

This text was adapted by The Saylor Foundation under a [Creative Commons Attribution-NonCommercial-ShareAlike 3.0 License](https://creativecommons.org/licenses/by-nc-sa/3.0/) without attribution as requested by the work's original creator or licensee.



Chapter 1

Zara: Fast Fashion from Savvy Systems

1.1 Introduction

LEARNING OBJECTIVE

After studying this section you should be able to do the following:

1. Understand how Zara's parent company Inditex leveraged a technology-enabled strategy to become the world's largest fashion retailer.

The poor, ship-building town of La Coruña in northern Spain seems an unlikely home to a tech-charged innovator in the decidedly ungeeky fashion industry, but that's where you'll find "The Cube," the gleaming, futuristic central command of the Inditex Corporation (Industrias de Diseño Textil), parent of game-changing clothes giant, Zara. The blend of technology-enabled strategy that Zara has unleashed seems to break all of the rules in the fashion industry. The firm shuns advertising and rarely runs sales. Also, in an industry where nearly every major player outsources manufacturing to low-cost countries, Zara is highly vertically integrated, keeping huge swaths of its production process in-house. These counterintuitive moves are part of a recipe for success that's beating the pants off the competition, and it has turned the founder of Inditex, Amancio Ortega, into Spain's wealthiest man and the world's richest fashion executive.

The firm tripled in size between 1996 and 2000, then its earnings skyrocketed from \$2.43 billion in 2001 to \$13.6 billion in 2007. By August 2008, sales edged ahead of Gap, making Inditex the world's largest fashion retailer. ^[1]Table 1.1 "Gap versus Inditex at a Glance" compares the two fashion retailers. While the firm supports eight brands, Zara is



unquestionably the firm's crown jewel and growth engine, accounting for roughly two-thirds of sales. [2]

Why Study Zara?

While competitors falter, Zara is undergoing one of the fastest global expansions the fashion world has ever seen, opening one store per day and entering new markets worldwide—seventy-three countries so far. The chain's profitability is among the highest in the industry. [3] The fashion director for luxury goods maker LVMH calls Zara “the most innovative and devastating retailer in the world.” [4]

Zara's duds look like high fashion but are comparatively inexpensive (average item price is \$27, although prices vary by country). [5] A Goldman analyst has described the chain as “Armani at moderate prices,” while another industry observer suggests that while fashions are more “Banana Republic,” prices are more “Old Navy.” [6] Legions of fans eagerly await “Z-day,” the twice-weekly inventory delivery to each Zara location that brings in the latest clothing lines for women, men, and children.

In order to understand and appreciate just how counterintuitive and successful Zara's strategy is, and how technology makes all of this possible, it's important to first examine the conventional wisdom in apparel retail. To do that we'll look at former industry leader—Gap.

Gap: An Icon in Crisis

Most fashion retailers place orders for a seasonal collection months before these lines make an appearance in stores. While overseas contract manufacturers may require hefty lead times, trying to guess what customers want months in advance is a tricky business. In retail in general and fashion in particular, there's a saying: inventory equals death. Have too much unwanted product on hand and you'll be forced to mark down or

write off items, killing profits. For years, Gap sold most of what it carried in stores. Micky Drexler, a man with a radar-accurate sense of style and the iconic CEO who helped turn Gap's button-down shirts and khakis into America's business casual uniform, led the way. Drexler's team had spot-on tastes throughout the 1990s, but when sales declined in the early part of the following decade, Drexler was left guessing on ways to revitalize the brand, and he guessed wrong—disastrously wrong. Chasing the youth market, Drexler filled Gap stores with miniskirts, low-rise jeans, and even a much-ridiculed line of purple leather pants. [7] The throngs of teenagers he sought to attract never showed up, and the shift in offerings sent Gap's mainstay customers to retailers that easily copied the styles that Gap had made classic.

The inventory hot potato Drexler was left with crushed the firm. Gap's same-store sales declined for twenty-nine months straight. Profits vanished. Gap founder and chairman Dan Fisher lamented, "It took us thirty years to get to \$1 billion in profits and two years to get to nothing." [8] The firm's debt was downgraded to junk status. Drexler was out and for its new head the board chose Paul Pressler, a Disney executive who ran theme parks and helped rescue the firm's once ailing retail effort. Pressler shut down hundreds of stores, but the hemorrhaging continued largely due to bad bets on colors and styles. [9] During one holiday season, Gap's clothes were deemed so off target that the firm scrapped its advertising campaign and wrote off much of the inventory. The marketing model used by Gap to draw customers in via big-budget television promotion had collapsed. Pressler's tenure saw same-store sales decline in eighteen of twenty-four months. [10] A *Fortune* article on Pressler's leadership was titled "Fashion Victim." *BusinessWeek* described his time as CEO as a "Total System Failure," [11] and Wall Street began referring to him as DMW for Dead Man Walking. In January 2007, Pressler resigned, with



Gap hoping its third chief executive of the decade could right the ailing giant.

Contract Manufacturing: Lower Costs at What Cost?

Conventional wisdom suggests that leveraging cheap contract manufacturing in developing countries can keep the cost of goods low. Firms can lower prices and sell more product or maintain higher profit margins—all good for the bottom line. But many firms have also experienced the ugly downside to this practice. Global competition among contract firms has led to race-to-the-bottom cost-cutting measures. Too often, this means that in order to have the low-cost bid, contract firms skimp on safety, ignore environmental concerns, employ child labor, and engage in other ghastly practices.

The apparel industry in particular has been plagued by accusations of employing sweatshop labor to keep costs down. Despite the fact that Gap audits contract manufacturers and has a high standard for partner conduct, the firm has repeatedly been taken to task by watchdog groups, the media, and its consumers, who have exposed unacceptable contract manufacturing conditions that Gap failed to catch. This negative exposure includes the October 2007 video showing Gap clothes made by New Delhi children as young as ten years old in what were described as “slave labor” conditions. ^[12]

Gap is not alone; Nike, Wal-Mart, and many other apparel firms have been tarnished in similar incidents. Big firms are big targets and those that fail to adequately ensure their products are made under acceptable labor conditions risk a brand-damaging backlash that may turn off customers, repel new hires, and leave current staff feeling betrayed. Today’s manager needs to think deeply not only about their own firm’s ethical practices, but also those of all of their suppliers and partners.

Tech for Good: The Fair Factories Clearinghouse

The problem of sweatshop labor has plagued the clothing industry for years. Managers often feel the pressure to seek ever-lower costs and all too often end up choosing suppliers with unacceptably poor practices. Even well-meaning firms can find themselves stung by corner-cutting partners that hide practices from auditors or truck products in from unmonitored off-site locations. The results can be tragic for those exploited, and can carry lasting negative effects for the firm. The sweatshop moniker continues to dog Nike years after allegations were uncovered and the firm moved aggressively to deal with its problems.

Nike rival Reebok (now part of Adidas) has always taken working conditions seriously. The firm even has a Vice President of Human Rights, and has made human dignity a key platform for its philanthropic efforts. Reebok invested millions in developing an in-house information system to track audits of its hundreds of suppliers along dimensions such as labor, safety, and environmental practices. The goal in part was to identify any bad apples, so that one division, sporting goods, for example, wouldn't use a contractor identified as unacceptable by the sneaker line.

The data was valuable to Reebok, particularly given that the firm has hundreds of contract suppliers. But senior management realized the system would do even more good if the whole industry could share and contribute information. Reebok went on to donate the system and provided critical backing to help create the nonprofit organization Fair Factories Clearinghouse. With management that includes former lawyers for Amnesty International, Fair Factories (FairFactories.org) provides systems where apparel and other industries can share audit information on contract manufacturers. Launching the effort wasn't as easy as sharing the technology. The U.S. Department of Justice needed to provide a special

exemption, and had to be convinced the effort wouldn't be used by buyers to collude and further squeeze prices from competitors (the system is free of pricing data).

Suppliers across industries now recognize that if they behave irresponsibly the Fair Factories system will carry a record of their misdeeds, notifying all members to avoid the firm. As more firms use the system, its database becomes broader and more valuable. To their credit, both Gap and Nike have joined the Fair Factories Clearinghouse.

KEY TAKEAWAYS

- Zara has used technology to dominate the retail fashion industry as measured by sales, profitability, and growth.
- Excess inventory in the retail apparel industry is the kiss of death. Long manufacturing lead times require executives to guess far in advance what customers will want. Guessing wrong can be disastrous, lowering margins through markdowns and write-offs.
- Contract manufacturing can offer firms several advantages, including lower costs and increased profits. But firms have also struggled with the downside of cost-centric contract manufacturing when partners have engaged in sweatshop labor and environmental abuse.

QUESTIONS AND EXERCISES

1. Has anyone shopped at Zara? If so, be prepared to share your experiences and observations with your class. What did you like about the store? What didn't you like? How does Zara differ from other clothing retailers in roughly the same price range? If you've visited Zara locations in different countries, what differences did you notice in terms of offerings, price, or other factors?
2. What is the "conventional wisdom" of the fashion industry with respect to design, manufacturing, and advertising?

3. What do you suppose are the factors that helped Gap to at one point rise to be first in sales in the fashion industry?
4. Who ran Gap in the 1990s? How did the executive perform prior to leaving Gap? Describe what happened to sales. Why?
5. Who was the Gap's second CEO of this decade? How did sales fare under him? Why?
6. Where do Gap clothes come from? Who makes them? Why? Are there risks in this approach?
7. Describe the Fair Factories Clearinghouse. Which firm thought of this effort? Why did they give the effort away? What happens as more firms join this effort and share their data?



1.2 Don't Guess, Gather Data

LEARNING OBJECTIVE

After studying this section you should be able to do the following:

1. Contrast Zara's approach with the conventional wisdom in fashion retail, examining how the firm's strategic use of information technology influences design and product offerings, manufacturing, inventory, logistics, marketing, and ultimately profitability.

Having the wrong items in its stores hobbled Gap for nearly a decade. But how do you make sure stores carry the kinds of things customers want to buy? Try asking them. Zara's store managers lead the intelligence-gathering effort that ultimately determines what ends up on each store's racks. Armed with personal digital assistants (PDAs)—handheld computing devices meant largely for mobile use outside an office setting—to gather customer input, staff regularly chat up customers to gain feedback on what they'd like to see more of. A Zara manager might casually ask, What if this skirt were in a longer length? Would you like it in a different color? What if this V-neck blouse were available in a round neck? Managers are motivated because they have skin in the game. The firm is keen to reward success—as much as 70 percent of salaries can come from commissions.^[1]

Another level of data gathering starts as soon as the doors close. Then the staff turns into a sort of investigation unit in the forensics of trendspotting, looking for evidence in the piles of unsold items that customers tried on but

didn't buy. Are there any preferences in cloth, color, or styles offered among the products in stock? [2]

PDA's are also linked to the store's point-of-sale (POS) system—a transaction process that captures customer purchase information—showing how garments rank by sales. In less than an hour, managers can send updates that combine the hard data captured at the cash register with insights on what customers would like to see. [3] All of this valuable data allows the firm to plan styles and issue rebuy orders based on feedback rather than hunches and guesswork. The goal is to improve the frequency and quality of decisions made by the design and planning teams.

Design

Rather than create trends by pushing new lines via catwalk fashion shows, Zara designs follow evidence of customer demand. Data on what sells and what customers want to see goes directly to “The Cube” outside La Coruña, where teams of some three hundred designers crank out an astonishing thirty thousand items a year versus two to four thousand items offered up at big chains like H&M (the world's third largest fashion retailer) and Gap. [4] While H&M has offered lines by star designers like Stella McCartney and Karl Lagerfeld, as well as celebrity collaborations with Madonna and Kylie Minogue, the Zara design staff consists mostly of young, hungry *Project Runway* types fresh from design school. There are no prima donnas in “The Cube.” Team members must be humble enough to accept feedback from colleagues and share credit for winning ideas. Individual bonuses are tied to the success of the team, and teams are regularly rotated to cross-pollinate experience and encourage innovation.

Manufacturing and Logistics

In the fickle world of fashion, even seemingly well-targeted designs could go out of favor in the months it takes to get plans to contract manufacturers, tool up production, then ship items to warehouses and eventually to retail locations. But getting locally targeted designs quickly onto store shelves is where Zara really excels. In one telling example, when Madonna played a set of concerts in Spain, teenage girls arrived to the final show sporting a Zara knock-off of the outfit she wore during her first performance. [5] The average time for a Zara concept to go from idea to appearance in store is fifteen days versus their rivals who receive new styles once or twice a season. Smaller tweaks arrive even faster. If enough customers come in and ask for a round neck instead of a V neck, a new version can be in stores with in just ten days. [6] To put that in perspective, Zara is *twelve times* faster than Gap despite offering roughly *ten times* more unique products! [7] At H&M, it takes three to five months to go from creation to delivery—and they're considered one of the best. Other retailers need an average of six months to design a new collection and then another three months to manufacture it. VF Corp (Lee, Wrangler) can take nine months just to design a pair of jeans, while J. Jill needs a year to go from concept to store shelves. [8] At Zara, most of the products you see in stores didn't exist three weeks earlier, not even as sketches. [9]

The firm is able to be so responsive through a competitor-crushing combination of vertical integration and technology-orchestrated coordination of suppliers, just-in-time manufacturing, and finely tuned logistics. Vertical integration is when a single firm owns several layers in its value chain. [10] While H&M has nine hundred suppliers and no factories, nearly 60 percent of Zara's merchandise is produced in-house, with an eye on leveraging technology in those areas that speed up complex tasks, lower cycle time, and reduce error. Profits from this clothing retailer come from

blending math with a data-driven fashion sense. Inventory optimization models help the firm determine how many of which items in which sizes should be delivered to each specific store during twice-weekly shipments, ensuring that each store is stocked with just what it needs. ^[11] Outside the distribution center in La Coruña, fabric is cut and dyed by robots in twenty-three highly automated factories. Zara is so vertically integrated, the firm makes 40 percent of its own fabric and purchases most of its dyes from its own subsidiary. Roughly half of the cloth arrives undyed so the firm can respond as any midseason fashion shifts occur. After cutting and dyeing, many items are stitched together through a network of local cooperatives that have worked with Inditex so long they don't even operate with written contracts. The firm does leverage contract manufacturers (mostly in Turkey and Asia) to produce staple items with longer shelf lives, such as t-shirts and jeans, but such goods account for only about one-eighth of dollar volume. ^[12]

All of the items the firm sells end up in a five-million-square-foot distribution center in La Coruña, or a similar facility in Zaragoza in the northeast of Spain. The La Coruña facility is some nine times the size of Amazon's warehouse in Fernley, Nevada, or about the size of ninety football fields. ^[13] The facilities move about two and a half million items every week, with no item staying in-house for more than seventy-two hours. Ceiling-mounted racks and customized sorting machines patterned on equipment used by overnight parcel services, and leveraging Toyota-designed logistics, whisk items from factories to staging areas for each store. Clothes are ironed in advance and packed on hangers, with security and price tags affixed. This system means that instead of wrestling with inventory during busy periods, employees in Zara stores simply move items from shipping box to store racks, spending most of their time on value-added functions



like helping customers find what they want. Efforts like this help store staff regain as much as three hours in prime selling time. [14]

Trucks serve destinations that can be reached overnight, while chartered cargo flights serve farther destinations within forty-eight hours. [15] The firm recently tweaked its shipping models through Air France–KLM Cargo and Emirates Air so flights can coordinate outbound shipment of all Inditex brands with return legs loaded with raw materials and half-finished clothes items from locations outside of Spain. Zara is also a pioneer in going green. In fall 2007, the firm’s CEO unveiled an environmental strategy that includes the use of renewable energy systems at logistics centers including the introduction of biodiesel for the firm’s trucking fleet.

Stores

Most products are manufactured for a limited production run. While running out of bestsellers might be seen as a disaster at most retailers, at Zara the practice delivers several benefits.

First, limited runs allow the firm to cultivate the exclusivity of its offerings. While a Gap in Los Angeles carries nearly the same product line as one in Milwaukee, each Zara store is stocked with items tailored to the tastes of its local clientele. A Fifth Avenue shopper quips, “At Gap, everything is the same,” while a Zara shopper in Madrid says, “you’ll never end up looking like someone else.” [16] Upon visiting a Zara, the CEO of the National Retail Federation marveled, “It’s like you walk into a new store every two weeks.” [17]

Second, limited runs encourage customers to buy right away and at full price. Savvy Zara shoppers know the newest items arrive on black plastic hangers, with store staff transferring items to wooden ones later on. Don’t bother asking when something will go on sale; if you wait three weeks the item you wanted has almost certainly been sold or moved out to make room

for something new. Says one twenty-three year-old Barcelona shopper, “If you see something and don’t buy it, you can forget about coming back for it because it will be gone.” [18] A study by consulting firm Bain & Company estimated that the industry average markdown ratio is approximately 50 percent, while Zara books some 85 percent of its products at full price. [19] The constant parade of new, limited-run items also encourages customers to visit often. The average Zara customer visits the store seventeen times per year, compared with only three annual visits made to competitors. [20] Even more impressive—Zara puts up these numbers with almost no advertising. The firm’s founder has referred to advertising as a “pointless distraction.” The assertion carries particular weight when you consider that during Gap’s collapse, the firm increased advertising spending but sales dropped. [21] Fashion retailers spend an average of 3.5 percent of revenue promoting their products, while ad spending at Inditex is just 0.3 percent. [22]

Finally, limited production runs allows the firm to, as Zara’s CEO once put it “reduce to a minimum the risk of making a mistake, and we do make mistakes with our collections.” [23] Failed product introductions are reported to be just 1 percent, compared with the industry average of 10 percent. [24] So even though Zara has higher manufacturing costs than rivals, Inditex gross margins are 56.8 percent compared to 37.5 percent at Gap. [25]

While stores provide valuable front-line data, headquarters plays a major role in directing in-store operations. Software is used to schedule staff based on each store’s forecasted sales volume, with locations staffing up at peak times such as lunch or early evening. The firm claims these more flexible schedules have shaved staff work hours by 2 percent. This constant refinement of operations throughout the firm’s value chain has helped reverse a prior trend of costs rising faster than sales. [26]



Even the store displays are directed from “The Cube,” where a basement staging area known as “Fashion Street” houses a Potemkin village of bogus storefronts meant to mimic some of the chain’s most exclusive locations throughout the world. It’s here that workers test and fine-tune the chain’s award-winning window displays, merchandise layout, even determine the in-store soundtrack. Every two weeks, new store layout marching orders are forwarded to managers at each location. [27]

Technology ≠ Systems. Just Ask Prada

Here’s another interesting thing about Zara. Given the sophistication and level of technology integration into the firm’s business processes, you’d think that Inditex would far outspend rivals on tech. But as researchers Donald Sull and Sefano Turconi discovered, “Whether measured by IT workers as a percentage of total employees or total spending as a percentage of sales, Zara’s IT expenditure is less than one-fourth the fashion industry average.” [28] Zara excels by targeting technology investment at the points in its value chain where it will have the most significant impact, making sure that every dollar spent on tech has a payoff. Contrast this with high-end fashion house Prada’s efforts at its flagship Manhattan location. The firm hired the Pritzker Prize–winning hipster architect Rem Koolhaas to design a location Prada would fill with jaw-dropping technology. All items for sale in the store would sport with radio frequency identification (RFID) tags (small chip-based tags that wirelessly emit a unique identifying code for the item that they are attached to). Walk into a glass dressing room and customers could turn the walls opaque, then into a kind of combination mirror and heads-up display. By wirelessly reading the tags on each garment, dressing rooms would recognize what was brought in and make recommendations of matching accessories as well as similar products that patrons might consider.

Customers could check inventory, and staff sporting PDAs could do the same. A dressing room camera would allow clients to see their front and back view side-by-side as they tried on clothes.

It all sounded slick, but execution of the vision was disastrous. Customers didn't understand the foot pedals that controlled the dressing room doors and displays. Reports surfaced of fashionistas disrobing in full view, thinking the walls went opaque when they didn't. Others got stuck in dressing rooms when pedals failed to work, or doors broke, unable to withstand the demands of the high-traffic tourist location. The inventory database was often inaccurate, regularly reporting items as out of stock even though they weren't. As for the PDAs, staff reported that they "don't really use them anymore" and that "we put them away so tourists don't play with them." The investment in Prada's in-store technology was also simply too high, with estimates suggesting the location took in just one-third the sales needed to justify expenses. [29]

The Prada example offers critical lessons for managers. While it's easy to get seduced by technology, an information system (IS) is actually made up of more than *hardware* and *software*. An IS also includes *data* used or created by the system, as well as the *procedures* and the *people* who interact with the system. [30] Getting the right mix of these five components is critical to executing a flawless information system rollout. Financial considerations should forecast the return on investment (ROI)—the amount earned from an expenditure—of any such effort (i.e., what will we get for our money and how long will it take to receive payback?). And designers need to thoroughly test the system before deployment. At Prada's Manhattan flagship store, the effort looked like tech chosen because it seemed fashionable rather than functional.

KEY TAKEAWAYS

- Zara store management and staff use PDAs and POS systems to gather and analyze customer preference data to plan future designs based on feedback, rather than on hunches and guesswork.
- Zara's combination of vertical integration and technology-orchestrated supplier coordination, just-in-time manufacturing, and logistics allows it to go from design to shelf in days instead of months.
- Advantages accruing to Inditex include fashion exclusivity, fewer markdowns and sales, lower marketing expenses, and more frequent customer visits.
- Zara's IT expenditures are low by fashion industry standards. The spectacular benefits reaped by Zara from the deployment of technology have resulted from targeting technology investment at the points in the value chain where it has the greatest impact, and not from the sheer magnitude of the investment. This is in stark contrast to Prada's experience with in-store technology deployment.

QUESTIONS AND EXERCISES

1. In what ways is the Zara model counterintuitive? In what ways has Zara's model made the firm a better performer than Gap and other competitors?
2. What factors account for a firm's profit margin? What does Gap focus on? What factors does Zara focus on to ensure a strong profit margin?
3. How is data captured in Zara stores? Using what types or classifications of information systems? How does the firm use this data?
4. What role does technology play in enabling the other elements of Zara's counterintuitive strategy? Could the firm execute its strategy without technology? Why or why not?
5. How does technology spending at Zara compare to that of rivals? Advertising spending? Failed product percentages? Markdowns?

6. What risks are inherent in the conventional practices in the fashion industry? Is Zara susceptible to these risks? Is Zara susceptible to different risks? If so, what are these?
7. Consider the Prada case mentioned in the sidebar “Technology ≠ Systems.” What did Prada fail to consider when it rolled out the technology in its flagship location? Could this effort have been improved for better results? If you were put in charge of this kind of effort, what factors would you consider? What would determine whether you’d go forward with the effort or not? If you did go forward, what factors would you consider and how might you avoid some of the mistakes made by Prada?

1.3 Moving Forward

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Detail how Zara’s approach counteracts specific factors that Gap has struggled with for over a decade.
2. Identify the environmental threats that Zara is likely to face, and consider options available to the firm for addressing these threats.

The holy grail for the strategist is to craft a sustainable competitive advantage that is difficult for competitors to replicate. And for nearly two decades Zara has delivered the goods. But that’s not to say the firm is done facing challenges.

Consider the limitations of Zara’s Spain-centric, just-in-time manufacturing model. By moving all of the firm’s deliveries through just two locations, both in Spain, the firm remains hostage to anything that

could create a disruption in the region. Firms often hedge risks that could shut down operations—think weather, natural disaster, terrorism, labor strife, or political unrest—by spreading facilities throughout the globe. If problems occur in northern Spain, Zara has no such fall back.

In addition to the operations (the organizational activities that are required to produce goods or services) vulnerabilities above, the model also leaves the firm potentially more susceptible to financial vulnerabilities as the Euro has strengthened relative to the dollar. Many low-cost manufacturing regions have currencies that are either pegged to the dollar or have otherwise fallen against the Euro. This situation means Zara's Spain-centric costs rise at higher rates compared to competitors, presenting a challenge in keeping profit margins in check. Rising transportation costs are another concern. If fuel costs rise, the model of twice-weekly deliveries that has been key to defining the Zara experience becomes more expensive to maintain.

Still, Zara is able to make up for some cost increases by raising prices overseas (in the United States, Zara items can cost 40 percent or more than they do in Spain). Zara reports that all North American stores are profitable, and that it can continue to grow its presence, serving forty to fifty stores with just two U.S. jet flights a week. ^[1] Management has considered a logistics center in Asia, but expects current capacity will suffice until 2013. ^[2] A center in the Maquiladora region of northern Mexico, for example, may also be able to serve the U.S. markets via trucking capacity similar to the firm's Spain-based access to Europe, while also providing a regional center to serve growth throughout Latin America should the firm continue its expansion across the Western Hemisphere. Rivals have studied the firm's secret recipe, and while none have attained the efficiency of Amancio Ortega's firm, many are trying to learn from the master. There is precedent for contract firms closing the cycle time gap with

vertically integrated competitors that own their own factories. Dell (a firm that builds its own PCs while nearly all its competitors use contract labor) has recently seen its manufacturing advantage from vertical integration fall as the partners that supply rivals have mimicked its techniques and have become far more efficient. [3] In terms of the number of new models offered, clothing is actually more complex than computing, suggesting that Zara's value chain may be more difficult to copy. Still, H&M has increased the frequency of new items in stores, Forever 21 and Uniqlo get new looks within six weeks, Renner, a Brazilian fast fashion rival, rolls out mini collections every two months, and Benetton, a firm that previously closed some 90 percent of U.S. stores, now replenishes stores as fast as once a week. [4] Rivals have a keen eye on Inditex, with the CFO of luxury goods firm Burberry claiming the firm is a "fantastic case study" and "we're mindful of their techniques." [5]

Finally, firm financial performance can also be impacted by broader economic conditions. When the economy falters, consumers simply buy less and may move a greater share of their wallet to less-stylish and lower-cost offerings from deep discounters like Wal-Mart. Zara is particularly susceptible to conditions in Spain, since the market accounts for nearly 40 percent of Inditex sales, [6] as well as to broader West European conditions (which with Spain make up 79 percent of sales). [7] Global expansion will provide the firm with a mix of locations that may be better able to endure downturns in any single region.

Zara's winning formula can only exist through management's savvy understanding of how information systems can enable winning strategies (many tech initiatives were led by José Maria Castellano, a "technophile" business professor who became Ortega's right-hand man in the 1980s). [8] It is technology that helps Zara identify and manufacture the clothes customers want, get those products to market quickly, and eliminate costs

related to advertising, inventory missteps, and markdowns. A strategist must always scan the state of the market as well as the state of the art in technology, looking for new opportunities and remaining aware of impending threats. With systems so highly tuned for success, it may be unwise to bet against “The Cube.”

KEY TAKEAWAY

- Zara’s value chain is difficult to copy; but it is not invulnerable, nor is future dominance guaranteed. Zara management must be aware of the limitations in its business model, and must continually scan its environment and be prepared to react to new threats and opportunities.

QUESTIONS AND EXERCISES

1. The Zara case shows how information systems can impact every single management discipline. Which management disciplines were mentioned in this case? How does technology impact each?
2. Would a traditional Internet storefront work well with Zara’s business model? Why or why not?

Chapter 2

Strategy and Technology: Concepts and Frameworks for Understanding What Separates Winners from Losers

2.1 Introduction

LEARNING OBJECTIVES

1. Define operational effectiveness and understand the limitations of technology-based competition leveraging this principle.
2. Define strategic positioning and the importance of grounding competitive advantage in this concept.
3. Understand the resource-based view of competitive advantage.
4. List the four characteristics of a resource that might possibly yield sustainable competitive advantage.

Managers are confused, and for good reason. Management theorists, consultants, and practitioners often vehemently disagree on how firms should craft tech-enabled strategy, and many widely read articles contradict one another. Headlines such as “Move First or Die” compete with “The First-Mover Disadvantage.” A leading former CEO advises, “destroy your business,” while others suggest firms focus on their “core competency” and “return to basics.” The pages of the *Harvard Business Review* have declared, “IT Doesn’t Matter,” while a *New York Times* bestseller hails technology as the “steroids” of modern business.



Theorists claiming to have mastered the secrets of strategic management are contentious and confusing. But as a manager, the ability to size up a firm’s strategic position and understand its likelihood of sustainability is one of the most valuable and yet most difficult skills to master. Layer on thinking about technology—a key enabler to nearly every modern business strategy, but also a function often thought of as easily “outsourced”—and it’s no wonder that so many firms struggle at the intersection where strategy and technology meet. The business landscape is littered with the corpses of firms killed by managers who guessed wrong.

Developing strong strategic thinking skills is a career-long pursuit—a subject that can occupy tomes of text, a roster of courses, and a lifetime of seminars. While this chapter can’t address the breadth of strategic thought, it is meant as a primer on developing the skills for strategic thinking about technology. A manager that understands issues presented in this chapter should be able to see through seemingly conflicting assertions about best practices more clearly; be better prepared to recognize opportunities and risks; and be more adept at successfully brainstorming new, tech-centric approaches to markets.

The Danger of Relying on Technology

Firms strive for sustainable competitive advantage, financial performance that consistently outperforms their industry peers. The goal is easy to state, but hard to achieve. The world is so dynamic, with new products and new competitors rising seemingly overnight, that truly sustainable advantage might seem like an impossibility. New competitors and copycat products create a race to cut costs, cut prices, and increase features that may benefit consumers but erode profits industry-wide. Nowhere is this balance more difficult than when competition involves technology. The fundamental strategic question in the Internet era is, “*How can I possibly compete when everyone can copy my technology and the competition is just a click away?*”



Put that way, the pursuit of sustainable competitive advantage seems like a lost cause.

But there are winners—big, consistent winners—empowered through their use of technology. How do they do it? In order to think about how to achieve sustainable advantage, it's useful to start with two concepts defined by Michael Porter. A professor at the Harvard Business School and father of the *value chain* and the *five forces* concepts (see the sections later in this chapter), Porter is justifiably considered one of the leading strategic thinkers of our time.

According to Porter, the reason so many firms suffer aggressive, margin-eroding competition is because they've defined themselves according to operational effectiveness rather than strategic positioning. Operational effectiveness refers to performing the same tasks better than rivals perform them. Everyone wants to be better, but the danger in operational effectiveness is “sameness.” This risk is particularly acute in firms that rely on technology for competitiveness. After all, technology can be easily acquired. Buy the same stuff as your rivals, hire students from the same schools, copy the look and feel of competitor Web sites, reverse engineer their products, and you can match them. The fast follower problem exists when savvy rivals watch a pioneer's efforts, learn from their successes and missteps, then enter the market quickly with a comparable or superior product at a lower cost.

Since tech can be copied so quickly, followers can be fast, indeed. Several years ago while studying the Web portal industry (Yahoo! and its competitors), a colleague and I found that when a firm introduced an innovative feature, at least one of its three major rivals would match that feature in, on average, only one and a half months. ^[1] Groupon CEO Andrew Mason claimed the daily deal service had spawned 500 imitators within two years of launch. ^[2] When



technology can be matched so quickly, it is rarely a source of competitive advantage. And this phenomenon isn't limited to the Web.

Consider TiVo. At first blush, it looks like this first mover should be a winner since it seems to have established a leading brand; TiVo is now a verb for digitally recording TV broadcasts. But despite this, TiVo has largely been a money loser, going years without posting an annual profit. By the time 1.5 million TiVos had been sold, there were over thirty million digital video recorders (DVRs) in use. [3] Rival devices offered by cable and satellite companies appear the same to consumers and are offered along with pay television subscriptions—a critical distribution channel for reaching customers that TiVo doesn't control.

The Flip video camera is another example of technology alone offering little durable advantage. The pocket-sized video recorders used flash memory instead of magnetic storage. Flip cameras grew so popular that Cisco bought Flip parent Pure Digital, for \$590 million. The problem was digital video features were easy to copy, and constantly falling technology costs (see [Chapter 5 "Moore's Law: Fast, Cheap Computing and What It Means for the Manager"](#)) allowed rivals to embed video into their products. Later that same year Apple (and other firms) began including video capture as a feature in their music players and phones. Why carry a Flip when one pocket device can do everything? The Flip business barely lasted two years, and by spring 2011 Cisco had killed the division, taking a more than half-billion-dollar spanking in the process. [4]

Operational effectiveness is critical. Firms must invest in techniques to improve quality, lower cost, and design efficient customer experiences. But for the most part, these efforts can be matched. Because of this, operational effectiveness is usually not sufficient enough to yield sustainable dominance over the competition. In contrast to operational effectiveness, strategic positioning refers to performing different activities

from those of rivals, or the same activities in a different way. Technology itself is often very easy to replicate, and those assuming advantage lies in technology alone may find themselves in a profit-eroding arms race with rivals able to match their moves step by step. But while technology can be copied, technology can also play a critical role in creating and strengthening strategic *differences*—advantages that rivals will struggle to match.

Different Is Good: FreshDirect Redefines the NYC Grocery Landscape

For an example of the relationship between technology and strategic positioning, consider FreshDirect. The New York City–based grocery firm focused on the two most pressing problems for Big Apple shoppers: selection is limited and prices are high. Both of these problems are a function of the high cost of real estate in New York. The solution? Use technology to craft an ultraefficient model that makes an end-run around stores.

The firm’s “storefront” is a Web site offering one-click menus, semiprepared specials like “meals in four minutes,” and the ability to pull up prior grocery lists for fast reorders—all features that appeal to the time-strapped Manhattanites who were the firm’s first customers. (The Web’s not the only channel to reach customers—the firm’s mobile apps are hugely popular, especially for repeat orders.) [5] Next-day deliveries are from a vast warehouse the size of five football fields located in a lower-rent industrial area of Queens. At that size, the firm can offer a fresh goods selection that’s over five times larger than local supermarkets. Area shoppers—many of whom don’t have cars or are keen to avoid the traffic-snarled streets of the city—were quick to embrace the model. The service is now so popular that apartment buildings in New York have begun to redesign common areas to include secure freezers that can accept FreshDirect deliveries, even when customers aren’t there. [6]



The FreshDirect model crushes costs that plague traditional grocers. Worker shifts are highly efficient, avoiding the downtime lulls and busy rush hour spikes of storefronts. The result? Labor costs that are 60 percent lower than at traditional grocers. FreshDirect buys and prepares what it sells, leading to less waste, an advantage that the firm claims is “worth 5 percentage points of total revenue in terms of savings.” [7] Overall perishable inventory at FreshDirect turns 197 times a year versus 40 times a year at traditional grocers. [8] Higher inventory turns mean the firm is selling product faster, so it collects money quicker than its rivals do. And those goods are fresher since they’ve been in stock for less time, too. Consider that while the average grocer may have seven to nine days of seafood inventory, FreshDirect’s seafood stock turns each day. Stock is typically purchased direct from the docks in order to fulfill orders placed less than twenty-four hours earlier. [9]

Artificial intelligence software, coupled with some seven miles of fiber-optic cables linking systems and sensors, supports everything from baking the perfect baguette to verifying orders with 99.9 percent accuracy. [10] Since it lacks the money-sucking open-air refrigerators of the competition, the firm even saves big on energy (instead, staff bundle up for shifts in climate-controlled cold rooms tailored to the specific needs of dairy, deli, and produce). The firm also uses recycled biodiesel fuel to cut down on delivery costs.

FreshDirect buys directly from suppliers, eliminating middlemen wherever possible. The firm also offers suppliers several benefits beyond traditional grocers, all in exchange for more favorable terms. These include offering to carry a greater selection of supplier products while eliminating the “slotting fees” (payments by suppliers for prime shelf space) common in traditional retail, cobranding products to help establish and strengthen supplier brand, paying partners in days rather than weeks, and sharing data to help improve supplier sales and operations. Add all these advantages together and the firm’s



big, fresh selection is offered at prices that can undercut the competition by as much as 35 percent. ^[11] And FreshDirect does it all with margins in the range of 20 percent (to as high as 45 percent on many semiprepared meals), easily dwarfing the razor-thin 1 percent margins earned by traditional grocers. ^[12] Today, FreshDirect serves a base of some 600,000 paying customers. That's a population roughly the size of metro-Boston, serviced by a single grocer with no physical store. The privately held firm has been solidly profitable for several years. Even in recession-plagued 2009, the firm's CEO described 2009 earnings as "pretty spectacular," ^[13] while 2010 revenues were estimated at roughly \$300 million. ^[14]

Technology is critical to the FreshDirect model, but it's the collective impact of the firm's differences when compared to rivals, this tech-enabled strategic positioning, that delivers success. Operating for more than half a decade, the firm has also built up a set of strategic assets that not only address specific needs of a market but are now extremely difficult for any upstart to compete against. Traditional grocers can't fully copy the firm's delivery business because this would leave them straddling two markets (low-margin storefront and high-margin delivery), unable to gain optimal benefits from either. Entry costs for would-be competitors are also high (the firm spent over \$75 million building infrastructure before it could serve a single customer), and the firm's complex and highly customized software, which handles everything from delivery scheduling to orchestrating the preparation of thousands of recipes, continues to be refined and improved each year. ^[15] On top of all this comes years of customer data used to further refine processes, speed reorders, and make helpful recommendations. Competing against a firm with such a strong and tough-to-match strategic position can be brutal. Just five years after launch there were one-third fewer supermarkets in New York City than when FreshDirect first opened for business. ^[16]

But What Kinds of Differences?

The principles of operational effectiveness and strategic positioning are deceptively simple. But while Porter claims strategy is “fundamentally about being different,” ^[17] how can you recognize whether your firm’s differences are special enough to yield sustainable competitive advantage?

An approach known as the resource-based view of competitive advantage can help. The idea here is that if a firm is to maintain sustainable competitive advantage, it must control a set of exploitable resources that have four critical characteristics. These resources must be (1) *valuable*, (2) *rare*, (3) *imperfectly imitable* (tough to imitate), and (4) *nonsubstitutable*. Having all four characteristics is key. Miss value and no one cares what you’ve got. Without rareness, you don’t have something unique. If others can copy what you have, or others can replace it with a substitute, then any seemingly advantageous differences will be undercut.

Strategy isn’t just about recognizing opportunity and meeting demand. Resource-based thinking can help you avoid the trap of carelessly entering markets simply because growth is spotted. The telecommunications industry learned this lesson in a very hard and painful way. With the explosion of the Internet it was easy to see that demand to transport Web pages, e-mails, MP3s, video, and everything else you can turn into ones and zeros, was skyrocketing.

Most of what travels over the Internet is transferred over long-haul fiber-optic cables, so telecom firms began digging up the ground and laying webs of fiberglass to meet the growing demand. Problems resulted because firms laying long-haul fiber didn’t fully appreciate that their rivals and new upstart firms were doing the exact same thing. By one estimate there was enough fiber laid to stretch from the Earth to the moon some 280 times! ^[18] On top of that, a technology called dense wave division multiplexing (DWDM) enabled



existing fiber to carry more transmissions than ever before. The end result—these new assets weren't rare and each day they seemed to be less valuable. For some firms, the transmission prices they charged on newly laid cable collapsed by over 90 percent. Established firms struggled, upstarts went under, and WorldCom became the biggest bankruptcy in U.S. history. The impact was also felt throughout all industries that supplied the telecom industry. Firms like Sun, Lucent, and Nortel, whose sales growth relied on big sales to telecom carriers, saw their value tumble as orders dried up. Estimates suggest that the telecommunications industry lost nearly \$4 trillion in value in just three years, ^[19] much of it due to executives that placed big bets on resources that weren't strategic.

KEY TAKEAWAYS

- Technology can be easy to copy, and technology alone rarely offers sustainable advantage.
- Firms that leverage technology for strategic positioning use technology to create competitive assets or ways of doing business that are difficult for others to copy.
- True sustainable advantage comes from assets and business models that are simultaneously valuable, rare, difficult to imitate, and for which there are no substitutes.

QUESTIONS AND EXERCISES

1. What is operational effectiveness?
2. What is strategic positioning?
3. Is a firm that competes based on the features of technology engaged in operational effectiveness or strategic positioning? Give an example to back up your claim.

4. What is the “resource-based” view of competitive advantage? What are the characteristics of resources that may yield sustainable competitive advantage?
5. TiVo has a great brand. Why hasn’t it profitably dominated the market for digital video recorders?
6. Examine the FreshDirect business model and list reasons for its competitive advantage. Would a similar business work in your neighborhood? Why or why not?
7. What effect did FreshDirect have on traditional grocers operating in New York City? Why?
8. Choose a technology-based company. Discuss its competitive advantage based on the resources it controls.
9. Use the resource-based view of competitive advantage to explain the collapse of many telecommunications firms in the period following the burst of the dot-com bubble.
10. Consider the examples of Barnes and Noble competing with Amazon, and Apple offering iTunes. Are either (or both) of these efforts straddling? Why or why not?



2.2 Powerful Resources

LEARNING OBJECTIVES

1. Understand that technology is often critical to enabling competitive advantage, and provide examples of firms that have used technology to organize for sustained competitive advantage.



2. Understand the value chain concept and be able to examine and compare how various firms organize to bring products and services to market.
3. Recognize the role technology can play in crafting an imitation-resistant value chain, as well as when technology choice may render potentially strategic assets less effective.
4. Define the following concepts: brand, scale, data and switching cost assets, differentiation, network effects, and distribution channels.
5. Understand and provide examples of how technology can be used to create or strengthen the resources mentioned above.

Management has no magic bullets. There is no exhaustive list of key resources that firms can look to in order to build a sustainable business. And recognizing a resource doesn't mean a firm will be able to acquire it or exploit it forever. But being aware of major sources of competitive advantage can help managers recognize an organization's opportunities and vulnerabilities, and can help them brainstorm winning strategies. And these assets rarely exist in isolation. Oftentimes, a firm with an effective strategic position can create an arsenal of assets that reinforce one another, creating advantages that are particularly difficult for rivals to successfully challenge.

Imitation-Resistant Value Chains

While many of the resources below are considered in isolation, the strength of any advantage can be far more significant if firms are able to leverage several of these resources in a way that makes each stronger and makes the firm's way of doing business more difficult for rivals to match. Firms that craft an imitation-resistant value chain have developed a way of doing business that others will struggle to replicate, and in nearly every successful effort of this kind, technology plays a key enabling role. The value chain is the set of



interrelated activities that bring products or services to market (see below). When we compare FreshDirect's value chain to traditional rivals, there are differences across every element. But most importantly, the elements in FreshDirect's value chain work together to create and reinforce competitive advantages that others cannot easily copy. Incumbents trying to copy the firm would be *straddled* across two business models, unable to reap the full advantages of either. And late-moving pure-play rivals will struggle, as FreshDirect's lead time allows the firm to develop brand, scale, data, and other advantages that newcomers lack (see below for more on these resources).

Key Framework: The Value Chain

The *value chain* is the “set of activities through which a product or service is created and delivered to customers.” ^[1] There are five primary components of the value chain and four supporting components. The primary components are:

- *Inbound logistics*—getting needed materials and other inputs into the firm from suppliers
- *Operations*—turning inputs into products or services
- *Outbound logistics*—delivering products or services to consumers, distribution centers, retailers, or other partners
- *Marketing and sales*—customer engagement, pricing, promotion, and transaction
- *Support*—service, maintenance, and customer support

The secondary components are:

- *Firm infrastructure*—functions that support the whole firm, including general management, planning, IS, and finance
- *Human resource management*—recruiting, hiring, training, and development
- *Technology / research and development*—new product and process design
- *Procurement*—sourcing and purchasing functions



While the value chain is typically depicted as it's displayed in the figure below, goods and information don't necessarily flow in a line from one function to another. For example, an order taken by the marketing function can trigger an inbound logistics function to get components from a supplier, operations functions (to build a product if it's not available), or outbound logistics functions (to ship a product when it's available). Similarly, information from service support can be fed back to advise research and development (R&D) in the design of future products.

Figure 2.2 The Value Chain



When a firm has an imitation-resistant value chain—one that's tough for rivals to copy in a way that yields similar benefits—then a firm may have a critical competitive asset. From a strategic perspective, managers can use the value chain framework to consider a firm's differences and distinctiveness compared to rivals. If a firm's value chain can't be copied by competitors without engaging in painful trade-offs, or if the firm's value chain helps to create and strengthen other strategic assets over time, it can be a key source for competitive advantage. Many of the examples used in this book, including FreshDirect, Amazon, and Zara, illustrate this point.

An analysis of a firm's value chain can also reveal operational weaknesses, and technology is often of great benefit to improving the speed and quality of execution. Firms can often buy software to improve things, and tools such as *supply chain management* (SCM; linking inbound and outbound logistics

with operations), *customer relationship management* (CRM; supporting sales, marketing, and in some cases R&D), and *enterprise resource planning* software (ERP; software implemented in modules to automate the entire value chain), can have a big impact on more efficiently integrating the activities within the firm, as well as with its suppliers and customers. But remember, these software tools can be purchased by competitors, too. While valuable, such software may not yield lasting competitive advantage if it can be easily matched by competitors as well.

There's potential danger here. If a firm adopts software that changes a unique process into a generic one, it may have co-opted a key source of competitive advantage particularly if other firms can buy the same stuff. This isn't a problem with something like accounting software. Accounting processes are standardized and accounting isn't a source of competitive advantage, so most firms buy rather than build their own accounting software. But using packaged, third-party SCM, CRM, and ERP software typically requires adopting a very specific way of doing things, using software and methods that can be purchased and adopted by others. During its period of PC-industry dominance, Dell stopped deployment of the logistics and manufacturing modules of a packaged ERP implementation when it realized that the software would require the firm to make changes to its unique and highly successful operating model and that many of the firm's unique supply chain advantages would change to the point where the firm was doing the same thing using the same software as its competitors. By contrast, Apple had no problem adopting third-party ERP software because the firm competes on product uniqueness rather than operational differences.

Dell's Struggles: Nothing Lasts Forever

Michael Dell enjoyed an extended run that took him from assembling PCs in his dorm room as an undergraduate at the University of Texas at Austin to



heading the largest PC firm on the planet. For years Dell's superefficient, vertically integrated manufacturing and direct-to-consumer model combined to help the firm earn seven times more profit on its own systems when compared with comparably configured rival PCs. [2] And since Dell PCs were usually cheaper, too, the firm could often start a price war and still have better overall margins than rivals.

It was a brilliant model that for years proved resistant to imitation. While Dell sold direct to consumers, rivals had to share a cut of sales with the less efficient retail chains responsible for the majority of their sales. Dell's rivals struggled in moving toward direct sales because any retailer sensing its suppliers were competing with it through a direct-sales effort could easily choose another supplier that sold a nearly identical product. It wasn't that HP, IBM, Sony, and so many others didn't see the advantage of Dell's model—these firms were wedded to models that made it difficult for them to imitate their rival without the inefficient burden of straddling two different models of doing business.

But then Dell's killer model, one that had become a staple case study in business schools worldwide, began to lose steam. Nearly two decades of observing Dell had allowed the contract manufacturers serving Dell's rivals to improve manufacturing efficiency. [3] Component suppliers located near contract manufacturers, and assembly times fell dramatically. And as the cost of computing fell, the price advantage Dell enjoyed over rivals also shrank in absolute terms. That meant savings from buying a Dell weren't as big as they once were. On top of that, the direct-to-consumer model also suffered when sales of notebook PCs outpaced the more commoditized desktop market. Notebooks can be considered to be more differentiated than desktops, and customers often want to compare products in person—lift them, type on keyboards, and view screens—before making a purchase decision.

In time, these shifts created an opportunity for rivals to knock Dell from its ranking as the world's number one PC manufacturer. Dell has even abandoned its direct-only business model and now also sells products through third-party brick-and-mortar retailers. Dell's struggles as computers, customers, and the product mix changed all underscore the importance of continually assessing a firm's strategic position among changing market conditions. There is no guarantee that today's winning strategy will dominate forever.

