Chapter 1

Principles of a Reporting Infrastructure

Businesses can't be run without some kind of reporting system. Management needs to know what is going on in the company in order to make accurate business decisions, so they need reports. Because you are reading this book, it is obvious that you are considering Reporting Services to get those reports out.

Before you can make a solid decision, you must know what you are looking into. What is Reporting Services and what can it do for you from a business perspective?

In this chapter we will:

- Define a reporting infrastructure
- Figure out why an enterprise-reporting infrastructure is necessary
- Help you decide if you need to use Reporting Services as your reporting infrastructure

Reporting Infrastructure Basics

Before you can decide whether to spend all of the time and money necessary to construct a reporting infrastructure, you need to know what one is and how it is used.

What Is a Reporting Infrastructure?

This sounds like a simple question, but it is more involved than it seems. You may think of a reporting infrastructure as any old type of reporting system that your business may use — from pen and paper to colossal database systems. But there is more to it than that.

A reporting infrastructure is not only the data source but also the mechanism that displays the report in a human-readable format, as shown in Figure 1.1. For instance, if you keep a grocery list in your head and you need someone to go to the store for you to pick some things up, you write down a list of what you need. The data source is the list you keep in your head, and the pen and paper are the mechanism used to display the report in human-readable format.

Of course, most modern reporting infrastructures are considerably more complex than that. Businesses use a variety of systems, depending on the data they keep and their reporting needs. Some may use a simple spreadsheet to tally sales at the end of the day, some may use a Microsoft Access database with custom reports, and others may use an enterprise-reporting system. That is where Reporting Services starts to make sense.

An enterprise-reporting infrastructure like Reporting Services is used to display information from disparate data sources in a format that people can understand and use in making business decisions. What does that mean? Let's consider an example.

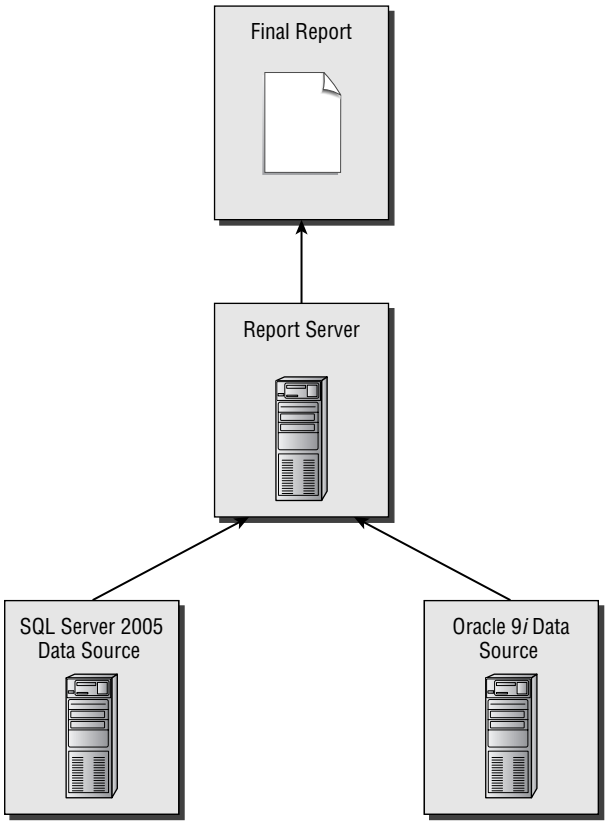

FIGURE 1.1
An enterprise-reporting infrastructure is used to combine data from different sources into a human-readable format.

Most large companies have a number of databases that contain data for different departments and different purposes. Human Resources may have a database for keeping track of employee records, the Maintenance department may have a database for keeping track of equipment, and so on.

Management needs to keep track of which employees have what equipment so that when someone gets promoted they can get new equipment needed for their new position. Or, if someone gets terminated, management can get all of the equipment back from that employee.

Maintenance does not have any records of who gets promoted or terminated in their database, and Human Resources does not keep track of equipment in their database, so by themselves these databases are not useful for tracking equipment based on employee status.

However, if you use an enterprise-reporting system, you can retrieve data from both of these data sources and combine them into a single report that helps management see which employees have what equipment and what needs to be done with that equipment.

