

PLUG-IN
B4

Enterprise Architectures

LEARNING OUTCOMES

1. Explain the three components of an enterprise architecture.
2. Describe how an organization can implement a solid information architecture.
3. List and describe the five qualities of an infrastructure architecture.
4. Compare web services and open systems.

Enterprise Architectures

A 66-hour failure of an FBI database that performed background checks on gun buyers was long enough to allow criminals to buy guns. The database failed at 1:00 p.m. on a Thursday and was not restored until 7:30 a.m. Sunday. The FBI must complete a gun check within three days; if it fails to do so, a merchant is free to make the sale. During this outage, any gun checks that were in progress were not finished, allowing merchants to complete those gun sales at their own discretion.

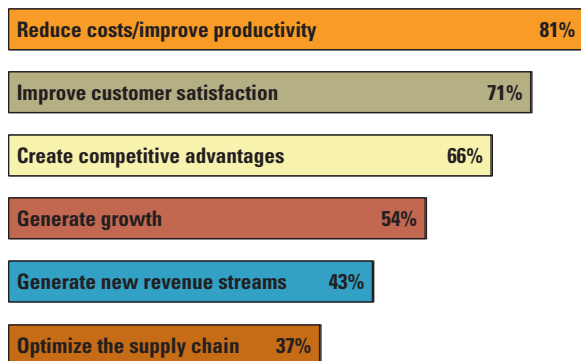


FIGURE B4.1

Primary Business Goals of Enterprise Architectures

To support the volume and complexity of today's user and application requirements, information technology needs to take a fresh approach to enterprise architectures by constructing smarter, more flexible environments that protect from system failures and crashes. **Enterprise architectures** include the plans for how an organization will build, deploy, use, and share its data, processes, and IT assets. A unified enterprise architecture will standardize enterprisewide hardware and software systems, with tighter links to the business strategy. A solid enterprise architecture can decrease costs, increase standardization, promote reuse of IT assets, and speed development

of new systems. The end result is that the right enterprise architecture can make IT cheaper, strategic, and more responsive. The primary business goals of enterprise architectures are displayed in Figure B4.1.

Enterprise architectures are never static; they continually change. Organizations use enterprise architects to help manage change. An **enterprise architect (EA)** is a person grounded in technology, fluent in business, a patient diplomat, who provides the important bridge between IT and the business. An EA is expensive and

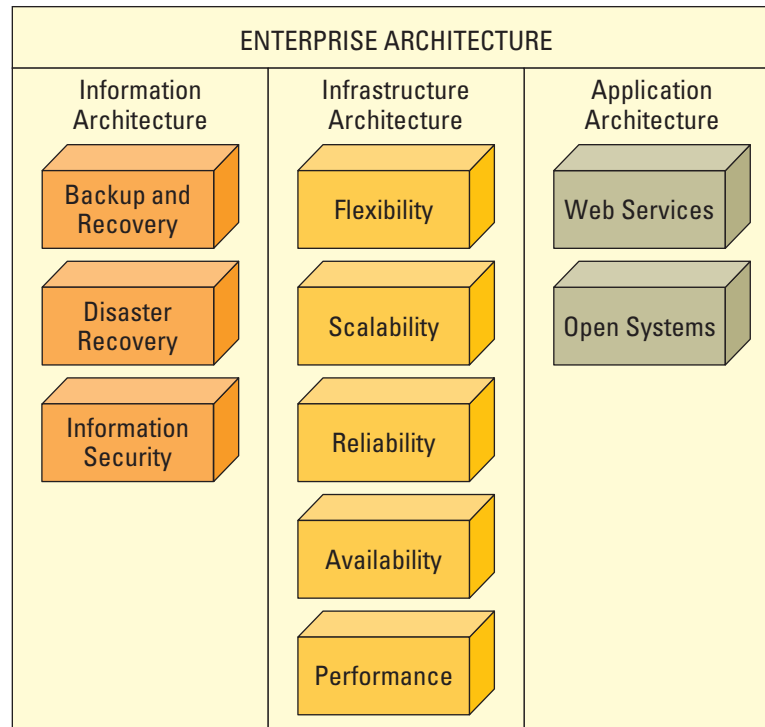


FIGURE B4.2
Three Components of Enterprise Architecture

generally receives a salary upward of \$150,000 per year. T-Mobile International’s enterprise architects review projects to ensure they are soundly designed, meet the business objectives, and fit in with the overall enterprise architecture. One T-Mobile project was to create software that would let subscribers customize the ring sounds on their cell phones. The project group assumed it would have to create most of the software from scratch. However, T-Mobile’s EAs found software already written elsewhere at T-Mobile that could be reused to create the new application. The reuse reduced the development cycle time by eight months, and the new application was available in less than six weeks.

Companies that have created solid enterprise architectures, such as T-Mobile, are reaping huge rewards in savings, flexibility, and business alignment. Basic enterprise architectures contain three components (see Figure B4.2).

1. **Information architecture** identifies where and how important information, like customer records, is maintained and secured.
2. **Infrastructure architecture** includes the hardware, software, and telecommunications equipment that, when combined, provide the underlying foundation to support the organization’s goals.
3. **Application architecture** determines how applications integrate and relate to each other.

