

 WILEY


 KIMBALL
GROUP


The Microsoft[®] Data Warehouse Toolkit

With SQL Server[™]
2005 and the
Microsoft[®] Business
Intelligence Toolset

Joy Mundy and
Warren Thornthwaite
with Ralph Kimball

Foreword by Bill Baker
General Manager, SQL Server Business Intelligence, Microsoft





The Microsoft® Data Warehouse Toolkit

**With SQL Server™ 2005 and the
Microsoft® Business Intelligence Toolset**

Joy Mundy and
Warren Thornthwaite
with Ralph Kimball



WILEY
Wiley Publishing, Inc.

**The Microsoft® Data Warehouse Toolkit:
With SQL Server™ 2005 and the Microsoft® Business Intelligence Toolset**

Published by
Wiley Publishing, Inc.
10475 Crosspoint Boulevard
Indianapolis, IN 46256
www.wiley.com

Copyright © 2006 by Wiley Publishing, Inc., Indianapolis, Indiana

Published simultaneously in Canada

ISBN-13: 978-0-471-26715-7

ISBN-10: 0-471-26715-5

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

1B/SW/QR/QW/IN

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600. Requests to the Publisher for permission should be addressed to the Legal Department, Wiley Publishing, Inc., 10475 Crosspoint Blvd., Indianapolis, IN 46256, (317) 572-3447, fax (317) 572-4355, or online at <http://www.wiley.com/go/permissions>.

Limit of Liability/Disclaimer of Warranty: The publisher and the author make no representations or warranties with respect to the accuracy or completeness of the contents of this work and specifically disclaim all warranties, including without limitation warranties of fitness for a particular purpose. No warranty may be created or extended by sales or promotional materials. The advice and strategies contained herein may not be suitable for every situation. This work is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional services. If professional assistance is required, the services of a competent professional person should be sought. Neither the publisher nor the author shall be liable for damages arising herefrom. The fact that an organization or Website is referred to in this work as a citation and/or a potential source of further information does not mean that the author or the publisher endorses the information the organization or Website may provide or recommendations it may make. Further, readers should be aware that Internet Websites listed in this work may have changed or disappeared between when this work was written and when it is read.

For general information on our other products and services or to obtain technical support, please contact our Customer Care Department within the U.S. at (800) 762-2974, outside the U.S. at (317) 572-3993 or fax (317) 572-4002.

Library of Congress Cataloging-in-Publication Data

Mundy, Joy, 1961–

The Microsoft data warehouse toolkit : with SQL Server 2005 and the Microsoft Business Intelligence toolset / Joy Mundy and Warren Thornthwaite with Ralph Kimball.

p. cm.

Includes index.

ISBN-13: 978-0-471-26715-7 (paper/website)

ISBN-10: 0-471-26715-5 (paper/website)

1. Data warehousing. 2. SQL server. I. Thornthwaite, Warren, 1957– II. Kimball, Ralph. III. Title.

QA76.9.D37M95 2006

005.74—dc22

2005032275

Trademarks: Wiley and the Wiley logo are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates, in the United States and other countries, and may not be used without written permission. Microsoft and SQL Server are trademarks or registered trademarks of Microsoft Corporation in the U.S. and/or other countries. All other trademarks are the property of their respective owners. Wiley Publishing, Inc., is not associated with any product or vendor mentioned in this book.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.



About the Authors

Joy Mundy, a member of the Kimball Group, has been developing, consulting on, and speaking and writing about business intelligence systems and technology since 1992. Joy began her career as a business analyst in banking and finance as one of the power users we talk about in business intelligence. In 1992 she joined the data warehouse team at Stanford University, an effort that was both educational and character building. She next co-founded InfoDynamics LLC, a data warehouse consulting firm, and then joined Microsoft WebTV to develop closed-loop analytic applications and a packaged business intelligence system.

From 2000 to 2004, Joy worked with the Microsoft SQL Server Business Intelligence product development team. She managed a team that developed the best practices for building business intelligence systems on the Microsoft platform. Joy graduated from Tufts University with a B.A. in Economics, and from Stanford with an M.S. in Engineering Economic Systems.

Warren Thornthwaite, a member of the Kimball Group, has been building decision support and data warehousing systems since 1980. Warren co-authored the best-selling *Data Warehouse Lifecycle Toolkit* (Wiley, 1998).

Warren worked at Metaphor Computer Systems for eight years starting in 1983, where he managed the consulting organization and implemented many major data warehouse systems. After Metaphor, Warren managed the enterprise-wide data warehouse development at Stanford University. He then co-founded InfoDynamics LLC, a data warehouse consulting firm. Warren joined up with WebTV to help build a world-class, multi-terabyte customer-focused data warehouse before returning to consulting.

In addition to designing data warehouses for a range of industries, Warren has extensive experience helping clients develop scalable, practical information access architectures. Warren holds an MBA in Decision Sciences from the University of Pennsylvania's Wharton School, and a BA in Communications Studies from the University of Michigan.

Ralph Kimball, Ph.D., founder of the Kimball Group, has been designing information systems and data warehouses since 1972.

Ralph wrote his Ph.D. dissertation in the Electrical Engineering department at Stanford University on the design of a man-machine system for tutoring mathematics students. In 1972 he joined the Xerox Palo Alto Research Center as a research scientist. Over the following ten years at Xerox, he became a development manager and the product marketing manager for the Xerox Star workstation, the first commercial product that used windows, icons, and the mouse. For this work at Xerox, he received the Alexander Williams Award from the IEEE Human Factors Society for user interface design.

Following his years at Xerox, Ralph was a vice president and member of the founding team at Metaphor Computer Systems, the first data warehousing company. Between 1982 and 1986, Metaphor installed many client-server data warehouse systems. In 1986 Ralph founded Red Brick Systems, which developed the first high-performance relational database for decision support. Since 1993, Ralph has designed data warehouse systems, written bestselling data warehouse books, and taught data warehousing skills to more than 10,000 IT professionals.

Defining Business Requirements

Building the foundation

Business requirements are the bedrock of the DW/BI system. Business requirements guide the development team in making the biggest strategic choices, such as prioritizing subject areas for implementation, and in making the smallest tactical design decisions, such as how to present key performance indicators on the users' screens. In this chapter, we cover the process of gathering business requirements and converting them into a DW/BI system strategy. We describe the process of interviewing business and IT representatives, categorizing their requirements into analytic themes, converting those themes into incremental projects, and working with senior management to prioritize those projects. We also include a partial set of example requirements for the Adventure Works Cycles business.

As Figure 1.1 illustrates, the Business Requirements Definition is the foundation of the Lifecycle methodology. Business requirements and their associated business value give you the guidance you need to make decisions in all three downstream tracks. As you'll see, they influence the project scope and plan as well.

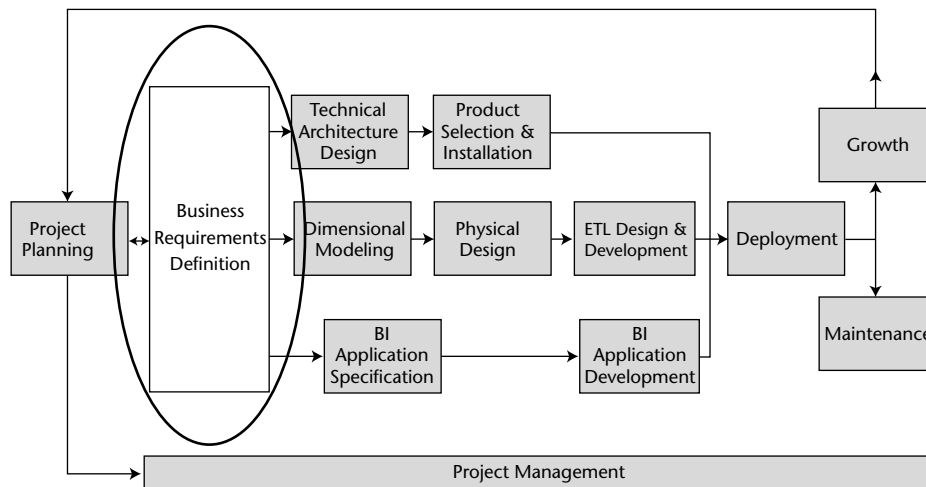


Figure 1.1 The Business Requirements Definition step of the Business Dimensional Lifecycle

This chapter is primarily about resisting temptation. Gathering business requirements is often outside a technical person's comfort zone. The overall success of the project is largely determined by your understanding of the business requirements and your relationships with the business people. Resist the temptation to just start loading data.

In this chapter you learn the following:

- The importance of understanding business requirements and securing solid business sponsorship
- How to define enterprise-level business requirements, including the interview process, developing analytic themes, linking themes to business processes, developing the data warehouse bus matrix, and prioritizing business processes with senior management
- How to plan the initial business process dimensional model implementation and gather project-level business requirements
- What a typical requirements summary document looks like and how it links to analytic themes and business process implementations

The Most Important Determinant of Long-Term Success

There is one common factor in successful business intelligence projects: delivering business value. Your DW/BI team must embrace the goal of enhancing business value as its primary purpose. This seems like an obvious statement, and we almost always get a chorus of agreements when we state this principle

to the DW/BI teams with which we work. But most DW/BI folks are technologists at heart. We like the certainty of computers and programming. It works or it doesn't; if it doesn't, we can debug it.

