Part 1

Requirements, Realities, and Architecture

Chapter 1: Defining Business Requirements
Chapter 2: Designing the Business Process Dimensional Model
Chapter 3: The Toolset
Chapter 4: System Setup

This first part of the lifecycle is where you lay the foundation for your success. Working with the business folks to understand and prioritize their requirements for analytics as we describe in Chapter 1 helps you set specific goals for your first pass through the lifecycle that are both valuable to the organization and achievable in a reasonable timeframe. Your understanding of the business requirements becomes the basis for designing a flexible, usable, high-performing dimensional model in Chapter 2.

What you learn in the first two chapters helps you tackle the architecture and technology track at the top of the Lifecycle. Your business understanding helps you determine what architectural components are important for your DW/BI system. Once you know the problem, you can identify the specific functionality you need, and where that functionality will come from in the Microsoft SQL Server toolset. This, in turn, allows you to make decisions on the server configurations and disk subsystems that will form the basic infrastructure of your DW/BI system.

Part 1 is about getting the lay of the land before you decide what you are going to build and where you will build it. Your primary focus here is on identifying the most promising business opportunities and designing the data structures and system architectures needed to deliver them. By the end of this section, you should have all the pieces in place for you to dig into the development work of creating the DW/BI system database. Skip this section at your peril.

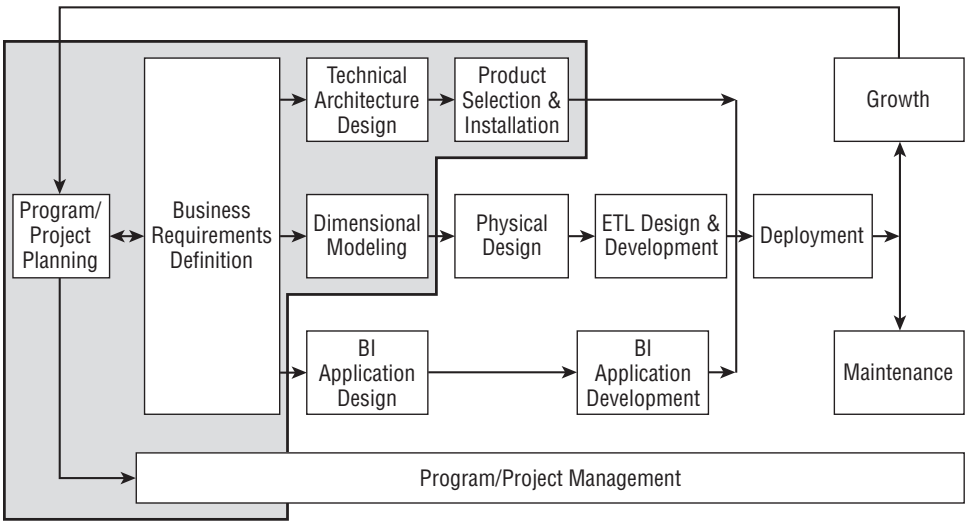

The Kimball Lifecycle steps covered in Part 1

CHAPTER 1

Defining Business Requirements

Building the foundation.

Business requirements are the bedrock of the successful data warehouse/business intelligence (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 and mapping their analytic requirements back to the core business processes (such as orders, page views, or account transactions) that generate the needed data. These business processes are the building blocks of the DW/BI system. After the requirements are documented, we offer a technique for working with senior management to prioritize the implementation of those business-process–based projects. We also illustrate these tasks with an example based on Microsoft's sample database business, Adventure Works Cycles.

As Figure 1-1 illustrates, the Business Requirements Definition step is the foundation of the Kimball 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, too.