Now let's go one step further. Human Resources may have a system that uses Microsoft SQL Server 2005 for data storage and Maintenance may use Oracle 9*i* for storing data. No problem. An enterprise-reporting system can combine data not only from different data sources but also from different vendors. That is what a reporting infrastructure is used for, but do you need one?

Do You Need a Reporting Infrastructure?

The short answer to this question is yes, you do. The real question is: What kind of reporting infrastructure do you need to put in place? To answer that question you need to consider a few basic things:

- How much data do you need to report on?
- Where is that data stored?
- How many reports do you need?
- How often do you need them?
- How complex are those reports?

Your answers to these questions will dictate whether or not you need an enterprise-reporting system. For example, if you have a small spreadsheet of weekly sales figures and you just need a total of the sales every week, you don't really need an enterprise-reporting infrastructure. In fact, it would be a waste of money in this case.

If, however, you have several database servers with multiple databases that you need to combine into multiple reports on a regular basis, then it is prudent for you to consider an enterprise-reporting infrastructure.

The most common business scenario is somewhere in between these two extremes. Many companies have a single database server that houses data for all of their departments in separate databases. They may have other small, external data sources, like spreadsheets, to report on as well. If your company falls into this category, you should consider an enterprise-reporting structure.

Take our employee/equipment report, for example. This report was based on sizable databases in two different departments that are stored on two different database systems. It would be difficult and costly to use a simple reporting tool to generate any meaningful reports from these two data sources. Sure, it could be done using something like Access or FileMaker, but these tools were not designed with this capability in mind. These applications were designed to report on small, local databases. So getting them to talk to large, remote databases is a time-consuming, costly task. And the reports are usually slow when they are complete.

In contrast, an enterprise-reporting infrastructure is designed specifically for combining data from multiple data sources, provided by different vendors, into reports that can be used by anyone. Because the enterprise-reporting system is designed specifically for this task, the reports are generated quickly, especially in comparison with the nonenterprise counterparts.

So if you have a lot of data on disparate systems, you definitely need an enterprise-reporting system to generate custom reports quickly. Now you just need to know why you should use Reporting Services to fill this need.

Real World Scenario

Case Study: Talksalot Communications

Talksalot (not their real name) were doing a lot of business with some of the big-name vendors in the communications market. They sold everything from home phone service to large business communication services like T1 lines.

Every month Talksalot received commissions from their vendors, along with accounting reports. The reports came in two formats — paper and electronic (on magnetic tape). The personnel in the company did not know how to use mag tape, so they tried to reconcile their commission checks using the paper copies of the accounting reports. This took so much time that they often did not complete the task before the next round of reports came in.

A consultant came in and wrote a small application to control a mag tape reader and read the tapes automatically. All the users had to do was load the tape. The application read the tape and transferred the data into a database for storage. The consultant then wrote a number of reports that showed revenue and credits for the different avenues of business. Talksalot was able to review the accounting reports from their vendors in less than a day.

What to Look For in a Reporting Infrastructure

There are several features that you and your users will likely expect from any enterprise-reporting infrastructure, so look for these features when making a decision:

- Quality reports — No matter how complex a report may be, users want them to be formatted in a professional manner that is easy to read.
- Personalization — Not all users have the same needs, so you should be able to deliver reports to the user in a format they can use with content they need.
- Localization — If you are in an international organization, you need to be able to display reports to users in their native language and cultural context.
- Interactivity — Your users should be able to navigate existing content easily and create new content on the fly.
- Private branding — Reports should blend seamlessly into existing applications and websites. End users should not know that they are looking at a report generated by a separate system.
- Extensible content tools — You need to be able to extend the infrastructure to include new data sources and reporting capabilities that may not be included out of the box.
- Browser-based report delivery — Everybody has a web browser on their computer these days, so you need to be able to deliver reports to the browser.
- High performance — The infrastructure has to be as fast as possible because users do not want to wait for their reports. They want their data now.
- Manageability — You need to be able to manage all aspects of the infrastructure easily, so make sure that it comes with the proper management tools and APIs.
- Scalability — When your company grows, you need to be certain that your reporting infrastructure can grow with it.

