

Chapter 4. Describing What Businesses Do and How They Do It

Robert J. Glushko and Tim McGrath

4.0 Introduction.....	2
4.1 Views of Business Organization.....	4
4.1.1 Physical Views of Business Organization	4
4.1.1.1 Organization Charts and Facilities Maps.....	4
4.1.1.2 Supply Chains.....	4
4.1.1.3 Distribution Channels	6
4.1.1.4 Marketplaces, Exchanges, and Auctions	7
4.1.2 Conceptual Views of Business Organization.....	8
SIDEBAR: Acronymology in Patterns of Business Organization.....	8
4.1.2.1 Conceptual Views of Supply Chains	10
4.1.2.2 Conceptual Views of Marketplaces, Exchanges, and Auctions.....	10
4.2 Views of Business Processes	11
4.2.1 Physical Views of Business Processes.....	12
4.2.2 Conceptual Views of Business Processes	12
4.2.2.1 Business Reference Models	13
4.2.2.2 Supply Chain Reference Model	13
4.2.2.3 RosettaNet.....	14
SIDEBAR: The Secret of RosettaNet's Success	15
4.2.2.4 Information Supply Chains	15
4.2.2.5 Demand Chains.....	16
4.2.2.6 Document Automation and Straight Through Processing	17
4.3 Views of Business Information.....	20
4.3.1 Physical Views of Business Information	20
4.3.1.1 Electronic Data Interchange.....	20
4.3.1.2 XMLification	22
4.3.1.3 Information Aggregations.....	23
4.3.2 Conceptual Views of Business Information	24
SIDEBAR: The Universal Business Language.....	25
4.4 Views of Business Architecture.....	26
4.4.1 Physical Views of Business Architecture	27
4.4.1.1 Technology Platforms and Infrastructure	27
4.4.1.2 Integration Architectures and Patterns.....	27
4.4.1.3 Web Services	29
4.4.2 Conceptual Views of Business Architecture.....	30
4.4.2.1 Conceptual Views of Integration Architecture	31
4.4.2.2 Service-Oriented Architectures.....	33
4.5 Key Points in Chapter 4	36
4.6 Notes	37

4.0 Introduction

Readers who came to Chapter 2, “XML Foundations,” knowing something about XML gained new insights and ways to apply that knowledge through our Document Engineering perspective. Likewise, readers with business backgrounds will be familiar with some of the material in this chapter but should benefit from seeing it from a Document Engineering perspective.

In Chapter 3 we deliberately used a general notion of *business pattern* because we wanted to emphasize the great extent to which businesses carry out their activities in regular and systematic ways. Now that we’ve made that basic point we will get more precise.

Historians, sociologists, business theorists and institutional economists have developed a rich set of categories for discussing variations in business organization and models.¹ This deep body of work has shaped our thinking, but we won't explicitly revisit much of it in this chapter. Instead, we will take a less formal and more pragmatic approach, adapting some of the categories and concepts as we discuss business models that use document exchanges and service-oriented architectures. So while some of the topics we'll discuss in this chapter will be familiar to anyone who has studied organizational design, supply chain management, or information technology management, the overall framework provided by Document Engineering is a new one.

We introduced this new perspective in Chapter 3 when we proposed the model matrix as a framework for understanding the relationship between organizational, process, and information patterns, which vary on a dimension of granularity. In this chapter we will take a more detailed look at each of these model layers to understand the orthogonal distinction between conceptual models and physical ones. We will then be ready to learn how to develop compatible and interconnected models from all three layers that describe both what the business wants to do and how it can do it.

We begin with models of how businesses organize their activities. *Business models* or *business reference models* are abstract descriptions of what businesses do. We will describe patterns like supply chains and marketplaces that capture complex sets of relationships within and between enterprises.

At the more granular view of business processes, *business process models* take a view that emphasizes the activities that create business value without focusing on the information exchanges that underlie them.

Only at the most granular level of *business information models* do we find patterns that reveal documents and their logical components. These are most commonly found when they are at the physical or implementation level in the form of XML schema libraries or EDI message standards.

We will introduce these different model layers using the distinction between physical and conceptual views we discussed in Chapter 3. Although the contrast isn't always perfect, one can describe most aspects of what businesses do in either way; for example, in highly physical terms of management reporting structures or facility locations or in highly conceptual terms such as whether it seeks efficiency through functional or cross-functional organization. Likewise, the information exchanged between organizations or systems can be described in highly physical terms by XML schemas or EDI implementation guidelines (that is, as *document implementation models*), or in highly conceptual terms by UML class diagrams² (as *document component* and *document assembly models*).

Even business processes, which may seem inherently abstract for processes that are information-intensive or computational, can be described from both physical and conceptual perspectives. It is certainly true that in contrast to observable processes like manufacturing, packaging, and transport of tangible goods, many business processes like accounting, scheduling, and payment are almost invisible. But even intangible or information-intensive processes need instructions about how they are carried out, and the documents that are the inputs and outputs of these processes also provide physical views of how the process works.

After a business has designed its organizational, process, and information models, many technology and architectural choices remain about how to implement them. And just like those models, the technology and architecture of a business can be described in physical or conceptual terms. Physical descriptions depict the specific computers, operating systems, and software applications that the business uses. In contrast,

conceptual and technology-neutral descriptions emphasize functional and topological characteristics, such as whether the solution embodies a service-oriented architecture and treats business functions as reusable components.

4.1 Views of Business Organization

One approach to describing a business is in terms of the organization, management, or control of its activities. These descriptions can explain the organization of a single firm or the organizational relationships between multiple firms.

4.1.1 Physical Views of Business Organization

4.1.1.1 Organization Charts and Facilities Maps

A common physical view of business organization is the *organization chart*.

Organization charts exhibit characteristic structural patterns that portray the arrangement of management and operational responsibilities within the firm and usually include specific people and their associated roles or titles. These patterns are explained in textbooks on organizational design or behavior and in a more mundane way are built into “enterprise definition” tables in HR and ERP applications and the “orgchart” templates in drawing programs like Visio, SmartDraw, or Powerpoint.

Organization charts and facilities maps are physical models of a business

The organizational chart for a business often closely mirrors the *facilities map*, another common physical model of business organization that shows the locations of offices, factories, distribution points, training centers, or other facilities.

The organizational chart for an enterprise is a highly specific and rich model of how it does business. For example, IBM uses its organizational charts as the core of a dynamic information resource called BluePeople.³ Starting with a name or email address, BluePeople makes explicit the network of links to coworkers, projects, publications, and other information to provide context for the name or address.

4.1.1.2 Supply Chains

A firm’s *supply chain* is the network of relationships, communication patterns, and distribution capabilities that provide raw materials, components, products, or services to a

firm so that it can make what it sells and deliver what it sells to its customers. Because the pattern of a supply chain is a highly abstract one that can be adapted to model any situation in which a product or service is created by bringing together different parts, it is an important part of the Document Engineering pattern repertoire.

Nevertheless, supply chains are often described in highly concrete or physical terms with details about assembly lines, warehouses, factories, and stores full of raw materials, partly finished or finished products, along with the equipment or modes of transport by which materials and products move between them. Likewise, because the perspective of a supply chain follows a product from raw material to consumption, another common and very concrete analogy is to the basin or drainage area for a large river: “A supply chain is much like a river system with raw materials at the headwaters and customers at the delta, with products floating down the river toward the customers.”⁴

A typical physical depiction of a supply chain model is shown in Figure 4-1. This graphically represents a global supply chain whose links are described using standard SCOR “process categories,” which we’ll discuss in Section 4.2.2.2.