RESOURCES **If you skipped the Introduction to this book, you should at least go back and read the overview of the Kimball Lifecycle because it is the organizing framework for this book and for implementing a successful DW/BI system.**

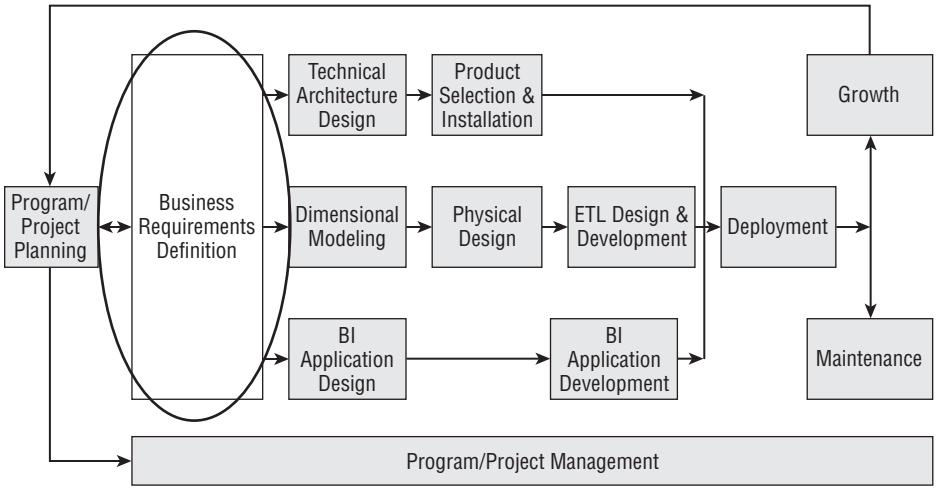

Figure 1-1: The Business Requirements Definition step of the Kimball Lifecycle

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
- The steps used to define enterprise-level business requirements, including the interview process, synthesizing requirements into their underlying business processes, developing the enterprise analytic data framework called 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 goes into a typical requirements summary document and how it links to business requirements for analytics and business process implementations

RESOURCES **Throughout the book, we provide specific references to the various titles in the *Kimball Toolkit* library to help you find more details on the concept or technique described. Each book in the *Toolkit* series is described in the Introduction to this book.**

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, but most DW/BI folks are technologists at heart. We like the certainty of computers and programming and shy away from the vague uncertainties of the business side.

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 points of view. You'll be working in a non-technical, highly ambiguous, politically sensitive environment. Are you feeling queasy yet? 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 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 learn more about the business.

NOTE **Perhaps your organization uses the "agile" development methodology. If so, then you have heard this story already! In the agile approach, projects are owned and driven by the business users. To learn more about this approach, see "agile software development" on www.wikipedia.org. For guidance relating agile to DW/BI system development, see *The Kimball Group Reader*, pp. 109–112.**

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. We understand you bought this book to learn about the SQL Server DW/BI toolset, but SQL Server is just a tool. Your success in using that tool in your organization depends on your understanding of the organization's unique requirements and priorities for business intelligence.

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.

Adventure Works Cycles Introduction

It always helps to see new concepts in the context of a specific example. Since everyone's organization is different, we'll use some of the business requirements for Microsoft's demo database company to illustrate the process of defining business requirements described in this chapter.

The current SQL Server sample business intelligence databases are based on a fictitious company called Adventure Works Cycles, a multinational manufacturer and seller of bicycles and accessories. The database and associated samples are not part of the software distribution set. Instead, you download them from the Microsoft code-sharing site called Codeplex (Search for "SQL Server Samples Database" at `http://www.codeplex.com`) or download the database from the Wiley web site at `www.wiley.com/go/MsftDWToolkit2E`. You will need to download and install the SQL Server 2008R2 version of the sample databases to follow the examples later in this book.

DOWNLOADS **You can find several detailed documents illustrating what the business requirements-gathering process might look like at a company such as Adventure Works Cycles on the book's web site (`http://kimballgroup.com/html/booksMDWTtools.html`). These include interview summaries and additional background information.**

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 Kimball 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? Well, it depends 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) department of the organization. The IT-driven DW/BI project gets cranked up 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, speaking to a room full of business users at the project kick-off, and getting spending approval for your new server. 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.