Brand

A firm's brand is the symbolic embodiment of all the information connected with a product or service, and a strong brand can also be an exceptionally powerful resource for competitive advantage. Consumers use brands to *lower search costs*, so having a strong brand is particularly vital for firms hoping to be the first online stop for consumers. Want to buy a book online? Auction a product? Search for information? Which firm would you visit first? Almost certainly Amazon, eBay, or Google. But how do you build a strong brand? It's *not* just about advertising and promotion. First and foremost, customer experience counts. A strong brand *proxies quality* and *inspires trust*, so if consumers can't rely on a firm to deliver as promised, they'll go elsewhere. As an upside, tech can play a critical role in rapidly and cost-effectively strengthening a brand. If a firm performs well, consumers can often be enlisted to promote a product or service (so-called viral marketing). Consider that while scores of dot-coms burned through money on Super Bowl ads and other costly promotional efforts, Google, Hotmail, Skype, eBay, Facebook, LinkedIn, Twitter, YouTube, and so many other dominant online properties built multimillion member followings before committing any significant spending to advertising.

Early customer accolades for a novel service often mean that positive press (a kind of free advertising) will also likely follow.

But show up late and you may end up paying much more to counter an incumbent's place in the consumer psyche. In recent years, Amazon has spent no money on television advertising, while rivals Buy.com and Overstock.com spent millions. Google, another strong brand, has become a verb, and the cost to challenge it is astonishingly high. Yahoo! and Microsoft's Bing each spent \$100 million on Google-challenging branding campaigns, but the early results of these efforts seemed to do little to grow share at Google's expense. [4] Branding is difficult, but if done well, even complex tech products can establish themselves as killer brands. Consider that Intel has taken an ingredient product that most people don't understand, the microprocessor, and built a quality-conveying name recognized by computer users worldwide.

Scale

Many firms gain advantages as they grow in size. Advantages related to a firm's size are referred to as scale advantages. Businesses benefit from economies of scale when the cost of an investment can be spread across increasing units of production or in serving a growing customer base. Firms that benefit from scale economies as they grow are sometimes referred to as being *scalable*. Many Internet and tech-leveraging businesses are highly scalable since, as firms grow to serve more customers with their existing infrastructure investment, profit margins improve dramatically.

Consider that in just one year, the Internet firm BlueNile sold as many diamond rings with just 115 employees and one Web site as a traditional jewelry retailer would sell through 116 stores. [5] And with lower operating costs, BlueNile can sell at prices that brick-and-mortar stores can't match, thereby attracting more customers and further fueling its scale advantages.



Profit margins improve as the cost to run the firm's single Web site and operate its one warehouse is spread across increasing jewelry sales.

A growing firm may also gain *bargaining power with its suppliers or buyers*. Apple's dominance of smartphone and tablet markets has allowed the firm to lock up 60 percent of the world's supply of advanced touch-screen displays, and to do so with better pricing than would be available to smaller rivals. [6] Similarly, for years eBay could raise auction fees because of the firm's market dominance. Auction sellers who left eBay lost pricing power since fewer bidders on smaller, rival services meant lower prices.

The scale of technology investment required to run a business can also act as a barrier to entry, discouraging new, smaller competitors. Intel's size allows the firm to pioneer cutting-edge manufacturing techniques and invest \$7 billion on next-generation plants. [7] And although Google was started by two Stanford students with borrowed computer equipment running in a dorm room, the firm today runs on an estimated 1.4 million servers. [8] The investments being made by Intel and Google would be cost-prohibitive for almost any newcomer to justify.

Switching Costs and Data

Switching costs exist when consumers incur an expense to move from one product or service to another. Tech firms often benefit from strong switching costs that cement customers to their firms. Users invest their time learning a product, entering data into a system, creating files, and buying supporting programs or manuals. These investments may make them reluctant to switch to a rival's effort.

Similarly, firms that seem dominant but that don't have high switching costs can be rapidly trumped by strong rivals. Netscape once controlled more than 80 percent of the market share in Web browsers, but when Microsoft began bundling Internet Explorer with the Windows operating system and (through

an alliance) with America Online (AOL), Netscape's market share plummeted. Customers migrated with a mouse click as part of an upgrade or installation. Learning a new browser was a breeze, and with the Web's open standards, most customers noticed no difference when visiting their favorite Web sites with their new browser.

Sources of Switching Costs

- Learning costs: Switching technologies may require an investment in learning a new interface and commands.
- Information and data: Users may have to reenter data, convert files or databases, or may even lose earlier contributions on incompatible systems.
- Financial commitment: Can include investments in new equipment, the cost to acquire any new software, consulting, or expertise, and the devaluation of any investment in prior technologies no longer used.
- Contractual commitments: Breaking contracts can lead to compensatory damages and harm an organization's reputation as a reliable partner.
- Search costs: Finding and evaluating a new alternative costs time and money.
- Loyalty programs: Switching can cause customers to lose out on program benefits. Think frequent purchaser programs that offer "miles" or "points" (all enabled and driven by software). [9]

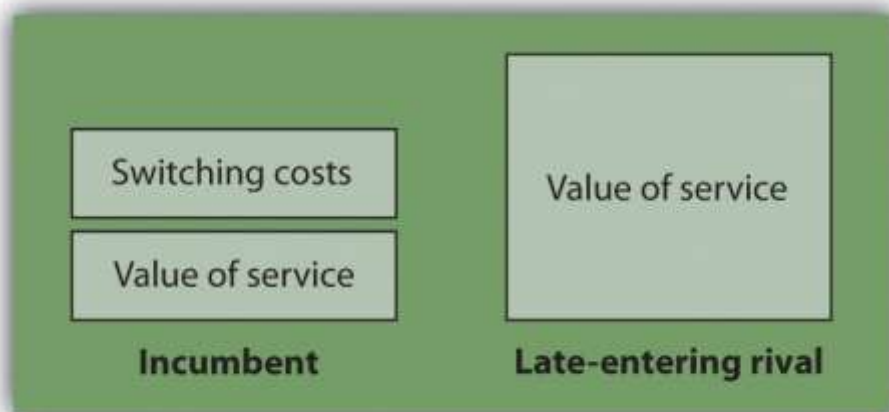
It is critical for challengers to realize that in order to win customers away from a rival, a new entrant must not only demonstrate to consumers that an offering provides more value than the incumbent, they have to ensure that their value added exceeds the incumbent's value *plus* any perceived customer switching costs (see [Figure 2.4](#)). If it's going to cost you and be inconvenient, there's no way you're going to leave unless the benefits are overwhelming.

Data can be a particularly strong switching cost for firms leveraging technology. A customer who enters her profile into Facebook, movie preferences into Netflix, or grocery list into FreshDirect may be unwilling to



try rivals—even if these firms are cheaper or offer more features—if moving to the new firm means she’ll lose information feeds, recommendations, and time savings provided by the firms that already know her well. Fueled by scale over time, firms that have more customers and have been in business longer can gather more data, and many can use this data to improve their value chain by offering more accurate demand forecasting or product recommendations.

Figure 2.4



In order to win customers from an established incumbent, a late-entering rival must offer a product or service that not only exceeds the value offered by the incumbent but also exceeds the incumbent’s value and any customer switching costs.


Competing on Tech Alone Is Tough: Gmail versus Rivals

Switching e-mail services can be a real a pain. You’ve got to convince your contacts to update their address books, hope that any message-forwarding from your old service to your new one remains active and works properly, and regularly check the old service to be sure nothing is caught in junk folder purgatory. Not fun. So when Google entered the market for free e-mail,

challenging established rivals Yahoo! and Microsoft Hotmail, it knew it needed to offer an overwhelming advantage to lure away customers who had used these other services for years. Google's offering? A mailbox with vastly more storage than its competitors. With 250 to 500 times the capacity of rivals, Gmail users were liberated from the infamous "mailbox full" error, and could send photos, songs, slideshows, and other rich media files as attachments.

A neat innovation, but one based on technology that incumbents could easily copy. Once Yahoo! and Microsoft saw that customers valued the increased capacity, they quickly increased their own mailbox size, holding on to customers who might otherwise have fled to Google. Four years after Gmail was introduced, the service still had less than half the users of each of its two biggest rivals.

Figure 2.5 *E-mail Market Share in Millions of Users* [10]



Differentiation

Commodities are products or services that are nearly identically offered from multiple vendors. Consumers buying commodities are highly price-focused since they have so many similar choices. In order to break the commodity trap, many firms leverage technology to *differentiate* their goods and services. Dell gained attention from customers not only because of its low prices, but also because it was one of the first PC vendors to build computers based on customer choice. Want a bigger hard drive? Don't need the fast graphics card? Dell will oblige.

Data is not only a switching cost, it also plays a critical role in differentiation. Each time a visitor returns to Amazon, the firm uses browsing records, purchase patterns, and product ratings to present a custom home page featuring products that the firm hopes the visitor will like. Customers value the experience they receive at Amazon so much that the firm received the highest score ever recorded on the University of Michigan's American Customer Satisfaction Index (ACSI). The score was not just the highest performance of any online firm, it was the highest ranking that any service firm in any industry had ever received.

Capital One has also used data to differentiate its offerings. The firm mines data and runs experiments to create risk models on potential customers. Because of this, the credit card firm aggressively pursued a set of customers that other lenders considered too risky based on simplistic credit scoring. Technology determined that a subset of underserved customers was not properly identified by conventional techniques and was actually a good bet. Finding profitable new markets that others ignored allowed Capital One to grow its EPS (earnings per share) 20 percent a year for seven years, a feat matched by less than 1 percent of public firms. ^[11]

Network Effects

Facebook is by far the most dominant social network worldwide. Microsoft Windows has a 90 percent market share in operating systems. EBay has an 80 percent share of online auctions. Why are these firms so dominant? Largely due to the concept of network effects (see [Chapter 6 "Understanding Network Effects"](#)). Network effects (sometimes called *network externalities* or *Metcalfe's Law*) exist when a product or service becomes more valuable as more people use it. If you're the first person with a Facebook account, then Facebook isn't very valuable. But with each additional user, there's one more person to communicate with. A firm with a big network of users might also see value added by third parties. Apple's iOS devices (the iPhone, iPod touch, and iPad) and Google's Android dominate rivals from Microsoft and HP in part because Apple and Google have tens of thousands more apps that run on and enhance these devices, and most of these apps are provided by firms other than Apple and Google. Third-party add-on products, books, magazines, or even skilled labor are all attracted to networks of the largest number of users, making dominant products even more valuable. Switching costs also play a role in determining the strength of network effects. Tech user investments often go far beyond simply the cost of acquiring a technology. Users spend time learning a product; they buy add-ons, create files, and enter preferences. Because no one wants to be stranded with an abandoned product and lose this additional investment, users may choose a technically inferior product simply because the product has a larger user base and is perceived as having a greater chance of being offered in the future. The virtuous cycle of network effects ^[12] doesn't apply to all tech products, and it can be a particularly strong asset for firms that can control and leverage a leading standard (think Apple's iPhone and iPad with their closed systems



versus the once-dominant but now rarely used Netscape browser, which was almost entirely based on open standards), but in some cases where network effects are significant, they can create winners so dominant that firms with these advantages enjoy a near-monopoly hold on a market.

Distribution Channels

If no one sees your product, then it won't even get considered by consumers. So distribution channels—the path through which products or services get to customers—can be critical to a firm's success. Again, technology opens up opportunities for new ways to reach customers.

Users can be recruited to create new distribution channels for your products and services (usually for a cut of the take). You may have visited Web sites that promote books sold on Amazon.com. Web site operators do this because Amazon gives them a percentage of all purchases that come in through these links. Amazon now has over 1 million of these “associates” (the term the firm uses for its affiliates), yet it only pays them if a promotion gains a sale. Google similarly receives some 30 percent of its ad revenue not from search ads, but from advertisements distributed within third-party sites ranging from lowly blogs to the *New York Times*. ^[13]

In recent years, Google and Microsoft have engaged in bidding wars, trying to lock up distribution deals that would bundle software tools, advertising, or search capabilities with key partner offerings. Deals with partners such as Dell, Nokia, and Verizon Wireless have been valued at up to \$1 billion each. ^[14] The ability to distribute products by bundling them with existing offerings is a key Microsoft advantage. But beware—sometimes these distribution channels can provide firms with such an edge that international regulators have stepped in to try to provide a more level playing field. Microsoft was forced by European regulators to unbundle the Windows Media Player, for fear that it



provided the firm with too great an advantage when competing with the likes of RealPlayer and Apple's QuickTime (see [Chapter 6 "Understanding Network Effects"](#)).

What about Patents?

Intellectual property protection can be granted in the form of a patent for those innovations deemed to be useful, novel, and nonobvious. In the United States, technology and (more controversially) even business models can be patented, typically for periods of twenty years from the date of patent application. Firms that receive patents have some degree of protection from copycats that try to identically mimic their products and methods.

The patent system is often considered to be unfairly stacked against start-ups. U.S. litigation costs in a single patent case average about \$5 million, ^[15] and a few months of patent litigation can be enough to sink an early stage firm.

Large firms can also be victims. So-called patent trolls hold intellectual property not with the goal of bringing novel innovations to market but instead in hopes that they can sue or extort large settlements from others. BlackBerry maker Research in Motion's \$612 million settlement with the little-known holding company NTP is often highlighted as an example of the pain trolls can inflict. ^[16] Litigation threats are pushing rivals to cooperate in patent portfolio acquisition. Apple, EMC, Ericsson, Microsoft, RIM, and Sony pooled resources for a \$4.5 billion purchase of some 6,000 patents formally held by the bankrupt telecom equipment firm Nortel. ^[17]

Even if an innovation is patentable, that doesn't mean that a firm has bulletproof protection. Some patents have been nullified by the courts upon later review (usually because of a successful challenge to the uniqueness of the innovation). Software patents are also widely granted, but notoriously difficult to defend. In many cases, coders at competing firms can write substitute

algorithms that aren't the same, but accomplish similar tasks. For example, although Google's PageRank search algorithms are fast and efficient, Microsoft, Yahoo! and others now offer their own noninfringing search that presents results with an accuracy that many would consider on par with PageRank. Patents do protect tech-enabled operations innovations at firms like Netflix and Caesars Entertainment Corporation (formerly known as Harrah's), and design innovations like the iPod click wheel. But in a study of the factors that were critical in enabling firms to profit from their innovations, Carnegie Mellon professor Wes Cohen found that patents were only the fifth most important factor. Secrecy, lead time, sales skills, and manufacturing all ranked higher. [18]

KEY TAKEAWAYS

- Technology can play a key role in creating and reinforcing assets for sustainable advantage by enabling an imitation-resistant value chain; strengthening a firm's brand; collecting useful data and establishing switching costs; creating a network effect; creating or enhancing a firm's scale advantage; enabling product or service differentiation; and offering an opportunity to leverage unique distribution channels.
- The value chain can be used to map a firm's efficiency and to benchmark it against rivals, revealing opportunities to use technology to improve processes and procedures. When a firm is resistant to imitation, a superior value chain may yield sustainable competitive advantage.
- Firms may consider adopting packaged software or outsourcing value chain tasks that are not critical to a firm's competitive advantage. A firm should be wary of adopting software packages or outsourcing portions of its value chain that are proprietary and a source of competitive advantage.

- Patents are not necessarily a sure-fire path to exploiting an innovation. Many technologies and business methods can be copied, so managers should think about creating assets like the ones previously discussed if they wish to create truly sustainable advantage.
- Nothing lasts forever, and shifting technologies and market conditions can render once strong assets as obsolete.

QUESTIONS AND EXERCISES

1. Define and diagram the value chain.
2. Discuss the elements of FreshDirect's value chain and the technologies that FreshDirect uses to give the firm a competitive advantage. Why is FreshDirect resistant to imitation from incumbent firms? What advantages does FreshDirect have that insulate the firm from serious competition from start-ups copying its model?
3. Which firm should adopt third-party software to automate its supply chain—Dell or Apple? Why? Identify another firm that might be at risk if it adopted generic enterprise software. Why do you think this is risky and what would you recommend as an alternative?
4. Identify two firms in the same industry that have different value chains. Why do you think these firms have different value chains? What role do you think technology plays in the way that each firm competes? Do these differences enable strategic positioning? Why or why not?
5. How can information technology help a firm build a brand inexpensively?
6. Describe BlueNile's advantages over a traditional jewelry chain. Can conventional jewelers successfully copy BlueNile? Why or why not?
7. What are switching costs? What role does technology play in strengthening a firm's switching costs?
8. In most markets worldwide, Google dominates search. Why hasn't Google shown similar dominance in e-mail, as well?

9. How can technology be a distribution channel? Name a firm that has tried to leverage its technology as a distribution channel.
10. Do you think it is possible to use information technology to achieve competitive advantage? If so, how? If not, why not?
11. What are network effects? Name a product or service that has been able to leverage network effects to its advantage.
12. For well over a decade, Dell earned above average industry profits. But lately the firm has begun to struggle. What changed?
13. What are the potential sources of switching costs if you decide to switch cell phone service providers? Cell phones? Operating systems? PayTV service?
14. Why is an innovation based on technology alone often subjected to intense competition?
15. Can you think of firms that have successfully created competitive advantage even though other firms provide essentially the same thing? What factors enable this success?
16. What role did network effects play in your choice of an operating system? Of a social network? Of a word processor? Of a mobile phone? Why do so many firms choose to standardize on Microsoft Windows for PCs and laptops?
17. What can a firm do to prepare for the *inevitable* expiration of a patent (patents typically expire after twenty years)? Think in terms of the utilization of other assets and the development of advantages through employment of technology.



2.3 Barriers to Entry, Technology, and Timing

LEARNING OBJECTIVES

1. Understand the relationship between timing, technology, and the creation of resources for competitive advantage.
2. Argue effectively when faced with broad generalizations about the importance (or lack of importance) of technology and timing to competitive advantage.
3. Recognize the difference between low barriers to entry and the prospects for the sustainability of new entrant's efforts.

Some have correctly argued that the barriers to entry for many tech-centric businesses are low. This argument is particularly true for the Internet where rivals can put up a competing Web site or deploy a rival app seemingly overnight. But it's absolutely critical to understand that market entry is *not* the same as building a sustainable business and just showing up doesn't guarantee survival.

Platitudes like “follow, don't lead” ^[1] can put firms dangerously at risk, and statements about low entry barriers ignore the difficulty many firms will have in matching the competitive advantages of successful tech pioneers. Should Blockbuster have waited while Netflix pioneered? In a year where Netflix profits were up sevenfold, Blockbuster lost more than \$1 billion, and today Blockbuster is bankrupt. ^[2] Should Sotheby's have dismissed seemingly inferior eBay? Sotheby's made \$171 million in 2011, but eBay earned over \$3.2 billion. Barnes & Noble waited seventeen months to respond to Amazon.com. 2011 was a down year for Amazon due to investments in the Kindle and other new businesses, but the firm still earned \$630 million in profits, while Barnes



& Noble lost over \$73 million. Today Amazon sports a market cap roughly 100 times greater than its one-time rival, and it has expanded far beyond being just a books and media business. [3] During that same year, Borders—once the nation’s number two book chain—declared bankruptcy and closed all of its stores. Today’s Internet giants are winners because in most cases, they were the first to move with a profitable model and they were able to quickly establish resources for competitive advantage. With few exceptions, established offline firms have failed to catch up to today’s Internet leaders.

Table 2.1 A Tale of Two Firms

	2007	2008	2009	2010
Amazon	\$476 million	\$645 million	\$902 million	\$1,152 million
Barnes & Noble	\$150 million	\$135 million	\$75 million	\$36 million

Barnes & Noble saw net income cut in half from 2007 to 2009 then fall half again in 2010. Over the same period Amazon’s profits are up nearly threefold—in a recession.

Timing and technology alone will not yield sustainable competitive advantage. Yet both of these can be *enablers* for competitive advantage. Put simply, it’s not the time lead or the technology; it’s what a firm *does* with its time lead and technology. True strategic positioning means that a firm has created differences that cannot be easily matched by rivals. Moving first pays off when the time lead is used to create critical resources that are valuable, rare, tough to imitate, and lack substitutes. Anything less risks the arms race of operational effectiveness. Build resources like brand, scale, network effects, switching costs, or other key assets and your firm may have a shot. But guess



wrong about the market or screw up execution and failure or direct competition awaits. It is true that most tech can be copied—there’s little magic in eBay’s servers, Intel’s processors, Oracle’s database software, or Microsoft’s operating systems that past rivals have not at one point improved upon. But the lead that each of these tech-enabled firms had was leveraged to create network effects, switching costs, data assets, and helped build solid and well-respected brands.

But Google Arrived Late! Why Incumbents Must Constantly Consider Rivals

Although its share is slowly eroding, Yahoo! has been able to hold onto its lead in e-mail for so long because the firm quickly matched and nullified Gmail’s most significant tech-based innovations before Google could inflict real damage. Perhaps Yahoo! had learned from prior errors. The firm’s earlier failure to respond to Google’s emergence as a credible threat in search advertising gave Sergey Brin and Larry Page the time they needed to build the planet’s most profitable Internet firm.

Yahoo! (and many Wall Street analysts) saw search as a commodity—a service the firm had subcontracted out to other firms including Alta Vista and Inktomi. Yahoo! saw no conflict in taking an early investment stake in Google or in using the firm for its search results. But Yahoo! failed to pay attention to Google’s advance. As Google’s innovations in technology and interface remained unmatched over time, this allowed the firm to build its brand, scale, and advertising network (distribution channel) that grew from network effects because content providers and advertisers attract one another. These are all competitive resources that rivals have never been able to match.

Now Google (and Apple, too) are once again running from this playbook—turning the smartphone software market into what increasingly looks like a

two-horse race. Many rivals, including Microsoft, had been trying to create a mobile standard for years, but their technical innovations offered little durable strategic value. It wasn't until app stores flourished, offered with a high-quality user experience, that dominant smartphone platforms emerged. Yes, Google and Apple arrived late, but nothing before them had created defensible strategic assets, and that left an opening.

Google's ability to succeed after being late to the search and mobile party isn't a sign of the power of the late mover; it's a story about the failure of incumbents to monitor their competitive landscape, recognize new rivals, and react to challenging offerings. That doesn't mean that incumbents need to respond to every potential threat. Indeed, figuring out which threats are worthy of response is the real skill here. Video rental chain Hollywood Video wasted over \$300 million in an Internet streaming business years before high-speed broadband was available to make the effort work. [4] But while Blockbuster avoided the balance sheet-cratering gaffes of Hollywood Video, the firm also failed to respond to Netflix—a new threat that had timed market entry perfectly (see [Chapter 4 "Netflix in Two Acts: The Making of an E-commerce Giant and the Uncertain Future of Atoms to Bits"](#)).

Firms that quickly get to market with the “right” model can dominate, but it's equally critical for leading firms to pay close attention to competition and innovate in ways that customers value. Take your eye off the ball and rivals may use time and technology to create strategic resources. Just look at Friendster—a firm that was once known as the largest social network in the United States but has become virtually irrelevant today.

KEY TAKEAWAYS

- It doesn't matter if it's easy for new firms to enter a market if these newcomers can't create and leverage the assets needed to challenge incumbents.
- Beware of those who say, "IT doesn't matter" or refer to the "myth" of the first mover. This thinking is overly simplistic. It's not a time or technology lead that provides sustainable competitive advantage; it's what a firm does with its time and technology lead. If a firm can use a time and technology lead to create valuable assets that others cannot match, it may be able to sustain its advantage. But if the work done in this time and technology lead can be easily matched, then no advantage can be achieved, and a firm may be threatened by new entrants

QUESTIONS AND EXERCISES

1. Does technology lower barriers to entry or raise them? Do low entry barriers necessarily mean that a firm is threatened?
2. Is there such a thing as the first-mover advantage? Why or why not?
3. Why did Google beat Yahoo! in search?
4. A former editor of the *Harvard Business Review*, Nick Carr, once published an article in that same magazine with the title "IT Doesn't Matter." In the article he also offered firms the advice: "Follow, Don't Lead." What would you tell Carr to help him improve the way he thinks about the relationship between time, technology, and competitive advantage?
5. Name an early mover that has successfully defended its position. Name another that had been superseded by the competition. What factors contributed to its success or failure?
6. You have just written a word processing package far superior in features to Microsoft Word. You now wish to form a company to market it. List and discuss the barriers your start-up faces.



7. What kinds of strategic assets are Google’s Android and Apple’s iOS seeking to create and exploit? Do you think these firms will be more successful than rivals? Why or why not?

2.4 Key Framework: The Five Forces of Industry Competitive Advantage

LEARNING OBJECTIVES

1. Diagram the five forces of competitive advantage.
2. Apply the framework to an industry, assessing the competitive landscape and the role of technology in influencing the relative power of buyers, suppliers, competitors, and alternatives.

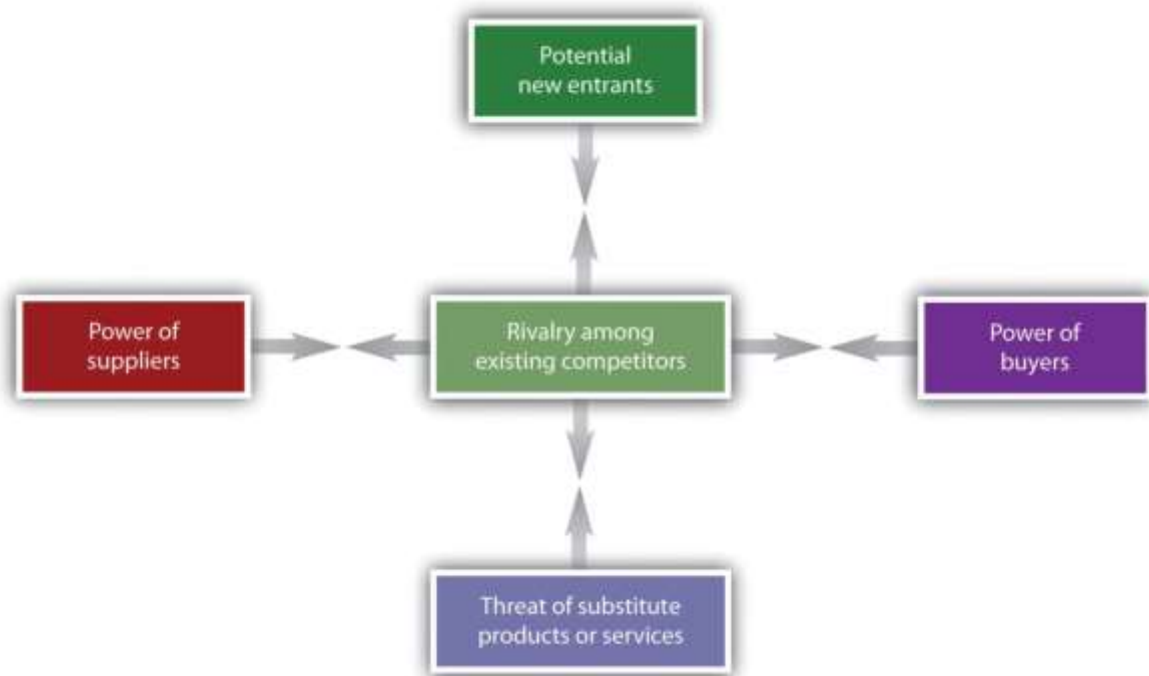
Professor and strategy consultant Gary Hamel once wrote in a *Fortune* cover story that “the dirty little secret of the strategy industry is that it doesn’t have any theory of strategy creation.” ^[1] While there is no silver bullet for strategy



creation, strategic frameworks help managers describe the competitive environment a firm is facing. Frameworks can also be used as brainstorming tools to generate new ideas for responding to industry competition. If you have a model for thinking about competition, it's easier to understand what's happening and to think creatively about possible solutions.

One of the most popular frameworks for examining a firm's competitive environment is Porter's five forces, also known as the *Industry and Competitive Analysis*. As Porter puts it, "analyzing [these] forces illuminates an industry's fundamental attractiveness, exposes the underlying drivers of average industry profitability, and provides insight into how profitability will evolve in the future." The five forces this framework considers are (1) the intensity of rivalry among existing competitors, (2) the threat of new entrants, (3) the threat of substitute goods or services, (4) the bargaining power of buyers, and (5) the bargaining power of suppliers (see [Figure 2.6 "The Five Forces of Industry and Competitive Analysis"](#)).

Figure 2.6 The Five Forces of Industry and Competitive Analysis



New technologies can create jarring shocks in an industry. Consider how the rise of the Internet has impacted the five forces for music retailers. Traditional music retailers like Tower and Virgin found that customers were seeking music online. These firms scrambled to invest in the new channel out of what is perceived to be a necessity. Their *intensity of rivalry* increases because they not only compete based on the geography of where brick-and-mortar stores are physically located, they now compete online as well. Investments online are expensive and uncertain, prompting some firms to partner with *new entrants* such as Amazon. Free from brick-and-mortar stores, Amazon, the dominant new entrant, has a highly scalable cost structure. And in many ways the online buying experience is superior to what customers saw in stores. Customers can hear samples of almost all tracks, selection is seemingly limitless (the *long tail* phenomenon—see this concept illuminated in [Chapter 4 "Netflix in Two Acts: The Making of an E-commerce Giant and the Uncertain Future of Atoms to Bits"](#)), and data is leveraged using *collaborative filtering* software to make product recommendations and assist in music

discovery. [2] Tough competition, but it gets worse because CD sales aren't the only way to consume music. The process of buying a plastic disc now faces *substitutes* as digital music files become available on commercial music sites. Who needs the physical atoms of a CD filled with ones and zeros when you can buy the bits one song at a time? Or don't buy anything and subscribe to a limitless library instead.

From a sound quality perspective, the *substitute good* of digital tracks purchased online is almost always inferior to their CD counterparts. To transfer songs quickly and hold more songs on a digital music player, tracks are encoded in a smaller file size than what you'd get on a CD, and this smaller file contains lower playback fidelity. But the additional tech-based market shock brought on by digital music players (particularly the iPod) has changed listening habits. The convenience of carrying thousands of songs trumps what most consider just a slight quality degradation. iTunes is now responsible for selling more music than any other firm, online or off. Apple can't rest on its laurels, either. The constant disruption of tech-enabled substitute goods has created rivals like Spotify and Pandora that don't sell music at all, yet have garnered tens of millions of users across desktop and mobile platforms. Most alarming to the industry is the other widely adopted substitute for CD purchases—theft. Illegal music “sharing” services abound, even after years of record industry crackdowns. And while exact figures on real losses from online piracy are in dispute, the music industry has seen album sales drop by 45 percent in less than a decade. [3] All this choice gives consumers (buyers) *bargaining power*. They demand cheaper prices and greater convenience. The *bargaining power of suppliers*—the music labels and artists—also increases. At the start of the Internet revolution, retailers could pressure labels to limit sales through competing channels. Now, with many of the major music retail chains in bankruptcy, labels have a freer hand to

experiment, while bands large and small have new ways to reach fans, sometimes in ways that entirely bypass the traditional music labels. While it can be useful to look at changes in one industry as a model for potential change in another, it's important to realize that the changes that impact one industry do not necessarily impact other industries in the same way. For example, it is often suggested that the Internet increases bargaining power of buyers and lowers the bargaining power of suppliers. This suggestion is true for some industries like auto sales and jewelry where the products are commodities and the price transparency of the Internet counteracts a previous information asymmetry where customers often didn't know enough information about a product to bargain effectively. But it's not true across the board.

In cases where network effects are strong or a seller's goods are highly differentiated, the Internet can strengthen supplier bargaining power. The customer base of an antique dealer used to be limited by how many likely purchasers lived within driving distance of a store. Now with eBay, the dealer can take a rare good to a global audience and have a much larger customer base bid up the price. Switching costs also weaken buyer bargaining power. Wells Fargo has found that customers who use online bill pay (where switching costs are high) are 70 percent less likely to leave the bank than those who don't, suggesting that these switching costs help cement customers to the company even when rivals offer more compelling rates or services.

Tech plays a significant role in shaping and reshaping these five forces, but it's not the only significant force that can create an industry shock. Government deregulation or intervention, political shock, and social and demographic changes can all play a role in altering the competitive landscape. Because we live in an age of constant and relentless change, managers need to continually visit strategic frameworks to consider any market-impacting shifts. Predicting the future is difficult, but ignoring change can be catastrophic.

KEY TAKEAWAYS

- Industry competition and attractiveness can be described by considering the following five forces: (1) the intensity of rivalry among existing competitors, (2) the potential for new entrants to challenge incumbents, (3) the threat posed by substitute products or services, (4) the power of buyers, and (5) the power of suppliers.
- In markets where commodity products are sold, the Internet can increase buyer power by increasing price transparency.
- The more differentiated and valuable an offering, the more the Internet shifts bargaining power to sellers. Highly differentiated sellers that can advertise their products to a wider customer base can demand higher prices.
- A strategist must constantly refer to models that describe events impacting their industry, particularly as new technologies emerge.

QUESTIONS AND EXERCISES

1. What are Porter's "five forces"?
2. Use the five forces model to illustrate competition in the newspaper industry. Are some competitors better positioned to withstand this environment than others? Why or why not? What role do technology and resources for competitive advantage play in shaping industry competition?
3. What is price transparency? What is information asymmetry? How does the Internet relate to these two concepts? How does the Internet shift bargaining power among the five forces?
4. How has the rise of the Internet impacted each of the five forces for music retailers?
5. In what ways is the online music buying experience superior to that of buying in stores?



6. What is the *substitute* for music CDs? What is the comparative sound quality of the substitute? Why would a listener accept an inferior product?
7. Based on Porter's five forces, is this a good time to enter the retail music industry? Why or why not?
8. What is the cost to the music industry of music theft? Cite your source.
9. Discuss the concepts of price transparency and information asymmetry as they apply to the diamond industry as a result of the entry of BlueNile. Name another industry where the Internet has had a similar impact.
10. Under what conditions can the Internet strengthen supplier bargaining power? Give an example.
11. What is the effect of switching costs on buyer bargaining power? Give an example.
12. How does the Internet impact bargaining power for providers of rare or highly differentiated goods? Why?

2.5 Key Framework: The Value Chain

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand the value chain concept.
2. Diagram and compare how various firms organize to bring products and services to market; recognizing components that provide operational effectiveness or strategic positioning.
3. Recognize when a firm might seek to purchase a third-party solution to automate its value chain, and when it may want to keep some functions in-house.

The value chain is the “set of activities through which a product or service is created and delivered to customers.” ^[1] By examining the activities in a

firm's value chain, managers are able to gain a greater understanding of how these factors influence a firm's cost structure and value delivery. There are five primary components of the value chain and four supporting components. The primary components are as follows:

- *Inbound logistics*—getting needed materials and other inputs into the firm from suppliers;
- *Operations*—turning inputs into products or services;
- *Outbound logistics*—delivering products or services to consumers, distribution centers, retailers, or other partners;
- *Marketing and sales*—customer engagement, pricing, promotion, transaction; and
- *Support*—service, maintenance, and customer support.

The secondary components are the following:

- *Firm infrastructure*—functions that support the whole firm, including general management, planning, IS, and finance;
- *Human resource management*—recruiting, hiring, training, and development;
- *Technology / research and development*—new product and process design; and
- *Procurement*—sourcing and purchasing functions.

While the value chain is typically depicted as it's displayed in [Figure 2.7 "The Value Chain"](#), goods and information don't necessarily flow in a line from one function to another. For example, an order taken by the marketing function can trigger an inbound logistics function to get components from a supplier, operations functions (to build a product if it's not available), or outbound logistics functions (to ship a product when it's available). Similarly, information from service support can be fed back to advise research and development (R&D) in the design of future products.



Figure 2.7 *The Value Chain*



An analysis of a firm's value chain can reveal operational weaknesses, and technology is often of great benefit to improving the speed and quality of execution. Software tools such as *supply chain management* (SCM: linking inbound and outbound logistics with operations), *customer relationship management* (CRM: supporting sales, marketing, and in some cases R&D), and *enterprise resource planning software* (ERP: software implemented in modules to automate the entire value chain), can have a big impact on more efficiently integrating the activities within the firm, as well as with its suppliers and customers. But remember, these software tools can be purchased by all competitors. Although they can cut costs and increase efficiency, if others can buy the same or comparable products then these technologies, while valuable, may not yield lasting competitive advantage. Even more important to consider, if a firm adopts software that changes a unique process into a generic one, it may have co-opted a key source of competitive advantage particularly if other firms can buy the same stuff. This isn't a problem with accounting software. Accounting processes are standardized and accounting isn't a source of competitive advantage, so most firms buy rather than build their own accounting software. But using packaged, third-party SCM, CRM, and ERP software typically requires adopting a very specific way of doing things, using software and methods that can be purchased and adopted by others. Dell stopped deployment of

the logistics and manufacturing modules of a packaged ERP implementation when it realized that the software would require the firm to make changes to its unique and highly successful operating model and that many of the firm's unique supply chain advantages would change to the point where the firm was doing the same thing using the same software as its competitors. By contrast, Apple had no problem adopting third-party ERP software because the firm competes on product uniqueness rather than operational differences.

From a strategic perspective, managers can also consider the firm's differences and distinctiveness compared to rivals. If a firm's value chain cannot be copied by competitors without engaging in painful tradeoffs, or if the firm's value chain helps to create and strengthen other strategic assets over time, it can be a key source for competitive advantage. Many of the cases covered in this book, including FreshDirect, Amazon, Zara, Netflix, and eBay, illustrate this point.

KEY TAKEAWAYS

- The value chain can be used to map a firm's efficiency and to benchmark it against rivals, revealing opportunities to use technology to improve processes and procedures. When these firms are resistant to imitation, a firm's value chain may yield sustainable competitive advantage.
- Firms may consider adopting packaged software or outsourcing value chain tasks that are not critical to a firm's competitive advantage.
- Firms should be wary of adopting software packages or outsourcing portions of its value chain that are proprietary and a source of competitive advantage.

QUESTIONS AND EXERCISES

1. What is the value chain?
2. List the five major components of the value chain.

3. List the four secondary components of the value chain.
4. Discuss the elements of FreshDirect's value chain and the technologies that FreshDirect uses to give them a competitive advantage.
5. Which firm should adopt third-party software to automate its supply chain—Dell or Apple? Why?
6. Identify two firms in the same industry that have different value components. Why do you think these firms have different value chains? What role do you think technology plays in the way that each firm competes? Do these differences enable strategic positioning or not? Why?
7. Under what circumstances might a firm adopt packaged software like SCM, CRM, and ERP? When might this not be such a good idea? Why?
8. Describe the sales and marketing component of Apple's value chain. How and where does Apple "sell" and generate revenue? In what ways is this the same as sales and marketing at HP and Dell? In what ways is it different? How has technology altered the way that sales and marketing take place?



Chapter 3

Netflix: David Becomes Goliath

3.1 Introduction

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand the basics of the Netflix business model.
2. Recognize the downside the firm may have experienced from an early IPO.
3. Appreciate why other firms found Netflix's market attractive, and why many analysts incorrectly suspected Netflix was doomed.

Entrepreneurs are supposed to want to go public. When a firm sells stock for the first time, the company gains a ton of cash to fuel expansion and its founders get rich. Going public is the dream in the back of the mind of every tech entrepreneur. But in 2007, Netflix founder and CEO Reed Hastings told *Fortune* that if he could change one strategic decision, it would have been to delay the firm's initial public stock offering (IPO). ^[1]“If we had stayed private for another two to four years, not as many people would have understood how big a business this could be.” Once Netflix was a public company, financial disclosure rules forced the firm to reveal that it was on a money-minting growth tear. Once the secret was out, rivals showed up.



Hollywood's best couldn't have scripted a more menacing group of rivals for Hastings to face. First in line with its own DVD-by-mail offering was Blockbuster, a name synonymous with video rental. Some forty million U.S. families were already card-carrying Blockbuster customers, and the firm's efforts promised to link DVD-by-mail with the nation's largest network of video stores. Following close behind was Wal-Mart—not just a big *Fortune* 500 company but *the* largest firm in the United States ranked by sales. In Netflix, Hastings had built a great firm, but let's face it, his was a dot-com, an Internetpure play without a storefront and with an overall customer base that seemed microscopic compared to these behemoths.

Before all this, Netflix was feeling so confident that it had actually raised prices. Customers loved the service, the company was dominating its niche, and it seemed like the firm could take advantage of a modest price hike, pull in more revenue, and use this to improve and expand the business. But the firm was surprised by how quickly the newcomers mimicked Netflix with cheaper rival efforts. This new competition forced Netflix to cut prices even lower than where they had been before the price increase. To keep pace, Netflix also upped advertising at a time when online ad rates were increasing. Big competitors, a price war, spending on the rise—how could Netflix possibly withstand this onslaught? Some Wall Street analysts had even taken to referring to Netflix's survival prospects as “The Last Picture Show.” [2]

Fast forward a year later and Wal-Mart had cut and run, dumping their experiment in DVD-by-mail. Blockbuster had been mortally wounded, hemorrhaging billions of dollars in a string of quarterly losses. And Netflix? Not only had the firm held customers, it grew bigger, recording record profits. The dot-com did it. Hastings, a man who prior to Netflix had already built and sold one of the fifty largest public software firms in the



United States, had clearly established himself as one of America's most capable and innovative technology leaders. In fact, at roughly the same time that Blockbuster CEO John Antioco resigned, Reed Hastings accepted an appointment to the Board of Directors of none other than the world's largest software firm, Microsoft. Like the final scene in so many movies where the hero's face is splashed across the news, *Time* named Hastings as one of the "100 most influential global citizens."

Why Study Netflix?

Studying Netflix gives us a chance to examine how technology helps firms craft and reinforce a competitive advantage. We'll pick apart the components of the firm's strategy and learn how technology played a starring role in placing the firm atop its industry. We also realize that while Netflix emerged the victorious underdog at the end of the first show, there will be at least one sequel, with the final scene yet to be determined. We'll finish the case with a look at the very significant challenges the firm faces as new technology continues to shift the competitive landscape.

Reed Hastings, a former Peace Corps volunteer with a Master's in Computer Science, got the idea for Netflix when he was late in returning the movie *Apollo 13* to his local video store. The forty-dollar late fee was enough to have bought the DVD outright with money left over. Hastings felt ripped off, and out of this initial outrage, Netflix was born. The model the firm eventually settled on was a DVD-by-mail service that charged a flat-rate monthly subscription rather than a per-disc rental fee. Customers don't pay a cent in mailing expenses, and there are no late fees.

Netflix offers nine different subscription plans, starting at less than five dollars. The most popular is a \$16.99 option that offers customers three movies at a time, and unlimited returns each month. Videos arrive in red Mylar envelopes. After tearing off the cover to remove the DVD, customers

reveal prepaid postage and return address. When done watching videos, consumers just slip the DVD back into the envelope, reseal it with a peel-back sticky-strip, and drop the disc in the mail. Users make their video choices in their “request queue” at Netflix.com.

If a title isn’t available, Netflix simply moves to the next title in the queue. Consumers use the Web site to rate videos they’ve seen, specify their movie preferences, get video recommendations, check out DVD details, and even share their viewing habits and reviews. In 2007, the firm added a “Watch Now” button next to those videos that could be automatically streamed to a PC. Any customer paying at least \$8.99 for a DVD-by-mail subscription plan can stream an unlimited number of videos each month at no extra cost.

KEY TAKEAWAYS

- Analysts and managers have struggled to realize that dot-com startup Netflix could actually create sustainable competitive advantage, beating back challenges from Wal-Mart and Blockbuster, among others.
- Data disclosure required by public companies may have attracted these larger rivals to the firm’s market.
- Netflix operates via a DVD subscription and video streaming model. Although sometimes referred to as “rental,” the model is really a substitute good for conventional use-based media rental.

QUESTIONS AND EXERCISES

1. How does the Netflix business model work?
2. Which firms are or have been Netflix most significant competitors? How do their financial results or performance of their efforts compare to Netflix’s efforts?
3. What recent appointment did Reed Hastings accept in addition to his job as Netflix CEO? Why is this appointment potentially important for Netflix?



4. Why did Wal-Mart and Blockbuster managers, as well as Wall Street analysts, underestimate Netflix? What issues might you advise analysts and managers to consider so that they avoid making these sorts of mistakes in the future?

3.2 Tech and Timing: Creating Killer Assets

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand how many firms have confused brand and advertising, why branding is particularly important for online firms, and the factors behind Netflix's exceptional brand strength.



2. Understand the “*long tail*” concept, and how it relates to Netflix’s ability to offer the customer a huge (the industry’s largest) selection of movies.
3. Know what *collaborative filtering* is, how Netflix uses collaborative filtering software to match movie titles with the customer’s taste, and in what ways this software helps Netflix garner sustainable competitive advantage.
4. List and discuss the several technologies Netflix uses in its operations to reduce costs and deliver customer satisfaction and enhance brand value.
5. Understand the role that scale economies play in Netflix’s strategies, and how these scale economies pose an entry barrier to potential competitors.
6. Understand the role that market entry timing has played in the firm’s success.

To understand Netflix strengths, it’s important to view the firm as its customers see it. And for the most part, what they see they like—a lot! Netflix customers are rabidly loyal and rave about the service. The firm repeatedly ranks at the top of customer satisfaction surveys. Ratings agency Forsee has named Netflix the number one e-commerce site in terms of customer satisfaction nine times in a row (placing it ahead of Apple and Amazon, among others). Netflix has also been cited as the best at satisfying customers by Nielsen and *Fast Company*, and was also named the Retail Innovator of the Year by the National Retail Federation.

Building a great brand, especially one online, starts with offering exceptional value to the customer. Don’t confuse branding with advertising. During the dot-com era, firms thought brands could be built through Super Bowl ads and expensive television promotion. Advertising can build awareness, but *brands are built through customer experience*. This is a particularly important lesson for online firms. Have a bad experience at a



burger joint and you might avoid that location but try another of the firm's outlets a few blocks away. Have a bad experience online and you're turned off by the firm's one and only virtual storefront. If you click over to an online rival, the offending firm may have lost you forever. But if a firm can get you to stay through quality experience; switching costs and data-driven value might keep you there for a long, long time, even when new entrants try to court you away.

If brand is built through customer experience, consider what this means for the Netflix subscriber. They expect the firm to offer a huge selection, to be able to find what they want, for it to arrive on time, for all of this to occur with no-brainer ease of use and convenience, and at a fair price. Technology drives all of these capabilities, so tech is at the very center of the firm's brand building efforts. Let's look at how the firm does it.

Selection: The Long Tail in Action

Customers have flocked to Netflix in part because of the firm's staggering selection. A traditional video store (and Blockbuster had some 7,800 of them) stocks roughly three thousand DVD titles on its shelves. For comparison, Netflix is able to offer its customers a selection of over one hundred thousand DVD titles, and rising! At traditional brick and mortar retailers, shelf space is the biggest constraint limiting a firm's ability to offer customers what they want when they want it. Just which films, documentaries, concerts, cartoons, TV shows, and other fare make it inside the four walls of a Blockbuster store is dictated by what the average consumer is most likely to be interested in. To put it simply, Blockbuster stocks blockbusters.

Finding the right product mix and store size can be tricky. Offer too many titles in a bigger storefront and there may not be enough paying customers to justify stocking less popular titles (remember, it's not just the cost of the

DVD—firms also pay for the real estate of a larger store, the workers, the energy to power the facility, etc.). You get the picture—there’s a breakeven point that is arrived at by considering the geographic constraint of the number of customers that can reach a location, factored in with store size, store inventory, the payback from that inventory, and the cost to own and operate the store. Anyone who has visited a video store only to find a title out of stock has run up against the limits of the physical store model. But many online businesses are able to run around these limits of geography and shelf space. Internet firms that ship products can get away with having just a few highly automated warehouses, each stocking just about all the products in a particular category. And for firms that distribute products digitally (think songs on iTunes), the efficiencies are even greater because there’s no warehouse or physical product at all (more on that later).