Information Architecture



Information architecture identifies where and how important information, like customer records, is maintained and secured. A single backup or restore failure can cost an organization more than time and money; some data cannot be re-created,

and the business intelligence lost from that data can be tremendous. Chief information officers should have enough confidence that they could walk around and randomly pull out cables to prove that the systems are safe. The CIO should also be secure enough to perform this test during peak business hours. If the thought of this test makes the CIO cringe, then the organization’s customers should be cringing also. Figure B4.3 depicts the three primary areas an enterprise information architecture should focus on:

1. Backup and recovery
2. Disaster recovery
3. Information security

BACKUP AND RECOVERY


Each year businesses lose time and money because of system crashes and failures. One way to minimize the damage of a system crash is to have a backup and recovery strategy in place. A **backup** is an exact copy of a system’s information. **Recovery** is the ability to get a system up and running in the event of a system crash or failure and includes restoring the information backup. Many different types of backup and recovery media are available, including redundant storage servers, tapes, disks, and even CDs and DVDs. All the different types of backup and recovery media are reliable; their primary differences are the speed and associated costs.

A chain of more than 4,000 franchise locations, 7-Eleven Taiwan uploads backup and recovery information from its central location to all its chain locations daily. The company implemented a new technology solution by Digital Fountain that could quickly and reliably download and upload backup and recovery information to all its stores. In addition, when a connection fails during the download or upload, the technology automatically resumes the download without having to start over, saving valuable time.

Organizations should choose a backup and recovery strategy that is in line with business goals. If the organization deals with large volumes of critical information, it will require daily backups, perhaps even hourly backups, to storage servers. If the organization deals with small amounts of noncritical information, then it might require only weekly backups to tapes, CDs, or DVDs. Deciding how often to back up information and what media to use is a critical business decision. If an organization decides to back up on a weekly basis, then it is taking the risk that, if a total system crash occurs, it could lose a week’s worth of work. If this risk is acceptable, then a weekly backup strategy will work. If this risk is unacceptable, then the organization needs to move to a daily backup strategy. Some organizations find the risk of losing a day’s worth of work too high and move to an hourly backup strategy.

Two techniques used to help in case of system failure are fault tolerance and failover. **Fault tolerance** is a computer system designed that in the event a component fails, a backup component or procedure can immediately take its place with no loss of service. Fault tolerance can be provided with software, or embedded

FIGURE B4.3
The Three Areas Enterprise Information Architecture Should Focus On



in hardware, or provided by some combination. **Failover** is a backup operational mode in which the functions of a computer component (such as a processor, server, network, or database) are assumed by secondary system components when the primary component becomes unavailable through either failure or scheduled down time. A failover procedure involves automatically offloading tasks to a standby system component so that the procedure is as seamless as possible to the end user. Used to make systems more fault tolerant, failover is typically an integral part of mission-critical systems that must be constantly available.

DISASTER RECOVERY

A northern Ohio power company, FirstEnergy, missed signs that there were potential problems in its portion of North America’s electrical grid. The events that followed left an estimated 50 million people in the Northeast and Canada in the dark. The failings are laid out in the widely reported findings of a joint U.S./Canada task force that investigated the causes of the blackout and recommended what to do to avoid big-scale outages in the future. The report detailed many procedures or best practices including:

- Mind the enterprise architectures.
- Monitor the quality of computer networks that provide data on power suppliers and demand.
- Make sure the networks can be restored quickly in the case of downtime.
- Set up disaster recovery plans.
- Provide adequate staff training, including verbal communication protocols “so that operators are aware of any IT-related problems that may be affecting their situational awareness of the power grid.”

Disasters such as power outages, floods, and even harmful hacking strike businesses every day. Organizations must develop a disaster recovery plan to prepare for such occurrences. A **disaster recovery plan** is a detailed process for recovering information or an IT system in the event of a catastrophic disaster such as a fire or flood. Spending on disaster recovery is rising worldwide among financial institutions (see Figure B4.4).

A comprehensive disaster recovery plan takes into consideration the location of the backup information. Many organizations store backup information in an off-site facility. StorageTek specializes in providing off-site information storage and disaster recovery solutions. A comprehensive disaster recovery plan also foresees the possibility that not only the computer equipment but also the building where employees work may be destroyed. A **hot site** is a separate and fully equipped facility where the company can move immediately after a disaster and resume business. A **cold site** is a separate facility that does not have any computer equipment, but is a place where employees can move after a disaster.

A **disaster recovery cost curve** charts (1) the cost to the organization of the unavailability of information and technology and (2) the cost to the organization of recovering from a disaster over time. Figure B4.5 displays a disaster recovery cost curve and shows that where the two lines intersect is the best recovery plan in terms of cost and time. Creating an organization’s disaster recovery cost curve is no small task. It must consider the cost of losing information and technology within each department or functional area, and the cost of losing information and technology across the whole enterprise. During the first few hours of a disaster, those costs will be low but become increasingly higher over time. With those costs in hand, an organization must then determine the costs of recovery. Cost of recovery during the first few hours of a disaster is exceedingly high and diminishes over time.

FIGURE B4.4
Financial Institutions
Worldwide Spending on
Disaster Recovery


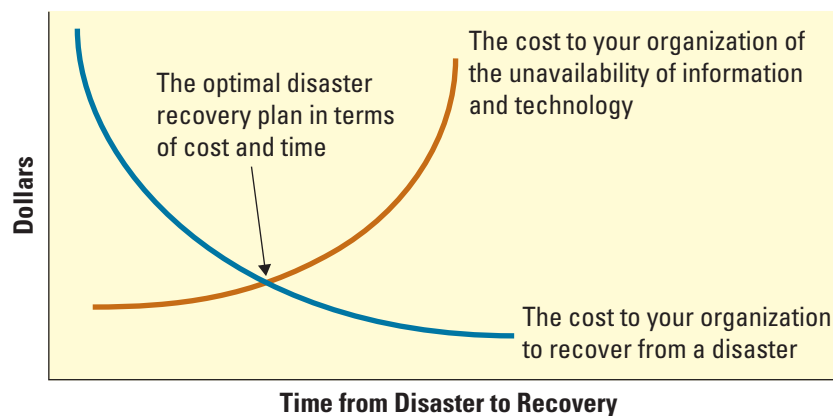


FIGURE B4.5

The Disaster Recovery Cost Curve



On April 18, 1906, San Francisco was rocked by an earthquake that destroyed large swathes of the city and claimed the lives of more than 3,000 inhabitants of the Bay Area. More than a century later, a bigger, bolder, rebuilt, and more resilient San Francisco is more important than ever. Now it serves as the heart of the global IT industry and a major world financial center. However, San Francisco remains well aware of the terrible potential danger that exists along the San Andreas fault.

