thread	78
Chapter 10. Applications	81
10.1 Knowledge Worker Requirements	82
10.2 End-User Perspective	83
10.3 Informational Application Function	83
10.4 Client-Server Operation	84
10.5 Visualizer	85
10.5.1 Visualizer Objects	86

10.5.2 Visualizer Templates	87
10.6 Information Warehouse Integration	87
Chapter 11. Conclusions	89
Appendix A. Models and Modeling	91
A.1 The Construction Model	91
A.1.1 Entity: Things	92
A.1.2 Entity: Agreements	92
A.2 The Annual Report As a Model	93
A.3 Information Warehouse and Modeling	93
List of Abbreviations	95
Index	97

Figures

1. Basic Set of Business Objects	6
2. IAA Framework: Simplified	34
3. Model Extents	37
4. Information Warehouse Architecture	46
5. Information Warehouse Framework	48
6. Access Enablers	49
7. Business Concept, Modeling Concept, and Operational Record	56
8. Data→Information→Chart	57
9. Organization Asset Data	62
10. Copy Tools	67
11. DataRefresher Sources and Targets	71
12. Workflow Management	73
13. Data→Information	79
14. Applications	81

Tables

1. Library of Information Warehouse: Products Covered	4
2. Meta-data Example: Policy Number in Claim Record	58
3. Meta-data Example: Policy Number in Premium Segment	58
4. Mapping Terms for Data	64

Special Notices

This publication is intended to help:

- Business analysts understand Information Warehouse (IW) architecture concepts
- IBM technical professionals understand industry environments
- Customer data processing personnel understand industry environments.

The information in this publication is not intended as the specification of any programming interfaces that are provided by a variety of products that perform functions described in the Information Warehouse architecture. See the PUBLICATIONS section of the IBM Programming Announcement for these products for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY 10594 USA.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM (VENDOR) products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

The following terms, which are denoted by an asterisk (*) in this publication, are trademarks of the International Business Machines Corporation in the United States and/or other countries:

BookManager	CICS/ESA
Common User Access	CUA
DataGuide	DataHub
DB2	DB2/2
DB2/6000	IBM
ImagePlus	IMS/ESA
MVS/ESA	PS/2
QMF	RISC System/6000

The following terms, which are denoted by a double asterisk (**) in this publication, are trademarks of other companies:

Apple	Apple Computer, Inc.
BRIDGE/FASTLOAD	Bridge Technology, Inc.
KnowledgeWare Application Development Workbench	KnowledgeWare, Inc.
MacIntosh	Apple Computer Company
Microsoft, Windows	Microsoft Corporation
Motif	Open Software Foundation
OSF	Open Software Foundation, Inc.
ObjectStore Database	Object Design, Inc.
1-2-3, Lotus, Freelance, Freelance Graphics	Lotus Development Corporation.
OMegamon	Candle Corporation
Fast Load	PLATINUM Technology Inc.
Fast Unload	PLATINUM Technology Inc.
Rapid Reorg	PLATINUM Technology Inc.
Quick Copy	PLATINUM Technology Inc.
In2itive	LEGENT

Other trademarks are trademarks of their respective companies.

Preface

This document is intended to merge industry analysis, industry architecture, the Information Warehouse architecture, new product discussion, and specific solutions to industry requirements. It contains discussion of specific industry issues, industry architecture for data processing, Information Warehouse architecture, and solutions.

This document is intended for business analysts and data processing professionals.

How This Document Is Organized

The document is organized as follows:

- Introduction

This part introduces the library within which this particular book is included.

- The Business View

This part establishes the business-oriented level set from which business requirements and Information Warehouse solutions are developed. The first chapter presents a perspective on the Insurance industry that includes trends and challenges, key systems, and information technology's position in the Insurance industry.

- The Technology View

This part presents the technology solutions for the business requirements established in the business view. It includes an overview of the industry application architecture and the Information Warehouse architecture. It then discusses the individual components of the Information Warehouse architecture and the solutions to those components, according to the needs of the industry.

Related Publications

The following publications are considered particularly suitable for a more detailed discussion of the topics covered in this document:

- “An Architecture for a Business and Information System,” B.A. Devlin and P.T. Murphy, *IBM Systems Journal*, Vol. 27, No. 1 (1988)
- “Building Business and Application Systems with the Retail Application Architecture,” P. Stecher, *IBM Systems Journal*, Vol. 32, No. 2 (1993)
- *Client-Server Computing: The Design and Coding of a Business Application*, GG24-3899-00.
- *DataGuide/2 V1: Using DataGuide/2*, SC26-3365
- *Delivering Data to the Information Warehouse*, Rob Goldring, InfoDB Summer 1992
- *Financial Application Architecture: FAA Concepts of Application and System Architectures*, LY38-4402-0
- *Information Technology and the Management Difference: A Fusion Map*, *IBM Systems Journal*, Vol. 32, No 1 (1993)
- *Financial Application Architecture Introduction*, GC31-3932-0
- “The Future of Health Care Information Systems,” Michael Carrigan, *Hospital Materiel Management Quarterly*, August 1993
- *Information Warehouse Architecture I*, SC26-3244
- *Insurance Industry Futures: Directions for the 21st Century*, Anderson Consulting and LOMA 1993
- “Loaning Banks Some Courage,” *Information Week*, August 12, 1993
- “The Model Business,” *IW Today*, August 12, 1993
- *Principles of Life and Health Insurance*, G. Morton, 1988 LOMA
- “The Spectrum of Data Delivery for Business Information Systems,” Rob Goldring, *DB/Expo92*

International Technical Support Organization Publications

- *Information Warehouse Architecture and Info. Catalog*, GG24-4019
- *Information Warehouse Storage Management Guidelines and Considerations*, GG24-4336
- *Information Warehouse in the Finance Industry*, GG24-4340
- *Information Warehouse in the Retail Industry*, GG24-4342
- *Library for System Solutions: Data Reference*, GG24-4103

A complete list of International Technical Support Organization publications, with a brief description of each, may be found in:

Bibliography of International Technical Support Organization Technical Bulletins, GG24-3070.

To get a catalog of ITSO technical publications (known as “redbooks”) online, VNET users may type:

```
TOOLS SENDTO WTSCPOK TOOLS REDBOOKS GET REDBOOKS CATALOG
```

How to Order ITSO Technical Publications

IBM employees in the USA may order ITSO books and CD-ROMs using PUBORDER. Customers in the USA may order by calling 1-800-879-2755 or by faxing 1-800-284-4721. Visa and Master Cards are accepted. Outside the USA, customers should contact their local IBM office.

Customers may order hardcopy ITSO books individually or in customized sets, called GBOFs, which relate to specific functions of interest. IBM employees and customers may also order ITSO books in online format on CD-ROM collections, which contain books on a variety of products.

Acknowledgments

The advisor for this project was:

Steve Schaffer
International Technical Support Organization, San Jose

The authors of this document were:

Normand Brin
IBM Canada

Wojeich Zagala
IBM Australia

Steve Schaffer
International Technical Support Organization, San Jose

This publication is the result of a residency conducted at the International Technical Support Organization, San Jose.

Thanks to the following people for the invaluable advice and guidance provided in the production of this document:

Paul Englefield, IBM Warwick
Rob Goldring, IBM SWS, Santa Teresa
Eileen Hiltbrand, IBM US
Jacques Labrie, IBM SWS, Santa Teresa
Bill Martin, IBM US
Mark Mauriello, IBM Charlotte Finance Industry
Rita Neuberg, IBM Charlotte
Bill Payne, IBM Charlotte Insurance Industry

Thanks to the following people for reviewing this document:

Thomas Bilfinger, IBM ITSO, San Jose
Don Cameron, IBM ITSO, San Jose
Don Murray, IBM US
Ralph Naegeli, IBM Switzerland
Tom Romeo, IBM US
Michele Schwartz, IBM SWS, Santa Teresa

Special thanks to Ueli Wahli for developing the tool to generate margin comments.

Part 1. Introduction

Chapter 1. Industry Library Introduction

This volume is one of three that look at Information Warehouse architecture and products in the finance, insurance, and retail industries. The three studies yielded somewhat different results but are presented in a standard structure. This introductory chapter describes the library, so that it may be used easily and most effectively by the individual.

1.1 Library at a Glance

The study of the finance, insurance, and retail industries has been made available as a library of three books. Each book takes the same approach to discussing the respective industry, though there is some variation in the aspects of the Information Warehouse architecture covered in each industry. Table 1 on page 4 gives an overview perspective of this library. The library's structure is based on a common set of topics that are addressed consistently across the industry studies, and different aspects of the Information Warehouse architecture being addressed in each study. This structure minimizes redundancy. Therefore, a complete review of Information Warehouse architecture function and products may require reading of all three industry studies.

The common set of topics presents the study flow from a high-level perspective of the industry down to the discussion of the Information Warehouse product technology. These topics are broken down into the business view and technical view as follows:

- Business view
 - Industry perspective
 - Business requirements
- Technology view
 - Industry application architecture
 - Information Warehouse architecture
 - Information Warehouse framework components.

Table 1. Library of Information Warehouse: Products Covered	
Industry	Product
Finance	DataPropagator Relational
Insurance	FlowMark Visualizer
Retail	DataGuide* S/390 Parallel Query Server

1.2 Terminology

The finance, insurance, and retail industry studies address general requirements of the respective industry. They are not studies of a real-world or a contrived enterprise. Rather, they are studies of industry issues and requirements put in the context of an industry enterprise. For this reason, we use the term financial enterprise, insurance enterprise, and retail enterprise, respectively, to reflect this approach.

The knowledge worker is the focus

We begin our study by identifying the knowledge worker as the primary focus of Information Warehouse technology. Knowledge workers are the individuals in an enterprise who make decisions. They exist at every organizational level and have one thing in common: they all need information to make decisions. They get the information through informational applications, also called executive information systems, decision support software, and decision support tools. Informational applications are used to display information provided by the data replication products in the Information Warehouse solution. Therefore, the objective of the studies is to describe knowledge worker's business need for information and the addressing of that need through the Information Warehouse architecture and products.

1.3 Introduction to Solution Threads

Solution thread: applying technology to a problem

A solution thread is a vehicle for applying architectures and products to a generic requirement of the industry. It is a generalized approach, in that the Information Warehouse architecture it uses is generalized for a given data processing function (for example, decision support) across industries, and the industry architecture it uses is generalized for enterprises within a given industry. The products used by the solution thread are generalized for that data processing function, rather than for a business requirement or hardware platform. The solution thread meets the business requirement by customizing the product to the needs of the business.

The Information Warehouse architecture is a generalized architectural approach to managing data and information in a complex business environment. The industry architectures—financial, insurance, and retail—discussed in the three books of this library represent structured approaches to analyzing specific business environments. This analysis leads to specification of requirements, expressed in business terms, which data processing technologies must address. This library presents the Information Warehouse architecture as an architected approach to data processing functions required across the industries, and across the lines of business within each industry. The value of using both the industry-specific architectures and the generalized data processing architecture (Information Warehouse architecture) is to leverage the Information Warehouse architecture functions across the lines of business and to leverage the resources already invested in the industry-specific architectures.

The Information Warehouse architecture is a generalized approach

Not all features of the Information Warehouse architecture nor every product is considered in each of the three industry studies chosen for this library. Review of the three studies as a set will provide information on most of the Information Warehouse architecture features and Information Warehouse framework products. We discuss IBM's published Information Warehouse architecture as a template for connecting business requirements to actual technical implementations, including products plugged into an Information Warehouse framework.

The following example taken from the insurance industry illustrates the leverage of the Insurance Application Architecture (IAA) together with the Information Warehouse architecture. Figure 1 on page 6 shows a set of business objects for the insurance industry. These business objects are used as examples of modeling entities—OBJECT, AGREEMENT, and DEMAND_FOR_DELIVERY—found in the IAA model. IAA defines a modeling entity called an OBJECT. This entity is used to symbolically represent anything that can be insured, such as a life. IAA defines another modeling entity called Agreement. We could use this modeling entity to represent the physical insurance entity called Policy in either personal or property and casualty insurance. This is an example of the generality of the IAA being leveraged across lines of business. We could further use another entity, called DEMAND_FOR_DELIVERY, to represent Claims and PREMIUMS for Premiums.

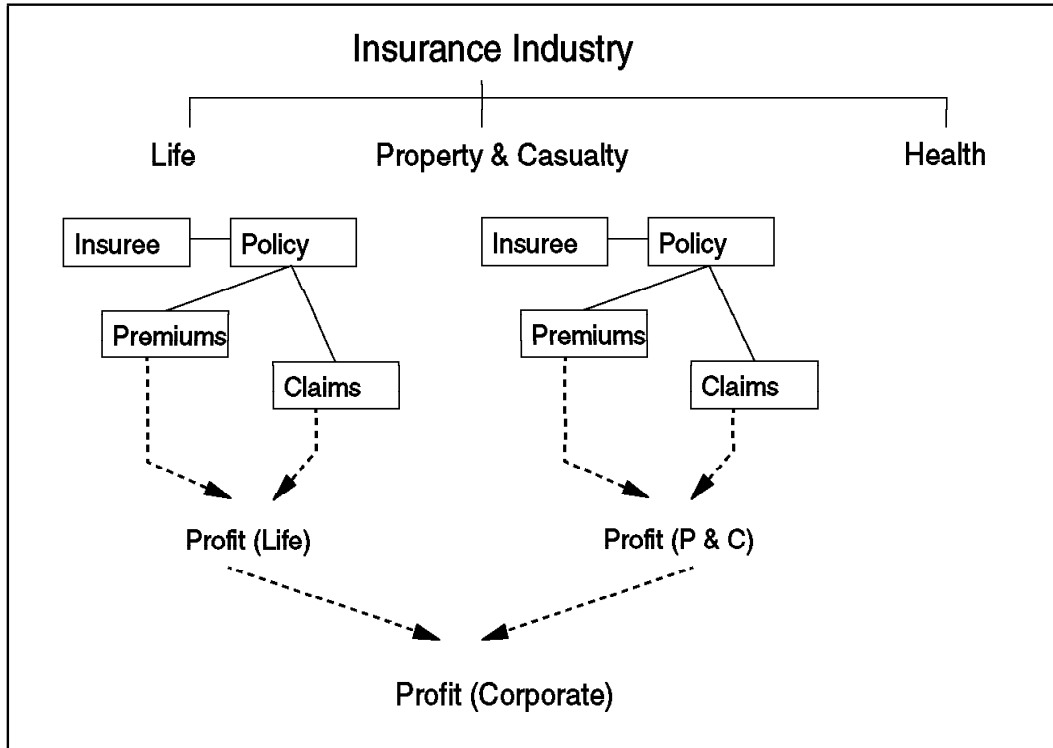


Figure 1. Basic Set of Business Objects. Dashed arrows indicate data flow.

These basic entities correspond to the real-world objects. The Claims entity corresponds to the many claims made by insurees. The Premiums entity corresponds to the many premium payments made by insurees. These claims and payments are implemented as records in an operational database, say, IMS/DB, DB2*, or VSAM. The prime concern of the insurance company is profitability. In this simple exercise, profitability is defined as total premiums minus total claims.

Assuming claims records and premiums records are kept in separate databases, there is a need to bring those two sets of data together. There is an additional need to reconcile this data as it is brought together. The Information Warehouse architecture defines the process by which this can be done in an architected manner. This architected approach to extracting data is called data replication. Data replication is generalized, so that it can be a solution for bringing together Claims and Premiums data for Life insurance, Property and Casualty, or any other insurance product that takes in premiums and pays out claims.

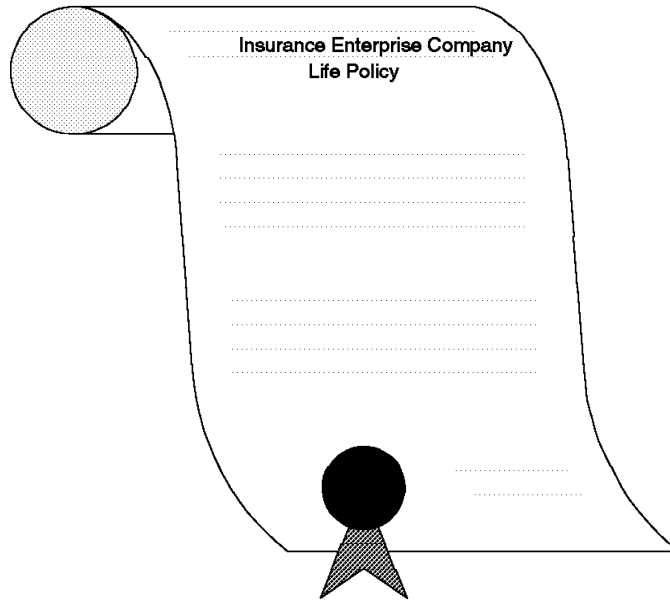
The Information Warehouse architecture provides guidance for data access as well as data replication functions. Data access applies to operational data or informational data copies of the operational data, generated through data replication. Access to operational data (*direct access*) or informational data shares certain requirements for implementation. These requirements can be best understood in terms of the business requirements for data access. The lines of business are assumed to have their own data stores containing records of insurees, sometimes called client files. The Life insurance line of business has an interest in using the client file owned by the property and casualty line of business for prospecting purposes, and vice versa. This requirement brings up two issues relevant to the Information Warehouse

framework: what data is available, and how to access it. The first requirement is resolved through the Information Catalog, while the second is resolved through access enablers or data replication. The point here is that both lines of business have the same needs to access data, and both needs can be resolved by solutions based on the Information Warehouse architecture.

Part 2. The Business View

Chapter 2. Insurance Industry Perspective	11
2.1 Insurance Industry Trends	13
2.1.1 Marketing	13
2.1.2 Distribution	13
2.1.3 Customer Service	14
2.1.4 Organizational Vitality	14
2.1.5 Competitive Forces	15
2.2 Information Technology in the Insurance Industry	15
2.2.1 Information Technology Potential	16
2.3 Industry Challenges	16
2.4 Key Systems	17
2.5 Key Requirements	17
Chapter 3. Insurance Industry Business Requirements	19
3.1 Insurance Industry Examples	20
3.1.1 Product-to-Market	20
3.1.2 The Health Industry	20
3.2 Technology Requirements	21
3.2.1 Data Replication Requirement	21
3.2.2 Workflow Management Requirement	22
3.2.2.1 Defining Processes	22
3.2.2.2 Complexity	23
3.2.2.3 Tracking Processes	23
3.2.2.4 Reusing Processes	23
3.2.2.5 Dynamics of the Industry	23
3.2.3 Workflow Management: Product-to-Market	23
3.3 Informational Application Requirement	24

Chapter 2. Insurance Industry Perspective



In terms of its age, the modern insurance business is an infant when compared to many other industries. In terms of its size, however, the industry is among the largest. Life insurance products were not widely offered until the 1800s, and health insurance products were not available generally until the early part of this century. Yet, through time, the amount of life insurance in force in the United States and Canada has grown to over \$5 trillion, and health insurance products cover most of the people in both countries.

Insurance is a young industry

The primary reason for this tremendous growth lies in the nature and purpose of insurance products. All insurance provides protection against some of the economic consequences of loss. Thus, insurance responds to the need of all persons for security. The insurance industry designs, revises, alters, and updates its insurance policies constantly to meet this need. However, despite these changes, the underlying purpose of these policies remains the same: providing economic protection against financial loss.

Insurance protects against loss

A policy is a contract

Essentially, an insurance policy is a contract—a legally enforceable agreement—under which the insurance company agrees to pay a certain amount of money, called the policy benefit, when specific losses occur, provided the insurer receives a specified amount of money, called the premium. In this way, the risk, or chance, of economic loss is transferred to the insurance company.

The insurance industry is generally divided into three subindustries: life, health, and property and casualty insurance. Life insurance provides a specified sum of money if the person who is insured dies while the policy is in effect. Health insurance pays specified benefits if the person who is insured becomes sick or has an accident. Property insurance provides a benefit should covered property be damaged or lost because of fire, theft, or other cause described in the policy. Liability insurance provides a benefit payable on behalf of a covered party who is held legally responsible (liable) for harming others or their property.¹

IW technology has much to offer the insurance industry

The interest of this study is the data processing systems that mirror and support the activities inherent in carrying on the insurance business. The study looks at a specific subset of the data processing activities, namely, those addressed by Information Warehouse technology. The insurance business, as an industry, is among the most advanced in using data processing technology to manage the day-to-day business, such as policy writing, premium collections, and claims processing. As an industry, however, it is far less advanced in the use of the resulting data for trend analysis and decision support. This represents a tremendous opportunity to apply Information Warehouse concepts and technology for the benefit of the industry.

This book takes a top-down approach to understanding the insurance industry and Information Warehouse technology's role in the industry. We start by reviewing the trends of the insurance business, because the trends indicate the future and separate more successful enterprises from the less successful. The trends indicate the competitive pressures and the business requirements that need to be met to survive and to excel. The data processing technology that is useful in meeting those pressures and requirements is a technology that is worthy of investment. We present the Information Warehouse framework as a data processing technology that does just that—it can be used to meet competitive and business pressures of the future.

We begin with an industry view of the business climate. This is followed by a review of information technology's current and potential role in the insurance industry. We identify key information systems of the insurance industry. Key systems take two basic forms: existing key systems that run the enterprise today and key systems that will differentiate the enterprise in the future. The overriding observation is that the existing key systems that run the enterprise today are information feeders to the differentiating key systems of tomorrow. The aim of this publication is to show how that feeding is done reliably and economically.

