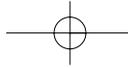


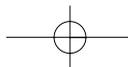
Part 1

Introduction



The term *Enterprise Architecture* refers to a collection of technology components and their interrelationships, which are integrated to meet the information requirements of an enterprise. This section introduces the concept of Enterprise IT Architectures with the intention of providing a framework for the various types of technologies used to meet an enterprise's computing needs.

Data warehousing technologies belong to just one of the many components in an IT architecture. This chapter aims to define how data warehousing fits within the overall IT architecture, in the hope that IT professionals will be better positioned to use and integrate data warehousing technologies with the other IT components used by the enterprise.



Chapter

1

The Enterprise IT Architecture

This chapter begins with a brief look at how changing business requirements have, over time, influenced the evolution of Enterprise Architectures. The *InfoMotion* (“Information in Motion”) Enterprise Architecture is introduced to provide IT professionals with a framework with which to classify the various technologies currently available.

► The Past: Evolution of Enterprise Architectures

The IT architecture of an enterprise at a given time depends on three main factors:

- the business requirements of the enterprise;
- the available technology at that time; and
- the accumulated investments of the enterprise from earlier technology generations.

The business requirements of an enterprise are constantly changing, and the changes are coming at an exponential rate. Business requirements have, over the years, evolved from the day-to-day clerical recording of transactions to the automation of business processes. Exception reporting has shifted from tracking and correcting daily transactions that have gone astray to the development of self-adjusting business processes.

Technology has likewise advanced by delivering exponential increases in computing power and communications capabilities. However, for all these advances in computing hardware, a significant lag exists in the realms of software development and architecture definition. Enterprise Architectures thus far have displayed a general inability to gracefully evolve in line with business requirements, without either compromising on prior technology investments or seriously limiting their own ability to evolve further.

In hindsight, the evolution of the typical Enterprise Architecture reflects the continuous, piecemeal efforts of IT professionals to take advantage of the latest technology to improve the support of business operations. Unfortunately, this piecemeal effort has often resulted in a morass of incompatible components.

► The Present: The IT Professional's Responsibility

Today, the IT professional continues to have a two-fold responsibility: Meet business requirements through Information Technology and integrate new technology into the existing Enterprise Architecture.

Meet Business Requirements

The IT professional must ensure that the enterprise IT infrastructure properly supports a myriad set of requirements from different business users, each of whom has different and constantly changing needs, as illustrated in Figure 1-1.



Figure 1-1 Different Business Needs

Take Advantage of Technology Advancements

At the same time, the IT professional must also constantly learn new buzzwords, review new methodologies, evaluate new tools, and maintain ties with technology partners. Not all the latest technologies are useful; the IT professional must first sift through the technology jigsaw puzzle (see Figure 1-2) to find the pieces that meet the needs of the en-

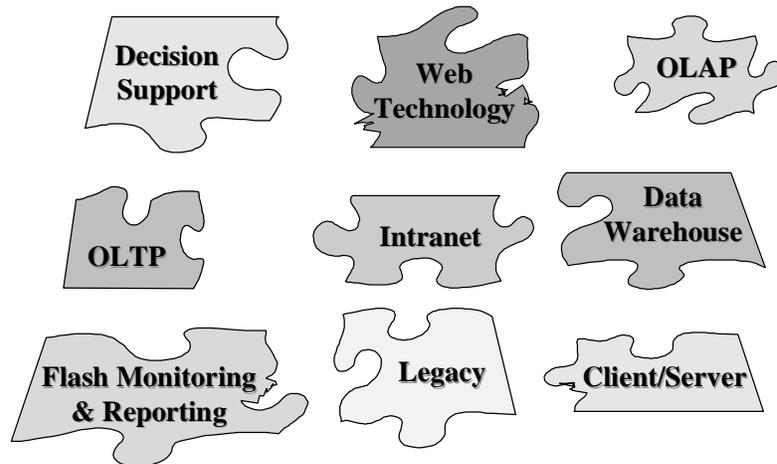


Figure 1-2 The Technology Jigsaw Puzzle

terprise, then integrate the newer pieces with the existing ones to form a coherent whole.

One of the key constraints the IT professional faces today is the current Enterprise IT Architecture itself. At this point, therefore, it is prudent to step back, assess the current state of affairs and identify the distinct but related components of modern Enterprise Architectures.

The two orthogonal perspectives of business and technology are merged to form one unified framework, as shown in Figure 1-3.

► Business Perspective

From the business perspective, the requirements of the enterprise fall into categories illustrated in Figure 1-4 and described below.

Operational

Technology supports the smooth execution and continuous improvement of day-to-day operations, the identification and correction of er-