The vast skyscrapers downtown may now be built to withstand huge pressures, but what about the infrastructure and the systems that keep modern business ticking—and the people who must be able to access them? **Business continuity planning (BCP)** is planning for how an organization will recover and restore partially or completely interrupted critical function(s) within a predetermined time after a disaster or extended disruption. Business continuity and disaster recovery are serious issues for all organizations in the Bay Area, including the Union Bank of California, which is based in the heart of downtown San Francisco.

Barry Cardoza, head of business continuity planning and disaster recovery at Union Bank of California, said, “You have disasters that you can see coming and you’ve got disasters that you can’t see coming and an earthquake is an example of [the latter]. And you don’t know how bad it’s going to be until it hits.”

Clearly, the bank must have processes in place ahead of such an event to mitigate the threat. Simply reacting is not a strategy. The continuity department must also understand every aspect of the business and weigh downtime for each in terms of financial and reputational damage. Union Bank of California has created a disaster recovery plan that includes multiple data centers in diverse locations, mirrored sites that can take over at the flick of a switch, hot sites where staff can walk in and start working exactly as they would if they were in their normal location, and a vast amount of redundancy. In addition, the bank has created real-time mirroring between data centers. It is now a matter of minutes not hours for Union Bank of California to be up and running in the event of a disaster.

INFORMATION SECURITY

Security professionals are under increasing pressure to do the job right and cost-effectively as networks extend beyond organizations to remote users, partners, and customers, and to cell phones, PDAs, and other mobile devices. Regulatory requirements to safeguard data have increased. Concerns about identity theft are at an all-time high. Hacking and other unauthorized access contribute to the approximately 10 million instances of identity theft each year, according to the Federal Trade Commission. A good information architecture includes a strong information security plan, along with managing user access and up-to-date antivirus software and patches.

Managing User Access

Managing user access to information is a critical piece of the information architecture. Passwords may still be the weakest link in the security chain. At Vitas Healthcare Corporation, with a workforce of 6,000 and operations across 15 states, authorized employees enter as many as a half-dozen passwords a day to access multiple systems. While it is important to maintain password discipline to secure customers' health care data, maintaining and managing the situation creates a drag on the IT department. "Our help desk spends 30 percent of their time on password management and provisioning," says John Sandbrook, senior IT director.

The company began using Fischer International Corporation's Identity Management Suite to manage passwords and comply with data-access regulations such as the Sarbanes-Oxley Act. The ID-management product includes automated audit, reporting, and compliance capabilities, plus a common platform for password management, provisioning, and self-service. With the software, Vitas can enforce stronger passwords with seven, eight, or nine characters, numbers, and capital letters that frequently change. The company anticipates curbing help-desk password time by 50 percent.


Up-to-Date Antivirus Software and Patches

Security is a top priority for business managers, regardless of the size of their company. Among Fortune 500 companies, more than 80 percent of those surveyed described updating security procedures, tools, and services as a key business priority. The same holds true for all small, midsize, or large companies and all IT managers and corporate managers.

The main focus for most managers is preventing hackers, spammers, and other malcontents from entering their networks, and nearly two-thirds are looking to enhance their network-security-management, intrusion-detection, content-filtering, and antispy software. More than half also plan to upgrade their encryption software.

Microsoft issues patches for its software on the second Tuesday of every month. These patches must be downloaded and installed on all systems across the entire enterprise if the company wants to keep its systems protected. At OMD, a media buying and planning subsidiary of Omnicom Group Inc., the network administrator had to manually install critical patches on all 100 servers, taking more than a week to deploy the patch across the company. Now, OMD uses automated installation software for patches and upgrades. The company purchased Altiris Management Suite for Dell servers, which let it move ahead with applying patches without taking down entire systems and balancing patch-deployment timing among servers so that all departments were not down at once during a patch install. Given everything else that security professionals need to think about, automated installation software is a welcome relief.

Infrastructure Architecture



Gartner Inc. estimates that the typical web application goes down 170 hours per year. At Illinois-based online brokerage OptionsXpress, application performance problems can have a serious impact on livelihoods. Nearly 7,000 options traders visit the OptionsXpress website at any given time, completing nearly 20,000 transactions a day. With all this online traffic, the brokerage's IT administrators were always up against the clock when re-creating troublesome applications offline in the

development environment. The company struggled to unlock the mystery behind a troublesome trading application that was forcing traders to resubmit orders. Sometimes the application would just die and then restart itself for no apparent reason.