¹ Adapted with permission from *Principles of Life and Health Insurance*, 2nd edition, LOMA © 1988. All rights reserved.

2.1 Insurance Industry Trends

Five major focus areas emerge from analysis of the survey data and associated executive insight: marketing, distribution, customer service, organizational vitality, and competitive forces.

2.1.1 Marketing

The insurance market is becoming increasingly crowded and is being attacked by nontraditional entrants with different rules of engagement. It is no longer adequate to flood a broad section of the market with product and hope that the channel can sell it, particularly as the market becomes more and more fragmented.

The right product to the right customer at the right time

Insurance enterprises need to carefully analyze the marketplace, identify explicit segments, select target markets, and then develop adaptable products that address the client's changing needs. Mechanisms must be developed that can track ongoing profitability, scrutinize competitive offerings, and define new opportunities. Distribution channel selection and service levels will need to be based on the amount of "value-added" the market segment will support. More robust and flexible analysis tools will be required to ensure that products are flexible and meet changing client needs over time.

2.1.2 Distribution

Greater integration of the sales process and home office administration will evolve as insurers seek greater efficiency and effectiveness of the distribution channel.

Bring the information and tools together for the agents

The need for greater differentiation will place new demands on insurers to add real measurable value to the marketing and customer service processes. New technology will enable existing distribution systems to improve both productivity and service quality. New marketing approaches will also be used—primarily to reach homogeneous subgroups such as professional women, ethnic groups, and seniors. Multiproduct pricing and product packaging strategies will drive the need for client-based information systems. Technology will allow home office client-oriented information and sales support systems to be better coordinated—both for marketing and customer service.

2.1.3 Customer Service

Fully integrate customer service departments

In the near future, customer service will need to be far more explicitly defined—in ways that will create sustainable competitive advantage and in terms relevant to the customer, not the enterprise. The ultimate service objective will be accomplished through a client view of the business.

Increase the overall access to information

Alternative access mechanisms will provide a significant increase in the availability of information to interested parties. The company, the client, the agent and broker, and the insured will be able to initiate customer service transactions. The result of these transactions will then automatically be reflected in the records of other involved parties with no overt action.

Through the application of technology, new dimensions of customer service will be realized, enabling the industry to make powerful image improvements.

2.1.4 Organizational Vitality

Provide service linked to marketing opportunities

Insurers will continue to invest in their people and encourage greater empowerment. Service team concepts will help organizations become more flexible to adapt to a rapidly changing environment. Employees will need to be trained more broadly so that they can perform a wider variety of functions. Companies will place greater focus on core competencies. Cross function processes will be further reengineered and integrated in order to improve business performance.

Combine administrative systems with client management systems

This focus on customer service must be driven by overall corporate strategies and supported with more integrated information systems. Organizational changes centered around the concepts of self-managed teams and employee empowerment require that the boundaries between current line-of-business systems come down. Strategies to overhaul legacy systems will be pursued, and future product development will benefit from some form of standardized “business architecture” so that offerings can be developed, modified, and serviced more easily. This kind of broadening will require a great deal more knowledge to be available throughout the organization—probably delivered by more sophisticated “just-in-time learning” tools. Continuous quality improvement will be an imperative. The availability of an organization to create a vital, exciting workplace will depend, at least in part, on its technology foundation. It is this foundation that will enable empowered teams to execute a wide range of sales and service functions.

2.1.5 Competitive Forces

In today's insurance world, both the rate of change and the intensity of competition are raising the bar in the area of service delivery. Insurers will look inside and outside the insurance industry for customer benchmarks. Never before have we seen change occurring at such an accelerated pace. The dynamics in the health segment of the business serve as a dramatic example.

All service providers are competitors

Technology will play a key role in an insurance enterprise's ability to compete profitably. Business process reengineering coupled with the strategic deployment of technology to develop and deliver innovative products and services will be the major determinants of future success in the insurance industry of tomorrow.

Direct marketing through technology is the standard

2.2 Information Technology in the Insurance Industry

Today, information technology has moved into a strategic role in insurance enterprises, providing critical support for achieving corporate goals and objectives. Used proactively and innovatively and aligned with corporate business objectives, it becomes one element in a strategy to gain and sustain competitive advantage.

Insurance enterprises are not achieving the full potential of information technology. There is a gap between the exploding potential of technology and the ability of individual insurance enterprises to achieve that potential; and this gap is growing. For most enterprises, narrowing this gap is a major challenge. Some reasons for not achieving the full potential include the following:

- Business and data processing goals have not been aligned.
Enterprises have tried to achieve information technology benefits but results have been less than desirable because business and information technology goals were not clearly aligned.
- Benefits have not been transferred to the bottom line.
Realizing the benefits does not automatically come from automating time-honored workflows. Rethinking how the business operates is required. There is a need to change roles and responsibilities across the organization and focus more consistently on what is important to achieve business goals.
- A range of technology was not considered.
Internal barriers have been built up over time. These include sunk investments in existing technology and systems, prohibitive reinvestment costs, and risks inherent in leading-edge positions. And the largest barrier is an enterprise's natural aversion to the gut-wrenching organizational change induced by technology innovation. An organization truly comfortable with change can take advantage of opportunities for compet-

itive advantage while its competitors continue to evaluate and analyze risks.

- The enterprise does not have the internal capability to achieve information technology benefits.

2.2.1 Information Technology Potential

Today, information technology can be used to achieve high-impact, visible shifts in how an insurance enterprise does business. For example, information technology, when used effectively, can:

- Drive business strategy
- Change industry dynamics
- Spawn new businesses
- Transform organizational roles and functions
- Alter competitive environments.

2.3 Industry Challenges

Over the past decade, we have seen enormous changes in underlying information technology and tools, with new complexity and sophistication in systems. Installed systems often act as barriers to the effective use of new information technology tools. At the same time, new requirements for cross-organizational consistency, integration, and sharing seem to create endless demands for new systems and enhancements to existing systems. These new systems and enhancements must interrelate with other automated systems and compete for constrained resources.

The volume and nature of these changes mean that planning is critical. Without it, scarce resources are at risk and business opportunities can be missed. For today, the insurance enterprise needs the following:

- A strategic focus for information technology planning and investment
- A cross-organizational view of information and systems
- A flexible architecture to define the business requirements and support future information technology initiatives.

Moving from today's realities to tomorrow's focus is hampered by a generation of planning tools that are a half step behind the evolving need. A new approach is needed that can represent your whole organization. The new approach needs to include the following:

- A method to focus information technology plans and investments more directly on strategic goals and long-term competitive advantage
- A process that results in a cross-organizational view including data consistency and sharing of information and the systems that supply it
- A stable, yet flexible architecture that provides a firm foundation to accommodate emerging technologies, changing business needs, and future generations of information systems

- A migration plan that includes projects and programs that capitalize on opportunities identified to use information technology for competitive advantage and move the organization toward the targets defined in the architecture.

2.4 Key Systems

The key business systems of the insurance industry emphasize the importance of the customer to the insurance enterprise. These systems are:

- Client administration
- Claim administration
- Product development
- Distribution management.

These key systems represent significant investments by the insurance enterprise and crucial assets of the insurance enterprise. They represent the data processing support of the customer, the product, and the selling of the product to the customer. The essence of an Information Warehouse implementation in the insurance industry is taking the primal data—claims data for claim administration and the client file data for client administration—and making it available for informational analysis. An emerging concept—the service office—merges client administration and claim administration. The strategic objective is to provide the best customer service in the industry.

2.5 Key Requirements

The key requirements in the insurance industry are essentially responses to economic and competitive pressures. The common thread among the solutions to these requirements is the need to access data already existing in the data processing environment of the insurance enterprise. The key is gaining access to data, and the facilitator of that access is the Information Warehouse framework. The key requirements identified are as follows:

- Product management
- Customer service
- Financial standing.

The economic and other environmental pressures are causing a shakeout of the insurance industry. The industry will see fewer and more specialized enterprises survive. The insurance enterprise will have to choose profitable products from nonprofitable ones and focus on the profitable ones to the limit of regulatory constraints. To do that, the insurance enterprise must have access to the data that reveals profitability. The enterprise compares the profitability of its products and focuses on those that have the greatest profitability. The Information Warehouse architecture is a structure for getting access to that data.

IW facilitates product management

One differentiator between insurance enterprises is their customer service. Customer service gives the enterprise representative access to more information about the customer for the benefit of the customer. The representative can give the customer a quicker and more thorough understanding of

the status of policies or other financial instruments held by the insurance enterprise on behalf of the customer. The information is buried in the operational systems, and those insurance enterprises that provide access to that information, using the Information Warehouse architecture, survive and excel.

Financial stability is a major selling point

An insurance enterprise is, after all, a financial institution. The profit of that enterprise is as much wrapped up in the investment side of the insurance enterprise as it is in the premiums of the insurance side. In effect, the financial status of an insurance enterprise becomes a study of a finance enterprise. The insurance industry has seen its share of shaky real estate investments. The insurance enterprise that can manage the data that tells the story of questionable investments and guides the enterprise toward the shedding of these investments will survive. The Information Warehouse architecture provides a structured approach to accessing and analyzing this data.

Chapter 3. Insurance Industry Business Requirements

Business requirements lead to needs for data processing solutions to the challenges in the insurance industry. These challenges derive from the trends and directions and pressures of the industry. We assume that these pressures persist over a relatively long period of time and that they are general in nature, rather than specific to an immediate narrow-focused need. Applications or application systems have been the usual response to specific, line-of-business, or project level requirements. These requirements have always been viewed as relevant to a specific project, product, or line of business and have been met by an independently developed application.

An architecture helps solve strategic problems

The solution to a long-term, strategic pressure must be executed in an architected context. This context assumes that a series of applications will be needed over the course of the business pressure's existence. The architecture presents a standard approach for all applications developed to meet different aspects of the business pressures. The relevant architectures are the Insurance Application Architecture (IAA) to accommodate specifics of the insurance industry and the Information Warehouse architecture to accommodate the cross-industry need to manage informational data and applications. These architectures are used as a guide for developing or acquiring software solutions to the business pressures.

Business pressures are strategic, as are architectures

We are particularly interested in the business challenge of competitive pressure to perform efficient product management. The information source for this initiative is the claim administration key system. Product management contains at least two subdisciplines: evaluating current products and developing new products in an expedient and efficient manner. Examples of new products in recent years include long-term care policies, preferred provider relationships, and mortgage insurance. Further analysis of these products reveals the nature of new products in the insurance industry and the demands placed on information technology to bring these products to market. Bringing new products to market in the insurance industry is a "soft" process, in that administrative support systems are the major investment for implementing the product; there is no manufacturing required, only capital for software development.

Product management is the key

3.1 Insurance Industry Examples

IW implementation plays a role in industry change

This section presents two examples of insurance industry systems having data processing requirements satisfied by an Information Warehouse implementation: product-to-market and the health industry. The economic environment forces the insurance industry to focus on the profitability of existing products. Profitable products are marketed and nonprofitable products are discontinued within the constraints of regulation. The few cases of new products being developed demand high efficiency for speed to market and for keeping administrative costs to a minimum. These objectives are a by-product of the competitive nature of the insurance industry.

For product-to-market, long-term care is an example of a new product being developed and brought to market. The health industry also has potential for using Information Warehouse architecture and products to meet the demands of change. This example explores unique issues of the health industry and the hospitals involved in that industry.

3.1.1 Product-to-Market

A changing world is an opportunity for new product

Long-term care insurance is a recently developed product brought to market in response to a change in the market—the growing senior population. Under traditional actuarial considerations, risk to the insurer and cost to the insuree rose dramatically at some age. The growth in the senior population indicated a change in the age at which this business was considered viable. Financial evidence indicated that insurance for this age group was creating an attractive business opportunity.

Information technology played a role in at least two specific facets of bringing this product to market: making information available upon which to design the product, and effectively and quickly analyzing the data to make the product parameters favorable to the insurer.

3.1.2 The Health Industry

Sharing patient information is a natural for IW implementation

Hospitals and hospital alternatives create a playground for shifting services and competition for customers.² The health care industry is becoming more patient-centered and less focused on the administrative and billing end of the business. Information about the customer—patient—becomes an asset that is shared by the owner with partners. For example, an X-ray taken at a hospital may be needed by a walk-in care center. A customer history may be needed by the same walk-in center or by the customer's primary doctor. In either case, there is a need for those data requesters to get the information from the provider. This translates to an Information Warehouse enterprise made up of a community hospital and multiple providers being knowledge workers requiring information from the organization asset data.

² *The Future of Health Care Information Systems*

The Information Warehouse architecture's access enablers component defines the function for delivering the information to data requesters. An alternative provider requests patient data kept in a relational database using SQL and DRDA. A more likely scenario is the use of an informational application to request data. The request is fulfilled electronically using a mapping function in the access enablers component or using electronic data interchange (EDI). This is very consistent with the overall objective in the health care industry of linking the provider with other industry participants, including health care insurers or other health care providers.

The health community needs an IW solution

3.2 Technology Requirements

Workflow management technology is a focus of the insurance study. The business requirement for workflow management grew out of an analysis of the requirements for data replication. Accordingly, data replication is discussed below as a prelude to the workflow management discussion.

Workflow Management manages complexity

3.2.1 Data Replication Requirement

Making the information available for analysis was dependent on bringing data, in appropriate form and level of aggregation, to the analyst. The data in question is basic actuarial data describing the current mortality rates of the population. This data may be owned by the enterprise or it may be purchased from outside sources. Another set of data may be profit and loss information accrued from the years of doing business in similar products for more traditional age groups. This would be a baseline for net profit, taking into account administrative overhead.

Data replication workflow management

The insurance enterprise needs to implement processes for bringing together new and existing data for the purposes of pricing the new product. The information systems department understood that the processes would be more effective and efficient if they were architected. That is, the easy approach would see each request for data as a set of applications that extract the data, enhance the data, and then present the data as information to the knowledge worker through an informational application. The preferred approach is to use the Information Warehouse architecture to establish a limited set of extract and enhancement tools that could be used and reused—with slight modification—for the specific request. The resultant information is then made available through informational applications for analysis by knowledge workers. The information systems department believed that the use of modular data replication tools would enable it to respond quickly to demands for variations of historical and purchased data.

An architectural approach is more flexible

Workflow management means control and discipline

3.2.2 Workflow Management Requirement

The first attempt at defining requirements for workflow management was based on a standard insurance process, the issuance of a new life insurance policy or rejection of the application. The process is based on a new application for insurance arriving at the head office by mail, FAX, or image copy. Major steps in the process—as modeled—are:

- Information gathering about the client and application
- Risk evaluation with the option to request a medical opinion
- Decision on acceptance or rejection of the application
- Document production and mailing.

The major requirements identified for workflow management are:

Workflow logic Workflow logic manipulation enables the graphic modeling of the overall data replication process in terms of individual process activities.

Definition point The point of definition must be the programmable workstation. The current platform of choice for administration is OS/2.

Client-server support The workflow solution must be integrated with other resources in client-server mode. That is, it should be able to launch the individual processes on the entire range of servers in the enterprise. The workflow solution does not itself require this remote launch capability but can achieve that capability by use of an administration and launch tool such as DataHub*.

Business processes in the insurance enterprise were becoming more and more complex. This trend translated to more difficulty in planning and managing the activities and resources involved in getting a job done. Controlling the workflow in an enterprise requires time and diverse skills and knowledge.

3.2.2.1 Defining Processes

Business processes are one of the insurance enterprise's most valuable assets. They are the result of putting skills and knowledge into operation; the better the processes are defined, the clearer the enterprise's direction is. However, defining the processes is a complicated exercise. Ideally, the effort in defining the processes is done once, and only maintenance and enhancements are necessary thereafter. More practically, process definition is an ongoing effort.

3.2.2.2 Complexity

Over time, business processes become more complex and harder to manage. Managing complex processes can require large amounts of funding, time, and expertise. Even the best-designed processes can falter when key personnel, with key knowledge of the process, are unavailable.

3.2.2.3 Tracking Processes

Tracking is a key part of managing processes. Enterprises desire to document the processes in detail, write procedures manuals for those processes, and train staff to monitor the processes for successful completion. This effort is in addition to the real work of the enterprise, its revenue-producing activities.

3.2.2.4 Reusing Processes

The processes developed are an enterprise asset, so it is reasonable to reuse the processes and take advantage of the investment in those processes. The enterprise needs a workflow management strategy to make process reuse possible.

3.2.2.5 Dynamics of the Industry

The insurance industry is a dynamic environment, as described in 2.1, “Insurance Industry Trends” on page 13. The need to react quickly applies to management of business processes as to all other aspects of the insurance enterprise’s information systems. The enterprise uses the existing processes as a foundation for refining other existing business processes, developing new processes, and trying them out. After using this process as an evaluation point for establishing requirements for workflow management, the information systems department revisited the overall process it had identified for product-to-market. It noted that the process contained at least two identifiable subprocesses: gathering of raw data, and presenting information in summarized, graphic format.

3.2.3 Workflow Management: Product-to-Market

The gathering and presenting activities are identified as steps in the “product-to-market” business process. The data processing equivalents of those business processes are the data copy (data replication) subprocess and the informational application subprocess, respectively. The information systems department recognized that the use of the term “subprocess” was purely arbitrary and made a commitment to adopt industry-standard terminology, if it existed.³ The data replication subprocess could be further divided into subprocesses, and it became clear that the terminology would be confusing.

Business processes translate to data processing processes

³ See 9.1.4, “Workflow Management Coalition” on page 76.

The subprocesses in this scenario included, for example:

- Extract premiums data
- Extract claims data
- Reconcile both premiums and claims data
- Aggregate both premiums and claims data by month
- Load into DB2 informational database
- Execute an informational application to generate a report object
- Print.

The strategy, then, is to manage the business activities, and map the individual steps to data processing activities. The information systems department chose to pursue workflow management technology to manage this multistep process. It contacted the Workflow Management Coalition, a cross-vendor coalition formed to establish standards in the area of workflow management and reviewed FlowMark as a tool to implement workflow management function. This review is presented in Chapter 9, "Workflow Management" on page 73.

3.3 Informational Application Requirement

Informational analysis is key to end users

The second dimension of the new strategy was the analysis process. Traditional underwriting approaches to massaging the data were based on third generation language programs. The applications would accumulate and massage the data, producing printed reports that gave pricing information. Later approaches included the use of relational DBMSs and SQL code to replace the traditional application. The information systems department felt that both the traditional and later approaches failed to recognize the significant role that informational applications play in the analysis of underwriting data.

Standardize informational applications

The solution to the informational application requirement should be based on a disciplined, standard set of informational application functions made available to the analysts to speed up analysis. It would also enable the analysis of new data for other insurance products without the lead-time of developing report applications. The standardization is defined in terms of the look and feel of the informational application and the use of the SQL API by the informational application. Both of these specifications are incorporated in the Information Warehouse architecture.

IW architecture benefited information users and providers

The use of these standards allows the business analysts to purchase informational applications independent of the information systems department and allows the information systems department to focus on managing informational data. The use of the SQL API allows the informational applications to talk to the relational informational data made available through data replication tools.

The strategic approach to solving the product-to-market pressure was expressed in terms of effective data replication and informational applications.


Product-to-market demands informational application capability

Part 3. The Technology View

Chapter 4. Insurance Application Architecture	29
4.1 IAA Modeling Terminology and Concepts	31
4.1.1 Modeling	31
4.2 IAA Framework	32
4.2.1 Models	32
4.2.1.1 The IAA Business Model	32
4.2.1.2 The IAA System Model	32
4.2.1.3 Other IAA Models	33
4.2.2 Dimensions	33
4.3 Data Model Commonality	34
4.4 IAA Data Model	36
4.4.1.1 Scope	37
4.4.1.2 Detail	37
4.4.1.3 Layer	38
Chapter 5. Information Warehouse Framework	41
5.1 Value of the Information Warehouse Framework	42
5.2 Why Data Replication	43
5.2.1 Operational Systems	43
5.2.2 Database Technology	44
5.2.3 Cost of Data Access	44
5.2.4 Historical Data	45
5.2.5 Ownership	45
5.2.6 Point-in-Time Data	45
5.2.7 Reconciliation	45
5.3 The Information Warehouse Architecture	45
5.4 Using the Information Warehouse Architecture	47
5.5 Access Enablers	48
5.5.1 Embedded SQL	50
5.5.2 SQL Call Level Interface	50
5.5.3 Distributed Relational Database Architecture	51
5.6 The Insurance Industry	51
5.7 Workflow Management Requirement	51
5.7.1 Workflow Management Function	52
5.7.2 Workflow Management Interface	52
5.8 Decision Support: Informational Applications	52
5.8.1 Informational Application Function	53
5.8.2 Informational Application Interfaces	53
Chapter 6. Insurance Solution Thread Overview	55
6.1 The Meta-data Requirement	57
6.2 Reconciliation	59
6.3 Solution Components	59
Chapter 7. Organization Asset Data	61
7.1 Business View	63

7.2 Mapping the Organization Asset Data	63
7.3 Managing the Organization Asset Data	64
Chapter 8. Data Replication Tools	67
8.1 Tool Selection	68
8.2 DataRefresher	69
8.2.1 Operating Environment	69
8.2.1.1 OS/2	69
8.2.1.2 MVS	70
8.3 Data Sources and Targets	70
8.3.1 Data Sources	71
8.3.2 Target Systems	72
Chapter 9. Workflow Management	73
9.1 Workflow Management Technology	74
9.1.1 Workflow Management Users	74
9.1.2 Terminology	74
9.1.3 Information Warehouse Framework Integration	76
9.1.4 Workflow Management Coalition	76
9.2 FlowMark	77
9.2.1 Definition and Documentation of Processes	77
9.2.2 Testing Processes	78
9.2.3 Running Processes	78
9.3 Insurance Solution Thread	78
Chapter 10. Applications	81
10.1 Knowledge Worker Requirements	82
10.2 End-User Perspective	83
10.3 Informational Application Function	83
10.4 Client-Server Operation	84
10.5 Visualizer	85
10.5.1 Visualizer Objects	86
10.5.2 Visualizer Templates	87
10.6 Information Warehouse Integration	87
Chapter 11. Conclusions	89

Chapter 4. Insurance Application Architecture



This chapter presents the Insurance Application Architecture as an example of an architecture that addresses strategic industry pressures. We also discuss the Insurance Application Architecture data model as an underlying model to the architecture. That is, the model serves as a basis for managing the things and activities in the insurance business that need to be administered by data processing systems. The overall goals are to respond to the pressures of the industry and minimize the cost of information systems.