Supply Chain Threads are Developed from the Geographic Product Flow

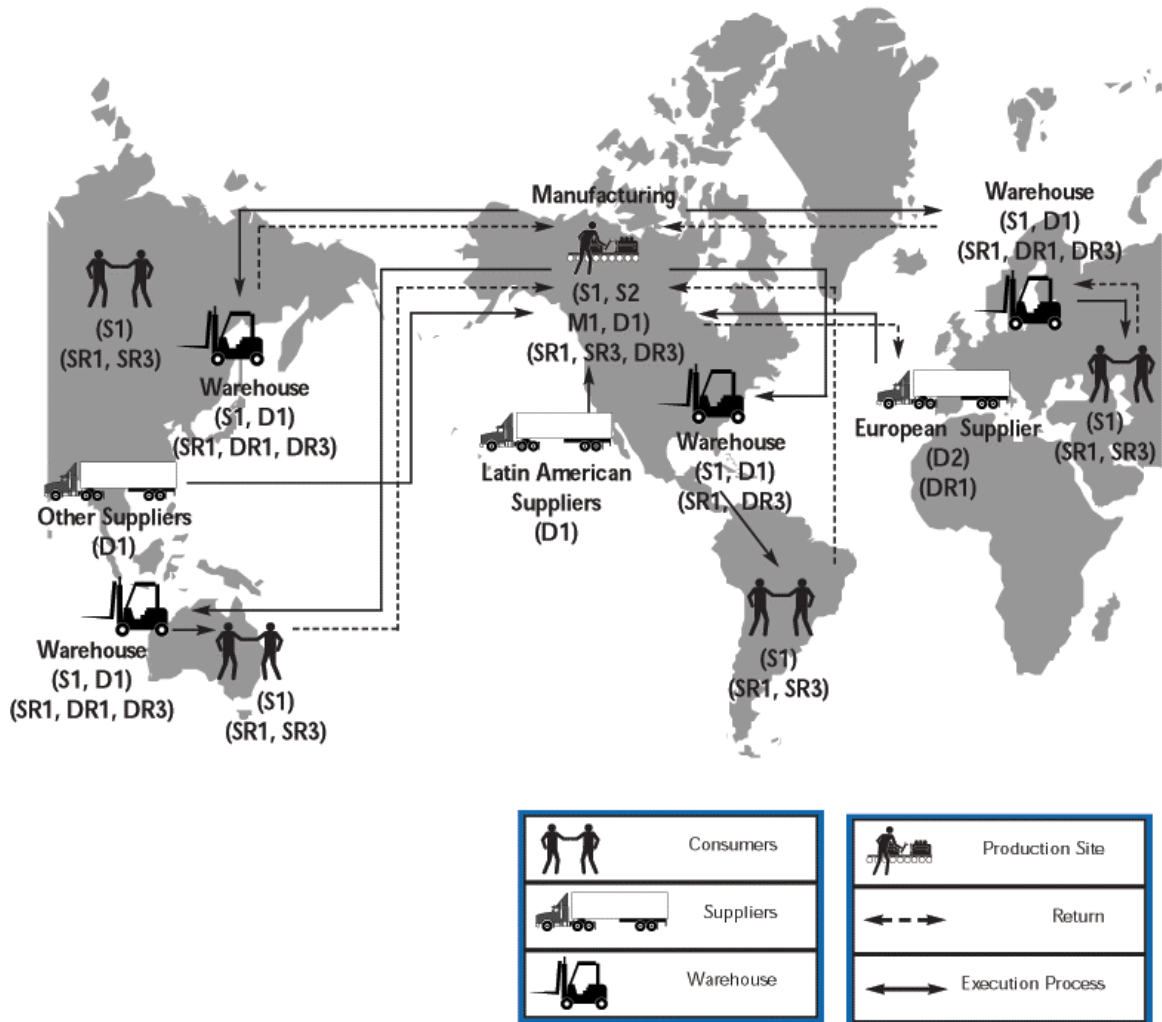


Figure 4-1. A Global Supply Chain Model

4.1.1.3 Distribution Channels

Getting finished goods to the purchaser is called *distribution* or *fulfillment*. And as with supply chains, distribution channels are often described in highly physical terms that detail the locations of warehouses or retail stores and the specific modes of transport between them.

The simplest distribution pattern is *direct* distribution, in which a company sells a product directly to the companies or consumers who buy it. However, most companies use an *indirect* strategy, selling their products through distributors, resellers, and retail outlets to increase their ability to reach customers. These distribution partners are called *intermediaries* or *channels* for the manufacturer; they may be organized according to sales territories, geographical regions, or customer segments. The Internet enabled many firms to shift from indirect to direct distribution, and this *disintermediation*—literally, cutting out the middleman—allowed them to increase their margins and learn more about their customers. A company can be tempted to sell the same products directly and through channels, but this can lead to *channel conflict* and alienated distributors.

Like supply chains, distribution channels are an abstract and general pattern. Applying the pattern involves choosing the number, roles, and locations of intermediaries and balancing the benefits of a larger network against the costs and delays of exchanging information within it.

4.1.1.4 Marketplaces, Exchanges, and Auctions

There are few business patterns that suggest more concrete and stereotyped depictions than marketplaces, exchanges, and auctions. We can all imagine and hear the crowded old town marketplace, the controlled frenzy of the stock exchange trading floor, and the insistent staccato of the auctioneer urging the bidders on.

These patterns have much in common, organizing their participants in characteristic ways to enable familiar business models. All embody the core ideas that bringing together a critical mass of buyers and sellers makes it easier to match them up and creates shared efficiencies and benefits that won't arise in interactions between a single buyer and a single seller. By eliminating the need for participants to be in the same physical location, the Internet allows more of them to take part, yielding much better matching between buyers and sellers. Consider that at any given time millions of items are offered on eBay in a set of categories nearly as broad as the web itself.

The differences between marketplaces, exchanges, and auctions are subtle. While almost any type of products might be offered for sale in a marketplace, an exchange is a type of marketplace for intangible goods like financial securities where price is the

essential attribute, and an auction is a method for establishing prices when market mechanisms don't work well, usually when goods are scarce for one reason or another.

Supply chains, distribution channels, markets and auctions are general business patterns that can be applied in novel contexts

Like supply chains and distribution channels, markets and auctions are very general patterns that can be applied in novel contexts. For example, an Internet marketplace called getloaded.com matches freight loads and trucks with excess capacity, attacking the costly problem of “deadheading” - when a truck returns without a backload on its return trip from delivering goods.

4.1.2 Conceptual Views of Business Organization

Physical views of business organizations are useful depictions of how they operate. In contrast, a conceptual perspective on how a business is organized explains why it exists and the kinds of activities it engages in to stay in business. In its most abstract, conceptual form, the “why” of a business is simply called its *business model*.

A business model is concerned with the nature and pattern of exchanges of one form of value for another

At the heart of every enterprise are trades or deals of some kind, exchanges of one form of value for another. A business model is concerned with the nature and pattern of these deals between businesses and their partners that ultimately yield the products or services it offers to its customers. A company's business model also addresses the roles played by other firms that work with and around it, such as suppliers; customers; stakeholders; intermediaries such as brokers, distributors, and agencies; and service providers of one sort or another. Viewed from the perspective of the enterprise at their intersection or common focus, this collection of parties and their organization is called the business *ecosystem*.⁵

SIDEBAR: Acronymology in Patterns of Business Organization

A very coarse level of describing patterns of business organization in a conceptual way emerged in the mid 1990s as a set of three-character acronyms beginning with B2B and B2C and still growing.

B2B, for “business to business,” was the first of these patterns and it is mentioned in millions of websites and domain names. It was used to describe business relationships in pre-Internet days, often in discussions of EDI document exchanges. For example, an industrial chemicals firm whose products are offered only to other businesses would be following the B2B pattern.

B2C, for “business to consumer,” emerged as a category label for Internet retail sites to contrast them with B2B ones. The number of B2C sites exploded with the popularity of the Web, and it is certainly a more visible category than B2B. Nevertheless, even if the breaking of the Internet bubble hadn’t caused a great many B2C sites to disappear, B2C as a sector would still be dwarfed in economic scale by B2B, since all B2C transactions depend on numerous B2B ones (recall our discussion in Chapter 1 of the B2B “drop shipment” pattern that underlies the Internet bookstore).

More recent variants of the B2B and B2C categories distinguish those that involve governments. B2G, for “business to government,” seems slightly more common than G2B, for “government to business,” but both have been showing steady growth as governments at both municipal and national levels introduce Web initiatives of various kinds. G2C, for “government to citizen,” is the dominant variant. None of these acronyms appears to stand a chance against the term *e-government*, even though a list of the “24 priority e-gov initiatives” in the United States sorts them into citizen and business categories.⁶

Many colleges and universities offer courses on the Internet directly to consumers (“e-learning”) but haven’t adopted the B2C category, perhaps because they aren’t for-profit businesses. Nor have they invented another acronym, although E2C, or “education to consumer” might fit. However, the for-profit distance- or lifelong-learning firms seem eager to embrace both the B2C and B2B labels.⁷

C2C, for “consumer to consumer,” had a brief appearance on the acronym stage to describe the organization of business relationships facilitated by auction sites like eBay, but this term didn’t seem to reach critical mass. In any case, Internet-facilitated business relationships between individuals are now almost universally described as P2P, for “peer-

to-peer.” This acronym is likely to have a long life because of its notoriety in file-sharing applications.

4.1.2.1 Conceptual Views of Supply Chains

Supply chains, especially those for heavy manufacturing industries like aerospace and automotive, are highly visible and physical. But when we want to design and analyze supply chains, it is less important to think in terms of buildings, vehicles, and pallets of goods and instead think from a more conceptual perspective.

Document Engineering treats supply chains as information flows

A conceptual view of supply chain design must deal with complex dependencies between the allocation of materials, production, and distribution responsibilities; the number and location of suppliers and distributors; the amount and location of material and product inventories; and the logistics of getting everything to its desired location at the right time.⁸ Most of this multidimensional design problem must be solved before applying Document Engineering to supply chain design. Document Engineering thinks of supply chains in terms of the information flows that accompany the movement of materials and goods, creating an abstract view of the physical events that trigger information exchanges and the reciprocal physical events resulting from those exchanges.

4.1.2.2 Conceptual Views of Marketplaces, Exchanges, and Auctions

A conceptual view of marketplaces, exchanges, and auctions defines them in terms of their participants and the services that they provide to each other. There needs to be a least one special participant who performs the role of the *operator*. The operator, sometimes called the *host* or *market maker*, must have the credibility or market power to attract the buyers and sellers and establish the governing rules. These rules define the terms and conditions for participation, the specifications for the information that participants will exchange, and the processes or services in which the exchanges will take place. The operator must provide a trusted environment, both in terms of technology considerations like security and reliability and in the business sense of trust about privacy and the honoring of transactional obligations.

A minimal marketplace or auction offers the “commodity” services related to buying and selling, but what attracts and keeps participants are other *value-added services* that create richer relationships between buyers and sellers and induce buyers to return. The services that are most useful depend on the industry, geography, and other characteristics of the context in which the marketplace or auction pattern is being adopted.⁹

By eliminating any need for physical presence the Internet has increased the feasibility and conceptual variety of business models

Auctions have been around since ancient times, but by eliminating any need for physical presence the Internet has increased the feasibility and conceptual variety of auctions. The many different types or patterns of auctions are distinguished by the extent of information exchange among the participants, and by the rules that govern the timing of offers, the selection of the winning offer, and the price the buyer pays.¹⁰

4.2 Views of Business Processes

We’ve talked about “business process” for three chapters without a precise definition because it is such a common phrase. It is obvious that functional business areas like engineering, manufacturing, and sales carry out systematic activities that are somehow interconnected, and we need a notion of business process to describe how this works. So we’ll define *business process* as a chain of related activities or events that take specified inputs, add value to the inputs, and yield a specific service or product that can be the input to another business process. The chain of business processes is maintained by the flow of information between them as the output of one process becomes the input to the next.

Business process models are the bridge between organizational models and business documents

Business process models are central to Document Engineering because they are the bridge between higher-level strategic expressions of what businesses do represented in organizational models and the lower-level operational concerns reflected in document and information models.

4.2.1 Physical Views of Business Processes

Physical views of business processes describe the way in which specific business activities are implemented by a firm. Most firms have a vast variety of policies and procedures governing how they hire, pay, train, evaluate, and terminate employees; how they approve, budget, staff, review, and learn from projects; how they conceive, design, manufacture, document, test, market, and sell products; how they procure needed goods and services and operate and maintain equipment; how they deal with business partners and customers; how they account for income and expenses and meet government reporting requirements—the list goes on and on.¹¹ All of these are physical views of business process models.