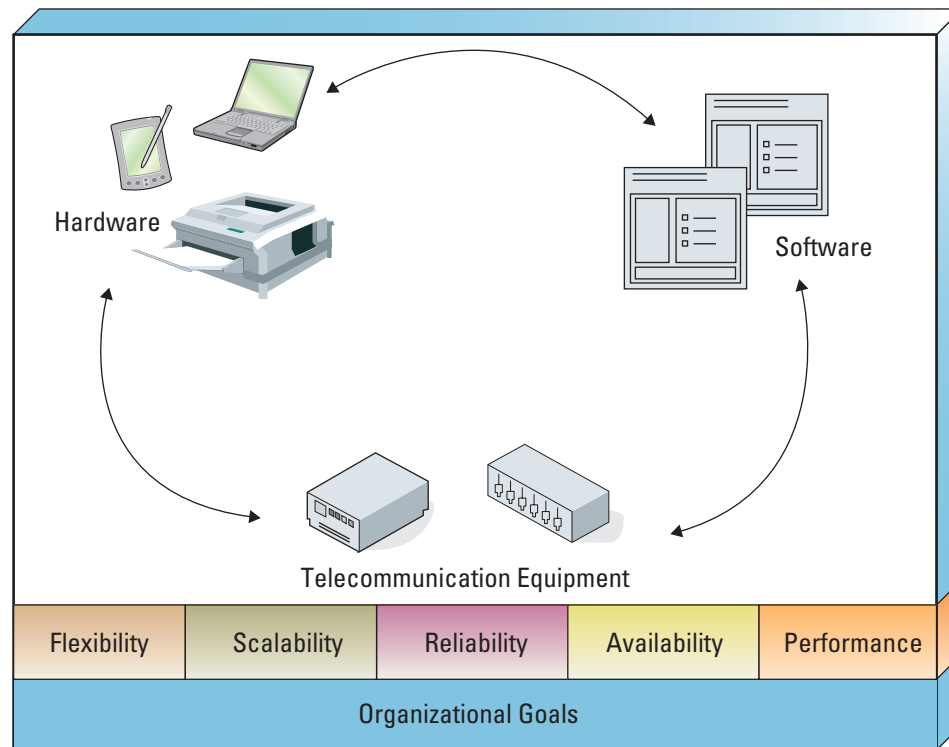
Infrastructure architecture includes the hardware, software, and telecommunications equipment that, when combined, provide the underlying foundation to support the organization’s goals (see Figure B4.6). As an organization changes, its systems must be able to change to support its operations. If an organization grows by 50 percent in a single year, its systems must be able to handle a 50 percent growth rate. Systems that cannot adapt to organizational changes can severely hinder the organization’s ability to operate. The future of an organization depends on its ability to meet its partners and customers on their terms, at their pace, any time of the day, in any geographic location. The following are the five primary characteristics of a solid infrastructure architecture:

1. Flexibility
2. Scalability
3. Reliability
4. Availability
5. Performance

FLEXIBILITY

Organizations must watch today’s business, as well as tomorrow’s, when designing and building systems. Systems must be flexible enough to meet all types of business changes. For example, a system might be designed to include the ability to handle multiple currencies and languages, even though the company is not currently performing business in other countries. When the company starts growing and performing business in new countries, the system will already have the flexibility to handle multiple currencies and languages. If the company failed to recognize that its business would someday be global, it would need to redesign all its systems to handle multiple currencies and languages, not easy once systems are up and running.

FIGURE B4.6
Infrastructure Architecture Characteristics



SCALABILITY

Estimating organizational growth is a challenging task. Growth can occur in a number of different forms including more customers and product lines and expansion into new markets. **Scalability** refers to how well a system can adapt to increased demands. A number of factors can create organizational growth including market, industry, and economy factors. If an organization grows faster than anticipated, it might experience all types of performance degradations, ranging from running out of disk space to a slowdown in transaction speeds. Anticipating expected—and unexpected—growth is key to building scalable systems that can support that growth.

MSNBC's website typically received moderate traffic. On September 11, 2001, the site was inundated with more than 91 million page views as its customers were trying to find out information about the terrorist attacks. Fortunately, MSNBC had anticipated this type of surging demand and built adaptable systems accordingly, allowing it to handle the increased page view requests.

Capacity planning determines the future IT infrastructure requirements for new equipment and additional network capacity. Performing a capacity plan is one way to ensure the IT infrastructure is scalable. It is cheaper for an organization to implement an IT infrastructure that considers capacity growth at the beginning of a system launch than to try to upgrade equipment and networks after the system has been implemented. Not having enough capacity leads to performance issues and hinders the ability of knowledge workers to perform their jobs. If 100 workers are using the Internet to perform their jobs and the company purchases bandwidth that is too small and the network capacity is too small, the workers will spend a great deal of time just waiting to get information from the Internet. Waiting for an Internet site to return information is not very productive.

Web 2.0 is driving demand for capacity planning. Delivering entertainment-grade video over the Internet poses significant challenges as service providers scale solutions to manage millions of users, withstand periods of peak demand, and deliver a superior quality of experience while balancing network capacity and efficient capital investment. Given the success of YouTube and the likelihood of similar video experiences, the bandwidth required to transport video services will continue to increase and the possibility of video degradation will become more challenging. Since video cannot tolerate packet loss (e.g., blocks of data lost), congestion due to overuse is not acceptable—admitting just one more stream to a network near peak capacity could degrade the video and broadcast quality for all users.

RELIABILITY

Reliability ensures all systems are functioning correctly and providing accurate information. Reliability is another term for accuracy when discussing the correctness of systems within the context of efficiency IT metrics. Inaccurate information processing occurs for many reasons, from the incorrect entry of data to information corruption. Unreliable information puts the organization at risk when making decisions based on the information.

AVAILABILITY

Availability (an efficiency IT metric) addresses when systems can be accessed by employees, customers, and partners. **High availability** refers to a system or component that is continuously operational for a desirably long length of time. Availability is typically measured relative to “100 percent operational” or “never failing.” A widely held but difficult-to-achieve standard of availability for a system or product is known as “five 9s” (99.999 percent) availability.