You can't deliver business value unless you work closely with business people. You need to understand their language and learn to see the world from their point of view. You'll be working in a non-technical, highly ambiguous, politically sensitive environment. Are you feeling queasy yet? Many of us went into the computer trade specifically to avoid such discomfort. But this unsettled environment is what the DW/BI system is all about. You must develop the business knowledge and people skills right along with your technical skills in order to meet the needs of your business users. We realize the entire team will not become smooth-talking MBAs. However, someone on the team must have strong business and communications skills, and everyone will be more effective if they work to develop some of these skills.

So, while many DW/BI teams and consultants pay lip service to business value, the reality of their day-to-day behavior is that technology rules. Do not let this happen to you. Technology is important; business value is mandatory.

As you read this book, you'll encounter recommendations that may seem unnecessarily complicated or just plain unnecessary. Every time you're tempted to dismiss the authors as overly fond of their design methodology, or just overzealous, consider whether your reactions are driven by your technical convenience, or by the business users' needs. Never lose sight of the business.

Uncovering Business Value

If you're going to be driven by business value, you need to go out and identify, understand, and prioritize the needs of the business. This is easier said than done if your focus has historically been on technology. Fortunately, the Business Dimensional Lifecycle provides the tools to work through an entire development iteration of a data warehouse, beginning with business requirements.

Where do you start with your business intelligence system? What is the first step? The consultant in us immediately blurts out the standard consulting answer: "It depends." In fact, it does depend on a host of factors, such as how your organization works, what you already know about the business, who is involved in the project at this point, what kinds of DW/BI efforts came before, and many other factors. Let's talk about the most common scenario first, and then we'll address a few exceptions.

More often than not, the DW/BI system starts as a project hosted by the Information Technology (IT) organization. There is generally some level of business interest; in fact, the business folks may be the source of inspiration. But they are pushing for information in a form they can use, not specifically for a DW/BI system (unless, of course, they had access to a well-built data warehouse in their last job and they really miss it).

Most often, the IT-driven DW/BI project gets started because the CIO decides the company needs a data warehouse, so people and resources are assigned to build one. This is a dangerous situation. Please refer to the first point in this chapter: Focusing on business value is the most important determinant of long-term success. The problem with the IT-driven DW/BI system is that it almost always centers on technology. The team has been assigned the task of building a “warehouse,” so that’s exactly what they do. They get some hardware and some software and start extracting data.

We know some of you are thinking, “Oops, I already bought the ETL server and the user reporting tools.” That’s probably okay, but put those tools aside for the moment. Step away from the keyboard. If you get sucked into the technology, you’re missing the whole point. You can build a technically great DW/BI system that provides very little business value. As a result, your project will fail. You have to start with business value, and identifying business value involves several major steps:

- Recruiting strong business sponsorship
- Defining enterprise-level business requirements
- Prioritizing business requirements
- Planning the project
- Defining project-level business requirements

We’ll run through each of these steps in the following sections.

Obtaining Sponsorship

Developing solid business sponsorship is the best place to start the DW/BI project. Your business sponsors (it is generally good to have more than one) will take a lead role in determining the purpose, content, and priorities of the DW/BI system. You will call on them to secure resources and to evangelize the DW/BI system to the rest of the organization. This includes activities such as arranging for a planning meeting with senior staff, or speaking to a room full of business users at the project kick-off. You need to find at least one person in the organization who scores well in each of the following areas:

- *Visionary*: Someone who has a sense for the value and potential of information and some clear, specific ideas on how to apply it.
- *Resourceful*: Someone who is able to obtain the necessary resources and facilitate the organizational change the data warehouse will bring about.
- *Reasonable*: Someone who can temper his or her enthusiasm with the understanding that it takes time and resources to build a major information system.

Often, if you've been with your company for a while, you already know who these people are. In this case, your task is to recruit them onto the project. However, if you're new to the company, or you have not been out of the IT group much, you'll need to go out and find your business sponsors. In either case, the best way to find and recruit these people is by conducting an enterprise business requirements gathering project. It's worth the effort. Good business sponsorship can provide the resources and support you need to deliver real business value.

Defining Enterprise-Level Business Requirements

From a technical point of view, one long-term goal of the DW/BI team is to build an enterprise information infrastructure. Clearly, you can't do this unless you understand business requirements from an enterprise level. It is especially important in larger organizations to begin with this broad understanding because it is rare for the DW/BI team to have such an enterprise-level perspective. It's also particularly important for organizations that are just starting their first DW/BI system (or starting over) because getting the enterprise perspective built into the initial project helps you avoid painful and costly redesign down the road.

In these cases, the high-level version of the Lifecycle presented in Figure 1.1 doesn't give you quite enough information about how best to get started. In particular, the little arrow that goes both ways between Project Planning and Business Requirements Definition actually breaks down into several sub-activities, as shown in Figure 1.2.

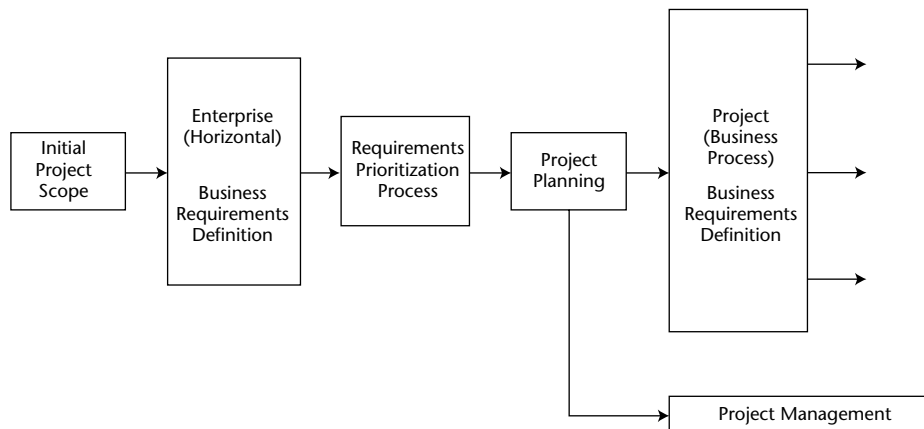


Figure 1.2 Detail of sub-activities between initial scope and project requirements

ENTERPRISE VERSUS PROJECT

Many organizations we've worked with have a clear understanding of which business process is their top priority right from the start. In these cases, we often combine the enterprise requirements definition step and the project requirements definition step into a single effort.

This does not lessen the importance of understanding the full range of enterprise requirements for information. In fact, we almost always go through the enterprise prioritization process with senior management. However, because the top priority is clear early on, we make sure we gather enough detailed information about that business process and the data it generates in the same interview set so we can create the design for the first business process in one pass instead of two.

In this subsection of the Lifecycle, defining the business requirements happens in several distinct steps. The rest of this section describes each of these steps in more detail.

NOTE We use the term "business process" throughout this section but do not describe it in detail for another few pages, when it makes more sense in the context of the requirements definition process. For now, think of it as a subject area or data source.

Establishing Initial Project Scope

Begin with the creation of an initial project scope based on the team's upfront knowledge of the organization's business needs and its experience in developing DW/BI systems, along with input from the business participants. The scope usually covers only the enterprise-level requirements definition and requirements prioritization steps in detail, leaving the initial project implementation plan for later when you have a much better idea of what the project needs to accomplish from a business perspective. The initial scope usually involves only user interviews, interview write-ups, a few meetings, and the creation of the final requirements document. It usually takes three to six weeks (or more) depending on how many interviews you do.

REFERENCE Additional information about project planning and management in the context of the Business Dimensional Lifecycle can be found in *The Data Warehouse Lifecycle Toolkit* (Wiley, 1998), Chapter 3.

Gathering and Documenting Enterprise-Level Business Requirements

The enterprise requirements definition step is designed to gather a broad, horizontal view of the organization from a business point of view. The process flow chart in Figure 1.3 breaks the Enterprise requirements definition box from Figure 1.2 down into its subtasks. As we see in Figure 1.3, the core part of defining business requirements involves gathering and documenting those requirements.

While the four steps that are circled on the left side of the figure are shown as separate subtasks, we usually do them in a pipeline fashion, conducting an interview, extracting its analytic themes, identifying their supporting business processes, and placing each business process in the initial bus matrix. We'll describe each of the core subtasks in Figure 1.3 in this section, leaving the senior management prioritization session for its own section.