Some of these policies and procedures exist as documents on employee desks and office shelves or on the company intranet. Others are embodied as business rules in software applications that range from electronic mail and spreadsheets to enterprise content management and ERP systems.

As we discussed in section 3.4.2, “Why Businesses Follow Patterns,” many business processes are dictated by laws, regulations, and standards. These may sometimes function as conceptual models that govern or guide many aspects of individual and corporate behavior and business processes. But sometimes they are highly prescriptive, specifying how things can and cannot be done, possibly even dictating the technology and manner of solution implementation. Prescriptive models of this sort are implemented using mundane document templates or software applications that create customized employee handbooks, procedure guides, and contracts.

4.2.2 Conceptual Views of Business Processes

A company’s business model or value chain shows the logical relationship between the functional areas in the enterprise. However, the granularity of functional areas often provides too coarse a perspective for analyzing what an enterprise does, what it needs to do better, and what it can do without. It is helpful to further decompose functional business areas into subareas and more specific business processes.

4.2.2.1 Business Reference Models

Because they are more stable descriptions of what an enterprise does, the highest level functional areas are categories for organizing models at lower levels, and the resulting hierarchy of business processes is called a *business reference model*. An important business reference model is the recently developed Federal Enterprise Architecture of the U.S. government,¹² which could be a pattern for other governments around the world.

A reference model consolidates the best practices of many companies

Business reference models exist in many industries and are most often created by industry associations or by consulting firms that have extensive industry experience. Almost by definition a single firm can't create a business reference model because a good reference model consolidates and abstracts from benchmarking or best practices analyses of many companies in the industry.

4.2.2.2 Supply Chain Reference Model

Many of the patterns in supply chain models can be seen in the Supply Chain Operations Reference Model (SCOR), a reference model developed by an industry group called the Supply Chain Council.¹³ SCOR provides standard patterns for describing supply chains in terms of five basic processes: plan, source, make, deliver, and return.

These patterns are organized as conceptual models whose two lower levels of detail refine the basic five processes to describe supply chain models for different industries and partner relationships. Figure 4-2 shows the top level view of the SCOR supply chain pattern.

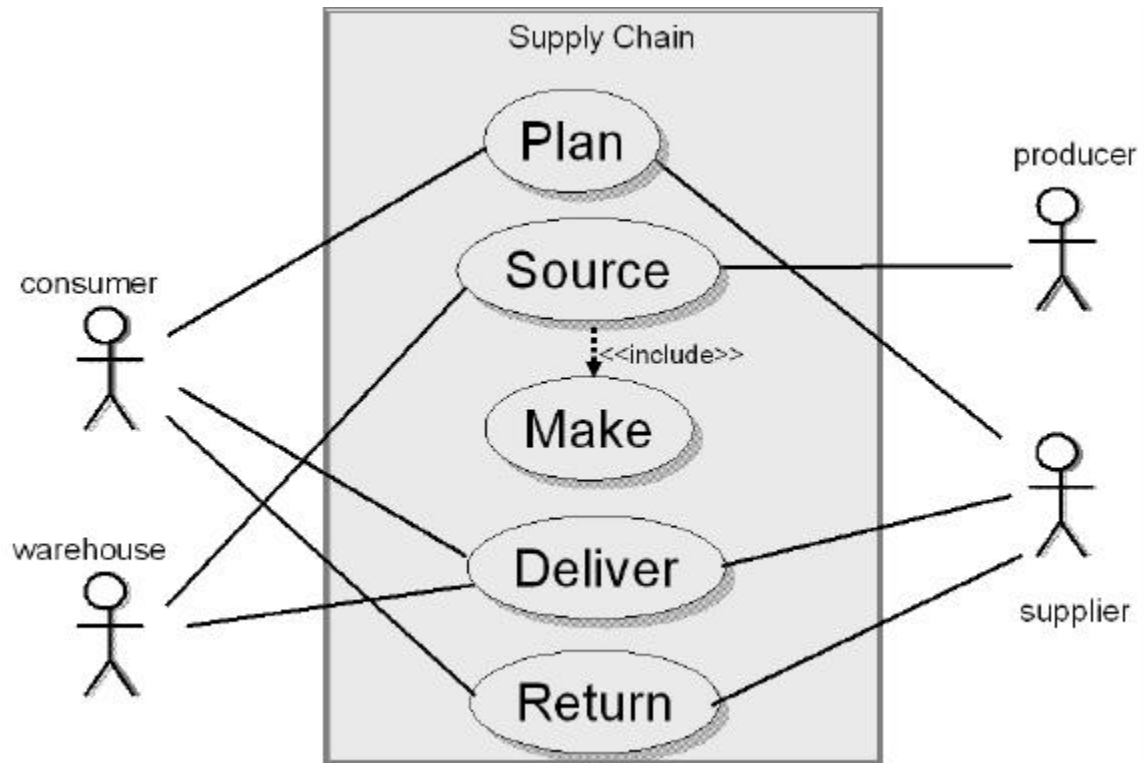


Figure 4-2. The SCOR Supply Chain Pattern

4.2.2.3 RosettaNet

The RosettaNet Consortium has developed standard specifications for processes in the global supply chain for the electronic components and IT industries. Approximately 100 detailed process models called partner interface processes (PIPs)¹⁴ are organized hierarchically by clusters and segments. The first PIPs that most firms implement are those in the Order Management cluster, which contains segments for Quote and Order Entry, Transportation and Distribution, Returns and Finance, and Product Configuration. The fourteen 14 PIPs in Quote and Order Entry, like PIP 3A4 for Request Purchase Order, define both the document models (as XML DTDs) and the collaboration of document exchanges between trading partners.

The implementation focus of PIPs means that they are physical views of business processes that we might have discussed in section 4.2.1. But in other respects the RosettaNet specifications represent a more conceptual view of business processes. In particular, the hierarchical arrangement of PIPs into clusters and segments provides a useful vocabulary for analyzing supply chains at different levels of abstraction. In

addition, the PIPs were developed using a common meta-model shared by all the PIPs facilitates its generalization to other industries. We will demonstrate the reuse of RosettaNet PIPs as business process patterns in Chapter 10.

SIDEBAR: The Secret of RosettaNet's Success

RosettaNet, founded in 1998, is a consortium of major information technology, electronic components, semiconductor manufacturing, telecommunications, and logistics companies that is creating and implementing business process patterns. RosettaNet began with the IT supply chain and has sought to expand its membership and scope to extend the coverage of its patterns beyond its current vertical market. RosettaNet stands apart from many standards efforts in its member commitment to implementing these common patterns. Maybe this follows from RosettaNet's steep annual dues—\$50,000 in 2004—which means that participation is a high-level strategic decision.

But participation in RosettaNet may be worth it. Intel, one of the founding members of the consortium, reported that more than 10 percent of its supplier and customer transactions in 2002 were based on RosettaNet, a total of about \$5 billion. Intel is using RosettaNet standards to work with more than 100 trading partners in more than 20 countries and is counting on RosettaNet to reach the aggressive goal of becoming a 100 percent e-corporation.¹⁵

4.2.2.4 Information Supply Chains

The flow of materials and goods in a supply chain or distribution channel has always been accompanied by the flow of information about it. When we unpack a box of something we've ordered, we often find that it contains a shipping label with our name and address, a packing slip or manifest that itemizes the contents, assembly or operating instructions, a payment receipt, an invoice, and other types of documents. But information about the processes is increasingly becoming separated from the physical flow of materials and goods, at which point it can be thought of in conceptual terms as an *information chain* or *information value chain* or *information supply chain*.¹⁶

Information about the business processes is distinct from the physical flow of materials and goods

An information flow specifies who exchanges information, what information they exchange, and the frequency with which they exchange it. The documents exchanged package the content of these information flows. And, while communication and information technology is what makes the information flows possible, the technology itself is less important than the abstract perspective of the patterns of information exchanges and processes.

The information flow of a supply chain differs in three critical ways from the physical supply chain:

- Information can flow qualitatively faster than materials and goods, which might spend weeks in trucks, trains, or shipping containers moving around the world.
- Information may flow in the opposite direction of the materials and goods, moving from customers and retailers back toward distributors, manufacturers, and their suppliers.
- Information can go many places at once so that supply chain participants can know about inventories, locations, sales, and so on without having to witness them.

These three characteristics of the information chain make it an essential adjunct to the physical supply chain and the key to keeping a business competitive and responsive to rapidly changing markets and customer requirements.

4.2.2.5 Demand Chains

When information flows in the opposite direction of the materials and goods, moving from customers and retailers back toward distributors, manufacturers, and their suppliers, the flow is sometimes called the *demand chain*. This backward (or feedback) flow of information isn't a new thing, but near real-time information about inventories and sales is profoundly more valuable than monthly reports. For example, websites allow a firm to capture implicit or explicit demand information from customers around the clock.

If retailers provide inventory information to suppliers, the suppliers can take responsibility for resupplying inventory and keeping the retailer's shelves stocked. In this *vendor managed inventory* (VMI) pattern, the supplier ships replacement goods directly to the retail store to keep inventories at agreed levels.

VMI is often the first stage of greater information chain collaboration because its benefits reinforce information sharing between retailers and their suppliers.¹⁷ If retailers are willing to share additional information, such as point of sale transaction data and customer information from loyalty programs, the suppliers and retailers can collaborate on business planning, sales forecasting, evaluations of pricing and promotions, and other opportunities for continuous improvement in their joint processes. This more comprehensive information chain pattern is called *collaborative forecasting, planning, and replenishment* (CPFR).¹⁸

VMI and CPFR patterns can be generalized to other information chain situations involving the delivery of services rather than goods for sale. An example is Otis Elevator's Remote Elevator Monitoring system, which monitors numerous elevator functions and initiates orders for service calls or maintenance parts. Remote monitoring of equipment, machinery, or facilities can be thought of as vendor-managed or outsourced asset management.¹⁹

4.2.2.6 Document Automation and Straight Through Processing

So far we have presented views of business organization and processes that mostly involve the movement of tangible things or information about the movement of tangible things. But a great deal of what businesses do involves even more abstract activities that can be described in terms of the movement of information, and sometimes the activities are so abstract that the only tangible things involved are artifacts that record the information. These kinds of business processes follow the related patterns of *document automation* and *straight through processing*.