IAA is a strategic approach to solving strategic problems

It is doubtful whether the insurance industry would ever be able to reduce the cost of its information systems unless a broad set of information models and standards are accepted by insurance enterprises and software vendors as a basis for industry applications. IBM has developed the Insurance Application Architecture as a comprehensive base for the insurance industry. IAA consists of a set of software architecture guidelines, interfaces, methods, and tools for insurance applications.

An architecture helps meet business pressures at minimal cost

IAA is a blueprint

The Insurance Application Architecture is a blueprint for developing information systems to meet the business requirements of the insurance industry. Emerging technology and improved understanding of the information systems development process make an industrywide response to these business requirements practical. The primary objective of the IAA is to provide the insurance industry with a cost-effective, low-risk means of implementing information systems to solve the broad range of application requirements. The primary interest of this study is to view the IAA as a blueprint for developing those key systems that feed or are a part of an Information Warehouse implementation. The blueprint of the business can help an enterprise respond quickly and effectively to new business opportunities and challenges.

IAA describes the operational data source systems

The IAA is comprised of the IAA models and the IAA business modeling method. These components represent the tools and processes by which information systems can be built. The information systems built from the models and the methods may be key systems that feed the informational applications that are part of the Information Warehouse framework. Systems that feed the Information Warehouse implementation are typically operational applications that manage the day-to-day business. They may be new operational applications developed to support new insurance products or they may be reengineered applications supporting existing products. In either case, they are being developed according to the IAA discipline and guidelines. IAA provides industry-oriented as well as data-processing-oriented guidelines for developing applications and is a directory into the enterprise's application portfolio as it grows.

IAA guides the data replication processes

IAA may also be used to develop applications to assist in delivering data to the Information Warehouse implementation as informational data. Delivering data in this context implies using the industry smarts of IAA to help understand the reconciliation and enhancement necessary to generate derived data for the Information Warehouse implementation. That is, IAA's broad view of the insurance enterprise helps the Information Warehouse implementation administrator understand operational data objects that appear in multiple physical data stores across lines of business or work groups but correspond to a single business object. Extensions to the IAA data models document the informational systems that are key to strategic systems.

IAA models describe the data and processes

IAA consists of a business model to describe the business and a system model to assist in the implementation. The three components that span the two models are as follows:

- Data model** Describes the information used in the enterprise
- Function model** Describes the activities that produce and update data
- Function flow model** Describes the triggers and prerequisite conditions for executing the functions.

An insurance enterprise uses the function and flow models of IAA to construct a process model of its own business processes.

The IAA business modeling method was used to develop the IAA models and can enable insurance enterprises to build models of their business by application or extension of the model to their particular business needs. IBM, IBM customers, and IBM business partners can address business requirements cooperatively and thus enable IAA to be adopted as an industrywide approach.

IAA is a common ground for software vendors and users

To satisfy the objectives of IAA, IBM provides a range of deliverable components: the models, the IAA method, and services in support of IAA. The IBM deliverable components are expected to be supplemented by offerings and services developed by other organizations and implementing insurance enterprises.

The IAA business modeling method is the essential link between the IAA objectives and their realization in the IAA business model. The method was the basis for development of the IAA business model. It is the recommended approach for organizations that want to create a business model or extend its scope or increase its detail.

4.1 IAA Modeling Terminology and Concepts

Some fundamental IAA modeling concepts and terms are described below, followed by more detailed description of the IAA method and its role. It may be of value to review the discussion of models and modeling in Appendix A, "Models and Modeling" on page 91. Models and modeling are generally accepted as tools that are beneficial to the enterprise and the industries, both insurance and data processing. Their benefit has been limited by the lack of commitment on the part of the data processing industry in general and the enterprise at the project level. For further discussion of these points, see *The Model Business*.

4.1.1 Modeling

Modeling is the act of creating, extending, or refining a model. It may be done primarily for analysis and understanding, or for communicating. In either case, it involves the application of a method and it may require the use of tools. Modeling can be limited in scope and level of detail so as to focus on specific objectives and produce a corresponding subset model or deliverable components. These components can be packaged in such a way that they constitute an architecture.

The primary output of the IAA method is a business model. Such a model is a symbolic representation of reality, a shorthand representation that is easier or faster to comprehend than the reality itself and is therefore well suited to the evaluation of business options. A model may represent a variety of things: a real object, another model, a part or aspect of reality, or a complete reality. The IAA models are used for different purposes, the sum of which is coverage of the bulk of requirements for modeling in the insurance industry.

4.2 IAA Framework

This section describes the models and the dimensions of those models included in the IAA framework.

4.2.1 Models

The IAA models present the structure of the IAA architectural components that comprise the submodels used to describe the information, activities, and rules used by organizations to conduct their business. The IAA models hold two fundamental views of the business in a prescribed relationship. One view represents the components relating directly to development and processing of the primary business. The second view represents procedures relating to the development and operation of information systems that support that business. Mapping the two views together is important to the effective communication between business and data processing staff. Thus the IAA models hold separately two related but distinct architectural components, as follows:

- A business model, which represents the business and is the outcome of an application of the method
- A system model, which represents the information systems that are needed to support the business.

Figure 2 on page 34 provides a three-dimensional representation of the IAA framework.

4.2.1.1 The IAA Business Model

The IAA business model represents the insurance business by showing business functions and their relationship to business data and business rules. It is not a financial model, an organization chart, or a human resources view of business. It represents the data and operational procedures of the business as they are or could be.

4.2.1.2 The IAA System Model

The IAA system model represents the technology interface and activities associated with the development of information systems to support the business activities represented in the business model. Other models can be included in the system model. For example, the process models for specific enterprise applications can be added to the system model to extend completeness of coverage.

4.2.1.3 Other IAA Models

The IAA data model, the IAA function model, and the IAA function flow model are views or submodels of the IAA business model and system model. Each concentrates on a particular aspect of the business: the information, functions, and flow rules, respectively.

The IAA models, collectively, are a framework that places the business and system models and the data, function, and function flow models in perspective. More precisely, the IAA models view the business and system models as an aggregate view of the data, function, and function flow models. This is compatible in approach with the enterprise and logical models described in generic modeling terminology. That is, models are easier to use if they start out with high-level models with few components, each component of which is resolved to more numerous subcomponents.

4.2.2 Dimensions

The IAA models are open-ended and therefore theoretically capable of covering every business and supporting information systems activity of interest to any organization over any period of time. However, since practical limits must be established to provide a focus for the modeling exercise, there is a consequent limit to the extent of the IAA models. These boundaries may vary from organization to organization. For example, an insurance enterprise that is also a bank will have a broader range of business activity than one that simply offers a single line of life insurance products.

Three dimensions are defined to describe the limits of the IAA framework:

- Scope** The scope indicates the breadth of business activity covered. A model covering only one line of business is of narrow scope, while a model covering all insurance-related activities would be regarded as broad in scope.
- Detail** The detail indicates the degree of detail or specificity to which the model is developed. A conceptual representation of the business has relatively little detail, whereas a complete specification of all information would include more detail.
- Layer** The layer indicates the intensification of focus as the perspective shifts, from conceptual business definition to specific business requirements to planning and design issues to specific system operational activities.

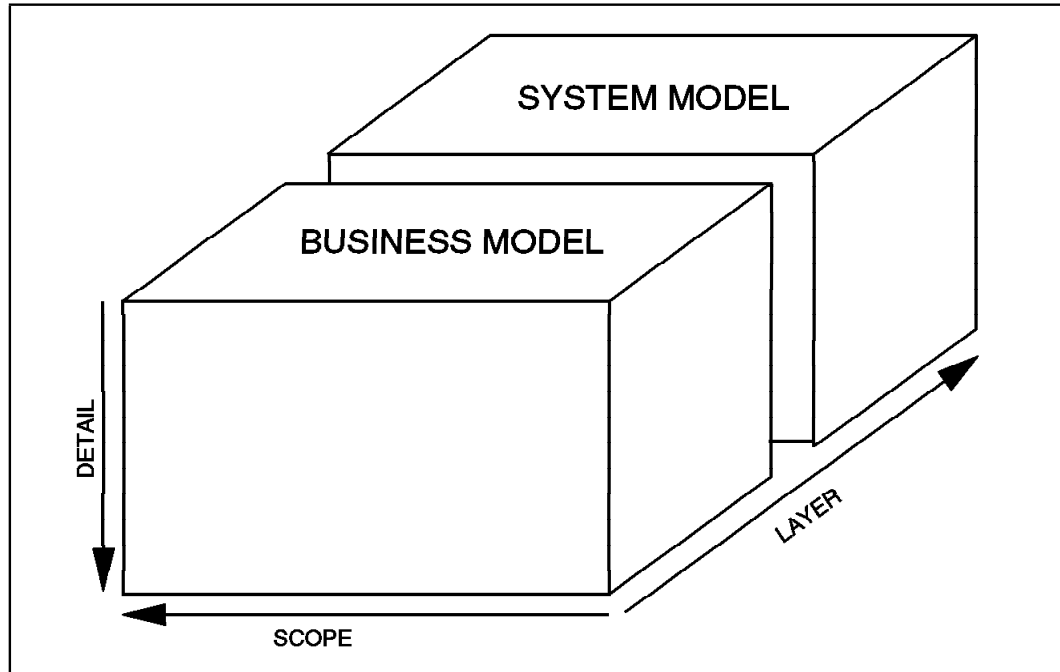


Figure 2. IAA Framework: Simplified

4.3 Data Model Commonality

IAA leverages the 80% overlap across the industry

An insurance industry data model is a set of diagrams and definitions that represents in a consistent way the data and data relationships an insurance enterprise must maintain to support its business function. The IAA data model is a generic model in that it represents the data and interrelationships that are found to be common throughout the insurance industry. This overlap—in the types of data retained and used by the business—has been estimated to be roughly 80% across insurance enterprises and lines of business.

Several points of commonality are recognized with respect to data models in the development of information systems for the insurance industry. These include the following:

- Process

The insurance industry is recognized as an industry group because it carries out similar processes on similar objects under broadly the same market pressures. Business objects such as perils, policies, premiums, agents, and clients are recognized in all insurance enterprises.
- Application development

The development of information systems and the application of information technology to the business is carried out in more or less the same way across insurance enterprises. The tasks of business analysis, system design, technology selection, and application development and implementation are universal information systems processes. Furthermore, the achievement of mutual savings by sharing the cost of potentially common activities, such as developing models and adopting

packaged software, among several organizations is practiced by most insurance enterprises.

- Business environment

The same market and competitive forces are felt by most insurance enterprises, so that the desire for easy and low-cost change is common. All enterprises have a need to know how their information is currently organized and what changes are needed to better organize the information.

- Communication

Clear communication within and between business and technology planning is crucial for successful change, irrespective of industry or organization.

Despite this commonality, each insurance enterprise has its own data model, either implicit in its databases or explicitly defined and published within the enterprise. The causes are partly historic and partly because no one insurance enterprise has the capacity to set a standard. The historic causes center on the fact that data modeling, and hence the recognition of the extent and nature of data and process commonality, is a relatively recent development. The potential to design data systems that are sufficiently flexible to be used by an entire industry group is also a recent development.

The absence of a standard is partly due to the need for enterprises to strike a balance between the economies of scale brought about by all parties—insurance enterprises, software vendors, and technology suppliers—adopting one standard and the need for individuality. Individuality is needed to accommodate existing disparate systems as well as allow the organization to differentiate itself in the insurance marketplace. The need for a single organization to bring a common standard into being is addressed by IBM being the catalyst for those insurance enterprises desiring a common model.

Information systems technology advances have overtaken the concept of differentiation by technology. The emphasis today is on differentiation by virtue of what an organization does with the applications and delivery systems made possible by the technology. This pressure is discussed in 2.1, “Insurance Industry Trends” on page 13, in that customer service will differentiate industry leaders from followers. Customer service in the insurance industry is based on the ability to provide accurate, up-to-date information on a customer’s relationship to the enterprise, at the customer’s demand.

4.4 IAA Data Model

The IAA data model is a multidimensioned structure, whose bounds are set by the extent of its scope, detail, and layer dimensions (see Figure 3 on page 37). The model recognizes four layers of transition from business map to implementation and operation for data. The goal of the IAA data model is to address the need for common model structures and terminology throughout the information systems. The goal is achieved through the data model's combination of stability and flexibility, and its position as a foundation for the insurance enterprise. The IAA data model's stability and flexibility are evident in its support for the development of portable and integrated insurance applications. This design includes a common view of information that is acceptable to both information systems and insurance professionals. It also assists in the mapping of existing data models into the IAA data model.

The IAA data model's value is realized by those insurance enterprises and software vendors that accept it as the standard data model for insurance applications. It therefore promotes itself to the insurance community worldwide as an unbiased, superior-quality, broad-scope, common view of the core operational activities of the insurance industry.

The IAA data model is intended to be independently useful as a basis for designing correct, consistent, flexible databases that may be shared across an enterprise or across the industry. It is, then, a foundation for a common architecture for insurance industry applications.

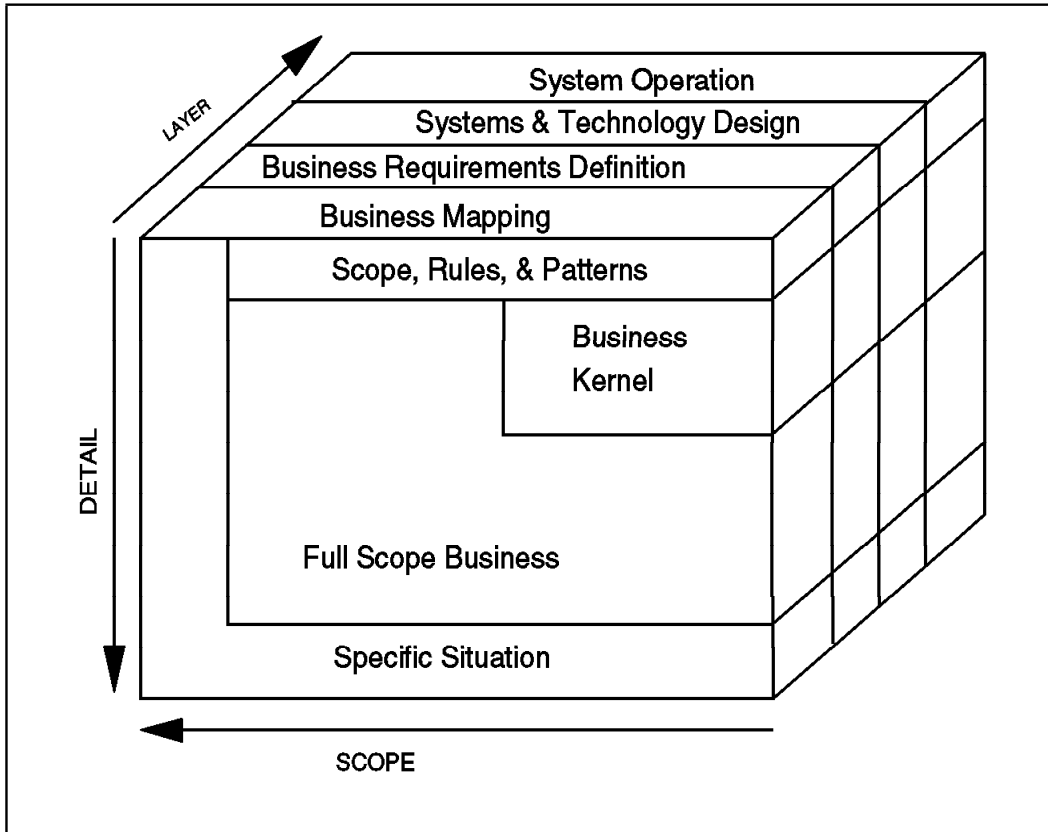


Figure 3. Model Extents

The entities and associated relationships defined in the map layer provide the control mechanism for developments in the business requirements definition, system and technology design, and system operation layers to produce a working information system. The IAA data model's scope, detail, and layer have been defined to suit the insurance industry as described below.

4.4.1.1 Scope

The scope is set to establish a practical breadth of insurance business activity to be included in the model. The Insurance Application Architecture's current scope embraces all of the major common activities of the insurance industry.

4.4.1.2 Detail

The detail is set to ensure that the extent to which entities are applied in the information systems is useful in practice. The detail of the IAA data model extends to the point where the commonality of the model content starts to be overshadowed by unique organizational or national requirements. Insurance enterprises can extend the generic IAA data model to accommodate the level of detail needed in their own business.

4.4.1.3 Layer

The layer is set to include the practical implementation requirements of most organizations. The IAA data model is primarily in the business mapping layer but also has elements that support the other layers (business requirements definition, systems and technology design, and system operation).

Models have a role in Information Warehouse solutions

Appendix A, “Models and Modeling” on page 91 discusses modeling and concludes that models have failed in meeting the promise of automating application development. This leaves open the question of what value there is to models and modeling. The Information Warehouse architecture addresses the need to access operational data through informational applications. The informational applications access informational data delivered from operational data sources through data replication techniques and products. This process makes two assumptions about the enterprise: the data replication designers and implementers know what operational data is available, and the knowledge worker knows what informational data and objects have become available through the data replication processes.

Object awareness is step one

Both assumptions about the data awareness of the data replication personnel and knowledge workers are addressed by the DataGuide family. DataGuide is a catalog that helps knowledge workers become aware of information and locate information. However, DataGuide requires population of its meta-data store. Population implies taking a massive amount of business terms and relating them to information systems objects. The information systems objects—for example, operational data sets—are within the realm of awareness of the information systems department. The information systems department as a whole should be familiar with every information systems object, but no one person can be familiar with every object. Therefore, the information systems department could benefit from DataGuide by making the collective awareness available to all information systems department members.

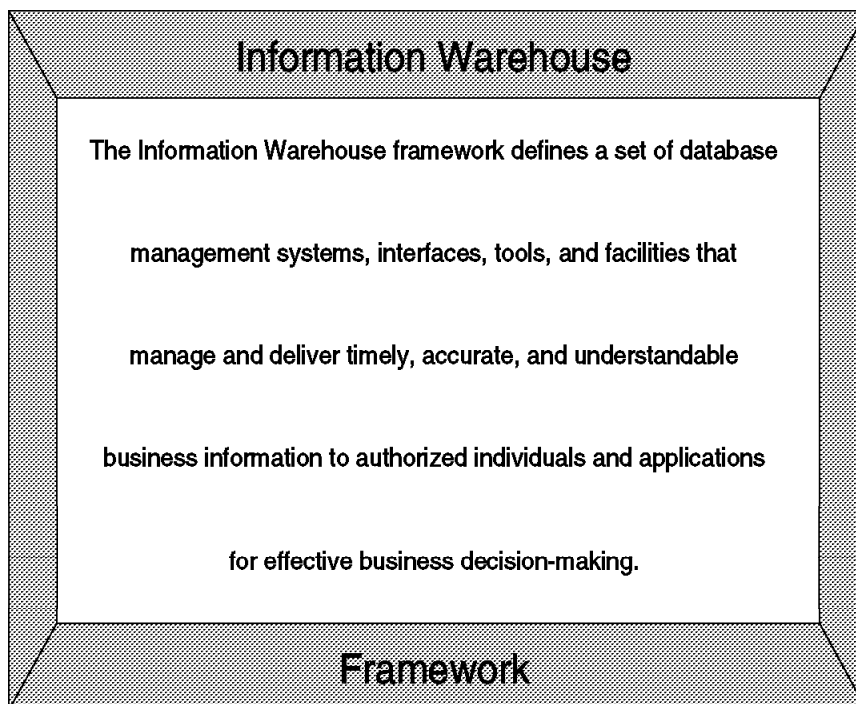
Business term awareness is the major challenge

It is more difficult to extend the DataGuide function to the business analyst than to the information systems department staff. If we assume that the knowledge worker is a member of the information systems department, then the process of building the meta-data for the operational data objects is primarily carried out within the information systems department. In the worst case, input is required from application development staff from the lines of business who are familiar with information systems concepts and terminology, so the meta-data gathering process is relatively easy. Building the business term meta-data is a greater challenge. This invariably requires input from line-of-business experts and analysts. These analysts know the business end but may not be able to relate the business terms to the informational objects being represented by meta-data in the DataGuide store. Furthermore, the meta-data gathering process implies crossing organizational lines and taking up the expensive time of business analysts. Therefore, gathering the business terms for the DataGuide store is difficult.

The IAA model contains a great amount of information in business terms; it is a great source of meta-data for populating DataGuide. This can be accomplished by building an extractor to generate a DataGuide import/export tag language file and then invoking the import function against that file. This approach would allow the bulk population of the DataGuide from an industry knowledge source. For more information on DataGuide tag language and import function, see *DataGuide/2 V1: Using DataGuide/2*.