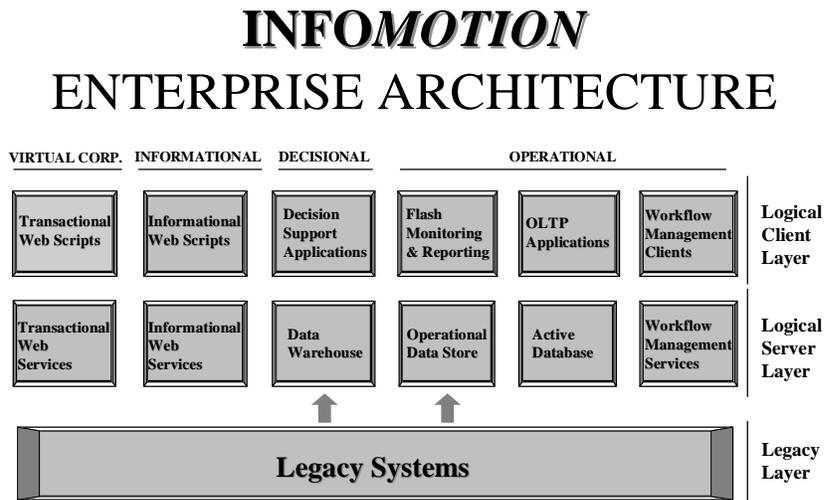


Figure 1-3 The InfoMotion Enterprise Architecture

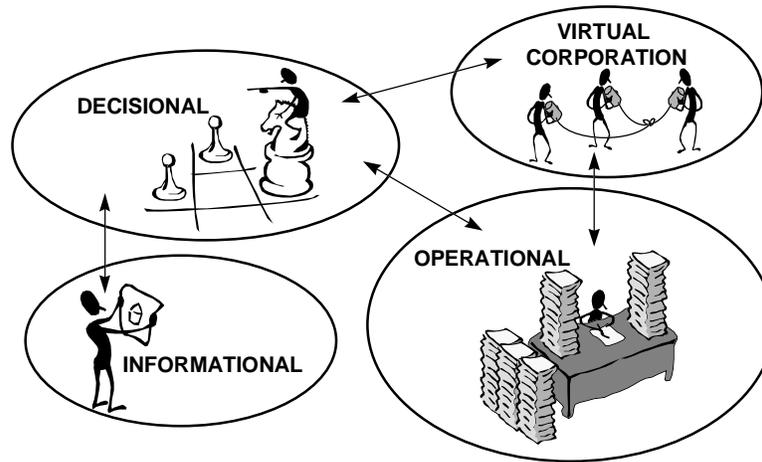


Figure 1-4 The Enterprise Architecture (Business Perspective)

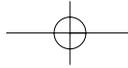
rors through exception reporting and workflow management, and the overall monitoring of operations. Information retrieved about the business from an operational viewpoint is used to either complete or optimize the execution of a business process.

Decisional

Technology supports managerial decision-making and long-term planning. Decision-makers are provided with views of enterprise data from multiple dimensions and in varying levels of detail. Historical patterns in sales and other customer behavior are analyzed. Decisional systems also support decision-making and planning through scenario-based modeling, what-if analysis, trend analysis, and rule discovery.

Informational

Technology makes current, relatively static information widely and readily available to as many people as need access to it. Examples include company policies, product and service information, organiza-



tional setup, office location, corporate forms, training materials, company profiles.

Virtual Corporation

Technology enables the creation of strategic links with key suppliers and customers to better meet customer needs. In the past, such links were feasible only for large companies because of economies of scale. Now, the affordability of Internet technology provides any enterprise with this same capability.

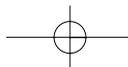
► Technology Perspective

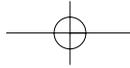
This section presents each architectural component from a technology standpoint and highlights the business need that each is best suited to support.

Operational Needs

Legacy Systems

The term *legacy system* refers to any information system currently in use that was built using previous technology generations. Most legacy systems are operational in nature, largely because the automation of transaction-oriented business processes had long been the priority of Information Technology projects.





OLTP Applications

The term *Online Transaction Processing* refers to systems that automate and capture business transactions through the use of computer systems. In addition, these applications traditionally produce reports that allow business users to track the status of transactions. OLTP applications and their related active databases compose the majority of client/server systems today.

Active Databases

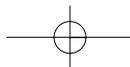
Databases store the data produced by Online Transaction Processing applications. These databases were traditionally passive repositories of data manipulated by business applications. It is not unusual to find legacy systems with processing logic and business rules contained entirely in the user interface or randomly interspersed in procedural code.

With the advent of client/server architecture, distributed systems, and advances in database technology, databases began to take on a more active role through database programming (e.g., stored procedures) and event management. IT professionals are now able to bullet-proof the application by placing processing logic in the database itself. This contrasts with the still-popular practice of replicating processing logic (sometimes in an inconsistent manner) across the different parts of a client application or across different client applications that update the same database. Through active databases, applications are more robust and conducive to evolution.

Operational Data Stores

An Operational Data Store or ODS is a collection of integrated databases designed to support the monitoring of operations. Unlike the databases of OLTP applications (that are function oriented), the Operational Data Store contains subject-oriented, volatile, and current enterprise-wide detailed information; it serves as a system of record that provides comprehensive views of data in operational systems.

Data are transformed and integrated into a consistent, unified whole as they are obtained from legacy and other operational systems to provide business users with an integrated and current view of operations (see Figure 1–5). Data in the Operational Data Store are constantly refreshed so that the resulting image reflects the latest state of operations.



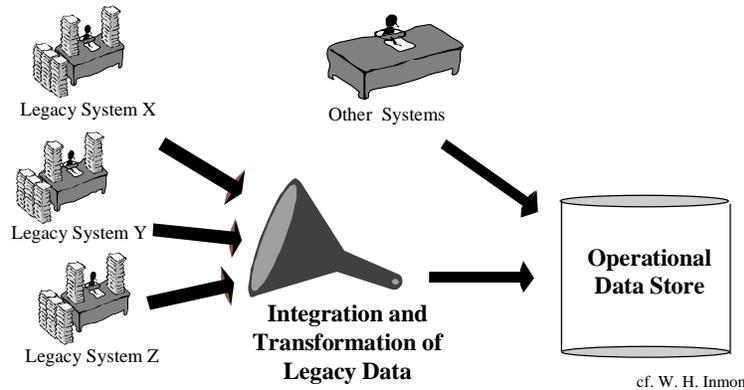


Figure 1-5 Legacy Systems and the Operational Data Store

Flash Monitoring and Reporting

These tools provide business users with a dashboard—meaningful online information on the operational status of the enterprise by making use of the data in the Operational Data Store. The business user obtains a constantly refreshed, enterprise-wide view of operations without creating unwanted interruptions or additional load on transaction processing systems.