Every significant business manages its money, files tax returns, and submits financial reports to various government agencies, often for multiple jurisdictions. In industries like healthcare, insurance, banking, real estate, financial services, and securities, the high business value activities centers around document processing for transactions. Many of these industries use some notion of a *financial value chain* as an analogue to the supply chain in industries with more tangible products.

Many of the information-intensive activities in these industries were once carried out using paper documents that moved from one organization or firm to another, with the

documents growing through the incremental addition of evidence, approvals, reconciliations and other information. Today businesses often make it a goal to use the Internet to capture and exchange documents from the moment they are created to the time they have served their purpose to complete, settle, or reconcile a transaction.

The exact definition of this end-to-end goal differs from industry to industry but is most often called *straight through processing* (STP). Such initiatives began in the securities industry²⁰ and document automation initiatives in insurance, real estate, and human resources have all adopted the STP label. And even though it isn't described as STP, an effort underway at the Florida State Senate to automate the end-to-end lifecycle of laws from their origins as draft bills all the way through their publication as printed and Web documents certainly fits the definition.²¹

Business processes vary greatly in how completely they can be automated. Those that require clerical functions of data entry, verification and calculation can often be totally automated. The business rules that need to be enforced can easily be encoded in XML schemas, spreadsheets, or in application logic. At the other extreme, those at the other end of the continuum that require expert analysis, tacit knowledge, and the interpretation of business policy with respect to competitors or customers can only be partially automated.

Indeed, the extent of automation in the latter context can sometimes be little more than more efficiently getting the computerized information fodder of the task to the knowledge worker who actually performs it. Nevertheless, even this limited degree of document automation can significantly improve productivity by more fairly distributing the workload in a group of such workers.

Document automation and STP efforts don't simply replace the physical workflow of paper documents with the logical flow of computerized ones. The electronic documents might all be stored in a centralized and shared document management system, which eliminates the need for documents to move from place to place or from system to system. Instead, all of the processing or approval transactions take place using a shared repository, with logical workflow and access privileges ensuring that the appropriate people interact with the documents at the desired time. The U.S. Army, which handles an

estimated 15 million copies annually of 100,000 different forms, hopes to save \$1.3 billion a year by implementing a centralized forms content management system.²²

Many STP efforts in the US are being driven by the Sarbanes-Oxley Act of 2002,²³ enacted to curb corrupt business activities and fraudulent accounting practices like those of Enron and WorldCom. Sarbanes-Oxley requires firms to implement adequate internal control structures and procedures and attest to their effectiveness. Informal or unautomated procedures don't enable sufficient auditing and tracing of information about where money came from, where it went, and why it went there.

Sarbanes-Oxley has inspired numerous efforts to create standard conceptual models for the information needed to conduct effective audits.²⁴ These models describe the relationships among business organization, processes, accounts, control procedures, types of risk, and so on. The overriding goal is to enable better electronic discovery and management of the documents needed by the audit and by the assurance reports. Sarbanes-Oxley is also driving increased spending on the enabling technologies of document and records management, business process automation, and security.²⁵

While not every document automation or STP effort is the same, they share some key characteristics or subgoals that define the pattern:

- They emphasize more efficient creation of the initial document or documents through the use of templates for different document types or guided assembly of a custom document from components.
- They seek to minimize manual intervention as the documents flow from process to process by transforming information for reuse in different contexts and by using business rules to automate routing, access control, and exception handling.
- They seek not just to automate existing processes, which would be akin to creating roads by paving cow paths, but to refine or reengineer them, possibly by adopting industry best practices or reference models.
- They view documents as dynamic rather than static, automatically propagating changed information into the processing pipeline so that it is current and available when needed.

- They take an end-to-end perspective that maximizes reuse and minimizes redundancy by extracting any sharable models or rules and making them available from a single logical repository.
- They emphasize XML standards for information and process models because those standards facilitate the other five subgoals.

The standards efforts in each industry have generally been led by industry associations or by firms that hope to prosper by offering the key services in the STP pipeline to firms who don't want to perform them. Notable exceptions are efforts in Denmark and Norway, where the strongest mandates and standards for document automation are coming from the government.²⁶

4.3 Views of Business Information

Views of business information are extremely important. We know them as the definitions of business documents and their components that are exchanged between different organizations or enterprises (or, more precisely, between their information systems or services).

4.3.1 Physical Views of Business Information

Physical views of documents have a long history of defining the interface a business presents to the world. Standard printed forms to initiate or record transactions, taxes, and other business activities have existed for centuries. By the mid-19th century, accounting practices and associated documents like balance sheets, statements of accounts, and business registrations were standard enough for the British government to mandate annual audits.²⁷

4.3.1.1 Electronic Data Interchange

Efforts to standardize electronic documents began more than 40 years ago in the trucking industry; spread to the banking, grocery, and retail sectors; and ultimately led in the 1980s to ongoing national and international standards activities for electronic data interchange (EDI).

Efforts to standardize electronic documents began more than 40 years ago

EDI was developed to automate the exchange of structured information in transactional documents such as orders, invoices, and payments between business applications. Initially these exchanges took place over dedicated leased telephone lines or over private networks in a batch store-and-forward fashion. By the 1980s, EDI had penetrated a variety of industries, especially automotive, aerospace, transportation, manufacturing, and retail, where relatively small numbers of firms are the dominant buyers from a large number of suppliers. The ANSI X12 U.S. standards and GTDI European standards began to emerge at this time, followed shortly by the ISO 9735 (EDIFACT) standard developed by the United Nations to consolidate numerous national EDI standards.²⁸

In theory, the EDI standards for documents and the business processes they support should be good starting points for relationships between trading partners. But EDI has fallen short of this promise. The competing X12 and EDIFACT standards are somewhat incompatible, and both syntaxes are brittle and encourage the overloading of semantics into opaque code lists. Furthermore, because the EDI standards process is formal and tedious, it takes a long time to create new standard documents, and the resulting standards are often a bloated laundry list of requirements in which almost everything is optional.

In practice the EDI standards are never used in standard ways

So in practice the EDI standards are never used in standard ways. The dominant trading partner typically selects a small subset of the information components from the standard document and imposes ad hoc “implementation guidelines” on the organizations who do business with it. For example, the EDI requirements imposed by Kroger,²⁹ a very large U.S. supermarket chain, specify extensive adaptations of the standard EDI document types to which its suppliers must conform. It is easy to understand how this subtractive customization approach makes EDI integration expensive for businesses that must adhere to the document formats imposed by the dominant trading partners they serve.

Despite these many limitations EDI remains an important technology. For firms in established business relationships that have made substantial investments to make EDI work, the sense is “if it ain’t broke, don’t fix it.” EDI is not the technology of choice

when setting up new document exchanges with business partners, especially when new document types must be developed for new business processes. But like mainframes and fax machines, EDI can still claim “I’m Not Dead Yet.”³⁰

4.3.1.2 XMLification

When XML emerged in the late 1990s as the preferred syntax for describing document formats, the EDI standards began to “XMLify,” and scores of XML business vocabularies emerged.³¹ As with early efforts in EDI, most of the latter were developed in specific vertical industries by trade associations or industry consortia to reduce the development and integration costs for small and medium-sized enterprises that could not afford to invest in EDI solutions.

New XML specifications often reinvent definitions of common information components

But while each new XML specification for a particular industry was a step forward for that industry, they have proliferated definitions of information components that cut across different industries. Each vocabulary reinvented descriptions of businesses and individuals, measurements, date and time, location, country codes, currencies, business classification codes, and basic business forms like catalogs, purchase orders, and invoices. As is often the case with new technologies, it was two steps forward and one step back.

The earliest effort to attack the problem of semantic overlap among XML vocabularies for business applications was the XML Common Business Library, whose first version was released in 1997. XCBL is a freely distributed set of XML business documents and a set of reusable components common to many business processes. XCBL, like many models of business information, is tied to specific technologies or syntaxes. We call them *document implementation models*. This means that they are typically published as libraries of XML schemas with the expectation that they will be reused at this physical level. The underlying semantics encoded in the vocabularies are only implicit or, at best, incompletely documented.

Because of the physical level of the models, syntax differences like those between X12 and UN/EDIFACT with EDI, or between either of these and an XML vocabulary,

can get in the way of doing electronic business, even if the concepts underlying the documents being exchanged are compatible. Communication usually requires a knowledgeable person to manually create a semantic map between corresponding syntactic components in the pair of models. This has given rise to a category of integration technology that attempts to reuse these semantic maps.³²

The reason physical level mapping is difficult is that it requires a common abstract view that defines the concepts involved rather than the implementation technology. So we need conceptual counterparts to our various physical models (see Section 4.3.2).

4.3.1.3 Information Aggregations

Information composites or aggregations occur where documents or data from numerous sources are brought together to create a consolidated information resource that is more valuable than the sum of the sources. In enterprise contexts this composite resource is typically called a *data warehouse* or *data mart*. Another common composite pattern is a *multivendor catalog* that includes product information from many manufacturers or suppliers. Other examples can be seen in daily shipping schedules and stock market trading tables.

Composite information sources can be created by extracting and transforming the original information and are usually built during “off hours” to minimize the impact on production systems, but as businesses become more global it is always “on hours” somewhere. So the challenge facing the enterprise to keep the composite repository accurate becomes more difficult as the source information becomes more volatile.

An alternative approach is to create a *virtual repository* or *virtual catalog* in which the metadata from each source is normalized and aggregated into the composite resource, not the content itself. This composite metadata enables the content information to be extracted from its source and dynamically transformed to the target schema and vocabulary when the information is requested.³³

Another information composition pattern is *syndication*, the consolidation and distribution of information products. This is widespread in traditional publishing with information like news events, articles, and editorial cartoons collated into a stream of

syndicated content in which items can be selected, routed, and managed using standard metadata on each piece of content.