Why Use Reporting Services

We've established that you need a reporting infrastructure of some kind. More complex reporting needs require a more robust reporting infrastructure. Why should you use Reporting Services as a reporting infrastructure, though? What can it do for you?

Although these may seem like trivial questions to the technically inclined, they are very important questions from a business perspective. This is especially true when you consider the expense involved in installing and configuring a reporting infrastructure.

Consider this from your boss's perspective. There is a lot involved in getting a reporting infrastructure up and running in a production environment:

- You have to purchase the software.
- Support personnel have to be trained.
- Developers will need training.
- The software will need to be installed across the enterprise.
- The software will then need to be configured and tuned.
- Developers will need to write new reports, plug-ins, supporting database objects, etc.

This can be an expensive undertaking that would be difficult, if not impossible, to back out of once you get started. So it is important to understand what Reporting Services can do for your organization and how it does so.

What Reporting Services Can Do for You

While it is true that Reporting Services is an enterprise-reporting infrastructure that can combine data from a variety of data sources into a report that you can make some sense out of, there is more to it than that.

Reporting Services supports the entire reporting lifecycle, beginning to end. Here is what the lifecycle entails:

Report Authoring One of the most important tasks that a reporting infrastructure must support is creating new reports. In Reporting Services, developers can use any tool that supports the Reporting Definition Language (RDL), which is an XML-based language used to define reports. Of course, Reporting Services has some basic tools for creating new reports built right in.

Report Management Once developed and released, reports can be accessed on a scheduled basis or on demand. The reports are usually cached on disk and an execution history is kept. Reporting Services has built-in tools for keeping track of the execution history, managing the cache, scheduling report delivery, and performing other management tasks. You can also access management functions via web services if you want to incorporate these functions into your own applications.

Report Delivery Users can receive reports in one of two ways. The report can be pushed to the user based on an event (a push report), or the user can request a report specifically (a pull report). The users can either get the report from a web page, which they can then print or save as a Microsoft Excel spreadsheet or a PDF, or they can have the report e-mailed to them.

Report Security You do not want everyone in the company to have access to every report available, so you can secure your reports. Reporting Services implements a role-based security model that you can extend to use custom security models if necessary.

These capabilities allow you to design reports for a number of different reporting scenarios:

Enterprise Reporting This is the primary function of Reporting Services. Developers can design reports at any location, on any supported data source, and deploy them to a central

location for users anywhere in the organization to access. For example, a developer in California can develop a report and deploy it to a server in Denver for a user in Bangkok to access.

Ad Hoc Reporting SQL Server 2005 Reporting Services comes with a fancy new tool called Report Builder, which allows users to create their own reports. This can be handy for users who need a report for that 8:00 meeting when the developers do not get in until 9:00.

Web-Based Reporting Everyone in business today has to work with other businesses to some extent, which means that they need to share data. Because Reporting Services delivers reports via the Web, it is an excellent tool for sharing data with vendors over the Internet or extranet.

All of this makes Reporting Services an excellent choice for your reporting infrastructure from a business perspective. But you are probably wondering how this works. So let's get a quick overview of how Reporting Services does its job.

Inside Reporting Services

SQL Server Reporting Services is a collection of several components. The main component is the Report Server. We call it the main component because this is the workhorse that processes incoming report requests and makes reports available to end users. Further, the Report Server is actually two different components combined into one: a Windows service and a web service. Let's see what these two main components do.

Report Server Components

The Windows service provides the necessary framework for getting Reporting Services up and running, and keeping it that way. The Windows service initializes the instance of Reporting Services, making sure the necessary databases exist and are ready for use, ensuring that all of the service components are available, etc. This service provides scheduling and delivery services for getting reports to end users. It also provides maintenance functions to keep the services running at peak performance.

The web service exposes a set of programmatic interfaces that developers can use to access the Report Server from their custom applications. Quite a few commands are available to you as a developer that you can access through your application, which makes the web service a very powerful tool.

These two components combine to form the Report Server, but there is still more. The Report Server does its work by using a number of subcomponents.