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 don't get out of the IT group much, you'll

need to investigate 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. Obtaining business sponsorship is fairly easy and well worth the effort. Good business sponsorship can provide the resources and support you need to deliver real business value.

RESOURCES **Learn more about developing sponsorship in *The Data Warehouse Lifecycle Toolkit, Second Edition*, pages 16–21.**

Defining Enterprise-Level Business Requirements

The successful DW/BI effort is an ongoing program guiding the multiple, iterative projects that build out the DW/BI system. Before you concentrate your efforts on specific projects within the larger data warehouse endeavor, you need a broad enterprise perspective that will help you set priorities and make better, more flexible implementation decisions. 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.

We almost always preface the first iteration of the Lifecycle with an enterprise business requirements definition project. This project is essentially a set of interviews, documentation, and a prioritization session with senior management. It provides a clear implementation plan for the DW/BI system, and it can be done in three to six weeks for most organizations.

Every DW/BI program has to keep the enterprise requirements context in mind. Larger organizations need to begin by establishing this broad understanding because it is rare for the DW/BI team to have such an enterprise-level perspective. Even in a smaller company or a departmental effort, the enterprise perspective will help build in flexibility and resilience. It is 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. By the same token, DW/BI systems that are well into their implementation will gain by taking a little time to validate their understanding of the enterprise business requirements. This usually results in significant changes to the DW/BI system strategy. Better late than too late.

Given this need for an enterprise perspective, it's best to preface your first Lifecycle iteration with a narrow-scoped enterprise requirements definition project as shown in Figure 1-2.

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.

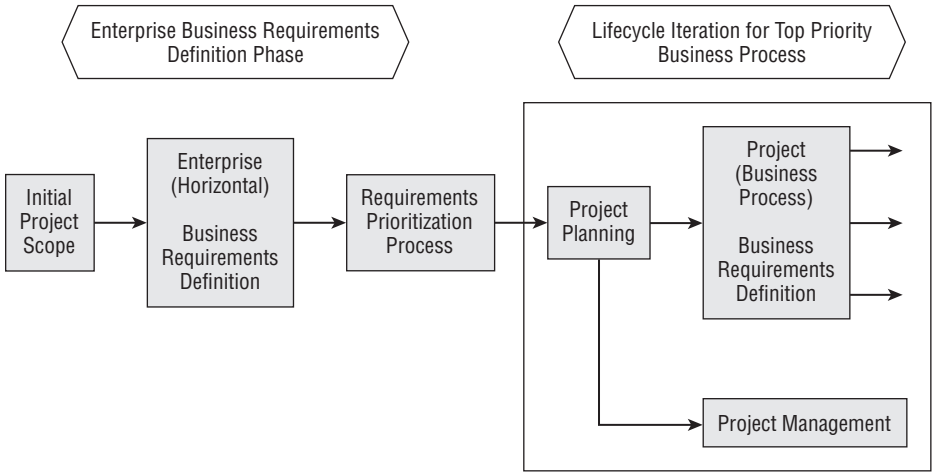

Figure 1-2: Prefacing the core Lifecycle with an Enterprise Requirements phase

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.

Establishing Initial Enterprise Requirements Project Scope

The initial scope usually covers only the enterprise-level requirements definition and requirements prioritization steps, leaving the detailed project implementation

plan for later when you have a much better idea of what the project needs to accomplish from a business perspective. The requirements and prioritization usually involve user interviews, interview write-ups, a few meetings, and the creation of the final requirements document. It typically takes three to six weeks (or more) depending on how many interviews you do.

COMBINING ENTERPRISE AND PROJECT REQUIREMENTS GATHERING

Some 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.

RESOURCES **Learn more about DW/BI project planning and management in Chapter 2 of *The Data Warehouse Lifecycle Toolkit, Second Edition.***

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 you see in Figure 1-3, the bulk of the work 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 and documenting an interview, and extracting its analytic requirements. As the interviews progress, we begin synthesizing what we've learned, identifying the business processes that support the users' analytic needs. At the same time, it's important to conduct initial data profiling to match analytic needs with data realities. At the end of the interview process, we build an initial bus matrix to summarize the business processes we've heard about during the interviews. We describe the bus matrix and each of the core subtasks in Figure 1-3 in this section, leaving the senior management prioritization session for its own section.