4.3.2 Conceptual Views of Business Information

We noted in Section 4.3.1 that implementation models of business information have a long history and are quite common. In contrast, models that embody a more conceptual view are a more recent development.

This is hardly surprising. Conceptual views are intellectually more challenging to develop and less beneficial in the short-term than physical ones. Even though models based on conceptual views pay off over time in greater robustness and adaptability, the investment it takes to develop an understanding of the semantic concepts in a domain is often seen as delaying the “real work” of implementation.

Conceptual views are more challenging to develop than physical ones

A notable attempt to develop conceptual models of business information is David Hay’s *Data Model Patterns*, whose subtitle *Conventions of Thought* emphasizes the abstractness and implementation-independence of good models. Hay’s models cover the basic subject areas of people and organizations, products and inventory, procedures and activities, and accounting. A similar book that organizes conceptual models by industry is Len Silverston’s *Data Model Resource Book*.³⁴

The ebXML initiative, launched in 1999 as a joint venture of EDI and XML standards organizations, was the first serious attempt to create conceptual views of business information that could be used in document implementation models in any syntax. The resulting document exchanges would be interoperable because of their common semantic foundations called the *core components*.³⁵ Unfortunately, the ebXML effort was not entirely successful at delivering on its promise to create standard business components, but more because of organizational and political squabbling between the standards groups than for technical problems it couldn’t overcome.

Nevertheless, ebXML paved the way for the Universal Business Language effort, which seems to be succeeding in its goal of creating a standard XML vocabulary for business that is based on a conceptual document component model.

SIDEBAR: The Universal Business Language

The Universal Business Language (UBL) effort began in late 2001 with the extremely ambitious goals of building on the ebXML core components, synthesizing the leading XML and EDI vocabularies for business, and creating standard business documents that would be nonproprietary and royalty free. In effect, it is attempting to provide the equivalent of HTML for e-Business use. It took over two years, but UBL met these goals with the release of version 1 of the UBL library in May 2004.

The UBL Library consists of various document implementation models defined using reusable XSD types. These are based on the UBL document component model for common business components like Party, Address, and Item. These components are reused in document assembly models for basic procurement, including Order, Order Response, Order Change, Order Cancellation, Despatch Advice, Receipt Advice, and Invoice—with many more documents on the way.

A formal set of rules can be applied to transform these document assembly models into document implementation models. In UBL, this means encoding them using XML schema. The UBL Naming and Design Rules define best practices for transforming the assembly model into the implementation model. These rules specify the use of elements and attributes, naming conventions, namespaces, modularity, versioning, and other considerations about how best to exploit XML schemas. These rules have been embedded into various computer applications that automate the generation of XML schemas.

For UBL to succeed as a standard global document format, it must deal with the challenge that most companies are part of numerous supply chains or trading partner relationships that require slightly different documents. It is simply impossible to create semantic components and documents that will work in all situations without customization. Instead, UBL aims to make 80 percent of the library directly useful as is, with the remainder requiring some customization. Making most of the library generic invokes the corollary to the 80/20 rule that the remaining 20 percent customization causes 80 percent of the complexity. For this reason, the UBL initiative is now developing a context methodology to support controlled customization in the document implementation model. Other areas of customization include localization of UBL into

different business regions and languages such as Chinese, Japanese, Spanish, and Korean. This last issue suggests yet another challenge for UBL—the need to fit in with other XML business information and messaging standards. UBL recognizes that no one vocabulary can express all the relevant semantics for business. So UBL has based its models on the ebXML core components metamodel, making it easier to align conceptually with vocabularies also based on that metamodel.

UBL has attracted worldwide interest from industry associations and governments and is on track to be both an OASIS standard and an international standard for trade through ISO Technical Committee 154.³⁶

4.4 Views of Business Architecture

Another important way to describe businesses is in terms of the information technology or systems they use. This method is especially common for businesses to which the Internet is strategic or essential; recall how popular the term “e-Business” was a few years ago. Firms like Amazon, eBay, or Google, none of whom could exist without the Internet, often tout their technology innovations.

Organizations can't have a business relationship if they can't efficiently share information

When different organizations within an enterprise or different firms want to do business with each other, they would prefer not to have to know anything about the systems or technologies each uses to carry out their respective activities. Nevertheless, they can't have a business relationship if they can't efficiently share information, so someone always needs to be concerned with how the business systems fit together. We call this the *business architecture*. In this context we define *architecture* as an abstract specification of a system or thing that describes its components and their relationships with each other, using hierarchical and compositional structure to define component boundaries.

4.4.1 Physical Views of Business Architecture

A description of a system and its components as a physical model is a *systems architecture*. A systems architecture describes a business in terms of its computing platforms, operating systems, databases, and software applications.

4.4.1.1 Technology Platforms and Infrastructure

Sometimes we characterize the business architecture of an enterprise in terms of its dominant software architectures or technology suppliers; this is often called its *platform*. We contrast Microsoft or SAP shops with J2EE or Linux or PeopleSoft ones. As XML takes hold as an implementation technology for document-intensive and Internet-based business systems, XML-centric system architectures have evolved to promote what works best in their design and implementation.³⁷

Companies that have implemented enterprise resource planning (ERP) systems often have a similar technology-centered perspective on how they are organized. ERP systems connect manufacturing control, production planning, inventory, procurement, finance, and human resources systems through a single database, or through a set of linked databases using middleware of some kind. The common data and associations among applications have been described as the “enterprise nervous system.”³⁸

Physical system architectures are often depicted using topology diagrams that show the key information repositories (like databases), computing resources (server farms), and dedicated communications links and networks needed to move data and documents around. These models are often closely related to or overlaid on facilities plans like those described in Section 4.1.1.1. The locations of company headquarters, data warehouses, call centers, and other computing or communications convergence points can be represented in a systems architecture diagram to create an organizational technology “wiring diagram” for the business.

4.4.1.2 Integration Architectures and Patterns

An important corollary to the systems architecture, which shows the interconnections between software systems or applications, is the architecture by which this integration is achieved. *Integration* is defined as the controlled sharing of data and business processes between any connected applications or data sources.³⁹

Integration is the controlled sharing of data and business processes between any connected applications or data sources

The number of potential integration points multiplies with the number of architectural components on each side; simply put, if each side followed the classic three-tier architecture with data-application-presentation layers, there would be nine possible categories of integration techniques. The specific techniques for getting information from one system or application to another also vary immensely to deal with numerous generations of software architectures.

Integration approaches that depend on implementation details or other characteristics at the physical level are said to be *tightly coupled*. At one extreme are “screen scraping” or database extraction approaches that extract data from legacy mainframe databases that were not designed to share information, techniques that require detailed analysis of the screen layout or internal record and table structures. More modern applications are often integrated within an enterprise through a shared data store or warehouse, or by synchronously invoking application program interfaces. Application layer to application layer coupling through APIs is used when the interconnected systems must exchange data at high transaction rates.

Tight coupling is used to exchange data at high transaction rates

Too often, however, the APIs may be very fine grained while carving up the application functionality in incompatible ways. Exchanging information using APIs in this situation requires many small method invocations that extract and set only one or two data values at a time, making the process cumbersome and brittle with all the liabilities of tight coupling and few of the benefits.

Loose coupling is necessary for integration across enterprise boundaries because interfaces might change

Tightly-coupled approaches generally aren’t suitable for integration across enterprise boundaries because of the likelihood of uncontrolled or unexpected changes to interfaces. Instead, cross-enterprise integration approaches try to avoid relying on implementation details, making them more *loosely-coupled*. Loosely-coupled techniques, which we

discuss further in the next two sections, raise the level of abstraction of the integration problem.

4.4.1.3 Web Services

Web services have emerged in the last few years as an important physical architectural idea especially for business-to-business relationships where looser coupling through document exchange is required or desirable (see Section 1.3.3).

Because almost anything can be turned into a service by “wrapping” it in XML document interfaces, there has been enormous hype about Web Services. A typical claim is “What the Web did for program-to-user interactions, Web Services are poised to do for program-to-program interactions.”⁴⁰ Some disappointment may set in when we realize that the essence of Web Services is a few simple specifications for using XML and messaging for application integration.

We can explain the concepts embodied by the primary Web Services specifications with a simple analogy to mailing a letter. First, if we don’t already know the party to whom we should send the letter requesting the service we want, we need a business directory in which we can find service providers. Next we need a service description that tells us the address of the service provider, what kind of message to send, and what kind of response to expect. Finally we need to know how to address the envelope and how to put the letter inside it.

More formally, a Web Service is defined as a platform-independent implementation of functionality that conforms to published specifications for the XML documents it sends and receives as its public interfaces (Web Service Description Language), the messaging protocol used to send and receive XML documents through those interfaces (SOAP), and a searchable directory of services (UDDI). Since these three initial specs were proposed in late 2000, many so-called “standards” have proliferated for other components implied by a completely service-oriented architecture but none of the basic ideas has substantially changed.⁴¹

Web Services enable a more loosely-coupled integration approach than previous integration technologies

Because they can wrap a hodgepodge of legacy technologies and hide proprietary data models and protocols with XML document interfaces, Web Services provide a layer of abstraction and enable a more loosely-coupled integration approach than previous integration technologies. However, this doesn't immediately solve the integration problem. Security, reliable delivery, performance, scalability, and other critical issues for deploying enterprise-level Web Services aren't completely handled by current specifications and vendors.

But there is a more fundamental reason why Web Services alone don't solve the integration problem. While a Web Service's technical specifications dictate how to reveal the interfaces and message definitions for the XML documents that it sends and receives, they say nothing about the conceptual design of those services and their enabling documents. They tell us how to package information into documents and where to put them, but they don't tell us what any of it means.

4.4.2 Conceptual Views of Business Architecture

In contrast to physical systems architectures, the architecture of a business can be described in more abstract terms, sometimes called a *logical architecture*. A logical architecture doesn't concern itself with specific implementation technologies but instead emphasizes topological or structural relationships between the functional components of business systems. Vendor and technology-neutral concepts like "N-tier," "middleware," "gateways," and "service networks" are used in logical architectures to describe the conceptual arrangement of computing and communications resources.