Report Server Subcomponents

The Report Server is made up of several subcomponents, as seen in Figure 1.2. Some of the subcomponents are divided into smaller components still. Let's look at an overview of what these subcomponents are and what they do.

Processors The Report Server actually uses two processors to perform the tasks of preparing and delivering reports. These cannot be modified or extended by developers.

Report Processor When a user makes a request for a report, this processor retrieves the report definition from the Report Server database, initializes any parameters and variables in the report, and performs some preliminary processing to get the report ready for delivery.

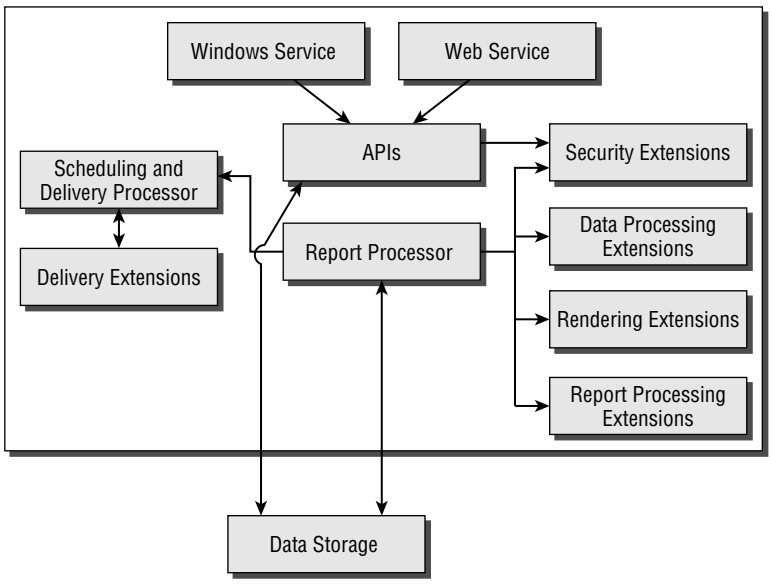

FIGURE 1.2 Reporting Services comprises several components and subcomponents.

Scheduling and Delivery Processor Reports can be scheduled for delivery to e-mail inboxes and shared folder destinations. This processor handles that scheduling functionality. To do this, the component maintains a queue of scheduled events and notifications. When an event fires, this component calls the report processor to execute a report, process a subscription, or perform some scheduled maintenance. After the report is generated, the scheduling processor calls the appropriate delivery extension to get the report to the user.

Data Storage Everything that users see in Reporting Services is stored in a database. The database contains metadata, report definitions, and the folder hierarchy displayed in the Report Manager.

Extensions Extensions are used to perform a variety of functions, such as rendering reports and processing data. These differ from processors in that developers can create new extensions to meet their own needs.

Security Extensions Security extensions are used to authenticate and authorize users who try to access reports on the server. The default method is Windows Authentication. If that does not meet your needs, you can create a custom security extension. A custom extension cannot run alongside the default extension because there can only be one security extension in place at a time.

Data Processing Extensions When a report is processed, this extension type is used to query the underlying data source and return a flattened rowset. The basic tasks needed to accomplish this are as follows:

1. Open a connection to the data source.
2. Analyze the query and return a list of field names.

3. Run the query against the data source and return a rowset.
4. Pass any required parameters to the query.
5. Iterate through the rowset and retrieve data.

There are several types of data processing extensions available right out of the box, including SQL Server, Oracle, OLE DB, ODBC, and Analysis Services. If these do not meet your needs, you can write your own. Unlike with security extensions, you can have multiple data processing extensions running side by side.

Rendering Extensions These extensions transform data from the Report Processor into a format that users can read and understand. There are six of these extensions by default, and you can add more as you see fit. Let's look at the default rendering extensions.

HTML Whenever a report is requested through a web browser, this extension is used to render the report. All HTML from this extension is generated using UTF-8 encoding.

Excel This extension is used to generate reports that can be viewed and changed using Microsoft Excel 97 or later. This extension generates files in the native file format for Excel, called the Binary Interchange File Format (BIFF).

CSV The CSV rendering extension creates plain-text files that have values separated by commas. There is no formatting applied. CSV files can be opened using most spreadsheet applications as well as standard text processor applications.