Workflow Management and Groupware

Workflow management systems are tools that allow groups to communicate and coordinate their work. Early incarnations of this technology supported group scheduling, e-mail, online discussions, and resource sharing. More advanced implementations of this technology are integrated with OLTP applications to support the execution of business processes.

Decisional Needs

Data Warehouse

The data warehouse concept developed as IT professionals increasingly realized that the structure of data required for transaction reporting was significantly different from the structure required to analyze data.



The data warehouse was originally envisioned as a separate architectural component that converted and integrated masses of raw data from legacy and other operational systems and from external sources. It was designed to contain summarized, historical views of data in production systems. This collection provides business users and decision-makers with a cross-functional, integrated, subject-oriented view of the enterprise.

The introduction of the Operational Data Store has now caused the data warehouse concept to evolve further. The data warehouse now contains summarized, historical views of the data in the Operational Data Store. This is achieved by taking regular “snapshots” of the contents of the Operational Data Store and using these snapshots as the basis for warehouse loads.

In so doing, the enterprise obtains the information required for long-term and historical analysis, decision-making, and planning.

Decision Support Applications

Also known as OLAP (Online Analytical Processing), these applications provide managerial users with meaningful views of past and present enterprise data. User-friendly formats, such as graphs and charts are frequently employed to quickly convey meaningful data relationships.

Decision support processing typically does not involve the update of data; however, some OLAP software allows users to enter data for budgeting, forecasting, and “what-if” analysis.

Informational Needs

Informational Web Services and Scripts

Web browsers provide their users with a universal tool or front-end for accessing information from web servers. They provide users with a

new ability to both explore and publish information with relative ease. Unlike other technologies, web technology makes any user an instant publisher by enabling the distribution of knowledge and expertise, with no more effort than it takes to record the information in the first place.



By its very nature, this technology supports a paperless distribution process. Maintenance and update of information is straightforward since the information is stored on the web server.

Virtual Corporation Needs

Transactional Web Services and Scripts

Several factors now make Internet technology and electronic commerce a realistic option for enterprises that wish to use the Internet for business transactions.



- **Cost:** The increasing affordability of Internet access allows businesses to establish cost-effective and strategic links with business partners. This option was originally open only to large enterprises through expensive, dedicated wide-area networks or metropolitan area networks.
- **Security.** Improved security and encryption for sensitive data now provide customers with the confidence to transact over the Internet. At the same time, improvements in security provide the

enterprise with the confidence to link corporate computing environments to the Internet.

- **User-friendliness.** Improved user-friendliness and navigability from web technology make Internet technology and its use within the enterprise increasingly popular.

Figure 1–6 recapitulates the architectural components for the different types of business needs.

The majority of the architectural components support the enterprise at the operational level. However, separate components are now clearly defined for decisional and information purposes, and the virtual corporation becomes possible through Internet technologies.

Other Components

Other architectural components are so pervasive that most enterprises have begun to take their presence for granted. One example is the group of applications collectively known as office productivity tools (such as Microsoft Office or Lotus SmartSuite). Components of this type can and should be used across the various layers of the Enterprise Architecture and, therefore, are not described here as a separate item.

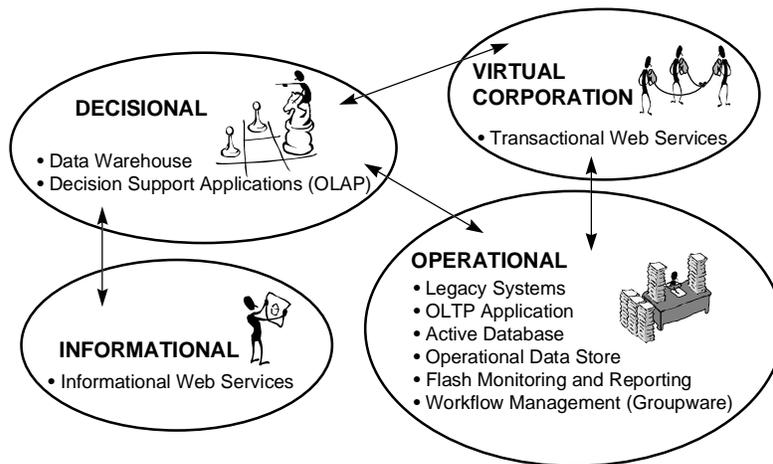
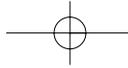


Figure 1–6 InfoMotion Enterprise Architecture Components (Applicability to Business Needs)



► Architecture Migration Scenarios

Given the typical path that most Enterprise Architectures have followed, an enterprise will find itself in need of one or more of the following six migration scenarios. We offer recommendations for fulfilling those needs.

Legacy Integration

The Need

The integration of new and legacy systems is a constant challenge because of the architectural templates upon which legacy systems were built. Legacy systems often attempt to meet all types of information requirements through a single architectural component; consequently, these systems are brittle and resistant to evolution.

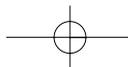
Despite attempts to replace them with new applications, many legacy systems remain in use because they continue to meet a set of business requirements: they represent significant investments that the enterprise cannot afford to scrap, or their massive replacement would result in unacceptable levels of disruption to business operations.