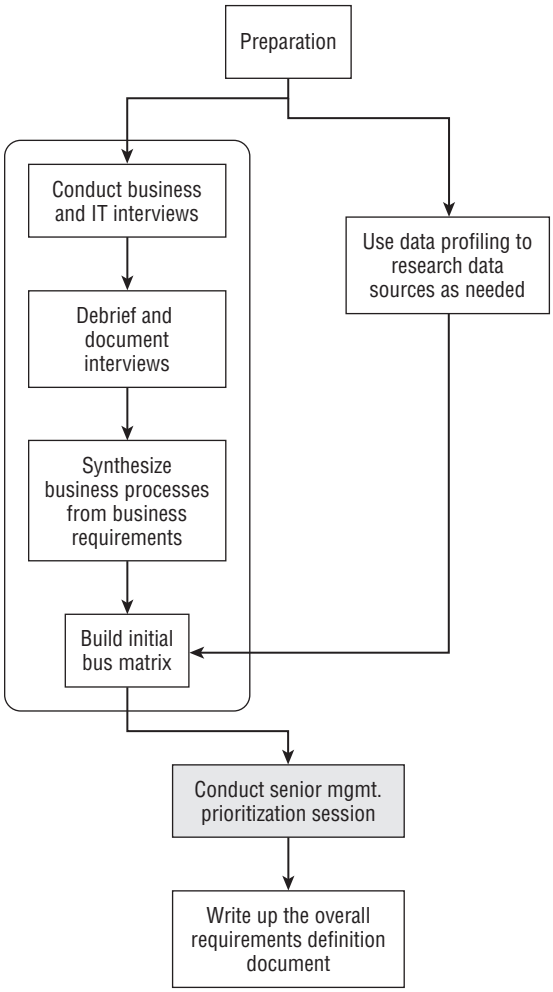

Figure 1-3: The enterprise requirements definition process flow chart

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. For the enterprise requirements pass, start the interview list with the CEO and 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 IT the most. If you've been at your company 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 and the quality of the data they collect.

NOTE **You are just making the list of whom to interview at this point, not the interview schedule. When you do start the schedule, begin 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 hinge 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.

RESOURCES **You can find additional information about preparing for interviews in *The Data Warehouse Lifecycle Toolkit, Second Edition*, pages 68–80.**

Adventure Works Example: Preparation

Much of preparation is about knowing your business. Not being a real business, Adventure Works does not have much strategic or market information available. However, you can get a good sense for the nature of any business with SQL Management Studio and a basic knowledge of SQL, or even better, a query tool like Report Builder 3.0. You also need permission to query the source system (or a copy of it, so you don't cause any transaction problems).

A few simple `SELECT`s and `SUM`s can tell you what products are selling well, where they are selling, who is selling them, and how this has all changed over time. At Adventure Works, you could quickly find out that over 80 percent of their business is bicycles, versus clothing or accessories, and almost half their business is overseas. From a time series perspective, the company has been growing rapidly and the internet sales channel is a major contributor to this

growth. Figure 1-4 shows a pivot table of the results of a query against the AdventureWorks2008R2 transaction database.

Sales Amount Row Labels	Column Labels FY 2006	FY 2007	FY 2008	% of 2008 Total
Bikes	$22,091	$28,180	$44,350	84.1%
Components	$1,167	$4,629	$6,003	11.4%
Clothing	$66	$751	$1,283	2.4%
Accessories	$37	$124	$1,077	2.0%
Grand Total	**$23,361**	**$33,684**	**$52,714**	**100.0%**
Sales in US dollars net of taxes and shipping				

Figure 1-4: Adventure Works sales by category by year

If you have a couple of hours, you can dig into other parts of the business, such as customer support, manufacturing, and finance.