XML Like the name implies, this extension renders reports into an Extensible Markup Language (XML) format, which can then be used in other applications. This extension can also use an Extensible Stylesheet Language Transformations (XSLT) schema to transform the resultant XML into any schema you may need. All of the XML is generated using UTF-8 encoding.

Image This extension can render reports in a variety of image formats, including TIFF, BMP, EMF, GIF, JPEG, and PNG. This is very useful in a few scenarios. Using this extension ensures that the resultant report looks the same on every client, no matter what application they use to view it. This extension can also be used to send a report directly to a printer or fax machine for transmission to users.

PDF This extension renders reports in a PDF format that can be read by Adobe Acrobat 6 and later.

Report Processing Extensions Reporting Services features a variety of items that can be included in your reports for processing and displaying data. Out of the box you can include tables, charts, matrices, text boxes, images, and more. However, if you want to include something more exotic, such as a custom map from a mapping application, you would need to create a custom report processing extension. You can have many of these installed on a single server at the same time.

Delivery Extensions These extensions are used to deliver reports to their destination so users can access them. There are two of these by default.

E-mail The e-mail delivery extension is used to send reports to users via Simple Mail Transport Protocol (SMTP). It can send the entire report in the body of the e-mail, or just a URL with a link to the report on a web server. Short notices without the link or

report can even be sent to Short Message Service (SMS) devices like pagers and cell phones.

File Share As you might expect, this extension delivers reports to a file share on a network drive that users can access. You can specify the location, rendering format, filename, and overwrite options to control this extension.

Now that you know what Reporting Services can do for you and how it does what you need, all you have to know is how much it costs. The actual price changes from vendor to vendor, so we can't discuss specific prices here. We can, however, help you understand the licensing requirements.

Licensing Reporting Services

Before you can make a valid business decision to start using Reporting Services in your organization, you need to know how much the software is going to cost. You should contact your vendor about specific prices, but if you understand the licensing model you will be better equipped to make a decision.

The good news is, SQL Server 2005 comes with the license to run Reporting Services. You do not need to buy a separate license to run the software or have users access it. There is a caveat, though. You need to have a valid SQL Server 2005 license for the system that reporting services is running on. There are three scenarios to consider when looking at licensing.

Scenario 1: Single-Server Environment If you run SQL Server 2005 and Reporting Services on the same computer, you do not need to purchase any extra licenses, as illustrated in Figure 1.3. You are fully licensed at that point.

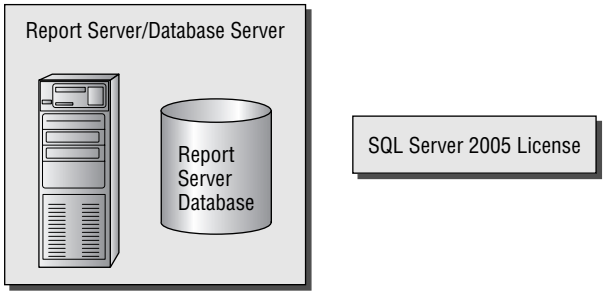

FIGURE 1.3 Only one license is required for a single-server environment.

Scenario 2: Multiple-Server Environment If you need to run SQL Server 2005 and Reporting Services on separate computers, as shown in Figure 1.4, you will need to purchase a SQL Server 2005 license for both systems.

Scenario 3: Web Farm A web farm, from a Reporting Services perspective, is a number of computers that have Reporting Services installed and that all connect to a separate server that has a single Reporting Services database shared by all instances, as illustrated in Figure 1.5. In this case you will need to purchase a SQL Server 2005 license for all of the front-end systems (the ones running Reporting Services) and a license for the database server.

Now all you need to do is decide which environment you need to run and purchase the correct number of SQL Server licenses and you will be in business.