Offer a nearly limitless selection and something interesting happens: there’s actually *more money* to be made selling the obscure stuff than the hits. Music service Rhapsody makes more from songs outside of the top ten thousand than it does from songs ranked above ten thousand. At Amazon.com, roughly 60 percent of books sold are titles that aren’t available in even the biggest Borders or Barnes & Noble Superstores. ^[1] And at Netflix, roughly 75 percent of DVD titles shipped are from back-catalog titles, not new releases (at Blockbuster outlets the equation is nearly flipped, with some 70 percent of business coming from new releases). ^[2] Consider that Netflix sends out forty-five thousand different titles each day. That’s *fifteen times* the selection available at your average video store! Each quarter, roughly 95 percent of titles are viewed—that means that every few weeks Netflix is able to find a customer for nearly *every* DVD title that has *ever* been commercially released.



This phenomenon whereby firms can make money by selling a near-limitless selection of less-popular products is known as the long tail. The term was coined by Chris Anderson, an editor at *Wired* magazine, who also wrote a best-selling business book by the same name. The “tail” (see [Figure 3.2 "The Long Tail"](#)) refers to the demand for less popular items that aren't offered by traditional brick and mortar shops. While most stores make money from the area under the curve from the vertical axis to the dotted line, long tail firms can also sell the less popular stuff. Each item under the right part of the curve may experience less demand than the most popular products, but someone somewhere likely wants it. And as demonstrated from the examples above, the total demand for the obscure stuff is often much larger than what can be profitably sold through traditional stores alone. While some debate the size of the tail (e.g., whether obscure titles collectively are more profitable for most firms), two facts are critical to keep above this debate: (1) selection attracts customers and (2) the Internet allows large-selection inventory efficiencies that offline firms can't match.

Figure 3.2 The Long Tail



The long tail works because the cost of production and distribution drop to a point where it becomes economically viable to offer a huge selection. For Netflix, the cost to stock and ship an obscure foreign film is the same as sending out the latest Will Smith chartbuster. The long tail gives the firm a selection advantage (or one based on scale) that traditional stores simply cannot match.

For more evidence that there is demand for the obscure stuff, consider Bollywood cinema—a term referring to films produced in India. When ranked by the number of movies produced each year, Bollywood is actually bigger than Hollywood, but in terms of U.S. demand, even the top-grossing Hindi film might open in only one or two American theaters, and few video stores carry many Bollywood DVDs. Again, we see the limits that geography and shelf space impose on traditional stores. As Anderson puts it, when it comes to traditional methods of distribution, “an audience too thinly spread is the same as no audience at all.” [3] While there are roughly 1.7 million South Asians living in the United States, Bollywood fans are geographically disbursed, making it difficult to offer content at a physical storefront. Fans of foreign films would often find the biggest selection at an ethnic grocery store, but even then, that wouldn’t be much. Enter Netflix. The firm has found the U.S. fans of South Asian cinema, sending out roughly one hundred thousand Bollywood DVDs a month. As geographic constraints go away, untapped markets open up!

The power of Netflix can revive even well-regarded work by some of Hollywood’s biggest names. In between *The Godfather* and *The Godfather Part II*, director Francis Ford Coppola made *The Conversation*, a film starring Gene Hackman that, in 1975, was nominated for a Best Picture Academy Award. Coppola has called “The Conversation” the finest film he has ever made, [4] but it was headed for obscurity as the ever-growing pipeline of new releases pushed the film off of video store shelves. Netflix



was happy to pick up *The Conversation* and put it in the long tail. Since then, the number of customers viewing the film has tripled, and on Netflix, this once underappreciated gem became the thirteenth most watched film from its time period.

For evidence on Netflix's power to make lucrative markets from nonblockbusters, visit the firm's "Top 100 page." [5] You'll see a list loaded with films that were notable for their *lack* of box office success. As of this writing the number one rank had been held *for nearly four years in a row*, not by a first-run mega-hit, but by the independent film *Crash* (an Oscar winner, but box office weakling). [6]

Netflix has used the long tail to its advantage, crafting a business model that creates close ties with film studios. Studios love Netflix because in most cases they earn a percentage of the subscription revenue for every disk sent out to a Netflix customer. In exchange, Netflix gets to buy the studio's DVDs at cost. The movie business is characterized by large fixed costs up front. Studio marketing budgets are concentrated on films when they first appear in theaters, and when they're first offered on DVD. After that, studios are done promoting a film, focusing instead on its most current titles. But Netflix is able to find an audience for a film without the studios spending a dime on additional marketing. Since so many of the titles viewed on Netflix are in the long tail, revenue sharing is all gravy for the studios—additional income they would otherwise be unlikely to get. It's a win-win for both ends of the supply chain. These supplier partnerships grant Netflix a sort of soft bargaining power that's distinctly opposite the strong-arm price bullying that giants like Wal-Mart are often accused of.

The VCR, the Real "Killer App"?

Netflix's coziness with movie studios is particularly noteworthy, given that the film industry has often viewed new technologies with a suspicion

bordering on paranoia. In one of the most notorious incidents, Jack Valenti, the former head of the Motion Picture Association of American (MPAA) once lobbied the U.S. Congress to limit the sale of home video recorders, claiming “the VCR is to the American film producer and the American public as the Boston strangler is to the woman home alone.” [7] Not only was the statement over the top, Jack couldn’t have been more wrong. Revenue from the sale of VCR tapes would eventually surpass the take from theater box offices, and today, home video brings in about two times box office earnings.

Cinematch: Technology Creates a Data Asset that Delivers Profits

Netflix proves there’s both demand and money to be made from the vast back catalog of film and TV show content. But for the model to work best, the firm needed to address the biggest inefficiency in the movie industry—“audience finding,” that is, matching content with customers. To do this, Netflix leverages some of the industry’s most sophisticated technology, a proprietary recommendation system that the firm calls Cinematch. Each time a customer visits Netflix after sending back a DVD, the service essentially asks “so, how did you like the movie?” With a single click, each film can be rated on a scale of one to five stars. If you’re new to Netflix, the service can prompt you with a list of movies (or you can search out and rate titles on your own). Love *Rushmore* but hate *The Life Aquatic*? Netflix wants to know.

The magic of Cinematch happens not by offering a gross average user rating—user tastes are too varied and that data’s too coarse to be of significant value. Instead, Cinematch develops a map of user ratings and steers you toward titles preferred by people with tastes that are most like yours. Techies and marketers call this trick collaborative filtering. The term

refers to a classification of software that monitors trends among customers and uses this data to personalize an individual customer's experience. Input from collaborative filtering software can be used to customize the display of a Web page for each user so that an individual is greeted only with those items the software predicts they'll most likely be interested in. The kind of data mining done by collaborative filtering isn't just used by Netflix; other sites use similar systems to recommend music, books, even news stories. While other firms also employ collaborative filtering, Netflix has been at this game for years, and is constantly tweaking its efforts. The results are considered the industry gold standard.

Collaborative filtering software is powerful stuff, but is it a source of competitive advantage? Ultimately it's just math. Difficult math, to be sure, but nothing prevents other firms from working hard in the lab, running and refining tests, and coming up with software that's as good, or perhaps one day even better than Netflix's offering. But what the software has created for the early-moving Netflix is an enormous data advantage that is valuable, results yielding, and impossible for rivals to match. Even if Netflix gave Cinematch to its competitors, they'd be without the over-two-billion ratings that the firm has amassed (according to the firm, users add about a million new ratings to the system each day). More ratings make the system seem smarter, and with more info to go on, Cinematch can make more accurate recommendations than rivals.

Evidence suggests that users trust and value Cinematch. Recommended titles make up over 60 percent of the content users place in their queues—an astonishing penetration rate. Compare that to how often you've received a great recommendation from the sullen teen behind the video store counter. While data and algorithms improve the service and further strengthen the firm's brand, this data is also a switching cost. Drop Netflix for Blockbuster and the average user abandons the two hundred or more



films they've rated. Even if one is willing to invest the time in recreating their ratings on Blockbuster's site, the rival will still make less accurate recommendations because there are fewer users and less data to narrow in on similarities across customers.

One way to see how strong these switching costs are is to examine the Netflix churn rate. Churn is a marketing term referring to the rate at which customers leave a product or service. A low churn is usually key to profitability because it costs more to acquire a customer than to keep one. And the longer a customer stays with the firm, the more profitable they become and the less likely they are to leave. If customers weren't completely satisfied with the Netflix experience, many would be willing to churn out and experiment with rivals offering cheaper service. However, the year after Blockbuster and Wal-Mart launched with copycat efforts, the rate at which customers left Netflix actually *fell* below 4 percent, an all-time low. And the firm's churn rates have continued to fall over time. By the middle of 2008, rates for customers in Netflix most active regions of the country were below 3 percent, meaning fewer than three in one hundred Netflix customers canceled their subscriptions each year. ^[8] To get an idea of how enviable the Netflix churn rates are, consider that in 2007 the mobile phone industry had a churn rate of 38.6 percent, while roughly one in four U.S. banking customers defected that year. ^[9]

All of this impacts marketing costs, too. Happy customers refer friends (free marketing from a source consumers trust more than a TV commercial). Ninety-four percent of Netflix subscribers say they have recommended the service to someone else, and 71 percent of new subscribers say an existing subscriber has encouraged them to sign up. It's no wonder subscriber acquisition costs have been steadily falling, further contributing to the firm's overall profitability.

The Netflix Prize

Netflix isn't content to stand still with its recommendation engine. Recognizing that there may be useful expertise outside its Los Gatos, California headquarters, the firm launched a crowdsourcing effort known as The Netflix Prize (for more on crowdsourcing, see [Chapter 6 "Peer Production, Social Media, and Web 2.0"](#)). The goal was simple: Offer one million dollars to the first group or individual who can improve Cinematch's ratings accuracy by 10 percent. In order to give developers something to work with, the firm turned over a large ratings database (with customer-identifying information masked, of course). The effort has attracted over 30,000 teams from 170 countries. Not bad when you consider that one million dollars would otherwise fund just four senior Silicon Valley engineers for about a year. And the effort earned Netflix a huge amount of PR, as newspapers, magazines, and bloggers chatted up the effort.

While Netflix gains access to any of the code submitted as part of the prize, it isn't exclusive access. The Prize underscores the value of the data asset. Even if others incorporate the same technology as Netflix, the firm still has user data (and attendant customer switching costs) that prevent rivals with equal technology from posing any real threat. Results incorporating many innovations offered by contest participants were incorporated into Cinematch, even before the prize was won.

As the contest dragged on, many participants wondered if the 10 percent threshold could ever be reached. While many teams grew within striking distance, a handful of particularly vexing titles thwarted all algorithms. Perhaps the most notorious title was Napoleon Dynamite. The film is so quirky, and Netflix customers so polarized, that there's little prior indicator to suggest if you're in the love it or hate it camp. One contestant claimed

that single film was responsible for 15 percent of the gap between his team's effort and the million dollars. ^[10]

The eventual winner turned out to be a coalition of four teams from four countries—prior rivals who sought to pool their noggins and grab fame and glory (even if their individual prize split was less). BellKor's Pragmatic Chaos, the first team to cross the 10 percent threshold, included a pair of coders from Montreal; two U.S. researchers from AT&T Labs; a scientist from Yahoo! Research, Israel; and a couple of Austrian consultants. ^[11] It's safe to say that without the Netflix Prize, these folks would likely never have met, let alone collaborated.

Patron Saint of the Independent Film Crowd

Many critically acclaimed films that failed to be box office hits have gained a second life on Netflix, netting significant revenue for the studios, with no additional studio marketing. *Babel*, *The Queen*, and *The Last King of Scotland* are among the films that failed to crack the top twenty in the box office, but ranked among the most requested titles on Netflix during the year after their release. Netflix actually delivered more revenue to Fox from *The Last King of Scotland* than it did from the final X-Men film. ^[12] In the true spirit of the long tail, Netflix has begun acquiring small market titles for exclusive distribution. One of its first efforts involved the Oscar-nominated PBS documentary, *Daughters from Danang*. PBS hadn't planned to distribute the disc after the Academy Awards; it was simply too costly to justify producing a run of DVDs that almost no retailer would carry. But in a deal with PBS, Netflix assumed all production costs in exchange for exclusive distribution rights. For months after, the film repeatedly ranked in the Top 15 most requested titles in the documentary category. Cost to PBS—nothing. ^[13]

A Look at Operations

Tech also lies at the heart of the warehouse operations that deliver customer satisfaction and enhance brand value. As mentioned earlier, brand is built through customer experience, and a critical component of customer experience is for subscribers to get their DVDs as quickly as possible. In order to do this, Netflix has blanketed the country with a network of fifty-eight ultrahigh-tech distribution centers that collectively handle in excess of 1.8 million DVDs a day. These distribution centers are purposely located within driving distance of 119 U.S. Postal Service (U.S.P.S.) processing and distribution facilities.

By 4:00 a.m. each weekday, Netflix trucks collect the day's DVD shipments from these U.S.P.S. hubs and returns the DVDs to the nearest Netflix center. DVDs are fed into custom-built sorters that handle disc volume on the way in and the way out. That same machine fires off an e-mail as soon as it detects your DVD was safely returned (now rate it for Cinematch).

Most DVDs never hit the restocking shelves. Scanners pick out incoming titles that are destined for other users and place these titles into a sorted outbound pile with a new, appropriately addressed red envelope. Netflix not only helps out the postal service by picking up and dropping off the DVDs at its hubs, it presorts all outgoing mail for faster delivery. This extra effort has a payoff—Netflix gets the lowest possible postal rates for first-class mail delivery. And despite the high level of automation, 100 percent of all discs are inspected by hand so that cracked ones can be replaced, and dirty ones given a wipe down. ^[14] Total in and out turnaround time for a typical Netflix DVD is just eight hours! ^[15]

First-class mail takes only one day to be delivered within a fifty-mile radius, so the warehouse network allows Netflix to service over 97 percent of its customer base within a two-day window—one day is allotted for receipt;



early the next morning the next item in their queue is processed; and the new title arrives at the customer's address by that afternoon. And in 2009, the firm added Saturday processing. All this means a customer with the firm's most popular "three disc at a time" plan could watch a movie a day and never be without a fresh title.

Warehouse processes don't exist in a vacuum; they are linked to Cinematch to offer the firm additional operational advantages. The software recommends movies that are likely to be in stock so users aren't frustrated by a wait.

Everyone on staff is expected to have an eye on improving the firm's processes. Every warehouse worker gets a free DVD player and Netflix subscription so that they understand the service from the customer's perspective and can provide suggestions for improvement. Quality management features are built into systems supporting nearly every process at the firm, allowing Netflix to monitor and record the circumstances surrounding any failures. When an error occurs, a tiger team of quality improvement personnel swoops in to figure out how to prevent any problems from recurring. Each phone call is a cost, not a revenue enhancement, and each error increases the chance that a dissatisfied customer will bolt for a rival.

By paying attention to process improvements and designing technology to smooth operations, Netflix has slashed the number of customer representatives even as subscriptions ballooned. In the early days, when the firm had one hundred and fifteen thousand customers, Netflix had one hundred phone support reps. By the time the customer base had grown thirtyfold, errors had been reduced to so that only forty-three reps were needed. ^[16] Even more impressive, because of the firm's effective use of technology to drive the firm's operations, fulfillment costs as a percentage

of revenue have actually dropped even though postal rates have increased and Netflix has cut prices.

Killer Asset Recap: Understanding Scale

Netflix executives are quite frank that the technology and procedures that make up their model can be copied, but they also realize the challenges that any copycat rival faces. Says the firm's VP of Operations Andy Rendich, "Anyone can replicate the Netflix operations if they wish. It's not going to be easy. It's going to take a lot of time and a lot of money." [17]

While we referred to Netflix as David to the goliaths of Wal-Mart and Blockbuster, within the DVD-by-mail segment Netflix is now the biggest player by far, and this size gives the firm significant scale advantages. The yearly cost to run a Netflix-comparable nationwide delivery infrastructure is about three hundred million dollars. [18] Think about how this relates to economies of scale. In the [Chapter 2 "Strategy and Technology"](#) we said that firms enjoy *scale economies* when they are able to leverage the cost of an investment across increasing units of production. Even if rivals have identical infrastructures, the more profitable firm will be the one with more customers (see [Figure 3.7](#)). And the firm with better scale economies is in a position to lower prices, as well as to spend more on customer acquisition, new features, or other efforts. Smaller rivals have an uphill fight, while established firms that try to challenge Netflix with a copycat effort are in a position where they're straddling markets, unable to gain full efficiencies from their efforts.

Figure 3.7





Running a nationwide sales network costs an estimated \$300 million a year. But Netflix has over 3.5 times more subscribers than Blockbuster. Which firm has economies of scale?

For Blockbuster, the arrival of Netflix plays out like a horror film where it is the victim. For several years now, the in-store rental business has been a money loser. Things got worse in 2005 when Netflix pressure forced Blockbuster to drop late fees, costing it about four hundred million dollars. [19] The Blockbuster store network once had the advantage of scale, but eventually its many locations were seen as an inefficient and bloated liability. Between 2006 and 2007, the firm shuttered over 570 stores. [20] By 2008, Blockbuster had been in the red for ten of the prior eleven years. During a three-year period that included the launch of its Total Access DVD-by-mail effort, Blockbuster lost over four billion dollars. [21] The firm tried to outspend Netflix on advertising, even running Super Bowl ads for Total Access in 2007, but a money loser can't outspend its more profitable rival for long, and it has since significantly cut back on promotion. Blockbuster also couldn't sustain subscription rates below Netflix's, so it has given up its price advantage. In early 2008, Blockbuster even briefly pursued a merger with another struggling giant, Circuit City, a strategy that has left industry experts scratching their heads. A Viacom executive said about the firm, "Blockbuster will certainly not survive and it will not be

missed.” [22] This assessment has to sting, given that Viacom was once Blockbuster’s parent (the firm was spun off in 2004).

For Netflix, what delivered the triple scale advantage of the largest selection; the largest network of distribution centers; the largest customer base; and the firm’s industry-leading strength in brand and data assets? Moving first. Timing and technology don’t always yield sustainable competitive advantage, but in this case, Netflix leveraged both to craft what seems to be an extraordinarily valuable pool of assets that continue to grow and strengthen over time. To be certain, competing against a wounded giant like Blockbuster will remain difficult. The latter firm has few options and may spend itself into oblivion, harming Netflix in its collapsing gasp. And as we’ll see in the next section, while technology shifts helped Netflix attack Blockbuster’s once-dominant position, even newer technology shifts may threaten Netflix. As they like to say in the mutual fund industry “past results aren’t a guarantee of future returns.”

KEY TAKEAWAYS

- Durable brands are built through customer experience, and technology lies at the center of the Netflix top satisfaction ratings and hence the firm’s best-in-class brand strength.
- Physical retailers are limited by shelf space and geography. This limitation means that expansion requires building, stocking, and staffing operations in a new location.
- Internet retailers serve a larger geographic area with comparably smaller infrastructure and staff. This fact suggests that Internet businesses are more scalable. Firms providing digital products and services are potentially far more scalable, since physical inventory costs go away.
- The ability to serve large geographic areas through lower-cost inventory means Internet firms can provide access to the long tail of products,

potentially earning profits from less popular titles that are unprofitable for physical retailers to offer.

- Netflix technology revitalizes latent studio assets. Revenue sharing allows Netflix to provide studios with a costless opportunity to earn money from back catalog titles: content that would otherwise not justify further marketing expense or retailer shelf space.
- The strategically aligned use of technology by this early mover has allowed Netflix to gain competitive advantage through the powerful resources of brand, data and switching costs, and scale.
- Collaborative filtering technology has been continually refined, but even if this technology is copied, the true exploitable resource created and leveraged through this technology is the data asset.
- Technology leveraged across the firm’s extensive distribution network offers an operational advantage that allows the firm to reach nearly all of its customers with one-day turnaround.

QUESTIONS AND EXERCISES

1. What are Netflix’s sources of competitive advantage?
2. Does Netflix have a strong brand? Offer evidence demonstrating why the firm’s brand is or isn’t strong. How is a strong brand built?
3. Scale advantages are advantages related to size. In what key ways is Netflix “bigger” than the two major competitors who tried to enter the DVD-by-mail market?
4. What is the long tail? How “long” is the Netflix tail compared to traditional video stores?
5. What “class” of software does Netflix use to make movie recommendations? Think about [Chapter 2 "Strategy and Technology"](#): Which key competitive resource does this software “create”? What kinds of benefits does this provide to the firm? What benefits does it provide to Netflix’s suppliers?



6. Could a new competitor match Netflix's recommendation software? If it did, would this create a threat to Netflix? Why or why not?
7. What is the Netflix churn rate and what are the reasons behind this rate?
8. Netflix uses technology to coordinate the process of sorting and dropping off DVDs for the U.S. Postal service. This application of technology speeds delivery. What other advantage does it give the firm?
9. How has Netflix improved its customer service operation? What results reflect this improvement?

3.3 From Atoms to Bits: Opportunity or Threat?

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand the shift from atoms to bits, and how this is impacting a wide range of industries.
2. Recognize the various key issues holding back streaming video models.
3. Know the methods that Netflix is using to attempt to counteract these challenges.

Nicholas Negroponte, the former head of MIT's Media Lab, wrote a now classic essay on the shift from atoms to bits. Negroponte pointed out that most media products are created as bits—digital files of ones and zeros that begin their life on a computer. Music, movies, books, and newspapers are all created using digital technology. When we buy a CD, DVD, or even a

“dead tree” book or newspaper, we’re buying physical atoms that are simply a container for the bits that were created in software—a sound mixer, a video editor, or a word processor.

The shift from atoms to bits is realigning nearly every media industry. Newspapers struggle as readership migrates online and once-lucrative classified ads and job listings shift to the bits-based businesses of Craigslist, Monster.com, and LinkedIn. Apple dominates music sales, selling not a single “atom” of physical CDs, while most of the atom-selling “record store” chains of a decade ago are bankrupt. Amazon has even begun delivering digital books, developing the Kindle digital reader. Who needs to kill a tree, spill ink, fill a warehouse, and roll a gas-guzzling truck to get you a book? Kindle can slurp your purchases through the air and display them on a device lighter than any college textbook. When Amazon CEO Bezos unveiled the Kindle DX at a press event at Pace University in Spring 2009, he indicated that Kindle book sales were accounting for 35 percent of sales for the two hundred and seventy-five thousand titles available for the device ^[1]—a jaw-dropping impact for a device many had thought to be an expensive, niche product for gadget lovers.

Video is already going digital, but Netflix became a profitable business by handling the atoms of DVDs. The question is, will the atoms to bits shift crush Netflix and render it as irrelevant as Hastings did Blockbuster? Or can Reed pull off yet another victory and recast his firm for the day that DVDs disappear?

Concerns over the death of the DVD and the relentless arrival of new competitors are probably the main cause for Netflix’s stock volatility these past few years. Through 2009, the firm’s growth, revenue, and profit graphs all go up and to the right, but the stock has experienced wild swings as pundits have mostly guessed wrong about the firm’s imminent demise (one well-known Silicon Valley venture capitalist even referred to the firm as “an



ice cube in the sun,” [2] a statement Netflix countered with four years of record-breaking growth and profits). The troughs on the Netflix stock graph have proven great investment opportunities for the savvy. NFLX was up some 40 percent for the first half of 2009, a time when the sub-prime crisis hammered major exchanges. The firm continued to enjoy its most successful quarters as a public company, and subscriber growth rose even as DVD sales fall. But even the most bullish investor knows there’s no stopping the inevitable shift from atoms to bits, and the firm’s share price swings continue. When the DVD dies, the high-tech shipping and handling infrastructure that Netflix has relentlessly built will be rendered worthless. Reed Hastings clearly knows this, and he has a plan. “We named the company Netflix for a reason; we didn’t name it DVDs-by-mail.” [3] But he also prepared the public for a first-cut service that was something less than we’d expect from the long tail poster child. When speaking about the launch of the firm’s Internet video streaming offering in January 2007, Hastings said it would be “underwhelming.” The two biggest limitations of this initial service? As we’ll see below, not enough content, and figuring out how to squirt the bits to a television.

Access to Content

First the content. Two years after the launch of Netflix streaming option (enabled via a “Watch Now” button next to movies that can be viewed online), only 12,000 videos were offered, just 12 percent of the firm’s long tail. And not the best 12 percent. Why so few titles? It’s not just studio reluctance or fear of piracy. There are often complicated legal issues involved in securing the digital distribution rights for all of the content that makes up a movie. Music, archival footage, and performer rights may all hold up a title from being available under “Watch Now.” The 2007 Writers



Guild strike occurred largely due to negotiations over digital distribution, showing just how troublesome these issues can be.

Add to that the exclusivity contracts negotiated by key channels, in particular the so-called “premium” television networks. Film studios release their work in a system called windowing. Content is available to a given distribution channel (in theaters, through hospitality channels like hotels and airlines, on DVD, via pay-per-view, via pay cable, then broadcast commercial TV) for a specified time window, usually under a different revenue model (ticket sales, disc sales, license fees for broadcast). Pay television channels in particular have negotiated exclusive access to content as they strive to differentiate themselves from one another. This exclusivity means that even when a title becomes available for streaming by Netflix, it may disappear when a pay TV window opens up. If HBO or Showtime has an exclusive for a film, it’s pulled from the Netflix streaming service until the exclusive pay TV time window closes. A 2008 partnership with the Starz network helped provide access to some content locked up inside pay television windows, and deals with Disney and CBS allow for streaming of current-season shows. [4] But the firm still has a long way to go before the streaming tail seems comparably long when compared against its disc inventory.

And there’s also the influence of the king of DVD sales: Wal-Mart. The firm accounts for about 40 percent of DVD sales—a scale that delivers a lot of the bargaining power it has used to “encourage” studios to hold content from competing windows or to limit offering titles at competitive pricing during the peak new release DVD period. [5] Taken together it’s clear that shifting the long tail from atoms to bits will be significantly more difficult than buying DVDs and stacking them in a remote warehouse.



But How Does It Get to the TV?

The other major problem lies in getting content to the place where most consumers want to watch it: the living room TV. Netflix's "Watch Now" button first worked only on Windows PCs. Although the service was introduced in January 2007, the months before were fueled with speculation that the firm would partner with TiVo. Just one month later, TiVo announced its partner—Amazon.com. At that point Netflix found itself up against a host of rivals that all had a path to the television: Apple had its own hardware solution in Apple TV (not to mention the iPod and iPhone for portable viewing), the cable companies delivered OnDemand through their set-top boxes, and now Amazon had TiVo.

An internal team at Netflix developed a prototype set top box that Hastings himself supported offering. But most customers aren't enthusiastic about purchasing yet another box for their set top, the consumer electronics business is brutally competitive, and selling hardware would introduce an entirely new set of inventory, engineering, marketing, distribution, and competitive complexities.

The solution Netflix eventually settled on was to think beyond one hardware alternative and instead recruit others to provide a wealth of choice. The firm developed a software platform and makes this available to firms seeking to build Netflix access into their devices. Today Netflix streaming is baked into televisions and DVD players from LG, Samsung, Sony, and Vizio, among others. It's also available on all major video game consoles—either as a direct feature or third-party software add-on. Even TiVo now streams Netflix. And that internally developed Netflix set-top box? The group was spun out to form Roku, an independent firm that launched their own \$99 Netflix streamer.

The switch to Blu-ray DVDs may offer the most promise. Blu-ray players are on the fast track to commoditization. If consumer electronics firms



incorporate Netflix access into their players as a way to attract more customers with an additional, differentiating feature, Hastings's firm could end up with more living room access than either Amazon or Apple. There are seventy-three million households in the United States that have a DVD player and an Internet connection. Should a large portion of these homes end up with a Netflix-ready Blu-ray player, Hastings will have built himself an enviable base through which to grow the video streaming business. Who's going to win the race for delivering bits to the television is still very much an uncertain bet. The models all vary significantly. Apple's early efforts were limited, with the firm offering only video purchases for Apple TV, but eventually moving to online "rentals" that can also play on iPods. Movie studios are now all in Apple's camp, although the firm did temporarily lose NBC's television content in a dispute over pricing. Amazon and Microsoft also have online rentals and purchase services, and can get their content to the television via TiVo and Xbox, respectively (yes, this makes Microsoft both a partner and a sort of competitor, a phenomenon often referred to as cooptation, or *frenemies* [6]). Hulu, a joint venture backed by NBC, Fox, and other networks, is free, earning money from ads that run like TV commercials. While Hulu has also received glowing reviews, the venture has lagged in offering a method to get streaming content to the television. Netflix pioneered "all-you-can-eat" subscription streaming. Anyone who has at least the \$8.99 subscription plan can view an unlimited amount of video streams. And Blockbuster isn't dead yet. It also streams over TiVo and has other offerings in the works. There's a clear upside to the model when it shifts to streaming; it will eliminate a huge chunk of costs associated with shipping and handling. Postage represents one-third of the firm's expenses. A round-trip DVD mailing, even at the deep discounts Netflix receives from the U.S. Postal Service, runs about eighty cents. The bandwidth and handling costs to send



bits to a TV set are around a nickel. [7] At some point, if postage goes away, Netflix may be in a position to offer even greater profits to its studio suppliers, and to make more money itself, too.

Wrangling licensing costs presents a further challenge. Estimates peg Netflix 2009 streaming costs at about one hundred million dollars, up 250 percent in three years. But these expenses still deliver just a fraction of the long tail. Streaming licensing deals are tricky because they're so inconsistent even when titles are available. Rates vary, with some offered via a flat rate for unlimited streams, a per-stream rate, a rate for a given number of streams, and various permutations in between. Some vendors have been asking as much as four dollars per stream for more valuable content [8]—a fee that would quickly erase subscriber profits, making any such titles too costly to add to the firm's library. Remember, Netflix doesn't charge more for streaming—it's built into the price of its flat-rate subscriptions.

Any extra spending doesn't come at the best time. The switch to Blu-ray DVDs means that Netflix will be forced into the costly proposition of carrying two sets of video inventory: standard and high-def. Direct profits may not be the driver. Rather, the service may be a feature that attracts new customers to the firm and helps prevent subscriber flight to rival video-on-demand efforts. The stealth arrival of a Netflix set-top box, in the form of upgraded Blu-ray DVD players, might open even more customer acquisition opportunities to the firm. Bought a Blu-ray player? For just nine dollars per month you can get a ticket to the all-you-can-eat Netflix buffet. And more customers ready to watch content streamed by Netflix may prime the pump for studios to become more aggressive in licensing more of their content. Many TV networks and movie studios are leery of losing bargaining power to a dominant firm, having witnessed how Apple now



dictates pricing terms to music labels. The good will Netflix has earned over the years may pay off if it can become the studios' partner of first choice. While one day the firm will lose the investment in its warehouse infrastructure, nearly all assets have a limited lifespan. That's why corporations depreciate assets, writing their value down over time. The reality is that the shift from atoms to bits won't flick on like a light switch; it will be a hybrid transition that takes place over several years. If the firm can grab long-tail content, grow its customer base, and lock them in with the switching costs created by Cinematch (all big "ifs"), it just might emerge as a key player in a bits-only world.

Is the hybrid strategy a dangerous straddling gambit or a clever ploy to remain dominant? Netflix really doesn't have a choice but to try. Hastings already has a long history as one of the savviest strategic thinkers in tech. As the networks say, stay tuned!

KEY TAKEAWAYS

- The shift from atoms to bits is impacting all media industries, particularly those relying on print, video, and music content. Content creators, middlemen, retailers, consumers, and consumer electronics firms are all impacted.
- Netflix's shift to a streaming model (from atoms to bits) is limited by access to content and in methods to get this content to televisions.
- Windowing and other licensing issues limit available content, and inconsistencies in licensing rates make profitable content acquisitions a challenge.

QUESTIONS AND EXERCISES

1. What do you believe are the most significant long-term threats to Netflix? How is Netflix trying to address these threats? What obstacles does the firm face in dealing with these threats?

2. Who are the rivals to Netflix's "Watch Now" effort? Do any of these firms have advantages that Netflix lacks? What are these advantages?
3. Why would a manufacturer of DVD players be motivated to offer the Netflix "Watch Now" feature in its products?
4. Describe various revenue models available as video content shifts from atoms to bits. What are the advantages and disadvantages to each—for consumers, for studios, for middlemen like television networks and Netflix?
5. Wal-Mart backed out of the DVD-by-mail industry. Why does the firm continue to have so much influence with the major film studios? What strategic asset is Wal-Mart leveraging?
6. Investigate the firm Red Box. Do you think they are a legitimate threat to Netflix? Why or why not?



Chapter 4

Moore's Law and More: Fast, Cheap Computing and What It Means for the Manager

4.1 Introduction

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Define Moore's Law and understand the approximate rate of advancement for other technologies, including magnetic storage (disk drives) and telecommunications (fiber-optic transmission).
2. Understand how the price elasticity associated with faster and cheaper technologies opens new markets, creates new opportunities for firms and society, and can catalyze industry disruption.
3. Recognize and define various terms for measuring data capacity.
4. Consider the managerial implication of faster and cheaper computing on areas such as strategic planning, inventory, and accounting.

Faster and cheaper—those two words have driven the computer industry for decades, and the rest of the economy has been along for the ride. Today it's tough to imagine a single industry not impacted by more powerful, less expensive computing. Faster and cheaper puts mobile phones in the hands of peasant farmers, puts a free video game in your Happy Meal, and drives the drug discovery that may very well extend your life.

Some Definitions

This phenomenon of “faster, cheaper” computing is often referred to as Moore's Law, after Intel cofounder, Gordon Moore. Moore didn't show up one day, stance wide, hands on hips, and declare “behold my law,” but



he did write a four-page paper for *Electronics Magazine* in which he described how the process of chip making enabled more powerful chips to be manufactured at cheaper prices. [1]

Moore's friend, legendary chip entrepreneur and CalTech professor, Carver Mead, later coined the "Moore's Law" moniker. That name sounded snappy, plus as one of the founders of Intel, Moore had enough geek cred for the name to stick. Moore's original paper offered language only a chip designer would love, so we'll rely on the more popular definition: *chip performance per dollar doubles every eighteen months* (Moore's original paper assumed two years, but many sources today refer to the *eighteen-month* figure, so we'll stick with that).

Moore's Law applies to chips—broadly speaking, to *processors*, or the electronics stuff that's made out of silicon. [2] The microprocessor is the brain of a computing device. It's the part of the computer that executes the instructions of a computer program, allowing it to run a Web browser, word processor, video game, or virus. For processors, Moore's Law means that next generation chips should be twice as fast in *eighteen months*, but cost the same as today's models (or from another perspective, in a year and a half, chips that are same speed as today's models should be available for half the price).

Random-access memory (RAM) is chip-based memory. The RAM inside your personal computer is volatile memory, meaning that when the power goes out, all is lost that wasn't saved to nonvolatile device (i.e., a more permanent storage media like a hard disk or flash memory). Think of RAM as temporary storage that provides fast access for executing computer programs and files. When you "load" or "launch" a program, it usually moves from your hard drive to those RAM chips, where it can be more quickly executed by the processor.

Cameras, MP3 players, USB drives, and mobile phones often use flash memory (sometimes called flash RAM). It's not as fast as the RAM used in most traditional PCs, but holds data even when the power is off (so flash memory is also nonvolatile memory). You can think of flash memory as the chip-based equivalent of a hard drive. In fact, flash memory prices are falling so rapidly that several manufactures including Apple and the One Laptop per Child initiative (see the "Tech for the Poor" sidebar later in this section) have begun offering chip-based, nonvolatile memory as an alternative to laptop hard drives. The big advantage? Chips are solid state electronics (meaning no moving parts), so they're less likely to fail, and they draw less power. The solid state advantage also means that chip-based MP3 players like the iPod nano make better jogging companions than hard drive players, which can skip if jostled. For RAM chips and flash memory, Moore's Law means that in *eighteen* months you'll pay the same price as today for twice as much storage.

Computer chips are sometimes also referred to as semiconductors (a substance such as silicon dioxide used inside most computer chips that is capable of enabling as well as inhibiting the flow of electricity). So if someone refers to the *semiconductor industry*, they're talking about the chip business. [3]

Strictly speaking, Moore's Law does not apply to other technology components. But other computing components are also seeing their price versus performance curves skyrocket exponentially. Data storage doubles every twelve months. Networking speed is on a tear, too. With an equipment change at the ends of the cables, the amount of data that can be squirt over an optical fiber line can double every nine months. [4]



Get Out Your Crystal Ball

Faster and cheaper makes possible the once impossible. As a manager, your job will be about predicting the future. First, consider how the economics of Moore's Law opens new markets. When technology gets cheap, price elasticity kicks in. Tech products are highly *price elastic*, meaning consumers buy more products as they become cheaper. [5] And it's not just that existing customers load up on more tech; entire *new markets* open up as firms find new uses for these new chips.

Just look at the *five waves of computing* we've seen over the previous five decades. [6] In the *first wave* in the 1960s, computing was limited to large, room-sized mainframe computers that only governments and big corporations could afford. Moore's Law kicked in during the 1970s for the *second wave*, and minicomputers were a hit. These were refrigerator-sized computers that were as speedy as or speedier than the prior generation of mainframes, yet were affordable by work groups, factories, and smaller organizations. The 1980s brought *wave three* in the form of PCs, and by the end of the decade nearly every white-collar worker in America had a fast and cheap computer on their desk. In the 1990s *wave four* came in the form of Internet computing—cheap servers and networks made it possible to scatter data around the world, and with more power, personal computers displayed graphical interfaces that replaced complex commands with easy-to-understand menus accessible by a mouse click. At the close of the last century, the majority of the population in many developed countries had home PCs, as did most libraries and schools. Now we're in *wave five*, where computers are so fast and so inexpensive that they have become ubiquitous—woven into products in ways few imagined years before. Silicon is everywhere! It's in the throwaway RFID tags that track your luggage at the airport. It provides the smarts in the



world's billion-plus mobile phones. It's the brains inside robot vacuum cleaners, next generation Legos, and the table lamps that change color when the stock market moves up or down. These digital shifts can rearrange entire industries. Consider that today the firm that sells more cameras than any other is Nokia, a firm that offers increasingly sophisticated chip-based digital cameras as a giveaway as part of its primary product, mobile phones. This shift has occurred with such sweeping impact that former photography giants Pentax, Konica, and Minolta have all exited the camera business.

Ambient Devices and the Fifth Wave

Carl Yankowski almost never gets caught in the rain without his umbrella. That's because Yankowski's umbrella regularly and wirelessly checks weather reports on its own. If the umbrella gets word it will rain in the next few hours, the handle blinks with increasing urgency, warning its owner with a signal that seems to declare "you will soon require my services." Yankowski is CEO of "fifth wave" firm Ambient Devices, a Massachusetts startup that's embedding computing and communications technology into everyday devices in an attempt to make them "smarter" and more useful. Ambient's ability to pull off this little miracle is evidence of how quickly innovative thinkers are able to take advantage of new opportunities and pioneer new markets on the back of Moore's Law. The firm's first product, the Orb, is a lamp that can be set up to change color in real time in reaction to factors such as the performance of your stock portfolio or the intensity of the local pollen count. In just six months, the ten refugees from MIT's Media Lab that founded Ambient Devices took the idea for the Orb, designed the device and its software, licensed wireless spectrum from a pager firm that had both excess capacity and a footprint to cover over 90



percent of the United States, arranged for manufacturing, and began selling the gizmo through Brookstone and Nieman Marcus. [7]

Ambient has since expanded the product line to several low-cost appliances designed to provide information at a glance. These include the Ambient Umbrella, as well as useful little devices that grab and display data ranging from sports scores to fluctuating energy prices (so you'll put off running the dishwasher until evening during a daytime price spike). The firm even partnered with LG on a refrigerator that can remind you of an upcoming anniversary as you reach for the milk.

The original iPod has also become the jumping off point for new business lines including the iPhone, Apple TV, and iTunes store. As an online store, iTunes is always open. iTunes regularly sells tens of millions of songs on Christmas Day alone, a date when virtually all of its U.S.-based competition is closed for the holiday. In a short five years after its introduction, iTunes has sold over four billion songs and has vaulted past retail giants Wal-Mart, Best Buy, and Target to become the number one music retailer in the United States. Today's iTunes is a digital media powerhouse, selling movies, TV shows, games, and other applications. And with podcasting, Apple's iTunes University even lets students at participating schools put their professors' lectures on their gym playlist for free. Surfing the fifth wave has increased the value of Apple stock sixteenfold six years after the iPod's launch. Ride these waves to riches, but miss the power and promise of Moore's Law and you risk getting swept away in its riptide. Apple's rise occurred while Sony, a firm once synonymous with portable music, sat on the sidelines unwilling to get on the surfboard. Sony's stock stagnated, barely moving in six years. The firm has laid off thousands of workers while ceding leadership in digital music (and video) to Apple.

Table 4.1 Top U.S. Music Retailers



1992	2005	2006	2008
1. Musicland	1. Wal-Mart	1. Wal-Mart	1. iTunes
2. The Handleman	2. Best Buy	2. Best Buy	2. Wal-Mart
3. Tower Records	3. Target	3. Target	3. Best Buy
4. Trans World Music	7. iTunes	4. iTunes, Amazon tie	4. Amazon, Target tie
<p>Moore's Law restructures industries. The firms that dominated music sales when you were born are now bankrupt, while one that had never sold a track now sells more than anyone else. [8]</p>			

While the change in hard drive prices isn't directly part of Moore's Law (hard drives are magnetic storage, not silicon chips), as noted earlier, the faster and cheaper phenomenon applies to storage, too. Look to Amazon as another example of jumping onto a once-impossible opportunity courtesy of the price/performance curve. When Amazon.com was founded in 1995, the largest corporate database was one terabyte or TB (see "Bits and Bytes" below) in size. In 2003, the firm offered its "Search Inside the Book" feature, digitizing the images and text from thousands of books in its catalog. "Search Inside the Book" lets customers peer into a book's contents in a way that's both faster and more accurate than browsing a physical bookstore. Most importantly for Amazon and its suppliers, titles featured in "Search Inside the Book" enjoyed a 7 percent sales increase over nonsearchable books. When "Search Inside the Book" launched, the database to support this effort was 20 TB in size. In just eight years, the firm found that it made good business sense to launch an effort that was a full *twenty times* larger than anything used by *any* firm less than a decade earlier. And of course, all of these capacities seem laughably small by



today's standards (see [Chapter 11 "The Data Asset: Databases, Business Intelligence, and Competitive Advantage"](#)). For Amazon, the impossible had not just become possible; it became good business. By 2009, digital books weren't just for search; they were for sale. Amazon's Kindle reader (a Moore's Law marvel sporting a microprocessor and flash storage) became the firm's top-selling product in terms of both unit sales and dollar volume.

Bits and Bytes

Computers express data as bits that are either one or zero. Eight bits form a byte (think of a byte as being a single character you can type from a keyboard). A kilobyte refers to roughly a thousand bytes, or a thousand characters, megabyte = 1 million, gigabyte = 1 billion, terabyte = 1 trillion, petabyte = 1 quadrillion, exabyte = 1 quintillion bytes.

While storage is most often listed in bytes, telecommunication capacity (bandwidth) is often listed in bits per second (bps). The same prefixes apply (Kbps = kilobits, or one thousand bits, per second, Mbps = megabits per second, Gbps = gigabits per second, Tbps = terabits per second).

These are managerial definitions, but technically, a kilobyte is 2^{10} or 1,024 bytes. Mega = 2^{20} , tera = 2^{30} , giga = 2^{40} , peta = 2^{50} , and exa = 2^{60} . To get a sense for how much data we're talking about, see [Table 4.2 "Bytes Defined"](#) below. ^[9]

Table 4.2 Bytes Defined

	Managerial Definition	Exact Amount	To put it in perspective
1 Byte	One keyboard character	8 bits	1 byte = a letter or number
1 Kilobyte	One thousand	2^{10} bytes	1 typewritten page = 2 KB

	Managerial Definition	Exact Amount	To put it in perspective
(KB)	bytes		1 digital book (Kindle) = approx. 500–800 KB
1 Megabyte (MB)	One million bytes	2^{20} bytes	1 digital photo (7 megapixels) = 1.3 MB
			1 MP3 song = approx. 3 MB
			1 CD = approx. 700 MB
1 Gigabyte (GB)	One billion bytes	2^{30} bytes	1 DVD movie = approx. 4.7 GB
1 Terabyte (TB)	One trillion bytes	2^{40} bytes	Printed collection of the Library of Congress = 20 TB
1 Petabyte (PB)	One quadrillion bytes	2^{50} bytes	Wal-Mart data warehouse (2008), 2.5 PB
1 Exabyte (EB)	One quintillion bytes	2^{60} bytes	
1 Zettabyte (ZB)	One sextillion bytes	2^{70} bytes	Amount of data consumed by U.S. households in 2008 = 3.6 ZB

Here's another key implication—if you are producing products with a significant chip-based component, the chips inside that product rapidly fall in value. That's great when it makes your product cheaper and opens up new markets for your firm, but it can be deadly if you overproduce and have excess inventory sitting on shelves for long periods of time. Dell claims its inventory depreciates as much as a single percentage point in value each



week. ^[10] That's a big incentive to carry as little inventory as possible, and to unload it, fast!

While the strategic side of tech may be the most glamorous, Moore's Law impacts mundane management tasks, as well. From an accounting and budgeting perspective, as a manager you'll need to consider a number of questions: How long will your computing equipment remain useful? If you keep upgrading computing and software, what does this mean for your capital expense budget? Your training budget? Your ability to make well-reasoned predictions regarding tech's direction will be key to answering these questions.

Tech for the Poor

Nicholas Negroponte, the former head of MIT's Media Lab, is on a mission. His OLPC (One Laptop per Child) project aims to create a one-hundred-dollar laptop for distribution to the world's poor. A rubberized keyboard and entirely solid-state design (flash RAM rather than hard drive) help make the machine durable. The laptop's ultrabright screen is readable in daylight and can be flipped to convert into an e-book reader. And a host of open source software and wiki tools for courseware development all aim to keep the costs low (the first generation of the product, called the XO, was priced around \$175). Mesh networking allows laptops within a hundred feet or so to communicate with each other, relaying a single Internet connection for use by all. And since the XO is targeted at the world's poorest kids in communities where power generation is unreliable or nonexistent, several battery-charging power generation schemes have been tested including a hand crank, pull chord, and foot pedal. In its first two years, the OLPC foundation delivered 825,000 laptops to children in twenty-four countries. ^[11] The XO is a product made possible by the rapidly falling price of computing.