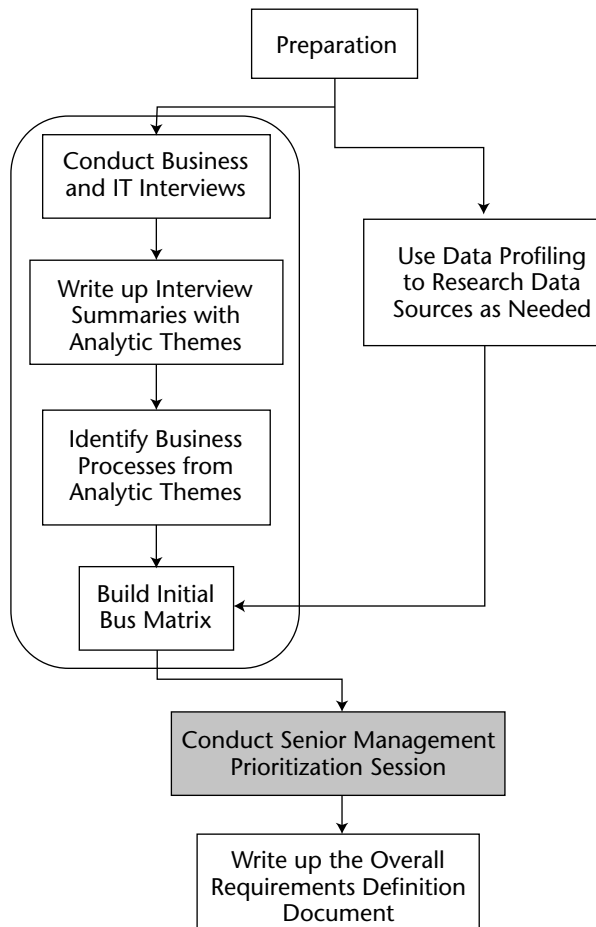


Figure 1.3 The enterprise requirements definition process flow chart

Preparation

Requirements definition is largely a process of interviewing business and technical people, but before you get started, you need to do a little preparation. Learn as much as you can about your business, your competitors, your industry, and your customers. Read your organization's annual report; track down any internal strategy documents; go online and see what's said about your organization, the competition, and the industry in the press. Find out what the big challenges are. Learn the terms and terminology the business people use to describe what they do. In short, do your homework.

Part of the preparation process is figuring out whom you should actually interview. This usually involves carefully examining an org chart with your sponsor and other key supporters. There are typically four groups of people you need to talk to early on: senior management responsible for making the strategic decisions for the organization; mid-management and business analysts responsible for exploring strategic alternatives and implementing decisions; source systems experts (the folks who really know what kinds of data issues are out there for you to trip over); and finally, people you need to interview for political reasons. This last group may not add any value, but if they're omitted, they could cause problems. Increasingly, there is a fifth group of people who may need the data warehouse—operational decision makers needing real-time or near real-time access to historical and integrated data. Your initial interviews should sample the operational community to see if such interest exists, and if so, a representative group of these users should be included in the interviews.

It's easy to start the interview list with the CEO and his senior staff. Add the analysts and managers who are known as leaders in the business intelligence area—folks whom senior management and co-workers turn to when they need information. They are also usually the folks who bug the IT organization the most. If you've been at your organization more than 12 months, you know who these people are. They have their own Access databases, they write SQL against the transaction system, and they create reports and charts with whatever tools they have available (mostly Excel and Access). Finally, add on a couple of the key IT folks who can educate you about the nature of the source systems.

NOTE You are just making the list of whom to interview at this point, not the interview schedule. Start the schedule with a few people you know and trust before you turn to senior management. At the same time, make sure you get the elusive executives on the calendar as early as possible. Some of these folks can be tough to pin down.

A major goal during the interviews is to build positive working relationships with the business folks. These relationships will build on your understanding of the business issues your organization faces. In short, be prepared. Fortunately, gathering this information is not as difficult as it used to be, thanks to the Internet. However, you still have to read it.

Interviewing Business and IT

The next step in requirements gathering is actually talking to business users. Interview individuals or small groups (two to five people at most), rather than hold a large group design session. The individual interview allows each person to present his or her views on what the challenges are and what is required for success. It also gives the DW/BI team more opportunities to capture and clarify critical information. Finally, individual interviews are less work for the business folks. They take only an hour or so, rather than a whole day or two for group design sessions. On the downside, individual interviews collectively take longer and are more work for the DW/BI team.

In this first pass at gathering requirements, you will interview more senior level folks across the different departments and get a comprehensive list of the major challenges and opportunities your organization faces. We call these challenges and opportunities *analytic themes*. These themes often (but not always) line up with the strategic goals and initiatives of the organization. We talk about themes in more detail later in this section.

When you speak to the business users, you will often hear business requirements that sound really important and valuable, but end up being impossible to meet. A request such as “Sales data for our competitors by SKU by retail store by day would allow us to increase our market share by at least 5 percent” sounds great. The catch is that, for most industries, the data is not available. Even internal data might not be available, or it might be so dirty or poorly structured that it will take a major effort to make it presentable.

This is precisely why you need to interweave some interviews with source systems experts to understand the structure and content of the source systems and the nature of any data problems that might be lurking out there. You must also perform data profiling on all candidate data sources. (Data profiling is described in an upcoming section.)

At the end of each interview the interview team must take a few minutes to debrief. Review your notes, fill in the blanks, make sure you understand the terms you heard, and capture the key issues. The longer you wait to do this, the less you will remember. We’ve found ourselves staring at a sentence that reads, “The most important factor in our business is . . .” with no idea what came next. Debrief as soon as possible.

TIP Never ask the end user “What do you want in your data warehouse?” That puts them in the position of designing the system. That’s your job. Besides, there is only one right answer to this question: “Everything.” Instead, ask questions that help you learn what the end user does, and then translate this into what needs to go into the system. A question such as “How do you know when you have done a great job?” can get you started in the right direction.

Interview Summaries and Analytic Themes

It’s a good idea to start writing up summaries of the individual interviews as you work through the interview schedule. This takes a bit of work because you don’t want to simply transcribe the interview. Instead, summarize the various requirements and group them according to category or theme. This is where the analytic themes begin to appear.

Creating the list of analytic themes requires good categorization skills. It’s like being able to look at each tree in the forest and group it by genus (and maybe even species). It gets easier with experience, but it’s never simple. Start by grouping related requirements as you write up the individual interview summaries. Often, these groups are actually higher-level themes that include several sub-themes. In some cases, they’re not really analytic themes at all, but represent a need for a new or improved operational process.

These are some example themes from a subscription business, such as MSN, grouped by the strategic goals they support:

- A. Improve the effectiveness of the customer acquisition programs.
 1. Advertising
 2. Promotions
 3. Direct mail
 4. Online
- B. Increase sales per customer by offering additional items of potential interest (cross-sell), or similar, higher value items (up-sell).
- C. Improve margins by negotiating better prices and terms with suppliers.
- D. Identify desirable profile(s) based on long-term customers and use them to create targeted acquisition programs.

This is a simple, enterprise list of goals and themes. Yours will need to be more detailed and extensive. Each theme should have examples of the specific analytic requirement you heard in the interview. It should also include some sense of the business value of meeting the requirement. In other words, how much is negotiating better prices and terms worth?

You may create an initial list of 20 or more major themes. Some of the themes on the list may represent new ways of doing business and will require new transaction systems, or at least significant changes to existing transaction systems. But even a quick review of the simple list provided begins to bring out ideas about how the DW/BI system can address some of these needs in the short term.

These analytic themes are a key input for the next step in the enterprise requirements process—identifying the supporting business processes. We'll get to this in a bit.

Meanwhile, the last step in creating the interview summaries is to validate what you heard. This can be done by sharing the interview summaries with the interviewees and asking for comments. Do this before you start working on the final overall requirements document. That way, any misunderstandings can be fixed and additions included. This validation step sends a positive message to your business users. It tells them you listened and understood their issues. In many cases, this is the first step in repairing damaged organizational relationships. It may be helpful to hold a findings review meeting where you invite everyone who participated in the interviews and have your business sponsor do the introductions.

Data Auditing/Data Profiling

At the same time you are interviewing people and creating the summaries, you will also need to do some queries against the source system data to get a firsthand understanding of the data issues. This kind of querying has come to be known as *data profiling* or *data auditing*, and there are several tools on the market designed to support it. There are three major points in the Lifecycle where data profiling is helpful. The first is here, during the requirements definition process where you should do a simple red light/green light assessment of your organization's data assets. You aren't looking for nuances at this point, but if a data source needs to be disqualified, now is the time. The second place to do data profiling is during the design of the dimensional model, and the third is during the design and implementation of the ETL process. You will want to do more in-depth data profiling once you select a specific business process and begin defining project-level business requirements. We describe data profiling in more detail when we discuss the dimensional modeling process in Chapter 2.

Identifying Business Processes That Support the Analytic Themes

As you extract the analytic themes for the interview summaries, you need to dig into each analytic theme to identify the business process (or processes) that generate the data needed to perform the desired analyses. You convert from themes to business processes because business processes are the units of work in building the DW/BI system. Each business process is usually measured by a single source system, which translates into a single pass through the Business Dimensional Lifecycle process. (See the related sidebar for more information.)

BUSINESS PROCESS: THE DW/BI SYSTEM UNIT OF WORK

We use the term “business process” to mean an operational activity the organization engages in to accomplish its primary goals. You can think of business processes as the links in the organization’s value chain. Each business process typically has its own operational system or module that enables it, such as the order entry system, or the call tracking system, or the inventory management system. The information generated by these business processes measures only the business process itself, but that information usually has value well beyond the boundaries of the individual business process. Information from a single business process, such as orders information, could be of great interest to Sales, Marketing, Customer Service, and other groups across the organization.