*Industry
models can
feed
DataGuide*

Chapter 5. Information Warehouse Framework



Enterprises have long recognized the opportunities that would be available if they could make better use of their data. Data is typically stored in many locations, in different formats, and managed by products from many different vendors. It is usually difficult to access and use the data across locations and vendor products. The Information Warehouse framework is a solution to this long-standing problem. IBM and other vendors are working to make access to data across vendor products and geographic locations easier by enabling their products to work together. The products and rules by which the products can work together constitute a framework. The Information Warehouse framework is designed to provide open access to data across vendor products and hardware platforms.

The framework: better use of enterprise data

IBM and vendor solutions fit into the framework

IBM and its business partners have been delivering database, data access, and decision support products which fit into the framework and make it possible for our customers to build effective integrated Information Warehouse solutions.

The framework enables access to all enterprise data

Enterprises benefit from the framework because they can increase the value of the investment that they have in current databases and files. Enterprises can get to the data that they need to effectively manage their businesses.

End users spend more time using data, less time finding it

The Information Warehouse framework of products includes decision support products. These applications can be used by end users to analyze and report business data from many parts of the enterprise. The Information Warehouse framework of products gives the end user access to data from other workstations, LANs, and host databases. The end users spend less time gathering and accessing the data, and more time in analysis and reporting of data.

5.1 Value of the Information Warehouse Framework

The Information Warehouse framework is a comprehensive solution having value beyond the simple collection of software products. The following characteristics differentiate the Information Warehouse framework from other approaches:

- **Published architecture**
Products that implement the published Information Warehouse architecture can work together with consistent user interfaces for easier operation.
- **Cross-platform coverage**
Database products reside on multiple platforms, and access is enabled to database products from a variety of software vendors on platforms from a variety of hardware vendors.
- **Architected connectivity**
Distributed relational database connectivity is architecture-based. IBM has products which support this architecture today.
- **Vendor support**
The Information Warehouse framework has the support of other vendors in the industry. These vendors are working with IBM to bring a wide variety of software solutions to market.
- **Architected systems management**
DataHub* is a cornerstone of the Information Warehouse framework. It is an architected solution for systems management, with the intent to cooperate with systems management products from other vendors.

- Integrated database tools

Tools for database design work with tools for application development, saving customers time in application modeling and development.

We can address the data delivery requirement to generate informational data from operational data by either iteratively writing applications to extract, enhance, and load information, or by using the architecture-based Information Warehouse approach.

5.2 Why Data Replication

There are two approaches to accessing data for informational analysis: accessing the operational data directly and accessing a replication of the data created by extraction, enhancement, and loading into the informational database. The consensus favors accessing informational copies of operational data rather than direct access of operational data. Some of the considerations leading to this preference are as follows:

- Operational systems
- Database technology
- Cost of data access
- Historical data
- Ownership
- Point-in-time data
- Reconciliation.

5.2.1 Operational Systems

Operational systems manage the day-to-day business activities. As such, they are critical to the ongoing viability of the enterprise. These systems often perform at the limit of the hardware and software with which they are implemented. They are often a key part of customer service, a major factor in the success of an individual retail enterprise in the very competitive retail industry. Therefore, the ongoing operation of the operational systems is the highest priority. The reasons for using copy-based informational systems with respect to protecting existing operational applications fall primarily into the areas of data accountability and application performance.

Operational systems operate at technology's limit

Data accountability includes security and audit considerations. Operational systems are designed from the beginning with a specific user community in mind. The data created and manipulated by those applications are carefully managed with respect to who has access to the data. The management is accomplished through operational policies and security function in the software. It includes allowing access of specific sets of data by specific users or by certain user group identifiers and associating specific users with those group identifiers. Specifying every combination of user and data is a considerable effort, so the group identifier approach is more practical. However, this may result in defining group identifiers for broad groups of users with diverse profiles to access operational data. At the very least, allowing informational access by knowledge workers would be an incremental burden on the security administrator. It could possibly be a major burden on the security software and a complication for the security policies of the enterprise. It

An isolated copy is less complicated to secure

is easier to have a copy of the data in an isolated informational environment where the access security can be handled independent of existing operational systems policies.

Data location impacts security

Data placement is also an issue in security. Copying data could increase the security exposure to the enterprise; once the copy is made, it is up to the possessor to manage security. However, controlled copying is more secure than allowing broad groups of users with diverse access profiles to access operational data. At the very least, fresh copies of data are controlled.

Performance of operational systems must be protected

Decision support activity is difficult to predict, but certain characteristics of this class of applications are well understood. Decision support queries tend to access large volumes of data; they tend to apply complicated, long-running manipulations against that data; and they may retain claims on that data for extended periods of human think time. The queries, then, can be expected to interfere with transaction systems because of extensive locking, heavy I/O activity, and high demands on CPU and buffer pool resources. Informational analysis against copies of the operational data rather than the operational data itself prevents this interference.

5.2.2 Database Technology

Database software and hardware technology favors copies

Mainframe database technology supports a wide range of operational function in terms of concurrent access and data transfer rate. The Information Warehouse architecture is platform-independent and is compatible with LAN-based operational databases and application strategies. The flexibility of the mainframe platform makes it an ideal location for data copies. The mainframe platform can support a wide range of data volume and user populations. It also leverages data by being a central location for data to be accessed across work groups or lines of business.

The LAN platform has value for certain informational uses

The LAN lags behind the mainframe in I/O interface technology and ability to process data in memory. Recoverability, availability, and security functions on the mainframe tend to have fuller function. This has historical reasons as well as reasons based on the sheer volume of data inherently found on the larger systems. There are, however, many situations where specific subsets of data can be processed effectively in the LAN environment.

5.2.3 Cost of Data Access

Bring the data copy to the knowledge worker

Communication costs are reduced and response time improved if a copy is created in one or more locations. The general rule is to bring copies of data as close to the knowledge worker as possible. The store structure in retail is an example of this strategy. The network topology includes LAN-based branches and central host databases. Frequently accessed data is accessed in a most cost-effective way on a LAN server rather than running queries on the host. Factors that influence placement of data copies on the LAN or the large server include the cost of data processing (host versus LAN server) and of sending the answer set from the host (LAN communication cost is lower).

5.2.4 Historical Data

Operational systems usually do not allow for historical data analysis, yet this is a major concern to business analysts.

5.2.5 Ownership

Legacy system history along with security and availability requirements have placed data remotely from business activity. Copying data allows for data placement closer to knowledge workers responsible for making decisions. Ownership implies identifying an individual who is responsible for the quality and currency of the informational object.

5.2.6 Point-in-Time Data

Operational data tends to change over time. A point-in-time picture of information may be necessary for comparisons and understanding trends. Point-in-time information is part of a historical database strategy.

5.2.7 Reconciliation

Reconciliation of operational data is impractical as a real-time operation. Staging of data and data replication techniques are necessary to perform the reconciliation required for informational analysis without impacting the operational environment.

5.3 The Information Warehouse Architecture

The primary goal of the Information Warehouse architecture is to define a basis giving end users and applications easy access to data. The Information Warehouse architecture (see Figure 4 on page 46) defines a structure, formats, protocols, and interfaces as the basis for implementing Information Warehouse solutions. This architected approach creates an environment wherein solutions are leveraged. The leveraging is realized by reusing individual component solutions across implementations and by integrating off-the-shelf components in those implementations. This is not to say that an Information Warehouse solution cannot be implemented without the architecture. Rather, the architecture contributes to the leveraging of effort and resources.

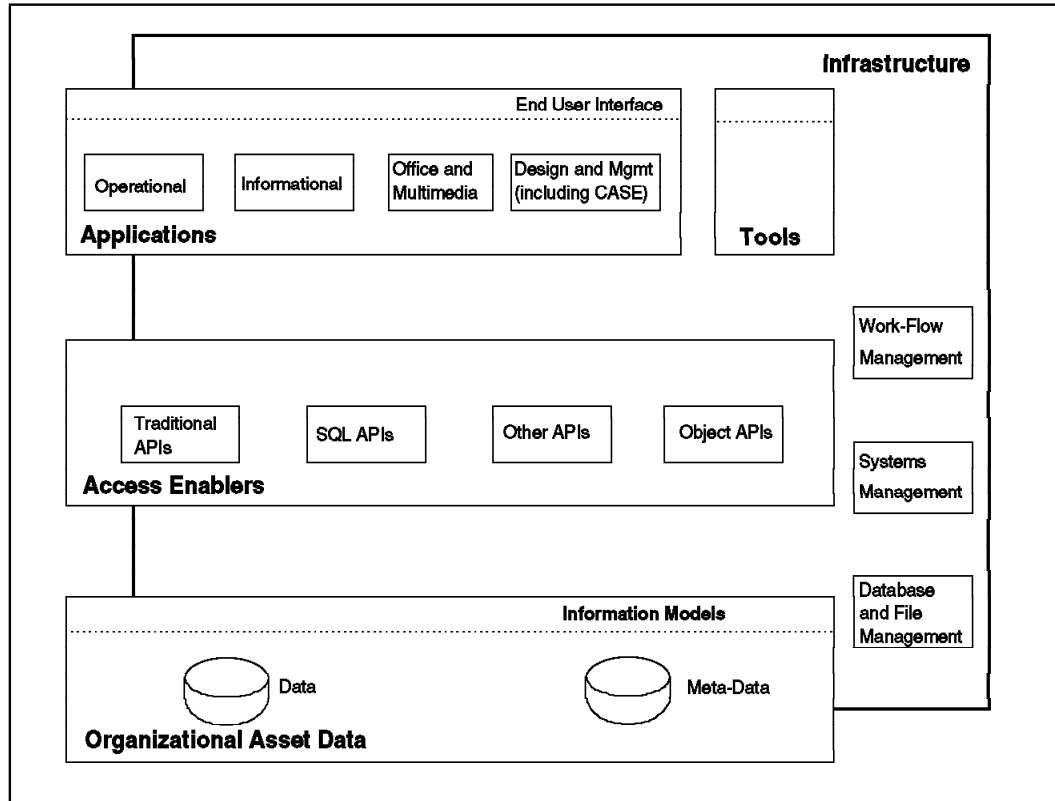


Figure 4. Information Warehouse Architecture

The long-term goal of the Information Warehouse framework is to provide access to data of all types in all stores in any environment, and the architecture is designed to accommodate the goal. The Information Warehouse architecture is open in that the interfaces are published and extensible in that software tools and data volume can be added without regressing the existing implementation.

The Information Warehouse architecture defines interfaces, protocols, and formats for accessing information in an Information Warehouse implementation. These interfaces are, as follows, grouped by focus area:

- Informational applications
 - End-user interface to informational applications
- Information Catalog and its access
 - Information Catalog API
 - Import/export interface to the Information Catalog
- Access to data
 - Embedded SQL API
 - Callable SQL API
 - Distributed Relational Database Architecture
- Data replication
 - Interface to the object handler meta-data
 - Tool invocation to tool interface
 - Interface to workflow management
 - Data staging interface.

The interfaces are identified and described in *Information Warehouse Architecture 1* for the public domain. The Information Warehouse architecture

approach, using a component structure with interfaces defined between the components, and its openness regarding system platforms makes it easier for an enterprise to implement an Information Warehouse solution on its own or with the help of software vendors and service providers.

5.4 Using the Information Warehouse Architecture

Figure 5 on page 48 shows the three fundamental components of the Information Warehouse framework: the architecture, products, and services. The three components work together to build a foundation of an extensible, flexible, and scalable Information Warehouse implementation. The products and services are requirements, whereas the architecture is a recommended participant in the Information Warehouse framework strategy.

Information Warehouse framework for access to data

The products are considered requirements because the Information Warehouse framework is a software solution. The products component encompasses any software solution to the Information Warehouse framework function requirement, not just purchased software. Off-the-shelf software has its advantages in speed of implementation but costs real money and may require some investment in resources to customize and integrate into the enterprise's environment.

Off-the-shelf software for productivity

The services component refers to resources expended by people, be they enterprise knowledge administrators or personnel hired from outside the enterprise. In either case, people must do the work of designing, developing, and implementing software solutions.

Services: implementing the solution

The Information Warehouse architecture is a structured approach for building a solution. Information Warehouse implementations built without the architecture will solve a problem, but they may not be easily extended. The architecture-based, leveraged approach allows for reuse of software for similar functions across lines of business or specific requests. It allows the incremental addition of new function without disruption to the existing implementation. It also allows the growth in usage and data volumes without disruption of the existing implementation. The Information Warehouse architecture-based implementation is the recommended approach, though it is not the only approach.

The architecture for a flexible solution

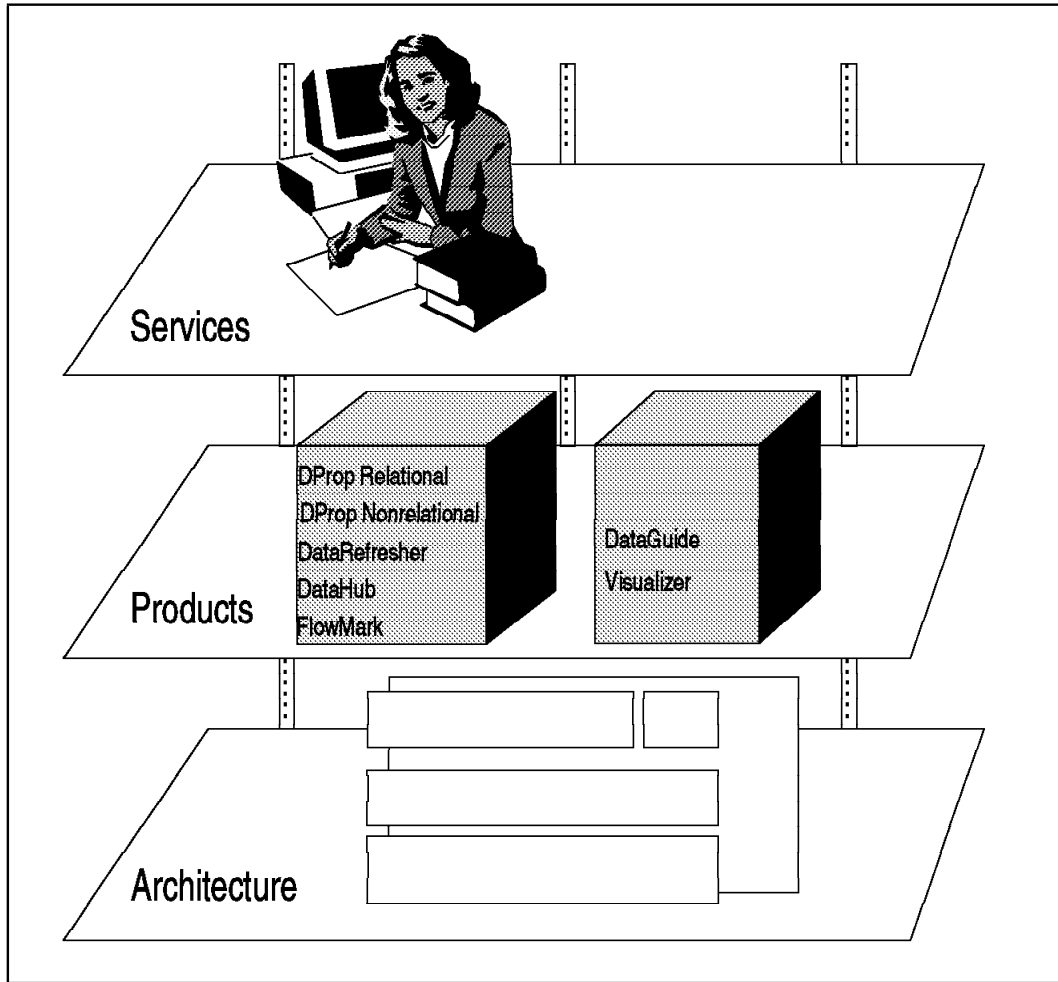


Figure 5. Information Warehouse Framework

5.5 Access Enablers

Access enablers connect applications and data

Access enablers is the layer in the Information Warehouse architecture between the application and the data (see Figure 6 on page 49). The data includes meta-data as well as real-time, changed, reconciled, and derived data. In the *Information Warehouse Architecture I*, the focus is on informational applications using SQL. These applications access relational databases locally with SQL or remotely with SQL and for example, DRDA. The value in using SQL and DRDA is that the application uses the same data access language to access local or remote relational databases. Nonrelational databases can be accessed by using SQL mappers in the implementation of the access enablers layer.

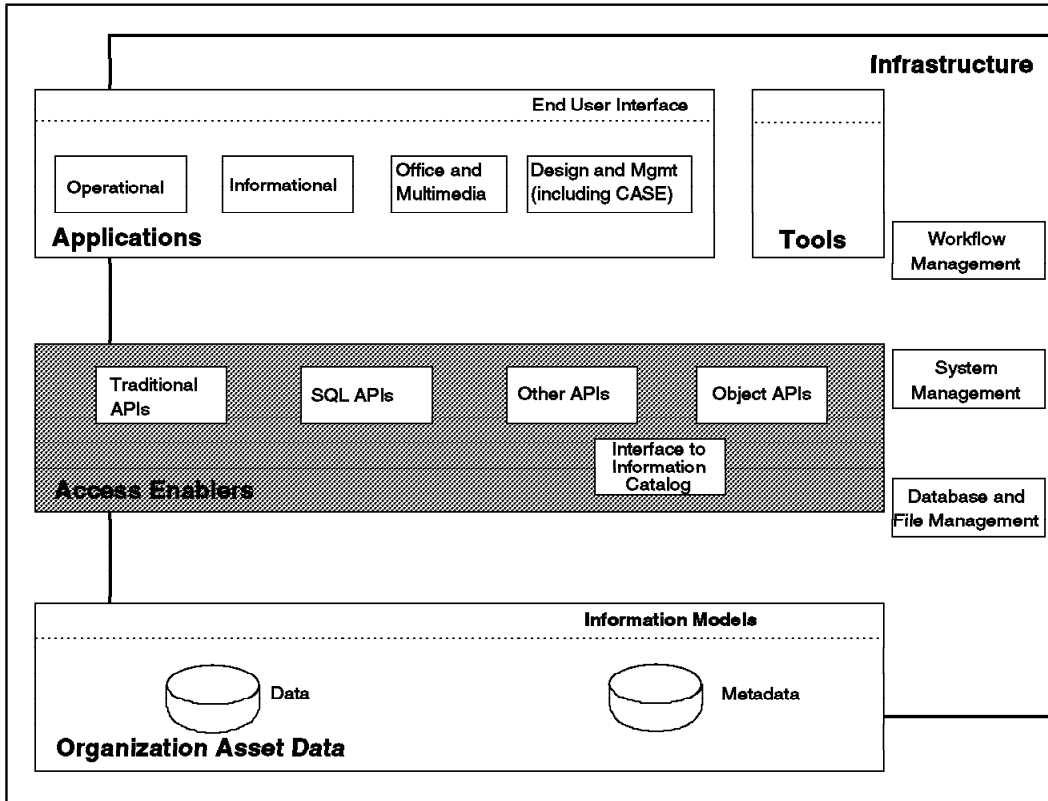


Figure 6. Access Enablers

The four focus areas of *Information Warehouse Architecture I* limit the discussion of the access enablers to the SQL application program interface and the interface to the Information Catalog. The concepts of *enabling* products and *deploying* products are key to understanding the use of the access enabler layer APIs and interfaces:

Interfaces are enabled, then deployed

Enabling An enabling product is a software program that accepts the commands defined in the API or interface and executes their defined function against a resource. DB2 is an enabling product for the SQL API, and the DataGuide products are enabling products for the interface to the Information Catalog.

Deploying A deploying product is a software program that submits requests for data resource in the form of commands defined in the API or interface. The commands are submitted to the enabling product for execution. Visualizer is a deploying product of the SQL API, and the DataGuide knowledge worker end-user interface is a deploying product for the interface to the Information Catalog.

An informational application could include commands defined in the interface to the Information Catalog and would become a deployer of both the interface to the Information Catalog and the SQL API. The advantage of this approach is that the knowledge worker continues to use the familiar environment of the informational application. The informational application would be enhanced by using business terminology stored in DataGuide.

SQL is used to access the informational and operational data categories and the interface to the Information Catalog is used to access meta-data. SQL is an industry standard data access language for performing relational operations, normally against data in a relational database. The access enablers layer also includes the definition of SQL mappers that allow SQL operations to be executed against nonrelational databases. Three key interfaces are included in the Access to Data focus area in *Information Warehouse Architecture I* and are related to the SQL data access language:

- Embedded SQL
- Callable SQL (commonly referred to as SQL call level interface)
- Distributed Relational Database Architecture.

The embedded SQL interface is an example of an interface taken from the public domain and is recognized by several standards bodies. The SQL call level interface is not as yet a standard; its prime focus is to enable software vendors to market shrink-wrap informational applications. The Distributed Relational Database Architecture was developed by IBM and is gaining recognition and support from a range of software vendors.

5.5.1 Embedded SQL

Embedded SQL refers to commands included in informational applications source code. The informational application must undergo special processing prior to normal compilation. It is this special preprocessing requirement that has fostered the development of the SQL call level interface.