DOWNLOADS **See the web content for more examples of these kinds of queries and results for the Adventure Works business.**

Conduct Business and IT Interviews

The key to success in requirements interviews is to remember your overall mission. You are designing a system to add significant long-term business value. The most common mistake in these interviews is to ask the business person what they want (or need). Asking this question is the equivalent of abdicating your design responsibility. You are saying, "Tell me what you want and I'll build it." At best, you will get a limited description of what the person wants to solve today's problem. For example, you may get a request to have all the data provided for a given report in an Excel format. This may sound great to you because it is easy to understand and execute. So, you go off to extract, clean, and load that data into a data warehouse, then set up a distribution system to put it into Excel and email it to the analyst every night. Once you are finished, you show it to the analyst, and they say, "It's nice, but what I really want is this other report in an Excel format." This is when we hear statements like "The business people don't know what they want!" Or, "The business people don't understand business intelligence!"

This lack of understanding is not their fault, and it's an easy mistake to avoid by changing the questions you ask in the interview. Remember, it's your job to design the system. In order to do that, you have to understand the business; you

need to know what your users do and how they use information to do it. Once you know this, the required system designs, data models, and BI applications all become clear.

The easy way to find this out is to start out with the simple question: "What do you do? Tell me about your roles and responsibilities." Explore each of the areas they describe in terms of the information they need, and the value they provide to the organization (or could provide with better, more accessible information). Be flexible and follow the leads provided by the interviewee. Avoid sticking to a predetermined script planned in too much detail. Interviewing is a very valuable skill.

RESOURCES **For more in-depth interviewing tips and techniques, see the articles in the *Kimball Group Reader* on pages 113 and 117.**

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. These challenges and opportunities often (but not always) line up with the strategic goals and initiatives of the organization.

You will also interview some of the 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.)

Debrief and Document Interviews

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. As you go through your notes, highlight and add comments to fully describe the following items:

- Common, repetitive business requirements themes
- Business processes (data sources) needed
- Business requirements for specific reports and analyses
- Misunderstandings or incomplete notes (the lead interviewer should keep a list of open issues)
- Data or other feasibility issues known to the team
- Success criteria

The individual interviews will yield a wealth of information including descriptions of the analytic requirements and their associated business processes, a starting point for the organization's overall information architecture, and a list of any feasibility issues such as poor data quality.

Identifying the business requirements for analytics is the hardest part. Depending on who you are talking to, similar analytic opportunities may be described broadly or specifically. As the interviews progress, you'll see common requirements repeated over time. For example, the marketing person responsible for internet promotions might describe an opportunity to improve promotion response and conversion rates by better targeting certain geographic and demographic subsets of the population. The person responsible for product promotions might describe an opportunity to improve conversions by offering product recommendations based on customer behaviors. Both of these opportunities could be grouped together under a broader heading called Improve Customer Acquisition.

Each of these broader analytic themes should have brief descriptions of the kinds of reports or analyses you heard in the interview. It should also include some sense of the business value of meeting the requirements. In other words, how much is improving customer conversion rates, or negotiating better prices and terms worth? Look for action words to identify these opportunities. Words like *improve*, *reduce*, *increase*, and *enhance* all lead to a business requirement we'd like to know about.

Do not put this review off. After a day of interviews, you will have a hard time remembering who you spoke with, let alone the details of what they said. Be cautious about scheduling too many interviews in one day. Our rule of thumb is four interviews and four debriefing periods per day.

It's a good idea to write up a summary document of each interview based on the annotated set of interview notes as soon as you can. This is more work because you need to summarize the various analytic areas covered in the interview, but it is a good communication and relationship-building tool with the business folks. Share this summary with the interviewee and ask for feedback; it shows that you listened to what they had to say and have an interest in helping. It also gives them a chance to clarify any misunderstandings and add any relevant items they overlooked.