Each business process is a unique, coherent measurement system implemented as an operational system. If you need data from a given business process, you need to extract that data in its business context. In other words, you need to pull the measures and all of the associated descriptors in a careful, systematic fashion. This makes the business process the fundamental unit of work for the DW/BI system. Unless you have unlimited resources, your DW/BI team will concentrate on designing and loading data from one business process at a time. The business process is the DW/BI system unit of work.

While many analytic themes require information only from a single business process, one challenge you will face is that some themes require data from multiple business processes to meet the overall analytic needs. Customer and product profitability themes are good examples. They sound like a single analysis (the customer scorecard), but they actually require data from many separate business processes. We call these *consolidated* themes because they cannot be completely built until all prerequisite business processes are in place.

Converting from analytic themes to business processes helps determine the level of effort needed to support a given theme. If a theme must have data from more than one business processes, it will take more than one pass through the Lifecycle to enable that theme. While these passes can happen in parallel, there is no way around doing the work.

Converting from analytic themes to business processes involves thinking about which business processes are required to support each theme. For example, it may be possible to support the cross-sell and up-sell theme (B) from the example list based on data from a single business process: orders. With a data mining tool and orders data that indicates which products are bought together (the order, or “market basket”), you can identify the common relationships, which you can then use to influence sales. (“Would you like fries to go with that shake?”)

On the other hand, a theme such as customer profiles (D) might draw on data from most of the major business processes. A profitable customer uses the web to get product information (low marketing costs), buys a lot (high sales), keeps what he buys (low returns), and doesn't need much help (few customer service calls). In this case, you'd need data from at least four business processes (or rows on the bus matrix, as described in the next section) to support theme (D). This is an example of a consolidated theme because it requires that data from more than one business process, such as orders or returns, be implemented first.

The worst case theme is often called a scorecard or executive dashboard. This deceptively simple application draws on data from almost all business processes in the organization. You can't create the entire dashboard until you've built the whole data warehouse foundation. Or worse, you end up building the dashboard by hand every day, manually extracting, copying, and pasting data from all those sources to make it work. It can be difficult to get business folks to understand the magnitude of the effort involved in creating this "simple" report.

Building the Initial Data Warehouse Bus Matrix

As you identify the business processes needed to support each analytic theme, you will also add those business processes to an enterprise data roadmap called the Data Warehouse Bus Matrix. This matrix maps your organizational business processes to the entities or objects that participate in those processes.

Each row in the matrix is a business process. Figure 1.4 shows a simplified example bus matrix for a retail company. Notice how the business processes down the left side of the matrix follow the organization's value chain. In this case, the company buys goods from their vendors and stores them in distribution centers. Then, as goods are demanded by consumers, they are moved out to the retail stores where they're held on shelves until the customer buys them and the goods leave the company's value chain. These business processes generally correspond to individual source systems or modules in the overall Enterprise Resource Planning (ERP) system.

	Business Processes	Dimensions						
		Date	Product	Vendor	Shipper	Dist Ctr	Store	Promo
Value Chain ↓	Purchase Orders	X	X	X		X		
	Dist Ctr Deliveries	X	X	X	X	X		
	Dist Ctr Inventory	X	X			X		
	Store Deliveries	X	X		X	X	X	
	Store Inventory	X	X				X	
	Store Sales	X	X				X	X

Figure 1.4 Example enterprise bus matrix for a retail company

The columns in the bus matrix are the descriptive objects that participate in the various business processes, such as Store, Product, and Date. They contrast with the measurement-driven business processes that label the rows of the matrix. We call these objects *dimensions* in the dimensional model. Each dimension participates in one or more business process—we indicate this by placing an X in the intersecting cell in the matrix. For example, the Vendor dimension is involved in both the Purchasing and Delivery processes. The Store Sale business process, on the other hand, does not involve the Vendor or Distribution Center.

The bus matrix is essentially your enterprise dimensional data architecture. For each business process (row), you can see exactly which dimensions (columns) you need to implement. And for each dimension, you can see which business processes it must support. This dimension-oriented view is the visual representation of *conformed dimensions*—a concept we define in the next chapter.

The business processes in the bus matrix, and the themes they support (and the value those themes represent) become the major inputs to the next step in the requirements definition process: a prioritization session with senior management.

Creating the Overall Requirements Document

It's best to wait until after the prioritization session to write the overall requirements document so you can include the resulting list of prioritized business processes. The overall requirements document includes the business process summaries, the bus matrix, and the prioritized results. You might want to include the interview summaries as an appendix for those readers who want all the detail.

Once you've completed the overall requirements document, the conceptual foundation of the DW/BI system is in place. The rest of the Lifecycle depends on what you learned in these initial steps to make decisions and set priorities for all three tracks that follow, and on into the deployment, maintenance, and growth phases.

The Prioritization Process

If you're a technical person, it's safe to say the prioritization process is one of the most powerful business tools you'll ever use. This is a bold statement, but we have used this tool many times and have been repeatedly successful. We've conducted a few prioritization sessions where the client decided not to move forward with the DW/BI project right away. This decision is usually reached because the prioritization process helped senior management better understand the nature of the commitment or the size of the data problems. This is

also a success because it means they will work to fix the problems rather than try to build a DW/BI system on shaky ground.

The prioritization process is a planning meeting involving the DW/BI team, the DW/BI project business sponsors, and other key senior managers from across the organization.

In this meeting, you describe the business processes you identified in the enterprise requirements gathering process so everyone has an understanding of the full list of possibilities. Go into this session armed with a PowerPoint presentation that describes each business process, gives a few examples of the associated analyses it will support along with a feel for the business value of those analyses, and includes an initial sense of level of effort needed to implement the business process (its feasibility). Be as crisp and clear as possible. Try to keep this presentation under two hours. As you describe each business process, you also describe the relative effort involved in supplying the needed data. Once everyone has an understanding of the business processes and terminology, take a break.

The second half of the session involves prioritizing the business processes. Lead the group in placing a sticky note for each business process onto a large version of a two-by-two grid like the one shown in Figure 1.5. This is an interesting exercise in negotiation and education and can easily take another hour and a half to two hours.

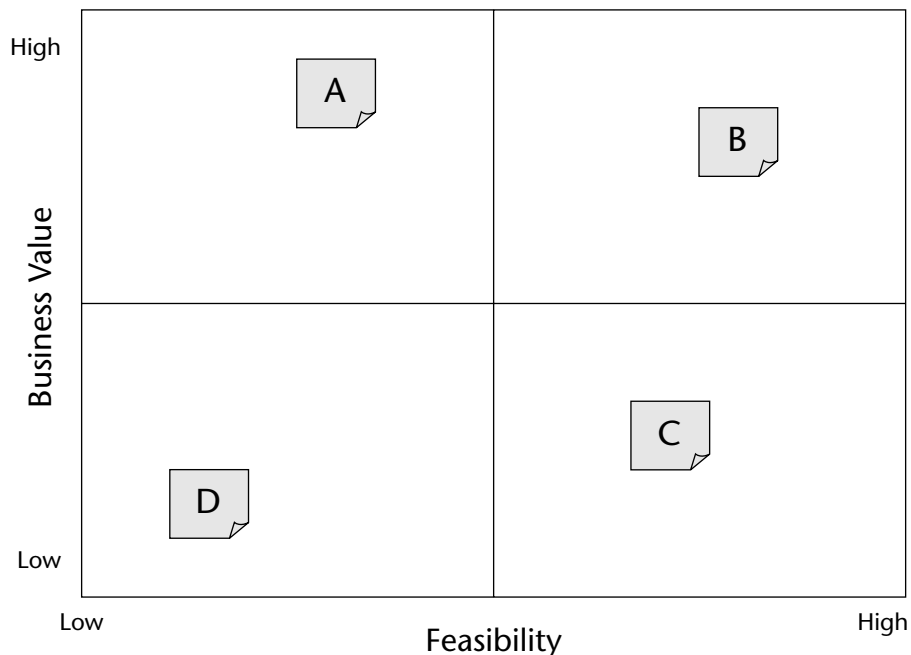


Figure 1.5 Example prioritization grid

The prioritization grid is deceptively simple: Study it carefully. The Y axis is about business value. The group needs to reach consensus on the relative impact of implementing each business process. The participants need to remember to take an organizational approach to assigning business value. There will always be someone who thinks any given business process is the absolute top priority. Gently remind them that there's more to the business than their little slice.

It helps to remind participants that the goal is to assign relative, not absolute, values on both axes. In Figure 1.5, we know that business process D is more difficult than business process A, but we don't need to know how many person-days it will take to implement either business process. The same goes for value. Figure 1.5 shows that business process A has more business value than business process B, but it's significantly less feasible to implement.

The X axis is about the level of effort each business process will take to implement. It is stated in terms of feasibility so the easier business processes go to the right (high feasibility) and the harder business processes go to the left (low feasibility). The DW/BI team leads the assignment of feasibility because team members have a better sense about the technical difficulties involved in each business process (although feasibility is not just technical—there are often organizational and political difficulties as well).