A logical architecture can portray the boundaries or interconnections among business systems and represent the extent to which systems are centralized or distributed within an enterprise. Architectural patterns reflect different requirements for system communication or integration. An architectural description can reveal the extent and direction of information exchanged between systems. It can also identify systems that are isolated islands or silos of functionality because they can't easily exchange information with other ones.

IBM's patterns for e-business⁴² are a rich source of conceptual models of business architectures. The IBM patterns grew out of an internal IBM effort to systematize the best

practices of its consulting division and identify feasible architectures for large-scale e-business applications.

At the top of the conceptual model hierarchy are what IBM calls the Business Patterns, which describe at the most conceptual level the ways in which users and businesses interact with information. There are four Business Patterns: Self-Service (also known as “user-to-business” or B2C), Collaboration (also known as “user-to-user” or C2C), Information Aggregation (also known as “user-to-data”), and Extended Enterprise (also known as “business-to-business” or B2B). These basic Business Patterns can be combined to create more complex patterns. One example is the “e-marketplace” pattern, which enables buyers and sellers to trade goods and services on a public website by combining the Self-Service and Information Aggregation patterns.

Similar conceptual patterns have been proposed by Weill and Vitale.⁴³ They describe eight “atomic” business patterns, each of which describes a distinct but irreducible business function, such as Content Provider, Direct to Consumer, and Intermediary. According to this approach, businesses compound the atomic patterns into more sophisticated business architectures.

Many of these business architecture patterns rely heavily on documents as user or service interfaces, but the patterns have never before been organized in a way that makes this explicit. In Chapter 15 we introduce our own framework, which emphasizes document interfaces, exchanges, and the management of components and metadata that they require.

4.4.2.1 Conceptual Views of Integration Architecture

It is preferable for many of the participants in a business relationship to take a technology-independent and conceptual view of the integration architecture and focus on the more abstract goal of *interoperability*.

Interoperability is a more abstract goal than integration

Interoperability means that the recipient can extract the required information from the sender’s document even if the sender’s implementation is not immediately compatible with the recipient’s business systems. This might require some reverse engineering of the

underlying conceptual model from the physical model in which the sender's information is encoded. Then the recipient must establish that the extracted conceptual model is what it needs to carry out the intended process. If this is established, transforming a different implementation to an encoding from which the needed information can be extracted is a necessary, but often trivial thing to do.

In Chapter 6, "When Models and Patterns Don't Match: Interoperability Challenges" we describe a range of examples that illustrate interoperability problems.

While it is easy to understand why interoperability challenges can arise when systems from different technology generations must be integrated, technology is neither the primary cause of this incompatibility nor the primary means of eliminating it. The best way to facilitate interoperability is often for the participants in the exchange to jointly define a conceptual model for the shared information, or for both of them to adopt the same industry standard. This approach allows them to use the same information model without any constraints on their implementation of it.

The best way to facilitate interoperability is for the participants to share the same conceptual model

There is no precise point when reducing the assumptions and dependencies between the participants turns the physical view required by a tightly coupled relationship into the more conceptual one implied by a loosely coupled relationship.⁴⁴ But a loosely coupled approach generally means that information is exchanged asynchronously rather than synchronously, and in larger, document-sized chunks governed by an explicit schema or model (as it would be by an industry standard) rather than as fine-grained information pieces whose semantic definition is implicit only in the integration code.

The benefits of a loosely coupled approach mean that for the interorganizational and interenterprise applications that are at the core of Document Engineering, the most practical integration architecture is often messaging. Applications communicate by sending messages to a "channel" that ensures the reliable asynchronous delivery to the recipient while vastly reducing how much the sender and recipient must know about each other's technology.⁴⁵ Messaging systems or messaging-oriented middleware must still be configured to fit the addressing, packaging, security, and delivery requirements of each

situation, but bringing all these concerns together substantially reduces the complexity of the integration challenge.

4.4.2.2 Service-Oriented Architectures

Web services allow a business to take a more abstract view of implementation and integration, and it is reasonable to deploy them in an incremental, point-to-point, and bottom-up manner to integrate systems two at a time. However, an even more abstract view of services in a business architecture is the top-down and strategic one that considers everything a business does as (potentially) realized by business service components that are combined and recombined as needed. This perspective defines a *service-oriented architecture* or *SOA*.

A SOA imposes a very abstract perspective on supply chains, marketplaces, drop shipment, and other processes because it deemphasizes technology and platform considerations and views them all as combinations of services. The emergence of *SOA* as an industry buzzword in recent years has been accompanied by other new terms like *enterprise ecosystem*, *enterprise service bus*, and *business service network* that likewise imply more generic approaches for enabling the interconnection of business services.⁴⁶

For example, a service-oriented view of marketplaces defines them entirely in terms of their participants and the set of services that they offer each other. (see Section 4.1.2.3). The drop shipment pattern followed by our hypothetical GMBooks.com bookstore (see Section 1.1) could be realized using a SOA that combines component business services like the Amazon.com catalog,⁴⁷ UPS package delivery and tracking functions, and Visa payment processing. All of these are available as document-based Web Services for easy integration into other business systems.

Furthermore, a SOA perspective highlights the principles of discovery and transparent substitutability of service providers because their roles and functional responsibilities are strictly defined by the XML documents that they produce and consume when providing a service. This is elegantly demonstrated by a Silicon Valley firm called Talaris, which hosts a procurement application for employee business services like travel, package shipping, conferencing, mobile communications, ground transportation, and other services consumed directly by end users. The Talaris application is built natively using

Web Services and SOA principles, enabling it to describe each class of end user services abstractly in an XML vocabulary called the Services Business Language.⁴⁸ Each SBL document harmonizes the APIs or functionality from multiple providers of the same service into a single interface. Each service provider receives exactly the same service request, and suppliers can be added or dropped without any changes to the SBL or the user experience.

An essential and emergent benefit of an SOA is that once some application functionality is re-packaged as a service, new *composite applications* can be developed by combining them. Furthermore, because of the abstraction provided by document interfaces and the Web Services standards, composite applications can be created with vastly less effort than required by tightly-coupled integration approaches. Another Silicon Valley firm called Above All Software⁴⁹ has developed visual tools that enable non-programmers like business analysts to create user interfaces that unify the inputs and outputs to separate services. For example, a Web Service that looks up customer details in a customer database can be combined with one that knows about orders in an ERP system, creating a combined service that locates the current orders for any specified customer.

Applied to the GMBooks.com user interface, this composite services architectural approach means that the Amazon catalog could be transparently replaced by one from Barnes and Noble, UPS by FedEx, and Visa by American Express. Figure 4-3 illustrates the idea of composite services with transparent substitutability of the service providers.

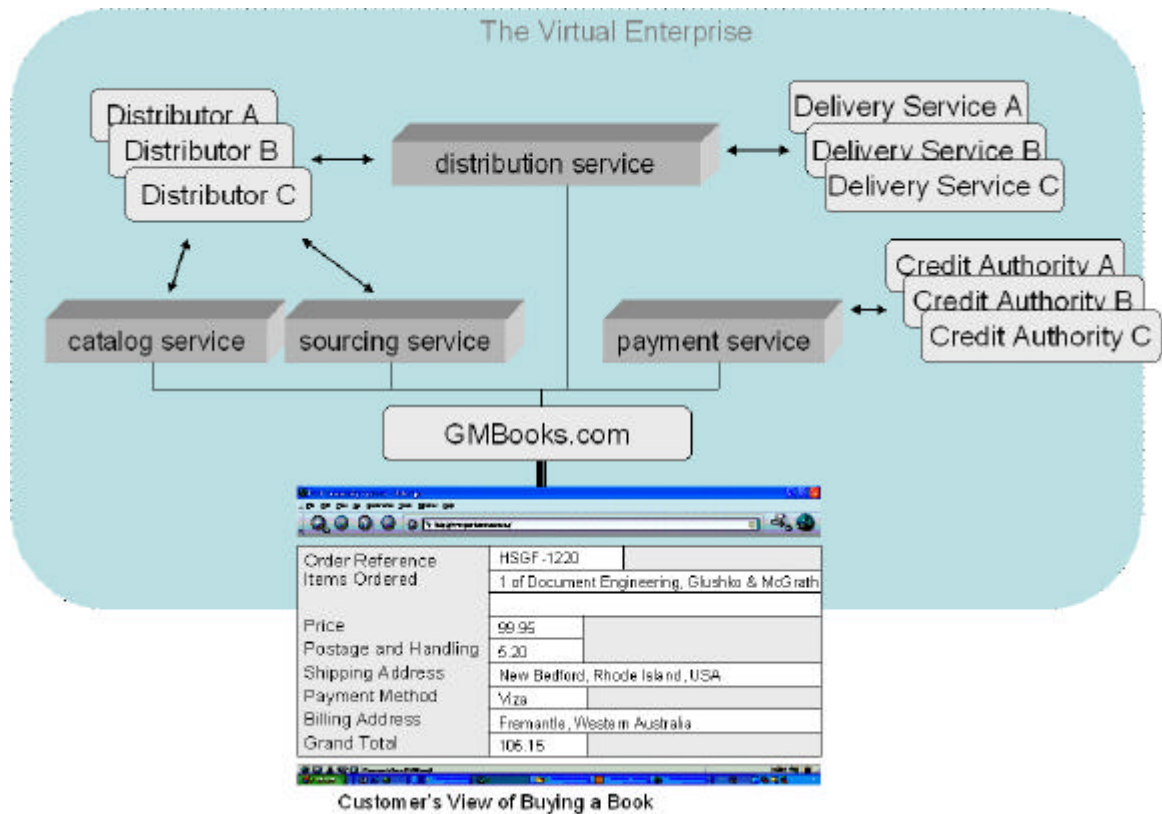


Figure 4-3. Composite Services and Transparent Substitutability

And when General Motors transforms the different APIs for the inventory systems in each of their 80 factories into Web Services using the same interfaces, this lets the firm abstract what's going on in each factory.⁵⁰ Any web-services-enabled application can then get inventory information from any factory whenever it's needed even if the underlying inventory application is changed.