The Recommended Approach

The integration of legacy systems with the rest of the architecture is best achieved through the Operational Data Store and/or the data warehouse. Figure 1-7 modifies Figure 1-5 to show the integration of legacy systems.

Legacy programs that produce and maintain summary information are migrated to the data warehouse. Historical data are likewise migrated to the data warehouse. Reporting functionality in legacy systems is moved either to the flash reporting and monitoring tools (for operational concerns), or to decision support applications (for long-term planning and decision-making). Data required for operational monitoring are moved to the Operational Data Store. Table 1-1 summarizes the migration avenues.

The Operational Data Store and the data warehouse present IT professionals with a natural migration path for legacy migration. By migrating legacy systems to these two components, enterprises can gain a



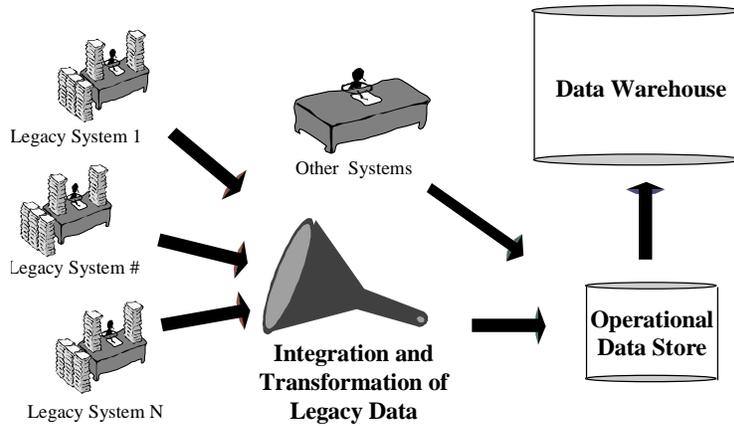


Figure 1-7 Legacy Integration

measure of independence from legacy components that were designed with old, possibly obsolete, technology. Figure 1-8 highlights how this approach fits into the Enterprise Architecture.

Operational Monitoring

The Need

Today's typical legacy systems are not suitable for supporting the operational monitoring needs of an enterprise. Legacy systems are typically structured around functional or organizational areas, in contrast

Table 1-1 Migration of Legacy Functionality to the Appropriate Architectural Component

Functionality in Legacy Systems	Should be Migrated to . . .
Summary Information	Data Warehouse
Historical Data	Data Warehouse
Operational Reporting	Flash Monitoring and Reporting Tools
Data for Operational Monitoring	Operational Data Store
Decisional Reporting	Decision Support Applications

INFOMOTION

LEGACY INTEGRATION

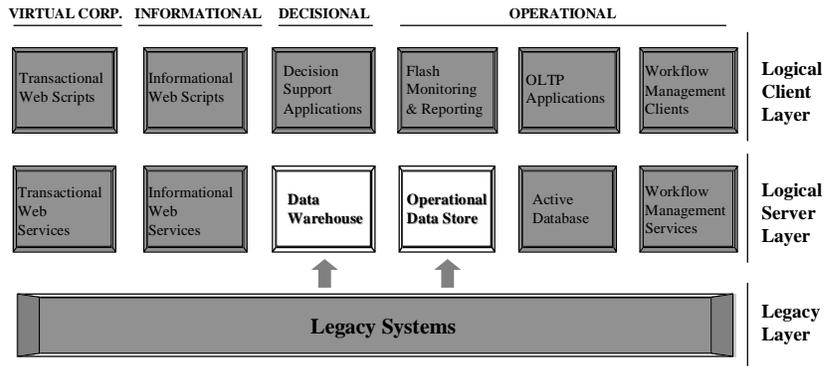


Figure 1-8 Legacy Integration: Architectural View

to the cross-functional view required by operations monitoring. Different and potentially incompatible technology platforms may have been used for different systems. Data may be available in legacy databases but are not extracted in the format required by business users. Or data may be available but may be too raw to be of use for operational decision-making (further summarization, calculation, or conversion is required). And lastly, several systems may contain data about the same item but may examine the data from different viewpoints or at different time frames, therefore requiring reconciliation.

The Recommended Approach

An integrated view of current, operational information is required for the successful monitoring of operations. Extending the functionality of legacy applications to meet this requirement would merely increase the enterprise's dependence on increasingly obsolete technology. Instead, an Operational Data Store, coupled with flash monitoring and reporting tools, as shown in Figure 1-9, meets this requirement without sacrificing architectural integrity.

Like a dashboard on a car, flash monitoring and reporting tools keep business users apprised of the latest cross-functional status of opera-

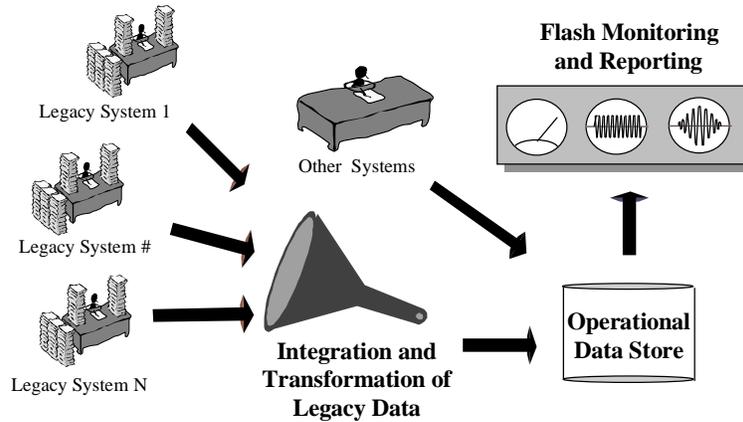


Figure 1-9 Operational Monitoring

tions. These tools obtain data from the Operational Data Store, which is regularly refreshed with the latest information from legacy and other operational systems.