The true feasibility is not fully understood at this point. If you have someone on the team who's been in the organization long enough, she should have a good sense for the level of effort required to implement each business process. One obvious factor is when business processes must be implemented together to support a high value, consolidated theme.

The priority session is a good opportunity to educate the business folks about how bad things really are. You don't want to sound negative, but it's important to explain the level of effort it takes to gather the data and make it useful. For example, integrating customer IDs from two different source systems is a grind.

A CREDIBILITY BOOSTER

The prioritization process uses a common business school tool called the two-by-two matrix. This matrix was popularized in the early 1970s by the Boston Consulting Group. BCG used a "Growth-Share Matrix" to compare different business units in a portfolio by comparing Relative market share with industry sales growth rates. A business unit with high market share in an industry with high growth rate was called a "Star." By contrast, a business unit with low market share in a low-growth industry was a "Pet" (later referred to as a "Dog").

The great thing about the matrix is the positive impression the DW/BI team makes by cleverly adapting a classic MBA tool.

Sources: The Boston Consulting Group, *Perspectives on Experience and The Product Portfolio* (Boston, MA: The Boston Consulting Group, 1968)

Once all the business processes have been placed and everyone agrees on their relative locations, convert the matrix to a prioritized list of projects. One way to do this is to start in the upper-right corner of the prioritization grid and move to the lower-left corner, numbering the business processes as you encounter them. The two-dimensional nature of the matrix makes this a little difficult. Use the concept of concentric circles to establish a priority order, like ripples on a pond, centered in the upper-right corner.

The output of the prioritization process is a list of business processes in priority order. This list is your DW/BI roadmap; it tells you which row on the matrix, and which dimensions, to implement first. Less tangible, but equally important outcomes of the prioritization process are senior management consensus around the DW/BI roadmap, and a general improvement in the relationships between IT and the business.

In most cases, you will make only one pass at the enterprise requirements. Once the priorities are in place, the next pass and all subsequent passes will be at the level of the individual row on the bus matrix, the business process. Each row essentially becomes a project in the overall DW/BI program. From here on out, you will update enterprise business requirements and revisit priorities as the business changes, but most requirements definition efforts will be at the business process project level.

Revisiting Project Planning

Now that you have a clear idea of your business priorities and how they relate to the business processes (data sources), you can lay out a more detailed and precise project plan. This process is not much different from project planning for any major information technology project. The DW/BI perspective on project planning is described in detail in the Lifecycle Toolkit book and does not merit repeating here.

The plan will continue to evolve as you get more detail about the business requirements in the next step. There is a two-way arrow between project planning and project requirements definition in Figure 1.1, but the backward flow is not as major because you gained significant understanding of the nature of the opportunity in the enterprise requirements gathering and narrowed your scope in the prioritization process.

Gathering Project Requirements

Gathering project requirements follows the same basic process as the enterprise requirements gathering process described earlier. The difference is that now you have selected a particular business process on the bus matrix to implement. The enterprise requirements definition process should be a solid foundation for the project requirements. You now will deepen your understanding of the chosen business process.

The project requirements gathering step is about pulling together the information you need to be successful in the three tracks that follow. Specifically, you need enough detail to create real, practical, flexible data models that will support a broad range of analytic needs. You need a solid understanding of the technical issues around data volumes, data cleansing, data movement, user access, and a host of other issues so you can create a capable, flexible technical architecture to support the warehouse now and in the future. Finally, you need a clear understanding of the business analysis requirements to build the initial set of business intelligence applications to demonstrate value from the very start.

The same three steps you followed in the enterprise requirements process apply to the project requirements process: preparation, interviews, and documentation.

As we described in the enterprise requirements section, preparation is the critical first step. If you haven't already, do your homework. Study the particular business process in detail. Figure out as much as you can about how it works before you begin the interviews. Learn the business terminology, the steps in the business process, and how it is measured.

The goal with this round of interviews is to drill down on the selected business process in detail to understand the analyses, data models, and technologies required to make it work. This time you may take a more vertical slice of the organization, depending on the business process (some business processes have broader organizational appeal than others). Talk to the analysts, managers, report developers, and source systems people who can help you understand the intricacies of the business process in question. The actual interview process itself is generally the same as before.

ALTERNATIVES TO INDIVIDUAL INTERVIEWS

If interviews won't work in your situation, we have had success with group requirements gathering sessions, but they are more risky. If you must do group sessions, here are a few tips:

- ◆ **Preparation is even more important. You have to know the business, and you also have to know what you want to accomplish and how you are going to go about it.**
- ◆ **Have a clear agenda with times listed for each section, breaks, and food and drink. Reserve a good room with plenty of space and comfortable chairs. Make sure you have all the tools you need—flip charts, markers, white boards, overheads, computers, and a projector—whatever makes sense for your plan.**
- ◆ **Get a strong, experienced design meeting leader to run the meetings. You have only a short time. If someone takes the meeting off course, you won't get what you need.**

Depending on the business process selected, consider whether to interview your customers and suppliers. They are, or could be, business users of information in the DW/BI system. In fact, the need to offer information outside the organization is common enough that many of the front-end tool vendors include extranet access functionality as part of their product line. Listen carefully during the interviews to see if this is a likely source of significant business value for your organization.

Interviews with key source system people and data profiling play a bigger role in the project requirements gathering process. Strive to inform the dimensional modeling process as much as possible about both the business requirements and the data realities.

The documentation process for the project requirements is similar to that of the enterprise definition process, except it is more detailed. Where the analytic themes at the enterprise level ranged across all the business processes, at the project level, they should all be focused on the initial business process.

Although the project requirements definition task sounds a bit abbreviated here, it is actually the definition task you will repeat over and over, every time you iterate through the Lifecycle to bring the next priority business process into the DW/BI system. Let's hope you need to do the enterprise-level task only once, and then keep it updated.

REFERENCE There is much more to the requirements definition process. For additional detail, please refer to the following sources:

- *The Data Warehouse Lifecycle Toolkit* (Wiley, 1998).
- Search kimballgroup.com for the topics "Business Requirements" and "Business Acceptance" for several related articles.

Business Requirements Example: Adventure Works Cycles

The sample business intelligence databases in SQL Server 2005 are based on a fictitious company called Adventure Works Cycles (AWC), a multinational manufacturer and seller of bicycles and accessories. According to the descriptions in Books Online, the company is based in Bothell, Washington, USA and has regional sales offices in several countries. In this section we provide an example of the business requirements gathering process based on the Adventure Works Cycles business. Assuming the DW/BI system already has good business sponsorship, the first major step in the Lifecycle is requirements definition, and the first task in requirements definition is preparation.

Interview Preparation at Adventure Works Cycles

Typically, you'd carefully review all the information about Adventure Works Cycles that you could find, reading through strategy documents, annual reports, marketing plans, competitive analyses, and presentations from senior management's annual offsite planning meeting. Because Adventure Works Cycles is a fictitious company, you can't really do the kind of research you should. SQL Server Books Online (BOL) provides some background information about Adventure Works Cycles in the section "Sample Databases and Business Scenarios." You may want to review the materials in BOL to get a general sense for Adventure Works Cycles and the AdventureWorks transaction system database. In this section, we have enhanced the Books Online content by doing some analyses of data from the database itself. We also provide you with business requirements information from our imaginations.

Adventure Works Cycles Basic Business Information

There is an incredible amount of information about your organization buried in your transaction systems. Buried is the key word here—not many people can get at it, which is why you are building a DW/BI system in the first place. However, as a competent systems professional, you should be able to get at this data. A few queries against the transaction database can reveal much about the dynamics of the business and the nature of the data at the same time.

The Adventure Works Cycles transaction database is called AdventureWorks and is installed as part of the SQL Server samples. It holds data from January, 2001 through June 30, 2004. A few queries on this data reveal that Adventure Works Cycles is doing well, at least in terms of orders and growth. The Grand Total line in Table 1.1 shows that orders have been increasing rapidly, tripling in 2002 and increasing almost 50 percent in 2003. Based on the first six months, 2004 looks like another banner year. Table 1.1 also shows that Adventure Works Cycles sells products in four major product categories: bikes, components, clothing, and accessories. Bicycles account for more than 80 percent of orders, with clothing and accessories making up about 4 percent.

THE HISTORY OF ADVENTURE WORKS

The Adventure Works database has been evolving at Microsoft for several years. It's been used in various forms as a demo database for Microsoft's CRM solution, as a training database for .NET architecture courses, as an application to demonstrate Microsoft Transaction Server, and as an example to demonstrate online shopping using Commerce Server. In these last two incarnations, Adventure Works was a camping gear retailer.

Table 1.1 Adventure Works Cycles Product Orders by Category

CATEGORY	2001	2002	2003	2004 (YTD)
Bikes	10,985	28,854	38,026	24,160
Components	708	4,230	6,418	2,477
Clothing	35	501	1,031	589
Accessories	19	88	549	531
Grand Total	11,746	33,673	46,023	27,757

All figures in U.S. dollars

Adventure Works Cycles sells a lot of bikes. Additional database queries tell us where their bike orders come from. Table 1.2 reveals that AWC has sales in six countries, with about 60 percent of orders coming from the United States. The percentage of orders coming from outside the U.S. has increased from about 25 percent in 2001 to close to 40 percent so far in 2004.