While the success of the OLPC effort will reveal itself over time, another tech product containing a microprocessor is already transforming the lives of some of the world's most desperate poor—the cell phone. There are three billion people worldwide that don't yet have a phone, but they will, soon. In the ultimate play of Moore's Law opening up new markets, mobiles from Vodafone and Indian telecom provider Spice sell for \$25 or less. While it took roughly twenty years to sell a billion mobile phones worldwide, the second billion sold in four years, and the third billion took just two years. Today some 80 percent of the world's population lives within cellular network range, double the 2000 level. Today the vast majority of mobile subscriptions are in developing countries. [12]

Why such demand? Mobiles change lives for the better. According to Columbia economist Jeffrey Sachs, “The cell phone is the single most transformative technology for world economic development.” [13] Think about the farmer who can verify prices and locate buyers before harvesting and transporting perishable crops to market; the laborer who was mostly unemployed but with a mobile is now reachable by those who have day-to-day work; the mother who can find out if a doctor is in and has medicine before taking off work to make the costly trek to a remote clinic with her sick child; or the immigrant laborer serving as a housekeeper who was “more or less an indentured servant until she got a cell phone” enabling new customers to call and book her services. [14]

As an example of impact, look to poor fishermen in the Indian state of Kerala. By using mobile phones to find the best local marketplace prices for sardines, these fishermen were able to increase their profits by an average of eight percent even though consumer prices for fish *dropped* 4 percent. The trends benefiting both buyer and seller occurred because the fishermen no longer had to throw away unsold catch previously lost by sailing into a port after all the buyers had left. A London Business School study found

that for every ten mobile phones per one hundred people, a country's GDP bumps up 0.5 percent. [15]

Bangladeshi economist Mohammed Yunus won the Nobel Peace Prize based on his work in the microfinance movement, an effort that provides very small loans to the world's poorest entrepreneurs. Microfinance loans grew the market for Grameen Phone Ltd., a firm that has empowered over two hundred and fifty thousand Bangladeshi "phone ladies" to start businesses that helped their communities become more productive. Phone ladies buy a phone and special antenna on microcredit for about \$150 each. These special long-life battery phones allow them to become a sort of village operator, charging a small commission for sending and receiving calls. Through phone ladies, the power of the mobile reaches even those too poor to afford buying one outright. Grameen Phone now has annual revenues of over one billion dollars and is Bangladesh's largest telecom provider. In another ingenious scheme, phone minutes become a proxy for currency. The *New York Times* reports that a person "working in Kampala, for instance, who wishes to send the equivalent of five dollars back to his mother in a village will buy a five-dollar prepaid airtime card, but rather than entering the code into his own phone, he will call the village phone operator and read the code to her. [The operator] then uses the airtime for her phone and completes the transaction by giving the man's mother the money, minus a small commission." [16]

South Africa's Wizzit and GCash in the Philippines allow customers to use mobile phones to store cash credits sent from another phone or purchased through a post office or kiosk operator. When phones can be used as currency for purchases or payments, who needs Visa? Vodafone's Kenyan-based M-Pesa mobile banking program landed 200,000 new customers in a month—they'd expected it would take a year to hit that mark. With 1.6 million customers by that time, the service is spreading throughout Africa.

The “mobile phone as bank” may bring banking to a billion unserved customers in a few years.

KEY TAKEAWAYS

- Moore’s law applies to the semiconductor industry. Moore’s Law states that for the same money, roughly eighteen months from now you should be able to purchase computer chips that are twice as fast or store twice as much information. Or over that same time period chips with the speed or storage of today’s chips should cost half as much as they do now.
- Nonchip-based technology also advances rapidly. Disk drive storage doubles roughly every twelve months, while equipment to speed transmissions over fiber-optic lines has doubled every nine months.
- These trends influence inventory value, depreciation accounting, employee training, and other managerial functions. They also help improve productivity and keep interest rates low.
- From a strategic perspective, these trends suggest that what is impossible from a cost or performance perspective today may be possible in the future. This fact provides an opportunity to those who recognize and can capitalize on the capabilities of new technology. As technology advances, new industries, business models, and products are created, while established firms and ways of doing business can be destroyed.
- Managers must regularly study trends and trajectory in technology to recognize opportunity and avoid disruption.

QUESTIONS AND EXERCISES

1. What is Moore’s Law? What does it apply to?
2. Are other aspects of computing advancing as well? At what rates?
3. What is a microprocessor? What devices do you or your family own that contain microprocessors (and hence are impacted by Moore’s Law)?

4. What is a semiconductor? What is the substance from which most semiconductors are made?
5. How does flash memory differ from the memory in a PC? Are both solid state?
6. Which of the following are solid state: an iPod shuffle, a TiVo DVR, a typical laptop PC?
7. Why is Moore's Law important for managers? How does it influence managerial thinking?
8. What is price elasticity? How does Moore's Law relate to this concept? What's special about falling chip prices compared to price drops for products like clothing or food?
9. Give examples of firms that have effectively leveraged the advancement of processing, storage, and networking technology.
10. What are the five waves of computing? Give examples of firms and industries impacted by the fifth wave.
11. As Moore's Law advances, technology becomes increasingly accessible to the poor. Give examples of how tech has benefited those who likely would not have been able to afford the technology of a prior generation.
12. How have cheaper, faster chips impacted the camera industry? Give an example of the leadership shifts that have occurred in this industry.
13. What has been the impact of "faster, cheaper" on Apple's business lines?
14. How did Amazon utilize the steep decline in magnetic storage costs to its advantage?
15. How does Moore's Law impact production and inventory decisions?



4.2 The Death of Moore's Law?

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Describe why Moore's Law continues to advance, and discuss the physical limitations of this advancement.
2. Name and describe various technologies that may extend the life of Moore's Law.
3. Discuss the limitations of each of these approaches.

Moore simply observed that we're getting better over time at squeezing more stuff into tinier spaces. Moore's Law is possible because the distance between the pathways inside silicon chips gets smaller with each successive generation. While chip plants (semiconductor fabrication facilities, or fabs) are incredibly expensive to build, each new generation of fabs can crank out more chips per silicon wafer. And since the pathways are closer together, electrons travel shorter distances. If electrons now travel half the distance to make a calculation, that means the chip is twice as fast.

But the shrinking can't go on forever, and we're already starting to see three interrelated forces—*size*, *heat*, and *power*—threatening to slow down the Moore's Law gravy train. When you make processors smaller, the more tightly packed electrons will heat up a chip—so much so that unless today's most powerful chips are cooled down, they will melt inside their packaging. To keep the fastest computers cool, most PCs, laptops, and video game consoles need fans, and most corporate data centers have elaborate and expensive air conditioning and venting systems to prevent a meltdown. A trip through the Facebook data center during its recent rise would show that the firm was a “hot” startup in more ways than one. The firm's servers ran so hot that the Plexiglass sides of the firm's server racks were warped



and melting! ^[1] The need to cool modern data centers draws a lot of power and that costs a lot of money.

The chief eco officer at Sun Microsystems has claimed that computers draw four to five percent of the world's power. Google's chief technology officer has said that the firm spends more to power its servers than the cost of the servers themselves. ^[2] Microsoft, Yahoo! and Google have all built massive data centers in the Pacific Northwest, away from their corporate headquarters, specifically choosing these locations for access to cheap hydroelectric power. Google's location in The Dalles, Oregon, is charged a cost per kilowatt hour of two cents by the local power provider, less than one-fifth of the eleven-cent rate the firm pays in Silicon Valley. ^[3] This difference means big savings for a firm that runs more than a million servers.

And while these powerful shrinking chips are getting hotter and more costly to cool, it's also important to realize that chips can't get smaller forever. At some point Moore's Law will run into the unyielding laws of nature. While we're not certain where these limits are, chip pathways certainly can't be shorter than a single molecule, and the actual physical limit is likely larger than that. Get too small and a phenomenon known as quantum tunneling kicks in, and electrons start to slide off their paths. Yikes!

Buying Time

One way to overcome this problem is with multicore microprocessors, made by putting two or more lower power processor cores (think of a core as the calculating part of a microprocessor) on a single chip. Philip Emma, IBM's Manager of Systems Technology and Microarchitecture, offers an analogy. Think of the traditional fast, hot, single-core processors as a three hundred-pound lineman, and a dual-core processor as two 160-pound guys. Says Emma, "A 300-pound lineman can generate a lot of power, but two



160-pound guys can do the same work with less overall effort.” [4] For many applications, the multicore chips will outperform a single speedy chip, while running cooler and drawing less power. Multicore processors are now mainstream.

By 2007, most PCs and laptops sold had at least a two-core (dual-core) processor. The Microsoft Xbox 360 has three cores. The PlayStation 3 includes the so-called *cell processor* developed by Sony, IBM, and Toshiba that runs nine cores. By 2008, quad-core processors were common in high-end desktops and low-end servers. By 2010, AMD plans a 12-core PC chip. Intel has even demonstrated chips with upward of 48 cores.

Multicore processors can run older software written for single-brain chips. But they usually do this by using only one core at a time. To reuse the metaphor above, this is like having one of our 160-pound workers lift away, while the other one stands around watching. Multicore operating systems can help achieve some performance gains. Versions of Windows or the Mac OS that are aware of multicore processors can assign one program to run on one core, while a second application is assigned to the next core. But, in order to take full advantage of multicore chips, applications need to be rewritten to split up tasks so that smaller portions of a problem are executed simultaneously inside each core.

Writing code for this “divide and conquer” approach is not trivial. In fact, developing software for multicore systems is described by Shahrokh Daijavad, software lead for next-generation computing systems at IBM, as “one of the hardest things you learn in computer science.” [5] Microsoft’s chief research and strategy officer has called coding for these chips “the most conceptually different [change] in the history of modern computing.” [6] Despite this challenge, some of the most aggressive adaptors of multicore chips have been video game console manufacturers. Video game applications are particularly well-suited for multiple cores since, for



example, one core might be used to render the background, another to draw objects, another for the “physics engine” that moves the objects around, and yet another to handle Internet communications for multiplayer games. Another approach to breathing life into Moore’s Law is referred to as stacked or three-dimensional semiconductors. In this approach, engineers slice a flat chip into pieces, then reconnect the pieces vertically, making a sort of “silicon sandwich.” The chips are both faster and cooler since electrons travel shorter distances. What was once an end-to-end trip on a conventional chip might just be a tiny movement up or down on a stacked chip. But stacked chips present their own challenges. In the same way that a skyscraper is more difficult and costly to design and build than a ranch house, 3D semiconductors are tougher to design and manufacture. IBM has developed stacked chips for mobile phones, claiming the technique improves power efficiency by up to 40 percent.

Quantum Leaps, Chicken Feathers, and the Indium Gallium Arsenide Valley?

Think about it—the triple threat of size, heat, and power means that Moore’s Law, perhaps the greatest economic gravy train in history, will likely come to a grinding halt in your lifetime. Mutlicore and 3D semiconductors are here today, but what else is happening to help stave off the death of Moore’s Law?

Every once in a while a material breakthrough comes along that improves chip performance. A few years back researchers discovered that replacing a chip’s aluminum components with copper could increase speeds up to 30 percent. Now scientists are concentrating on improving the very semiconductor material that chips are made of. While the silicon used in chips is wonderfully abundant (it has pretty much the same chemistry found in sand), researchers are investigating other materials that might

allow for chips with even tighter component densities. Researchers have demonstrated that chips made with supergeeky-sounding semiconductor materials such as indium gallium arsenide, indium aluminum arsenide, germanium, and bismuth telluride can run faster and require less wattage than their silicon counterparts. [7] Perhaps even more exotic (and downright bizarre), researchers at the University of Delaware have experimented with a faster-than-silicon material derived from chicken feathers! Hyperefficient chips of the future may also be made out of carbon nanotubes, once the technology to assemble the tiny structures becomes commercially viable. Other designs move away from electricity over silicon. Optical computing, where signals are sent via light rather than electricity, promises to be faster than conventional chips, if lasers can be mass produced in miniature (silicon laser experiments show promise). Others are experimenting by crafting computing components using biological material (think a DNA-based storage device).

One yet-to-be-proven technology that could blow the lid off what's possible today is quantum computing. Conventional computing stores data as a combination of bits, where a bit is either a one or a zero. Quantum computers, leveraging principles of quantum physics, employ qubits that can be both one *and* zero at the same time. Add a bit to a conventional computer's memory and you double its capacity. Add a bit to a quantum computer and its capacity increases exponentially. For comparison, consider that a computer model of serotonin, a molecule vital to regulating the human central nervous system, would require 10^{94} bytes of information. Unfortunately there's not enough matter in the universe to build a computer that big. But modeling a serotonin molecule using quantum computing would take just 424 qubits. [8]

Some speculate that quantum computers could one day allow pharmaceutical companies to create hyperdetailed representations of the

human body that reveal drug side effects before they're even tested on humans. Quantum computing might also accurately predict the weather months in advance or offer unbreakable computer security. Ever have trouble placing a name with a face? A quantum computer linked to a camera (in your sunglasses, for example) could recognize the faces of anyone you've met and give you a heads up to their name and background. [9] Opportunities abound. Of course, before quantum computing can be commercialized, researchers need to harness the freaky properties of quantum physics wherein your answer may reside in another universe, or could disappear if observed (Einstein himself referred to certain behaviors in quantum physics as "spooky action at a distance"). Pioneers in quantum computing include IBM, HP, NEC, and a Canadian startup named D-Wave. At a 2007 at the Computer History Museum in Mountain View, California, D-Wave demonstrated Orion, a sixteen-qubit computer that could find a protein in a database, figure out the optimal wedding guest seating arrangements, and solve a Sudoku puzzle. Scientific opinion varied widely as to the significance of the D-Wave advance. The Orion was built using a chip cooled to minus 273 degrees Celsius in a bath of liquid helium and tasks were performed at speeds about 100 times slower than conventional PCs. Not exactly commercial stuff. But it was the most advanced quantum computing demonstration to date. If or when quantum computing becomes a reality is still unknown, but the promise exists that while Moore's Law may run into limits imposed by Mother Nature, a new way of computing may blow past anything we can do with silicon, continuing to make possible the once impossible.

KEY TAKEAWAYS

- As chips get smaller and more powerful, they get hotter and present power-management challenges. And at some point Moore's Law will stop



because we will not longer be able to shrink the spaces between components on a chip.

- Multicore chips use two or more low-power calculating “cores” to work together in unison, but to take optimal advantage of multicore chips, software must be rewritten to “divide” a task among multiple cores.
- 3D or Stackable semiconductors can make chips faster and run cooler by shortening distances between components, but these chips are harder to design and manufacture.
- New materials may extend the life of Moore’s Law, allowing chips to get smaller, still. Entirely new methods for calculating, such as quantum computing, may also dramatically increase computing capabilities far beyond what is available today.

QUESTIONS AND EXERCISES

1. What three interrelated forces threaten to slow the advancement of Moore’s Law?
2. Which commercial solutions, described in the section above, are currently being used to counteract the forces mentioned above? How do these solutions work? What are the limitations of each?
3. Will multicore chips run software designed for single-core processors?
4. As chips grow smaller they generate increasing amounts of heat that needs to be dissipated. Why is keeping systems cool such a challenge? What are the implications for a firm like Google or Yahoo!? For a firm like Apple or Dell?
5. What are some of the materials that may replace the silicon that current chips are made of?
6. What kinds of problems might be solved if the promise of quantum computing is achieved? How might individuals and organizations leverage quantum computing? What sorts of challenges could arise from the widespread availability of such powerful computing technology?



4.3 Bringing Brains Together: Supercomputing and Grid Computing

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Give examples of the business use of supercomputing and grid computing.
2. Describe grid computing and discuss how grids transform the economics of supercomputing.
3. Understand the characteristics of problems that are and are not well suited for supercomputing and grid computing.

As Moore's Law makes possible the once impossible, businesses have begun to demand access to the world's most powerful computing technology. Supercomputers are computers that are among the fastest of any in the world at the time of their introduction.^[1] Supercomputing was once the domain of governments and high-end research labs, performing tasks such as simulating the explosion of nuclear devices, or analyzing large-scale weather and climate phenomena. But it turns out with a bit of tweaking, the algorithms used in this work are profoundly useful to business. Consider perhaps the world's most well-known supercomputer,



IBM's Deep Blue, the machine that rather controversially beat chess champion Gary Kasparov. While there is not a burning need for chess-playing computers in the world's corporations, it turns out that the computing algorithms to choose the best among multiple chess moves are similar to the math behind choosing the best combination of airline flights. One of the first customers of Deep Blue technologies was United Airlines, which gained an ability to examine three hundred and fifty thousand flight path combinations for its scheduling systems—a figure well ahead of the previous limit of three thousand. Estimated savings through better yield management? Over fifty million dollars! Finance found uses, too. An early adopter was CIBC (the Canadian Imperial Bank of Commerce), one of the largest banks in North America. Each morning CIBC uses a supercomputer to run its portfolio through Monte Carlo simulations that aren't all that different from the math used to simulate nuclear explosions. An early adopter of the technology, at the time of deployment, CIBC was the only bank that international regulators allowed to calculate its own capital needs rather than use boilerplate ratios. That cut capital on hand by hundreds of millions of dollars, a substantial percentage of the bank's capital, saving millions a year in funding costs. Also noteworthy: the supercomputer-enabled, risk-savvy CIBC was relatively unscathed by the sub-prime crisis. Modern supercomputing is typically done via a technique called massively parallel processing (computers designed with many microprocessors that work together, simultaneously, to solve problems). The fastest of these supercomputers are built using hundreds of microprocessors, all programmed to work in unison as one big brain. While supercomputers use special electronics and software to handle the massive load, the processors themselves are often of the off-the-shelf variety that you'd find in a typical PC. Virginia Tech created what at the time was the world's third-fastest supercomputer by using chips from 1,100 Macintosh



computers lashed together with off-the-shelf networking components. The total cost of the system was just \$5.2 million, far less than the typical cost for such burly hardware. The Air Force recently issued a request-for-proposal to purchase 2,200 PlayStation 3 systems in hopes of crafting a supercheap, superpowerful machine using off-the-shelf parts.

Another technology, known as grid computing, is further transforming the economics of supercomputing. With grid computing, firms place special software on its existing PCs or servers that enables these computers to work together on a common problem. Large organizations may have thousands of PCs, but they're not necessarily being used all the time, or at full capacity. With grid software installed on them, these idle devices can be marshaled to attack portions of a complex task as if they collectively were one massively parallel supercomputer. This technique radically changes the economics of high-performance computing. *BusinessWeek* reports that while a middle-of-the-road supercomputer could run as much as thirty million dollars, grid computing software and services to perform comparable tasks can cost as little as twenty-five thousand dollars, assuming an organization already has PCs and servers in place.

An early pioneer in grid computing is the biotech firm Monsanto. Monsanto enlists computers to explore ways to manipulate genes to create crop strains that are resistant to cold, drought, bugs, pesticides, or that are more nutritious. Previously with even the largest computer Monsanto had in-house, gene analysis was taking six weeks and the firm was able to analyze only ten to fifty genes a year. But by leveraging grid computing, Monsanto has reduced gene analysis to less than a day. The fiftyfold time savings now lets the firm consider thousands of genetic combinations in a year. [2] Lower R&D time means faster time to market—critical to both the firm and its customers.



Grids are now everywhere. Movie studios use them to create special effects and animated films. Proctor & Gamble has used grids to redesign the manufacturing process for Pringles potato chips. GM and Ford use grids to simulate crash tests, saving millions in junked cars and speeding time to market. Pratt and Whitney test aircraft engine designs on a grid. And biotech firms including Aventis, GlaxoSmithKlein, and Pfizer push their research through a quicker pipeline by harnessing grid power. JP Morgan Chase even launched a grid effort that mimics CIBC's supercomputer, but at a fraction of the latter's cost. By the second year of operation, the JPMorgan Chase grid was saving the firm five million dollars per year. Massive clusters of computers running software that allows them to operate as a unified service also enable new service-based computing models, such as software as a service (SaaS) and cloud computing. In these models, organizations replace traditional software and hardware that they would run in-house with services that are delivered online. Google, Microsoft, Salesforce.com, and Amazon are among the firms that have sunk billions into these Moore's Law-enabled server farms, creating entirely new businesses that promise to radically redraw the software and hardware landscape while bringing the gargantuan computing power to the little guy (see [Chapter 10 "Software in Flux: Partly Cloudy and Sometimes Free"](#)). You can join a grid, too. SETI@Home turns your computer screen saver into a method to help "search for extraterrestrial intelligence," analyzing data from the Arecibo radio telescope system in Puerto Rico (no E.T. spotted yet). FightAids@Home will enlist your PC to explore AIDS treatments. And Folding@Home is an effort by Stanford researchers to understanding the science of protein-folding within diseases such as Alzheimer's, cancer, and cystic fibrosis. A version of Folding@Home software for the PlayStation 3 had enlisted over half a million consoles within months of release. Having access to these free resources is an



enormous advantage for researchers. Says the director of Folding@Home, “Even if we were given all of the NSF supercomputing centers combined for a couple of months, that is still fewer resources than we have now.” [3]

Multicore, massively parallel, and grid computing are all related in that each attempts to lash together multiple computing devices so that they can work together to solve problems. Think of multicore chips as having several processors in a single chip. Think of massively parallel supercomputers as having several chips in one computer, and think of grid computing as using existing computers to work together on a single task (essentially a computer made up of multiple computers). While these technologies offer great promise, they’re all subject to the same limitation: software must be written to divide existing problems into smaller pieces that can be handled by each core, processor, or computer, respectively. Some problems, such as simulations, are easy to split up, but for problems that are linear (where, say, step two can’t be completed until the results from step one are known), the multiple-brain approach doesn’t offer much help.

Moore’s Law will likely hit its physical limit in your lifetime, but no one really knows if this “Moore’s Wall” is a decade away or more. What lies ahead is anyone’s guess. Some technologies, such as still-experimental quantum computing, could make computers that are more powerful than all the world’s conventional computers combined. Think strategically—new waves of innovation might soon be shouting “surf’s up!”

KEY TAKEAWAYS

- Most modern supercomputers use massive sets of microprocessors working in parallel.
- The microprocessors used in most modern supercomputers are often the same commodity chips that can be found in conventional PCs and servers.

- Moore’s Law means that businesses as diverse as financial services firms, industrial manufacturers, consumer goods firms, and film studios can now afford access to supercomputers.
- Grid computing software uses existing computer hardware to work together and mimic a massively parallel supercomputer. Using existing hardware for a grid can save a firm the millions of dollars it might otherwise cost to buy a conventional supercomputer, further bringing massive computing capabilities to organizations that would otherwise never benefit from this kind of power.
- Massively parallel computing also enables the vast server farms that power online businesses like Google and Facebook, and which create new computing models, like software as a service (SaaS) and cloud computing.
- The characteristics of problems best suited for solving via multicore systems, parallel supercomputers, or grid computers are those that can be divided up so that multiple calculating components can simultaneously work on a portion of the problem. Problems that are linear—where one part must be solved before moving to the next and the next—may have difficulty benefiting from these kinds of “divide and conquer” computing. Fortunately many problems such as financial risk modeling, animation, manufacturing simulation, and gene analysis are all suited for parallel systems.

QUESTIONS AND EXERCISES

1. What is the difference between supercomputing and grid computing?
How is each phenomenon empowered by Moore’s Law?
2. How does grid computing change the economics of supercomputing?
3. Which businesses are using supercomputing and grid computing?
Describe these uses and the advantages they offer their adopting firms.
Are they a source of competitive advantage? Why or why not?



4. What are the characteristics of problems that are most easily solved using the types of parallel computing found in grids and modern day supercomputers? What are the characteristics of the sorts of problems not well suited for this type of computing?
5. Visit the SETI@Home Web site (<http://seti.ssl.berkeley.edu>). What is the purpose of the SETI@Home project? How do you participate? Is there any possible danger to your computer if you choose to participate? (Read their rules and policies.)
6. “Big Ben” has won two Super Bowls with the Pittsburgh Steelers. “Big Ben” is also a Cray XT3 supercomputer located at the Pittsburgh Supercomputing Center (PSC). Visit the PSC Web site (<http://www.psc.edu>) and type “big ben” in the search window. How many processors does “Big Ben” have? How many cores per processor?
7. What is “Moore’s Wall”?
8. What is the advantage of using grid computing to simulate an automobile crash test as opposed to actually staging a crash?

4.4 E-waste: The Dark Side of Moore’s Law

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand the magnitude of the environmental issues caused by rapidly obsolete, faster and cheaper computing.



2. Explain the limitations of approaches attempting to tackle e-waste.
3. Understand the risks firms are exposed to when not fully considering the lifecycle of the products they sell or consume.
4. Ask questions that expose concerning ethical issues in a firm or partner's products and processes, and that help the manager behave more responsibly.

We should celebrate the great bounty Moore's Law and the tech industry bestow on our lives. Costs fall, workers become more productive, innovations flourish, and we gorge at a buffet of digital entertainment that includes music, movies, and games. But there is a dark side to this faster and cheaper advancement. A PC has an expected lifetime of three to five years. A cell phone? Two years or less. Rapid obsolescence means the creation of ever-growing mountains of discarded tech junk, known as electronic waste or e-waste. According to the U.S. Environmental Protection Agency (EPA), in 2007 the United States alone generated over 2.5 million tons of e-waste, ^[1] and the results aren't pretty. Consumer electronics and computing equipment can be a toxic cocktail that includes cadmium, mercury, lead, and other hazardous materials. Once called the "effluent of the affluent," e-waste will only increase with the rise of living standards worldwide.

The quick answer would be to recycle this stuff. Not only does e-waste contain mainstream recyclable materials we're all familiar with, like plastics and aluminum, it also contains small bits of increasingly valuable metals such as silver, platinum, and copper. In fact, there's more gold in one pound of discarded tech equipment than in one pound of mined ore. ^[2] But as the sordid record of e-waste management shows, there's often a disconnect between consumers and managers who *want* to do good and those efforts that are *actually* doing good. The complexities of the modern

value chain, the vagaries of international law, and the nefarious actions of those willing to put profits above principle show how difficult addressing this problem will be.

The process of separating out the densely packed materials inside tech products so that the value in e-waste can be effectively harvested is extremely labor intensive, more akin to reverse manufacturing than any sort of curbside recycling efforts. Sending e-waste abroad can be ten times cheaper than dealing with it at home, [3] so it's not surprising that up to 80 percent of the material dropped off for recycling is eventually exported. [4] Much of this waste ends up in China, South Asia, or sub-Saharan Africa, where it is processed in dreadful conditions.

Consider the example of Guiyu, China, a region whose poisoning has been extensively chronicled by organizations such as the Silicon Valley Toxics Coalition, the Basel Action Network (BAN), and Greenpeace. Workers in and around Guiyu toil without protective equipment, breathing clouds of toxins generated as they burn the plastic skins off of wires to get at the copper inside. Others use buckets, pots, or wok-like pans (in many cases the same implements used for cooking) to sluice components in acid baths to release precious metals—recovery processes that create even more toxins. Waste sludge and the carcasses of what's left over are most often dumped in nearby fields and streams. Water samples taken in the region showed lead and heavy metal contamination levels some four hundred to six hundred times greater than what international standards deem safe. [5] The area is so polluted that drinking water must be trucked in from eighteen miles away. Pregnancies are six times more likely to end in miscarriage, and 70 percent of the kids in the region have too much lead in their blood. [6]

China cares about its environment. The nation has banned the importing of e-waste since 2000. [8] But corruption ensures that e-waste continues to

flow into the country. According to one exporter, all that's required to get e-waste past customs authorities is to tap a one-hundred-dollar bill on the side of the container. [9] Well-meaning U.S. recyclers, as well as those attempting to collect technology for reuse in poorer countries, are often in the dark as to where their products end up

The trade is often brokered by middlemen who mask the eventual destination and fate of the products purchased. BAN investigators in Lagos, Nigeria, documented mountains of e-waste with labels from schools, U.S. government agencies, and even some of the world's largest corporations. And despite Europe's prohibition on exporting e-waste, many products originally labeled for repair and reuse end up in toxic recycling efforts. Even among those products that gain a second or third life in developing nations, the inevitable is simply postponed, with e-waste almost certain to end up in landfills that lack the protective groundwater barriers and environmental engineering of their industrialized counterparts. The reality is e-waste management is extraordinarily difficult to monitor and track, and loopholes are rampant.

Thinking deeply about the ethical consequences of a firm's business is an imperative for the modern manager. A slip up (intentional or not) can, in seconds, be captured by someone with a cell phone, uploaded to YouTube, or offered in a blog posting for the world to see. When Dell was caught using Chinese prison labor as part of its recycling efforts, one blogger chastised the firm with a tweak of its marketing tagline, posting "Dude, you're getting a cell." [10] The worst cases expose firms to legal action and can tarnish a brand for years. Big firms are big targets, and environmentalists have been quick to push the best-known tech firms and retailers to take back their products for responsible recycling and to eliminate the worst toxins from their offerings.



Consider that even Apple (where Al Gore sits on the firm's Board of Directors), has been pushed by a coalition of environmental groups on all of these fronts. Critics have shot back that signaling out Apple is unfair. The firm was one of the first computer companies to eliminate lead-lined glass monitors from its product line, and has been a pioneer of reduced-sized packaging that leverage recyclable materials. But if the firm that counts Al Gore among its advisors can get tripped up on green issues, all firms are vulnerable.

Environmentalists see this pressure to deal with e-waste as yielding results: Apple and most other tech firms have continually moved to eliminate major toxins from their manufacturing processes. All this demonstrates that today's business leaders have to be far more attuned to the impact not only of their own actions, but also to those of their suppliers and partners. How were products manufactured? Using which materials? Under what conditions? What happens to items when they're discarded? Who provides collection and disposal? It also shows the futility of legislative efforts that don't fully consider and address the problems they are meant to target.

Finding Responsible E-waste Disposers

A recent government sting from the U.S. Government Accountability Office (U.S. GAO) found that forty-three American recyclers were willing to sell e-waste illegally to foreign countries, without gaining EPA or foreign country approval. Appallingly, at least three of them held Earth Day electronics-recycling events. ^[11]

So how can firms and individuals choose proper disposal partners? Several certification mechanisms can help shed light on whether the partner you're dealing with is a responsible player. The Basel Action Network e-Stewards program certifies firms via a third-party audit, with compliant participants committing to eliminating e-waste export, land dumping, incineration, and

toxic recycling via prison labor. The International Association of Electronics Recyclers (IAER) also offers audited electronics recycler certification. And firms certified as ISO 9001 and ISO 14001 compliant attest to quality management and environmental processes. Standards, techniques, and auditing practices are constantly in flux, so consult these organizations for the latest partner lists, guidelines, and audit practices. ^[12]

Which brings us back to Gordon Moore. To his credit, Moore is not just the founder of the world's largest microprocessor firm and first to identify the properties we've come to know as Moore's Law, he has also emerged as one of the world's leading supporters of environmental causes. The generosity of the Gordon and Betty Moore foundation includes, among other major contributions, the largest single gift to a private conservation organization. Indeed, Silicon Valley, while being the birthplace of products that become e-waste, also promises to be at the forefront of finding solutions to modern environmental challenges. The Valley's leading venture capitalists, including Sequoia and Kleiner Perkins (where Al Gore is now a partner), have started multimillion-dollar green investment funds, targeted at funding the next generation of sustainable, environmental initiatives.

KEY TAKEAWAYS

- E-waste may be particularly toxic since many components contain harmful materials such as lead, cadmium, and mercury.
- Managers must consider and plan for the waste created by their products, services, and technology used by the organization. Consumers and governments are increasingly demanding that firms offer responsible methods for the disposal of their manufactured goods and the technology used in their operations.

- Managers must audit disposal and recycling partners with the same vigor as their suppliers and other corporate partners. If not, an organization's equipment may end up in environmentally harmful disposal operations.

QUESTIONS AND EXERCISES

1. What is e-waste? What is so dangerous about e-waste?
2. What sorts of materials might be harvested from e-waste recycling?
3. Many well-meaning individuals thought that recycling was the answer to the e-waste problem. But why hasn't e-waste recycling yielded the results hoped for?
4. What lessons do the challenges of e-waste offer the manager? What issues will your firm need to consider as it consumes or offers products that contain computing components?
5. Why is it difficult to recycle e-waste?
6. Why is e-waste exported abroad for recycling rather than processed domestically?
7. What part does corruption play in the recycling and disposal of e-waste?
8. What part might product design and production engineering play in the reduction of the impact of technology waste on the environment?
9. What are the possible consequences should a U.S. firm be deemed "environmentally irresponsible"?
10. Name two companies that have incurred the wrath of environmental advocates. What might these firms have done to avoid such criticism?



Chapter 5

Understanding Network Effects

5.1 Introduction

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Define network effects.
2. Recognize products and services that are subject to network effects.
3. Understand the factors that add value to products and services subject to network effects.

Network effects are sometimes referred to as “Metcalfe’s Law” or “Network Externalities.” But don’t let the dull names fool you—this concept is rocket fuel for technology firms. Bill Gates leveraged network effects to turn Windows and Office into virtual monopolies and in the process became the wealthiest man in America. Mark Zuckerberg of Facebook, Pierre Omidiyar of eBay, Caterina Fake and Stewart Butterfield of Flickr, Kevin Rose of



Digg, Evan Williams and Biz Stone of Twitter, Chris DeWolfe and Tom Anderson—the MySpace guys—all of these entrepreneurs have built massive user bases by leveraging the concept. When network effects are present, *the value of a product or service increases as the number of users grows*. Simply, more users = more value. Of course, most products aren't subject to network effects—you probably don't care if someone wears the same socks, uses the same pancake syrup, or buys the same trash bags as you. But when network effects are present they're among *the most important* reasons you'll pick one product or service over another. You may care very much, for example, if others are part of your social network, if your video game console is popular, if the Wikipedia article you're referencing has had prior readers. And all those folks who bought HD-DVD players sure were bummed when the rest of the world declared Blu-ray the winner. In each of these examples, network effects are at work.

Not *That* Kind of Network

The term “network” sometimes stumps people when first learning about network effects. In this context, a network doesn't refer to the physical wires or wireless systems that connect pieces of electronics. It just refers to a common user base that is able to communicate and share with one another. So Facebook users make up a network. So do owners of Blu-ray DVD players, traders that buy and sell stock over the NASDAQ, or the sum total of hardware and outlets that support the BS 1363 electrical standard.

KEY TAKEAWAY

- Network effects are among the most powerful strategic resources that can be created by technology-based innovation. Many category-dominating organizations and technologies owe their success to network effects, including Microsoft, Apple, NASDAQ, eBay, Facebook, and Visa.

Network effects are also behind the establishment of most standards, including Blu-ray DVD, wi-fi, and Bluetooth.

QUESTIONS AND EXERCISES

1. What are network effects? What are the other names for this concept?
2. List several products or services subject to network effects. What factors do you believe helped each of these efforts achieve dominance?
3. Which firm do you suspect has stronger end-user network effects: Google's online search tool or Microsoft's Windows operating system? Why?
4. Network effects are often associated with technology, but tech isn't a prerequisite for the existence of network effects. Name a product, service, or phenomenon that is not related to information technology that still dominates due to network effects.



5.2 Where's All That Value Come From?

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Identify the three primary sources of value for network effects.
2. Recognize factors that contribute to the staying power and complementary benefits of a product or service subject to network effects.
3. Understand how firms like Microsoft and Apple each benefit from strong network effects.

The value derived from network effects comes from three sources: exchange, staying power, and complementary benefits.

Exchange

Facebook for one person isn't much fun, and the first guy in the world with a fax machine doesn't have much more than a paperweight. But as each new Facebook friend or fax user comes online, a network becomes more valuable because its users can potentially communicate with more people. These examples show the importance of *exchange* in creating value. Every product or service subject to network effects fosters some kind of exchange. For firms leveraging technology, this might include anything you can represent in the ones and zeros of a data stream, such as movies, music, money, video games, and computer programs. And just about any standard that allows things to plug into one another, interconnect, or otherwise communicate will live or die based on its ability to snare network effects.



Exercise: Graph It

Some people refer to network effects by the name Metcalfe's Law. It got this name when, toward the start of the dot-com boom, Bob Metcalfe (the inventor of the Ethernet networking standard) wrote a column in *InfoWorld* magazine stating that the value of a network equals its number of users squared. What do you think of this formula? Graph the law with vertical access labeled "value" and horizontal access labeled "users." Do you think the graph is an accurate representation of what's happening in network effects? If so, why? If not, what do you think the graph really looks like?

Staying Power

Users don't want to buy a product or sign up for a service that's likely to go away, and a number of factors can halt the availability of an effort: a firm could go bankrupt, fail to attract a critical mass of user support, a rival may successfully invade its market and draw away current customers. Networks with greater numbers of users suggest a stronger staying power. The staying power, or long-term viability, of a product or service is particularly important for consumers of technology products. Consider that when someone buys a personal computer and makes a choice of Windows, Mac OS, or Linux, their investment over time usually greatly exceeds the initial price paid for the operating system. One invests in learning how to use a system, buying and installing software, entering preferences or other data, creating files—all of which mean that if a product isn't supported anymore, much of this investment is lost.

The concept of staying power (and the fear of being stranded in an unsupported product or service) is directly related to switching costs (the cost a consumer incurs when moving from one product to another) and switching costs can strengthen the value of network effects as a strategic

asset. The higher the value of the user's overall investment, the more they're likely to consider the staying power of any offering before choosing to adopt it. Similarly, the more a user has invested in a product, the less likely he or she is to leave.

Switching costs also go by other names. You might hear the business press refer to products (particularly Web sites) as being "sticky" or creating "friction." Others may refer to the concept of "lock-in." And the elite Boston Consulting Group is really talking about a firm's switching costs when it refers to how well a company can create customers who are "barnacles" (that are tightly anchored to the firm) and not "butterflies" (that flutter away to rivals). The more friction available to prevent users from migrating to a rival, the greater the switching costs. And in a competitive market where rivals with new innovations show up all the time, that can be a very good thing!

How Important Are Switching Costs to Microsoft?

"It is this switching cost that has given our customers the patience to stick with Windows through all our mistakes, our buggy drivers, our highTCO [total cost of ownership], our lack of a sexy vision at times, and many other difficulties [...] Customers constantly evaluate other desktop platforms, [but] it would be so much work to move over that they hope we just improve Windows rather than force them to move. [...] In short, without this exclusive franchise [meaning Windows] we would have been dead a long time ago."

Comments from a Microsoft General Manager in a memo to Bill Gates ^[1]

Complementary Benefits

Complementary benefits are those products or services that add additional value to the network. These products might include "how-to" books, software add-ons, even labor. You'll find more books on auctioning over

eBay, more virtual storefronts in Second Life, and more accountants that know Excel, than on any of their rivals. Why? Book authors, Second Life partners, and accountants invest their time where they're likely to reach the biggest market and get the greatest benefit. In auctions, virtual worlds, and spreadsheet software, eBay, Second Life, and Excel each dwarf their respective competition.

Products and services that encourage others to offer complementary goods are sometimes called platforms. [2] Allowing other firms to contribute to your platform can be a brilliant strategy, because those firms will spend *their* time and money to enhance *your* offerings. Consider the billion-dollar hardware ecosystem that Apple has cultivated around the iPod. There are over ninety brands selling some 280 models of iPod speaker systems. [3] Thirty-four auto manufacturers now trumpet their cars as being iPod-ready, many with in-car docking stations and steering wheel iPod navigation systems. Each add-on enhances the value of choosing an iPod over a rival like the Microsoft Zune. And now with the App Store for the iPhone and iPod Touch, Apple is doing the same thing with software add-ons. In less than a year after its introduction, the iTunes App store boasted over fifty thousand applications, collectively downloaded over one billion times.

The iPod Economy

These three value-adding sources—*exchange*, *staying power*, and *complementary benefits*—often work together to reinforce one another in a way that makes the network effect even stronger. When users *exchanging* information attract more users, they can also attract firms offering *complementary* products. When developers of complementary products invest time writing software—and users install, learn, and customize these products—switching costs are created that

enhance the *staying power* of a given network. From a strategist's perspective this can be great news for dominant firms in markets where network effects exist. The larger your network, the more difficult it becomes for rivals to challenge your leadership position.

KEY TAKEAWAYS

- Products and services subject to network effects get their value from exchange, perceived staying power, and complementary products and services. Tech firms and services that gain the lead in these categories often dominate all rivals.
- Many firms attempt to enhance their network effects by creating a platform for the development of third-party products and services that enhance the primary offering.

QUESTIONS AND EXERCISES

1. What are the factors that contribute to the value created by network effects?
2. Why is staying power particularly important to many technology products and services?
3. Think about the kinds of technology products that you own that are subject to network effects. What sorts of exchange do these products leverage (e.g., information, money, software, or other media)?
4. Think about the kinds of technology projects you own. What sorts of switching costs are inherent in each of these? Are these strong switching costs or weak switching costs? What would it take for you to leave one of these services and use a rival? How might a competitor try to lessen these switching costs to persuade you to adopt their product?
5. Which other terms are sometimes used to describe the phenomenon of switching costs?

6. Think about the kinds of technology products that you own that are subject to network effects. What sorts of complementary benefits are available for these products? Are complementary benefits strong or weak (meaning, do people choose the product primarily based on these benefits, or for some other reason)?
7. Identify firms that you believe have built a strong platform. Can you think of firms that have tried to develop a platform, but have been less successful? Why do you suppose they have struggled?

5.3 One-Sided or Two-Sided Markets?

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Recognize and distinguish between one-sided and two-sided markets.
2. Understand same-side and cross-side exchange benefits.

Understanding Network Structure

To understand the key sources of network value, it's important to recognize the structure of the network. Some networks derive most of their value from a single class of users. An example of this kind of network is instant messaging (IM). While there might be some add-ons for the most popular IM tools, they don't influence most users' choice of an IM system. You



pretty much choose one IM tool over another based on how many of your contacts you can reach. Economists would call IM a one-sided market (a market that derives most of its value from a single class of users), and the network effects derived from IM users attracting more IM users as being same-side exchange benefits (benefits derived by interaction among members of a single class of participant).

But some markets are comprised of two distinct categories of network participant. Consider video games. People buy a video game console largely based on the number of really great games available for the system.

Software developers write games based on their ability to reach the greatest number of paying customers, and so they're most likely to write for the most popular consoles, first. Economists would call this kind of network a two-sided market (network markets comprised of two distinct categories of participant, both of which that are needed to deliver value for the network to work). When an increase in the number of users on one side of the market (say console owners) creates a rise in the other side (software developers), that's called a cross-side exchange benefit.

The Positive Feedback Loop of Network Effects

IM is considered a one-sided market, where the value-creating, positive-feedback loop of network effects comes mostly from same-side benefits from a single group (IM members who attract other IM members who want to communicate with them). Video game consoles, however, are considered a two-sided network, where significant benefits come from two distinct classes of users that add value from cross-side benefits by attracting their opposite group. In the game console market, more users of a console attract more developers who write more software for that console, and that attracts more users. Game availability is the main reason the Sony PlayStation 2 dominated over the original Xbox.



It is possible that a network may have both same-side and cross-side benefits. Xbox 360 benefits from cross-side benefits in that more users of that console attract more developers writing more software titles and vice versa. However, the Xbox Live network that allows users to play against each other has same-side benefits. If your buddies use Xbox Live and you want to play against them, you're more likely to buy an Xbox.

KEY TAKEAWAYS

- In one-sided markets, users gain benefits from interacting with a similar category of users (think instant messaging, where everyone can send and receive messages to one another).
- In two-sided markets, users gain benefits from interacting with a separate, complementary class of users (e.g., in the video game industry console owners are attracted to platforms with the most games, while innovative developers are attracted to platforms that have the most users).

QUESTIONS AND EXERCISES

1. What is the difference between same-side exchange benefits and cross-side exchange benefits?
2. What is the difference between a one-sided market and a two-sided market?
3. Give examples of one-sided and two-sided markets.
4. Identify examples of two-sided markets where both sides pay for a product or service. Identify examples where only one side pays. What factors determine who should pay? Does paying have implications for the establishment and growth of a network effect? What might a firm do to encourage early network growth?
5. The Apple iPhone Developer Program provides developers access to the App Store where they can distribute their free or commercial applications

to millions of iPhone and iPod Touch customers. Would the iPhone market be considered a one or two-sided market?

5.4 How Are These Markets Different?

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand how competition in markets where network effects are present differ from competition in traditional markets.
2. Understand the reasons why it is so difficult for late-moving, incompatible rivals to compete in markets where a dominant, proprietary standard is present.

When network effects play a starring role, competition in an industry can be fundamentally different than in conventional, nonnetwork industries. First, network markets experience *early, fierce competition*. The positive-feedback loop inherent in network effects—where the biggest networks become even bigger—causes this. Firms are very aggressive in the early stages of these industries because once a leader becomes clear, *bandwagons* form, and new adopters begin to overwhelmingly favor the leading product over rivals, tipping the market in favor of one dominant firm or standard. This tipping can be remarkably swift. Once the majority of major studios and retailers began to back Blu-ray over HD-DVD, the latter effort folded within weeks.

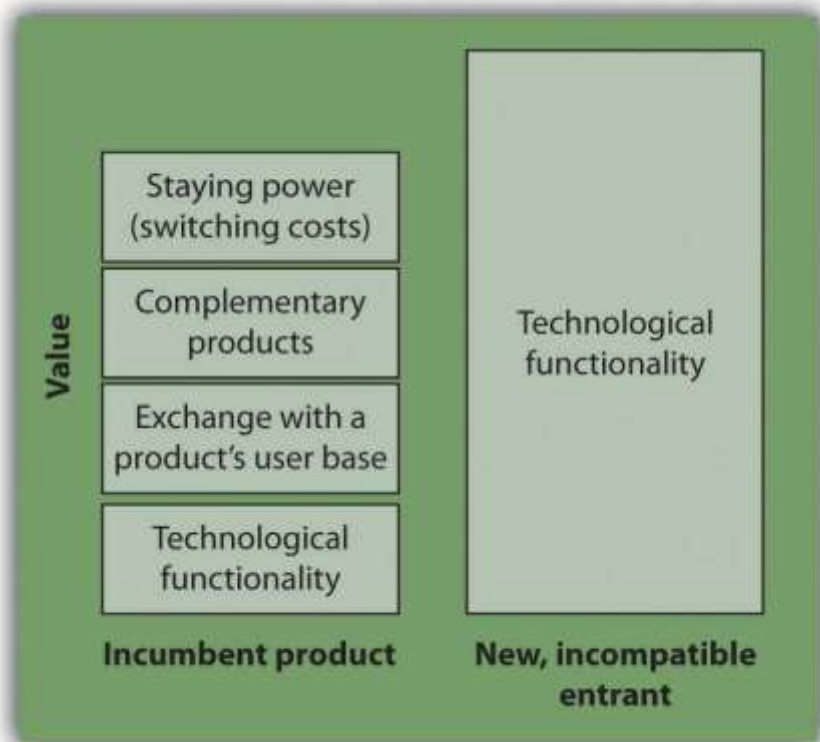


These markets are also often winner-take-all or winner-take-most, *exhibiting monopolistic tendencies* where one firm dominates all rivals. Look at all of the examples listed so far—in nearly every case the dominant player has a market share well ahead of all competitors. When, during the U.S. Microsoft antitrust trial, Judge Thomas Penfield Jackson declared Microsoft to be a monopoly (a market where there are many buyers but only one dominant seller), the collective response should have been “of course.” Why? The *natural state* of a market where network effects are present (and this includes operating systems and Office software) is for there to be one major player. Since bigger networks offer more value, they can charge customers more. Firms with a commanding network effects advantage may also enjoy substantial bargaining power over partners. For example, Apple, which controls over 75 percent of digital music sales, for years was able to dictate song pricing, despite the tremendous protests of the record labels. ^[1] In fact, Apple’s stranglehold was so strong that it leveraged bargaining power even though the “Big Four” record labels (Universal, Sony, EMI, and Warner) were themselves an oligopoly (a market dominated by a small number of powerful sellers) that together provide over 85 percent of music sold in the United States.

Finally, it’s important to note that the best product or service doesn’t always win. PlayStation 2 dominated the video console market over the original Xbox, despite the fact that nearly every review claimed the Xbox was hands-down a more technically superior machine. Why were users willing to choose an inferior product (PS2) over a superior one (Xbox)? The power of network effects! PS2 had more users, which attracted more developers offering more games.

Figure 5.3 Battling a leader with network effects is tough. ^[2]





This last note is a critical point to any newcomer wishing to attack an established rival. Winning customers away from a dominant player in a network industry isn't as easy as offering a product or service that is better. Any product that is incompatible with the dominant network has to exceed the value of the technical features of the leading player, plus (since the newcomer likely starts without any users or third-party product complements) the value of the incumbent's exchange, switching cost, and complementary product benefit (see [Figure 5.3 "Battling a leader with network effects is tough."](#)). And the incumbent must not be able to easily copy any of the newcomer's valuable new innovations; otherwise the dominant firm will quickly match any valuable improvements made by rivals. As such, technological leap-frogging, or competing by offering a superior generation of technology, can be really tough. [3]

Is This Good for Innovation?