RESOURCES **You can find additional information about conducting the interviews and debriefing on pages 80–85 of *The Data Warehouse Lifecycle Toolkit, Second Edition.***

Adventure Works Example: Interview Documentation

Adventure Works has close to 300 employees. The CEO has seven direct reports, and the vast majority of employees are manufacturing workers. In most small- to medium-sized organizations like this, you could build a solid set of business requirements for analytics by speaking to most of the senior staff and a subset of managers and analysts; maybe 15 people, plus or minus five.

The primary content of each interview summary write up will be a list of business analysis requirements. Each analytic requirement should also include a list of the business processes that generate the data needed to support the analysis, and any associated issues or concerns, such as data quality or availability problems. Figure 1-5 shows a summary of the business requirements identified by the VP of Sales.

DOWNLOADS **You can find an Adventure Works organizational chart and an example interview summary from the Adventure Works VP of Sales in the web content downloads.**

It should come as no surprise that most of the VP of Sales' business requirements for analysis are based on data from the orders business process. However, there are several other business processes that inform decision making for the sales department. Note that the VP has described a requirement for a customer satisfaction dashboard. By decomposing it to its underlying data sources, it becomes clear that you will need data from three major source systems: orders, call tracking, and returns. This means three iterations of the Lifecycle with three sets of ETL code, BI applications, testing, and deployment. Now is a good time to start educating and setting expectations.

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.

Business Requirements Category	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	Orders, Orders 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	Customer Satisfaction Dashboard	Multiple	This is a compound requirement based on several underlying business processes
	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 1-5: Business requirements and supporting business processes from the interview summary

Creating the Program Requirements Findings Document

The overall findings document for the enterprise-level requirements 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.

The bulk of the requirements document will be a list of the business processes and the business requirements they support. Each business process section should include some sense for the business value it would generate, the data quality and other feasibility issues associated with it, and the parts of the organization that will benefit from it. Some of the requirements 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.

Synthesize Around Business Processes

The goal here is to tie the business requirements back to the underlying data needed to make them happen. This is the primary factor in determining the level of effort required to deliver a solution to a given business requirement.

As you extract the business requirements from the interview summaries, you need to dig into each opportunity to identify the business process (or processes) that generate the data needed to perform the desired analyses. For example, the requirement to reduce purchasing costs through better contract negotiations can be supported by historical data from the purchasing transaction system. You convert from requirements 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 module, which translates into a single pass through the Kimball Lifecycle process. (Refer back to the related sidebar, "Business Process: the DW/BI System Unit of Work," for more information.)

Although many business requirements need information only from a single business process, one challenge you will face is that some requirements need data from multiple business processes to meet the overall analytic needs. Customer and product profitability analyses are good examples. The "customer scorecard" may sound like a single analysis, but it actually requires data from many separate business processes. We call these *consolidated requirements* (or, occasionally, *second level business processes*) because they cannot be completely built until data from all prerequisite business processes have been loaded into the data warehouse.

Tying business requirements to the underlying business processes helps determine the level of effort needed to support a given requirement. If a requirement

must have data from more than one business process, it will take more than one iteration of the Lifecycle. While these passes can happen in parallel with additional resources, there is no way around doing the work.

Converting from business requirements to business processes involves thinking about which business processes are required to support each requirement. For example, you can get a start at improving promotion response rates with data from a single business process: orders. As long as the order data captures a promotion code, you can calculate response rates and purchase amounts associated with each promotion. Better analysis would need to include data from the promotions business process; specifically, what was the total population of prospects who received a promotion. Now you can see who responded from orders, and who didn't from promotions. In the best case, you may want to bring in demographic data for prospects and use a data mining model to help identify the characteristics of prospects that are more likely to respond to a given promotion. Ultimately, complete promotions analysis would need data from at least three business processes (at least one, demographic data, comes from an external source). If improving promotion response rates is a top business priority, this decomposition process will help you show why it will take three iterations to complete the required data set. It will also show that basic analysis can begin once the first business process (orders) is loaded.