Service Oriented Architecture is a design philosophy; Web Services are a set of standards and techniques

The examples of services we've described in this section illustrate the contrast between SOA as a business design philosophy and Web Services as a set of standards and techniques for platform-independent integration. A SOA perspective drives a business to ask strategic questions like these as it systematically structures its business capabilities as self-contained resources or processes:

- What patterns of service combination are required to meet our business objectives?

- How can we design what each service does so that as a set they will be sufficient and flexible enough as business conditions change?
- Which of these services can we “carve out” of existing applications by changing their implementations or APIs?
- Which services should we build ourselves, and which should we obtain from others?
- Should we offer any of our services to other firms?

Because it makes decisions about the design of services depending on business considerations, a SOA approach tends to yield services that are more process oriented and that provide “coarser” or larger units of functionality with greater business value than the services that emerge from the more technical perspective of Web Services. This is not to denigrate useful Web Services like those that provide current weather conditions, that decode a coded value (for example, “What country has ‘IS’ as its ISO 3166 code?”), or that return the author and title of a book given its ISBN.

But the functionality of these bottom-up and more opportunistically provided Web Services is relatively small, low impact, and not likely to create compelling or competitive advantages for their providers or users. The Web Services that are more likely to do so are those that produce and consume entire business documents like catalogs, orders, invoices, and payments.

4.5 Key Points in Chapter 4

- A business model is concerned with the nature and pattern of exchanges of one form of value for another.
- Organizations can’t have a business relationship if they can’t efficiently share information.
- By eliminating any need for physical presence the Internet has increased the feasibility and conceptual variety of business models.
- A reference model consolidates the best practices of many companies.
- Supply chains, distribution channels, markets and auctions are general business patterns that can be applied in novel contexts.

- Document Engineering treats supply chains as information flows.
- Business process models are the bridge between organizational models and business documents.
- Information about the business processes is distinct from the physical flow of materials and goods.
- Efforts to standardize electronic documents began more than 40 years ago.
- New XML specifications often reinvent definitions of common information components.
- Integration is the controlled sharing of data and business processes between any connected applications or data sources
- Interoperability is a more abstract goal than integration
- The best way to facilitate interoperability is for the participants to share the same conceptual model.
- Tight coupling is used to exchange data at high transaction rates.
- Loose coupling is necessary for integration across enterprise boundaries because interfaces might change.
- Web Services enable a more loosely-coupled integration approach than previous integration technologies.
- Service Oriented Architecture is a design philosophy; Web Services are a set of standards and techniques.

4.6 Notes

¹ Some of the classic works here are Alfred D. Chandler, *The Visible Hand: The Managerial Revolution in American Business* (Cambridge University Press, 1977), William McKelvey, *Organizational Systematics: Taxonomy, Evolution, Classification* (University of California Press, 1982), Oliver Williamson, *Markets and Hierarchies: Analysis and Antitrust Implications* (Free Press, 1975), and Oliver Williamson, *The Economic Institutions of Capitalism*, (Free Press, 1985). Carl Shapiro and Hal Varian, *Information Rules:*

A Strategic Guide to the Network Economy, (Harvard Business School Press), 1999 isn't quite old enough to be a classic but it is on its way.

² Object-oriented programming techniques usually use UML class diagrams in highly physical ways to describe models that have a one-to-one correspondence to implementing code, but in Document Engineering we use them more conceptually (see Section 3.2).

³ Demo of BluePeople seen in May 2004 at IBM Watson Research Center.

⁴ Robert Haugen and William E. McCarthy, "REA, a semantic model for Internet supply chain collaboration," <http://jeffsutherland.org/oopsla2000/mccarthy/mccarthy.htm>. (last visited 20 October 2004).

⁵ This use of a word more commonly associated with biology to describe an Internet-based network of a firm's relationships with other entities originated in Jay Tenenbaum, Tripatinder Chowdry, and Kevin Hughes, "Eco System: An Internet Commerce Architecture," *Computer*, 30 (5), May 1997, 48-55. Scores of articles on ecosystem topics can be found at Internet newsmagazine Line56's E-business Ecosystem section at http://www.line56.com/articles/ebiz_ecosys_index.asp. (last visited 18 October 2004).

⁶ Federal Enterprise Architecture Program Office, "24 presidential priority e-gov initiatives," http://www.feapmo.gov/resources/24_PPE-Gov_Init_Rev_1.pdf (last visited 14 February 2004).

⁷ "Education anytime, anywhere. How do you spell B2B and B2C?" <http://www.trendsreport.net/2000/education/3.html> (last visited 15 February 2004).

⁸ This problem is increasingly attacked with linear or constraint-based mathematical programming techniques that evaluate huge numbers of feasible designs for a supply chain network against some objective function like minimizing cost, maximizing customer service levels, or minimizing cycle time. Some of these capabilities are built into the planning and scheduling components of ERP systems. See Jim Shepard and Larry Lapide. "Supply Chain Optimization: Just the Facts." *ASCET*, Volume 1, 15 April 1999. http://www.ascet.com/documents.asp?d_ID=217# (last visited 20 October 2004).

⁹ Some generally useful services include those that provide information about potential business partners (like credit or customer satisfaction ratings), that facilitate financial or accounting processes (like tax calculation, payment, factoring), or that ensure the delivery of goods (like escrow, trade facilitation, freight forwarding, and shipment).

¹⁰ The most common auction pattern is the *forward* or *English auction* typified by eBay, in which one seller offers something to many potential buyers. When information about price is continuously exchanged, the offered price moves up as buyers bid against each other. Another auction pattern involving aggregated buyers is the *uniform price* or *Dutch auction* in which multiple identical products are available, and the price moves down until there are enough buyers willing to pay that clearing price. This kind of auction has been used in stock IPOs to set an offering price that maximizes the money going to the newly public firm while ensuring that all the new shares are sold. Auction patterns can also involve aggregated sellers, most often in direct procurement where the buyer provides product specifications and the sellers bid against each other in what is called a *reverse auction*. The price the seller pays goes down as the sellers bid against each other. For a serious discussion of auctions, see Vijay Krishna, *Auction Theory* (Academic Press, 2002), or

Lawrence Ausubel, "Auction theory for the new economy," in Derek Jones (Ed.), *New Economy Handbook* (Academic Press, 2003).

¹¹ For example, there are hundreds of administrative policies and procedures at the University of California, Berkeley. One of them guarantees a free reserved parking place for life to any **employee** who wins a Nobel Prize. See <http://campuspol.chance.berkeley.edu/directory.htm> (last visited 15 February 2004).

¹² Federal Enterprise Architecture Program Management Office. Business Reference Model (BRM) Version 2.0. <http://www.feapmo.gov/feaBrm2.asp> (last visited 20 October 2004). Denmark has a similar e-government effort underway: see the Openness Initiative at <http://www.oio.dk/english> (last visited 20 November 2004).

¹³ Supply-Chain Council. Supply-Chain Operations Reference Model. <http://www.supply-chain.org/public/scor.asp>. (last visited 20 October 2004).

¹⁴ RosettaNet PIP Directory, [http://www.rosettanet.org/RosettaNet/Rooms/DisplayPages/LayoutInitial?Container=com.webridge.entity.Entity\[OID\]9A6EEA233C5CD411843C00C04F689339](http://www.rosettanet.org/RosettaNet/Rooms/DisplayPages/LayoutInitial?Container=com.webridge.entity.Entity[OID]9A6EEA233C5CD411843C00C04F689339)] (last visited 20 October 2004).

¹⁵ Tom Krazit, "Intel conducts \$5b in transactions via RosettaNet," *InfoWorld*, 10 December 2002, <http://archive.infoworld.com/articles/hn/xml/02/12/10/021210hntelrose.xml?s=IDGNS> (last visited 14 February 2004). See also "e-Business at Intel." <https://supplier2.intel.com/B2Bi/> (last visited 14 February 2004).

¹⁶ "Value chain" is usually attributed to Michael E. Porter, *Competitive Advantage: Creating and Sustaining Superior Performance* (Free Press, 1998). Treating the flow of information in a value chain as an independent source of value is discussed in Chapter 4, "The Information Supply Chain," in Larry Downes. *The Strategy Machine*. (Harper-Collins, 2002).

¹⁷ Christopher Koch. "It all began with Drayer," *CIO*, 1 August 2002, <http://www.cio.com/archive/080102/drayer.html>. (last visited 20 October 2004). Both parties in a VMI relationship benefit from the increased efficiency of procurement and logistics: the retailer no longer loses sales because goods are out of stock and no longer has to maintain inventory in warehouses, and the supplier can control inventory and transportation costs while providing better service. VMI works best for consumer packaged goods, consumables, and other merchandise that is purchased regularly and in large volumes, such as clothing, cosmetics and groceries.

¹⁸ See <http://www.cpfr.org> (last visited 20 October 2004) for a rich archive of specifications and case studies. See also Dirk Seifert, *Collaborative Planning, Forecasting and Replenishment: How to Create a Supply Chain Advantage* (SAP Press, 2003).

¹⁹ "Remote elevator monitoring," http://www.otis.com/innovationdetail/0.1416.CL11_IID805_RES1.00.html (last visited 20 October 2004). Jay Miller, "Keeping Tabs," *Manufacturer* May 2003, http://www.themanufacturer.com/content_detail.html?header=article&contents_id=1236&t=manufacturer_us#. (last visited 20 October 2004).

²⁰ Much of the work is in response to a mandate by the U.S. Securities and Exchange Commission that firms settle trades in just one day (called T+1, following a T+3 initiative in 1995). Settlement involves getting information from the front office of the selling entity to its back office and then to the back and front office of the buying entity. The goal is to reduce the time to “settle” securities trades to one day after the trade takes place.