Table 1.2 Adventure Works Cycles Product Orders by Country/Region

COUNTRY	2001	2002	2003	2004 (YTD)
United States	8,980	23,717	27,265	14,631
United Kingdom	467	2,348	5,683	4,232
Canada	1,340	4,359	4,721	2,278
France	27	1,542	3,915	2,576
Australia	814	1,186	2,142	1,917
Germany	119	521	2,298	2,123
Grand Total	11,746	33,673	46,023	27,757

All figures in U.S. dollars

In terms of sales channels, Table 1.3 shows that the bulk of orders come from bicycle stores and distributors, also known as Resellers. This channel accounts for about 70 percent of orders and is handled by a direct sales force of 18 people. Adventure Works Cycles broadened its business during the late 1990s by opening up a direct sales channel to consumers on the Internet. Oddly, Internet orders have been fairly constant at around 30 percent of sales across the four years of data we have available. While we don't know why this is so, we predict the VP of Sales will take credit for being able to grow the direct sales channel as fast as the Internet has grown.

Table 1.3 Adventure Works Cycles Product Orders by Sales Channel

SALES CHANNEL	2001	2002	2003	2004 (YTD)
Reseller	9,119	27,992	37,318	18,715
Internet	2,627	5,681	8,705	9,041
Total	11,746	33,673	46,023	27,757

All figures in U.S. dollars

Now you have a sense for what Adventure Works Cycles sells, where they sell it, and how they sell it. The next question is to whom do they sell? Table 1.4 gives a snapshot of customers broken down by sales channel. This split is vital to understanding AWC's customers because the two channels are very different.

Table 1.4 Adventure Works Cycles Customers by Sales Channel Snapshot

	TOTAL CUSTOMERS	ACTIVE CUSTOMERS	\$US SALES (000)	ORDER COUNT	AVG \$ PER ORDER
Reseller	701	467	18,715	901	23,751
Internet	18,484	11,377	9,041	13,050	827
Total	19,185	11,844	27,757	13,951	2,308

The Reseller channel has a total of 701 customers, 467 of whom are active (defined as having placed an order in the first six months of 2004). The Internet channel is much larger, with a total of 18,484 customers, 11,377 of whom have placed an order in 2004. The Reseller channel justifies a direct sales force in part because the average order is close to \$24,000 compared with an average of \$827 for the Internet. On the other hand, the Internet business should be significantly more profitable because the price would be closer to retail rather than wholesale, and the cost of selling would be much lower without a direct sales force. You may hear more about this during the business requirements interviews.

As the DW/BI system manager, you should continue from here, researching top reseller customers and their historical buying patterns (seasonality, product lifecycle, and so on). Investigate the Internet customers as well because you have demographic information on them that tells you who are they, what they buy, and where they come from.

These few reports demonstrate the power of analysis. They give you a general sense for the size and shape of Adventure Works Cycles. However, they do not give you enough information to build a DW/BI system. There is much more to learn in terms of the business, its strategies and plans, competitive

environment, and key players. This is where you depart from the facts and venture into the fuzzy-edged land of organization and politics.

Interview Planning

The Adventure Works Cycles documentation doesn't include an org chart, but the HumanResources.Employee table in the AdventureWorks database lists a total of 290 employees. The Employee table is self-referencing with a ManagerID field that allows us to generate the company's org chart. The 290 employees are distributed across departments according to the org chart based on the Employee table shown in Figure 1.6. The org chart shows seven direct reports to the CEO.

The org chart should raise a few questions for you, such as why James Hamilton, the VP of Production, has only one direct report, and why Facilities and Maintenance (with seven people) is the only group reporting to him. Meanwhile, why does Peter Krebs, the Production Control Manager (not a VP) have all the rest of the 200-person manufacturing organization reporting to him? Why is David Bradley, the Marketing Manager, only a manager when Brian Welker is the Sales VP? If this were your organization, these would be important questions to get answered. This is why examining the org chart early on is generally a good idea.

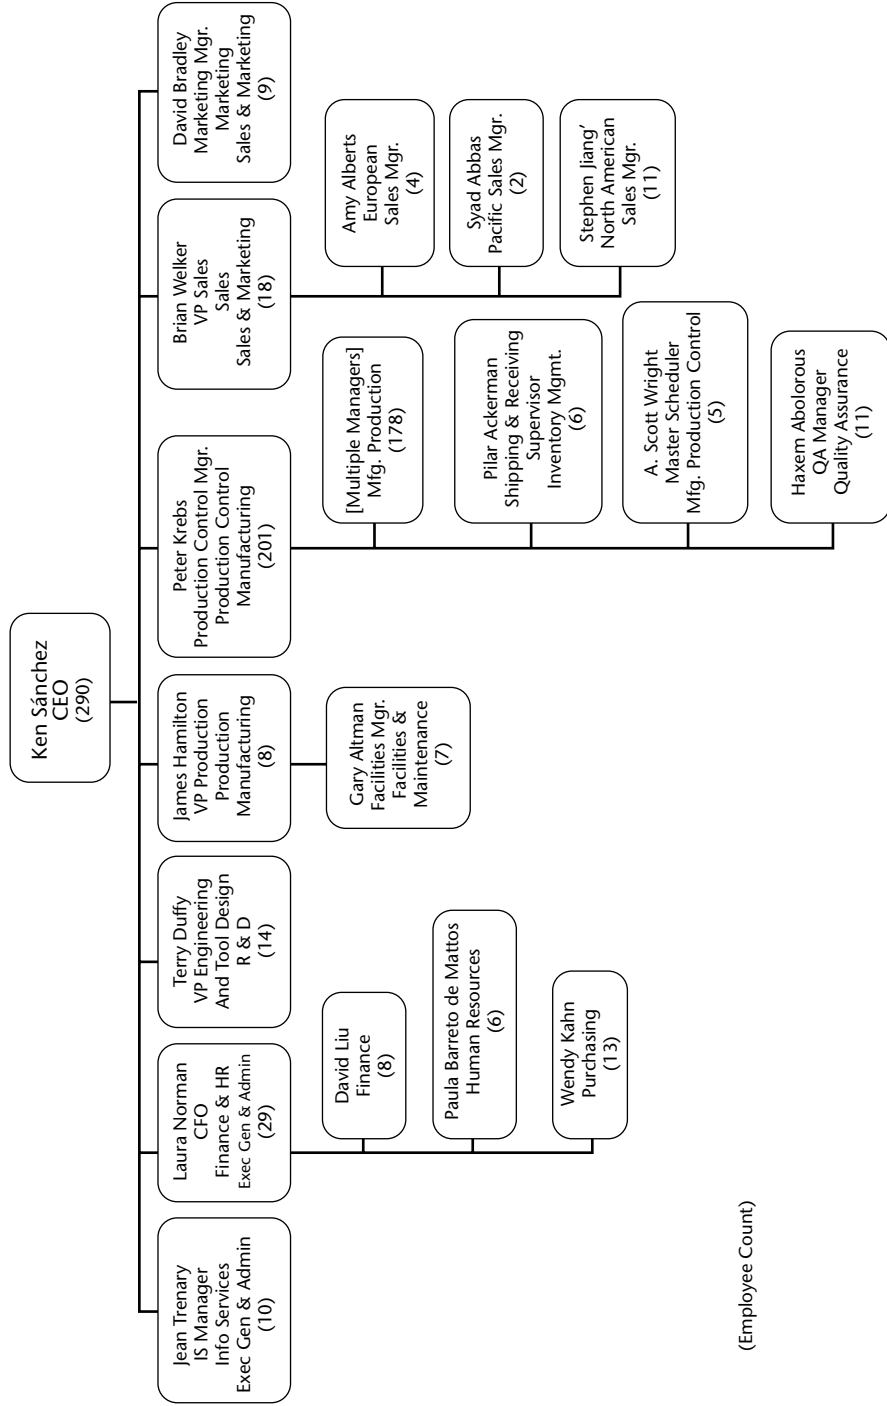
Given the number of senior executives and the size of the organization, you would probably plan on more than 10 but fewer than 20 interviews at Adventure Works Cycles. Depending on availability, expect this to take a minimum of a week, more likely one and a half to two weeks.

The rest of this chapter highlights key parts of the requirements definition process in the context of the Adventure Works Cycles case study.

Once you've completed the project requirements definition step, you will be ready to begin designing your dimensional model, as covered in Chapter 2.

Adventure Works Cycles Enterprise Business Requirements

Begin with a series of enterprise requirements interviews to get a broad sense of the important business processes and their business value, and to identify potential sponsors. Summarize each interview by grouping the individual requirements into common analytic themes. This brings a useful structure to what is often a fairly scattered interaction. While the same analytic requirement may come up several times in the conversation, you need to summarize it only once. The typical summary for an hour-long interview takes about three to four pages. Each summary should include the business narrative along with example analyses and potential data problems. The narrative should capture an estimate of the value of each business process.