The most demanding type of business requirement 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.

Adventure Works Example: Enterprise Requirements Documentation

This section walks you through the steps to create a requirements findings document for the Adventure Works example. The easiest way to create the requirements summary is to start with the requirements from one interview, such as those shown in Figure 1-5, and add in requirements from subsequent interviews. Often requirements from different interviews will fall into existing business requirements categories.

To understand what analyses will be enabled by each iteration of the Lifecycle, you need to re-sort the requirements by business process. Table 1-1 shows a subset of the analyses enabled by each individual business process. The Letter column is meant to serve as a shorthand reference you may use later in the prioritization process.

This table includes analytic requirements from across the Adventure Works enterprise. The VP of Sales' requirements are underlined. You should keep track of where the individual requirement for analysis came from within the organization. This will help you track back to the originator later on, after you re-sort the requirements by underlying business process.

Table 1-1: A subset of business processes derived from the requirements interviews

LETTER	BUSINESS PROCESS	SUPPORTED BUSINESS ANALYSES
A	Orders	Orders reporting and analysis, orders forecasting, advertising effectiveness, customer satisfaction, production forecasting, product profitability, customer profitability
B	Orders forecast	Sales performance, business planning, production forecasting
C	Call tracking	Call center performance, customer satisfaction, product quality, call center resource planning, customer profitability, product profitability
D	Returns	Customer satisfaction, product quality, customer profitability, product profitability, net sales

Building the Initial Data Warehouse Bus Matrix

As you identify the business processes needed to support each analytic requirement, you will also add those business processes to an enterprise data framework 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 1-6 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.

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 processes — 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.

Value Chain ↓

Business Processes	Dimensions						
	Date	Product	Vendor	Shipper	Dist Ctr	Store	Promo
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 1-6: Example enterprise bus matrix for a retail company

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 analytic requirements they support (and the value those requirements represent) become the major inputs to the next step in the requirements definition process: a prioritization session with senior management.

Adventure Works Example: Bus Matrix

As you go through the interview process, you may be surprised to discover a many-to-many relationship between people and data. That is, people need access to data from multiple business processes, and many people often want to look at data from the same business process, but from their own business perspective. For instance, people in marketing might be interested in orders data by product over time, while folks in sales might be interested in the same orders data, only by sales rep and region. As a reminder, this means you would design the orders data model at the atomic level so the same data set could be used to support both sales and marketing. It does NOT mean you should have two separate data marts, one for each department.

Figure 1-7 shows the start of a bus matrix for Adventure Works based on the interview with the VP of Sales that is included on the book's web site.

Adventure Works Cycles: Enterprise Data Warehouse Bus Matrix Entries for Sales Department											
	Dimensions										
Business Process	Date	Product	Employee	Customer (Reseller)	Customer (Internet)	Sales Territory	Currency	Channel	Promotion	Call Reason	Facility
Orders Forecasting	X	X	X	X	X	X	X	X			
Orders	X	X	X	X	X	X	X	X	X		
Inventory	X	X	X								X
Call tracking	X	X	X	X	X	X				X	
Returns	X	X		X	X	X	X		X		X

Figure 1-7: Bus Matrix business processes referred to in the VP of Sales interview

Here you see that one person, the VP of Sales, is interested in data from five different business processes.

Once you've got the requirements documented, it will become clear that you can't deliver them all at once. The prioritization process will help you and your organization figure out the appropriate order of events.

Prioritizing the Business Requirements

If you're a technical person, it's safe to say the prioritization process we describe here 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 90 minutes. 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 Adventure Works example shown in Figure 1-8. This is an interesting exercise in negotiation and education and can easily take another hour and a half or more.