Critics of firms that leverage proprietary standards for market dominance often complain that network effects are bad for innovation. But this statement isn't entirely true. While network effects limit competition *against* the dominant standard, innovation *within* a standard may actually blossom. Consider Windows. Microsoft has a huge advantage in the desktop operating system market, so few rivals try to compete with it. Apple and the open-source Linux are the firm's only credible rivals, and both have tiny market shares. But the dominance of Windows is a magnet for developers to innovate within the standard. Programmers with novel ideas are willing to make the investment in learning to write software for Windows because they're sure that a Windows version can be used by the overwhelming majority of computer users.

By contrast, look at the mess in the mobile phone market. With so many different handsets offering different screen sizes, running different software, having different key layouts, and working on different carrier networks, writing a game that's accessible by the majority of users is nearly impossible. Glu Mobile, a maker of online games, launched fifty-six reengineered builds of Monopoly to satisfy the diverse requirements of just one telecom carrier. [4] As a result, entrepreneurs with great software ideas for the mobile market are deterred because writing, marketing, and maintaining multiple product versions is both costly and risky.

KEY TAKEAWAYS

- Unseating a firm that dominates with network effects can be extremely difficult, especially if the newcomer is not compatible with the established leader. Newcomers will find their technology will need to be so good that it must leapfrog not only the value of the established firm's tech, but also the perceived stability of the dominant firm, the exchange

benefits provided by the existing user base, and the benefits from any product complements. For evidence, just look at how difficult it's been for rivals to unseat the dominance of Windows.

- Because of this, network effects might limit the number of rivals that challenge a dominant firm. But the establishment of a dominant standard may actually encourage innovation within the standard, since firms producing complements for the leader have faith the leader will have staying power in the market.

QUESTIONS AND EXERCISES

1. How is competition in markets where network effects are present different from competition in traditional markets?
2. What are the reasons it is so difficult for late-moving, incompatible rivals to compete in markets where a dominant, proprietary standard is present? What is technological leap-frogging and why is it so difficult to accomplish?
3. Does it make sense to try to prevent monopolies in markets where network effects exist?
4. Are network effects good or bad for innovation? Explain.
5. What is the relationship between network effects and the bargaining power of participants in a network effects “ecosystem”?
6. Cite examples where the best technology did not dominate a network effects-driven market.

5.5 Competing When Network Effects Matter

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Plot strategies for competing in markets where network effects are present, both from the perspective of the incumbent firm and the new market entrant.
2. Give examples of how firms have leveraged these strategies to compete effectively.

Why do you care whether networks are one-sided, two-sided, or some sort of hybrid? Well, when crafting your plan for market dominance, it's critical to know if network effects exist, how strong they might be, where they come from, and how they might be harnessed to your benefit. Here's a quick rundown of the tools at your disposal when competing in the presence of network effects.

Strategies for Competing in Markets with Network Effects (Examples in Parentheses)

- Move early (Yahoo! Auctions in Japan)
- Subsidize product adoption (PayPal)
- Leverage viral promotion (Skype; Facebook feeds)
- Expand by redefining the market to bring in new categories of users (Nintendo Wii) or through convergence (iPhone).
- Form alliances and partnerships (NYCE vs. Citibank)
- Establish distribution channels (Java with Netscape; Microsoft bundling Media Player with Windows)
- Seed the market with complements (Blu-ray; Nintendo)
- Encourage the development of complementary goods—this can include offering resources, subsidies, reduced fees, market research, development kits, venture capital (Facebook fbFund).



- Maintain backward compatibility (Apple’s Mac OS X Rosetta translation software for PowerPC to Intel)
- For rivals, be compatible with larger networks (Apple’s move to Intel; Live Search Maps)
- For incumbents, constantly innovate to create a moving target and block rival efforts to access your network (Apple’s efforts to block access to its own systems)
- For large firms with well-known followers, make preannouncements (Microsoft)

Move Early

In the world of network effects, this is a biggie. Being first allows your firm to start the network effects snowball rolling in your direction. In Japan, worldwide auction leader eBay showed up just five months after Yahoo! but was never able to mount a credible threat and ended up pulling out of the market. Being just five months late cost eBay billions in lost sales, and the firm eventually retreated, acknowledging it could never unseat Yahoo!’s network effects lead.

Another key lesson from the loss of eBay Japan? Exchange depends on the ability to communicate! eBay’s huge network effects in the United States and elsewhere didn’t translate to Japan because most Japanese aren’t comfortable with English, and most English speakers don’t know Japanese. The language barrier made Japan a “greenfield” market with no dominant player, and Yahoo!’s early move provided the catalyst for victory.

Timing is often critical in the video game console wars, too. Sony’s PlayStation 2 enjoyed an eighteen-month lead over the technically superior Xbox (as well as Nintendo’s Game Cube). That time lead helped to create what for years was the single most profitable division at Sony. By contrast, the technically superior PS3 showed up months after Xbox 360 and at



roughly the same time as the Nintendo Wii, and has struggled in its early years, racking up multibillion-dollar losses for Sony. ^[1]

What If Microsoft Threw a Party and No One Showed Up?

Microsoft launched the Zune media player with features that should be subject to network effects—the ability to share photos and music by wirelessly “squirting” content to other Zune users. The firm even promoted Zune with the tagline “Welcome to the Social.” Problem was the Zune Social was a party no one wanted to attend. The late-arriving Zune garnered a market share of just 3 percent, and users remained hard pressed to find buddies to leverage these neat social features. ^[2] A cool idea does not make a network effect happen.

Subsidize Adoption

Starting a network effect can be tough—there’s little incentive to join a network if there’s no one in the system to communicate with. In one admittedly risky strategy, firms may offer to subsidize initial adoption in hopes that network effects might kick in shortly after. Subsidies to adopters might include a price reduction, rebate, or other giveaways. PayPal, a service that allows users to pay one another using credit cards, gave users a modest rebate as a sign-up incentive to encourage adoption of its new effort (in one early promotion, users got back fifteen dollars when spending their first thirty dollars). This brief subsidy paid to early adopters paid off handsomely. eBay later tried to enter the market with a rival effort, but as a late mover its effort was never able to overcome PayPal’s momentum. PayPal was eventually purchased by eBay for \$1.5 billion, and the business unit is now considered one of eBay’s key drivers of growth and profit.

When Even Free Isn't Good Enough

Subsidizing adoption after a rival has achieved dominance can be an uphill battle, and sometimes even offering a service for free isn't enough to combat the dominant firm. When Yahoo! introduced a U.S. auction service to compete with eBay, it initially didn't charge sellers at all (sellers typically pay eBay a small percentage of each competed auction). The hope was that with the elimination of seller fees, enough sellers would jump from eBay to Yahoo! helping the late-mover catch up in the network effect game.

But eBay sellers were reluctant to leave for two reasons. First, there weren't enough buyers on Yahoo! to match the high bids they earned on much-larger eBay. Some savvy sellers played an arbitrage game where they'd buy items on Yahoo!'s auction service at lower prices and resell them on eBay, where more users bid prices higher.

Second, any established seller leaving eBay would give up their valuable "seller ratings," and would need to build their Yahoo! reputation from scratch. Seller ratings represent a critical switching cost, as many users view a high rating as a method for reducing the risk of getting scammed or receiving lower-quality goods.

Auctions work best for differentiated goods. While Amazon has had some success in peeling away eBay sellers who provide commodity products (a real danger as eBay increasingly relies on fixed-price sales), eBay's dominant share of the online auction market still towers over all rivals. [3] While there's no magic in the servers used to create eBay, the early use of technology allowed the firm to create both network effects and switching costs—a dual strategic advantage that has given it a hammerlock on auctions even as others have attempted to mimic its service and undercut its pricing model.

Leverage Viral Promotion

Since all products and services foster some sort of exchange, it's often possible to leverage a firm's customers to promote the product or service. Internet calling service Skype has over 300 million registered users yet has spent almost nothing on advertising. Most Skype users were recruited by others who shared the word on free and low-cost Internet calls. Within Facebook, feeds help activities to spread virally (see [Chapter 7 "Facebook: Building a Business from the Social Graph"](#)). Feeds blast updates on user activities on the site, acting as a catalyst for friends to join groups and load applications that their buddies have already signed up for.

Expand by Redefining the Market

If a big market attracts more users (and in two-sided markets, more complements), why not redefine the space to bring in more users? Nintendo did this when launching the Wii. While Sony and Microsoft focused on the graphics and raw processing power favored by hard-core male gamers, Nintendo chose to develop a machine to appeal to families, women, and age groups that normally shunned alien shoot 'em ups. By going after a bigger, redefined market, Nintendo was able to rack up sales that exceeded the Xbox 360, even though it followed the system by twelve months. ^[4]

Seeking the Blue Ocean

Reggie Fils-Aimé, the President of Nintendo of America, describes the Wii Strategy as a Blue Ocean effort. ^[5] The concept of blue ocean strategy was popularized by European Institute of Business Administration (INSEAD) professors W. Chan Kim and Renée Mauborgne (authors of a book with the same title). ^[6] The idea—instead of competing in *blood-red* waters where

the sharks of highly competitive firms vie for every available market scrap, firms should seek the *blue waters* of uncontested, new market spaces. For Nintendo, the granny gamers, moms, and partygoers who flocked to the Wii represented an undiscovered feast in the Blue Ocean. Talk about new markets! Consider that the best-selling video game at the start of 2009 was Wii Fit—a genre-busting title that comes with a scale so you can weigh yourself each time you play! That’s a far cry from Grand Theft Auto IV, the title ranking fifth in 2008 sales, and trailing four Wii-only exclusives. Blue ocean strategy often works best when combined with strategic positioning described in [Chapter 2 "Strategy and Technology"](#). If an early mover into a blue ocean can use this lead to create defensible assets for sustainable advantage, late moving rivals may find markets unresponsive to their presence.

Market expansion sometimes puts rivals who previously did not compete on a collision course as markets undergo convergence (when two or more markets, once considered distinctly separate, begin to offer similar features and capabilities). Consider the market for portable electronic devices. Separate product categories for media players, cameras, gaming devices, phones, and global positioning systems (GPS) are all starting to merge. Rather than cede its dominance as a media player, Apple leveraged a strategy known as envelopment, where a firm seeks to make an existing market a subset of its product offering. Apple deftly morphed the iPod into the iPhone, a device that captures all of these product categories in one device. But the firm went further; the iPhone is wi-fi capable, offers browsing, e-mail, and an application platform based on a scaled-down version of the same OS X operating system used in Macintosh computers. As a “Pocket Mac,” the appeal of the device broadened beyond just the phone or music player markets, and within two quarters of launch, iPhone became the second leading smart phone in North America—outpacing



Palm, Microsoft, Motorola and every other rival, except RIM's Blackberry. [7]

Alliances and Partnerships

Firms can also use partnerships to grow market share for a network. Sometimes these efforts bring rivals together to take out a leader. In a classic example, consider ATM networks. Citibank was the first major bank in New York City to offer a large ATM network. But the Citi network was initially proprietary, meaning customers of other banks couldn't take advantage of Citi ATMs. Citi's innovation was wildly popular and being a pioneer in rolling out cash machines helped the firm grow deposits fourfold in just a few years. Competitors responded with a partnership. Instead of each rival bank offering another incompatible network destined to trail Citi's lead, competing banks agreed to share their ATM operations through NYCE (New York Cash Exchange). While Citi's network was initially the biggest, after the NYCE launch a Chase bank customer could use ATMs at a host of other banks that covered a geography far greater than Citi offered alone. Network effects in ATMs shifted to the rival bank alliance, Citi eventually joined NYCE and today, nearly every ATM in the United States carries a NYCE sticker.

Google has often pushed an approach to encourage rivals to cooperate to challenge a leader. Its Open Social standard for social networking (endorsed by MySpace, LinkedIn, Bebo, Yahoo! and others) is targeted at offering a larger alternative to Facebook's more closed efforts (see [Chapter 7 "Facebook: Building a Business from the Social Graph"](#)), while its Android open source mobile phone operating system has gained commitments from many handset makers that collectively compete with Apple's iPhone.



Share or Stay Proprietary?

Defensive moves like the ones above are often meant to diffuse the threat of a proprietary rival. Sometimes firms decide from the start to band together to create a new, more open standard, realizing that collective support is more likely to jumpstart a network than if one firm tried to act with a closed, proprietary offering. Examples of this include the coalitions of firms that have worked together to advance standards like Bluetooth and wi-fi. While no single member firm gains a direct profit from the sale of devices using these standards, the standard's backers benefit when the market for devices expands as products become more useful because they are more interoperable.

Leverage Distribution Channels

Firms can also think about novel ways to distribute a product or service to consumers. Sun faced a challenge when launching the Java programming language—no computers could run it. In order for Java to work, computers need a little interpreter program called the Java Virtual Machine (JVM). Most users weren't willing to download the JVM if there were no applications written in Java, and no developers were willing to write in Java if no one could run their code. Sun broke the log jam when it *bundled* the JVM with Netscape's browser. When millions of users downloaded Netscape, Sun's software snuck in, almost instantly creating a platform of millions for would-be Java developers. Today, even though Netscape has failed, Sun's Java remains one of the world's most popular programming languages. Indeed, Java was cited as one of the main reasons for Oracle's 2009 bid to acquire Sun, with Oracle's CEO saying the language represented “the single most important software asset we have ever acquired.” [8]



As mentioned in [Chapter 2 "Strategy and Technology"](#), Microsoft is in a particularly strong position to leverage this approach. The firm often bundles its new products into its operating systems, Office suite, Internet Explorer browser, and other offerings. The firm used this tactic to transform once market-leader Real Networks into an also-ran in streaming audio. Within a few years of bundling Windows Media Player (WMP) with its other products, WMP grabbed the majority of the market, while Real's share had fallen to below 10 percent. ^[9]

Caution is advised, however. Some international regulatory efforts consider product bundling by dominant firms to be anticompetitive. European regulators have forced Microsoft to unbundle Windows Media Player from its operating system and to provide a choice of browsers alongside Internet Explorer.

Antitrust: Real Versus Microsoft

From October 2001 to March 2003, Microsoft's bundling of Windows Media Player in versions of its operating system ensured that the software came preinstalled on nearly all of the estimated 207 million new PCs shipped during that period. By contrast, Real Networks' digital media player was preinstalled on less than 2 percent of PCs. But here's the kicker that got to regulators (and Real): Microsoft's standard contract with PC manufacturers "prevented them not only from removing the Windows Media Player, but even [from] providing a desktop icon for Real Networks."^[10] While network effects create monopolies, governments may balk at allowing a firm to leverage its advantages in ways that are designed to deliberately keep rivals from the market.

Seed the Market

When Sony launched the PS3, it subsidized each console by selling at a price estimated at three hundred dollars below unit cost. ^[11] Subsidizing



consoles is a common practice in the video game industry—game player manufacturers usually make most of their money through royalties paid by game developers. But Sony’s subsidy had an additional benefit for the firm—it helped sneak a Blu-ray DVD player into every home buying a PS3. Since Sony is also a movie studio and manufacturer of DVD players and other consumer electronics, it had a particularly strong stake in encouraging the adoption of Blu-ray over rival HD-DVD.

Giving away products for half of a two-sided market is an extreme example of this kind of behavior, but it’s often used. In two-sided markets, you charge the one who will pay. Adobe gives away the Acrobat reader to build a market for the sale of software that creates Acrobat files. Firms with Yellow Page directories give away countless copies of their products, delivered straight to your home, in order to create a market for selling advertising. And Google does much the same by providing free, ad-supported search.

Encourage the Development of Complementary Goods

There are several ways to motivate others to create complementary goods for your network. These efforts often involve some form of developer subsidy or other free or discounted service. A firm may charge lower royalties or offer a period of royalty-free licensing. It can also offer free software development kits (SDKs), training programs, co-marketing dollars, or even startup capital to potential suppliers. Microsoft and Apple both allow developers to sell their products online through Xbox Live Marketplace and iTunes, respectively. This channel lowers developer expenses by eliminating costs associated with selling physical inventory in brick and mortar stores and can provide a free way to reach millions of potential consumers without significant promotional spending.

Venture funds can also prompt firms to create complementary goods. Facebook announced it would spur development for the site in part by



administering the fbFund, which initially pledged ten million dollars in startup funding (in allotments of up to \$250,000 each) to firms writing applications for its platform.

Leverage Backward Compatibility

Those firms that control a standard would also be wise to ensure that new products have backward compatibility with earlier offerings. If not, they reenter a market at installed-base zero and give up a major source of advantage—the switching costs built up by prior customers. For example, when Nintendo introduced its 16-bit S.N.E.S. system, it was incompatible with the firm’s highly successful prior generation 8-bit model. Rival Sega, which had entered the 16-bit market two years prior to Nintendo, had already built up a large library of 16-bit games for its system. Nintendo entered with only its debut titles, and no ability to play games owned by customers of its previous system, so there was little incentive for existing Nintendo fans to stick with the firm. ^[12]

Backward compatibility was the centerpiece of Apple’s strategy to revitalize the Macintosh through its move to the Intel microprocessor. Intel chips aren’t compatible with the instruction set used by the PowerPC processor used in earlier Mac models. Think of this as two incomprehensible languages—Intel speaks French, PowerPC speaks Urdu. To ease the transition, Apple included a free software-based adaptor, called Rosetta, that automatically emulated the functionality of the old chip on all new Macs (a sort of Urdu to French translator). By doing so, all new Intel Macs could use the base of existing software written for the old chip; owners of PowerPC Macs were able to upgrade while preserving their investment in old software; and software firms could still sell older programs while they rewrote applications for new Intel-based Macs.



Even more significant, since Intel is the same standard used by Windows, Apple developed a free software adaptor called Boot Camp that allowed Windows to be installed on Macs. Boot Camp (and similar solutions by other vendors) dramatically lowered the cost for Windows users to switch to Macs. Within two years of making the switch, Mac sales skyrocketed to record levels. Apple now boasts a commanding lead in notebook sales to the education market, [13] and a survey by Yankee Group found that 87 percent of corporations were using at least some Macintosh computers, up from 48 percent at the end of the PowerPC era two years earlier. [14]

Rivals: Be Compatible with the Leading Network

Companies will want to consider making new products compatible with the leading standard. Microsoft's Live Maps and Virtual Earth 3D arrived late to the Internet mapping game. Users had already put in countless hours building resources that meshed with Google Maps and Google Earth. But by adopting the same keyhole markup language (KML) standard used by Google, Microsoft could, as TechCrunch put it, "drink from Google's milkshake." Any work done by users for Google in KML could be used by Microsoft. Voilà, an instant base of add-on content!

Incumbents: Close Off Rival Access and Constantly Innovate

Oftentimes firms that control dominant networks will make compatibility difficult for rivals who try to connect with their systems. America Online has been reluctant to open up its instant messaging tool to rivals, and Skype for years had been similarly closed to non-Skype clients.

Firms that constantly innovate make it particularly difficult for competitors to become compatible. Again, we can look to Apple as an example of these concepts in action. While Macs run Windows, Windows computers can't



run Mac programs. Apple has embedded key software in Mac hardware, making it tough for rivals to write a software emulator like Boot Camp that would let Windows PCs drink from the Mac milkshake. And if any firm gets close to cloning Mac hardware, Apple sues. The firm also modifies software on other products like the iPhone and iTunes each time wily hackers tap into closed aspects of its systems. Even if firms create adaptors that emulate a standard, a firm that constantly innovates creates a moving target that's tough for others to keep up with.

Apple has been far more aggressive than Microsoft in introducing new versions of its software. Since the firm never stays still, would-be cloners never get enough time to create an emulator that runs the latest Apple software.

Large, Well-Known Followers: Preannouncements

Large firms that find new markets attractive but don't yet have products ready for delivery might *preannounce* efforts in order to cause potential adaptors to sit on the fence, delaying a purchasing decision until the new effort rolls out. Preannouncements only work if a firm is large enough to pose a credible threat to current market participants. Microsoft, for example, can cause potential customers to hold off on selecting a rival because users see that the firm has the resources to beat most players (suggesting staying power). Statements from startups however often lack credibility to delay user purchases. The tech industry acronym for the impact firms try to impart on markets through preannouncements is *FUD* for fear, uncertainty, and doubt.

The Osborne Effect

Preannouncers beware. Announce an effort too early and a firm may fall victim to what's known as "The Osborne Effect." It's been suggested that

portable computer manufacturer Osborne Computer announced new models too early. Customers opted to wait for the new models, so sales of the firm's current offerings plummeted. While evidence suggests that Osborne's decline had more to do with rivals offering better products, the negative impact of preannouncements has hurt a host of other firms.^[15] Among these, Sega, which exited the video game console market entirely after preannouncements of a next-generation system killed enthusiasm for its Saturn console. ^[16]

Too Much of a Good Thing?

When network effects are present, more users attract more users. That's a good thing as long as a firm can earn money from this virtuous cycle. But sometimes a network effect attracts too many users and a service can be so overwhelmed it becomes unusable. These so-called congestion effects occur when increasing numbers of users lower the value of a product or service. This most often happens when a key resource become increasingly scarce. Users of the game Ultima were disappointed in an early online version that launched without enough monsters to fight or server power to handle the crush of fans. Facebook users with a large number of friends may also find their attention is a limited resource, as feeds push so much content that it becomes difficult to separate interesting information from the noise of friend actions.

And while network effects can attract positive complementary products, a dominant standard may also be the first place where virus writers and malicious hackers choose to strike.

Feel confident! Now you've got a solid grounding in network effects, the key resource leveraged by some of the most dominant firms in technology. And these concepts apply beyond the realm of tech, too. Network effects can explain phenomena ranging from why some stock markets are more

popular than others to why English is so widely spoken, even among groups of nonnative speakers. On top of that, the strategies explored in the last half of the chapter show how to use these principles to sniff out, create, and protect this key strategic asset. Go forth, tech pioneer—opportunity awaits!

KEY TAKEAWAYS

- Moving early matters in network markets—firms that move early can often use that time to establish a lead in users, switching costs, and complementary products that can be difficult for rivals to match.
- Additional factors that can help a firm establish a network effects lead include subsidizing adoption; leveraging viral marketing, creating alliances to promote a product or to increase a service’s user base; redefining the market to appeal to more users; leveraging unique distribution channels to reach new customers; seeding the market with complements; encouraging the development of complements; and maintaining backward compatibility.
- Established firms may try to make it difficult for rivals to gain compatibility with their users, standards, or product complements. Large firms may also create uncertainty among those considering adoption of a rival by preannouncing competing products.

QUESTIONS AND EXERCISES

1. Is market entry timing important for network effects markets? Explain and offer an example to back up your point.
2. How might a firm subsidize adoption? Give an example.
3. Give an example of a partnership or alliance targeted at increasing network effects.
4. Is it ever advantageous for firms to give up control of a network and share it with others? Why or why not? Give examples to back up your point.



5. Do firms that dominate their markets with network effects risk government intervention? Why or why not? Explain through an example.
6. How did Sony seed the market for Blu-ray DVD players?
7. What does backward compatibility mean and why is this important? What happens if a firm is not backward compatible?
8. What tactic did Apple use to increase the acceptability of the Mac platform to a broader population of potential users?
9. How has Apple kept clones at bay?
10. What are preannouncements? What is the danger in announcing a product too early? What is the term for negative impacts from premature product announcements?
11. How did PayPal subsidize adoption?
12. Name two companies that leveraged viral promotion to compete.
13. Name a product that is the result of the *convergence* of media players, cameras, and phones.
14. What is bundling? What are the upsides and downsides of bundling?
15. Why does Adobe allow the *free* download of Acrobat Reader?
16. What tactic might an established firm employ to make it impossible, or at least difficult, for a competitor to gain access to, or become compatible with, their product or service?
17. How do Apple, Microsoft, and Facebook encourage the development of complementary products?
18. What is the “congestion effect”? Give an example.
19. Do network effects apply in nontech areas? Give examples.

Chapter 6

Peer Production, Social Media, and Web 2.0



6.1 Introduction

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Recognize the unexpected rise and impact of social media and peer production systems, and understand how these services differ from prior generation tools.
2. List the major classifications of social media services.

Over the past few years a fundamentally different class of Internet services has attracted users, made headlines, and increasingly garnered breathtaking market valuations. Often referred to under the umbrella term “Web 2.0,” these new services are targeted at harnessing the power of the Internet to empower users to collaborate, create resources, and share information in a distinctly different way from the static Web sites and transaction-focused storefronts that characterized so many failures in the dot-com bubble. Blogs, wikis, social networks, photo and video sharing sites, and tagging systems all fall under the Web 2.0 moniker, as do a host of supporting technologies and related efforts.

The term Web 2.0 is a tricky one because like so many popular technology terms there’s not a precise definition. Coined by publisher and pundit Tim O’Reilly in 2003, techies often joust over the breadth of the Web 2.0 umbrella and over whether Web 2.0 is something new or simply an extension of technologies that have existed since the creation of the Internet. These arguments aren’t really all that important. What is significant is how quickly the Web 2.0 revolution came about, how unexpected it was, and how deeply impactful these efforts have become. Some of the sites and services that have evolved and their Web 1.0 origins are listed in [Table 6.1 "Web 1.0 versus Web 2.0"](#). ^[1]



Table 6.1 Web 1.0 versus Web 2.0

Web 1.0		Web 2.0
DoubleClick	→	Google AdSense
Ofoto	→	Flickr
Akamai	→	BitTorrent
mp3.com	→	Napster
Britannica Online	→	Wikipedia
personal Web sites	→	blogging
evite	→	upcoming.org and E.V.D.B.
domain name speculation	→	search engine optimization
page views	→	cost per click
screen scraping	→	Web services
publishing	→	participation
content management systems	→	wikis
directories (taxonomy)	→	tagging (“folksonomy”)
stickiness	→	syndication
instant messaging	→	Twitter
Monster.com	→	LinkedIn

To underscore the speed with which Web 2.0 arrived on the scene, and the impact of leading Web 2.0 services, consider the following efforts:

- According to a spring 2008 report by Morgan Stanley, Web 2.0 services ranked as seven of the world’s top ten most heavily trafficked Internet sites



(YouTube, Live.com, MySpace, Facebook, Hi5, Wikipedia, and Orkut); only one of these sites (MySpace) was on the list in 2005. [2]


- With only seven full-time employees and an operating budget of less than one million dollars, Wikipedia has become the Internet's fifth most visited site on the Internet. [3] The site boasts well over fourteen million articles in over two hundred sixty different languages, all of them contributed, edited, and fact-checked by volunteers.
- Just two years after it was founded, MySpace was bought for \$580 million by Rupert Murdoch's News Corporation (the media giant that owns the *Wall Street Journal* and the Fox networks, among other properties). By the end of 2007, the site accounted for some 12 percent of Internet minutes and has repeatedly ranked as the most-visited Web site in the United States. [4]
- At rival Facebook, users in the highly sought after college demographic spend over thirty minutes a day on the site. A fall 2007 investment from Microsoft pegged the firm's overall value at fifteen billion dollars, a number that would make it the fifth most valuable Internet firm, despite annual revenues at the time of only one hundred fifty million dollars. [5]
- Just twenty months after its founding, YouTube was purchased by Google for \$1.65 billion. While Google struggles to figure out how to make profitable what is currently a money-losing resource hog (over twenty hours of video are uploaded to YouTube each minute) [6] the site has emerged as the Web's leading destination for video, hosting everything from apologies from JetBlue's CEO for service gaffes to questions submitted as part of the 2008 U.S. presidential debates. Fifty percent of YouTube's roughly three hundred million users visit the site at least once a week. [7]

•

Table 6.2 Major Web 2.0 Tools



	Description	Features	Technology Providers	Use Case Examples
Blogs	<p>Short for “Web log”—an online diary that keeps a running chronology of entries. Readers can comment on posts. Can connect to other blogs through blog rolls or trackbacks.</p> <p>Key uses: Share ideas, obtain feedback, mobilize a community.</p>	<p>Ease of use</p> <p>Reverse chronology</p> <p>Comment threads</p> <p>Persistence</p> <p>Searchability</p> <p>Tags</p> <p>Trackbacks</p>	<p>Blogger (Google)</p> <p>WordPress</p> <p>SixApart (TypePad and Movable Type)</p>	<p>Corporate Users:</p> <p>News outlets</p> <p>Sun</p> <p>Microsoft</p> <p>G.M.</p> <p>Kaiser Permanente</p>
Wikis	<p>A Web site that anyone can edit directly from within the browser. Key uses: Collaborate on common tasks</p>	<p>All changes are attributed</p> <p>A complete revision history is maintained</p>	<p>SocialText</p> <p>PBWorks</p> <p>Google Sites</p> <p>WetPaint</p> <p>Microsoft</p> <p>SharePoint</p> <p>Apple OS X Server</p>	<p>Corporate Users:</p> <p>Wikipedia.com</p> <p>Intuit</p> <p>Amazon</p> <p>eBay</p> <p>Pixar</p>



	Description	Features	Technology Providers	Use Case Examples
	or to create a common knowledge base.	Automatic notification of updates Searchability Tags Monitoring		
Electronic Social Network	Online community that allows users to establish a personal profile, link to other profiles (i.e., friends), and browse the connections of other, and communicate with members via messaging, posts, et cetera. Key Uses: Discover and reinforce affiliations. Identify experts.	Detailed personal profiles using multimedia Affiliations with groups Affiliations with individuals Messaging and public discussions Media sharing “Feeds” of recent activity among	MySpace Facebook LinkedIn SelectMinds LiveWorld IBM/Lotus Connections	Corporate Users: Deloitte Consulting Goldman-Sachs Reuters IBM Starbucks Dell



	Description	Features	Technology Providers	Use Case Examples
	Message individuals or groups. Share media.	members		
Microblogging	Short, asynchronous messaging system. Users send messages to “followers.” Key Uses: distributing time, sensitive information, sharing opinions, virally spreading ideas, running contests and promotions, soliciting feedback, providing customer support, tracking commentary on	140 character messages sent and received from mobile device. ability to respond publicly or privately can specify tags to classify discussion topics for easy searching and building comment threads follower	Twitter (overwhelming leader) SocialText Signals (organizational use) Yammer Salesforce.com	Corporate Users: Dell Starbucks Intuit Small businesses



	Description	Features	Technology Providers	Use Case Examples
	firms/products/issues, organizing protests.	lists		

Millions of users, billions of dollars, huge social impact, and these efforts weren't even on the radar of most business professionals when today's graduating college seniors first enrolled as freshmen. The trend demonstrates that even some of the world's preeminent thought leaders and business publications can be sideswiped by the speed of the Internet. Consider that when management guru Michael Porter wrote a piece titled, "Strategy and the Internet" at the end of the dot-com bubble, he lamented the high cost of building brand online, questioned the power of network effects, and cast a skeptical eye on ad-supported revenue models. Well, it turns out Web 2.0 efforts challenged *all* of these assumptions. Among the efforts above, all built brand on the cheap with little conventional advertising, and each owes their hypergrowth and high valuation to their ability to harness the network effect. In June, 2008, *BusinessWeek* also confessed to having an eye off the ball. In a cover story on social media, the magazine offered a mea culpa, confessing that while blogging was on their radar, editors were blind to the bigger trends afoot online, and underestimated the rise and influence of social networks, wikis, and other efforts. [8]

While the Web 2.0 moniker is a murky one, we'll add some precision to our discussion of these efforts by focusing on peer production, perhaps Web 2.0's most powerful feature, where users work, often collaboratively, to create content and provide services online. Web-based efforts that foster



peer production are often referred to as social media or *user-generated content* sites. These sites include blogs; wikis; social networks like Facebook and MySpace; communal bookmarking and tagging sites like Del.icio.us; media sharing sites like YouTube and Flickr; and a host of supporting technologies. And it's not just about media. Peer-produced services like Skype and BitTorrent leverage users' computers instead of a central IT resource to forward phone calls and video. This ability saves their sponsors the substantial cost of servers, storage, and bandwidth. Techniques such as crowdsourcing, where initially undefined groups of users band together to solve problems, create code, and develop services, are also a type of peer production. These efforts will be expanded on below, along with several examples of their use and impact.

KEY TAKEAWAYS

- A new generation of Internet applications is enabling consumers to participate in creating content and services online. Examples include Web 2.0 efforts such as social networks, blogs, and wikis, as well as efforts such as Skype and BitTorrent, which leverage the collective hardware of their user communities to provide a service.
- These efforts have grown rapidly, most with remarkably little investment in promotion. Nearly all of these new efforts leverage network effects to add value and establish their dominance and viral marketing to build awareness and attract users.
- Experts often argue whether Web 2.0 is something new or merely an extension of existing technologies. The bottom line is the magnitude of the impact of the current generation of services.
- Network effects play a leading role in enabling Web 2.0 firms. Many of these services also rely on ad-supported revenue models.

QUESTIONS AND EXERCISES

1. What distinguishes Web 2.0 technologies and services from the prior generation of Internet sites?
2. Several examples of rapidly rising Web 2.0 efforts are listed in this section. Can you think of other dramatic examples? Are there cautionary tales of efforts that may not have lived up to their initial hype or promise? Why do you suppose they failed?
3. Make your own list of Web 1.0 and Web 2.0 services and technologies. Would you invest in them? Why or why not?
4. In what ways do Web 2.0 efforts challenge the assumptions that Michael Porter made regarding Strategy and the Internet?



6.2 Blogs

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Know what blogs are and how corporations, executives, individuals, and the media use them.
2. Understand the benefits and risks of blogging.
3. Appreciate the growth in the number of blogs, their influence, and their capacity to generate revenue.

Blogs (short for Web logs) first emerged almost a decade ago as a medium for posting online diaries. (In a perhaps apocryphal story, *Wired* magazine claimed the term “Web log” was coined by Jorn Barger, a sometimes homeless, yet profoundly prolific, Internet poster). From humble beginnings, the blogging phenomenon has grown to a point where the number of public blogs tracked by Technorati (the popular blog index) has surpassed one hundred million. [1] This number is clearly a long tail phenomenon, loaded with niche content that remains “discoverable” through search engines and blog indexes. Trackbacks (third-party links back to original blog post), and blog rolls (a list of a blogger’s favorite sites—a sort of shout-out to blogging peers) also help distinguish and reinforce the reputation of widely read blogs.

The most popular blogs offer cutting-edge news and commentary, with postings running the gamut from professional publications to personal



diaries. While this cacophony of content was once dismissed, blogging is now a respected and influential medium. Consider that the political blog The Huffington Post is now more popular than all but eight newspaper sites and has a valuation higher than many publicly traded papers. ^[2]Keep in mind that this is a site without the sports, local news, weather, and other content offered by most papers. Ratings like this are hard to achieve—most bloggers can't make a living off their musings. But among the elite ranks, killer subscriber numbers are a magnet for advertisers. Top blogs operating on shoestring budgets can snare several hundred thousand dollars a month in ad revenue. ^[3] Most start with ad networks like Google AdSense, but the most elite engage advertisers directly for high-value deals and extended sponsorships.

Top blogs have begun to attract well-known journalists away from print media. The Huffington Post hired a former *Washington Post* editor Lawrence Roberts to head the site's investigative unit. The popular blog TechCrunch hired Erick Schonfeld away from Time Warner's business publishing empire, while Schonfeld's cohort Om Malik founded another highly ranked tech industry blog, GigaOM. And sometimes this drift flows the other direction. Robert Scoble, a blogger who made his reputation as the informal online voice for Microsoft, was hired to run social media at *Fast Company* magazine.

Senior executives from many industries have also begun to weigh in with online ruminations, going directly to the people without a journalist filtering their comments. Sun Microsystems's Jonathan Schwartz, GM's Bob Lutz, and Paul Levy (CEO of healthcare quality leader Beth Israel Deaconess Medical Center) use their blogs for purposes that include a combination of marketing, sharing ideas, gathering feedback, press response, and image shaping. Blogs have the luxury of being more topically focused than traditional media, with no limits on page size, word count, or



publication deadline. Some of the best examples engage new developments in topic domains much more quickly and deeply than traditional media. For example, it's not uncommon for blogs focused on the law or politics to provide a detailed dissection of a Supreme Court opinion within hours of its release—offering analysis well ahead of, and with greater depth, than via what bloggers call themainstream media (MSM). As such, it's not surprising that most mainstream news outlets have begun supplementing their content with blogs that can offer greater depth, more detail, and deadline-free timeliness.

Blogs

While the feature set of a particular blog depends on the underlying platform and the preferences of the blogger, several key features are common to most blogs:

- *Ease of use.* Creating a new post usually involves clicking a single button.
- *Reverse chronology.* Posts are listed in reverse order of creation, making it easy to see the most recent content.
- *Comment threads.* Readers can offer comments on posts.
- *Persistence.* Posts are maintained indefinitely at locations accessible by permanent links.
- *Searchability.* Current and archived posts are easily searchable.
- *Tags.* Posts are often classified under an organized tagging scheme.
- *Trackbacks.* Allows an author to acknowledge the source of an item in their post, which allows bloggers to follow the popularity of their posts among other bloggers.

The voice of the blogosphere can wield significant influence. Examples include leading the charge for Dan Rather's resignation and prompting the design of a new insulin pump. In an example of what can happen when a



firm ignores social media, consider the flare-up Ingersoll Rand faced when the online community exposed a design flaw in its Kryptonite bike lock. Online posts showed the thick metal lock could be broken with a simple ball-point pen. A video showing the hack was posted online. When Ingersoll Rand failed to react quickly, the blogosphere erupted with criticism. Just days after online reports appeared, the mainstream media picked up the story. The *New York Times* ran a story titled “The Pen Is Mightier Than the Lock” that included a series of photos demonstrating the ballpoint Kryptonite lock pick. The event tarnished the once-strong brand and eventually resulted in a loss of over ten million dollars.

Concern over managing a firm’s online reputation by monitoring blog posts and other social media commentary has led to the rise of an industry known as online reputation management. Firms specializing in this field will track a client firm’s name, brand, executive names, or other keywords and report online activity and whether it is positive or negative

Like any Web page, blogs can be public, tucked behind a corporate firewall, or password protected. Most blogs offer a two-way dialog, allowing users to comment on posts (sort of instant “letters to the editor,” posted online and delivered directly to the author). The running dialog can read like an electronic bulletin board, and can be an effective way to gather opinion when vetting ideas. Just as important, user comments help keep a blogger honest. Just as the “wisdom of crowds” keeps Wikipedia accurate, a vigorous community of commenters will quickly expose a blogger’s errors of fact or logic.

Despite this increased popularity, blogging has its downside. Blog comments can be a hothouse for spam and the disgruntled. Ham-handed corporate efforts (such as poor response to public criticism or bogus “praise posts”) have been ridiculed. Employee blogging can be difficult to control and public postings can “live” forever in the bowels of an Internet search



engine or as content pasted on other Web sites. Many firms have employee blogging and broader Internet posting policies to guide online conduct that may be linked to the firm. Bloggers beware—there are dozens of examples of workers who have been fired for what employers viewed as inappropriate posts.

Blogs can be hosted via third-party services (Google Blogger, WordPress.com, TypePad, Windows Live Spaces), with most offering a combination of free and premium features. Blogging features have also been incorporated into social networks such as Facebook, MySpace, and Ning, as well as wiki tools such as SocialText. Blogging software can also be run on third-party servers, allowing the developer more control in areas such as security and formatting. The most popular platform for users choosing to host their own blog server is the open source Word Press system (based on PHP and MySQL).

Blogs have become a fire hose of rapidly delivered information as tool providers have made it easier for would-be bloggers to capture and post content. For example, the Google Toolbar has a “BlogThis!” feature that allows anyone with a Google Blogger account to post links directly to their blogs. Apple ships blog hosting tools with its server products, and Microsoft offers the free LiveWriter blog editing tool and the Windows Spaces blogging service. Blogger and many social networking sites also offer easy features for posting content and photos from a mobile phone or devices like the iPod Touch.

In the end, the value of any particular blog derives from a combination of technical and social features. The technical features make it easy for a blogger and his or her community to engage in an ongoing conversation on some topic of shared interest. But the social norms and patterns of use that emerge over time in each blog are what determine whether technology features will be harnessed for good or ill. Some blogs develop norms of



fairness, accuracy, proper attribution, quality writing, and good faith argumentation, and attract readers that find these norms attractive. Others mix it up with hotly contested debate, one-sided partisanship, or deliberately provocative posts, attracting a decidedly different type of discourse.

KEY TAKEAWAYS

- Blogs provide a rapid way to distribute ideas and information from one writer to many readers.
- Ranking engines, trackbacks, and comments allow a blogger's community of readers to spread the word on interesting posts and participate in the conversation, and help distinguish and reinforce the reputations of widely read blogs.
- Well-known blogs can be powerfully influential, acting as flashpoints on public opinion.
- Firms ignore influential bloggers at their peril, but organizations should also be cautious about how they use and engage blogs, and avoid flagrantly promotional or biased efforts.
- Top blogs have gained popularity, valuations, and profits that far exceed those of many leading traditional newspapers, and leading blogs have begun to attract well-known journalists away from print media.
- Senior executives from several industries use blogs for business purposes including marketing, sharing ideas, gathering feedback, press response, and image shaping.

QUESTIONS AND EXERCISES

1. Visit Technorati and find out which blogs are currently the most popular. Why do you suppose the leaders are so popular?
2. How are popular blogs discovered? How is their popularity reinforced?

3. Are blog comment fields useful? If so, to whom or how? What is the risk associated with allowing users to comment on blog posts? How should a blogger deal with comments that they don't agree with?
4. Why would a corporation, an executive, a news outlet, or a college student want to blog? What are the benefits? What are the concerns?
5. Identify firms and executives that are blogging online. Bring examples to class and be prepared to offer your critique of their efforts.
6. How do bloggers make money? Do all bloggers have to make money?
7. According to your reading, how does the blog The Huffington Post compare with the popularity of newspaper Web sites?
8. What advantage do blogs have over the MSM? What advantage does the MSM have over the most popular blogs?
9. Start a blog using Blogger.com, WordPress.com, or some other blogging service. Post a comment to another blog. Look for the trackback field when making a post, and be sure to enter the trackback for any content you cite in your blog.

6.3 Wikis

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Know what wikis are and how they are used by corporations and the public at large.
2. Understand the technical and social features that drive effective and useful wikis.



3. Be able to suggest opportunities where wikis would be useful and consider under what circumstances their use may present risks.
4. Recognize how social media such as wikis and blogs can influence a firm's customers and brand.

A wiki is a Web site anyone can edit directly within a Web browser (provided the site grants the user edit access). Wikis derive their name from the Hawaiian word for “quick.” Ward Cunningham, the “wiki father” christened this new class of software with the moniker in honor of the wiki-wiki shuttle bus at the Honolulu airport. Wikis can indeed be one of the speediest ways to collaboratively create content online. Many popular online wikis serve as a shared knowledge repository in some domain. The largest and most popular wiki is Wikipedia, but there are hundreds of publicly accessible wikis that anyone can participate in. Each attempts to chronicle a world of knowledge within a particular domain, with examples ranging from Wine Wiki for oenophiles to Wookieepedia, the Star Wars wiki. But wikis can be used for any collaborative effort—from meeting planning to project management. And in addition to the hundreds of public wikis, there are many thousand more that are hidden away behind firewalls, used as proprietary internal tools for organizational collaboration. Like blogs, the value of a wiki derives from both technical and social features. The *technology* makes it easy to create, edit, and refine content; learn when content has been changed, how and by whom; and to change content back to a prior state. But it is the *social motivations of individuals* (to make a contribution, to share knowledge) that allow these features to be harnessed. The larger and more active a wiki community, the more likely it is that content will be up-to-date and that errors will be quickly corrected. Several studies have shown that large community wiki entries are as or more accurate than professional publication counterparts.



Want to add to or edit a wiki entry? On most sites you just click the Edit link. Wikis support what you see is what you get (WYSIWYG) editing that, while not as robust as traditional word processors, is still easy enough for most users to grasp without training or knowledge of arcane code or markup language. Users can make changes to existing content and can easily create new pages or articles and link them to other pages in the wiki. Wikis also provide a version history. Click the “history” link on Wikipedia, for example, and you can see when edits were made and by whom. This feature allows the community to roll back a wiki to a prior page, in the event that someone accidentally deletes key info, or intentionally defaces a page.

Vandalism is a problem on Wikipedia, but it’s more of a nuisance than a crisis. A *Wired* article chronicled how Wikipedia’s entry for former U.S. President Jimmy Carter was regularly replaced by a photo of a “scruffy, random unshaven man with his left index finger shoved firmly up his nose.” ^[1] Nasty and inappropriate, to be sure, but the Wikipedia editorial community is now so large and so vigilant that most vandalism is caught and corrected within seconds. Watch-lists for the most active targets (say the Web pages of political figures or controversial topics) tip off the community when changes are made. The accounts of vandals can be suspended, and while mischief-makers can log in under another name, most vandals simply become discouraged and move on. It’s as if an army of do-gooders follows a graffiti tagger and immediately repaints any defacement.

Wikis

As with blogs, a wiki’s features set varies depending on the specific wiki tool chosen, as well as administrator design, but most wikis support the following key features:

- All changes are *attributed*, so others can see who made a given edit.
- A complete *revision history* is maintained so changes can be compared against prior versions and rolled back as needed.
- There is automatic *notification* and *monitoring* of updates; users subscribe to wiki content and can receive updates via e-mail or RSS feed when pages have been changed or new content has been added.
- All the pages in a wiki are *searchable*.
- Specific wiki pages can be classified under an organized *tagging* scheme.

Wikis are available both as software (commercial as well as open source varieties) that firms can install on their own computers or as online services (both subscription or ad-supported) where content is hosted off site by third parties. Since wikis can be started without the oversight or involvement of a firm's IT department, their appearance in organizations often comes from grassroots user initiative. Many wiki services offer additional tools such as blogs, message boards, or spreadsheets as part of their feature set, making most wikis really more full-featured platforms for social computing.

Examples of Wiki Use

Wikis can be vital tools for collecting and leveraging knowledge that would otherwise be scattered throughout an organization; reducing geographic distance; removing boundaries between functional areas; and flattening preexisting hierarchies. Companies have used wikis in a number of ways:

- At Pixar, all product meetings have an associated wiki to improve productivity. The online agenda ensures that all attendees can arrive knowing the topics and issues to be covered. Anyone attending the meeting (and even those who can't make it) can update the agenda, post supporting materials, and make comments to streamline and focus in-person efforts.



- At European investment bank Dresdner Kleinwort Wasserstein, employees use wikis for everything from setting meeting agendas to building multimedia training for new hires. Six months after launch, wiki use had surpassed activity on the firm's established intranet. Wikis are also credited with helping to reduce Dresdner e-mail traffic by 75 percent. [2]
- Sony's PlayStation team uses wikis to regularly maintain one-page overviews on the status of various projects. In this way, legal, marketing, and finance staff can get quick, up-to-date status reports on relevant projects, including the latest projected deadlines, action items, and benchmark progress. Strong security measures are enforced that limit access to only those who must be in the know, since the overviews often discuss products that have not been released.
- Employees at investment-advisory firm Manning and Napier use a wiki to collaboratively track news in areas of critical interest. Providing central repositories for employees to share articles and update evolving summaries on topics such as universal health care legislation, enables the firm to collect and focus what would otherwise be fragmented findings and insight. Now all employees can refer to central pages that each serve as a lightning rod attracting the latest and most relevant findings.

When brought outside the firewall, corporate wikis can also be a sort of value-generation greenhouse, allowing organizations to leverage input from their customers and partners:

- Amazon leverages its Amapedia wiki, allowing customers to create user-generated product information and side-by-side review comparisons.
- eBay, a site that has regularly been shaped by user input through bulletin-board forums, now provides a wiki where users can collectively create articles on any topic of interest to the community, including sales strategy, information for collectable enthusiasts, or new user documentation.