(Employee Count)

Figure 1.6 The Adventure Works Cycles organization chart

We've included an example interview summary for Brian Welker, the VP of Sales. The following summary gives you much of the information you need to understand the business processes and the dimensional modeling decisions made in the next chapter. To get a more complete understanding, read through the additional abbreviated interview summaries at the book's web site, www.MsftDWTToolkit.com, for the following people:

- Ken Sanchez, CEO
- Peter Krebs, Production Control Manager
- David Bradley, Marketing Manager
- Mary Gibson, Internet Channel Analyst
- David Liu, Finance Manager

Adventure Works Cycles Vice President of Sales, Brian Welker, is a big believer in the power of information. He and his group have taken over one of the IT organization's analysts full time just to generate reports and analyses. The team decided to interview Brian early on to get some positive reinforcement.

After you've discussed roles and responsibilities, you might start the interview by asking Brian: How do you tell when you're doing a great job? If he's as smart as you think, his answer might be: "When my sales planning is accurate, when my sales grow, when I can leverage special offers, and when I have good customer satisfaction." That would be quite an answer, but it's a gold mine to drill down into his information needs. Your job is to react to his answer point by point, drawing out more detail. The following summary captures the results of that drill-down.

Example Interview Summary

Interviewee: Brian Welker, VP of Sales

Date: 7/25/2005

Interviewer: Joy Mundy

Scribe: Warren Thornthwaite

Additional attendees: Stuart Ozer, Carolyn Chau, Joy Byrd, Dave Wong

Roles and Responsibilities

Brian Welker is head of the sales organization. He's responsible for sales to Resellers, which was \$37 million last year, or about 70 percent of total sales. He has 17 people who report to him, including 3 regional sales managers. Brian is excited about his team and eager for them to be successful. They are all "bike freaks" who love to ride bikes and love to talk about them—perfect bike sales people. Brian is measured on achievement of the total Reseller sales target for the year.

Information Requirements

Brian is particularly frustrated with how difficult it is to get information out of the company's systems. When he asks for a report, it can take days or weeks to get the information. Often he's told "It can't be done." The major analytic areas that Brian works with are as follows:

- *Sales planning*: Planning for the year begins in the fall of the previous year with the Sales planning process. Sales territories are based on geography. All new customers are assigned to a sales territory when they place their first order based on where they are located. Sales planning includes looking at the following:
 - *Growth analysis*: Overall market, new products, new geographies, new sales people.
 - *Customer analysis*: Who are the top customers, how have they changed over the last year?
 - *Territory analysis*: Where are top customers located, what are the current sales territories, and how balanced are they? How does this map to sales regions?

Brian and Ramesh Meyyappan (the analysts who work in IT) also look at sales by sales regions, which are groupings of customers based on the state where they're located. Regions overlap sales territories and are based on seasonal buying patterns and regional preferences. Being able to group historical sales like this helps the sales team do a better job of forecasting monthly sales. They usually look at regional sales a lot during the sales planning cycle and then compare actuals to the forecast during the year. Every year, they change the regions a bit to line up with changes in buying patterns.

Ramesh does all of this data analysis for the Sales forecasting and quota assignment process in a spreadsheet. The spreadsheet includes territory growth factors, allocations, and manual adjustments. The planning process is totally manual and takes a couple of months in the fall, and maybe a week per month to do the reporting during the year.

During the annual planning process, Brian wants to be able to see reseller customer orders by year by customer territory, regardless of the sales rep assigned to the territory. In previous jobs, Brian would adjust the size of the territories by moving customers from one territory to another with the goal of making the territories more even. He has not done this at AWC yet, so all customers are still assigned to their original territory. Sales reps can be reassigned to different territories, usually when a sales rep leaves.

- *Sales performance*: Once the planning process is done, Brian wants to see sales according to the new territory assignments, all the way back

through history so he can compare with actuals as they come in. At any time, Sales must be able to re-create historical sales and commission reports based on what happened at the time of the order, not which territory gets credit today.

Brian also wants to look at orders from a sales rep perspective. The first thing he wants to see at the start of the week is how his sales reps are doing year to date. If Brian sees a problem in the higher level data, he wants to be able to drill down to detailed orders for individual reps. Of course, Brian has other reports he would like to see: for example, top 20 customers and orders by Reseller versus online.

- *Basic sales reporting:* Brian wants to take better advantage of the customer information buried in the orders transaction system. The sales reps would really appreciate it if they could get a list of the customers in their territory ranked by orders. Because most sales go to a small percentage of Resellers, the sales reps would concentrate on making sure those important customers are happy.

Beyond this, Brian knows that 17 percent of 2002 customers did not reorder in 2003. And to date in 2004, he still has not heard from an additional 17 percent or so. His sales people could use this information to bring the best of these customers back to the business.

- *Price lists:* The fact that the sales reps are out in the field most of the time makes it difficult for them to keep their price lists current. The price list changes fairly often, but only a few things on the list change. It would be great to get a report to the sales reps that flagged changes and special offers, and maybe even highlighted the relevant customers.
- *Special offers:* The special offers could be a great sales tool. AWC just finished an inventory clearance sale on the silver Mountain 500s. The color didn't sell as well as others, resulting in too many in stock at model changeover. Mary Gibson, the marketing assistant for mountain bikes and David Liu in Finance put their heads together to come up with ideas to stimulate demand and came up with a 40 percent off offer. This is something the sales people can work with, but it's a random process. They'd like a report that shows which of their customers bought a lot of the product that is on special offer to see if they're interested in more at a great price.

Brian would like his sales people to start with the biggest potential customers first and keep selling down the list until they run out. Actually, Brian thinks the business would be better served if they contacted the more profitable customers about special offers first. Some of the biggest customers are big because they scoop up specials, which don't make a lot of money for AWC. That's another thing: The sales reps need to know when out of stocks occur on special offers.

- *Customer satisfaction:* Brian would like to create some measures of customer satisfaction and has been trying to get more information out of the customer care system lately, with limited success. He would love to be able to track calls by complaint type, product, sales region, and customer to get a sense for customer satisfaction and product quality. He also thinks comparing order date and ship date in the sales data to identify late orders, and determining the percentage of returned items might be indicators of customer satisfaction. This would make a great start at a customer satisfaction scorecard.
- *International support:* The company has been growing internationally, but the transaction systems haven't kept up. The systems do take orders in multiple currencies, but none of the descriptions has been translated from English. This is a problem for the sales people, who have split up the product list and done the translations themselves. This doesn't work in the long run because no one knows if they've translated the information correctly. All materials must be bilingual to comply with Canadian law. The product tags and documents are already bilingual, but the sales materials are not.

Additional Issues

Brian expressed a frustration on the part of his sales force about the difficulty they have using existing reports. It is our sense that Brian would like an analytic system that provides his sales reps most of the information they need in a standard format with just a few keystrokes. The time zone differences make it hard for some of them to get live support from headquarters. If they need to get custom information, he would like it to be easy for them to get it themselves.

Success Criteria

Brian would like the system to provide him and his team with:

- Easy access to basic sales data for the whole field organization
- Flexible reporting and analysis tools
- All the data in one place (especially sales and forecast data)

Requirements Summaries: What's Missing?

The short answer to "what's missing" is "a lot." This particular summary doesn't include much indication of business value. And one summary will never provide the range of business requirements you get from a full set of interviews. In addition, you would need IT and business analyst representation to get a sense for the level of effort or even availability of some of the data requested. However, there is enough information here to help you make sense of the next few steps.

Analytic Themes and Business Processes

The interview summaries have already grouped similar requirements into common analytic themes. We've found that it's easiest to identify the business processes that support each analytic theme as you write the summary. For example, Sales Planning is an analytic theme that is supported by information from the orders business process. Figure 1.7 is an example of how you might distill Adventure Works Cycles' business processes from the analytic themes based on Brian Welker's interview summary. Your final list will include many more themes and business processes.

The Supporting Business Process column in Figure 1.7 allows you to group themes that rely on data from the same business process and to identify themes that require data from multiple business processes. Recall that themes that rely on data from multiple business processes, or second level themes, are more difficult to implement. Look for the highly leveraged opportunities where several themes can be delivered with data from a single business process. In most organizations, this opportunity is the sales business process, which Adventure Works Cycles calls *orders*.