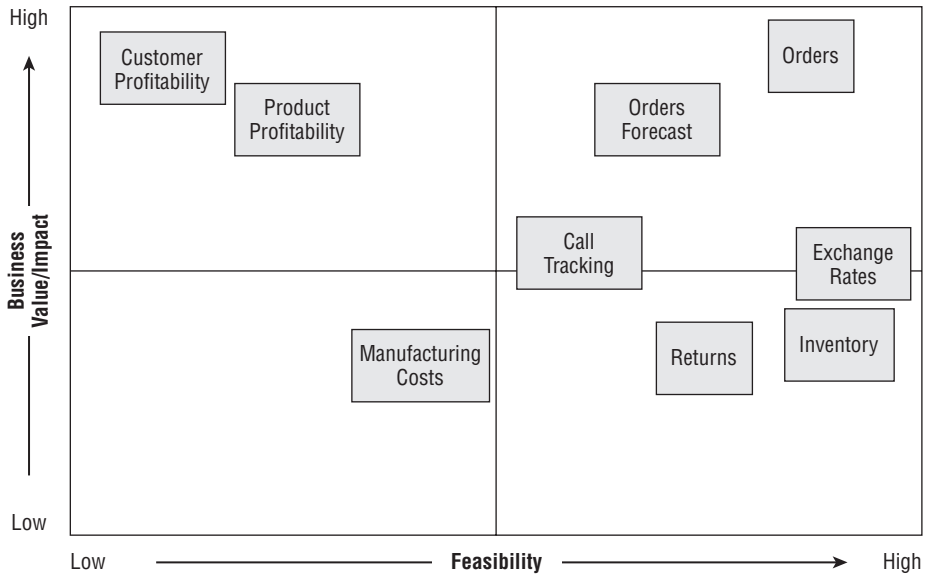

Figure 1-8: Example prioritization grid from Adventure Works Cycles

The prioritization grid is deceptively simple: Study it carefully. The Y axis is about relative 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.

The X axis represents the level of effort each business process will take to implement. It is stated in terms of relative 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, such as customer or product profitability.

The prioritization 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.

When reviewing Figure 1-8, note that there are two items on the grid that are not actually business processes. Customer profitability 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. The team should get to work on this right away!

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.

RESOURCES **Learn more about the requirements prioritization process on pages 91–93 of *The Data Warehouse Lifecycle Toolkit, Second Edition.***

Once you've completed the prioritization session you can finalize the overall requirements document by including the resulting list of prioritized business processes. At this point 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.

RESOURCES **You can find additional information about creating the requirements deliverables on pages 85–91 of *The Data Warehouse Lifecycle Toolkit, Second Edition.***

Revisiting the Project Planning

Now that you have a clear idea of your top priority business process, the data it generates, and the business requirements it supports, 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 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 the project planning and business requirements definition steps in Figure 1-1, but the backward flow is not as large 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-Level 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 provides 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 cleaning, 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.

Applying this interview approach to the Adventure Works example, the team will need to hold an additional set of interviews to drill down on orders-related analyses before it can start working designing the Orders business process dimensional model. The team needs to understand several issues that were raised in the enterprise requirements process. We'll look at the impact of a detailed understanding of questions like, "What is a customer?" and, "How do we determine the Sales Territory?" in the next chapter, which is on dimensional

modeling. The team should also get more specific about the kinds of new reports and analyses people want to see as input to the BI Application track.

In fact, all of the information gathered in this second pass becomes the grist for the Adventure Works Cycles business dimensional modeling process case study in Chapter 2.

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, 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 BI 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 learn 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 requirements 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.

RESOURCES **To learn more about defining project level requirements, see pages 93–101 of *The Data Warehouse Lifecycle Toolkit, Second Edition.***

Search kimballgroup.com for the topics "Business Requirements" and "Business Acceptance" for several related articles.

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.

We 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 tying the analytic requirements down to the business processes that provide the underlying information.

The chapter also summarized some of the business requirements that might be found at a company like Adventure Works Cycles. The VP of Sales provided a set of analytic business requirements that tied to the business processes that fed into the bus matrix and the prioritization process.

These upfront business-related phases of each DW/BI project iteration are 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 for the DW/BI system and for the business itself.