- Microsoft leveraged its customer base to supplement documentation for its Visual Studio software development tool. The firm was able to enter the Brazilian market with Visual Studio in part because users had created product documentation in Portuguese. [3]
- ABC and CBS have created public wikis for the television programs Lost, The Amazing Race, and CSI, among others, offering an outlet for fans, and a way for new viewers to catch up on character backgrounds and complex plot lines.
- Executive Travel, owned by American Express Publishing, has created a travel wiki for its more than one hundred and thirty thousand readers with the goal of creating what it refers to as “a digital mosaic that in theory is more authoritative, comprehensive, and useful” than comments on a Web site, and far more up-to-date than any paper-based travel guide. [4] Of course, one challenge in running such a corporate effort is that there may be a competing public effort already in place. Wikitravel.org currently holds the top spot among travel-based wikis, and network effects suggest it will likely grow and remain more current than rival efforts.

Jump-starting a wiki can be a challenge, and an underused wiki can be a ghost town of orphan, out-of-date, and inaccurate content. Fortunately, once users see the value of wikis, use and effectiveness often snowballs. The unstructured nature of wikis are also both a strength and weakness. Some organizations employ wikimasters to “garden” community content; “prune” excessive posts, “transplant” commentary to the best location, and “weed” as necessary. Often communities will quickly develop norms for organizing and maintaining content. Wikipatterns.com offers a guide to the stages of wiki adoption and a collection of community-building and content-building strategies.



Don't Underestimate the Power of Wikipedia

Not only is the nonprofit Wikipedia, with its enthusiastic army of unpaid experts and editors, an immediate threat to the three-hundred-year-old *Encyclopedia Britannica*, Wikipedia entries can impact nearly all large-sized organizations. Wikipedia is the go-to, first-choice reference site for a generation of “netizens,” and Wikipedia entries are invariably one of the top links, often the first link, to appear in Internet search results.

This position means that anyone from top executives to political candidates to any firm large enough to warrant an entry has to contend with the very public commentary offered up in a Wikipedia entry. In the same way that firms monitor their online reputations in blog posts and Twitter tweets, they've also got to keep an eye on wikis.

But firms that overreach and try to influence an entry outside of Wikipedia's mandated neutral point of view (NPOV), risk a backlash and public exposure. Version tracking means the wiki sees all. Users on computers at right-leaning Fox News were embarrassingly caught editing the wiki page of the lefty pundit and politician Al Franken (a nemesis of Fox's Bill O'Reilly); ^[5] Sony staffers were flagged as editing the entry for the Xbox game Halo 3; ^[6] and none other than Wikipedia founder Jimmy Wales was criticized for editing his own Wikipedia biography; ^[7] an act that some consider bad online form at best, and dishonest at worst.

One last point on using Wikipedia for research. Remember that according to its own stated policies, Wikipedia isn't an original information source; rather, it's a clearinghouse for verified information. So citing Wikipedia as a reference usually isn't considered good form. Instead, seek out original (and verifiable) sources via the links at the bottom of Wikipedia entries.

KEY TAKEAWAYS

- Wikis can be powerful tools for many-to-many content collaboration, and can be ideal for creating resources that benefit from the input of many such as encyclopedia entries, meeting agendas, and project status documents.
- The greater the number of wiki users, the more likely the information contained in the wiki will be accurate and grow in value.
- Wikis can be public or private.
- The availability of free or low-cost wiki tools can create a knowledge clearinghouse on topics, firms, products, and even individuals. Organizations can seek to harness the collective intelligence (wisdom of the crowds) of online communities. The openness of wikis also act as a mechanism for promoting organizational transparency and accountability.

QUESTIONS AND EXERCISES

1. Visit a wiki, either an established site like Wikipedia, or a wiki service like SocialText. Make an edit to a wiki entry or use a wiki service to create a new wiki for your own use (e.g., for a class team to use in managing a group project). Be prepared to share your experience with the class.
2. What factors determine the value of a wiki? Which key concept, first introduced in [Chapter 2 "Strategy and Technology"](#), drives a wiki's success?
3. If anyone can edit a wiki, why aren't more sites concerned about vandalism, inaccurate, or inappropriate content? Are there technical reasons not to be concerned? Are there "social" reasons that can alleviate concern?
4. Give examples of corporate wiki use, as well as examples where firms used wikis to engage their customers or partners. What is the potential payoff of these efforts? Are there risks associated with these efforts?



5. Do you feel that you can trust content in wikis? Do you feel this content is more or less reliable than content in print encyclopedias? Than the content in newspaper articles? Why?
6. Have you ever run across an error in a wiki entry? Describe the situation.
7. Is it ethical for a firm or individual to edit their own Wikipedia entry? Under what circumstances would editing a Wikipedia entry seem unethical to you? Why? What are the risks a firm or individual is exposed to when making edits to public wiki entries? How do you suppose individuals and organizations are identified when making wiki edits?
8. Would you cite Wikipedia as a reference when writing a paper? Why or why not?



6.4 Electronic Social Networks

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Know what social networks are, be able to list key features, and understand how they are used by individuals, groups, and corporations.
2. Understand the difference between major social networks MySpace, Facebook, and LinkedIn.
3. Recognize the benefits and risks of using social networks.
4. Be aware of trends that may influence the evolution of social networks.

Social networks have garnered increasing attention as established networks grow and innovate, new networks emerge, and value is demonstrated.

MySpace signed a billion-dollar deal to carry ads from Google's AdSense network. Meanwhile, privately held Facebook has moved beyond its college roots, and opened its network to third-party application developers.

LinkedIn, which rounds out the Big Three U.S. public social networks, has grown to the point where its influence is threatening recruiting sites like Monster.com and CareerBuilder. ^[1] It now offers services for messaging, information sharing, and even integration with the *BusinessWeek* Web site.

Media reports often mention MySpace, Facebook, and LinkedIn in the same sentence. However, while these networks share some common features, they serve very different purposes. MySpace pages are largely for public consumption. The site was originally started by musicians as a tool to help users discover new music and engage with bands. Today, MySpace



members leverage the service to discover people with similar tastes or experiences (say fans of a comic, or sufferers of a condition).

Facebook, by contrast, is more oriented towards reinforcing existing social ties between people who already know each other. This difference leads to varying usage patterns. Since Facebook is perceived by users as relatively secure, with only invited “friends” seeing your profile, over a third of Facebook users post their mobile phone numbers on their profile pages. LinkedIn was conceived from the start as a social network for business users. The site’s profiles act as a sort of digital Rolodex that users update as they move or change jobs. Users can pose questions to members of their network, engage in group discussions, ask for introductions through mutual contacts, and comment on others’ profiles (e.g., recommending a member). Active members find the site invaluable for maintaining professional contacts, seeking peer advice, networking, and even recruiting. Carmen Hudson, Starbucks manager of enterprise staffing, states LinkedIn is “one of the best things for finding midlevel executives.” [2] Such networks are also putting increasing pressure on firms to work particularly hard to retain top talent. While once HR managers fiercely guarded employee directories for fear that a list of talent may fall into the hands of rivals, today’s social networks make it easy for anyone to gain a list of a firm’s staff, complete with contact information.

While these networks dominate in the United States, the network effect and cultural differences work to create islands where other social networks are favored by a particular culture or region. The first site to gain traction in a given market is usually the winner. Google’s Orkut, Bebo (now owned by AOL), and Cyworld have small U.S. followings, but are among the largest sites in Brazil, Europe, and South Korea. Research by Ipsos Insight also suggests that users in many global markets, including Brazil, South Korea,



and China, are more active social networkers than their U.S. counterparts. [3]

Perhaps the most powerful (and controversial) feature of most social networks is the feed(or newsfeed). Pioneered by Facebook but now adopted by most services, feeds provide a timely update on the activities of people or topics that an individual has an association with. Facebook feeds can give you a heads up when someone makes a friend, joins a group, posts a photo, or installs an application.

Feeds are inherently viral. By seeing what others are doing on a social network, feeds can rapidly mobilize populations and dramatically spread the adoption of applications. Leveraging feeds, it took just ten days for the Facebook group Support the Monks' Protest in Burma to amass over one hundred and sixty thousand Facebook members. Feeds also helped music app iLike garner three million Facebook users just two weeks after its launch. [4] Its previous Web-based effort took eight months to reach those numbers.

But feeds are also controversial. Many users react negatively to this sort of public broadcast of their online activity, and feed mismanagement can create public relations snafus, user discontent, and potentially open up a site to legal action. Facebook initially dealt with a massive user outcry at the launch of feeds, and faced a subsequent backlash when its Beacon service broadcast user purchases without first explicitly asking their permission (see [Chapter 7 "Facebook: Building a Business from the Social Graph"](#) for more details).

Social Networks

The foundation of a social network is the user profile, but utility goes beyond the sort of listing found in a corporate information directory. Typical features of a social network include support for the following:

- Detailed *personal profiles*

- Affiliations with *groups*, such as alumni, employers, hobbies, fans, health conditions)
- Affiliations with *individuals* (e.g., specific “friends”)
- Private *messaging* and public *discussions*
- Media *sharing* (text, photos, video)
- “*Feeds*” of recent activity among members (e.g., status changes, new postings, photos, applications installed)
- The ability to install and use third-party applications tailored to the service (games, media viewers, survey tools, etc.), many of which are also social and allow others to interact.

Corporate Use of Social Networks

The use of public social networks within private organizations is becoming widespread. Many employees have organized groups using publicly available social networking sites because similar tools are not offered by their firms. *Workforce Management* reported that MySpace had over forty thousand groups devoted to companies or coworkers, while Facebook had over eight thousand. [5] Assuming a large fraction of these groups are focused on internal projects, this demonstrates a clear pent up demand for corporate-centric social networks.

Many firms are choosing to meet this demand by implementing internal social network platforms that are secure and tailored to firm needs. At the most basic level, these networks have supplanted the traditional employee directory. Social network listings are easy to update and expand. Employees are encouraged to add their own photos, interests, and expertise to create a living digital identity.

Firms such as Deloitte, Dow Chemical, and Goldman Sachs have created social networks for “alumni” who have left the firm or retired. These networks can be useful in maintaining contacts for future business leads,



rehiring former employees (20 percent of Deloitte’s experienced hires are so called “boomerangs,” or returning employees), or recruiting retired staff to serve as contractors when labor is tight. [6] Maintaining such networks will be critical in industries like I.T. and health care that are likely to be plagued by worker shortages for years to come.

Social networking can also be important for organizations like IBM, where some 42 percent of employees regularly work from home or client locations. IBM’s social network makes it easier to locate employee expertise within the firm, organize virtual work groups, and communicate across large distances. [7] As a dialogue catalyst, a social network transforms the public directory into a front of knowledge sharing that promotes organization flattening and value-adding expertise sharing.

As another example of corporate social networks, Reuters has rolled out Reuters Space, a private online community for financial professionals. Profile pages can also contain a personal blog and news feeds (from Reuters or external services). Every profile page is accessible to the entire Reuters Space community, but members can choose which personal details are available to whom. While IBM and Reuters have developed their own social network platforms, firms are increasingly turning to third-party vendors like SelectMinds (adopted by Deloitte, Dow Chemical, and Goldman Sachs) and LiveWorld (adopted by Intuit, eBay, the NBA, and Scientific American).

Firms have also created their own online communities to foster brainstorming and customer engagement. Dell’s IdeaStorm.com forum collects user feedback and is credited with prompting line offerings, such as the firm’s introduction of a Linux-based laptop. [8] At MyStarbucksIdea.com, the coffee giant has leveraged user input to launch a series of innovations ranging from splash sticks that prevent spills in to-go cups, to new menu items. Both IdeaStorm and MyStarbucksIdea run on a



platform offered by Salesforce.com's that not only hosts these sites, but provides integration into Facebook and other services. This platform has allowed Starbucks to leverage Facebook to promote free coffee on election day, a free cup on inauguration day for those pledging five hours of community service, and awareness of the firm's AIDS-related (RED) campaign. This latter effort garnered an astonishing three hundred ninety million "viral impressions," including activities such as joining a group to support the cause, messaging others, or making wall posts. [9]

Concerns

As with any type of social media, content flows in social networks are difficult to control. Embarrassing disclosures can emerge from public systems or insecure internal networks. Employees embracing a culture of digital sharing may err and release confidential or proprietary information. Networks could serve as a focal point for the disgruntled (imagine the activity on a corporate social network after a painful layoff). Publicly declared affiliations, political or religious views, excessive contact, declined participation, and other factors might lead to awkward or strained employee relationships. Users may not want to add a coworker as a friend on a public network if it means they'll expose their activities, lives, persona, photos, sense of humor, and friends as they exist outside of work. And many firms fear wasted time as employees surf the musings and photos of their peers.

All are advised to be cautious in their social media sharing. Employers are trawling the Internet, mining Facebook, and scouring YouTube for any tip-off that a would-be hire should be passed over. A word to the wise: those Facebook party pics, YouTube videos of open mic performances, or blog postings from a particularly militant period might not age well, and may haunt you forever in a Google search. Think twice before clicking the



upload button! As Socialnomics author Erik Qualman puts it, “What happens in Vegas stays on YouTube (and Flickr, Twitter, Facebook...).” Despite these concerns, trying to micromanage a firm’s social network is probably not the answer. At IBM, rules for online behavior are surprisingly open. The firm’s code of conduct reminds employees to remember privacy, respect, and confidentiality in all electronic communications. Anonymity is not permitted on IBM’s systems, making everyone accountable for their actions. As for external postings, the firm insists that employees don’t disparage competitors or reveal customers’ names without permission, and asks that employee posts from IBM accounts or that mention the firm also include disclosures indicating that opinions and thoughts shared publicly are the individual’s and not Big Blue’s.

Trends to Watch

Global positioning system (GPS) and location-based services in devices are appearing in social software that allows colleagues and friends to find one another when nearby, ushering in a host of productivity and privacy issues along with them. At the launch of Apple’s iPhone 3G, Loopt CEO Sam Altman showed how the firm’s touch app for iPhones and iPods could display a map with dots representing anyone on his or her contact list who elected to broadcast their whereabouts to him. Claims Altman, “You never have to eat lunch alone again.”

Other efforts seek to increase information exchange and potentially lower switching costs between platforms. Facebook Connect allows users to access affiliated services using their Facebook I.D. and to share information between Facebook and these sites. OpenSocial, sponsored by Google and embraced by MySpace, LinkedIn, and many others, is a platform allowing third-party programmers to build widgets that take advantage of personal data and profile connections across social networking sites. Similarly,



DataPortability.org, is an effort, supported by all major public social network platforms, to allow users to migrate data from one site for reuse elsewhere. The initiative should also allow vendors to foster additional innovation and utility through safe, cross-site data exchange.

KEY TAKEAWAYS

- Electronic social networks help individuals maintain contacts, discover and engage people with common interests, share updates, and organize as groups.
- Modern social networks are major messaging services, supporting private one-to-one notes, public postings, and broadcast updates or “feeds.”
- Social networks also raise some of the strongest privacy concerns, as status updates, past messages, photos, and other content linger, even as a user’s online behavior and network of contacts changes.
- Network effects and cultural differences result in one social network being favored over others in a particular culture or region.
- Information spreads virally via news feeds. Feeds can rapidly mobilize populations, and dramatically spread the adoption of applications. The flow of content in social networks is also difficult to control and sometimes results in embarrassing public disclosures.
- Feeds have a downside and there have been instances where feed mismanagement has caused user discontent, public relations problems, and the possibility of legal action.
- The use of public social networks within private organizations is growing, and many organizations are implementing their own, private, social networks.
- Firms are also setting up social networks for customer engagement and mining these sites for customer ideas, innovation, and feedback.



- Emerging trends related to social networks that bear watching include: the use of e-mail to add contacts; the use of GPS; and the development of platforms to facilitate single sign-on, and sharing of information between sites.

QUESTIONS AND EXERCISES

1. Visit the major social networks (MySpace, Facebook, LinkedIn). What distinguishes one from the other? Are you a member of any of these services? Why or why not?
2. How are organizations like Deloitte, Goldman Sachs, and IBM using social networks? What advantages do they gain from these systems?
3. What factors might cause an individual, employee, or firm to be cautious in their use of social networks?
4. How do you feel about the feed feature common in social networks like Facebook? What risks does a firm expose itself to if it leverages feeds? How might a firm mitigate these kinds of risks?
5. What sorts of restrictions or guidelines should firms place on the use of social networks or the other Web 2.0 tools discussed in this chapter? Are these tools a threat to security? Can they tarnish a firm's reputation? Can they enhance a firm's reputation? How so?
6. Why do information and applications spread so quickly within networks like Facebook? What feature enables this? What key promotional concept (described in [Chapter 2 "Strategy and Technology"](#)) does this feature foster?
7. Why are some social networks more popular in some nations than others?
8. Investigate social networks on your own. Look for examples of their use for fostering political and social movements; for their use in healthcare, among doctors, patients, and physicians; and for their use among other professional groups or enthusiasts. Identify how these networks might be



used effectively, and also look for any potential risks or downside. How are these efforts supported? Is there a clear revenue model, and do you find these methods appropriate or potentially controversial? Be prepared to share your findings with your class.

6.5 Twitter and the Rise of Microblogging

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Appreciate the rapid rise of Twitter—its scale, scope, and broad appeal.
2. Understand how Twitter is being used by individuals, organizations, and political movements.
3. Contrast Twitter and microblogging with Facebook, conventional blogs, and other Web 2.0 efforts.
4. Consider the commercial viability of the effort, its competitive environment, and concerns regarding limited revenue.

Spawned in 2006 as a side project at the now-failed podcasting startup Odeo (an effort backed by Blogger.com founder Evan Williams), Twitter had a breakout year in 2009. The actor Ashton Kutcher beat out CNN as owner of the first account to net one million followers. Oprah's first tweet is credited with boosting traffic by 43 percent. And by April 2009, Compete.com tracked the site as having thirty-two million visitors, more than the monthly totals for Digg (twenty-three million), the *New York*



Times (seventeen and a half million), or LinkedIn (sixteen million). Reports surfaced of rebuffed buyout offers as high as five hundred million dollars. ^[1] Pretty impressive numbers for a firm which that month was reported to have only forty-nine employees. ^[2]

Twitter is a microblogging service that allows users to post 140 character tweets via the Web, SMS, or a variety of third-party applications. The microblog moniker is a bit of a misnomer. The service actually has more in common with Facebook's status update field than traditional blogs. But unlike Facebook, where most users must approve "friends" before they can see status updates, Twitter's default setting allows for asymmetrical following (although it is possible to set up private Twitter accounts and to block followers).

Sure there's a lot of inane "tweeting" going on. Lots of meaningless updates that read, "at the coffee shop" or "in line at the airport." But many find Twitter to be an effective tool for quickly blasting queries to friends, colleagues, or strangers with potentially valuable input. Says futurist Paul Saffo, "Instead of creating the group you want, you send it and the group self-assembles." ^[3]

Consider the following commercial examples:

- Starbucks uses Twitter in a variety of ways. It has tweeted corrections of bogus claims, such as the rumor the firm was dropping coffee shipments to U.S. troops to protest U.S. war efforts. The firm has also run Twitter-based contests, and promoted free samples of new products, such as its VIA instant coffee line. Starbucks has even recruited and hired off of Twitter.
- Dell got an early warning sign on poor design on its Mini 9 Netbook PC. After a series of early tweets indicated that the apostrophe and return keys were positioned too closely together, the firm dispatched design change orders quickly enough to correct the problem when the Mini 10 was launched just three months later. By June '09, Dell also claimed to have

netted two million dollars in outlet store sales referred via the Twitter account @DellOutlet (more than six hundred thousand followers), and another one million dollars from customers who have bounced from the outlet to the new products site. [4]

- The True Massage and Wellness Spa in San Francisco tweets when last-minute cancellations can tell customers of an opening. With Twitter, appointments remain booked solid.

Surgeons and residents at Henry Ford Hospital have even tweeted during brain surgery (the teaching hospital sees the service as an educational tool).

Some tweets are from those so young they've got "negative age."

Twitter.com/kickbee is an experimental fetal monitor band that sends tweets when motion is detected: "I kicked Mommy at 08:52." And savvy hackers are embedding "tweeting" sensors into all sorts of devices.

Botanicalls (<http://www.botanicalls.com>), for example, offers a sort of electronic flowerpot stick that detects when plants need water and sends Twitter status updates to a mobile phone.

Twitter has provided early warning on earthquakes, terror attacks, and other major news events. And the site was seen as such a powerful force that the Iranian government blocked access to Twitter (and other social media sites) in an attempt to quell the June 2009 election protests. Users quickly assemble comments on a given topic using hash tags (keywords preceded by the # or "hash" symbol), allowing others to quickly find related tweets (e.g., #iranelection, #mumbai, #swineflu, #sxsw).

Organizations are well advised to monitor Twitter activity related to the firm, as it can act as a sort of canary-in-a-coal mine uncovering emerging events. Users are increasingly using the service as a way to form flash protest crowds. Amazon.com, for example, was caught off guard over a spring 2009 holiday weekend when thousands used Twitter to rapidly



protest the firm’s reclassification of gay and lesbian books (hash tag #amazonfail).

For all the excitement, many wonder if Twitter is overhyped. All the adoption success mentioned above came to a firm with exactly zero dollars in revenue—none! Many wonder how the site will make money, if revenues will ever justify initially high valuations, and if rivals could usurp Twitter’s efforts with similar features. Another issue—many loyal Twitter users rarely visit the site. Most active users post and read tweets using one of the dozens of free applications provided by third parties, such as Seesmic, TweetDeck, Tweetie, and Twirl. If users don’t visit Twitter.com, that lessens the impact of any ads running on the site (although Twitter was not running any ads on its site as of this writing). This creates what is known as the “free rider problem,” where users benefit from a service while offering no value in exchange. And some reports suggest that many Twitter users are curious experimenters who drop the service shortly after signing up. [5] Microblogging is here to stay and the impact of Twitter has been deep, broad, stunningly swift, and at times humbling in the power it provides. But whether or not Twitter will be a durable, profit-gushing powerhouse remains to be seen.

KEY TAKEAWAYS

- While many public and private microblogging services exist, Twitter remains by far the dominant service.
- Unlike status updates found on services like Facebook and LinkedIn, Twitter’s default supports asymmetric communication, where someone can follow updates without first getting their approval. This function makes Twitter a good choice for anyone cultivating a following—authors, celebrities, organizations, and brand promoters.

- Twitter hash tags (keywords preceded by the # character) are used to organize “tweets” on a given topic. Users can search on hash tags, and many third-party applications allow for Tweets to be organized and displayed by tag.
- Firms are leveraging Twitter in a variety of ways, including: promotion, customer response, gathering feedback, and time-sensitive communication.
- Like other forms of social media, Twitter can serve as a hothouse that attracts opinion and forces organizational transparency and accountability.
- Activists have leveraged the service worldwide, and it has also served as an early warning mechanism in disasters, terror, and other events.
- Despite its rapid growth and impact, significant questions remain regarding the firm’s durability, revenue prospects, and enduring appeal to initial users.

QUESTIONS AND EXERCISES

1. If you don’t already have one, set up a Twitter account and “follow” several others. Follow a diverse group—corporations, executives, pundits, or other organizations. Do you trust these account holders are who they say they are? Why? Which examples do you think use the service most effectively? Which provide the weaker examples of effective Twitter use? Why? Have you encountered Twitter “spam” or unwanted followers? What can you do to limit such experiences? Be prepared to discuss your experiences with class.
2. If you haven’t done so, install a popular Twitter application such as TweetDeck, Seesmic, or a Twitter client for your mobile device. Why did you select the product you chose? What advantages does your choice offer over simply using Twitter’s Web page? What challenges do these

- clients offer Twitter? Does the client you chose have a clear revenue model? Is it backed by a viable business?
3. Which Twitter hash tags are most active at this time? What do you think of the activity in these areas? Is there legitimate, productive activity happening?
 4. Why would someone choose to use Twitter over Facebook's status update, or other services? Which (if either) do you prefer and why?
 5. What do you think of Twitter's revenue prospects? Is the firm a viable independent service or simply a feature to be incorporated into other social media activity? Advocate where you think the service will be in two years, five, ten. Would you invest in Twitter? Would you suggest that other firms do so? Why?
 6. Assume the role of a manager for your firm. Advocate how the organization should leverage Twitter and other forms of social media. Provide examples of effective use, and cautionary tales, to back up your recommendation.
 7. Some instructors have mandated Twitter for classroom use. Do you think this is productive? Would your professor advocate tweeting during lectures? What are the pros and cons of such use? Work with your instructor to discuss a set of common guidelines for in-class and course use of social media.

6.6 Other Key Web 2.0 Terms and Concepts



LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Know key terms related to social media, peer production, and Web 2.0, including RSS, folksonomies, mash-ups, virtual worlds, and rich media.
2. Be able to provide examples of the effective business use of these terms and technologies.

RSS

RSS (an acronym that stands for both “really simple syndication” and “rich site summary”) enables busy users to scan the headlines of newly available content and click on an item’s title to view items of interest, thus sparing them from having to continually visit sites to find out what’s new. Users begin by subscribing to an RSS feed for a Web site, blog, podcast, or other data source. The title or headline of any new content will then show up in an RSS reader. Subscribe to the *New York Times* Technology news feed, for example, and you will regularly receive headlines of tech news from the *Times*. Viewing an article of interest is as easy as clicking the title you like. Subscribing is often as easy as clicking on the RSS icon appearing on the home page of a Web site of interest.

Many firms use RSS feeds as a way to manage information overload, opting to distribute content via feed rather than e-mail. Some even distribute corporate reports via RSS. RSS readers are offered by third-party Web sites such as Google and Yahoo! and they have been incorporated into all popular browsers and most e-mail programs. Most blogging platforms provide a mechanism for bloggers to automatically publish a feed when each new post becomes available. Google’s FeedBurner is the largest publisher of RSS blog feeds, and offers features to distribute content via e-mail as well.



Folksonomies

Folksonomies (sometimes referred to as social tagging) are keyword-based classification systems created by user communities as they generate and review content. (The label is a combination of “folks” and “sonomy,” meaning a people-powered taxonomy). Bookmarking site Del.icio.us and photo-sharing site Flickr (both owned by Yahoo!) make heavy use of folksonomies.

With this approach, classification schemes emerge from the people most likely to understand them—the users. By leveraging the collective power of the community to identify and classify content, objects on the Internet become easier to locate, and content carries a degree of recommendation and endorsement.

Flickr cofounder Stewart Butterfield describes the spirit of folksonomies, saying, “The job of tags isn’t to organize all the world’s information into tidy categories, it’s to add value to the giant piles of data that are already out there.” ^[1] The Guggenheim Museum in New York City and the San Francisco Museum of Modern Art, among other museums, are taking a folksonomic approach to their online collections, allowing user-generated categories to supplement the specialized lexicon of curators. Amazon.com has introduced a system that allows readers to classify books, and most blog posts and wiki pages allow for social tagging, oftentimes with hot topics indexed and accessible via a “tag cloud” in the page’s sidebar.

Mash-up

Mash-ups are combinations of two or more technologies or data feeds into a single, integrated tool. Some of the best known mash-ups leverage Google’s mapping tools. HousingMaps.com combines Craigslist.org listings with Google Maps for a map-based display for apartment hunters. IBM linked together job feeds and Google Maps to create a job-seeker service for victims of Hurricane Katrina. SimplyHired links job listings with Google



Maps, LinkedIn listings, and salary data from PayScale.com. And Salesforce.com has tools that allow data from its customer relationship management (CRM) system to be combined with data feeds and maps from third parties.

Mash-ups are made easy by a tagging system called XML (for extensible markup language). Site owners publish the parameters of XML data feeds that a service can accept or offer (e.g., an address, price, product descriptions, images). Other developers are free to leverage these public feeds using application programming interfaces (APIs), published instructions on how to make programs call one another, to share data, or to perform tasks. Using APIs and XML, mash-up authors smooch together seemingly unrelated data sources and services in new and novel ways. Lightweight, browser-friendly software technologies like Ajax can often make a Web site interface as rich as a desktop application, and rapid deployment frameworks like Ruby on Rails will enable and accelerate mash-up creation and deployment.

Virtual Worlds

In virtual worlds, users appear in a computer-generated environment in the form of an avatar, or animated character. Users can customize the look of their avatar, interact with others by typing or voice chat, and can travel about the virtual world by flying, teleporting, or more conventional means. The most popular general-purpose virtual world is Second Life by Linden Labs, although many others exist. Most are free, although game-oriented worlds, such as World of Warcraft (with ten million active subscribers) charge a fee. Many corporations and organizations have established virtual outposts by purchasing “land” in the world of Second Life, while still others have contracted with networks to create their own, independent virtual worlds.



Even grade schoolers are heavy virtual world users. Many elementary school students get their first taste of the Web through Webkinz, an online world that allows for an animated accompaniment with each of the firm's plush toys. Webkinz's parent, privately held Ganz doesn't release financial figures, but according to Compete.com, by year-end 2008 Webkinz.com had roughly the same number of unique visitors as FoxNews. The kiddie set virtual world market is considered so lucrative that Disney acquired ClubPenguin for three hundred fifty million dollars with agreements to pay another potential three hundred fifty million if the effort hits growth incentives. [2]

Most organizations have struggled to commercialize these Second Life forays, but activity has been wide-ranging in its experimentation. Reuters temporarily "stationed" a reporter in Second Life, presidential candidates have made appearances in the virtual world, organizations ranging from Sun Microsystems to Armani have set up virtual storefronts, and there's a significant amount of virtual mayhem. Second Life "terrorists" have "bombed" virtual outposts run by several organizations, including ABC News, American Apparel, and Reebok.

YouTube, Podcasting, and Rich Media

Blogs, wikis, and social networks not only enable sharing text and photos, they also allow for the creation and distribution of audio and video. Podcasts are digital audio files (some also incorporate video), provided as a series of programs. Podcasts range from a sort of media blog, archives of traditional radio and television programs, and regular offerings of original online content. While the term podcast derives from Apple's wildly successful iPod, podcasts can be recorded in audio formats such as MP3 that can be played on most portable media players. (In perhaps the ultimate concession to the market leader, even the iPod rival Microsoft Zune refers to serialized audio files as podcasts on its navigation menu).



There are many podcast directories, but Apple's iTunes is by far the largest. Anyone who wants to make a podcast available on iTunes can do so for free. A podcast publisher simply records an audio file, uploads the file to a blog or other hosting server, then sends the RSS feed to Apple (copyrighted material cannot be used without permission, with violators risking banishment from iTunes). Files are discovered in the search feature of the iTunes music store, and listings seamlessly connect the user with the server hosting the podcast. This path creates the illusion that Apple serves the file even though it resides on a publisher's servers.

While blogs have made stars of some unknowns, the most popular podcasts are from mainstream media outlets. A recent visit to the podcasting section of iTunes showed that eight of the top ten most popular podcasts were high-quality productions of mainstream media programs, including offerings from CBS, Comedy Central, NPR, and PBS. Podcasts are also revolutionizing education, with scores of universities "open sourcing" their classrooms and offering lectures for public consumption via Apple's iTunesU.

In contrast to iTunes, YouTube actually hosts video on its own servers, so all you need to do is shoot a video and upload it to the site. YouTube is a bastion of amateur video, with most clips shot and uploaded by nonprofessionals. It's also become a protest site (examples: "United Breaks Guitars" and "A Comcast Technician Sleeping on my Couch"). However, YouTube has also become a go-to distribution platform for professional content such as ad clips, customer support guides, music videos, TV shows, movies, and more. Much of this rich media content can be distributed or streamed within another Web site, blog, or social network profile.



KEY TAKEAWAYS

- RSS fosters the rapid sharing and scanning of information, including updates from Web 2.0 services such as blogs, wikis, and social networks. RSS feeds can be received via Web browsers, e-mail, cell phones, and special RSS readers.
- Folksonomies allow users to collaboratively tag and curate online media, making it easy for others to find useful content. Since folksonomies are created by users themselves, they are often more easily understood and embraced than classification schemes imposed by site owners.
- Mash-ups promote the useful combination of different Web services, such as maps and other information feeds. Mash-up authors leverage technologies such as APIs and XML to combine seemingly unrelated data sources and services in new and novel ways.
- Virtual Worlds allow users to interact with and within a computer-generated alternate reality.
- Internet media is increasingly becoming “richer,” leveraging audio, video, and animation. Organizations and users are creating and distributing rich media online, with interesting content spreading virally.

QUESTIONS AND EXERCISES

1. What is RSS and an RSS reader? Why would an individual use one? Why would a firm use RSS?
2. Use an RSS reader like Google Reader, or the features built into your e-mail program or browser, and subscribe to RSS feeds. Discuss your experience with the class. Which feeds did you subscribe to? What did you like or not like about using an RSS reader?
3. Visit Second Life or another virtual world. Create an avatar and look for examples of corporate/commercial involvement. Be prepared to discuss your experience—both positive and negative.



4. Investigate some of the many virtual worlds targeted at children, including Webkinz, Club Penguin, and Whyville. What are the revenue models for these efforts? How do these sites ensure they are safe for children? Assume the role of a parent—what are the pros and cons of these sites? Which (if any) would you allow your children to participate in? Why? Would you invest in any of these efforts or advise corporations to enter the kid-focused virtual world space? Why or why not?
5. If you don't already own it, download iTunes for free, go to the iTunes music store, and explore the free podcast section. Alternatively, find podcasts from another service. Which podcasts seem to be the most popular? Why? Do you use podcasts or other rich media? Why or why not?
6. Visit YouTube. Identify examples of corporations using the service. Identify examples of customer use. Be prepared to discuss your findings with the class. Do you think Google is making much money with YouTube? Why or why not?
7. Which firm do you think spends more on the technology infrastructure that powers its service—Apple with iTunes podcasts or Google with YouTube? Explain your answer.
8. Why would a firm make its data available via XML to use as a mash-up? What can it gain? Are there any risks involved in providing programming hooks that allow the creation of mash-ups?
9. Give examples of efforts that take advantage of folksonomies. Why are folksonomies considered to be useful?
10. Do you spend time in rich media Web sites? Which ones? How much time do you or your friends spend on these sites? How would you describe the quality of rich media content found online?
11. How might a firm use rich media online? What concerns does a firm or individual face with respect to rich media?



12. Why do you suppose that the most popular podcasts come from established media firms (e.g., Comedy Central, NPR) rather than amateurs, while the top bloggers emerged outside the professional journalist / writer community?

6.7 Prediction Markets and the Wisdom of Crowds

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand the concept of the wisdom of crowds as it applies to social networking.
2. List the criteria necessary for a crowd to be smart.

Many social software efforts leverage what has become to be known as the wisdom of crowds. In this concept, a group of individuals (the crowd often consists mostly of untrained amateurs), collectively has more insight than a single or small group of trained professionals. Made popular by author James Surowiecki (whose best-selling book carried the same name), the idea of crowd wisdom is at the heart of wikis, folksonomy tagging systems, and many other online efforts. An article in the journal *Nature* positively comparing Wikipedia to Encyclopedia Britannica lent credence to social software's use in harnessing and distilling crowd wisdom. ^[1]



The crowd isn't always right, but in many cases where topics are complex, problems are large, and outcomes are uncertain, a large, diverse group may bring collective insight to problem solving that one smart guy or a professional committee lacks. One technique for leveraging the wisdom of the crowds is a prediction market, where a diverse crowd is polled and opinions aggregated to form a forecast of an eventual outcome. The concept is not new. The stock market is arguably a prediction market, with a stock price representing collective assessment of the discounted value of a firm's future earnings. But Internet technologies are allowing companies to set up predictive markets for exploring all sorts of problems.

Consider Best Buy, where employees are encouraged to leverage the firm's TagTrade prediction market to make forecasts, and are offered small gifts as incentives for participation. The idea behind this incentive program is simple: the "blue shirts" (Best Buy employees) are closest to customers. They see traffic patterns and buying cycles, can witness customer reactions first hand, and often have a degree of field insight not available to senior managers at the company's Minneapolis headquarters. Harness this collective input and you've got a group brain where, as wisdom of crowds proponents often put it, "the we is greater than the me." When Best Buy asked its employees to predict gift card sales, the "crowd's" collective average answer was 99.5 percent accurate; experts paid to make the prediction were off by 5 percent. Another experiment predicting holiday sales was off by only 1/10 of 1 percent. The experts? Off by 7 percent! [2] In an article in the *McKinsey Quarterly*, Surowiecki outlined several criteria necessary for a crowd to be "smart." [3] The crowd must be the following things:

- *diverse*, so that participants are bringing different pieces of information to the table
- *decentralized*, so that no one at the top is dictating the crowd's answer



- *summarize* people's opinions into one collective verdict
- *independent*, so that each focuses on information rather than the opinions of others

Google, which runs several predictive markets, underscored these principles when it found that predictions were less accurate when users were geographically proximate, meaning folks in the same work group who sat near one another typically thought too much alike.^[4] Poorer predictive outcomes likely resulted because these relatively homogeneous clusters of users brought the same information to the table (yet another reason why organizations should hire and cultivate diverse teams).

Many firms run predictive markets to aid in key forecasts, and with the potential for real financial payoff. But University of Chicago law professor Todd Henderson warns predictive markets may also hold legal and ethical challenges. The Securities and Exchange Commission may look askance at an employee who gets a heads-up in a predictive market that says a certain drug is going to be approved or fail clinical trials. If she trades on this information is she an insider, subject to prosecution for exploiting proprietary data? Disclosure issues are unclear. Gambling laws are also murky, with Henderson uncertain as to whether certain predictive markets will be viewed as an unregulated form of betting. ^[5]

Publicly accessible prediction markets are diverse in their focus. The Iowa Electronic Market attempts to guess the outcome of political campaigns, with mixed results. Farecast (now part of Microsoft's Bing knowledge engine) claims a 75 percent accuracy rate for forecasting the future price of airline tickets. ^[6] The Hollywood Stock Exchange allows participants to buy and sell prediction shares of movies, actors, directors, and film-related options. The exchange, now owned by investment firm Cantor Fitzgerald, has picked Oscar winners with 90 percent accuracy. ^[7] And at HedgeStreet.com, participants can make microbets, wagering as little as ten



dollars on the outcome of economic events, including predictions on the prices of homes, gold, foreign currencies, oil, and even the economic impact of hurricanes and tropical storms. HedgeStreet is considered a market and is subject to oversight by the Commodity Futures Trading Commission. ^[8]

KEY TAKEAWAYS

- Many Web 2.0 efforts allow firms to tap the wisdom of the crowds, identifying collective intelligence.
- Prediction markets tap crowd opinion with results that are often more accurate than the most accurate expert forecasts and estimates.
- Prediction markets are most accurate when tapping the wisdom of a diverse and variously skilled and experienced group, and are least accurate when participants are highly similar.

QUESTIONS AND EXERCISES

1. What makes for a “wise” crowd? When might a crowd not be so wise?
2. Find a prediction market online and participate in the effort. Be prepared to share your experience with your class, including any statistics of predictive accuracy, participant incentives, business model of the effort, and your general assessment of the appeal and usefulness of the effort.
3. Brainstorm on the kinds of organizations that might deploy prediction markets. Why might you think the efforts you suggest and advocate would be successful?
4. In what ways are legal issues of concern to prediction market operators?

6.8 Crowdsourcing

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand the value of crowdsourcing.
2. Identify firms that have used crowdsourcing successfully.

The power of Web 2.0 also offers several examples of the democratization of production and innovation. Need a problem solved? Offer it up to the crowd and see if any of their wisdom offers a decent result. This phenomenon, known as crowdsourcing, has been defined by Jeff Howe, founder of the blog crowdsourcing.com and an associate editor at *Wired*, as “the act of taking a job traditionally performed by a designated agent (usually an employee) and outsourcing it to an undefined, generally large group of people in the form of an open call.” ^[1]

Can the crowd really do better than experts inside a firm? At least one company has literally struck gold using crowdsourcing. As told by Don Tapscott and Anthony Williams in their book *Wikinomics*, mining firm Goldcorp was struggling to gain a return from its 55,000-acre Canadian property holdings. Executives were convinced there was gold “in them thar hills,” but despite years of efforts, the firm struggled to strike any new pay dirt. CEO Rob McEwen, a former mutual fund manager without geology experience who unexpectedly ended up running Goldcorp after a takeover battle, then made what seemed a Hail Mary pass—he offered up all the firm’s data, on the company’s Web site. Along with the data, McEwen ponied up \$575,000 from the firm as prize money for the Goldcorp Challenge to anyone who came up with the best methods and estimates for reaping golden riches. Releasing data was seen as sacrilege in the intensely



secretive mining industry, but it brought in ideas the firm had never considered. Taking the challenge was a wildly diverse group of “graduate students, consultants, mathematicians, and military officers.” Eighty percent of the new targets identified by entrants yielded “substantial quantities of gold.” The financial payoff? In just a few years a one-hundred-million-dollar firm grew into a nine-billion-dollar titan. For Goldcorp, the crowd coughed up serious coin.

Netflix followed Goldcorp’s lead, offering anonymous data to any takers, along with a one-million-dollar prize to the first team that could improve the accuracy of movie recommendations by 10 percent. Top performers among the over fifty thousand entrants included research scientists from AT&T Labs, researchers from the University of Toronto, a team of Princeton undergrads, and the proverbial “guy in a garage” (and yes, was his team name). Frustrated for nearly three years, it took a coalition of four teams from Austria, Canada, Israel, and the United States to finally cross the 10 percent threshold. The winning team represented an astonishing brain trust that Netflix would never have been able to harness on its own. [2] Other crowdsourcers include Threadless.com, which produces limited run T-shirts with designs users submit and vote on. Marketocracy runs stock market games and has created a mutual fund based on picks from the 100 top-performing portfolios. Just under seven years into the effort, the firm’s m100 Index reports a 75 percent return versus 35 percent for the S&P 500. The St. Louis Cardinals baseball team is even crowdsourcing. The club’s One for the Birds contest calls for the fans to submit scouting reports on promising players, as the team hopes to broaden its recruiting radar beyond its classic recruiting pool of Division I colleges.

There are several public markets for leveraging crowdsourcing for innovation, or as an alternative to standard means of production. Waltham, Massachusetts–based InnoCentive allows “seekers” to offer cash prizes

ranging from ten to one hundred thousand dollars. Over one hundred twenty thousand “solvers” have registered to seek solutions for tasks put forward by seekers that include Dow Chemical, Eli Lilly, and Procter & Gamble. Among the findings offered by the InnoCentive crowd are a biomarker that measures progression of ALS. Amazon.com has even created an online marketplace for crowdsourcing called Mechanical Turk. Anyone with a task to be completed or problem to be solved can put it up for Amazon, setting their price for completion or solution. For its role, Amazon takes a small cut of the transaction. And alpha geeks looking to prove their code chops can turn to TopCoder, a firm that stages coding competitions that deliver real results for commercial clients such as ESPN. By 2009, TopCoder contests had attracted over 175,000 participants from 200 countries. [3]

Not all crowdsourcers are financially motivated. Some benefit by helping to create a better service. Facebook leveraged crowd wisdom to develop versions of its site localized in various languages. Facebook engineers designated each of the site’s English words or phrases as a separate translatable object. Members were then invited to translate the English into other languages, and rated the translations to determine which was best. Using this form of crowdsourcing, fifteen hundred volunteers cranked out Spanish Facebook in a month. It took two weeks for two thousand German speakers to draft Deutsch Facebook. How does the Facebook “poking” concept translate? The majority of users decided on “dar un toque” in Spanish, “anklopfen” in German, and “envoyer un poke” in French. [4] Vive le crowd!

KEY TAKEAWAYS

- Crowdsourcing tackles challenges through an open call for to a broader community of potential problem solvers. Examples include Gold Corp’s

discovering of optimal mining locations in land it already held, Facebook's leverage of its users to create translations of the site for various international markets, and Netflix solicitation of improvements to its movie recommendation software.

- Several firms run third-party crowdsourcing forums, among them InnoCentive for scientific R&D, TopCoder for programming tasks, and Amazon's Mechanical Turk for general work.

QUESTIONS AND EXERCISES

1. What is crowdsourcing? Give examples of organizations that are taking advantage of crowdsourcing and be prepared to describe these efforts.
2. What ethical issues should firms be aware of when consider crowdsourcing? Are there other concerns firms may have when leveraging this technique?
3. Assume the role of a manager or consultant. Recommend a firm and a task that would be appropriate for crowdsourcing. Justify your choice, citing factors such as cost, breadth of innovation, time, constrained resources, or other factors. How would you recommend the firm conduct this crowdsourcing effort?

Chapter 7

Facebook: Building a Business from the Social Graph

7.1 Introduction

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Be familiar with Facebook's origins and rapid rise.

2. Understand how Facebook's rapid rise has impacted the firm's ability to raise venture funding and its founder's ability to maintain a controlling interest in the firm.

Here's how much of a Web 2.0 guy Mark Zuckerberg is: during the weeks he spent working on Facebook as a Harvard sophomore, he didn't have time to study for a course he was taking, "Art in the Time of Augustus," so he built a Web site containing all of the artwork in class and pinged his classmates to contribute to a communal study guide. Within hours, the wisdom of crowds produced a sort of custom CliffsNotes for the course, and after reviewing the Web-based crib sheet, he aced the test. Turns out he didn't need to take that exam, anyway. Zuck (that's what the cool kids call him) ^[1] dropped out of Harvard later that year.

Zuckerberg is known as both a shy, geeky, introvert who eschews parties, as well as for his brash Silicon Valley bad-boy image. After Facebook's incorporation, Zuckerberg's job description was listed as "Founder, Master and Commander [and] Enemy of the State." ^[2] An early business card read "I'm CEO...Bitch." ^[3] And let's not forget that Facebook came out of drunken experiments in his dorm room, one of which was initially to have compared classmates to farm animals (Zuckerberg, threatened with expulsion, later apologized). For one meeting with Sequoia Capital, the venerable Menlo Park venture capital firm that backed Google and YouTube, Zuckerberg showed up in his pajamas. ^[4]

By the age of twenty-three, Mark Zuckerberg had graced the cover of *Newsweek*, been profiled on *60 Minutes*, and was discussed in the tech world with a reverence previously reserved only for Steve Jobs and the Google guys, Sergey Brin and Larry Page. But Mark Zuckerberg's star rose much faster than any of these predecessors. Just two weeks after Facebook launched, the firm had four thousand users. Ten months later it was up to



one million. The growth continued, and the business world took notice. In 2006, Viacom (parent of MTV) saw that its core demographic was spending a ton of time on Facebook and offered to buy the firm for three quarters of a billion dollars. Zuckerberg passed. [5] Yahoo! offered up a cool one billion dollars (twice). Zuck passed again, both times.

As growth skyrocketed, Facebook built on its stranglehold of the college market (85 percent of four-year college students are Facebook members), opening up first to high schoolers, then to everyone. Web hipsters started selling shirts emblazoned with “I Facebooked your Mom!” Even Microsoft wanted some of Facebook’s magic. In 2006, the firm locked up the right to broker all banner ad sales that run on the U.S. version of Facebook, guaranteeing Zuckerberg’s firm one hundred million dollars a year through 2011. In 2007, Microsoft came back, buying 1.6 percent of the firm for two hundred forty million dollars and securing the rights to sell banner ads on all Facebook sites worldwide (70 percent of Facebook users are outside the United States).