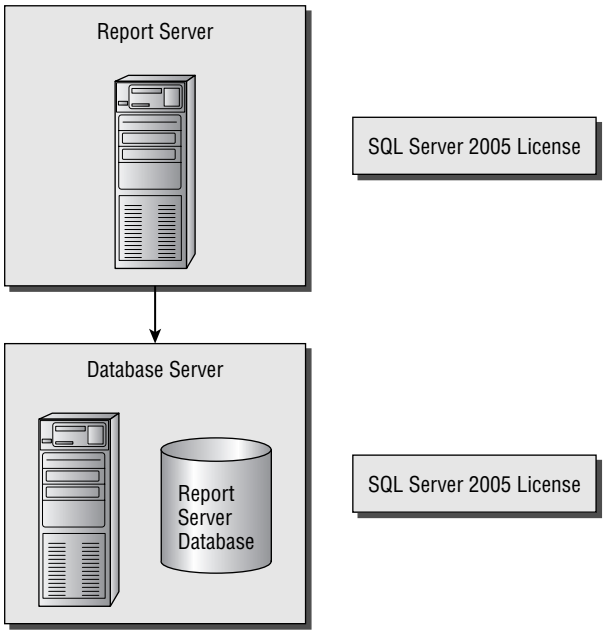

FIGURE 1.4
You will need a license for all servers in a multiple-server environment.

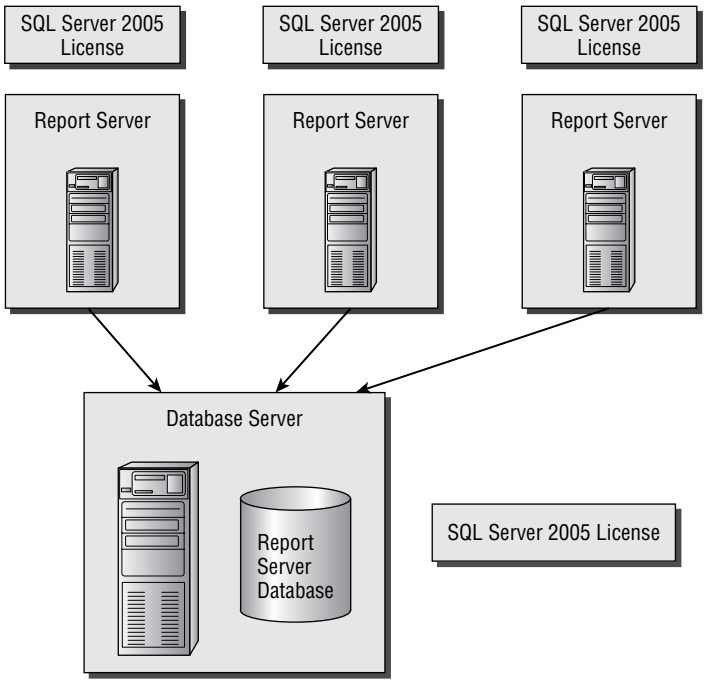

FIGURE 1.5
You will need a license for each system running Reporting Services and the database server in a web farm.

The Bottom Line

Define a reporting infrastructure. Simply put, a reporting infrastructure combines one or more data sources with a mechanism for displaying the data in a format that makes sense to end users.

> **Master It** What are the two primary components of a reporting infrastructure?

Figure out why an enterprise-reporting infrastructure is necessary. Just because you have data does not mean that you need an enterprise-reporting infrastructure. You may only have a small amount of data that you report on infrequently, and the reports may be fairly simple. This scenario may not require an enterprise-reporting infrastructure. If you have a lot of data from a number of data sources that you require complex reports from, then you should seriously consider an enterprise-reporting infrastructure.

> **Master It** You work for a mid-sized company that has a Human Resources database on an Oracle 9*i* system and an Accounting database on SQL Server 2005. You need to generate salary and bonus reports for upper management based on both of these databases. Does this require an enterprise-reporting infrastructure?

Decide if you need to use Reporting Services as your reporting infrastructure. Reporting Services can easily handle any reporting task you can throw at it, so functionality is not the question here. The real question is how much Reporting Services will cost to implement. Talk with your vendor for actual pricing, but you will find that, with the right licensing, Reporting Services is a definite contender.

> **Master It** You work for a mid-sized company that has an inventory management database on a SQL Server 2005 server. You need to create a report for your customers that shows your current inventory. Your customers need to access this report over the Internet. Can you use Reporting Services to meet your needs?