²¹ Jonathan Parsons, “Legislation, deliberation, and documents: XML and the legislative process,” *Idealliance XML 2004 Conference*, November 2004, http://www.idealliance.org/proceedings/xml04/papers/179/XML_and_Legislative_Process.html (last visited 20 November 2004).

²² Eric Auchard, “U.S. Army aims to halt paperwork with IBM system.” *Computerworld*, 17 December 2004. <http://www.computerworld.com/printthis/2004/0,4814,98358,00.html> (last visited 5 January 2005).

²³ See US Securities and Exchange Commission, “Spotlight on Sarbanes-Oxley Rulemaking and Reports.” <http://www.sec.gov/spotlight/sarbanes-oxley.htm> (last visited 20 October 2004).

²⁴ Nigel King, “Web services to support Sarbanes Oxley activities,” *Idealliance XML 2004 Conference*, November 2004, <http://www.idealliance.org/proceedings/xml04/papers/16/XML2004.html> (last visited 20 November 2004).

²⁵ Jim Ericson, “Technology rising for SOX,” *Line 56*, 23 November 2004, <http://www.line56.com/articles/default.asp?ArticleID=6172> (last visited 5 January 2005).

²⁶ Mikkel Brun and Brian Nielsen, “Naming and design rules for e-government - The Danish approach,” *Idealliance XML 2003 Conference*, December 2003, http://www.idealliance.org/papers/dx_xml03/papers/05-06-04/05-06-04.html (last visited 20 November 2004).

²⁷ Not surprisingly the British Companies Act of 1844, was soon followed by the creation of a slew of accounting firms in London by people whose names were Deloitte, Price, Waterhouse, Coopers, and Peat. Two centuries later these names remain associated with the largest global accounting firms.

²⁸ The US EDI standards are maintained by the Associated Standards Committee at <http://www.x12.org>. The international EDI standards are maintained by the United Nations Centre for Trade Facilitation and Electronic Business at <http://www.unece.org/cefact/>.

²⁹ Kroger, “EDI Programs & Requirements,” http://edi.kroger.com/edi/programs_001.htm (last visited 14 February 2004).

³⁰ John Edwards, “I’m not dead yet,” *Line56*, May 2001, <http://www.line56.com/articles/default.asp?NewsID=2563> (last visited 14 February 2004).

³¹ Many XML vocabularies are published at web sites whose domain names are the vocabulary acronym and “.org,” like <http://www.xcbl.org/> (last visited 20 October 2004) for the XML Common Business Library. Lists of XML vocabularies can be found in the Cover Pages at <http://xml.coverpages.org/xmlApplications.html> and at <http://xml.org>. (last visited 20 October 2004).

-
- ³² John Edwards, "Doing it with meaning," *CIO Magazine*, 15 August 2002, http://www.cio.com/archive/081502/et_article.html (last visited 20 October 2004). For a more theoretical treatment see Joshua Fox, "Know what your schemas mean," *Idealliance XML 2003 Conference*, December 2003, http://www.idealliance.org/papers/dx_xml03/papers/04-03-04/04-03-04.html (last visited 20 October 2004).
- ³³ See, for example, Narinder Singh, "Unifying heterogeneous information models." *Communications of the ACM*, 41 (1998): 37-44, or Michael Stonebraker and Joseph Hellerstein, "Content Integration for E-Business," *ACM SIGMOD 2001*, 552-560.
- ³⁴ David Hay, *Data Model Patterns: Conventions of Thought*. (Dorset House, 1996).
Len Silverston, *The Data Model Resource Book*, 2 vols. (John Wiley, 2001).
- ³⁵ Electronic Business XML, <http://ebxml.org/>. The original Core Component Dictionary created by the ebXML working group is at <http://ebxml.org/specs/ccDICT.pdf> (last visited 20 October 2004). The Core Components work is now being carried out under the auspices of UN/CEFACT and a more recent version of the specification is at http://www.unece.org/cefact/ebxml/CCTS_V2-01_Final.pdf (last visited 20 October 2004).
- ³⁶ The UBL home page is Organization for the Advancement of Structured Information Standards (OASIS) Universal Business Language. http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=ubl. (last visited 20 October 2004). See Mark Crawford (Editor). OASIS Universal Business Language (UBL) Naming and Design Rules. 15 September 2004. <http://www.oasis-open.org/committees/download.php/9236/cd-UBL-NDR-1.0.pdf> (last visited 20 October 2004). See also Eve Maler, "Schema design rules for UBL...and maybe for you," *Idealliance XML 2002 Conference*, December 2002, http://www.idealliance.org/papers/xml02/dx_xml02/papers/05-01-02/05-01-02.pdf. (last visited 20 October 2004). The UBL customization methodology is described by Arofan Gregory and Eduardo Gutentag in "UBL and object-oriented XML: Making type-aware systems work." *Idealliance XML 2003 Conference*, December 2003, http://www.idealliance.org/papers/dx_xml03/papers/04-04-04/04-04-04.html. (last visited 20 October 2004).
- ³⁷ See, for example, David S. Lithicum, *Next Generation Application Integration* (Addison-Wesley, 2004), or W. Scott Means, *Strategic XML* (SAMS, 2002).
- ³⁸ This phrase originated with the Gartner Group. See Gian Trotta, "Get a grip, with enterprise nervous systems," *ebizQ*, 23 September 2003, http://www.ebizq.net/topics/real_time_enterprise/features/2807.html (last visited 20 October 2004).
- ³⁹ David S. Linthicum, *B2B Application Integration* (Addison-Wesley, 2001).
- ⁴⁰ Heather Kreger, "Web services conceptual architecture," IBM, <http://www-4.ibm.com/software/solutions/webservices/pdf/WSCA.pdf>, page 6. We single out this author only because her report was one of the first clear explanations of Web Services and undoubtedly encouraged others to adopt a similarly enthusiastic perspective.

⁴¹ The idea of using XML specifications for services and the documents they exchanged had emerged a few years earlier in a 1997 proposal titled “XML in Component-Based Commerce” to the U.S. Department of Commerce Advanced Technology Program by several Silicon Valley firms. See Brad Meltzer and Robert Glushko, “XML and electronic commerce: Enabling the network economy,” *ACM SIGMOD* 27 (1998), and Robert J. Glushko, Jay M. Tenenbaum, and Bart Meltzer, “An XML framework for agent-based commerce,” *Communications of the ACM*, 42 (1999): 106-114. This work also inspired a quasi-standards effort called the eCo Framework that in 1998-1999 developed a set of specifications two years before the Web Services “standards stack.” See “eCo Architecture for Electronic Commerce Interoperability,” <http://www.commerce.net/docs/ecoframework.pdf> (last visited 20 October 2004). The ebXML initiative, begun in 1999, had also begun to cover some of the same ground a year before the Web Services specifications emerged. The most important work today on specifications for Web Services is being conducted under the auspices of the Web Services Interoperability Organization, <http://www.ws-i.org/> (last visited 20 October 2004). For an impartial compendium of articles and news about WS-I efforts see <http://xml.coverpages.org/ws-i.html> (last visited 20 October 2004).

⁴² Jonathan Adams, Srinivas Koushik, Guru Vasudeva, and George Galambos, *Patterns for E-Business: A Strategy for Reuse* (IBM Press, 2001). See also IBM Patterns for E-Business, <http://www-106.ibm.com/developerworks/patterns/>.

⁴³ Peter Weill and Michael R. Vitale, *Place to Space*, (Harvard Business School Press, 2001)

⁴⁴ David Kaye, author of *Loosely Coupled: The Missing Pieces of Web Services*, (RDS Associates, 2003), <http://www.rds.com/doug/weblogs/webServicesStrategies/2002/11/18.html> (last visited 20 October 2004), says “Loose coupling is like pornography. Everyone talks about it, but when challenged, few can tell you what it is.”

⁴⁵ Gregor Hoppe and Bobby Wolfe, *Enterprise Integration Patterns: Designing, Building, and Deploying Messaging Systems*, (Addison-Wesley, 2004).

⁴⁶ See note 5 in this chapter. See also Atul Saini, “Demystifying the enterprise service bus,” *Business Integration Journal* (September 2003): 24-27, <http://bijonline.com/Article.asp?ArticleID=764&DepartmentID=9> (last visited 22 October 2004) and Marty Tenenbaum, “CommerceNet’s vision: Millions of interoperable business services,” *CommerceNet Whitepaper*, http://www.commerce.net/docs/BSN_vision.pdf (last visited 22 October 2004). It is also enlightening to study the changes in product positioning of software vendors like CommerceOne, Ariba, Web Methods, and BEA since the B2B bubble began to burst in 2001. CommerceOne developed the first XML-based marketplace platform in 1999, followed by other vendors who offered similar software that was highly functional and complex and came bundled with a suite of marketplace, supply chain, and auction services. Today the surviving companies have reimplemented and repositioned their software to have a much lighter footprint and to function as more generic service integration platforms; none offers completely packaged “marketplace” software anymore.

⁴⁷ Nick Wingfield, “In latest strategy shift amazon is offering a home to retailers,” *Wall Street Journal*, 24 September 2003. See also Amazon Web Services at <http://www.amazon.com/gp/browse.html/104-5602037-8299135?node=3435361> (last visited 22 October 2004).

⁴⁸ Talaris Corp., “Services business language (SBL): Supplier integration using SBL,” http://www.talaris.com/technology/SBL_Whitepaper.pdf (last visited 22 October 2004).

⁴⁹ Above All Software, “Above All Studio.” <http://www.aboveallsoftware.com/products/studio.asp> (last visited 22 October 2004). Above All Software was founded by Roger Sippl, a Silicon Valley serial entrepreneur who also founded Informix, Vantive, and Visigenix. Just as Informix did nearly 30 years ago with SQL for relational databases and Visigenix did a decade ago with application servers, Above All is making composite services that latest step in an evolutionary trend to raise the level of abstraction to increase the reuse information assets and business logic.

⁵⁰ Eric Knorr, “Enterprises sketch out service-oriented architectures.” *Infoworld.com*, 26 November 2003, http://reviews.infoworld.com/article/03/11/26/47FEwsretrofit_1.html?s=feature. (last visited 22 October 2004).