Some companies have systems available 24x7 to support business operations and global customer and employee needs. With the emergence of the web, companies expect systems to operate around the clock. A customer who finds that a website closes at 9:00 p.m. is not going to be a customer long.


Systems, however, must come down for maintenance, upgrades, and fixes. One challenge organizations face is determining when to schedule system downtime if the system is expected to operate continually. Exacerbating the negative impact of scheduled system downtime is the global nature of business. Scheduling maintenance during the evening might seem like a great idea, but the evening in one city is the morning somewhere else in the world, and global employees may not be able to perform their jobs if the system is down. Many organizations overcome this problem by having redundant systems, allowing the organization to take one system down by switching over to a redundant, or duplicate, system.

PERFORMANCE

Performance measures how quickly a system performs a certain process or transaction (in terms of efficiency IT metrics of both speed and throughput). Not having enough performance capacity can have a devastating, negative impact on a business. A customer will wait only a few seconds for a website to return a request before giving up and moving on to another website. To ensure adaptable systems performance, capacity planning helps an organization determine future IT infrastructure requirements for new equipment and additional network capacity. It is cheaper for an organization to design and implement an IT infrastructure that envisions performance capacity growth than to update all the equipment after the system is already operational.

Abercrombie & Fitch (A&F) uses the Internet to market its distinctive image of being a fashion trendsetter to one of its largest customer segments, college students. The company designed its enterprise architecture with the help of IBM, which ensured www.bercrombie.com paralleled the same sleek but simple design of *A&F Quarterly*, the company's flagship magazine. Abercrombie & Fitch knew that its website had to be accessible, available, reliable, and scalable to meet the demands of its young customers. Young customers tend to be Internet savvy, and their purchasing habits vary from customers who only shop for sale items at midnight to customers who know exactly what they want immediately. The highly successful website gives customers not only an opportunity to shop online, but also a taste of the Abercrombie & Fitch lifestyle through downloadable MP3s, calendars, and desktop accessories.

Application Architecture



Gartner Inc. research indicates that application problems are the single largest source of downtime, causing 40 percent of annual downtime hours and 32 percent of average downtime costs. **Application architecture**

determines how applications integrate and relate to each other. Advances in integration technology—primarily web services and open systems—are providing new ways for designing more agile, more responsive enterprise architectures that provide the kind of value businesses need. With these new architectures, IT can build new business capabilities faster, cheaper, and in a vocabulary the business can understand.

WEB SERVICES

Web services promise to be the next major frontier in computing. **Web services** contain a repertoire of web-based data and procedural resources that use shared protocols and standards permitting different applications to share data and services. The major application of web services is the integration among different applications (refer to Figure B4.7). Before web services, organizations had trouble with

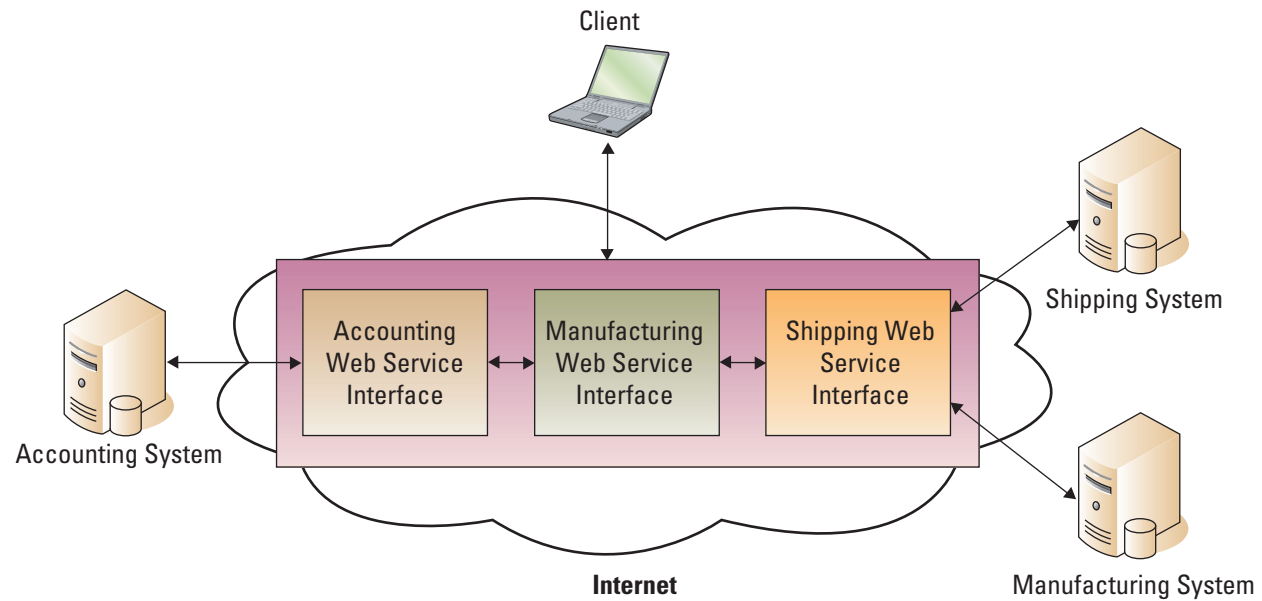


FIGURE B4.7
Web Service Architecture

interoperability. **Interoperability** is the capability of two or more computer systems to share data and resources, even though they are made by different manufacturers. If a supply chain management (SCM) system can talk to (share information with) a customer relationship management (CRM) system, interoperability exists between the two systems. The traditional way that organizations achieved interoperability was to build integrations. Now, an organization can use web services to perform the same task.