Business users are consequently able to step in and correct problems in operations while they are still small or, better still, prevent problems from occurring altogether. Once alerted of a potential problem, the business user can manually intervene or make use of automated tools (i.e., control panel mechanisms) to fine-tune operational processes. Figure 1-10 highlights how this approach fits into the Enterprise Architecture.

Process Implementation

The Need

In the early 90s, the popularity of business process reengineering (BPR) caused businesses to focus on the implementation of new and redefined business processes.

Raymond Manganelli and Mark Klein, in their book *The Reengineering Handbook* (AMACOM, 1994, ISBN: 0-8144-0236-4) define BPR as “the rapid and radical redesign of *strategic, value-added business processes*—and the systems, policies, and organizational structures

INFOMOTION

OPERATIONAL MONITORING

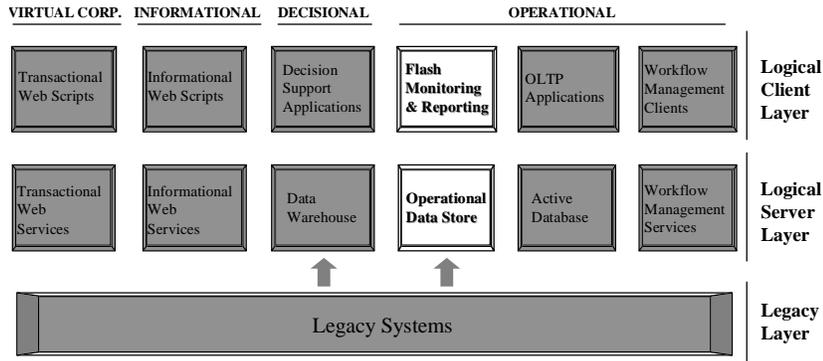


Figure 1-10 Operational Monitoring: Architectural View

that support them—to optimize the work flow and productivity in an organization.” Business processes are redesigned to achieve desired results in an optimum manner.

The Recommended Approach

With BPR, the role of Information Technology shifted from simple automation to enabling radically redesigned processes. Client/server technology, such as OLTP applications serviced by active databases, is particularly suited to supporting this type of business need. Technology advances have made it possible to build and modify systems quickly in response to changes in business processes. New policies, procedures, and controls are supported and enforced by the systems.

In addition, workflow management systems can be used to supplement OLTP applications. A workflow management system converts business activities into a goal-directed process that flows through the enterprise in an orderly fashion (see Figure 1-11). The workflow management system alerts users through the automatic generation of notification messages or reminders and routes work so that the desired business result is achieved in an expedited manner.

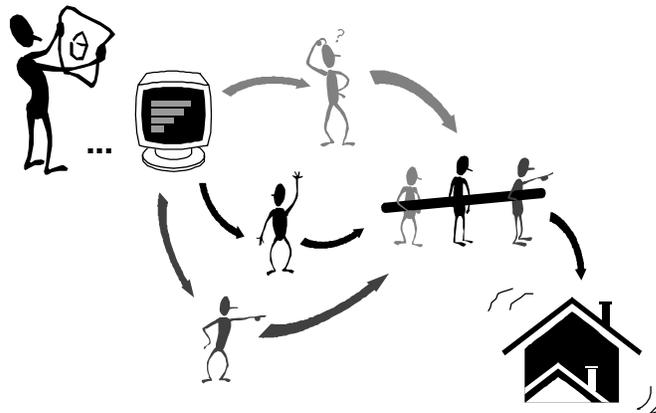


Figure 1-11 Process Implementation

Figure 1-12 highlights how this approach fits into the Enterprise Architecture.

Decision Support

The Need

It is not possible to anticipate the information requirements of decision-makers for the simple reason that their needs depend on the business situation that they face. Decision-makers need to review enterprise data from different dimensions and at different levels of detail to find the source of a business problem before they can attack it. They likewise need information for detecting business opportunities to exploit.

Decision-makers also need to analyze trends in the performance of the enterprise. Rather than waiting for problems to present themselves, decision-makers need to proactively mobilize the resources of the enterprise in anticipation of a business situation.

Since these information requirements cannot be anticipated, the decision-maker often resorts to reviewing predesigned inquiries or reports in an attempt to find or derive needed information. Alternatively, the IT professional is pressured to produce an ad hoc report from legacy sys-

INFOMOTION

PROCESS IMPLEMENTATION

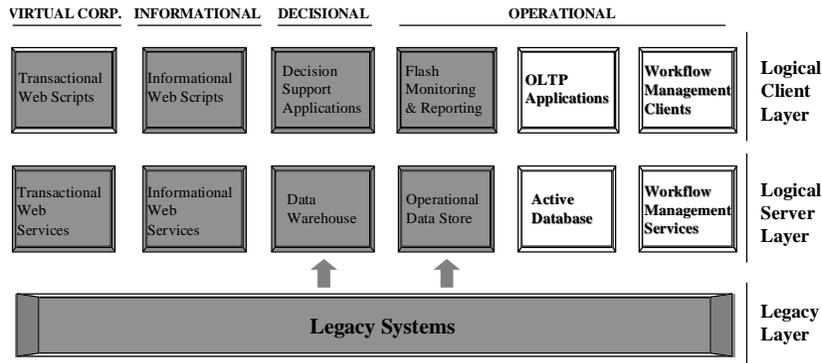


Figure 1-12 Process Implementation: Architectural View

tems as quickly as possible. If unlucky, the IT professional will find the data needed for the report are scattered throughout different legacy systems. An even unluckier may find that the processing required to produce the report will have a toll on the operations of the enterprise.