5.5.2 SQL Call Level Interface

The SQL call level interface (CLI) is an alternative mechanism for invoking SQL from programs. The objective of CLI is to provide additional language commands (verbs) to extend the function of SQL. The most desired extension to SQL function is support of “shrink-wrap” applications. Software vendors would like to market applications utilizing the SQL but have experienced difficulty with the preprocessing and BIND requirements of embedded SQL. The informational applications are targeted for knowledge workers with little data processing skill. Requiring knowledge workers to go through the precompile, compile, and bind steps would diminish the acceptance of the informational applications by these users. The CLI introduces extensions to the embedded SQL command set which allow run-time precompile and BIND. The informational application can be used out of the box and does not require systems or database administration resource for the SQL portion of the application.

5.5.3 Distributed Relational Database Architecture

Distributed Relational Database Architecture (DRDA) is a communications vehicle for issuing SQL statements to a remote relational database and returning the results. The statements are executed against a remote relational database rather than a local database. Though there may be differences in the relational databases, such as the form and content of catalogs, the SQL statements themselves normally should not have to be changed to reflect the new location. It is an evolving architecture: DRDA level 2 introduces distributed two-phase commit protocols. IBM has developed and published the DRDA for use by software developers throughout the industry.

The overall objective of the access enablers component is to insulate the application from the enterprise data format and location. SQL mappers manage the mapping from SQL in the applications to nonrelational enterprise data when necessary. DRDA allows for specification of the location of the relational enterprise data. That is, an application using SQL can execute against a local database on that programmable workstation's DB2/2 database. That same application can execute against a relational database on the LAN server running DB2/2 by moving the database and causing the application to connect to DB2/2 on the LAN server. That same application can execute against a relational database on a remote server running any DB2 family member by moving the database, using DRDA and causing the application to connect to the DB2 database on the server.

5.6 The Insurance Industry

The pressing business requirements for the insurance industry fall into the areas of workflow management and informational applications. These requirements are addressed in the Information Warehouse architecture's Workflow Management and Applications components, respectively. This chapter presents these two components from an Information Warehouse architecture point of view, and sets the stage for the insurance solution thread that follows.

5.7 Workflow Management Requirement

The Workflow Management component supports the objective of process integration, which can be defined as coordinating the execution of different tools or applications to complete a task. Workflow management includes facilities to define a complex of activities that represent tools and applications to be coordinated in the execution of large tasks. It also manages tasks to be performed by business staff independent of data processing systems. These are considered activities in the overall business flow and are represented in the workflow manager.

*Workflow
integrates DP
and nonDP
activities*

5.7.1 Workflow Management Function

Simple activities may not need a workflow manager

As an example, consider the task of extracting a subset of real-time data and loading that data into a derived data store. Both the source and target data are on a single system, are clean—free of data integrity and consistency errors. The process integration is trivial and is achieved by a single local activity possibly containing multiple steps. In this case, it may not be necessary to develop a complex definition of activities and to execute it using the facilities of workflow management.

Most activities are complex, DO need a workflow manager

If the source and target data are on dissimilar, remote systems, if the source data is on a different data store type from the target data, if the source data is not clean, and if the execution request is based on some external trigger such as time, then process integration becomes more demanding and complex. Several independent applications and tools have to be coordinated into the larger task which can be triggered without human intervention. Workflow management components are needed to define this task and to coordinate its execution.

5.7.2 Workflow Management Interface

Information Warehouse Architecture I defines an interface for workflow management as one of the four data replication interfaces. This interface allows the workflow management software tool to be a plug-in participant in data replication. It allows the data replication administrator to define data replication processes independently of the steps performed by the copy tools. The data replication administrator defines simple or complex, multistep processes using the end user interface of the workflow management tool. The workflow management tool uses the workflow management interface to interact with the tools specified in the workflow processes, thus making the overall effort highly modular, and more manageable. As this interface is between tools that the insurance enterprise plans to purchase, it is not of immediate concern to the information systems department, but it is more a value to the overall strategy of the insurance enterprise's Information Warehouse implementation effort. That is, the enterprise may purchase a variety of copy tools that need to be integrated into existing or new workflow definitions. The interface facilitates the integration of dissimilar copy tools without affecting the workflow manager.

5.8 Decision Support: Informational Applications

Informational applications are software tools that present informational data to knowledge workers as informational objects rather than raw informational data. Informational objects include spreadsheets, charts, reports, queries, and images.

5.8.1 Informational Application Function

Informational applications help knowledge workers with such tasks as querying informational data and preparing reports and spreadsheets, and also more complex tasks such as business planning and modeling and project management. Informational application function varies to accommodate the needs of the end user. For example, executives are relatively results-oriented. They do not get involved with the creation of the informational object, but would want to view one already defined. The informational data behind the informational object tends to be aggregated at a high level with respect to the organization and the need for detailed information is minimal. These characteristics describe the area of executive information systems (EIS). EIS users prefer canned queries and reports with minimal customization.

Informational application requirements vary by end user

Business analysts tend to be more geared toward the development of the informational application request. They get involved in the formulation of the SQL query, and in the design of the graphic representation. They would decide whether a spreadsheet, pie chart, surface chart, or histogram is the best presentation style for a particular request. They may be involved in the specification of the informational data to be accessed as input to the report. These characteristics describe the area of decision support (DSS). EIS and DSS systems do overlap because the end users vary in their capabilities and interests.

Business analysts are more active in report formulation

The *Informational Applications* component of the Information Warehouse architecture is an architectural approach to functions that meet the needs of both EIS and DSS communities. The total set of functions provided by the informational applications may be mapped, function-by-function, to those communities. However, the Information Warehouse architecture approach provides a basic set of functions from which DSS or EIS requirements can be met.

IW architecture unites DSS and EIS

5.8.2 Informational Application Interfaces

Information Warehouse Architecture's interface for informational applications identifies specifications for how informational applications should look and feel to the knowledge worker, as follows:

- CUA**'91
- OSF**/Motif**
- Apple** Desktop Interface

DS interfaces optimize knowledge worker skill

The inclusion of the specifications recognizes the variety of environments present in the marketplace. The objective is to optimize knowledge workers' skills by limiting the variety of appearances informational applications have. Knowledge workers can move from among informational applications without learning a whole new way of using program function keys or help screens.

Chapter 6. Insurance Solution Thread Overview

The insurance industry has historically made excellent use of transaction systems to run its business. As an industry, insurance enterprises have lagged behind those in other industries in the area of management information systems (MIS), also called executive information systems (EIS) and decision support systems (DSS). MIS take operational, transaction system data and make it available for decision support processing. For example, an operational transaction system might create a record to represent the payment of a premium. The record would be created in a VSAM data set, whose business term name is Premium. The data processing name of this file might be DSID1.DSID2.PREMIUM, but the business analyst would know it as the Premium file, and the modeling analyst would know it as the PREMIUM entity. This set of concepts, comprised of:

The industry needs to improve MIS capability

Data records	Data processing record of a premium paid
Business term	The business analyst's perception of the premiums paid
Entity	A modeler's perception of the business activity, premium paid

is a unifying example of using modeling to bring together business analyst, data processing professional, and operational support. Figure 7 on page 56 shows the relationship of business concept to modeling entity to operational record.

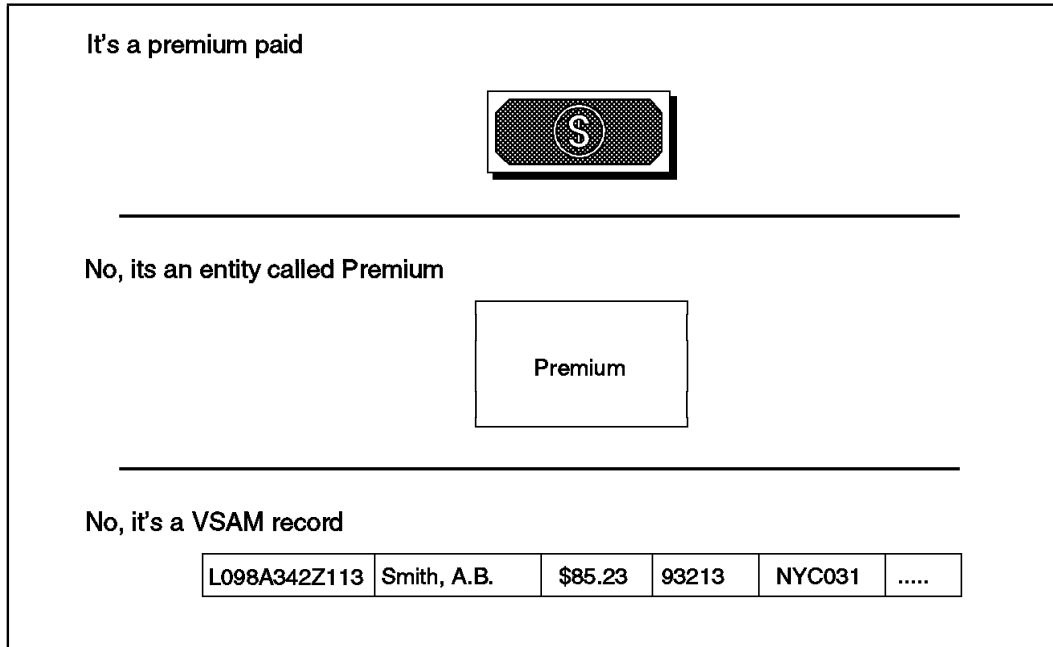


Figure 7. Business Concept, Modeling Concept, and Operational Record. They are different views of the same thing.

The variety of informational analysis demands data replication

The activity of paying a premium is a day-to-day operational occurrence with a very narrow focus—a single transaction. The MIS objective is to take some business-defined set of these transactions and aggregate them in such a way as to focus on more data presented as less data. For example, one MIS analysis might sum up all of the premiums paid on a given day or paid against a given set of insurance policies. The variety of analyses is unlimited and clearly defines the need for flexibility in delivering informational data. For every business analyst's new aggregation or subsetting of data, a new set of information must be delivered by the information systems department. This requirement for flexibility is a justification for a data replication strategy.

Need to deliver dispersed data in reconciled form

The MIS strategy, however, is not so simple. A primary challenge in the insurance industry, as in the finance industry, is to manage profit. Profit in the insurance industry might be loosely defined in terms of premiums received minus claims paid. This is not the complete definition, but it will serve for the purposes of a discussion of the requirement for meta-data. To perform profit analysis, the analyst must have access to profit (premium) and loss (claim) information. The revenue information is based on premium data in the premium file, and the claim data in the claim file. The challenge, then, is for information systems to deliver both sets of data in a reconciled form acceptable for merging into one single report. Figure 8 on page 57 shows the progression from data to information as referred to in the meta-data discussion.

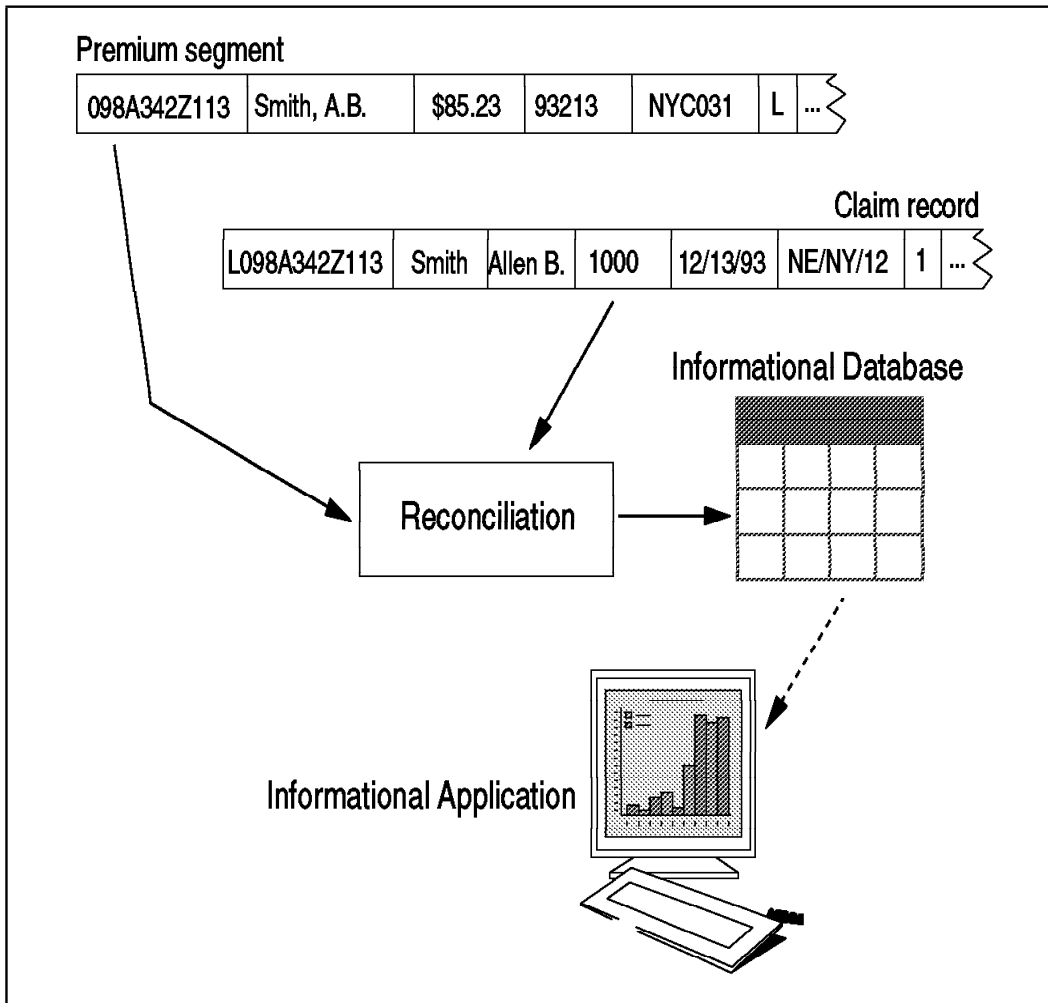


Figure 8. Data→Information→Chart

6.1 The Meta-data Requirement

The meta-data requirement is a prerequisite to reconciliation. One type of reconciliation is the manipulation of two or more sets of data so that inconsistencies between them are resolved. Understanding the individual sets of data is an obvious prerequisite to resolving these inconsistencies and meta-data is the vehicle for understanding the data. The record and segment represented in Figure 8 serve as an example for discussing meta-data

Meta-data explains our data

The two transaction systems, CLMPRCG and PREMPCG, are legacy systems crucial to the enterprise's business. The PREMPCG application was developed using IMS/DB as a data store, COBOL as the programming language, and IMS/TM as a transaction manager; the CLMPRCG application was developed using VSAM as a data store, COBOL as a programming language, and CICS/MVS as a transaction manager. They were also developed before the need for meta-data was recognized.

Legacy systems require meta-data creation

Meta-data can be elusive

The knowledge workers try to reconcile fields in the segment and the record that have the same business meaning but different data processing implementations. The programmers who wrote the applications and designed the record and segment are currently in other jobs, and the applications and record layouts are undocumented.

Go to the business analysts for lost meta-data

The knowledge administrators interview business analysts from the line of business to get a business perspective and thereby understand the applications and data. This understanding is now implemented electronically as meta-data and is managed electronically by DataGuide. This process is necessary rework because the original systems were not documented. Table 2 shows the meta-data for the Claim record's first field, called Policy_Number. The meta-data was generated as a result of the interviews with the business analysts.

Table 2. Meta-data Example: Policy Number in Claim Record	
Entity Name	Entity Description
Policy Number	Policy number is a 12-character field whose first character indicates the type of policy and the second, third, and fourth characters indicate the sales office.

The interviews also yielded meta-data for the premium segment's first field, Policy_num, as shown in Table 3.

Table 3. Meta-data Example: Policy Number in Premium Segment	
Entity Name	Entity Description
Policy Number	Policy number is a 12-character field whose first, second, and third characters indicate the sales office. The type of policy is indicated by the sixth character, called policy type.

Integrating disparate data requires reconciliation

It is clear that the merging of these data sources requires some sort of reconciliation and cleansing before the data can be stored in one informational database. Meta-data plays a key role in reconciling data. It provides the knowledge of what the data means and gives direction on reconciliation that is needed for informational analysis. There is also a need to store the newly identified meta-data for later use. This requirement is addressed in the Information Warehouse architecture through the Information Catalog function. For more information on the Information Catalog, see the following:

- *Information Warehouse Architecture I*
- *Information Warehouse Architecture and Information Catalog Overview*
- *Information Warehouse in the Retail Industry.*

6.2 Reconciliation

Reconciliation and merging are neither straightforward nor simple. Profit analysis might be desired at a policy level, group-of-policies level, sales level, or sales representative level. This unpredictable nature in how business analysts want data presented requires a flexible method for making the informational data available. Data replication is the discipline by which operational data can be delivered for MIS utilization in a flexible manner.

Reconciliation and merging are not simple

Once the information is delivered, the analyst must be able to visualize the data. Report formats, with rows and columns of data, are the traditional approach to visualizing and analyzing data. Recent developments have trended toward spreadsheets, pie charts, and a variety of other functions for presenting information to knowledge workers. The process of choosing a preferred presentation function is key and establishes a preferred, organizationwide set of information presentation tools for all business analysts.

Presentation trends toward graphics

6.3 Solution Components

This brief tutorial on the insurance environment is a prelude to the discussion of the insurance solution thread, product management. Product management includes at least two disciplines, as follows:

- Profit analysis of existing products
- New product-to-market activity.

We assume that both profit analysis and new product-to-market activity are based on the revenue and cost information described in the premiums and claims solution thread. That recognition allows us to focus on bringing the operational data to the business analyst and on that analyst's use of an informational application. A verification of these assumptions is based on the recent product-to-market experience of the long-term-care product. Development of this new type of insurance product utilized existing policies and actuarial information for the pricing of the product. The challenge is to make this data and information readily available without the lead-time required for information systems to develop a traditional extract program.

The final piece of the insurance solution thread is based on the complexity of delivering data to the knowledge worker. The claims and premium data is assumed to be designed and implemented for the needs of a high-performance transaction system. By inference, the data is not suitable for MIS consumption. The data, furthermore, requires multiple steps of manipulation to arrive at a state acceptable for MIS utilization. The requirement to manage the multistep transition process is incorporated under the workflow management discipline.

Data replication and data complexity are key factors

The chapters discussing the Information Warehouse architecture components are presented in an order that follows a reasonable Information Warehouse implementation project plan, as follows:

1. Organization asset data (OAD)
2. Data replication tools

3. Workflow management
4. Informational applications.

This order supports the data-driven approach to Information Warehouse implementation by first addressing the operational data sources and the processes that are necessary to turn the sources into informational data. This first step establishes the business process of turning operational data into informational data as requiring multiple data processing processes. The next discussion covers the data replication tools included in the tools component of the Information Warehouse architecture. The data replication tools implement the individual steps identified in the organization asset data step. Workflow management is then discussed as the discipline that manages the complexity of the multistep process. Finally, informational applications are discussed as the tools that are used by business analysts to access the informational data generated by the previous steps.

Chapter 7. Organization Asset Data

Organizations consider their data in all forms to be a vital asset of the enterprise. For historical and organizational reasons, very few enterprises have a master plan for managing their data. Data exists in databases and files and is identified only as data objects by their data processing technology name. There is little categorization of the data objects, no enterprise-level directory telling who the data belongs to, what it means from a business point of view, or what role it plays in the applications that manipulate it. The Information Warehouse architecture defines five categories of data with respect to how they are used by applications and specifies four configurations, or collections of the five categories. The categorization and configurations are a first step toward developing a data management system for the enterprise. DataGuide is a catalog for describing what data means from a business point of view. Information Warehouse architecture's *Organization Asset Data* component incorporates the categorization and configuration methodology.

IW architecture brings order to the data chaos

The organization asset data component, highlighted in Figure 9 on page 62, of the Information Warehouse architecture serves as a guideline for conceptually managing the raw data that is input to the product-to-market strategy.

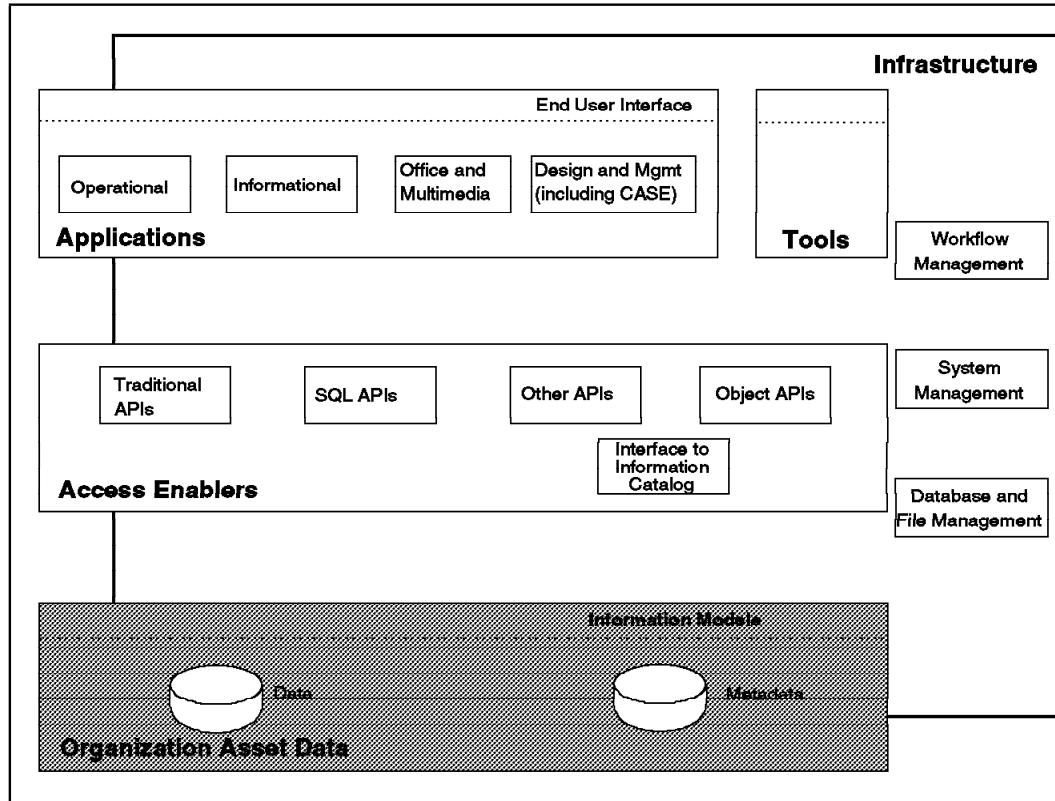


Figure 9. Organization Asset Data

An industry model presents the business view

The organization asset data is composed of two elements: data and meta-data. The business view of the data is achieved through the modeling process, whether it is informal or tool-based. The discussion in Appendix A, "Models and Modeling" on page 91 lays the groundwork for that view. The business analyst understands the enterprise's business along with the business objects and activities that contribute to that business. The model is a translation of that view to the data processing view of the data processing objects and processes that mirror the business objects and activities.