Verizon's massive enterprise architecture includes three different companies GTE, Bell Atlantic, and Nynex, each with its own complex systems. To find a customer record in any of the three companies' systems, Verizon turns to its search engine, called Spider. Spider is Verizon's version of Google, and it's helping Verizon's business to thrive.

Spider contains a vital customer information web service that encapsulates Verizon's business rules, which help it to access the correct data repository when looking for customer information. Whenever a new system is built that needs to link to customer information, all the developer has to do is reuse the web service that will link to the customer records. Because Verizon has the web service in place as part of its enterprise architecture, development teams can build new applications within a month, as opposed to six months.

Web services encompass all the technologies that are used to transmit and process information on and across a network, most specifically the Internet. It is easiest to think of an individual web service as software that performs a specific task, with that task being made available to any user who needs its service. For example, a "Deposit" web service for a banking system might allow customers to perform the task of depositing money to their accounts. The web service could be used by a bank teller, by the customer at an ATM, and/or by the customer performing an online transaction through a web browser.

The "Deposit" web service demonstrates one of the great advantages of using the web service model to develop applications. Developers do not have to reinvent the wheel every time they need to incorporate new functionality. A web service is really a piece of reusable software code. A software developer can quickly build a new application by using many of these pieces of reusable code. The two primary parts of web services are events and services.

Events

Events are the eyes and ears of the business expressed in technology—they detect threats and opportunities and alert those who can act on the information. Pioneered by telecommunication and financial services companies, this involves using IT systems to monitor a business process for events that matter—a stock-out in the warehouse or an especially large charge on a consumer’s credit card—and automatically alert the people best equipped to handle the issue. For example, a credit monitoring system automatically alerts a credit supervisor and shuts down an account when the system processes a \$7,000 charge on a credit card with a \$6,000 limit.

Services

Services are more like software products than they are coding projects. They must appeal to a broad audience, and they need to be reusable if they are going to have an impact on productivity. Early forms of services were defined at too low a level in the architecture to interest the business, such as simple “Print” and “Save” services. The new services are being defined at a higher level; they describe such things as “Credit Check,” “Customer Information,” and “Process Payment.” These services describe a valuable business process. For example, “Credit Check” has value not just for programmers who want to use that code in another application, but also for businesspeople who want to use it across multiple products—say, auto loans and mortgages—or across multiple businesses.

The trick to building services is finding the right level of granularity. T-Mobile builds services starting at the highest level and then works its way down to lower levels, helping to ensure it does not build services that no one uses. The company first built a “Send Message” web service and then built a “Send SMS Message” web service that sends messages in special formats to different devices such as cell phones and pagers.

Lydian Trust’s enterprise architects designed a web service called “Get Credit” that is used by several different business units for loan applications. “Get Credit” seeks out credit ratings over the Internet from the major credit bureaus. One day, one of the credit bureaus’ web servers crashed, and Lydian Trust’s “Get Credit” web service could not make a connection. Since the connection to the server was loosely linked, the system did not know what to do. “Get Credit” was not built to make more than one call. So, while it waited for a response, hundreds of loan applications sat idle.

Lydian Trust’s loan officers had to work overnight to ensure that all of the applications were completed within 24 hours as promised by the company. Fortunately, Lydian Trust’s customers never felt the pain; however, its employees did. Systems must be designed to deal with the existence of certain events, or the lack of an event, in a way that does not interrupt the overall business. The “Get Credit” web service has been modified to include an automatic email alert to a supervisor whenever the web service encounters a delay.

OPEN SYSTEMS

Microsoft Internet Explorer’s share of the web browser market has dipped below 90 percent because of Mozilla’s Firefox, an open source web browser. According to WebSideStory, which has been tracking the Firefox versus Internet Explorer numbers, the Mozilla-made open source browser had captured 5 percent of the U.S. market in January 2005, an increase of almost a full percentage point in a month. Firefox claimed more than 25 million copies of the browser had been downloaded in its first 15 weeks of release.

An **open system** is a broad, general term that describes nonproprietary IT hardware and software made available by the standards and procedures by which their products work, making it easier to integrate them. Amazon.com embraced open

source technology converting from Sun's proprietary operating system to Linux. The switch to an open source operating system, such as Linux, is simplifying the process by which Amazon.com associates can build links to Amazon.com applications into their websites.

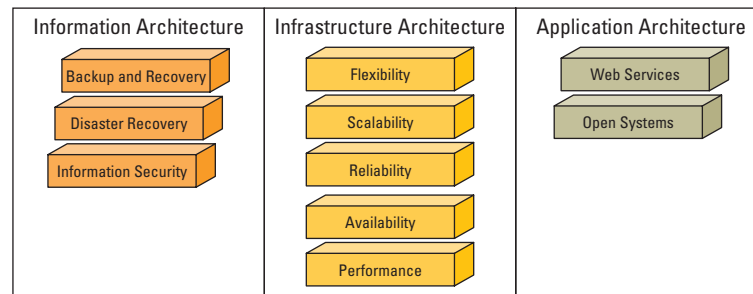
The designs of open systems allow for information sharing. In the past, different systems were independent of each other and operated as individual islands of control. The sharing of information was accomplished through software drivers and devices that routed data allowing information to be translated and shared between systems. Although this method is still widely used, its limited capability and added cost are not an effective solution for most organizations. Another drawback to the stand-alone system is it can communicate only with components developed by a single manufacturer. The proprietary nature of these systems usually results in costly repair, maintenance, and expansion because of a lack of competitive forces. On the other hand, open system integration is designed to:

- Allow systems to seamlessly share information. The sharing of information reduces the total number of devices, resulting in an overall decrease in cost.
- Capitalize on enterprise architectures. This avoids installing several independent systems, which creates duplication of devices.
- Eliminate proprietary systems and promote competitive pricing. Often a sole-source vendor can demand its price and may even provide the customer with less than satisfactory service. Utilization of open systems allows users to purchase systems competitively.