These delays are not only frustrating both for the decision-maker and the IT professional, they are dangerous for the enterprise. The information that eventually reaches the decision-maker may be inconsistent, inaccurate, or worse, obsolete.

The Recommended Approach

Decision support applications (or OLAP) that obtain data from the data warehouse are recommended for this particular need. The data warehouse holds transformed and integrated enterprise-wide operational data appropriate for strategic decision-making, as shown in Figure 1-13. The data warehouse also contains data obtained from external sources, whenever this data is relevant to decision-making.

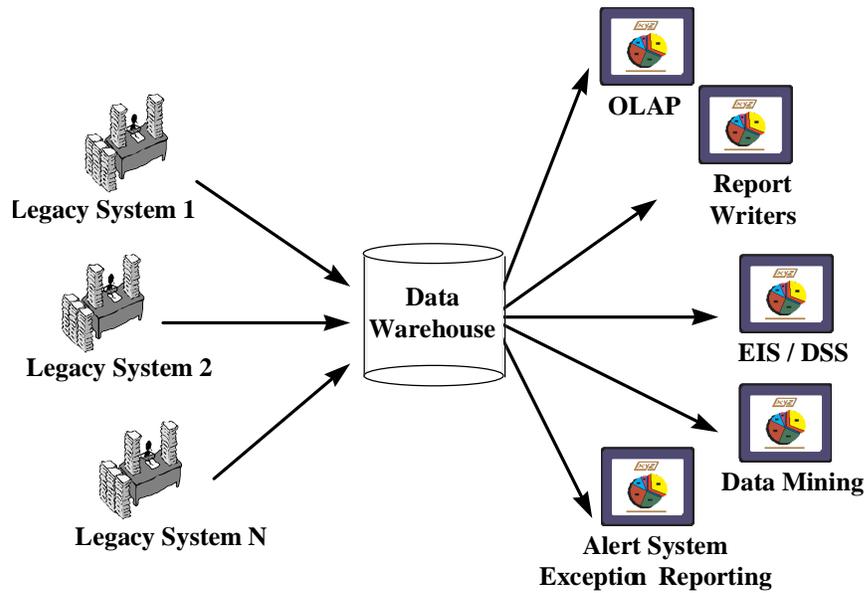


Figure 1-13 Decision Support

Decision support applications analyze and make data warehouse information available in formats that are readily understandable by decision-makers. Figure 1-14 highlights how this approach fits into the Enterprise Architecture.

Hyperdata Distribution

The Need

Past informational requirements were met by making data available in physical form through reports, memos, and company manuals. This practice resulted in an overflow of documents providing much data and not enough information.

Paper-based documents also have the disadvantage of becoming dated. Enterprises encountered problems in keeping different versions of related items synchronized. There was a constant need to update, republish, and redistribute documents.

INFOMOTION DECISION SUPPORT

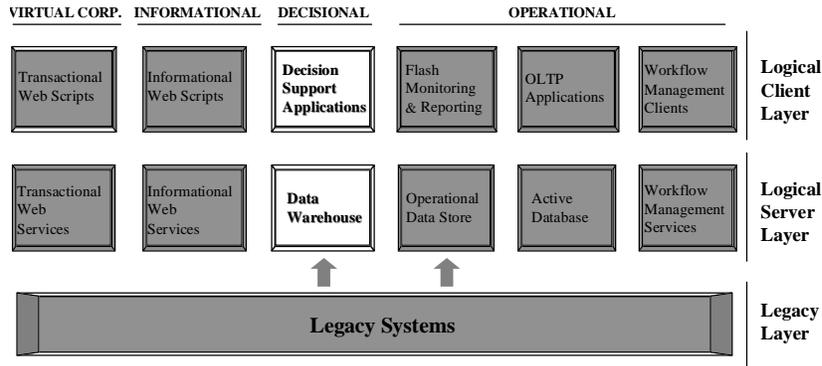


Figure 1-14 Decision Support: Architectural View

In response to this problem, enterprises made data available to users over a network to eliminate the paper. It was hoped that users could selectively view the data that they needed whenever they needed it. This approach likewise proved to be insufficient because users still had to navigate through a sea of data to locate the specific item of information that was needed.

The Recommended Approach

Users need the ability to browse through nonlinear presentations of data. Web technology is particularly suitable to this need because of its extremely flexible and highly visual method of organizing information (see Figure 1-15).

Web technology allows users to display charts and figures; navigate through large amounts of data; visualize the contents of database files; seamlessly navigate across charts, data, and annotation; and organize charts and figures in a hierarchical manner. Users are therefore able to locate information with relative ease.

Figure 1-16 highlights how this approach fits into the Enterprise Architecture.