The investment was a shocker. Do the math and a 1.6 percent stake for two hundred forty million dollars values Facebook at fifteen billion dollars (more on that later). That meant that a firm that at the time had only five hundred employees, one hundred fifty million dollars in revenues, and was helmed by a twenty-three-year-old college dropout in his first “real job,” was more valuable than General Motors. Then in May 2008 Facebook hit another major milestone: it passed MySpace to become the world’s largest social network. Rupert Murdoch, whose News Corporation owns MySpace, engaged in a little trash talk, referring to Facebook as “the flavor of the month.” [6]

Watch your back Rupert. Or on second thought, watch Zuckerberg’s. By spring 2009, Facebook had more than twice MySpace’s monthly unique visitors worldwide; [7] by June, Facebook surpassed MySpace in the United



States; [8] and by July, Facebook announced that it surpassed 350 million users and was cash-flow positive. [9] Murdoch, the media titan who stood atop an empire that includes the *Wall Street Journal* and Fox, had been outmaneuvered by “the kid.”

Why Study Facebook?

Looking at the “flavor of the month” and trying to distinguish the reality from the hype is a critical managerial skill. In Facebook’s case, there are a lot of folks with a vested interest in figuring out where the firm is headed. If you want to work there, are you signing on to a firm where your *stock options* and *401k* contributions are going to be worth something or worthless? If you’re an investor and Facebook goes public, should you short the firm or increase your holdings? Would you invest in or avoid firms that rely on Facebook’s business? Should your firm rush to partner with the firm? Would you extend the firm credit? Offer it better terms to secure its growing business, or worse terms because you think it’s a risky bet? Is this firm the next Google (underestimated at first, and now wildly profitable and influential), the next GeoCities (Yahoo! paid three billion dollars for it—no one goes to the site today), or the next Skype (deeply impactful with over four hundred forty million accounts worldwide, but not much of a profit generator)? The jury is still out on all this, but let’s look at the fundamentals with an eye to applying what we’ve learned. No one has a crystal ball, but we do have some key concepts that can guide our analysis. And there are a lot of broadly applicable managerial lessons that can be gleaned by examining Facebook’s successes and missteps.

Zuckerberg Rules!

Many entrepreneurs accept startup capital from venture capitalists (VCs), investor groups that provide funding in exchange for a stake in the firm,

and often, a degree of managerial control (usually in the form of a voting seat or seats on the firm's board of directors). Typically, the earlier a firm accepts VC money, the more control these investors can exert (earlier investments are riskier, so VCs can demand more favorable terms). VCs usually have deep entrepreneurial experience, a wealth of contacts, and can often offer important guidance and advice, but strong investor groups can oust a firm's founder and other executives if they're dissatisfied with the firm's performance.

At Facebook, however, Zuckerberg owns an estimated 20 percent to 30 percent of the company, and controls three of five seats on the firm's board of directors. That means that he's virtually guaranteed to remain in control of the firm, regardless of what investors say. Maintaining this kind of control is unusual in a startup, and his influence is a testament to the speed with which Facebook expanded. By the time Zuckerberg reached out to VCs, his firm was so hot that he could call the shots, giving up surprisingly little in exchange for their money.

KEY TAKEAWAYS

- Facebook was founded by a nineteen-year-old college sophomore and eventual dropout.
- It is currently the largest social network in the world, boasting more than three hundred fifty million members and usage rates that would be the envy of most media companies. The firm is now larger than MySpace in both the United States and worldwide.
- The firm's rapid rise is the result of network effects and the speed of its adoption placed its founder in a particularly strong position when negotiating with venture firms. As a result, Facebook founder Mark Zuckerberg retains significant influence over the firm.



- While revenue prospects remain sketchy, some reports have valued the firm at fifteen billion dollars, based largely on an extrapolation of a Microsoft stake.

QUESTIONS AND EXERCISES

1. Who started Facebook? How old was he then? Now? How much control does the founding CEO have over his firm? Why?
2. Which firms have tried to acquire Facebook? Why? What were their motivations and why did Facebook seem attractive? Do you think these bids are justified? Do you think the firm should have accepted any of the buyout offers? Why or why not?
3. As of late 2007, Facebook boasted an extremely high “valuation.” How much was Facebook allegedly “worth?” What was this calculation based on?
4. Why study Facebook? Who cares if it succeeds?



7.2 What's the Big Deal?

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Recognize that Facebook's power is allowing it to encroach on and envelop other Internet businesses.
2. Understand the concept of the "dark Web" and why some feel this may one day give Facebook a source of advantage vis-à-vis Google.
3. Understand the basics of Facebook's infrastructure, and the costs required to power the effort.

The prior era's Internet golden boy, Netscape founder Marc Andreessen, has said that Facebook is "an amazing achievement one of the most significant milestones in the technology industry in this decade." [1] Why is Facebook considered such a big deal?

First there's the growth: between December 2008 and 2009, Facebook was adding between six hundred thousand and a million users a day. It was as if every twenty-four hours, a group as big or bigger than the entire city of Boston filed into Facebook's servers to set up new accounts. Roughly half of Facebook users visit the site every single day, [2] with the majority spending thirty minutes or more getting their daily Facebook fix. [3] And it seems that Mom really is on Facebook (Dad, too); users thirty-five years and older



account for more than half of Facebook's daily visitors and its fastest growing population.^[4]

Then there's what these users are doing: Facebook isn't just a collection of personal home pages and a place to declare your allegiance to your friends. The integrated set of Facebook services encroaches on a wide swath of established Internet businesses. It's not just that the site offers tools for messaging and chat; it's the *first-choice* messaging tool for a generation. E-mail is for your professors, but Facebook is for friends. In photos, Google, Yahoo! and MySpace all spent millions to acquire photo sharing tools (Picasa, Flickr, and Photobucket, respectively). But Facebook is now the biggest photo-sharing site on the Web, taking in some twenty-eight million photos each day. ^[5] And watch out YouTube. Facebookers share eight million videos each month. YouTube will get you famous, but Facebook is a place most go to share clips they only want friends to see.

As for search, Facebook's got designs on that, too. Google and Bing index some Facebook content, but since much of Facebook is private, accessible only among friends, this represents a massive blind spot for Google search. Sites that can't be indexed by Google and other search engines are referred to as the dark Web. While Facebook's partnership with Microsoft currently offers web search results through Bing.com, Facebook has announced its intention to offer its own search engine with real-time access to up-to-the-minute results from status updates, links, and other information made available to you by your friends. If Facebook can tie together standard Internet search with its dark Web content, this just might be enough for some to break the Google habit.

Facebook is a kingmaker, opinion catalyst, and traffic driver. While in the prior decade news stories would carry a notice saying "copyright: do not distribute," today major news outlets, including the *New York Times*, offer Facebook icons alongside every copyrighted story, encouraging users to



“share” content on in their profile pages. Like digital photos, video, and instant messaging, link sharing is a Facebook sharp elbow to the competition. Suddenly Facebook gets space on a page along side Digg.com and Del.icio.us, even though those guys showed up first.

And Facebook is political—in big, regime-threatening ways. The site is considered such a powerful tool in the activist’s toolbox that China, Iran, and Syria are among nations that have, at times, attempted to block Facebook access within their borders. Egyptians have used the site to protest for democracy. Saudi women have used it to lobby for driving privileges. ABC News cosponsored U.S. presidential debates with Facebook. And Facebook cofounder Chris Hughes was even recruited by the Obama campaign to create my.barackobama.com, a social media site considered vital in the 2008 U.S. presidential victory. [6]

So What’s It Take to Run This Thing?

The Facebook cloud (the big group of connected servers that power the site) is scattered across multiple facilities, including server farms in San Francisco, Santa Clara, and northern Virginia. [7] The innards that make up the bulk of the system aren’t that different from what you’d find on a high-end commodity workstation. Standard hard drives and eight core Intel processors—just a whole lot of them lashed together through networking and software.

Much of what powers the site is open source software (OSS). A good portion of the code is in PHP (a scripting language particularly well-suited for Web site development), while the databases are in MySQL (a popular open source database). The object cache that holds frequently accessed objects in chip-based RAM instead of on slower hard drives is managed through the open source product called Memcache.



Other code components are written in a variety of languages, including C++, Java, Python, and Ruby, with access between these components managed by a proprietary code layer the firm calls Thrift. Facebook also developed its own media serving solution, called Haystack. Haystack coughs up photos 50 percent faster than more expensive, proprietary solutions, and since it's done in-house, it saves Facebook costs that other online outlets spend on third-party content delivery networks (CDN) like Akamai. Facebook receives some fifty million requests per second, [8] yet 95 percent of data queries can be served from a huge, distributed server cache that lives in over fifteen terabytes of RAM (objects like video and photos are stored on hard drives). [9]

Hot stuff (literally), but it's not enough. The firm raised several hundred million dollars more in the months following the fall 2007 Microsoft deal, focused largely on expanding the firm's server network to keep up with the crush of growth. The one hundred million dollars million raised in May 2008 was "used entirely for servers." [10] Facebook will be buying them by the thousands for years to come. And it'll pay a pretty penny to keep things humming. Estimates suggest the firm spends one million dollars a month on electricity, another half million a month on telecommunications bandwidth, and at least fifteen million dollars a year in office and data center rental payments. [11]

KEY TAKEAWAYS

- Facebook's position as the digital center of its members' online social lives has allowed the firm to envelop related businesses such as photo and video sharing, messaging, bookmarking, and link sharing.
- Much of the site's content is in the dark Web, unable to be indexed by Google or other search engines. Some suggest this may create an opportunity for Facebook to challenge Google in search.



- Facebook's can be a vital tool for organizers—presenting itself as both opportunity and threat to those in power, and an empowering medium for those seeking to bring about change.
- Facebook's growth requires a continued and massive infrastructure investment. The site is powered largely on commodity hardware, open source software, and proprietary code tailored to the specific needs of the service.

QUESTIONS AND EXERCISES

1. What is Facebook? How do people use the site? What do they “do” on Facebook?
2. What markets has Facebook entered? What factors have allowed the firm to gain share in these markets at the expense of established firms? In what ways does it enjoy advantages that a traditional new entrant in such markets would not?
3. What is the “dark Web” and why is it potentially an asset to Facebook? Why is Google threatened by Facebook's dark Web? What firms might consider an investment in the firm, if it provided access to this asset? Do you think the dark Web is enough to draw users to a Facebook search product over Google? Why or why not?
4. As Facebook grows, what kinds of investments continue to be necessary? On a per-user basis, what are the trends in these costs over time? Do you think Facebook should wait in making these investments? Why or why not?
5. Investments in servers and other capital expenses typically must be depreciated over time. What does this imply about how the firm's profitability is calculated?
6. How have media attitudes toward their copyrighted content changed over the past decade? Why is Facebook a potentially significant partner for firms like the *New York Times*? What does the *Times* stand to gain by



encouraging “sharing” its content? What do newspapers and others sites really mean when they encourage sites to “share?” What actually is being passed back and forth? Do you think this ultimately helps or undermines the *Times* and other newspaper and magazine sites? Why?

7.3 The Social Graph

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand the concept of feeds, why users rebelled against Facebook feeds, and why users eventually embraced this feature.
2. Recognize the two strategic resources that are most critical to Facebook’s competitive advantage and why Facebook was able to create these resources while MySpace has fallen short.



3. Appreciate that while Facebook’s technology can be easily copied, barriers to sustain any new entrant are extraordinarily high, and the likelihood that a firm will win significant share from Facebook by doing the same thing is considerably remote.

At the heart of Facebook’s appeal is a concept Zuckerberg calls the social graph, which refers to Facebook’s ability to collect, express, and leverage the connections between the site’s users, or as some describe it, “the global mapping of everyone and how they’re related.” [1] Think of all the stuff that’s on Facebook as a node or endpoint that’s connected to other stuff. You’re connected to other users (your friends), photos about you are tagged, comments you’ve posted carry your name, you’re a member of groups, you’re connected to applications you’ve installed—Facebook links them all. [2]

While MySpace and Facebook are often mentioned in the same sentence, from their founding these sites were conceived differently. It goes beyond the fact that Facebook, with its neat, ordered user profiles, looks like a planned community compared to the garish, Vegas-like free-for-all of MySpace. MySpace was founded by musicians seeking to reach out to unknown users and make them fans. It’s no wonder the firm, with its proximity to LA and ownership by News Corporation, is viewed as more of a media company. It has cut deals to run network television shows on its site, and has even established a record label. It’s also important to note that from the start anyone could create a MySpace identity, and this open nature meant that you couldn’t always trust what you saw. Rife with bogus profiles, even News Corporation’s Rupert Murdoch has had to contend with the dozens of bogus Ruperts who have popped up on the service! [3] Facebook, however, was established in the relatively safe cocoon of American undergraduate life, and was conceived as a place where you

could *reinforce* contacts among those who, for the most part, you already knew. The site was one of the first social networks where users actually identified themselves using their real names. If you wanted to be distinguished as working for a certain firm or as a student of a particular university, you had to verify that you were legitimate via an e-mail address issued by that organization. It was this “realness” that became Facebook’s distinguishing feature—bringing along with it a degree of safety and comfort that enabled Facebook to become a true social utility and build out a solid social graph consisting of verified relationships. Since “friending” (which is a link between nodes in the social graph) required both users to approve the relationship, the network fostered an incredible amount of trust. Today, many Facebook users post their cell phone numbers, their birthdays, offer personal photos, and otherwise share information they’d never do outside their circle of friends. Because of trust, Facebook’s social graph is stronger than MySpace’s.

There is also a strong network effect to Facebook (see [Chapter 5 "Understanding Network Effects"](#)). People are attracted to the service because others they care about are more likely to be there than anywhere else online. Without the network effect Facebook wouldn’t exist. And it’s because of the network effect that another smart kid in a dorm can’t rip off Zuckerberg in any market where Facebook is the biggest fish. Even an exact copy of Facebook would be a virtual ghost town with no social graph (see “It’s Not the Technology” below).

The switching costs for Facebook are also extremely powerful. A move to another service means recreating your entire social graph. The more time you spend on the service, the more you’ve invested in your graph and the less likely you are to move to a rival.



It's Not the Technology

Does your firm have Facebook envy? KickApps, an eighty-person startup in Manhattan, will give you the technology to power your own social network. All KickApps wants is a cut of the ads placed around your content. In its first two years, the site has provided the infrastructure for twenty thousand “mini Facebooks,” registering three hundred million page views a month. [4] NPR, ABC, AutoByTel, Harley Davidson, and Kraft all use the service (social networks for Cheez Whiz?).

There's also Ning, run by former Goldman Sachs analyst Gina Bianchini (Netscape founder Mark Andreessen is her CTO). Ning has over one million mini networks organized on all sorts of topics; families, radio personalities, church groups, vegans, diabetes sufferers; and is adding 1,500 to 2,000 a day.

Or how about the offering from Agriya Infoway, based in Chennai, India? The firm will sell you Kootali, a software package that lets developers replicate Facebook's design and features, complete with friend networks, photos, and mini-feeds. They haven't stolen any code, but they have copied the company's look and feel. Those with Zuckerberg ambitions can shell out the four hundred bucks for Kootali. Sites with names like Faceclub.com and Umicity.com have done just that—and gone nowhere.

Mini networks that extend the conversation (NPR) or make it easier to find other rapidly loyal product fans (Harley Davidson) may hold a niche for some firms. And Ning is a neat way for specialized groups to quickly form in a secure environment that's all their own (it's just us, no “creepy friends” from the other networks). While every market has a place for its niches, none of these will grow to compete with the dominant social networks. The value isn't in the technology; it's in what the technology has created over time. For Facebook, it's a huge user base that (for now at least) is not going anywhere else.



KEY TAKEAWAYS

- The social graph expresses the connections between individuals and organizations.
- Trust created through user verification and friend approval requiring both parties to consent encouraged Facebook users to share more and helped the firm establish a stronger social graph than MySpace or other social networking rivals.
- Facebook's key resources for competitive advantage are network effects and switching costs. These resources make it extremely difficult for copycat firms to steal market share from Facebook.

QUESTIONS AND EXERCISES

1. Which is bigger, Facebook or MySpace? How are these firms different? Why would a person or organization be attracted to one service over another?
2. What is the social graph? Why is Facebook's social graph considered to be stronger than the social graph available to MySpace users?
3. In terms of features and utility, how are Facebook and MySpace similar? How are they different? Why would a user choose to go to one site instead of another? Are you a member of either of these sites? Both? Why? Do you feel that they are respectively pursuing lucrative markets? Why or why not? If given the opportunity, would you invest in either firm? Why or why not?
4. If you were a marketer, which firm would you target for an online advertising campaign—Facebook or MySpace? Why?
5. Does Facebook have to worry about copycat firms from the United States? In overseas markets? Why or why not? If Facebook has a source (or sources) of competitive advantage, explain these. If it has no advantage, discuss why.

7.4 Facebook Feeds—Ebola for Data Flows

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand the concept of feeds, why users rebelled, and why users eventually embraced this feature.
2. Recognize the role of feeds in viral promotions, catalyzing innovation, and supporting rapid organizing.

While the authenticity and trust offered by Facebook was critical, offering News Feeds concentrated and released value from the social graph. With feeds, each time a user performs an activity in Facebook—makes a friend, uploads a picture, joins a group—the feed blasts this information to all of your friends in a reverse chronological list that shows up right when they next log on. An individual user’s activities are also listed within a mini Feed that shows up on their profile. Get a new job, move to a new city, read a great article, have a pithy quote—post it to Facebook—the feed picks it up, and the world of your Facebook friends will get an update.

Feeds are perhaps the linchpin of Facebook’s ability to strengthen and deliver user value from the social graph, but for a brief period of time it looked like feeds would kill the company. News Feeds were launched on September 5, 2006, just as many of the nation’s undergrads were arriving



on campus. Feeds reflecting any Facebook activity (including changes to the relationship status) became a sort of gossip page splashed right when your friends logged in. To many, feeds were first seen as a viral blast of digital nosiness—a release of information they hadn’t consented to distribute widely.

And in a remarkable irony, user disgust over the News Feed ambush offered a whip crack demonstration of the power and speed of the feed virus.

Protest groups formed, and every student who, for example, joined a group named Students Against Facebook News Feed, had this fact blasted to their friends (along with a quick link where friends, too, could click to join the group). Hundreds of thousands of users mobilized against the firm in just twenty-four hours. It looked like Zuckerberg’s creation had turned on him, Frankenstein style.

The first official Facebook blog post on the controversy came off as a bit condescending (never a good tone to use when your customers feel that you’ve wronged them). “Calm down. Breathe. We hear you,” wrote Zuckerberg on the evening of September 5. The next post, three days after the News Feed launch, was much more contrite (“We really messed this one up,” he wrote). In the 484-word open letter, Zuckerberg apologized for the surprise, explaining how users could opt out of feeds. The tactic worked, and the controversy blew over. ^[1] The ability to stop personal information from flowing into the feed stream was just enough to stifle critics, and as it turns out, a lot of people really liked the feeds and found them useful. It soon became clear that if you wanted to use the Web to keep track of your social life and contacts, Facebook was the place to be. Not only did feeds not push users away, by the start of the next semester subscribers had nearly doubled!



KEY TAKEAWAYS

- Facebook feeds foster the viral spread of information and activity.
- Feeds were initially unwanted by many Facebook users. Feeds themselves helped fuel online protests against the feed feature.
- Today feeds are considered one of the most vital, value-adding features to Facebook and other social networking sites.
- Users often misperceive technology and have difficulty in recognizing an efforts value (as well as its risks). They have every right to be concerned and protective of their privacy. It is the responsibility of firms to engage users on new initiatives and to protect user privacy. Failure to do so risks backlash.

QUESTIONS AND EXERCISES

1. What is the “lynchpin” of Facebook’s ability to strengthen and deliver user-value from the social graph?
2. How did users first react to feeds? What could Facebook have done to better manage the launch?
3. How do you feel about Facebook feeds? Have you ever been disturbed by information about you or someone else that has appeared in the feed? Did this prompt action? Why or why not?
4. Visit Facebook and experiment with privacy settings. What kinds of control do you have over feeds and data sharing? Is this enough to set your mind at ease? Did you know these settings existed before being prompted to investigate features?
5. What other Web sites are leveraging features that mimic Facebook feeds? Do you think these efforts are successful or not? Why?



7.5 F8—Facebook as a Platform

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand how Facebook created a platform, and the potential value this offers the firm.
2. Recognize that running a platform also presents a host of challenges to the platform operator.

In May 2007, Facebook followed News Feeds with another initiative that set it head and shoulders above its competition. At a conference called F8 (pronounced “fate”), Mark Zuckerberg stood on stage and announced that he was opening up the screen real estate on Facebook to other application developers. Facebook published a set of application programming interfaces (APIs) that specified how programs could be written to run within and interact with Facebook. Now any programmer could write an application that would run inside a user’s profile. Geeks of the world, Facebook’s user base could be yours! Just write something good.

Developers could charge for their wares, offer them for free, and even run ads. And Facebook let developers keep what they made. This was a key distinction; MySpace initially restricted developer revenue on the few products designed to run on their site, at times even blocking some applications. The choice was clear, and developers flocked to Facebook. To promote the new apps, Facebook would run an Applications area on the site where users could browse offerings. Even better, News Feed was a viral injection that spread the word each time an application was installed. Your best friend just put up a slide show app? Maybe you’ll check it out, too. The predictions of one billion dollars in social network ad spending were geek



catnip, and legions of programmers wanted a little bit for themselves. Apps could be cobbled together on the quick, feeds made them spread like wildfire, and the early movers offered adoption rates never before seen by small groups of software developers. People began speaking of the Facebook Economy. Facebook was considered a platform. Some compared it to the next Windows, Zuckerberg the next Gates (hey, they both dropped out of Harvard, right?).

And each application potentially added more value and features to the site without Facebook lifting a finger. The initial event launched with sixty-five developer partners and eighty-five applications. There were some missteps along the way. Some applications were accused of spamming friends with invites to install them. There were also security concerns and apps that violated the intellectual property of other firms (see the “Scrabulous” sidebar below), but Facebook worked to quickly remove errant apps, improve the system, and encourage developers. Just one year in, Facebook had marshaled the efforts of some four hundred thousand developers and entrepreneurs, twenty-four thousand applications had been built for the platform, 140 new apps were being added each day, and 95 percent of Facebook members had installed at least one Facebook application. As Sarah Lacy, author of *Once You’re Lucky, Twice You’re Good*, put it, “with one masterstroke, Zuck had mobilized all of Silicon Valley to innovate for him.”

With feeds to spread the word, Facebook was starting to look like the first place to go to launch an online innovation. Skip the Web, bring it to Zuckerberg’s site first. Consider iLike: within the first three months, the firm saw installs of its Facebook app explode to seven million, more than doubling the number of users the firm was able to attract through the Web site it introduced the previous year. iLike became so cool that by September, platinum rocker KT Tunstall was debuting tracks through the



Facebook service. A programmer named Mark Pincus wrote a Texas hold 'em game at his kitchen table. [1] Today his social gaming firm, Zynga, is a powerhouse, with over three dozen apps, over sixty million users, and more than one hundred million dollars in annual revenue. [2] Zynga games include MafiaWars, Vampires, and the wildly successful Farmville, which boasts some twenty times the number of actual farms in the United States. Seth Goldstein, the CEO of SocialMedia, which developed the Food Fight app, claimed that a buffalo wing chain offered the firm eighty thousand dollars to sponsor virtual wings that Facebook users could hurl at their friends. [3] Lee Lorenzen, founder of Altura Ventures, an investment firm exclusively targeting firms creating Facebook apps, said, "Facebook is God's gift to developers. Never has the path from a good idea to millions of users been shorter." [4]

I Majored in Facebook

Once Facebook became a platform, Stanford professor B. J. Fogg thought it would be a great environment for a programming class. In ten weeks his seventy-three students built a series of applications that collectively received over sixteen million installs. By the final week of class, several applications developed by students, including KissMe, Send Hotness, and Perfect Match, had all received millions of users, and class apps collectively generated more than half a million dollars in ad revenue. At least three companies were formed from the course.

Some of the folks developing Facebook applications have serious geek cred. Max Levchin, who (along with Peter Theil) founded PayPal, has widely claimed that his widget company, Slide, would be even bigger than his previous firm (for the record, PayPal sold to eBay for \$1.5 billion in 2002). Legg Mason and Fidelity have each already invested at levels giving the firm a five-hundred-million-dollar valuation. [5] But is there that much money to



be made from selling ads and mining data on SuperPoke? Levchin's firm is working to craft new advertising measures to track and leverage engagement. Just what are users doing with the applets? With whom? And how might advertisers use these insights to build winning campaigns? Levchin thinks Slide can answer these questions and capture a ton of business along the way.

But legitimate questions remain. Are Facebook apps really a big deal? Just how important will apps be to adding sustained value within Facebook? And how will firms leverage the Facebook framework to extract their own value? A chart from FlowingData showed the top category, Just for Fun, was larger than the next four categories combined. That suggests that a lot of applications are faddish time wasters. Yes, there is experimentation beyond virtual Zombie Bites. Visa has created a small business network on Facebook (Facebook had some eighty thousand small businesses online at the time of Visa's launch). Educational software firm Blackboard offered an application that will post data to Facebook pages as soon as there are updates to someone's Blackboard account (new courses, whether assignments or grades have been posted, etc.). We're still a long way from Facebook as a Windows rival, but F8 helped push Facebook to #1, and it continues to deliver quirky fun (and then some) supplied by thousands of developers off its payroll.

Scrabulous

Rajat and Jayant Agarwalla, two brothers in Kolkata, India, who run a modest software development company, decided to write a Scrabble clone as a Facebook application. The app, named Scrabulous, was social—users could invite friends to play, or they could search for new players looking for an opponent. Their application was a smash, snagging three million registered users and seven hundred thousand players a day after just a few months. Scrabulous was featured in *PC World's* 100 best products of 2008,

received coverage in the *New York Times*, *Newsweek*, and *Wired*, and was pulling in about twenty-five thousand dollars a month from online advertising. Way to go, little guys! [6]

There is only one problem: the Agarwalla brothers didn't have the legal rights to Scrabble, and it was apparent to anyone that from the name to the tiles to the scoring—this was a direct rip off of the well-known board game. Hasbro owns the copyright to Scrabble in the United States and Canada; Mattel owns it everywhere else. Thousands of fans joined Facebook groups with names like “Save Scrabulous” and “Please God, I Have So Little: Don't Take Scrabulous, Too.” Users in some protest groups pledged never to buy Hasbro games if Scrabulous was stopped. Even if the firms wanted to succumb to pressure and let the Agarwalla brothers continue, they couldn't. Both Electronic Arts and RealNetworks have contracted with the firms to create online versions of the game.

While the Facebook Scrabulous app is long gone, the tale shows just one of the challenges of creating a platform. In addition to copyright violations, app makers have crafted apps that annoy, raise privacy and security concerns, purvey pornography, or otherwise step over the boundaries of good taste. Firms from Facebook to Apple (through its iTunes app store) have struggled to find the right mix of monitoring, protection, and approval while avoiding cries of censorship.

KEY TAKEAWAYS

- Facebook's platform allows the firm to further leverage the network effect. Developers creating applications create complementary benefits that have the potential to add value to Facebook beyond what the firm itself provides to its users.
- There is no revenue-sharing mandate among platform partners—whatever an application makes can be kept by its developers.

- Most Facebook applications are focused on entertainment. The true, durable, long-term value of Facebook’s platform remains to be seen.
- Despite this, some estimates claimed Facebook platform developers would earn more than Facebook itself in 2009.
- Running a platform can be challenging. Copyright, security, appropriateness, free speech tensions, efforts that tarnish platform operator brands, privacy, and the potential for competition with partners, all can make platform management more complex than simply creating a set of standards and releasing this to the public.

QUESTIONS AND EXERCISES

1. Why did developers prefer to write apps for Facebook than MySpace?
2. What competitive asset does the F8 initiative help Facebook strengthen? For example, how does F8 make Facebook stronger when compared to rivals?
3. What’s Scrabulous? Did the developers make money? What happened to the firm and why?
4. Have you used Facebook apps? Which are your favorites? What makes them successful?
5. Leverage your experience or conduct additional research—are there developers who you feel have abused the Facebook app network? Why? What is Facebook’s responsibility (if any) to control such abuse?
6. How do most app developers make money? Have you ever helped a Facebook app developer earn money? How or why not?
7. How do Facebook app revenue opportunities differ from those leveraged by a large portion of iTunes app store apps?

7.6 Advertising and Social Networks: A Work in Progress

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Describe the differences in the Facebook and Google ad models.
2. Explain the Hunt versus Hike metaphor, contrast the relative success of ad performance on search compared to social networks, and understand the factors behind the latter's struggles.
3. Recognize how firms are leveraging social networks for brand and product engagement, be able to provide examples of successful efforts, and give reasons why such engagement is difficult to achieve.

If Facebook is going to continue to give away its services for free, it needs to make money somehow. Right now that means advertising. Fortunately for the firm, online advertising is hot. For years, online advertising has been the only major media category that has seen an increase in spending (see [Chapter 8 "Google: Search, Online Advertising, and Beyond..."](#)). And



while 20 percent of media consumption happens online, the Internet attracts just 6 percent of advertising dollars, suggesting there's plenty of growth still ahead. [1] Firms spend more advertising online than they do on radio ads, and the Internet will soon beat out spending on cable TV and magazine ads. [2] But not all Internet advertising is created equal. And there are signs that social networking sites are struggling to find the right ad model.

In early 2008, Google founder Sergey Brin stated, "I don't think we have the killer best way to advertise and monetize social networks yet." Brin went on to state that social networking ad inventory as a whole was proving problematic and that the "monetization work we were doing there didn't pan out as well as we had hoped." [3] When Google ad partner Fox Interactive Media (the News Corporation division that contains MySpace) announced that revenue would fall one hundred million dollars short of projections, News Corporation's stock tumbled 5 percent, analysts downgraded the company, and the firm's chief revenue officer was dismissed. [4]

Why aren't social networks having the success of Google and other sites? Problems advertising on these sites include content adjacency, and user attention. The *content adjacency* problem refers to concern over where a firm's advertisements will run. Look at all of the creepy titles in social networking news groups. Do advertisers really want their ads running alongside conversations that are racy, offensive, illegal, or that may even mock their products? This potential juxtaposition is a major problem with any site offering ads adjacent to free-form social media. Summing up industry wariness, one P&G manager said, "What in heaven's name made you think you could monetize the real estate in which somebody is breaking up with their girlfriend?" [5] An IDC report suggests that it's because of



content adjacency that “brand advertisers largely consider user-generated content as low-quality, brand-unsafe inventory” for running ads. [6]

Now let’s look at the user attention problem.

Attention Challenges: The Hunt Versus The Hike

In terms of revenue model, Facebook is radically different from Google and the hot-growth category of search advertising. Users of Google and other search sites are on a *hunt*—a task-oriented expedition to collect information that will drive a specific action. Search users want to learn something, buy something, research a problem, or get a question answered. To the extent that the hunt overlaps with ads, it works. Just searched on a medical term? Google will show you an ad from a drug company. Looking for a toy? You’ll see Google ads from eBay sellers and other online shops. Type in a vacation destination and you get a long list of ads from travel providers aggressively courting your spending. Even better, Google only charges text advertisers when a user clicks through. No clicks? The ad runs at no cost to the firm. From a return on investment perspective, this is extraordinarily efficient. How often do users click on Google ads? Enough for this to be the single most profitable activity among *any* Internet firm. In 2008, Google revenue totaled nearly twenty-two billion dollars. Profits exceeded \$4.2 billion, almost all of this from pay-per-click ads (see [Chapter 8 "Google: Search, Online Advertising, and Beyond..."](#) for more details).

While users go to Google to hunt, they go to Facebook as if they were going on a *hike*—they have a rough idea of what they’ll encounter, but they’re there to explore and look around, enjoy the sights (or site). They’ve usually allocated time for fun and they don’t want to leave the terrain when they’re having conversations, looking at photos or videos, and checking out updates from friends.



These usage patterns are reflected in click-through rates. Google users click on ads around 2 percent of the time (and at a much higher rate when searching for product information). At Facebook, click-throughs are about 0.04 percent. [7]

Most banner ads don't charge per click, but rather via something called CPM—that's cost per thousand impressions (an impression being each time an ad appears on someone's screen). The holy grail of Internet advertising is *targeting*. The more focused an audience, the more likely a firm is to attract advertisers willing to pay a higher CPM to reach those users. Rates vary widely. In 2008, MySpace lowered its banner ad rate from \$3.25 CPM to less than two dollars. By contrast, TechTarget, a Web publisher focusing on technology professionals, is able to command CPM rates of one hundred dollars and above (an ad inventory that valuable helped the firm go public in 2007). *Technology Review* magazine boasts a CPM of seventy dollars. The social networking blog Mashable has CPM rates ranging between seven and thirty-three dollars. But Facebook ads go for much, much less. Lookery, a one-time ad network that bought ad space on Facebook in bulk, had been reselling inventory at a CPM of 7.5 cents (note that Facebook does offer advertisers pay-per-click as well as impression-based, or CPM, options). Even Facebook ads with a bit of targeting are poorly priced (Facebook's Social Ads, which allow advertisers to target users according to location and age, have a floor price of fifteen cents CPM). [8]

Getting Creative with Promotions: Does It Work?

Facebook and other social networks are still learning what works. Ad inventory displayed on high-traffic home pages have garnered big bucks for firms like Yahoo! With MySpace and Facebook offering advertisers greater audience reach than most network television programs, there's little reason



to suggest that chunks of this business won't eventually flow to the social networks. But even more interesting is how Facebook and widget sites have begun to experiment with relatively new forms of advertising. Many feel that Facebook has a unique opportunity to get consumers to engage with their brand, and some initial experiments point where this may be heading. Many firms have been leveraging so-called engagement ads by making their products part of the Facebook fun. COO Sheryl Sandberg discussed Ben and Jerry's promotion for the ice cream chain's free cone day event. To promote the upcoming event, Ben and Jerry's initially contracted to make two hundred and fifty thousand "gift cones" available to Facebook users; they could click on little icons that would gift a cone icon to a friend, and that would show up in their profile. Within a couple of hours, customers had sent all two hundred and fifty thousand virtual cones. Delighted, Ben and Jerry's bought another two hundred and fifty thousand cones. Within eleven hours, half a million people had sent cones, many making plans with Facebook friends to attend the real free cone day. The day of the Facebook promotion, Ben and Jerry's Web site registered fifty-three million impressions, as users searched for store locations and wrote about their favorite flavors.^[9] The campaign dovetailed with everything Facebook was good at: it was viral, generating enthusiasm for a promotional event and even prompting scheduling.

In other promotions, Paramount Pictures gave away two hundred and fifty thousand virtual fedoras for "Indiana Jones and the Kingdom of the Crystal Skull." They sold out within hours, with users rallying friends, family, and colleagues to see the movie. Honda gave away three quarters of a million hearts during Valentine's Day 2009. ^[10] The Dr. Pepper Snapple Group offered two hundred and fifty thousand virtual Sunkist sodas, which earned the firm one hundred thirty million brand impressions in twenty-two



hours. Says Sunkist's brand manager, "A Super Bowl ad, if you compare it, would have generated somewhere between six to seven million." [11]

Of course, even with this business, Facebook may find that it competes with widget makers. Slide also offers wacky advertising programs through its own Facebook apps (the firm sits atop the most popular application developer rankings). Fox Searchlight went to Slide to allow friends to throw orange Tic Tacs at each other as part of a promotion for the movie "Juno." Coke used Slide to distribute virtual Vitamin Water. By some estimates, in 2009, Facebook app developers took in well over half a billion dollars—exceeding Facebook's own haul. [12] Unlike Apple's app store (where much of developer-earned revenue comes from selling apps), the vast majority of Facebook apps are free and ad-supported. That means Facebook and its app providers are in competition to monetize Facebook users, and both are running at a finite pot of recession-constrained advertising dollars.

While these efforts might be quirky and fun, are they even effective? Some of these programs are considered successes; others, not so much. Jupiter Research surveyed marketers trying to create a viral impact online and found that only about 15 percent of these efforts actually caught on with consumers. [13] While the Ben and Jerry's gift cones were used up quickly, a visit to Facebook in the weeks after this campaign saw CareerBuilder, Wide Eye Caffeinated Spirits, and Coors Light icons lingering days after their first appearance. Brands seeking to deploy their own applications in Facebook have also struggled. *New Media Age* reported that applications rolled out by top brands such as MTV, Warner Brothers, and Woolworths were found to have as little as five daily users. Congestion may be setting in for all but the most innovative applications, as standing out in a crowd of over twenty-three thousand applications becomes increasingly difficult. [14]

To its credit, consumer products giant P&G has been relentlessly experimenting with leveraging social networks for brand engagement, but

the results show what a tough slog this can be. The firm did garner fourteen thousand Facebook “fans” for its Crest Whitestrips product, but those fans were earned while giving away free movie tickets and other promos.

The *New York Times* quipped that with those kinds of incentives, “even a hemorrhoid cream” could have attracted a similar group of “fans,” and when the giveaways stopped, thousands promptly “unfanned” Whitestrips.

Results for Proctor and Gamble’s “2X Ultra Tide” fan page were also pretty grim. P&G tried offbeat appeals for customer-brand bonding, including asking Facebookers to post “their favorite places to enjoy stain-making moments.” But a check eleven months after launch had garnered just eighteen submissions, two from P&G, two from staffers at spoof news site *The Onion*, and a bunch of short posts such as “Tidealicious!” ^[15]

Efforts around engagement opportunities like events (Ben and Jerry’s) or products consumer are anxious to identify themselves with (a band or a movie) may have more success than trying to promote consumer goods that otherwise offer little allegiance, but efforts are so new that metrics are scarce, impact is tough to gauge, and best practices are still unclear.

KEY TAKEAWAYS

- Content adjacency and user attention make social networking ads less attractive than search and professionally produced content sites.
- Google enjoys significantly higher click through rates than Facebook.
- Display ads are often charged based on impression. Social networks also offer lower CPM rates than many other, more targeted Web sites.
- Social networking has been difficult to monetize, as users are online to engage friends, not to hunt for products or be drawn away by clicks.
- Many firms have begun to experiment with engagement ads. While there have been some successes, engagement campaigns often haven’t yielded significant results.

QUESTIONS AND EXERCISES

1. How are most display ads billed? What acronym is used to describe pricing of most display ads?
2. How are most text ads billed?
3. Contrast Facebook and Google click through rates. Contrast Facebook CPMs with CPMs at professional content sites. Why the discrepancy?
4. What is the content adjacency problem? Search for examples of firms that have experienced embracement due to content adjacency—describe them, why they occurred, and if site operators could have done something to reduce the likelihood these issues could have occurred.
5. What kinds of Web sites are most susceptible to content adjacency? Are news sites? Why or why not? What sorts of technical features might act as breeding grounds for content adjacency problems?
6. If a firm removed user content because it was offensive to an advertiser, what kinds of problems might this create? When (if ever) should a firm remove or take down user content?
7. How are firms attempting to leverage social networks for brand and product engagement?
8. What are the challenges that social networking sites face when trying to woo advertisers?
9. Describe an innovative marketing campaign that has leveraged Facebook or other social networking site. What factors made this campaign work? Are all firms likely to have this sort of success? Why or why not?
10. Have advertisers ever targeted you when displaying ads on Facebook? How were you targeted? What did you think of the effort?

7.7 Beacon Busted



LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand the difference between opt-in and opt-out efforts
2. Recognize how user issues and procedural implementation can derail even well-intentioned information systems efforts.
3. Recognize the risks in being a pioneer associated with new media efforts, and understand how missteps led to Facebook and its partners being embarrassed (and in some cases sued) by Beacon's design and rollout issues.

Conventional advertising may grow into a great business for Facebook, but the firm was clearly sitting on something that was unconventional compared to prior generations of Web services. Could the energy and virulent nature of social networks be harnessed to offer truly useful, consumer information to its users? Word of mouth is considered the most persuasive (and valuable) form of marketing, ^[1] and Facebook was a giant word of mouth machine. What if the firm worked with vendors and grabbed consumer activity at the point of purchase to put into the News Feed and post to a user's profile? If you then rented a video, bought a cool product, or dropped something in your wish list, your buddies could get a heads up and they might ask you about it. The person being asked feels like an expert, the person with the question gets a frank opinion, and the vendor providing the data just might get another sale. It looked like a home run.

This effort, named Beacon, was announced in November 2007. Some forty e-commerce sites signed up, including Blockbuster, Fandango, eBay, Travelocity, Zappos, and the *New York Times*. Zuckerberg was so confident of the effort that he stood before a group of Madison Avenue ad executives and declared that Beacon would represent a "once-in-a-hundred-years" fundamental change in the way media works.

Like News Feeds, user reaction was swift and brutal. The commercial activity of Facebook users began showing up without their consent. The biggest problem with Beacon was that it was “opt out” instead of “opt in.” Facebook (and its partners) assumed users would agree to sharing data in their feeds. A pop-up box did appear briefly on most sites supporting Beacon, but it disappeared after a few seconds. [2] Many users, blind to these sorts of alerts, either clicked through or ignored the warnings. And well...there are some purchases you might not want to broadcast to the world.

“Facebook Ruins Christmas for Everyone!” screamed one headline from MSNBC.com. Another from *U.S. News and World Report* read “How Facebook Stole Christmas.” *The Washington Post* ran the story of Sean Lane, a twenty-eight-year-old tech support worker from Waltham, Massachusetts, who got a message from his wife just two hours after he bought a ring on Overstock.com. “Who is this ring for?” she wanted to know. Facebook had not only posted a feed that her husband had bought the ring, but also that he got it for a 51 percent discount! Overstock quickly announced that it was halting participation in Beacon until Facebook changed its practice to opt in. [3]

MoveOn.org started a Facebook group and online petition protesting Beacon. The Center for Digital Democracy and the U.S. Public Interest Research Group asked the Federal Trade Commission to investigate Facebook’s advertising programs. And a Dallas woman sued Blockbuster for violating the Video Privacy Protection Act (a 1998 U.S. law prohibiting unauthorized access to video store rental records).

To Facebook’s credit, the firm acted swiftly. Beacon was switched to an opt-in system, where user consent must be given before partner data is sent to the feed. Zuckerberg would later say regarding Beacon: “We’ve made a lot of mistakes building this feature, but we’ve made even more with how we’ve



handled them. We simply did a bad job with this release, and I apologize for it.” [4] Despite the Beacon fiasco, new users continued to flock to the site, and loyal users stuck with Zuck. Perhaps a bigger problem was that many of those forty A-list e-commerce sites that took a gamble with Facebook now had their names associated with a privacy screw up that made headlines worldwide. A manager so burned isn’t likely to sign up first for the next round of experimentation. Several months later, some thirty firms were using the opt-in system, but uptake was limited and that hundred-year change Zuckerberg talked about looked like it had been put on hold.

From the Prada example in Chapter 1 we learned that savvy managers look beyond technology and consider complete information systems—not just the hardware and software of technology but also the interactions among the data, people, and procedures that make up (and are impacted by) information systems. Beacon’s failure is a cautionary tale of what can go wrong if users fail to broadly consider the impact and implications of an information system on all those it can touch. Technology’s reach is often farther, wider, and more significantly impactful than we originally expect.

Reputation Damage and Increased Scrutiny—The Facebook TOS Debacle

Facebook also suffered damage to its reputation, brand, and credibility, further reinforcing perceptions that the company acts brazenly, without considering user needs, and is fast and loose on privacy and user notification. Facebook worked through the feeds outrage, eventually convincing users of the benefits of feeds. But Beacon was a fiasco. And now users, the media, and watchdogs were on the alert.

When the firm modified its terms of service (TOS) policy in Spring 2009, the uproar was immediate. As a cover story in *New York* magazine summed it up, Facebook’s new TOS appeared to state, “We can do anything we want

with your content, forever,” even if a user deletes their account and leaves the service. [5] Yet *another* privacy backlash!

Activists organized, the press crafted juicy, attention-grabbing headlines, and the firm was forced once again to backtrack. But here’s where others can learn from Facebook’s missteps and response. The firm was contrite and reached out to explain and engage users. The old TOS was reinstated, the firm posted a proposed new version that gave the firm broad latitude in leveraging user content without claiming ownership. And the firm renounced the right to use this content if a user closed their Facebook account. This new TOS was offered in a way that solicited user comments, and it was submitted to a community vote, considered binding if 30 percent of Facebook users participated. Zuckerberg had turned Facebook into a democracy and helped empower users to determine the firm’s next step. Despite the uproar, only about 1 percent of Facebook users eventually voted on the measure, but the firm 74 percent to 26 percent ruling in favor of the change gave Facebook some cover to move forward. [6] This event also demonstrates that a tempest can be generated by a relatively small number of passionate users. Firms ignore the vocal and influential at their own peril!

In Facebook’s defense, the broad TOS was probably more a form of legal protection than any nefarious attempt to exploit all user posts ad infinitum. The U.S. legal environment does require that explicit terms be defined and communicated to users, even if these are tough for laypeople to understand. But a “trust us” attitude toward user data doesn’t work, particularly for a firm considered to have committed ham-handed gaffes in the past. Managers must learn from the freewheeling Facebook community. Your actions are now subject to immediate and sustained review. Violate the public trust and expect the equivalent of a high-powered investigative

microscope examining your every move, and a very public airing of the findings.

KEY TAKEAWAYS

- World of mouth is the most powerful method for promoting products and services, and Beacon was conceived as a giant word-of-mouth machine with win-win benefits for firms, recommenders, recommendation recipients, and Facebook.
- Beacon failed because it was an opt-out system that was not thoroughly tested beforehand, and because user behavior, expectations, and system procedures were not completely taken into account.
- Partners associated with the rapidly rolled out, poorly conceived, and untested effort were embarrassed. Several faced legal action.
- Facebook also reinforced negative perceptions regarding the firm's attitudes toward users, notification, and their privacy. This attitude only served to focus a continued spotlight on the firm's efforts, and users became even less forgiving.
- Activists and the media were merciless in criticizing the firm's Terms of Service changes. Facebook's democratizing efforts demonstrate lessons other organizations can learn from, regarding user scrutiny, public reaction, and stakeholder engagement.

QUESTIONS AND EXERCISES

1. What is Beacon? Why was it initially thought to be a good idea? What were the benefits to firm partners, recommenders, recommendation recipients, and Facebook? Who were Beacon's partners and what did they seek to gain through the effort?
2. Describe "the biggest problem with Beacon?" Would you use Beacon? Why or why not?



3. How might Facebook and its partners have avoided the problems with Beacon? Could the effort be restructured while still delivering on its initial promise? Why or why not?
4. Beacon shows the risk in being a pioneer—are there risks in being too cautious and not pioneering with innovative, ground-floor marketing efforts? What kinds of benefits might a firm miss out on? Is there a disadvantage in being late to the party with these efforts, as well? Why or why not?
5. Why do you think Facebook changed its Terms of Service? Did these changes concern you? Were users right to rebel? What could Facebook have done to avoid the problem? Did Facebook do a good job in follow up? How would you advise Facebook to apply lessons learned from the TOS controversy?

7.8 Predators and Privacy

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand the extent and scope of the predator problem on online social networks.
2. Recognize the steps firms are taking to proactively engage and limit these problems.

While spoiling Christmas is bad, sexual predators are far worse, and in October 2007, Facebook became an investigation target. Officials from the New York State Attorney General’s office had posed as teenagers on Facebook and received sexual advances. Complaints to the service from investigators posing as parents were also not immediately addressed. These



were troubling developments for a firm that prided itself on trust and authenticity.

In a 2008 agreement with forty-nine states, Facebook offered aggressive programs, many of which put it in line with MySpace. MySpace had become known as a lair for predators and after months of highly publicized tragic incidents, had become very aggressive about protecting minors. To get a sense of the scope of the problem, consider that MySpace claimed that it had found and deleted some twenty-nine thousand accounts from its site after comparing profiles against a database of convicted sex offenders. [1] Following MySpace's lead, Facebook agreed to respond to complaints about inappropriate content within twenty-four hours and to allow an independent examiner to monitor how it handles complaints. The firm imposed age-locking restrictions on profiles, reviewing any attempt by someone under the age of eighteen to change their date of birth. Profiles of minors were no longer searchable. The site agreed to automatically send a warning message when a child is at risk of revealing personal information to an unknown adult. And links to explicit material, the most offensive Facebook groups, and any material related to cyberbullying were banned.