Once you fill in the complete list of themes, re-sorting it by business process will reveal that many of the business requirements rely on data from a few business processes. For Adventure Works Cycles, data from the orders business process is all that is needed to enable a wide range of analytic requirements, many of which are not on Brian's list in Figure 1.7. As the full set of requirements would reveal, Sales is not the only department interested in orders data. This is important because it adds to the overall business value of implementing the orders business process. All or most of the following requirements could be met with data from the orders business process:

- Sales planning input
- Basic sales reporting
- Special offers
- Production forecast input
- Product planning and monitoring
- Internet customer demographics (with a well-designed customer dimension)
- Customer profiling
- Customer loyalty program
- Standardized currency reporting

Analytic Theme	Inferred or Requested Analyses	Supporting Business Process	Comments
Sales Planning	- Reseller historical orders analyses	- orders	By customer, by territory, by sales region (from state)
	- Sales forecast	- orders	Forecast is a business process that uses orders data as an input
Sales Performance	- Orders by current territory	- orders	
	- Orders by original territory	- orders	
	- Sales rep performance report	- forecast	Orders and forecast by sales rep
Sales Reporting	- Resellers ranked by orders in a given territory	- orders	
	- Churned customer list	- orders	Customers who have not ordered in X months
Price Lists	- Current price list	- orders	This is a connectivity issue, not a data warehouse issue
Special Offers	- Relevant customers by territory based on orders history	- orders	
	- Inventory status (out of stock)	- inventory	
Customer (Reseller) Satisfaction	- Calls by complaint type, product and customer attributes	- call tracking	
	- Order metrics of satisfaction	- orders	e.g. due date versus ship date
	- Returns by reseller by return reason	- returns	
International Support	- Local language translations of Product descriptions	- n/a (product dimension)	This is a transaction system problem. We need to make sure we can handle multiple languages in the DW/BI system, but the source system has to capture them when new products are created.

Figure 1.7 Analytic themes and supporting business processes from the interview summary

THE NAME GAME IN ACTION

The business processes listed in Figure 1.7 do not necessarily match up with the terms we heard from Brian Welker in the interview. He's the VP of Sales, and tends to use the word "sales" to describe the information he'd like to track. As it turns out, the term "sales" has a specific meaning carefully defined by the accounting department based on order dates, ship dates, and accounts receivables. Brian is not actually measured on sales; he is measured on orders. The process of determining business process, table, and attribute names starts here. You need to be as precise as possible in your use of terms. In fact, as soon as you understand the distinction between sales and orders, you should rewrite the interview summary to reflect it. We left the ambiguity in our summary to make the point.

Brian also discussed a few consolidated analytic themes that require data from more than one business process (from Figure 1.7):

- *Sales performance*: Orders and forecast
- *Customer satisfaction*: Call tracking, orders, and returns

The prioritization process focuses on business processes because these are the coherent units of work for the DW/BI system. In describing each business process, you need to tie it back to the analytic themes it supports, so senior management can assess its business value. The business processes from Brian's interview are shown in Table 1.5. Brian's supported analytic themes are underlined; the rest came from other interviews.

Table 1.5 Business Processes Derived from Brian Welker's Interview

LETTER	BUSINESS PROCESS	SUPPORTED ANALYTIC THEMES
A	Orders	<u>Orders reporting and analysis</u> , <u>orders forecasting</u> , advertising effectiveness, <u>customer satisfaction</u> , production forecasting, product profitability, customer profitability
B	Orders forecast	<u>Sales performance</u> , business planning, production forecast
C	Call tracking	Call center performance, customer satisfaction, product quality, customer profitability, product profitability
D	Returns	<u>Customer satisfaction</u> , product quality, customer profitability, product profitability, net sales

By the time you're finished with the interviews and write-ups, you should have a list of business processes that represents the major activities of your organization: 25 to 50 is a reasonable count at this point. You'll leave some items off the list, even though they will be discussed in the document and with senior management. In Brian's case, price lists and international support are important issues to his organization, but they are transaction system problems because they involve enhancing transaction systems or building new IT infrastructure. You can help with better reporting, but you shouldn't be dealing with connectivity and data capture issues if you can avoid it.

Adventure Works Cycles Bus Matrix

Business processes are the units of work in building the information layer of the BI system. The bus matrix shows the business processes and the dimensions needed to support them. The more you know about how the business and its supporting systems work the easier it is to define the bus matrix. Your full version of Table 1.5 would be the starting point for an initial version of your bus matrix. Figure 1.8 shows a high-level version of the Adventure Works Cycles bus matrix for the business processes from Table 1.5.

The bus matrix gets another level of detail during the dimensional modeling process when you add in the different types of fact tables. You'll get more on this in Chapter 2.

Business Process	Dimensions										
	Date	Product	Employee	Customer (Reseller)	Customer (Internet)	Sales Territory	Currency	Channel	Promotion	Call Reason	Facility
Sales Forecasting	X	X	X	X	X	X	X				
Orders	X	X	X	X	X	X	X	X	X		
Call tracking	X	X	X	X	X	X				X	
Returns	X	X		X	X	X	X		X		X

Figure 1.8 The Adventure Works Cycles bus matrix

The Adventure Works Cycles Prioritization Process

As we described earlier in this chapter, the prioritization process is a facilitated session where the DW/BI team presents the preliminary project findings and facilitates discussion with key business representatives. The primary goals of the prioritization process are to ensure a common understanding of the requirements findings and business processes, and prioritize the business processes based on business value and feasibility. The top priority business process becomes the focus of the initial DW/BI project.

Figure 1.9 shows the resulting grid from the Adventure Works Cycles prioritization session. It includes a few more business processes than we listed in Table 1.5, but it's still not a full set.

Note that there are two items on the grid that are not actually business processes: Customer and product profitability are consolidated themes that senior management has expressed significant interest in analyzing. These have been included on the grid to show their importance, but they are far over to the left to indicate the difficulty involved in building all the needed business processes. Given the number of analyses supported by data from the orders business process, it should come as no surprise that orders is the top priority theme. The team should get to work on this right away!

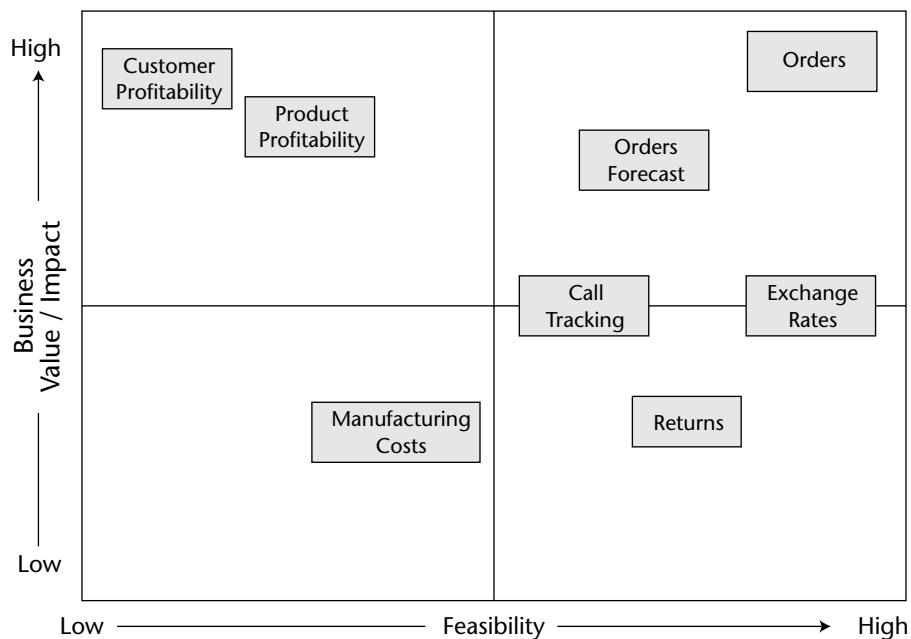


Figure 1.9 The Adventure Works Cycles prioritization grid

Business Requirements for the Orders Project

Getting to work on the orders business process requires holding an additional set of interviews to drill down on orders-related analyses. The team needs to understand several issues that were raised in the enterprise requirements process, such as what the different kinds of regions that people described are. The team would get more specific about the kinds of new reports and analyses people want to see as input to the BI Application track.

All of this information becomes the grist for the Adventure Works Cycles business dimensional modeling process case study in Chapter 2.

Summary

This chapter concentrated on the early tasks in the Lifecycle involving business requirements gathering, prioritization, and project planning. We gave special emphasis to the importance of understanding and documenting the business requirements.

The first part of the chapter described a process for gaining sponsorship, defining and documenting the enterprise-level business requirements, prioritizing the opportunities with senior business people, and gathering project requirements related to the top priority business process. This process also included the challenging task of distilling the analytic themes down to the business processes that provide the underlying information.

The second part of the chapter went through an abbreviated example of the business requirements gathering process at Adventure Works Cycles. The interview summary for the VP of Sales provided the analytic themes that tied to the business processes that fed into the bus matrix and the prioritization process.

These upfront business-related phases of the DW/BI project are by far the most important. Unfortunately, they can be intimidating for technologists. Do not resist or avoid the requirements gathering phase of the project. The resulting understanding of the business issues, their priorities, and the data that supports their solution is priceless for the DW/BI team. The requirements document will be your reference point for all major decisions from here on out. You get huge value just from the content of the document alone.

But wait, there's more! The requirements gathering process also helps you build positive working relationships with the business people. As the business people participate in the requirements process, they see that you've done your homework. You understand them, you speak their language, you want to help solve the problem—in short, you get it.

If that's not enough to convince you, there are even more benefits to this process. Not only do you get documented requirements and better relationships, you gain active user support. As the business folks begin to understand your vision for an information solution, they see how your success ultimately leads to their success. They begin to see how their involvement will improve the chances of success.