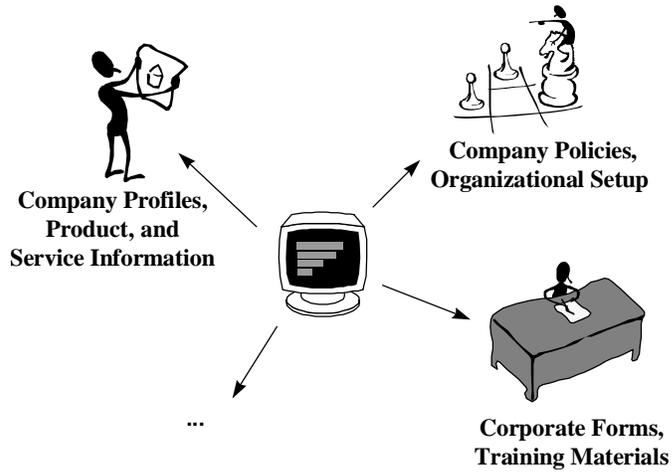


Figure 1-15 Hyperdata Distribution

INFOMOTION HYPERDATA DISTRIBUTION

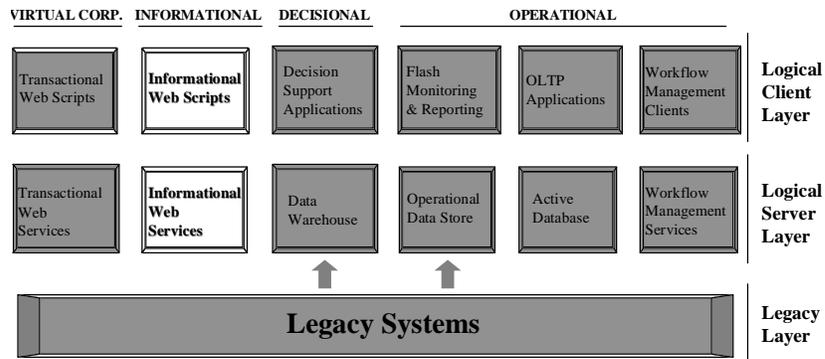


Figure 1-16 Hyperdata Distribution: Architectural View

Virtual Corporation

The Need

A virtual corporation is an enterprise that has extended its business processes to encompass both its key customers and suppliers. Its business processes are newly redesigned; its product development or service delivery are accelerated to better meet customer needs and preferences; its management practices promote new alignments between management and labor, as well as new linkages between enterprise, supplier and customer. A new level of cooperation and openness is created and encouraged between the enterprise and its key business partners.

The Recommended Approach

Partnerships at the enterprise level translate into technological links between the enterprise and its key suppliers or customers (see Figure 1–17). Information required by each party is identified, and steps are taken to ensure that this data crosses organizational boundaries properly.

Some organizations seek to establish a higher level of cooperation with their key business partners by jointly redesigning their business processes to provide greater value to the customer.

Internet and web technologies are well suited to support redesigned, transactional processes, thanks to decreasing Internet costs, improved

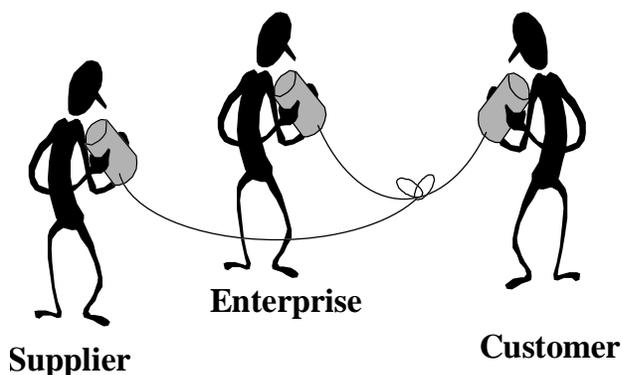


Figure 1–17 Virtual Corporation

INFOMOTION VIRTUAL CORPORATION

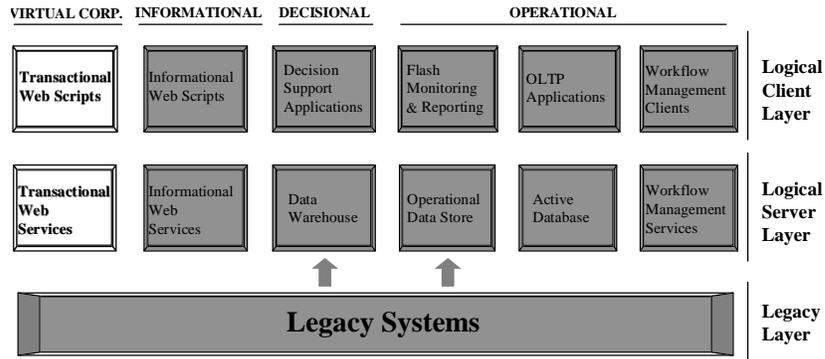


Figure 1-18 Virtual Corporation: Architectural View

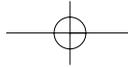
security measures, improved user-friendliness, and navigability. Figure 1-18 highlights how this approach fits into the Enterprise Architecture.

► Migration Strategy: How Do We Move Forward?

The strategies presented in the previous section enable organizations to move from their current technology architectures into the *InfoMotion* Enterprise Architecture. This section describes the tasks for any migration effort.

Review the Current Enterprise Architecture

As simplistic as this may sound, the starting point is a review of the current Enterprise Architecture. It is important to have an idea of what we have now before we can plan for where we want to be.



The IT department or division should have this information readily available, although it may not necessarily be expressed in terms of the architectural components identified above. A short and simple exercise of mapping the current architecture of your enterprise to the architecture described above should quickly highlight any gaps in the current architecture.

Identify Information Architecture Requirements

Knowing that the Enterprise IT Architecture has gaps is not sufficient. It is important to know whether these can be considered real gaps when viewed within the context of the enterprise's requirements. Gaps should cause concern only if the absence of an architectural component prevents the IT infrastructure from meeting present requirements or from supporting long-term strategies.

For example, if transactional web scripts are not critical to your enterprise given its current needs and strategies, there should be no cause for concern.