The OAD categories complement the industry model

The model is a way of managing the categories of data defined in *Information Warehouse Architecture I* and creating a communication path between business analysts and the information systems department staff. The model starts from a business perspective and progresses toward the implementation of an equivalent data processing perspective. The data categories in the Information Warehouse architecture are more geared toward how the data is used by applications. The Information Warehouse architecture view helps to make the data available to informational applications from the operational source from which it is extracted. The business model view provides the business semantics of the data. Both views are necessary to maintain an understanding and management control of the data.

The data categories defined in *Information Warehouse Architecture I* are as follows:

- Real-time data** is created and manipulated by operational applications to run the day-to-day business.
- Changed data** represents the changes that transactions make to the operational data.

Reconciled data is an informational copy of the operational data with basic conversions of codes, and resolution of inconsistencies of data stored in different parts of the enterprise.

Derived data is an aggregated version of the reconciled data.

Meta-data is descriptive data about the data in the other categories. It is used by knowledge workers searching for and trying to understand the data in those other categories. The meta-data is maintained in the Information Catalog.

7.1 Business View

The business analyst perspective on the input requirements uses the terminology “claims and premium” data. The data processing perspective uses the terminology DB2 tables and VSAM data sets. The IAA ties the business and data processing perspectives together. Use of the IAA would give both perspectives a consistent approach, and a documented methodology that could be reused later on. Variations that would later benefit from this approach were identified in terms of different demographic subsettings of the insuree by age, geographic location, and profession.

7.2 Mapping the Organization Asset Data

The first step in managing the organization asset data required mapping the data→information path to the data categories described in *Information Warehouse Architecture I*. This was done in two steps: describing the data→information path in business terms, then translating it to data processing terms.

The claims and premiums data generated and managed by the operational transaction systems was identified as the original source operational data, called real-time data in the data configurations of the Information Warehouse architecture. Coded fields in these records were recognized as essential for transaction performance but were seen as too esoteric for use by business analysts. These coded fields included a male/female indicator (1 or 2) and a policy type (whole life or term) indicator (L or T). These codes are to be translated to descriptive equivalents, and the resulting data would be stage 1 data. The translation of codes constitutes reconciliation in Information Warehouse architecture terms, so that stage 1 data is reconciled data in Information Warehouse architecture terms.

The stage 1 data would then be aggregated at the policy level, so that a paid out amount or taken in amount would represent sums for a given policy. This is stage 2 data in business terms and derived data in Information Warehouse architecture terms. The mapping of the business terms for the data to the Information Warehouse architecture terms is illustrated in Table 4 on page 64.

Coded fields are too esoteric for business knowledge workers

IW architecture data terms map to business terms for data

Table 4. Mapping Terms for Data	
Business Analyst	Information Warehouse Architecture
Original data (Claims and Premiums)	Real-time data
Stage 1	Reconciled data
Stage 2	Derived data

7.3 Managing the Organization Asset Data

The data configurations in the *Information Warehouse Architecture I* are guidelines for architecting the operational and informational data. The insurance enterprise recognized a need for real-time, changed, reconciled, derived, and meta-data; this led the insurance enterprise to a data configuration “D” environment. For a definition of data configuration “D,” see *Information Warehouse Architecture I*. The need for real-time data is obvious: this is the data that ran the claims and premiums processing systems. The need for changed data is more subtle.

Use changed data to recreate data at a historical point in time

The Information Warehouse architecture describes changed data as data that describes the change in the real-time due to a business transaction occurring. The changed data is a means of recreating a historical period of real-time. For example, changed data for the period January 1, 1990 through June 30, 1990 can be used starting with the before image of the first changed data image. The changed data images for the remaining period are applied until a period-in-time operational system is established. This system can be used for analysis as if it had been built by a live operational system. However, the claims processing application or the DBMS would have to be modified to generate the changed data. One implementation of changed data stores the before and after image of the real-time data in a log and saves the log as a history file.

Derived data is the basis for informational analysis

The need for derived data as a source for doing informational analysis is obvious, as most informational analysis is performed against common aggregations, such as daily, quarterly, branch, or geographical aggregations. It is the dynamic nature of decision support that justifies the reconciled data store. The knowledge workers may start out looking at reconciled—detailed—data. They perform informational analysis on this data and accumulate a portfolio of queries, reports, and charts that aggregate the detailed data to different levels on different attributes (columns in the table). If they all aggregated on the same column and to some common level, such as daily or by policy, the reconciled data could be kept only as an interim step toward the derived data. The common aggregation levels would be stored as the derived data. Reality indicated that this was not the case and that the detailed data was necessary to allow the variety of analyses.

Additionally, the knowledge workers saw a need to do prospecting analysis. This meant developing queries to discover trends in their business, rather than using the data to back up hypotheses on the trends. In this environment, the knowledge workers needed the reconciled data, and the freedom to perform unlimited analyses of the data. The requirement to keep reconciled data and derived data made configuration “D” the proper strategy for managing the organization asset data.

*Reconciled
plus derived
data means
configuration
“D”*

The next step is development of the data enhancement applications to make the transition from original data to stage 1 data to stage 2 data. The enhancement applications are called ORG2STG1 and STG1STG2. This analysis and establishment of terms are input to the discussion of workflow management in Chapter 9, “Workflow Management” on page 73. The requirement for enhancement and subsequent replication of data can be met through custom-written applications or through the use of off-the-shelf software. DataRefresher, a very attractive solution in the off-the-shelf category, is discussed in Chapter 8, “Data Replication Tools” on page 67.

Chapter 8. Data Replication Tools

Information Warehouse Architecture's tools component includes data replication tools (also known as copy tools) as well as other administration tools, such as Information Catalog administration tools. Data replication tools include a wide range of software functions used to manage the replication and enhancement of data from operational to informational data stores. These tools are one piece of the data replication strategy, which also includes workflow management, tool invocation, and a data staging interface. The data replication strategy has as its overall objective the automation of replication processes. Automation of replication processes, in turn, is the key to making an Information Warehouse solution feasible and reliable as a long-term component of the enterprise information systems department.

Data replication is the key to reliable informational systems

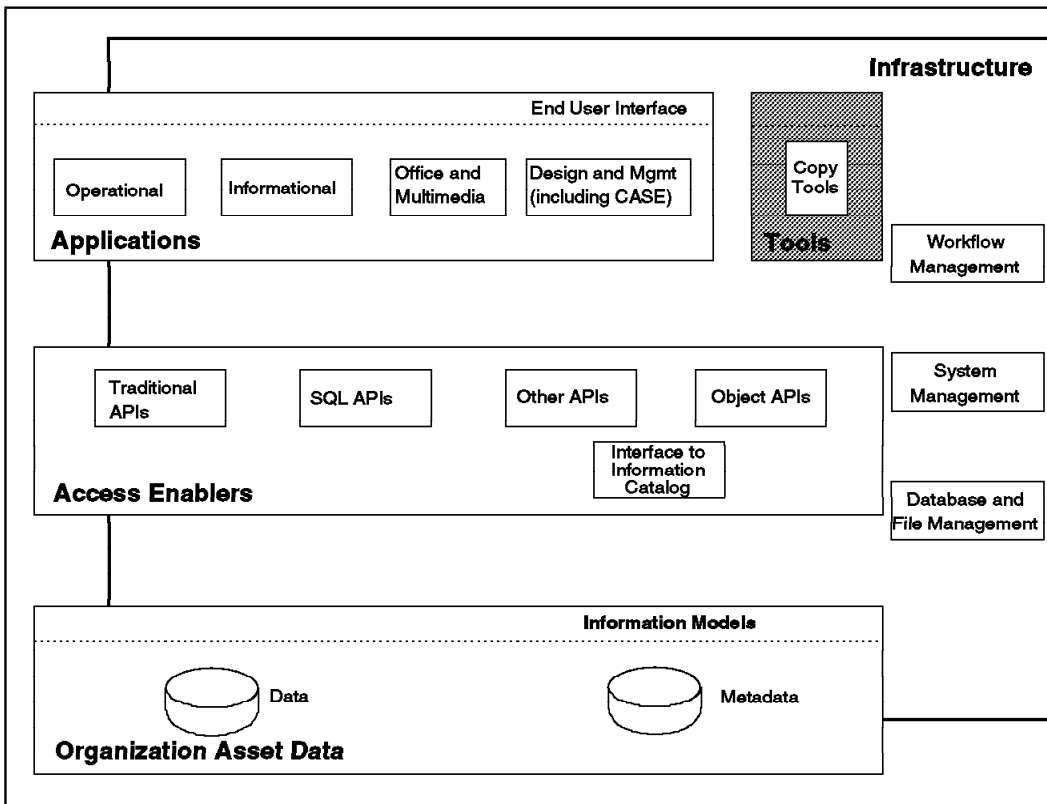


Figure 10. Copy Tools

Copy tools play in the path from operational to informational data

There are many copy tools available to the Information Warehouse implementer. They perform data extraction, conversion, cleansing, and transformation for the generation of informational data. They transport data between heterogeneous systems, load refresh data, and apply changed data in target informational datastores. Knowledge workers use informational applications against the informational data thus delivered to apply further analytical manipulation and present the information in graphic form. Informational applications overlap with data replication tools with respect to aggregation and derivation of data on its way to informational data.

8.1 Tool Selection

The insurance enterprise settled on the refresh propagation as its primary mode of data copy. This meant that its product-to-market strategy would copy point-in-time snapshots of claims data, and that selection or design of insurance products would be done against this data. It implied that update propagation was not necessary. The information systems department ruled out DataPropagator Relational and focused on DataHub and DataRefresher. For more information on update and refresh propagation, see *Information Warehouse in the Finance Industry*.

The information systems department established a need for DataHub as a single programmable-workstation-based point of control for its copy processes. It then faced the decision of using the copy function in DataHub or adding DataRefresher as an additional member of the copy tool suite. The decision was made that DataRefresher was necessary because it could extract from the VSAM files common to the insurance enterprise, whereas DataHub had only relational databases as sources. The enterprise recognized that tools may become available that allow DataHub to extract from other databases but chose to go with DataRefresher for the time being.

Information Warehouse architecture's data replication strategy is designed to provide organizations with advantages in the following areas:

- Synchronizing data between separate data stores
- Providing a central programmable-workstation-based point of administration for data replication across host and LAN databases
- Increasing the number of databases that can be accessed by data replication tools.

8.2 DataRefresher

DataRefresher is part of the family of tools that comprise IBM's data replication solution. DataRefresher provides the capabilities for copying, refining, and manipulating data from a source database or file on one system and formatting it for a target database or file on the same or different system. DataRefresher extracts data from one or more operational systems; the data is used to refresh the informational data. DataRefresher can extract data from multiple sources, combine the data, and send the resulting output to a single target system.

DataRefresher can extract data from a variety of systems and databases more easily than if individual applications were written to extract the data. The insurance enterprise saw value in DataRefresher to refresh the data stored in its:

- DB2 for MVS and DB2/2 databases with the data stored in its VSAM claims and premiums data sets
- Test databases with the data stored in production databases
- Corporate DB2 for MVS system with data from the several DB2 for MVS systems in the geographic data centers.

DataRefresher's ultimate goal in the insurance enterprise's Information Warehouse data replication strategy is to support the decision making process within the enterprise. It accomplishes this goal by making data available to informational application products. For example, Query Management Facility, Visualizer, Personal AS/2, and Application System are informational applications that massage and present informational data made available by DataRefresher. DataRefresher can be integrated with DataPropagator NonRelational and DataPropagator Relational and the data copies they produce.

8.2.1 Operating Environment

DataRefresher provides facilities for accessing the data stored in source systems and databases from the OS/2 and MVS operating environments.

8.2.1.1 OS/2

DataRefresher operating in an OS/2 environment provides a GUI to assist in identifying data sources and target databases, and in creating refresh requests. A refresh is a request to extract data from a data source and place it in a target system. The GUI was structured to match the skill level and frequency of use with the performance of a task. It reduces the skill level required to extract data from a data source and format the data for a target database. The more complex tasks of identifying the data sources and targets on the host system are separated from the less technical task of creating the extract.

*DataRefresher
provides a
GUI*

DataRefresher is consistent with the workstation-based strategy

DataRefresher, as part of Information Warehouse architecture's data replication solution, can be invoked from DataHub; this is a new feature of DataRefresher. This feature supports the strategy that makes the programmable workstation a single point of control for data replication. With DataRefresher, DataHub users can administer the nonrelational sources supported by DataRefresher. This expands the capabilities of DataHub, making it a single data replication control point for both relational and nonrelational databases.

8.2.1.2 MVS

DataRefresher operating under MVS provides the following set of facilities for generating and processing extracts on the full range of MVS DataRefresher data sources:

Structures Access Program: The Structures Access Program (SAP) lets the DataRefresher administrator generate data description statements and extract requests based on stored user-specified data structures. A data description is a description of a particular data source, which is to be used in an extract. The SAP creates the data descriptions and extract requests for IMS databases, VSAM files, and physical sequential files.

Dictionary Access Program: The Dictionary Access Program (DAP) is used in conjunction with the OS/VS DB/DC Data Dictionary* to create data descriptions. DAP is a program which uses the Program Access Facility of the Data Dictionary to generate the data descriptions for IMS databases and IMS program specification blocks.

Online DataRefresher Commands: DataRefresher provides a series of commands to let you create data descriptions and extract requests online.

DataRefresher Dialogs: DataRefresher provides two types of dialogs for generating data descriptions and extract requests based on a series of supplied models: application development and end-user dialogs. Application development dialogs are intended for a database administrator. They let the administrator build and maintain data descriptions, extract requests, and JCL/JCS models and files. End user dialogs help the less experienced DataRefresher user to build and submit extract requests.

System Editor: The system editor can be used to create data descriptions and extract requests based on the models supplied with DataRefresher.

8.3 Data Sources and Targets

Figure 11 on page 71 provides an overview of the different types of data sources that can be used with DataRefresher. It also shows the different types of target systems to which the extracted data can be sent.

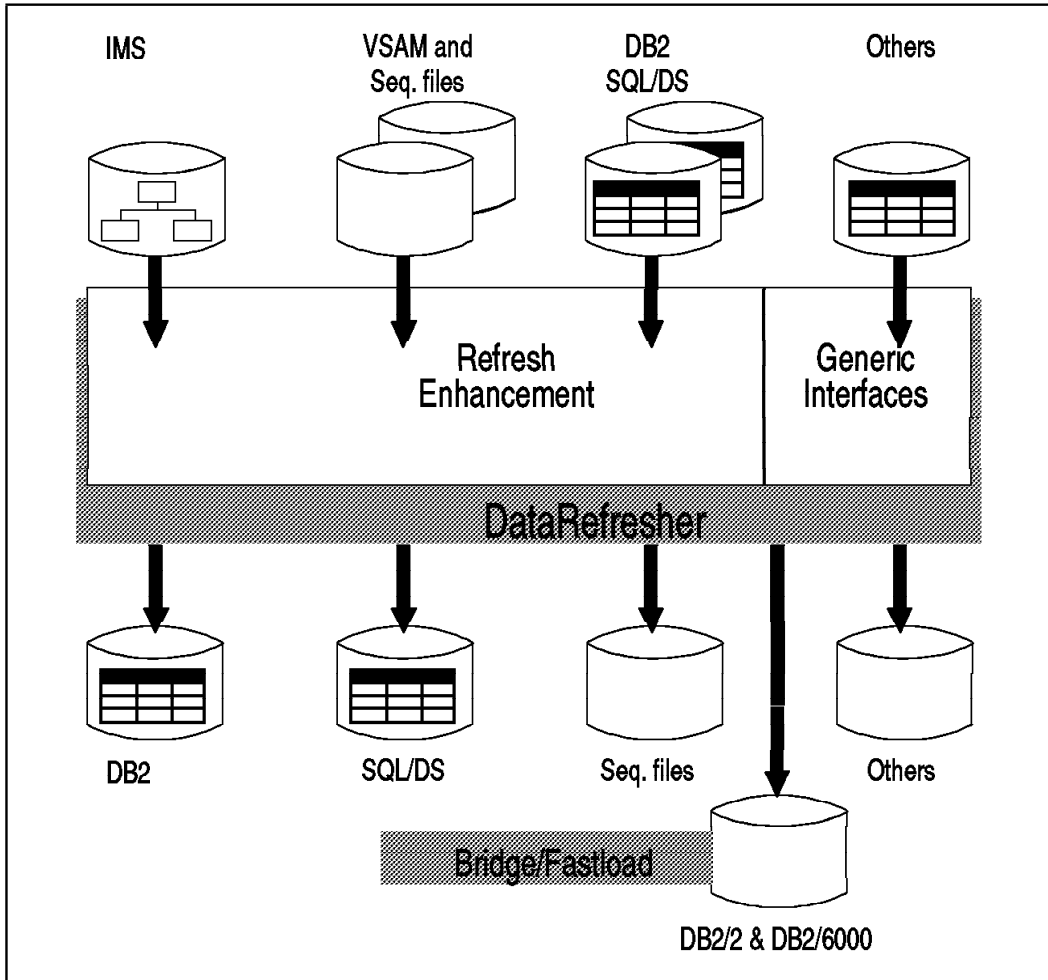


Figure 11. DataRefresher Sources and Targets

8.3.1 Data Sources

DataRefresher provides you with facilities for extracting data from any of the data sources shown in Figure 11:

- IMS databases
- VSAM data sets and physical sequential data sets (PSDSs)
- DB2 for MVS and SQL/DS for VM.

DataRefresher also provides you with facilities for registering exit routines, which are to be used by the extract routine to extend the range of data sources or format the extracted data for a target system. The following types of exit routines are used to extend the range of data sources:

Generic data interface (GDI) exit routines

GDI exit routines can be developed to access data that is stored in a data source which is not directly supported by DataRefresher. A GDI exit makes it possible to extract data from Integration Exchange Format (IXF) files, and from other IBM and non-IBM sources which are not supported by DataRefresher.

GDI exits are different from the other types of user exits, as DataRefresher does not access the source data; it is accessed by the routine.

Map capture exit routines

Map capture exit routines can be developed so that the DataRefresher definition and extract information for a specific extract can be saved, or used by DataPropagator Relational for data propagation. For more information about DataPropagator Relational see *Information Warehouse in the Finance Industry* or the relevant product documentation.

DataRefresher provides sample exit routines and interface control blocks for each exit routine.

8.3.2 Target Systems

DataRefresher can format extracted data for DB2 tables and PSDSs (including those in IXF) in MVS systems. DataRefresher can also be used with vendor supplied products, such as Bridge Technology, Inc.'s Bridge/Fastload** product, to support DB2/2 and DB2/6000 tables as targets. Support is provided for other IBM and non-IBM targets by the generic output interface (GOI) exit routine. GOI exit routines provide the flexibility to make changes to the extracted data, ensuring that the correct data is received by the target system. A GOI exit can perform the following operations:

- Format the data for a target system that is unsupported by DataRefresher
- Summarize the extracted data
- Provide totals for the extracted data
- Make changes to the extracted data; for example, adding or removing data
- Perform complex data enhancement.

The following exit routine interfaces are also supplied with DataRefresher and can be used to make enhancements to the extracted data before it is written to a target system:

Data exit routines

Data exit routines can be developed to let you make selected changes to the format of any extracted data, before DataRefresher loads it to a target.

Date/time conversion exit routines

Date/Time conversion exits can be developed to change data and time fields to the International Organization for Standardization (ISO) format used by DataRefresher.

User data type exit routines

User data type exits can be developed to convert data types to a format that DataRefresher supports.

Besides the exit routine interfaces listed above, and in 8.3.1, "Data Sources," DataRefresher provides an accounting exit interface. This interface can be used to track the resources used by the data extract manager (DEM). DataRefresher provides sample exit routines and interface control blocks for all of the exit routine interfaces supplied with the product.

Chapter 9. Workflow Management

Workflow management (WFM) is the technology chosen to manage the data→information flow requirement discussed in Chapter 7, “Organization Asset Data” on page 61 and Chapter 3, “Insurance Industry Business Requirements” on page 19. It is shown in Figure 12 to be part of the infrastructure of the Information Warehouse architecture.