* PLUG-IN SUMMARY

Companies that have created solid enterprise architectures are reaping huge rewards in savings, flexibility, and business alignment. Basic enterprise architectures contain three components:

1. Information architecture identifies where and how important information, like customer records, is maintained and secured.
2. Infrastructure architecture includes the hardware, software, and telecommunications equipment that, when combined, provide the underlying foundation to support the organization's goals.
3. Application architecture determines how applications integrate and relate to each other.



* KEY TERMS

Application architecture, B4.3, B4.10
 Availability, B4.9
 Backup, B4.4
 Business continuity planning, B4.6
 Capacity planning, B4.9
 Cold site, B4.5
 Disaster recovery cost curve, B4.5

Disaster recovery plan, B4.5
 Enterprise architect (EA), B4.2
 Enterprise architecture, B4.2
 Failover, B4.5
 Fault tolerance, B4.4
 High availability, B4.9
 Hot site, B4.5
 Information architecture, B4.3
 Infrastructure architecture, B4.3, B4.8

Interoperability, B4.11
 Open system, B4.12
 Performance, B4.10
 Recovery, B4.4
 Reliability, B4.9
 Scalability, B4.9
 Web service, B4.10

* CLOSING CASE ONE

Chicago Tribune's Server Consolidation a Success

The *Chicago Tribune* is the seventh-largest newspaper in the country. Overhauling its data center and consolidating servers was a difficult task; however, the payoff was tremendous. The *Chicago Tribune* successfully moved its critical applications from a mishmash of mainframes and older Sun Microsystems servers to a new dual-site enterprise architecture, which has resulted in lower costs and increased reliability throughout the company.

The paper's new enterprise architecture clustered its servers over a two-mile distance, lighting up a 1Gbps dark-fiber link—an optical fiber that is in place but not yet being used—between two data centers. This architecture lets the newspaper spread the processing load between the servers while improving redundancy and options for disaster recovery.

The transfer to the new architecture was not smooth. A small piece of software written for the transition contained a coding error that caused the *Tribune's* editorial applications to experience intermittent processing failures. As a result, the paper was forced to delay delivery to about 40 percent of its 680,000 readers and cut 24 pages from a Monday edition, costing the newspaper nearly \$1 million in advertising revenue.

After editorial applications were stabilized, the *Tribune* proceeded to migrate applications for operations—the physical production and printing of the newspaper—and circulation to the new enterprise architecture. “As we gradually took applications off the mainframe, we realized that we were incurring very high costs in maintaining underutilized mainframes at two different locations,” says Darko Dejanovic, vice president and CTO of the Tribune Co., which owned the *Chicago Tribune*, the *Los Angeles Times*, Long Island’s *Newsday*, and about a dozen other metropolitan newspapers. “By moving from two locations to one, we’ve achieved several million dollars in cost savings. There’s no question that server consolidation was the right move for us.”

The company is excited about its new enterprise architecture and is looking to consolidate software across its newspapers. Currently, each newspaper maintains its own applications for classified advertising and billing, which means the parent company must support about 10 billing packages and the same number of classified-ad programs. Most of the business processes can be standardized. So far, the company has standardized about 95 percent of classified-ad processes and about 90 percent of advertising-sales processes. Over three years, the company will replace the disparate billing and ad applications with a single package that will be used by all business units. The different newspapers will not necessarily share the same data, but they will have the same processes and the same systems for accessing them. Over time, that will allow some of the call centers to handle calls for multiple newspapers; East Coast centers will handle the early-morning calls and West Coast centers the late-day and evening calls.

The company is looking at a few additional projects including the implementation of hardware that will allow its individual applications to run on partial CPUs, freeing up processor power and making more efficient use of disk space.

Questions

1. Review the five characteristics of infrastructure architecture and rank them in order of their potential impact on the Tribune Co.’s business.
2. What is the disaster recovery cost curve? Where should the Tribune Co. operate on the curve?
3. Define backups and recovery. What are the risks to the Tribune’s business if it fails to implement an adequate backup plan?
4. Why is a scalable and highly available enterprise architecture critical to current operations and future growth?
5. Identify the need for information security at the Tribune Co.
6. How could the Tribune Co. use a classified ad web service across its different businesses?

* CLOSING CASE TWO

Fear the Penguin

Linux has proved itself the most revolutionary software of the past decade. Spending on Linux was reported to reach \$280 million by 2006. Linus Torvalds, who wrote the kernel (the core) of the Linux operating system at age 21, posted the operating system on the Internet and invited other programmers to improve his code and users to download his operating system for free. Since then, tens of thousands of people have, making Linux perhaps the single largest collaborative project in the planet’s history.

Today, Linux, if not its penguin mascot, is everywhere. You can find Linux inside a boggling array of computers, machines, and devices. Linux is robust enough to run the world's most powerful supercomputers, yet sleek and versatile enough to run inside consumer items like TiVo, cell phones, and handheld portable devices. Even more impressive than Linux's increasing prevalence in living rooms and pockets is its growth in the market for corporate computers.

Since its introduction in 1991, no other operating system in history has spread as quickly across such a broad range of systems as Linux, and it has finally achieved critical mass. According to studies by market research firm IDC, Linux is the fastest-growing server operating system, with shipments expected to grow by 34 percent per year over the next four years. With its innovative open source approach, strong security, reliability, and scalability, Linux can help companies achieve the agility they need to respond to changing consumer needs and stay ahead of the game.