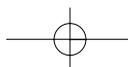
Develop a Migration Plan Based on Requirements

It is not advisable for an enterprise to use this list of architectural gaps to justify a dramatic overhaul of its IT infrastructure; such an undertaking would be expensive and would cause unnecessary disruption of business operations.

Instead, the enterprise would do well to develop a migration plan that consciously maps coming IT projects to the *InfoMotion* Enterprise Architecture.

The Natural Migration Path

While developing the migration plan, the enterprise should consider the natural migration path that the *InfoMotion* architecture implies, as illustrated in Figure 1-19.



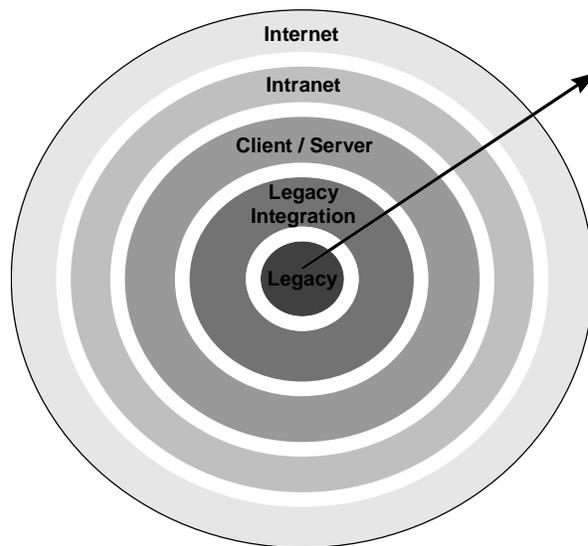


Figure 1-19 Natural Migration Roadmap

- At the very core of the Enterprise Architecture is the Legacy layer. For most companies, this core layer is where the majority of technology investments have been made. It should also be the starting point of any architecture migration effort, i.e., the enterprise should start from this core technology before focusing its attention on newer forms or layers of technology.
- The Legacy Integration layer insulates the rest of the Enterprise Architecture from the growing obsolescence of the Legacy layer. It also provides the succeeding technology layers with a more stable foundation for future evolution.
- Each of the succeeding technology layers (i.e., Client/Server, Intranet, Internet) builds upon its predecessors.
- At the outermost layer, the public Internet infrastructure itself supports the operations of the enterprise.

The Customized Migration Path

Depending on the priorities and needs of the enterprise, one or more of the migration scenarios described in the previous section will be help-

ful starting points. The scenarios provide generic roadmaps that address typical architectural needs.

The migration plan, however, must be customized to address the specific needs of the enterprise. Each project defined in the plan must individually contribute to the enterprise in the short term, while laying the groundwork for achieving long-term enterprise and IT objectives.

By incrementally migrating its IT infrastructure (one component and one project at a time), the enterprise will find itself slowly but surely moving towards a modern, resilient Enterprise Architecture, with minimal and acceptable levels of disruption in operations.

Monitor and Update the Migration Plan

The migration plan must be monitored, and the progress of the different projects fed back into the planning task. One must not lose sight of the fact that a modern Enterprise Architecture is a moving target; new technology renders continuous evolution of the Enterprise Architecture inevitable.

► In Summary

An enterprise has longevity in the business arena only when its products and services are perceived by its customers to be of value.

Likewise, Information Technology has value in an enterprise only when its cost is outweighed by its ability to increase and guarantee quality, improve service, cut costs or reduce cycle time, as depicted in Figure 1–20.

The Enterprise Architecture is the foundation for all Information Technology efforts. It therefore must provide the enterprise with the ability to:

$$\text{Value} = \frac{\text{Quality x Service}}{\text{Cost x Cycle Time}}$$

Figure 1–20 The Value Equation

- distill information of value from the data which surrounds it, which it continuously generates (information/data); and
- get that information to the right people and processes at the right time (motion).

These requirements form the basis for the *InfoMotion* equation, shown in Figure 1-21.

$$\text{InfoMotion} = \frac{\text{Information}}{\text{Data}} \times \text{Motion}$$

Figure 1-21 The *InfoMotion* Equation

By identifying distinct architectural components and their interrelationships, the *InfoMotion* Enterprise Architecture increases the capability of the IT infrastructure to meet present business requirements while positioning the enterprise to leverage emerging trends, such as data warehousing, in both business and technology. Figure 1-22 reprises the *InfoMotion* Enterprise Architecture, the elements of which we have discussed.

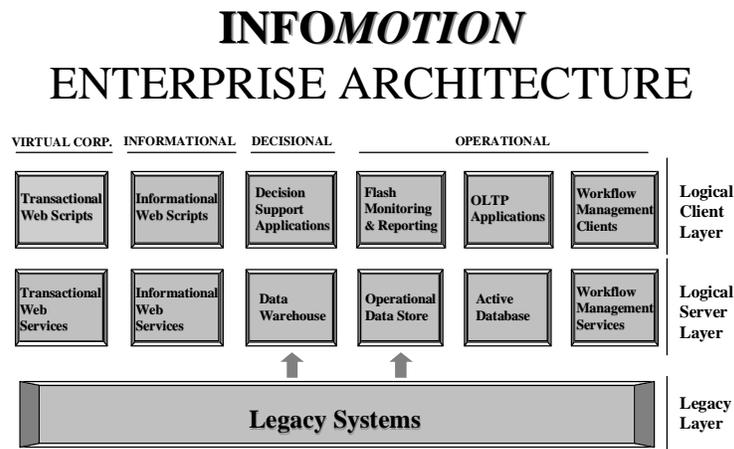


Figure 1-22 The *InfoMotion* Architecture