WFM to control data flow

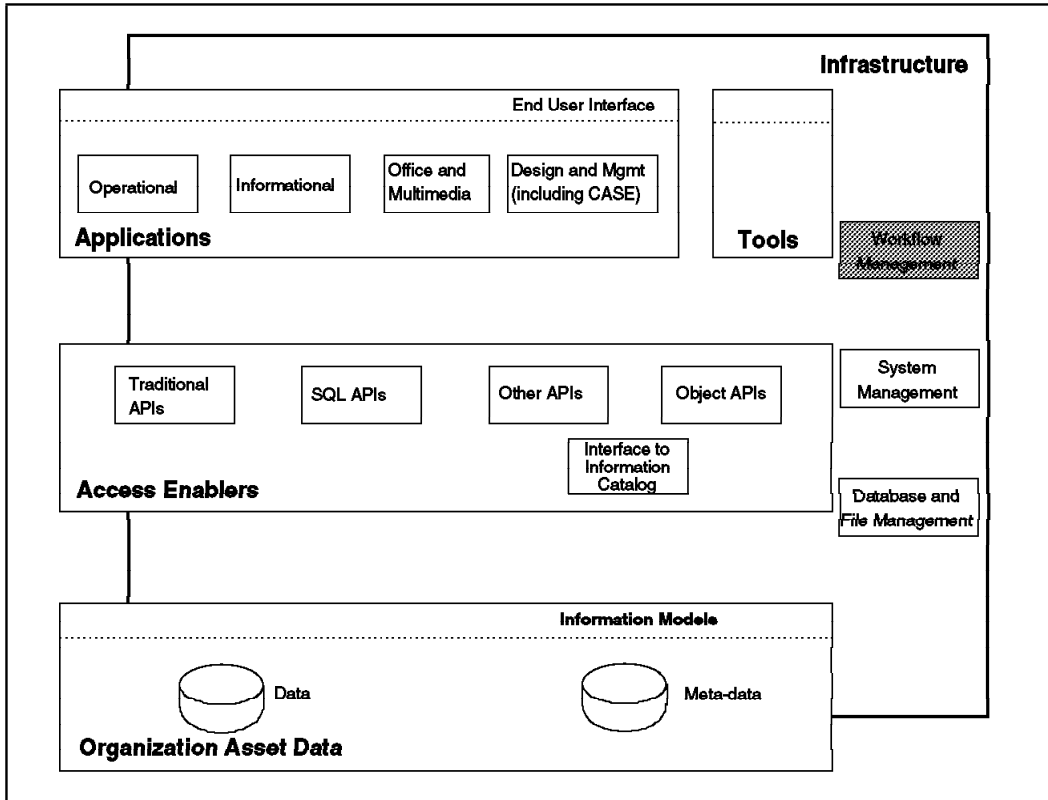


Figure 12. Workflow Management

WFM is a successor to traditional job control and job scheduling. Job control meant the conditional execution of individual job steps in a larger stream of jobs, performed by conditional control syntax on the individual job steps. Job scheduling was the management of data processing tasks based on time of day or other event and was performed by software operating at a higher level than the tasks themselves. It tended to be linear, in that each job in a jobstream followed from the one before it. This made it difficult to manage more complicated job mixes. Job control and job scheduling focused on

mainframe applications and were administrated using mainframe definitions. Workflow management improves on this traditional approach by being model-based and administered at the programmable workstation, and by recognizing that workflow includes activities that are performed by people as well as data processing applications.

Workflow Management focuses on integrating business processes

9.1 Workflow Management Technology

Information Warehouse architecture workflow management technology specifies the GUI for defining workflows from the highest level down to the individual processes. The process designers can keep their focus on the overall workflow and address the detailed processes as a separate step. The workflow management provides easy ways for representing activities and decision points, with the workflow management tool providing the control of the overall process. In summary, the requirements for workflow management function are as follows:

- Define, graphically depict, and document models of enterprise processes
- Designate the staff or application program required for the activities in the process
- Track the process and the status of tasks assigned to staff members.

WFM has two types of user groups

9.1.1 Workflow Management Users

Workflow management has two communities of users: workflow modelers and administrators, and workflow management users. This view implies software with two discrete sets of functions to serve these two groups. We use the term “build time” for the modelers and administrators and the term “run time” for the users. We assume a client-server environment, wherein modeling and administration of the process models are performed on a programmable workstation. The programmable workstation may be the server as well, for providing process services to the requester of process services. Alternatively, the process services’ server may be in another programmable workstation. The client-server environment also applies to the execution of the programs defined in the activities. That is, the process model may be defined on a programmable workstation, but it may initiate an application on a remote platform.

9.1.2 Terminology

Understanding terminology is important to managing workflow. This is particularly true because of the nested nature of units of work within larger units of work, or business processes. The terminology presented describes basic building blocks for managing workflows.

Process

A process is a sequence of activities that must be completed over time to accomplish a given piece of work. A process defines how work progresses from one activity to the next and whether each activity is performed by a person or by a data processing applica-

tion. Multiple instances of a process can run in parallel.

Workflow Model

A workflow model is a graphical representation of a process. Whereas a process uses words to describe a sequence of working procedures, a workflow model uses pictures to represent the activities and subprocesses that make up the process. The possible ways that work and data can flow through a model are represented graphically by arrows called connectors. A workflow model might contain processes, activities, programs, staff, connectors, containers for data, and conditions.

Activity

An activity is a step within a process. This step is a unit of work to be done by a person or a program. Information associated with an activity includes the following:

- The person or program, or both, responsible for completing the activity
- Data required to complete the activity
- Conditions to be fulfilled *before* the activity can start
- Conditions that indicate activity completion
- The time by which the activity must complete
- Escalation procedures.

This definition of an activity underscores the integration of data processing with human resources. An activity is defined as work that may be done by a person or by a data processing application. Workflow management software understands this and makes provision for invoking a data processing application or by posting a tickler for a person to perform some manual work. After completing the activity, the person executes an electronic function to inform the workflow management software of activity completion. Connectors are used to define the sequence of activities and transmission of data between activities.

Program

A program is a computer-based function or application that supports the work to be done in an activity. The program can perform the work automatically, for example, by starting a print job. Alternatively, it can enable a manual activity, for example, bringing up an online telephone directory so that a person can make a telephone call.

Staff

People can be assigned to perform activities in the process.

Connector

Connectors link activities in the workflow model. They control the sequence of activities and flow of data between activities. The activity from which the connector originates is the source activity. The destination of the connector is the target activity.

Containers for data Data is passed between activities in the process.

Conditions The flow of data and control along connectors is controlled by conditions. If these conditions are evaluated to be true, the flow of the process can continue. Conditions are means of enforcing rules and controls on a process.

9.1.3 Information Warehouse Framework Integration

FlowMark integrates with data replication tools

Information Warehouse Architecture I defines an interface to workflow management. It defines how users interact with tool and workflow management user interfaces and what copy tools must pass to the workflow management tool. FlowMark is integrated with DataHub through the key interfaces identified in *Information Warehouse Architecture I*. This integration allows the data replication administrator to define the complex business process in FlowMark. Each step in the business process is defined as an interaction with DataHub. DataHub invokes the copy tool to execute the necessary data replication activity. This integration allows the data replication administrator to separate the business process, the data processing activities of the process, and the execution of the data processing steps.

9.1.4 Workflow Management Coalition

Establish a standard for WFM

The Workflow Management Coalition is a consortium of major international WFM software vendors and WFM software exploiters whose aim is to develop specifications for software that will allow different WFM products to interoperate in various key areas. IBM is a founding member of the Workflow Management Coalition.

Today's software developers, businesses, and governments encounter difficulties in the management of processes when WFM software from multiple vendors is used. This is due primarily to the lack of specifications that describe how interoperation of that software should occur. The situation is further complicated because the business personnel must know multiple WFM product end-user interfaces and associated function.

The coalition will develop and promote specifications that, when implemented in multiple WFM software, will save significant time and effort for WFM software exploiters. The WFM areas to be addressed by the Coalition include:

- APIs which will allow a consistent method of access to WFM functions by application programs. When supported by multiple WFM vendors, these APIs will allow a WFM exploiter the capability of having a consistent WFM end user interface and front-end function even when multiple WFM products are used in a business.
- Specifications for formats and protocols to flow between WFM products which will provide interoperability between those products. This will allow a business to construct and execute single business processes that span multiple WFM products. In addition to providing interoperation definitions between WFM product engines, these formats and protocols

are intended to operate between WFM clients and WFM servers. This provides the capability of mixing WFM clients and WFM servers.

- Specifications for formats and protocols to flow between WFM products and the application invoking systems. This will provide the capability of executing work (software) across heterogeneous environments.
- WFM model interchange specifications which will allow process definitions created on one product to be interchanged and used to execute a process definition on another product.

Through WFM vendor implementation of the specifications, enterprises will benefit by reusing applications across WFM products and by having the ability to seamlessly connect different WFM products together for the purpose of managing cross company workflows.

9.2 FlowMark

FlowMark* is a workflow management tool that can be used to optimize business processes. It assists in the definition, refinement, documentation, and control of business processes. It shifts the focus of the enterprise to the business, while FlowMark operates the data processing processes. The results are:

*FlowMark
helps manage
business
processes*

- Better products and services
- Higher productivity through automation
- Easier interaction between people and computers.

FlowMark assists in the daily operations of the enterprise, in planning and management, and in the design of applications tailored to the business. FlowMark supports the following workflow management functions:

- Definition and documentation of business processes
- Testing of processes
- Execution of processes to:
 - Support the people doing the work
 - Fully automate activities that do not require human guidance
- Refine and update processes as the business changes.

9.2.1 Definition and Documentation of Processes

FlowMark combines client-server computing with two standard business concepts: to-do lists and procedures manuals. To-do lists become online work lists in FlowMark, and procedures manuals become online process models.

Define business processes graphically, easily

The process model is built by drawing it as a graph, using an advanced design tool that is part of FlowMark. The graph depicts business activities, the staff that performs them, the programs that support the people, and the flow of control and information between the activities. FlowMark stores the modeling information in an ObjectStore** database. The model—the diagram—can be documented with supporting online documentation. FlowMark includes facilities for printing process models.

9.2.2 Testing Processes

Testing the processes makes for better management

FlowMark provides an animation tool that steps through the process model, simulating real conditions. Bottlenecks, loops, and dead spots can be identified and fixed before processes are run.

9.2.3 Running Processes

FlowMark can execute the final process model. Activities to perform appear on the FlowMark work lists of assigned staff members. When a staff member selects an activity, the program that supports the work is started automatically, with the necessary information available. The work lists give the staff a continuously updated overview of their pending activities. As workflow management is implemented, people will use work lists as their primary entry to the computing environment.

9.3 Insurance Solution Thread

FlowMark manages data replication

FlowMark manages the two applications, ORIGSTG1 and STG1STG2, used for massaging claims and premium data. In Figure 13 on page 79, these applications are labeled CORGST1 and CST1ST2 for massaging claims data and PORGST1 and PST1ST2 for premiums. In a graphic manner, the workflow management tool defines the process of transforming the original data into derived data. This process is a network of two sets of the ORIGSTG1 and STG1STG2 applications: one for Claims and one for Premiums. Figure 13 on page 79 shows the network made up of applications to massage Claims data and Premiums data, arriving at information which can be used by decision support tools to perform MIS analysis. FlowMark provides the environment within which to define a business process in terms of the ordered execution of the data replication tools. FlowMark interfaces with DataHub to execute each tool for the corresponding step defined in the business process. It checks completion codes from each step and manages the conditional execution of subsequent steps.

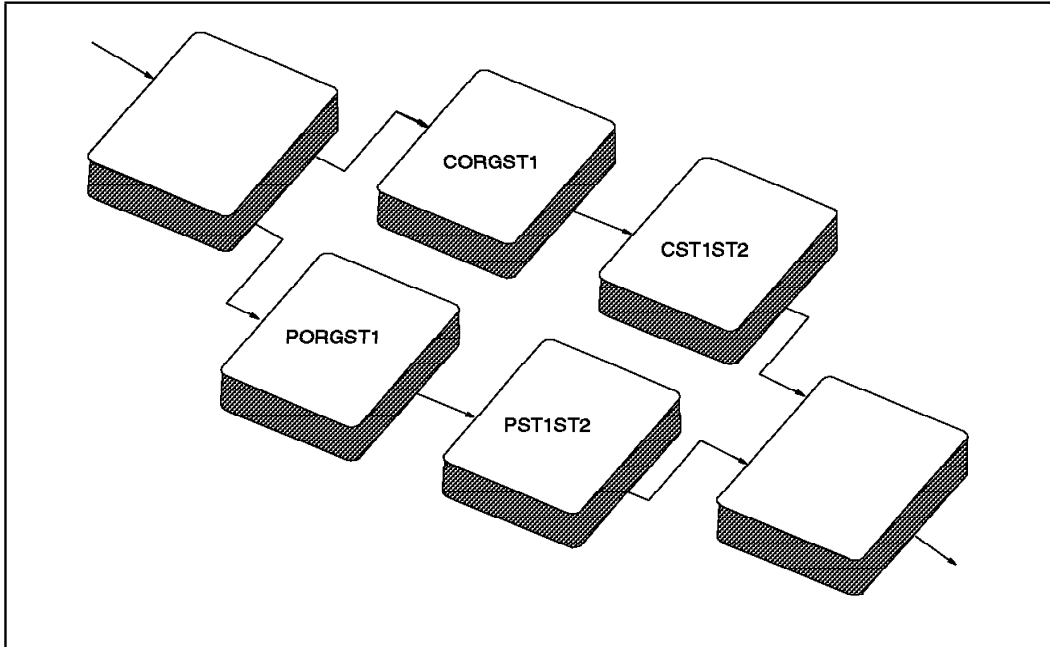


Figure 13. Data→Information

Chapter 10. Applications

In this chapter, we review the applications component of the Information Warehouse architecture. The objective is to establish a strategy and direction for the end-user set of tools used to analyze data delivered by the copy tools. Figure 14 highlights the applications component of the Information Warehouse architecture. *Information Warehouse Architecture I* focuses on one of the classes of applications—informational applications.

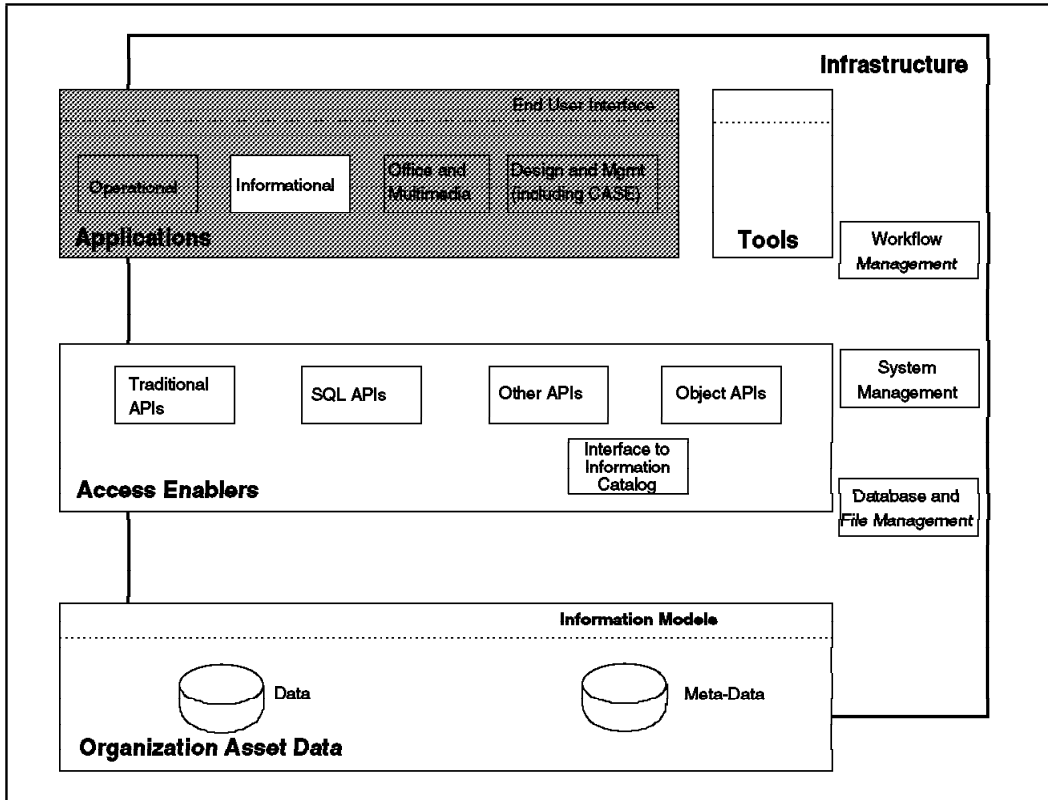


Figure 14. Applications

Informational applications includes both informational applications and data enhancement—aggregation and derivation—function. Tools specifically designed for data replication and enhancement for use at the work group level or above—where informational data is shared—are addressed in the tools component of the Information Warehouse architecture.

Informational applications are used by knowledge workers

Informational applications are those used directly by knowledge workers for information analysis; they enhance data for the needs of the individual knowledge worker. Enhancement algorithms may be elevated from the informational application to the data replication tools if they are seen to be common across knowledge workers. In the insurance solution thread, these are the tools used to analyze existing historical claims and premiums data and actuarial data. The objective of this analysis is “what-if” in nature, as it is being used to propose and analyze new products (see Chapter 6, “Insurance Solution Thread Overview” on page 55).

10.1 Knowledge Worker Requirements

Knowledge workers need DSS tools and good information

The knowledge worker depends on information and decision support software (DSS) tools to analyze and present that information. Accurate, consistent, and timely information, made available through data replication, is key to the successful use of DSS tools. DSS tools facilitate effective information analysis in support of the operational and strategic decisions that are essential for business success. Strategic decisions include business process reengineering which leads to redesigned operational systems, which in turn lead to new kinds of informational data.

Knowledge workers need to find the information

Knowledge workers have been denied access to important informational data for two reasons: the informational data it is not in a form suitable for their decision support tools and the knowledge workers lack a guide to the availability and currency of the data. Data replication addresses the issue of informational data format being incompatible with decision support tools; DataGuide addresses the requirement for a guide to the availability and currency of the informational data.

Locate and launch is the best operating mode

DataGuide provides a locate and launch environment that integrates the locate process with the decision support tool execution process. The knowledge workers begin their knowledge task with DataGuide. They use business terminology through the keyword search or navigational modes supported by the knowledge worker end-user interface to find the desired informational object.

The knowledge worker also needs access to a range of functions to analyze and present accessible data. Critical function is often lacking because it is contained in a product that is not compatible with function that has already been selected. End-user application creation, for tailoring decision support function, is necessary for the more repetitive types of analysis and presentation.

10.2 End-User Perspective

The decision support strategy makes two basic premises about knowledge workers, as follows:

- They are workstation-based
- Their focus is on objects, not applications.

Focus on objects, not applications

The workstation is the platform for future decision support. In this environment, decision support tools run on the workstation and access information anywhere in the environment. This requires the support and integration with the access enablers component of the Information Warehouse architecture. The access enablers component specifies SQL as the data access language to access information in relational DBMSs. It also specifies DRDA as the protocol to use when accessing relational DBMSs over a wide area network.

Knowledge workers need objects to perform their decision making tasks, they do not need applications. Obviously, applications are used to manipulate information and present them as informational objects—for instance, charts, query results, reports. However, knowledge workers want to think in terms of these business objects rather than the applications that generate the business object. The direction for the informational applications component of the Information Warehouse architecture is to interact with the knowledge worker based on the object, rather than the application.

Knowledge workers think in terms of objects

The desktop is the knowledge workers' platform for performing their tasks. The informational application strategy is to fill the desktop with icons that represent the informational objects that the knowledge workers need. Knowledge workers can then populate the desktop with the objects and group them in folders according to their business taskload. This orientation toward objects and object groupings is more task-oriented than one that focuses on software tools.

Desktop contains objects, not applications

10.3 Informational Application Function

Informational applications or DSS tools include a broad range of decision-assisting facilities—including query formulation and results presentation—for managers and professionals. It also provides specialist functions such as project management and statistics.

Other leading examples of functional areas include executive support (commonly called Executive Information Systems—EIS) and spreadsheet applications. The following are some typical examples of DSS uses and users:

- Ad hoc queries—by sales managers or auditors or database administrator
- Running predefined queries—by sales analysis departments
- Producing presentation materials—by planners or marketers
- Optimization—by transport or process managers
- Statistical analysis—by brand managers or forecasters or quality controllers
- Budget tracking and analysis—by finance departments or product managers

- Project management—by information systems managers or construction departments
- Concise and consistent views of the company—by managers or executives.

Decision support can assist many types of users, including but not limited to those above, to compare, analyze, assess impacts, determine alternatives, optimize profit, and minimize cost through the manipulation and presentation of data.

EIS products cover a segment of the whole DSS area. They generally include a specialized end user interface suitable for casual users (especially executives) and are used to present information output by DSS, external news and information sources and business applications. They are generally “read only” in that they are used to review information provided by others and not for the generation of new information.

Increasingly, they include “what if” as well as “what is” capabilities and are becoming used by managerial-level users as well as executives.

10.4 Client-Server Operation

Knowledge workers normally perform their data analysis on a programmable workstation. This implies that they execute the information location function on the programmable workstation. DataGuide/2, the solution to the information location requirement, provides the GUI on the programmable workstation to meet this requirement. This configuration leads to a question about where the decision support tool executes.