Thanks to its unique open source development process, Linux is reliable and secure. A "meritocracy," a team specifically selected for their competence by the technical developer community, governs the entire development process. Each line of code that makes up the Linux kernel is extensively tested and maintained for a variety of different platforms and application scenarios.

This open collaborative approach means the Linux code base continually hardens and improves itself. If vulnerabilities appear, they get the immediate attention of experts from around the world, who quickly resolve the problems. According to Security Portal, which tracks vendor response times, it takes an average of 12 days to patch a Linux bug compared to an average of three months for some proprietary platforms. With the core resilience and reliability of Linux, businesses can minimize downtime, which directly increases their bottom line.

The Spread of Open Systems

Businesses and governments are opting for open source operating systems like Linux instead of Windows. One attendee at the Linux Desktop Consortium in 2004 was Dr. Martin Echt, a cardiologist from Albany, New York. Dr. Echt, chief operating officer of Capital Cardiology Associates, an eight-office practice, discussed his decision to shift his business from Microsoft's Windows to Linux. Dr. Echt is not your typical computer geek or Linux supporter, and he is not the only one switching to Linux.

The State Council in China has mandated that all ministries install the local flavor of Linux, dubbed Red Flag, on their PCs. In Spain, the government has installed a Linux operating system that incorporates the regional dialect. The city of Munich, despite a personal visit from Microsoft CEO Steve Ballmer, is converting its 14,000 PCs from Windows to Linux.

"It's open season for open source," declared Walter Raizner, general manager of IBM Germany. One of the biggest corporate backers of Linux, IBM has more than 75 government customers worldwide, including agencies in France, Spain, Britain, Australia, Mexico, the United States, and Japan.

The move toward Linux varies for each country or company. For Dr. Echt, it was a question of lower price and long-term flexibility. In China, the government claimed national security as a reason to move to open source code because it permitted engineers to make sure there were no security leaks and no spyware installed on its computers. In Munich, the move was largely political. Regardless of the reason, the market is shifting toward Linux.

Microsoft versus Linux

Bill Gates has openly stated that Linux is not a threat to Microsoft. According to IDC analysts, Microsoft's operating systems ship with 93.8 percent of all desktops worldwide. Ted Schadler, IDC research principal analyst, states that despite the push of lower cost Linux players into the market, Microsoft will maintain its desktop market share for the following three reasons:

1. Linux adds features to its applications that most computer users have already come to expect.
2. Linux applications might not be compatible with Microsoft applications such as Microsoft Word or Microsoft Excel.

3. Microsoft continues to innovate, and the latest version of Office is beginning to integrate word processing and spreadsheet software to corporate databases and other applications.

The Future of Linux

IDC analyst AI Gillen predicts that an open source operating system will not enjoy explosive growth on the desktop for at least six or eight years. Still, even Gillen cannot deny that Linux's penetration continues to rise, with an estimated 18 million users. Linux's market share increased from 1.5 percent at the end of 2000 to 4.2 percent at the beginning of 2004. According to IDC, by the end of 2005 it surpassed Apple's Mac OS, which has 2.9 percent of the market, as the second most popular operating system. Gartner Dataquest estimates Linux's server market share will grow seven times faster than Windows.

Questions

1. How does Linux differ from traditional software?
2. Should Microsoft consider Linux a threat? Why or why not?
3. How is open source software a potential trend shaping organizations?
4. How can you use Linux as an emerging technology to gain a competitive advantage?
5. Research the Internet and discover potential ways that open source software might revolutionize business in the future.

* MAKING BUSINESS DECISIONS

1. Planning for Disaster Recovery

You are the new senior analyst in the IT department at Beltz, a large snack food manufacturing company. The company is located on the beautiful shoreline in Charleston, North Carolina. The company's location is one of its best and also worst features. The weather and surroundings are beautiful, but the threat of hurricanes and other natural disasters is high. Compile a disaster recovery plan that will minimize any risks involved with a natural disaster.

2. Comparing Backup and Recovery Systems

Research the Internet to find three different vendors of backup and recovery systems. Compare and contrast the three systems and determine which one you would recommend if you were installing a backup and recovery system for a medium-sized business with 3,500 employees that maintains information on the stock market. Compile your findings in a presentation that you can give to your class that details the three systems' strengths and weaknesses, along with your recommendation.

3. Ranking the -ilities

In a group, review the following list of IT infrastructure qualities and rank them in order of their impact on an organization's success. Use a rating system of 1 to 7, where 1 indicates the biggest impact and 7 indicates the least impact.

4. Designing an Enterprise Architecture

Components of a solid enterprise architecture include everything from documentation to business concepts to software and hardware. Deciding which components to implement and how to implement them can be a challenge. New IT components are released

IT Infrastructure Qualities	Business Impact
Availability	
Accessibility	
Reliability	
Scalability	
Flexibility	
Performance	
Capacity Planning	

daily, and business needs continually change. An enterprise architecture that meets your organization’s needs today may not meet those needs tomorrow. Building an enterprise architecture that is scalable, flexible, available, accessible, and reliable is key to your organization’s success.

You are the enterprise architect (EA) for a large clothing company called Xedous. You are responsible for developing the initial enterprise architecture. Create a list of questions you will need answered to develop your architecture. Below is an example of a few of the questions you might ask.

- What are the company’s growth expectations?
- Will systems be able to handle additional users?
- How long will information be stored in the systems?
- How much customer history must be stored?
- What are the organization’s business hours?
- What are the organization’s backup requirements?