Informational applications may be on any platform

DataGuide/2 has the ability to launch—from the programmable workstation—informational applications that support command-line interface invocation. It also can lead to the invocation of informational applications on remote servers but depends on client-server function in the informational application. For example, AS can be launched on a remote server by using the client-server function in AS Version 3. AS Version 3 is invoked in a client-server configuration by Personal AS/2 running on the programmable workstation. The flow for the knowledge worker invoking remote decision support function has the knowledge worker using DataGuide/2 and requesting an informational object to be displayed. DataGuide/2 invokes the programmable workstation component of the decision support software, which invokes the remote server component of the decision support software. The resultant informational object—for example, a chart or image—is returned to the knowledge worker in a desktop window.

The focus, then, is on the informational applications. A variety of informational applications in the enterprise were being used to do simple informational analysis, spreadsheets, reports, and query execution. The variety of end-user interfaces causes considerable waste in the end-user training budget. Redundant resource was being expended in teaching end users how to use different tools that performed basically the same function. This redundancy was narrowed down to the differences in end-user interface between the tools. Work groups using client-server environments are purchasing informational applications on their own without regard to consistency across the applications. A disciplined approach to informational applications is the

key to maximum benefit from those tools. The objectives of the disciplined approach are to:

- Achieve the consistency and interoperability of IBM and vendor informational applications
- Enable end users to select required DSS/EIS functions from IBM or other vendors on a “mix or match” basis
- Build a modular approach to providing decision support function to the end user
- Provide a competitive set of tools for the knowledge workers of the enterprise, whether they be executives or support staff, managers or business professionals.
- Provide analysis and presentation functions that enrich the value of information. Examples of analysis include statistics or business modeling; presentation functions range from simple reports to sophisticated business graphics.

The enterprise demands greater business value from the organization asset data investment. This value had to be extended to all knowledge workers—decision makers, business analysts, and brand managers—rather than a select subset of executives or high-level managers. The interconnection of knowledge workers to the organization asset data would be based on the untapped computing power of their intelligent workstations. The delivery of the information to the workstations is addressed separately as part of the data replication strategy.

The value to the enterprise is a greater return on existing information systems investment and a more rapid and efficient implementation of new end-user solutions for the knowledge worker.

10.5 Visualizer

In this section, we discuss Visualizer as a solution to the informational application component of the Information Warehouse architecture. Visualizer is intended for users who require desktop access to data from a variety of sources. It provides the user with the ability to manipulate and present data from these sources through the use of desktop objects.

Visualizer can access DB2/2, DB2/6000, DB2 for MVS, SQL/DS for VM, and DB2/400 through the support of the OS/2 client of DB2/2, DDCS/2 or Client Application Enabler. These data sources are represented as the database table object in Visualizer. Visualizer also supports a file structure for purposes such as the retention of intermediate results, for copies of database tables for mobile users, or for entering ad hoc data for analysis.

Visualizer is an informational application and a strategy

Visualizer supports client-server access of relational data

Visualizer's presentation is objects

Visualizer presents itself to knowledge workers as a series of objects that are fully integrated with the OS/2 desktop. Users have seamless access to remote databases and the ability to create, delete, and modify tables and views. The object approach is fully in keeping with the need of knowledge workers to think and work in terms of work objects rather than the applications that manipulate them.

Visualizer is optimized for ease of use

Visualizer has also focused on the knowledge worker's end-user interface. All objects have a standard interface with which the knowledge worker can manipulate the object. Each tool on the Visualizer tool bar has an accompanying function description that can be seen by a single right click of the mouse button. Both the tool icon and the description were chosen to be as simple and as obvious as possible.

10.5.1 Visualizer Objects

Visualizer provides a series of objects that represent the basic work activities of its users. The objects and their descriptions are as follows:

Database folder	A database folder is associated with a real database and is used to bring the contents of the database to the desktop.
Database table	A database table represents a table within a database and can be used to view the definition or the contents of the table. It can also be used to define a new database table.
Database view	A database view represents a view defined with the database and can be used to view the definition or the result of the view. It can also be used to define a database view.
Visualizer table	The Visualizer table object represents a table held in the Visualizer file system. It can be used to view the definition or the contents of the table. It can also be used to define a new table.
SQL statement	The SQL statement object corresponds to an object in the filing system that contains SQL language statements. These statements can be typed, edited, executed, and stored for reuse.
Query	The query icon represents an object that can be used to define queries on database tables and views. Visualizer tables are also supported by this object.
Report	The report icon represents the object used to format data for text base display or printing. It is used to manipulate report headings, column headings, row ordering, and other standard aspects of report appearance.
Form	This icon represents the object used to define a form attached to a database and to support the update of the database through the form.

The objects have inherent relationships between them: A report can be applied to a query and a query is applied to a database table object. Interactions between objects are well-defined and conform to drag-and-drop protocols. Visualizer also supports dynamic data exchange (DDE).

10.5.2 Visualizer Templates

Visualizer informational objects are created from templates. The templates can be copied and then modified to represent the instance of that object. Each template has settings that can be used to manage the object in a consistent manner.

10.6 Information Warehouse Integration

Visualizer is a solution to the informational application component of the Information Warehouse architecture. It can be invoked by DataGuide/2 to bring together the search and launching of informational objects. Visualizer can access relational tables and Visualizer tables. These data objects can be the target of Information Warehouse framework data replication tools. This integration can be viewed from the knowledge worker's perspective, as in the following process:

1. Start DataGuide/2
2. Search for Reports
3. Select Profit Report 1993
4. Start Program (DataGuide/2 function)
5. View report.

This scenario assumes that a knowledge administrator has defined the report using Visualizer definition function. The knowledge administrator has also defined Visualizer as a program that can be started by DataGuide/2. For more information on DataGuide/2, see *Information Warehouse in the Retail Industry* or the appropriate DataGuide/2 documentation.

The Information Warehouse framework integration can also be viewed from the data replication administrator's perspective. The profit report is based on data existing in the insurance enterprise in the Claims and Premium databases. Data replication administrators have analyzed the requirements for turning these operational sources into informational data. They have selected the appropriate data replication tools to extract and manipulate the data according to these requirements. They have also used FlowMark to define the multiple steps necessary to perform the business process. The data replication tools replicate the data into the DB2/2 target, from which Visualizer can present the information as a report.

Visualizer is part of the IW framework

IW framework integrates with data replication functions

Chapter 11. Conclusions

The insurance industry has been deeply affected by global business conditions and innovative use of information technology. The outlook for the future suggests more competitive pressures and a particular push toward specialization. Specialization implies choosing the most profitable products as a focus for the enterprise. Informational analysis of operational data plays a key role in that process and the Information Warehouse framework is a guide toward building informational analysis systems. It defines, in a generalized way, applications, access enablers, organization asset data, tools, and infrastructure to support that building process. It defines interfaces and data representation formats for openness and integration.

Strategies address a multitude of corporate needs

Key information systems emerge as critical to meeting insurance business objectives in the future. They are:

- Client information
- Claims management
- Product development
- Distribution management.

Building these systems challenges the present technology. These systems need an enterprise view of data, which requires integration and reconciliation of operational data across lines of business. The key systems are complex, involve large data volumes, and incur high development costs. Deployment is a major effort, possibly requiring a new technology, and may take many years to accomplish.

It is doubtful whether these systems can be built, deployed, and maintained at a reasonable cost, without a generalized, architected way of gaining access to information. The generalized approach calls for use of architectures, models, standards, and interfaces at both the enterprise and industry levels. The Information Warehouse architecture offers a structured approach to informational data access, and IAA offers it at the industry level.

A structured approach is necessary

At the enterprise level, this structured approach calls for understanding and organizing the organization asset data; building a master plan for data and information. It also implies a careful assessment of processes and source data needed to support the systems deployed for information delivery.

IAA is an industry specific approach to strategic solutions

The IAA information model offers the model of the data and business activities at the industry level. It provides a mapping into the enterprise's information technology and attempts to shield business logic from changes in rapidly changing technology. Use of the basic model with customization can save individual organizations time and money and help achieve data and systems integration. The question of how far to proceed with the model implementation, when to control the actual data representation (for example, database formats and program copybooks), has been deliberately left open and needs to be decided by individual organizations.

IW architecture is the generalized approach for information delivery

The Information Warehouse architecture satisfies the industry requirement for a generalized, structured approach for data access. Information is context dependent; it is more a process of being informed, than passive data storage or presentation. Key in the process is the end user environment, where data can be massaged and correlated. Today, workstation and work group environments are best suited for information acquisition. The Information Warehouse architecture defines standards and interfaces in this environment.

Openness is key to a strategic solution

The insurance industry needs open standards to facilitate interoperability with respect to diverse platforms, while keeping location and data representation transparent to the knowledge worker. Information Warehouse architecture gives the basis for bringing together products and tools from many vendors. For the insurance enterprise, it reduces the cost of gaining information, accommodates diverse requirements within one generic solution, and simplifies systems management.

Workflow Management is crucial in a complex business

The insurance industry is becoming more complex, not less so. This complexity is evident in the number of steps and the complexity of each step that make up what business analysts see as a single activity. Data replication has this characteristic of complexity, in the heterogeneity of operational data sources, and the number of sources necessary to build a single informational data object. The business complexity and data replication complexity can both be addressed through a workflow management product like FlowMark.

Data replication helps meet needs for diverse information

Meeting diverse information needs requires an effective data replication strategy. An effective data replication strategy has provisions for automated copying, a single workstation point-of-control, and both full and differential refresh. New Information Warehouse products that contribute to those strategy provisions include DataPropagator Relational, DataHub, and DataRefresher.

Desktop access

Information is context and people dependent. The most productive environment for information acquisition, for knowledge, is the programmable workstation running decision support software, with transparent access to a wide variety of data stores and navigation assistance through that maze of data stores. This is the direction of the Information Warehouse framework and its set of products.

Appendix A. Models and Modeling

The solution presented in this book devotes significant attention to models and modeling. Modeling has been given much publicity over time and has been considered crucial to a well-organized and efficient data processing organization. However, modeling has, in general, failed to deliver on the promised benefits of model-based application generation. Some of this failure is due to the sheer variety of modeling methodologies available, and some is due to the esoteric language and complexity of the concepts inherent in modeling. Perhaps the largest contribution to this failure is the lack of open interfaces and automation between the components and phases of the application development life-cycle. In this appendix, we present the concepts and benefits of modeling by a simple example and lay the groundwork for the role of the Financial Application Architecture* (FAA), Insurance Application Architecture* (IAA), and Retail Application Architecture* (RAA), in data processing in general and Information Warehouse implementation in specific.

Modeling organizes the data processing environment

A.1 The Construction Model

The building of a structure—a home, an office building, or other complex structure—serves as a platform for discussing the benefits of a model and modeling. A model is an abstract representation of a real world environment. Modeling classifies certain aspects of building into things called entities. The concept of entities immediately presents a challenge for relating modeling to a real-world environment. Entity is a term used to classify people, places, things, ideas, concepts, or events that are relevant to the business. The key here is to understand the motivation for entities. Entities provide a way to group things that have common characteristics or role with respect to the business.

Entity: classifying things

For example, within a construction company, entities include nails, boards, and windows; carpenters, plumbers, and electricians; contracts; trucks; and other components of the business. As a general categorization, entities is a convenient catch-all for anything with which the construction project leader—the general contractor (GC)—has to be concerned. The benefit is that the GC can look at one list of things needed to build the house.

Entities can be grouped to simplify their use

Further reflection reveals that these entities are not all the same, that some subsets of these entities have common characteristics. If it is of benefit to group all entities, then it is of more benefit to group them so that the common aspects of the subgroupings can be utilized.

A.1.1 Entity: Things

Entities make the overall process simpler

The next step is to create entities within the larger-scope Entity, based on these common aspects or ways of being used. For example, nails and boards have attributes in common: they are both things to be ordered, stored, and physically incorporated into the structure. We therefore create a specific type of entity called Materials. Materials is a grouping of things that are handled in a similar manner from a business perspective and typically have common attributes. The attributes describe the nature of the entity. In this case, the attributes are the color, size, and other physical aspects of the thing. The value here is that the GC can use one order sheet for all things that fit into the Entity category Materials. The GC can use another form for all subcontracts for services. The GC's job is now easier because the different parts of the job can be generalized.

Entities help generalize data processing processes

At this point, we have introduced two perspectives on things essential to the business: the way these things fit into the business—what the business does with the thing—and the attributes or descriptions of these things. The benefit then is that we can think of the things that are part of our business as a general group. We can also generalize the business activities performed against the things.

A.1.2 Entity: Agreements

The next set of entities is centered around the contract, or agreement, legal or otherwise, that is part of building the structure. Contracts are written, agreed to, and enforced. Contracts are different from nails and boards, which are ordered, installed, or are part of the physical structure. By separating these two sets, we can treat them appropriately from a business point of view. What is of more interest here is that they can be treated consistently from a data processing point of view. That is, application code can be written to consistently operate on data about things defined as being the same entity. Furthermore, application code written to operate on data about things used in building a house is consistent with application code written to operate on data about those same things used in building an office.

This approach suggests that contracts in the construction business, categorized as an entity called agreements, can be viewed from both a business and a data processing perspective as similar to contracts in the insurance industry, where they are policies.

A.2 The Annual Report As a Model

We can then extend these concepts to the corporate statement. The annual report presents the financial status of an enterprise in terms of its assets and liabilities. Both the assets and liabilities can be seen as entities, but this is still rather ambiguous with respect to common experience. A better example is the subheading Plant and Property under assets. This is a financial view of things the enterprise has as an asset. This perspective on the enterprise is similar to the points we stressed as being a benefit for modeling, entities, and the general categorization philosophy of modeling. Regardless of the industry within which the enterprise does business, the enterprise always has some type of plant and property.

The annual report is a model

From an industry perspective, we can treat all plant and property as something that has value, that exists. From a data processing perspective, we can expect to write applications that sum up the present value of that plant and property. Furthermore, we can expect to write applications that depreciate that plant and property over time. The treatment and expectations of these business objects are independent of the enterprise's industry. We have gained perspective on a component of the business and have gained an opportunity to leverage data processing resources by this generalization, which is wholly compatible with the objectives of modeling.

A.3 Information Warehouse and Modeling

Models contain representations of the enterprise's business in the form of Entities and other modeling constructs. These constructs contain technical and descriptive information about the business objects. The technical information includes data type and length and may in fact be used in limited ways by an application generator. The descriptive information is for reading and understanding purposes only for the model user. This descriptive information is called meta-data and is a crucial part of an the Information Warehouse environment.

Models are a formal representation

Meta-data explains the meaning of the object and the data processing object in business terms. It helps the administrator know whether the business/data processing object is needed for informational analysis. It also helps the knowledge worker understand the meaning of the object, once it is incorporated into the Information Warehouse implementation. This connection between modeling, the model, and the dictionary for the knowledge worker (the Information Catalog) is dependent on a subset of the model information in the Information Catalog.

Meta-data explains object meaning

The descriptive information also reflects the reconciliation and enhancement process. It describes the business/data processing object as it exists. The knowledge administrator and the business analyst know what the informational data needs to look like for use in an informational environment. The processes of decoding, reconciling, and enhancing operational data for use in an informational environment are based on these before and after descriptions.

Meta-data reflects data enhancement

List of Abbreviations

ATM	automated teller machine	GOI	generic output interface
CPU	central processing unit	GSA	General Store Application
DASD	direct access storage device	HHT	hand-held terminal
DBA	database administrator	IAA	Insurance Application Architecture
DIS	data interpretation system	IBM	International Business Machines Corporation
DRDA	Distributed Relational Database Architecture	ITSO	International Technical Support Organization
EDI	electronic data interchange	MIS	management information system
EFT	electronic funds transfer	PWS	programmable workstation
FAA	Financial Application Architecture	ROA	return on assets
GMROI	gross margin return on investments	RAA	Retail Application Architecture
		UPC	universal product code

Index

A

- access enablers
 - DRDA 51
 - health industry use 21
 - Information Warehouse architecture 48
 - purpose 51
 - SQL mappers 50
- accounting exit 72
- activity 75
- application
 - informational 4
 - narrow focus 19
 - series 19
- application development
 - commonality 34
- architecture
 - industry 5
 - Information Warehouse 5
- automation
 - value 67
- awareness 38

B

- barriers to change
 - installed systems 16
 - legacy systems 14
- business analysts 53
- business objects
 - business and technology views 63
 - standards 35
 - types 34
- business processes 22

C

- changed data implementation 64

- client-server
 - AS V3 84
 - FlowMark 77
 - informational application 84
 - modeling 74
 - Visualizer 85
 - work groups 84
 - workflow management 22
- Common User Access 53
- complex processes 52
- connectors 75
- customer service 17

D

- data
 - aggregation 53, 63
 - categorization 61
 - changed 64
 - clean 52
 - coded 63
 - derived 64
 - gathering and presenting 23
 - meta-data 38
 - real-time 64
- data access
 - common requirement 17
- data replication
 - automation 67
 - flexible 59
 - integration 69
 - justification 56
 - tool functions 68
- DataGuide
 - extractor 39
 - implementation 38
 - interface deploying 49
 - interface enabling 49
 - populating 39
 - role 61
 - storing meta-data 58

- DataHub
 - copy tool 68
 - data replication 90
 - DataRefresher 70
 - FlowMark 76, 78
 - systems management 42
 - workflow management 22
- DataRefresher
 - exit routines 72
 - role 69
- DB2
 - data processing terminology 63
 - interface enabling 49
 - platform flexibility 51
- deploying product 49
- Distributed Relational Database Architecture
 - See DRDA
- DRDA
 - in Access Enablers 51

E

- embedded SQL 50
- employee skills
 - in insurance 14
- enabling product 49
- entity
 - classify 91

F

- FlowMark 77

G

- generic data interface 71
- generic output interface 72
- GUI
 - DataRefresher 69
 - workflow management 74

H

- health insurance 12

I

- industry group 34
- information delivery
- Information Warehouse architecture
 - benefits 24
 - objective 45
- Information Warehouse framework
 - connectivity 42
 - DataHub 42, 90
 - definition 41
 - Visualizer 87
- informational application
 - function 52
 - interface deploying 49
 - Visualizer 69, 85
- informational object
 - types 52
- Insurance Application Architecture
 - business model 32
 - composition 29
 - data replication 30
 - goal 29
 - system model 32
- insurance industry
 - aversion to change 15
 - changing roles 15
 - customer service 35
 - DataGuide 38
 - focus areas 13
 - information technology requirements 16
 - information technology role 15
 - Information Warehouse architecture 17, 21
 - Information Warehouse's role 17
 - key systems 12
 - long-term care 20
 - management information systems 55
 - market pressures 13
 - new policy process 22
 - product management 19
 - profit 18
 - profit analysis 56
 - self-managed teams 14
 - survivors 17
- integration 76
- interfaces
 - informational applications 53
 - workflow management 52

J

- job control 73

justifying reconciled data 64

K

knowledge worker
definition 4

L

liability insurance 12

M

management information systems
definition 55
objective 56
meta-data
building 38
definition 63
example 58
gathering 58
Information Catalog 63
population 38
reconciliation 58
role 57
storage 58
model
-based application generation 91
-based workflow management 74
abstract 91
benefits 31
generic 34
high-level 33
individuality 35
practical limit 33
workflow 75
modeling
definition 31

O

organization asset data
data and meta-data 62
DataGuide 61
health industry 20
managing 64
organizing 61, 89
prospecting 65

P

people tasks 51
policy
definition 12
premium payment 56
presentation function
in insurance industry 59
process
definition 74
product management 59
profitability
in insurance 13
property insurance 12
prospecting 65

R

reconciliation
definition 57
multisource 58
refresh 69
refresh propagation 68

S

service office 17
solution thread
objective 4
SQL call level interface 50

T

tool bar 86

U

underwriting 24

V

visualization 59
Visualizer
database targets 85
interface deploying 49
objects 86

W

- workflow management
 - client-server support 22
 - DataHub 22
 - interface 52, 76
 - process integration 51
 - program 75
 - requirements 22, 74
- Workflow Management coalition 24, 76

ITSO Technical Bulletin Evaluation

RED000

Information Warehouse in The Insurance Industry

Publication No. GG24-4341-00

Your feedback is very important to help us maintain the quality of ITSO Bulletins. **Please fill out this questionnaire and return it using one of the following methods:**

- Mail it to the address on the back (postage paid in U.S. only)
- Give it to an IBM marketing representative for mailing
- Fax it to: Your International Access Code + 1 914 432 8246
- Send a note to REDBOOK@VNET.IBM.COM

Please rate on a scale of 1 to 5 the subjects below.
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction	_____		
Organization of the book	_____	Grammar/punctuation/spelling	_____
Accuracy of the information	_____	Ease of reading and understanding	_____
Relevance of the information	_____	Ease of finding information	_____
Completeness of the information	_____	Level of technical detail	_____
Value of illustrations	_____	Print quality	_____

Please answer the following questions:

- a) If you are an employee of IBM or its subsidiaries:
Do you provide billable services for 20% or more of your time? Yes____ No____
Are you in a Services Organization? Yes____ No____
- b) Are you working in the USA? Yes____ No____
- c) Was the Bulletin published in time for your needs? Yes____ No____
- d) Did this Bulletin meet your needs? Yes____ No____
- If no, please explain:

What other topics would you like to see in this Bulletin?

What other Technical Bulletins would you like to see published?

Comments/Suggestions: (THANK YOU FOR YOUR FEEDBACK!)

Name

Address

Company or Organization

Phone No.




Fold and Tape

Please do not staple

Fold and Tape



NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES



BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM International Technical Support Organization
Department 471, Building 070B
5600 COTTLE ROAD
SAN JOSE CA
USA 95193-0001



Fold and Tape

Please do not staple

Fold and Tape