Busted on Facebook

Chapter 6 "Peer Production, Social Media, and Web 2.0" warned that your digital life will linger forever, and that employers are increasingly plumbing the depths of virtual communities in order to get a sense of job candidates. And it's not just employers. Sleuths at universities and police departments have begun looking to Facebook for evidence of malfeasance. Oxford University fined graduating students more than £10,000 for their post-exam celebrations, evidence of which was picked up from Facebook. Police in Vermont have made underage drinking busts and issued graffiti warnings based on Facebook photos, too. Beware—the Web knows!



KEY TAKEAWAYS

- Thousands of sex offenders and other miscreants have been discovered on MySpace, Facebook, and other social networks. They are a legitimate risk to the community and they harm otherwise valuable services.
- A combination of firm policies, computerized and human monitoring, aggressive reporting and follow-up, and engagement with authorities can reduce online predator risks.
- Firms that fail to fully engage this threat put users and communities at risk and may experience irreparable damage to firms and reputations.

QUESTIONS AND EXERCISES

1. How big is the predator problem on MySpace? What efforts have social networks employed to cut down on the number of predators online?
2. Investigate the current policies regarding underage users on Facebook. Do you think the firm adequately protects its users? Why or why not?
3. What age is appropriate for users to begin using social networks? Which service are appropriate at which ages? Are there social networks targeted at very young children? Do you think that these safe places? Why or why not?



7.9 Walled Garden or Open Field?

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand the concept of the walled garden.
2. Describe Open Social, Facebook Connect, and the contrasting motivations of their backers.
3. List and discuss the potential benefits and risks of engaging in the kinds of inter-site sharing and collaboration efforts described in this section.

The approach of keeping users nestled in the confines of a single service is sometimes called the walled garden. In this approach, a closed network or single set of services is controlled by one dominant firm. The prior iteration of America Online was a walled garden, and the term is often used when referring to cell phone providers. While this is changing, U.S. phone carriers were typically able to decide which handsets worked on their networks, which features they could offer, and which services can be accessed (see [Chapter 8 "Google: Search, Online Advertising, and Beyond..."](#) to learn how the firm is working to change this, too). And at first, Facebook was a walled garden. Sitting atop the largest number of social network users, the largest set of online photos, and the largest number of gadgets and widgets, Facebook seems to have found the sweet



spot. Network effects suggest that users and application developers will show up first to the place where all the action is. As Facebook continues its influence, its control becomes more and more absolute.

But Google has a different vision. The firm is trying to establish a set of standards that rival Facebook's APIs. The effort, called OpenSocial, allows a developer to write an application for one service, then let it run on any site that adopts the OpenSocial framework. Google, along with other sites, is also pushing data portability. In the grand vision, users would be able to consolidate all of their social networking information regardless of where it was created. Friends on MySpace, LinkedIn, and Gmail could all be reachable through a single contact list and inbox. Feeds would be consolidated in one view regardless of where these actions were performed. A photo posted on Flickr, a Twitter tweet, an article gets a Digg or bookmark—in theory you could combine these in a single site so that you'd have just one place to look instead of many. And efforts would promote data portability, in theory making it easier to pack up your data and take it with you when you leave one site for another. Other efforts, like OpenID, promise a single sign-on—just one password that could grant access across all compliant sites. It's a neat vision, one where social networks could grow from anywhere, linked via a sort of connective tissue of openness that lets data and features flow, regardless of where they were created.

Then a strange thing happened—Facebook opened up too. Of course, it's not playing nice with Google's standards—it has its own, built around something it calls Facebook Connect—but the service is remarkably open, and getting more so.

According to Dave Morin, Facebook's Senior Platform Manager, "Facebook Connect has threefold aim: transporting your Facebook identity, making your friends lists portable, and seeing detailed activity feeds from what your friends do across the Web." [1]



The ability to transport your Facebook ID is fast realizing a goal that has thwarted firms from Microsoft to Sun to Yahoo!—enabling a single sign-on that can be used for different services across the Web. You can now visit a site like CBS.com, watch a video, and post a comment all by leveraging your Facebook account. Users like this because they can access content without the hurdle of creating a new account. Web sites like it because with the burden of sign-up out of the way, Facebook becomes a Web site experimentation lubricant “Oh, I can use my Facebook ID to sign in? Then let me try this out.” After you post that comment to CBS, it can then squirt back into your feed, to be shared with your friends. The kind of viral activity that spreads Facebook apps can now be harnessed to promote activities outside Facebook’s walls! And this movement can be good for Facebook, as it more firmly plants the site in the center of the user’s online universe. Perhaps leveraging Facebook’s identity makes the Web becomes even more social. When I see a Citysearch review from some random user I don’t know, I’m not sure how much I trust it. But when my Facebook buddy uses Citysearch to post a restaurant praise or pan, I might value that opinion more than any others I’ve read.

Facebook Connect can also be thought of as a way to integrate the Web into Facebook in a way that moves beyond what you can do with in-Facebook apps. Says Zuckerberg, “A lot of people saw [the F8 Facebook application platform] and asked ‘Why is Facebook trying to get all these applications inside Facebook when the Web is clearly the platform?’ And we actually agreed with that... As time goes on, we’re shifting away from Platform inside Facebook and shifting more towards Connection (outside of Facebook).” [2] Facebook Connect just might be a way for Facebook to colonize the Web, growing its sphere of influence as partner, hub, traffic driver, and standards setter.



Facebook Connect also allows third-party developers to create all sorts of apps. Facebook feeds are now streaming direct to Samsung, Vizio, and Sony TVs; through the Xbox 360 and Wii game consoles; and via Verizon's FIOS pay TV service. While Facebook might never have the time or resources to create apps that put its service on every gadget on the market, they don't need to. Developers using Facebook Connect will gladly pick up the slack. But there are major challenges with a more open approach, most notably a weakening of strategic assets, revenue sharing, and security. First, weakened assets. Mark Zuckerberg's geeks have worked hard to make their site the top choice for most of the world's social networkers and social network application developers. Right now everyone goes to Facebook because everyone else is on Facebook. As Facebook opens up access to users and content, it risks supporting efforts that undermine the firm's two most compelling sources of competitive advantage: network effects and switching costs. Even with Facebook Connect, it still remains more difficult to export contacts, e-mails, photos, video from Facebook than it does from sites supporting Google's Open Social, [3] and any effort that makes it easier to pack up your "social self" and move it elsewhere risks undermining vital competitive resources advantages. This situation also puts more pressure on Facebook to behave. Lower those switching costs at a time when users are disgusted with firm behavior, and it's not inconceivable that a sizable chunk of the population could bolt for a new rival (to Facebook's credit, the site also reached out to prior critics like MoveOn.org, showing Facebook Connect and soliciting input months before its official release). Along with asset weakening comes the issue of revenue sharing. As mentioned earlier, hosting content, especially photos and rich media, is a very expensive proposition. What incentive does a site have to store data if it will just be sent to a third-party site that will run ads around this content and not share the take? Too much data portability presents



afree rider problem where firms mooch off of Facebook's infrastructure without offering much in return. Consider services like TweetDeck. The free application allows users to access their Facebook feeds and post new status updates—alongside Twitter updates and more—all from one interface. Cool for the user, but bad for Facebook, since each TweetDeck use means Facebook users are “off site,” not looking at ads, and hence not helping Zuckerberg & Co. earn revenue. It's as if the site has encouraged the equivalent of an ad blocker, yet Facebook Connect lets this happen! Finally, consider security. Allowing data streams that contain potentially private posts and photographs to squirt across the Internet and land where you want them raises all sorts of concerns. What's to say an errant line of code doesn't provide a back door to your address book or friends list? To your messaging account? To let others see photos you'd hoped to only share with family? Security breaches can occur on any site, but once the data is allowed to flow freely, every site with access is, for hackers, the equivalent of a potential door to open or a window to crawl through.

Social Networking Goes Global

Facebook will eventually see stellar growth start to slow as the law of large numbers sets in. The shift from growth business to mature one can be painful, and for online firms it can occur relatively quickly. That doesn't mean these firms will become unprofitable, but to sustain growth (particularly important for keeping up the stock price of a publicly traded company), firms often look to expand abroad.

Facebook's crowdsourcing localization effort, where users were asked to look at Facebook phrases and offer translation suggestions for their local language (see [Chapter 6 "Peer Production, Social Media, and Web 2.0"](#)), helped the firm rapidly deploy versions in dozens of markets, blasting the firm past MySpace in global reach. But network effects are both quick and



powerful, and late market entry can doom a business reliant on the positive feedback loop of a growing user base.

And global competition is out there. Worldwide, Facebook wannabes include “Studiverzeichnis” (German for “student index”), a site with millions of users; Vkontakte (“in contact”), Russia’s most popular social networking site; and Xiaonei, Chinese for “in the school,” which is said to have registered 90 percent of China’s college students.

China’s Internet is booming. The nation now has the world’s biggest population of online users and is attracting venture capital to match. Xiaonei, which was sold to Oak Pacific Interactive in 2005, recently received a four-hundred-thirty-million dollar investment to help growth. ^[4] That’s more than any single group has invested in Facebook, and about as much as the firm has raised in total since its founding.

There are over a hundred social networks operating in China. ^[5] And if competing with clones weren’t bad enough, now there are clones of clones. Wang Xing, who created Xiaonei, has created another social networking service, Hainei, aimed at adults rather than students. Cover up the logo on Xiaonei, Hainei, and the Chinese Facebook (which launched in mid-2008), and you’d be hard pressed to tell the difference. China is proving a particularly difficult market for foreign Internet firms. Google, eBay, Yahoo! and MySpace have all struggled there (at one point, Rupert Murdoch even sent his wife, Wendi Deng Murdoch, to head up the MySpace China effort). And don’t be surprised to see some of these well-capitalized overseas innovators making a move on U.S. markets too.

While global growth can seem like a good thing, acquiring global users isn’t the same as making money from them. Free sites with large amounts of users from developing nations face real cost/revenue challenges. As the *New York Times* points out, there are 1.6 billion Internet users worldwide, but fewer than half of them have disposable incomes high

enough to interest major advertisers. [6] Worse still, telecommunications costs in these markets are also often higher, too. Bandwidth costs and dim revenue options caused video site Veoh to block access coming from Africa, Eastern Europe, Latin America, and some parts of Asia. MySpace already offers a stripped-down Lite option as its default in India. And execs at YouTube and Facebook haven't ruled out lowering the quality of streaming media, file size, or other options, discriminating by region or even by user. Making money in the face of this so-called "International Paradox" requires an awareness of "fast and cheap" tech trends highlighted in [Chapter 4 "Moore's Law and More: Fast, Cheap Computing and What It Means for the Manager"](#), as well as an ability to make accurate predictions regarding regional macroeconomic trends. Ignore a market that's unprofitable today and a rival could swoop in and establish network effects and other assets that are unbeatable tomorrow. But move too early and losses could drag you down.

KEY TAKEAWAYS

- Walled gardens attempt to keep users within limited experiences that can be controlled and exploited by service sponsors.
- Open Social is Google's attempt to break down walled gardens among sites and encourage intersite sharing. From Google's perspective, a successful Open Social would create a rival application development platform and a set of complementary services to combat Facebook. These platforms also represent areas where Google has a greater likelihood of competing for ad revenue.
- Facebook is offering its own intersite sharing platform, Facebook Connect. Features include support for leveraging Facebook User IDs for verification on other sites, integration of feed and other content across

sites, and allowing Facebook data to be leveraged by third-party applications.

- These efforts come with risks, including the potential of weaker network effects and switching costs for firms that had previously run “walled garden” or stand alone sites. Free riders that exploit the firms content without compensation. And the potential for privacy and security risks.
- Global growth is highly appealing to firms, but expensive bandwidth costs and low prospects for ad revenue create challenges akin to the free rider problem.

QUESTIONS AND EXERCISES

1. What is OpenSocial? What challenges does it face in attempting to become a dominant standard?
2. Should Facebook offer Connect? In what ways might this benefit the firm? In what ways is it a risk?
3. How did Facebook help limit activist criticisms of Facebook Connect? Was this a wise move? What are the risks in using this approach?
4. What is TweetDeck? Why is this a potential threat to Facebook for products like this?
5. Cite effective examples you’ve seen of Facebook Connect in action (or if you haven’t seen this, do some background research to uncover such efforts). Why do the efforts you’ve highlighted “work?” How do they benefit various parties? Does everyone benefit? Is anyone at risk and if so, explain the risks?
6. Facebook and MySpace now have global competitors. What determines the success of a social network within a given country? Why do network effects for social networks often fail to translate across national borders?
7. How did Facebook localize its site so quickly for various different regions of the world?



8. What factors encourage firms to grow an international user base as quickly as possible? Why is this a risk? What sorts of firms are at more risk than others?

7.10 Is Facebook Worth It?

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Question the fifteen-billion-dollar valuation so often cited by the media.
2. Understand why Microsoft might be willing to offer to invest in Facebook at a higher valuation rate.

It has often been said that the first phase of the Internet was about putting information online and giving people a way to find it. The second phase of

the Web is about connecting people with one another. The Web 2.0 movement is big and impactful, but is there much money in it? While the valuations of private firms are notoriously difficult to pin down due to a lack of financial disclosure, the often-cited fifteen-billion-dollar valuation from the fall of 2007 Microsoft investment is rich, even when made by such a deep-pocketed firm. Using estimates at the time of the deal, if Facebook were a publicly traded company, it would have a price to earnings ratio of five hundred; Google's at the time was fifty-three, and the average for the S&P 500 is historically around fifteen. But it's not as simple as a raw valuation. The deal was also done in conjunction with an agreement to let Microsoft manage the sale of Facebook's banner ads worldwide. And Microsoft's investment was done on the basis of preferred stock, granting the firm benefits beyond common stock, such as preference in terms of asset liquidation. ^[1] Both of these are reasons a firm would be willing to "pay more" to get in on a deal. Another argument can be made for Microsoft purposely inflating the value of Facebook in order to put it out of reach from most rival bidders. A fat valuation by Microsoft and a deal locking up ad rights makes the firm seem more expensive, less attractive, and out of reach for perhaps all but the richest and most committed suitors. Google may be the only firm with that could possibly launch a credible bid, and Zuckerberg is reported to be genuinely uninterested in being absorbed by the Search Sovereign. ^[2] Since the fall of 2007, several others have invested private money into Facebook as well, including the Founders Fund and Li Kai Shing, the Hong Kong billionaire behind Hutchison Whampoa. Press reports and court documents suggest that these deals were done at valuations that were lower than what Microsoft accepted. In May 2009 Russian firm Digital Sky paid two hundred million dollars for 1.96 percent of the firm, a ten-billion-dollar valuation (also in preferred stock). That's a one third haircut off the



Microsoft price, albeit without the Redmond-specific strategic benefits of the investment.^[3]

So despite the headlines, even at the time of the Microsoft investment, Facebook was almost certainly not valued at a pure fifteen billion dollars. This isn't to say definitively that Facebook won't be worth fifteen billion dollars (or more) someday, but even a valuation at "just" ten billion dollars is a lot to pay for a profitless firm with estimated 2009 revenues of five hundred million dollars. Of course, raising more capital enables Zuckerberg to go on the hunt as well. Facebook investor Peter Theil confirmed the firm had already made an offer to buy Twitter, a firm which at the time had zero dollars in revenues and no discernable business model, for a cool half billion dollars.^[4]

Much remains to be proven for any valuation to hold. Facebook is new. Its models are evolving, and it has quite a bit to prove. Consider efforts to try to leverage friend networks. According to Facebook's own research, "an average Facebook user with 500 friends actively follows the news on only forty of them, communicates with twenty, and keeps in close touch with about ten. Those with smaller networks follow even fewer."^[5] That might not be enough critical mass to offer real, differentiable impact. Also, the advantages of leveraging the friend network hinge on increased sharing and trust, a challenge for a firm that has had so many high-profile privacy stumbles. There is promise. Profiling firm RapLeaf found that targeting based on actions within a friend network can increase click-through rates threefold—that's an advantage advertisers are willing to pay for. But Facebook is still far from proving it can comfortably achieve the promise of these initial efforts.

Steve Rubel wrote the following on his Micro Persuasion blog: "The Internet amber is littered with fossilized communities that once dominated. These former stalwarts include AOL, Angelfire, theGlobe.com, GeoCities,

and Tripod.” Network effects and switching cost advantages can be strong, but not necessarily insurmountable if value is seen elsewhere and if an effort becomes more fad than must have. Time will tell if Facebook’s competitive assets and constant innovation are enough to help it avoid the fate of those that have gone before them.

KEY TAKEAWAYS

- Not all investments are created equal, and a simple calculation of investment dollars multiplied by the percentage of firm owned does not tell the whole story.
- Microsoft’s investment entitled the firm to preferred shares; it also came with advertising deal exclusivity.
- Microsoft may also benefit from offering higher valuations that discourage rivals from making acquisition bids for Facebook.
- Facebook has continued to invest capital raised in expansion, particularly in hardware and infrastructure. It has also pursued its own acquisitions, most notably it made a failed bid to acquire Twitter.
- The firm’s success will hinge on its ability to create sustainably profitable revenue opportunities. It has yet to prove that data from the friend network will be large enough and can be used in a way that is differentially attractive to advertisers. However, some experiments in profiling and ad targeting across a friend network have shown very promising results. Firms exploiting these opportunities will need to have a deft hand in offering consumer and firm value while quelling privacy concerns.

QUESTIONS AND EXERCISES

1. Who else might want to acquire Facebook? Is it worth it at current valuation rates?

2. Circumstances change over time. Research the current state of Facebook’s financials—how much is the firm “valued at?” How much revenue does it bring in? How profitable is it? Are these figures easy or difficult to find? Why or why not?
3. What motivation does Microsoft have in bidding so much for Facebook?
4. Do you think Facebook was wise to take funds from Digital Sky? Why or why not?
5. Do you think Facebook’s friend network is large enough to be leveraged as a source of revenue in ways that are notably different than conventional pay-per-click or CPM-based advertising? Would you be excited about certain possibilities? Creeped out by some? Explain possible scenarios that might work or might fail and justify your positions.
6. So you’ve had a chance to learn about Facebook, its model, growth, outlook, strategic assets, and competitive environment. How much do you think the firm is worth? Which firms do you think it should compare with in terms of value, influence, and impact? Would you invest in Facebook?
7. Which firms might make good merger partners with Facebook? Would these deals ever go through? Why or why not?



Chapter 8

Google: Search, Online Advertising, and Beyond...

8.1 Introduction

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand the extent of Google’s rapid rise and its size and influence when compared with others in the media industry.
2. Recognize the shift away from traditional advertising media to Internet advertising.
3. Gain insight into the uniqueness and appeal of Google’s corporate culture.

Google has been called a one-trick pony, ^[1] but as tricks go, it’s got an exquisite one. Google’s “trick” is matchmaking—pairing Internet surfers with advertisers and taking a cut along the way. This cut is substantial—over twenty-one billion dollars in 2008. In fact, as *Wired*’s Steve Levy puts it, Google’s matchmaking capabilities may represent “the most successful business idea in history.” ^[2] For perspective, consider that as a ten-year-old firm, and one that had been public for less than five years, Google had



already grown to earn more annual advertising dollars than *any* U.S. media company. No television network, no magazine group, no newspaper chain brings in more ad bucks than Google. And none is more profitable. While Google's stated mission is "to organize the world's information and make it universally accessible and useful," advertising drives profits and lets the firm offer most of its services for free.

As more people spend more time online, advertisers are shifting spending away from old channels to the Internet; and Google is swallowing the lion's share of this funds transfer.^[3] By some estimates Google has 76 percent of the search advertising business.^[4] Add to that Google's lucrative AdSense network that serves ads to sites ranging from small time bloggers to the *New York Times*, plus markets served by Google's acquisition of display ad leader DoubleClick, and the firm controls in the neighborhood of 70 percent of *all* online advertising dollars.^[5] Google has the world's strongest brand^[6] (its name is a verb—*just Google it*). It is regularly voted among the best firms to work for in America (twice topping *Fortune's* list). While rivals continue to innovate (see the box "Search: Google Rules, but It Ain't Over" in [Section 8.10 "The Battle Unfolds"](#)) through Q1 2009, Google's share of the search market has consistently grown while its next two biggest competitors have shrunk.

Wall Street has rewarded this success. The firm's market capitalization (market cap), the value of the firm calculated by multiplying its share price by the number of shares, makes Google the most valuable media company on the planet. By early 2009, Google's market cap was greater than that of News Corp (which includes Fox, MySpace, and the *Wall Street Journal*), Disney (including ABC, ESPN, theme parks, and Pixar), Time Warner (*Fortune*, *Time*, *Sports Illustrated*, CNN, and Warner Bros.), Viacom (MTV, VH1, and Nickelodeon), CBS, and the *New York Times*—combined! Not bad for a business started by two twenty-something

computer science graduate students. By 2007 that duo, Sergey Brin and Larry Page, were billionaires, tying for fifth on the *Forbes* 400 list of wealthiest Americans.

Genius Geeks and Plum Perks

Brin and Page have built a talent magnet. At the Googleplex, the firm's Mountain View, California headquarters, geeks are lavished with perks that include on-site laundry, massage, carwash, bicycle repair, free haircuts, state of the art gyms, and wi-fi equipped shuttles that ferry employees between Silicon Valley and the San Francisco Bay area. The Googleplex is also pretty green. The facility gets 30 percent of its energy from solar cells, representing the largest corporate installation of its kind. [9]

The firm's quirky tech-centric culture is evident everywhere. A T-Rex skeleton looms near the volleyball court. Hanging from the lobby ceiling is a replica of SpaceShipOne, the first commercial space vehicle. And visitors to the bathroom will find "testing on the toilet," coding problems or other brainteasers to keep gray matter humming while seated on one of the firm's \$800 remote-controlled Japanese commodes. Staff also enjoy an A-list lecture series attracting luminaries ranging from celebrities to heads of state.

And of course there's the food—all of it free. The firm's founders felt that no employee should be more than 100 feet away from nourishment, and a tour around Google offices will find espresso bars, snack nooks, and fully stocked beverage refrigerators galore. There are eleven gourmet cafeterias on site, the most famous being "Charlie's Place," first run by the former executive chef for the Grateful Dead.

CEO Eric Schmidt says the goal of all this is "to strip away everything that gets in our employees' way." [10] And the perks, culture, and sense of mission have allowed the firm to assemble one of the most impressive

rosters of technical talent anywhere. The Googleplex is like a well-fed Manhattan project, and employee ranks have included “Father of the Internet” Vint Cerf; Hal Varian, the former Dean of the U.C. Berkeley School of Information Management and Systems; Kai-Fu Lee, the former head of Microsoft Research in China; and Andy Hertzfeld, one of the developers of the original Macintosh user interface.

Engineers find Google a particularly attractive place to work, in part due to a corporate policy of offering “20 percent time,” the ability work the equivalent of one day a week on new projects that interest them. It’s a policy that has fueled innovation. Google Vice President Marissa Mayer (who herself regularly ranks among *Fortune*’s most powerful women in business) has stated that roughly half of Google products got their start in 20 percent time. ^[11]

Studying Google gives us an idea of how quickly technology-fueled market disruptions can happen, and how deeply these disruptions penetrate various industries. We’ll also study the underlying technologies that power search, online advertising, and customer profiling. We’ll explore issues of strategy, privacy, fraud, and discuss other opportunities and challenges the firm faces going forward.

KEY TAKEAWAYS

- Online advertising represents the only advertising category trending with positive growth.
- Google controls the lion’s share of the Internet search advertising business and online advertising dollars. The firm also earns more total advertising revenue than any other firm, online or off.
- Google’s market cap makes it the most valuable media company in the world; it has been rated as having the world’s strongest brand.



QUESTIONS AND EXERCISES

1. List the reasons why Google has been considered a particularly attractive firm to work for. Are all of these associated with perks?
2. Market capitalization and market share change frequently. Investigate Google's current market cap and compare it with other media companies. Do patterns suggested in this case continue to hold? Why or why not?
3. Search industry numbers presented are through year-end 2008, but in mid-year 2009, Microsoft introduced Bing. Research online to find out the most current Google versus Bing versus Yahoo! market share. Does Google's position seem secure to you? Why or why not?

8.2 Understanding Search



LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand the mechanics of search, including how Google indexes the Web and ranks its organic search results.
2. Examine the infrastructure that powers Google and how its scale and complexity offer key competitive advantages.

Before diving into how the firm makes money, let's first understand how Google's core service, search, works.

Perform a search (or query) on Google or another search engine, and the results you'll see are referred to by industry professionals as organic or natural search. Search engines use different algorithms for determining the order of organic search results, but at Google the method is called PageRank (a bit of a play on words, it ranks Web pages, and was initially developed by Google cofounder Larry Page). Google does not accept money for placement of links in organic search results. Instead, PageRank results are a kind of popularity contest. Web pages that have more pages *linking to them* are ranked higher.

The process of improving a page's organic search results is often referred to as search engine optimization (SEO). SEO has become a critical function for many marketing organizations since if a firm's pages aren't near the top of search results, customers may never discover its site.

Google is a bit vague about the specifics of precisely how PageRank has been refined, in part because many have tried to game the system. The less scrupulous have tried creating a series of bogus Web sites, all linking back to the pages they're trying to promote (this is called link fraud, and Google actively works to uncover and shut down such efforts). We do know that links from some Web sites carry more weight than others. For example, links from Web sites that Google deems as "influential," and links from



most “.edu” Web sites, have greater weight in PageRank calculations than links from run-of-the-mill “.com” sites.

Spiders and Bots and Crawlers—Oh My!

When performing a search via Google or another search engine, you’re not actually searching the Web. What really happens is that the major search engines make what amounts to a *copy* of the Web, storing and indexing the text of online documents on their own computers. Google’s index considers over one trillion URLs. ^[1] The upper right-hand corner of a Google query shows you just how fast a search can take place (in the example above, rankings from over eight million results containing the term “Toyota Prius” were delivered in less than two tenths of a second).

To create these massive indexes, search firms use software to crawl the Web and uncover as much information as they can find. This software is referred to by several different names: software robots, spiders, Web crawlers. They all pretty much work the same way. In order to make its Web sites visible, every online firm provides a list of all of the public, named servers on its network, known as Domain Name Service (DNS) listings. For example, Yahoo! has different servers that can be found at <http://www.yahoo.com>, sports.yahoo.com, weather.yahoo.com, finance.yahoo.com, et cetera. Spiders start at the first page on every public server and follow every available link, traversing a Web site until all pages are uncovered.

Google will crawl frequently updated sites, like those run by news organizations, as often as several times an hour. Rarely updated, less popular sites might only be reindexed every few days. The method used to crawl the Web also means that if a Web site isn’t the first page on a public server, or isn’t linked to from another public page, then it’ll never be



found. [2] Also note that each search engine also offers a page where you can submit your Web site for indexing.

While search engines show you what they've found on their *copy* of the Web's contents; clicking a search result will direct you to the actual Web site, not the copy. But sometimes you'll click a result only to find that the Web site doesn't match what the search engine found. This happens if a Web site was updated before your search engine had a chance to reindex the changes. In most cases you can still pull up the search engine's copy of the page. Just click the "Cached" link below the result (the term cache refers to a temporary storage space used to speed computing tasks).

But what if you want the content on your Web site to remain off limits to search engine indexing and caching? Organizations have created a set of standards to stop the spider crawl, and all commercial search engines have agreed to respect these standards. One way is to put a line of *HTML code* invisibly embedded in a Web site that tells all software robots to stop indexing a page, stop following links on the page, or stop offering old page archives in a cache. Users don't see this code, but commercial Web crawlers do. For those familiar with HTML code (the language used to describe a Web site), the command to stop Web crawlers from indexing a page, following links, and listing archives of cached pages looks like this:

```
<META NAME="ROBOTS" CONTENT="NOINDEX, NOFOLLOW, NOARCHIVE">
```

There are other techniques to keep the spiders out, too. Web site administrators can add a special file (called robots.txt) that provides similar instructions on how indexing software should treat the Web site. And a lot of content lies inside the "*dark Web*," either behind corporate firewalls or inaccessible to those without a user account—think of private Facebook updates no one can see unless they're your friend—all of that is out of Google's reach.



What's It Take to Run This Thing?

Sergey Brin and Larry Page started Google with just four scavenged computers. [3] But in a decade, the infrastructure used to power the search sovereign has ballooned to the point where it is now the largest of its kind in the world. [4] Google doesn't disclose the number of servers it uses, but by some estimates, it runs over 1.4 million servers in over a dozen so-called server farms worldwide. [5] In 2008, the firm spent \$2.18 billion on capital expenditures, with data centers, servers, and networking equipment eating up the bulk of this cost. [6] Building massive server farms to index the ever-growing Web is now the cost of admission for any firm wanting to compete in the search market. This is clearly no longer a game for two graduate students working out of a garage.

The size of this investment not only creates a barrier to entry, it influences industry profitability, with market-leader Google enjoying huge economies of scale. Firms may spend the same amount to build server farms, but if Google has two-thirds of this market (and growing) while Microsoft's search draws only about one-tenth the traffic, which do you think enjoys the better return on investment?

The hardware components that power Google aren't particularly special. In most cases the firm uses the kind of Intel or AMD processors, low-end hard drives, and RAM chips that you'd find in a desktop PC. These components are housed in rack-mounted servers about 3.5 inches thick, with each server containing two processors, eight memory slots, and two hard drives.

In some cases, Google mounts racks of these servers inside standard-sized shipping containers, each with as many as 1,160 servers per box. [7] A given data center may have dozens of these server-filled containers all linked together. Redundancy is the name of the game. Google assumes individual components will regularly fail, but no single failure should interrupt the



firm's operations (making the setup what geeks call fault-tolerant). If something breaks, a technician can easily swap it out with a replacement. Each server farm layout has also been carefully designed with an emphasis on lowering power consumption and cooling requirements. And the firm's custom software (much of it built upon open-source products) allows all this equipment to operate as the world's largest grid computer.

Web search is a task particularly well suited for the massively parallel architecture used by Google and its rivals. For an analogy of how this works, imagine that working alone, you need try to find a particular phrase in a hundred-page document (that's a one server effort). Next, imagine that you can distribute the task across five thousand people, giving each of them a separate sentence to scan (that's the multi-server grid). This difference gives you a sense of how search firms use massive numbers of servers and the divide-and-conquer approach of grid computing to quickly find the needles you're searching for within the Web's haystack (for more on grid computing, see [Chapter 4 "Moore's Law and More: Fast, Cheap Computing and What It Means for the Manager"](#), and for more on server farms, see [Chapter 10 "Software in Flux: Partly Cloudy and Sometimes Free"](#)).

Google will even sell you a bit of its technology so that you can run your own little Google in-house without sharing documents with the rest of the world. Google's line of search appliances are rack-mounted servers that can index documents within a corporation's Web site, even specifying password and security access on a per-document basis. Selling hardware isn't a large business for Google, and other vendors offer similar solutions, but search appliances can be vital tools for law firms, investment banks, and other document-rich organizations.



Trendspotting with Google

Google not only gives you search results, it lets you see aggregate trends in what its users are searching for, and this can yield powerful insights. For example, by tracking search trends for flu symptoms, Google's Flu Trends Web site can pinpoint outbreaks one to two weeks faster than the Centers for Disease Control and Prevention. [8] Want to go beyond the flu? Google's Trends, and Insights for Search services allow anyone to explore search trends, breaking out the analysis by region, category (image, news, product), date, and other criteria. Savvy managers can leverage these and similar tools for competitive analysis, comparing a firm, its brands, and its rivals.

KEY TAKEAWAYS

- Ranked search results are often referred to as organic or natural search. PageRank is Google's algorithm for ranking search results. PageRank orders organic search results based largely on the number of Web sites linking to them, and the "weight" of each page as measured by its "influence."
- Search engine optimization (SEO) is the process of using natural or organic search to increase a Web site's traffic volume and visitor quality. The scope and influence of search has made SEO an increasingly vital marketing function.
- Users don't really search the Web; they search an archived copy built by crawling and indexing discoverable documents.
- Google operates from a massive network of server farms containing hundreds of thousands of servers built from standard, off-the-shelf items. The cost of the operation is a significant barrier to entry for competitors. Google's share of search suggests the firm can realize economies of



scales over rivals required to make similar investments while delivering fewer results (and hence ads).

- Web site owners can hide pages from popular search-engine Web crawlers using a number of methods, including HTML tags, a no-index file, or ensuring that Web sites aren't linked to other pages and haven't been submitted to Web sites for indexing.

QUESTIONS AND EXERCISES

1. How do search engines discover pages on the Internet? What kind of capital commitment is necessary to go about doing this? How does this impact competitive dynamics in the industry?
2. How does Google rank search results? Investigate and list some methods that an organization might use to improve its rank in Google's organic search results. Are there techniques Google might not improve of? What risk does a firm run if Google or another search firm determines that it has used unscrupulous SEO techniques to try to game ranking algorithms?
3. Sometimes Web sites returned by major search engines don't contain the words or phrases that initially brought you to the site. Why might this happen?
4. What's a cache? What other products or services have a cache?
5. What can be done if you want the content on your Web site to remain off limits to search engine indexing and caching?
6. What is a "search appliance?" Why might an organization choose such a product?
7. Become a better searcher: Look at the advanced options for your favorite search engine. Are there options you hadn't used previously? Be prepared to share what you learn during class discussion.



8. Visit Google Trends and Google Insights for Search. Explore the tool as if you were comparing a firm with its competitors. What sorts of useful insights can you uncover? How might businesses use these tools?

8.3 Understanding the Increase in Online Ad Spending

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand how media consumption habits are shifting.
2. Be able to explain the factors behind the growth and appeal of online advertising.

For several years, Internet advertising has been the only major media ad category to show significant growth. There are three factors driving online ad growth trends: (1) increased user time online, (2) improved measurement and accountability, and (3) targeting.

American teenagers (as well as the average British, Australian, and New Zealander Web surfer) now spend more time on the Internet than watching television. ^[1] They're reading fewer print publications, and radio listening among the iPod generation is down 30 percent. ^[2] So advertisers are simply



following the market. Online channels also provide advertisers with a way to reach consumers at work—something that was previously much more difficult to do.

Many advertisers have also been frustrated by how difficult it's been to gauge the effectiveness of traditional ad channels such as TV, print, and radio. This frustration is reflected in the old industry saying, "I know that half of my advertising is working—I just don't know which half." Well, with the Internet, now you know. While measurement technologies aren't perfect, advertisers can now count ad impressions (each time an ad appears on a Web site), whether or not a user clicks on an ad, and the product purchases or other Web site activity that comes from those clicks. [3] And as we'll see, many online ad payment schemes are directly linked to ad performance.

Various technologies and techniques also make it easier for firms to target users based on how likely a person is to respond to an ad. In theory a firm can use targeting to spend marketing dollars only on those users deemed to be its best prospects. Let's look at a few of these approaches in action.

KEY TAKEAWAYS

- There are three reasons driving online ad growth trends: (1) increasing user time online, (2) improved measurement and accountability, and (3) targeting.
- Digital media is decreasing time spent through traditional media consumption channels (e.g., radio, TV, newspapers), potentially lowering the audience reach of these old channels and making them less attractive for advertisers.
- Measurement techniques allow advertisers to track the performance of their ads—indicating things such as how often an ad is displayed, how often an is clicked, where an ad was displayed when it was clicked, and

more. Measurement metrics can be linked to payment schemes, improving return on investment (ROI) and accountability compared to many types of conventional advertising.

- Advertising ROI can be improved through targeting. Targeting allows a firm to serve ads to specific categories of users, so firms can send ads to groups it is most interested in reaching, and those that are most likely to respond to an effort.

QUESTIONS AND EXERCISES

1. How does your media time differ from your parents? Does it differ among your older or younger siblings, or other relatives? Which media are you spending more time with? Less time with?
2. Put yourself in the role of a traditional media firm that is seeing its market decline. What might you do to address decline concerns? Have these techniques been attempted by other firms? Do you think they've worked well? Why or why not?
3. Put yourself in the role of an advertiser for a product or service that you're interested in. Is the Internet an attractive channel for you? How might you use the Internet to reach customers you are most interested in? Where might you run ads? Who might you target? Who might you avoid? How might the approach you use differ from traditional campaigns you'd run in print, TV, or radio? How might the size (money spent, attempted audience reach) and timing (length of time run, time between campaigns) of ad campaigns online differ from offline campaigns?
4. List ways in which you or someone you know has been targeted in an Internet ad campaign. Was it successful? How do you feel about targeting?



8.4 Search Advertising

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand Google's search advertising revenue model.
2. Know the factors that determine the display and ranking of advertisements appearing on Google's search results pages.
3. Be able to describe the uses and technologies behind geotargeting.

The practice of running and optimizing search engine ad campaigns is referred to as search engine marketing (SEM). ^[1] SEM is a hot topic in an increasingly influential field, so it's worth spending some time learning how search advertising works on the Internet's largest search engine.

Roughly two-thirds of Google's revenues come from ads served on its own sites, and the vast majority of this revenue comes from search engine ads. ^[2] During Google's early years, the firm actually resisted making money through ads. In fact, while at Stanford, Brin and Page even coauthored a paper titled "The Evils of Advertising." ^[3] But when Yahoo! and others balked at buying Google's search technology (offered for as little as \$500,000), Google needed to explore additional revenue streams. It wasn't until two years after incorporation that Google ran ads alongside organic search results. That first ad, one for "Live Mail Order Lobsters," appeared just minutes after the firm posted a link reading "See Your Ad Here"). ^[4] Google has only recently begun incorporating video and image ads into search. For the most part, the ads you'll see to the right (and sometimes



top) of Google’s organic search results appear as keyword advertising, meaning they’re targeted based on a user’s query. Advertisers bid on the keywords and phrases that they’d like to use to trigger the display of their ad. Linking ads to search was a brilliant move, since the user’s search term indicates an overt interest in a given topic. Want to sell hotel stays in Tahiti? Link your ads to the search term “Tahiti Vacation.”

Not only are search ads highly targeted, advertisers only pay for results. Text ads appearing on Google search pages are billed on a pay-per-click (PPC) basis, meaning that advertisers don’t spend a penny unless someone actually clicks on their ad. Note that the term pay-per-click is sometimes used interchangeably with the term cost-per-click (CPC).

Not Entirely Google’s Idea

Google didn’t invent pay-for-performance search advertising. A firm named GoTo.com (later renamed Overture) pioneered pay-per-click ads and bidding systems and held several key patents governing the technology. Overture provided pay-per-click ad services to both Yahoo! and Microsoft, but it failed to refine and match the killer combination of ad auctions and search technology that made Google a star. Yahoo! eventually bought Overture and sued Google for patent infringement. In 2004, the two firms settled, with Google giving Yahoo! 2.7 million shares in exchange for a “fully paid, perpetual license” to over sixty Overture patents. [5]

If an advertiser wants to display an ad on Google search, they can set up a Google AdWords advertising account in minutes, specifying just a single ad, or multiple ad campaigns that trigger different ads for different keywords. Advertisers also specify what they’re willing to pay each time an ad is clicked, how much their overall ad budget is, and they can control additional parameters, such as the timing and duration of an ad campaign.



If no one clicks on an ad, Google doesn't make money, advertisers don't attract customers, and searchers aren't seeing ads they're interested in. So in order to create a winning scenario for everyone, Google has developed a precise ad ranking formula that rewards top performing ads by considering two metrics: the maximum CPC that an advertiser is willing to pay, and the advertisement's quality score—a broad measure of ad performance. Create high quality ads and your advertisements might appear ahead of competition, even if your competitors bid more than you. But if ads perform poorly they'll fall in rankings or even drop from display consideration.

Below is the formula used by Google to determine the rank order of sponsored links appearing on search results pages.

Ad Rank = Maximum CPC × Quality Score

One factor that goes into determining an ad's quality score is the click-through rate (CTR) for the ad, the number of users who clicked an ad divided by the number of times the ad was delivered (the impressions). The CTR measures the percentage of people who clicked on an ad to arrive at a destination site. Also included in a quality score are the overall history of click performance for the keywords linked to the ad, the relevance of an ad's text to the user's query, and Google's automated assessment of the user experience on the landing page—the Web site displayed when a user clicks on the ad. Ads that don't get many clicks, ad descriptions that have nothing to do with query terms, and ads that direct users to generic pages that load slowly or aren't strongly related to the keywords and descriptions used in an ad, will all lower an ad's chance of being displayed. [6]

When an ad is clicked, advertisers don't actually pay their maximum CPC; Google discounts ads to just one cent more than the minimum necessary to maintain an ad's position on the page. So if you bid one dollar per click, but the ad ranked below you bids ninety cents, you'll pay just ninety-one cents

if the ad is clicked. Discounting was a brilliant move. No one wants to get caught excessively overbidding rivals, so discounting helps reduce the possibility of this so-called bidder's remorse. And with this risk minimized, the system actually encouraged higher bids! [7]

Ad ranking and cost-per-click calculations take place as part of an automated auction that occurs *every time* a user conducts a search. Advertisers get a running total of ad performance statistics so that they can monitor the return on their investment and tweak promotional efforts for better results. And this whole system is automated for self-service—all it takes is a credit card, an ad idea, and you're ready to go.

How Much Do Advertisers Pay per Click?

Google rakes in billions on what amounts to pocket change earned one click at a time. Most clicks bring in between thirty cents and one dollar.

However, costs can vary widely depending on industry, current competition, and perceived customer value. [Table 8.1 "10 Most Expensive Industries for Keyword Ads"](#) shows some of the highest reported CPC rates. But remember, any values fluctuate in real time based on auction participants.

Since rates are based on auctions, top rates reflect what the market is willing to bear. As an example, law firms, which bring in big bucks from legal fees, decisions, and settlement payments often justify higher customer acquisition costs. And firms that see results will keep spending. Los Angeles–based Chase Law Group has said that it brings in roughly 60 percent of its clients through Internet advertising. [8]

IP Addresses and Geotargeting

Geotargeting occurs when computer systems identify a user's physical location (sometimes called the *geolocation*) for the purpose of delivering tailored ads or other content. On Google AdWords, for example, advertisers

can specify that their ads only appear for Web surfers located in a particular country, state, metropolitan region, or a given distance around a precise locale. They can even draw a custom ad-targeting region on a map and tell Google to only show ads to users detected inside that space.

Ads in Google Search are geotargeted based on IP address. Every device connected to the Internet has a unique IP address assigned by the organization connecting the device to the network. Normally you don't see your IP address (a set of four numbers, from 0 to 255, separated by periods; e.g., 136.167.2.220). But the range of IP addresses "owned" by major organizations and Internet service providers (ISPs) is public knowledge. In many cases it's possible to make an accurate guess as to where a computer, laptop, or mobile phone is located simply by cross-referencing a device's current IP address with this public list.

For example, it's known that all devices connected to the Boston College network contain IP addresses starting with the numbers 136.167. If a search engine detects a query coming from an IP address that begins with those two numbers, it can be fairly certain that the person using that device is in the greater Boston area.

IP addresses will change depending on how and where you connect to the Internet. Connect your laptop to a hotel's wi-fi when visiting a new city, and you're likely to see ads specific to that location. That's because your Internet service provider has changed, and the firm serving your ads has detected that you are using an IP address known to be associated with your new location.

Geotargeting via IP address is fairly accurate, but it's not perfect. For example, some Internet service providers may provide imprecise or inaccurate information on the location of their networks. Others might be so vague that it's difficult to make a best guess at the geography behind a set of numbers (values assigned by a multinational corporation with many



locations, for example). And there are other ways locations are hidden, such as when Internet users connect to proxy servers, third-party computers that pass traffic to and from a specific address without revealing the address of the connected users.

What's My IP Address?

While every operating system has a control panel or command that you can use to find your current IP address, there are also several Web sites that will quickly return this value (and a best guess at your current location). One such site is <http://ip-address.com> (note the spelling has only one “d”). Visit this or a similar site with a desktop, laptop, and mobile phone. Do the results differ and are they accurate? Why?

Geotargeting Evolves Beyond the IP Address

There are several other methods of geotargeting. Firms like Skyhook Wireless can identify a location based on its own map of wi-fi hotspots and nearby cell towers. Many mobile devices come equipped with global positioning system (GPS) chips (identifying location via the GPS satellite network). And if a user provides location values such as a home address or zip code to a Web site, then that value might be stored and used again to make a future guess at a user's location.

KEY TAKEAWAYS

- Roughly two-thirds of Google's revenues come from ads served on its own sites, and the vast majority of this revenue comes from search engine ads.
- Advertisers choose and bid on the keywords and phrases that they'd like to use to trigger the display of their ad.



- Advertisers pay for cost-per-click advertising only if an ad is clicked on. Google makes no money on CPC ads that are displayed but not clicked.
- Google determines ad rank by multiplying CPC by Quality Score. Ads with low ranks might not display at all.
- Advertisers usually don't pay their maximum CPC. Instead, Google discounts ads to just one cent more than the minimum necessary to maintain an ad's position on the page—a practice that encourages higher bids.
- Geotargeting occurs when computer systems identify a user's physical location (sometimes called geolocation) for the purpose of delivering tailored ads or other content.
- Google uses IP addresses to target ads.

QUESTIONS AND EXERCISES

1. Which firm invented pay-per-click advertising? Why does Google dominate today and not this firm?
2. How are ads sold via Google search superior to conventional advertising media such as TV, radio, billboard, print, and yellow pages? Consider factors like the available inventory of space to run ads, the cost to run ads, the cost to acquire new advertisers, and the appeal among advertisers.
3. Are there certain kinds of advertising campaigns and goals where search advertising wouldn't be a good fit? Give examples and explain why.
4. Can a firm buy a top ad ranking? Why or why not?
5. List the four factors that determine an ad's quality score.
6. How much do firms typically pay for a single click?
7. Sites like SpyFu.com and KeywordSpy.com provide a list of the keywords with the highest cost per click. Visit the Top Lists page at SpyFu, KeywordSpy, or a comparable site, to find estimates of the current



highest paying cost per click. Which key words pay the most? Why do you think firms are willing to spend so much?

8. What is bidder's remorse? How does Google's ad discounting impact this phenomenon?
9. Visit <http://www.ip-adress.com/using> a desktop, laptop, and mobile phone (work with a classmate or friend if you don't have access to one of these devices). How do results differ? Why? Are they accurate? What factors go into determining the accuracy of IP-based geolocation?
10. List and briefly describe other methods of geotargeting besides IP address, and indicate the situations and devices where these methods would be more and less effective.
11. The field of search engine marketing (SEM) is relatively new and rising in importance. And since the field is so new and constantly changing, there are plenty of opportunities for young, knowledgeable professionals. Which organizations, professional certification, and other resources are available to SEM professionals? Spend some time searching for these resources online and be prepared to share your findings with your class.

8.5 Ad Networks—Distribution beyond Search

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand ad networks, and how ads are distributed and served based on Web site content.
2. Recognize how ad networks provide advertiser reach and support niche content providers.
3. Be aware of content adjacency problems and their implications.
4. Know the strategic factors behind ad network appeal and success.



Google runs ads not just in search, but also across a host of Google-owned sites like Gmail, Google News, and Blogger. It will even tailor ads for its map products and for mobile devices. But about 30 percent of Google's revenues come from running ads on Web sites that the firm doesn't even own. [1]

Next time you're surfing online, look around the different Web sites that you visit and see how many sport boxes labeled "Ads by Google." Those Web sites are participating in Google's AdSense ad network, which means they're running ads for Google in exchange for a cut of the take.

Participants range from small-time bloggers to some of the world's most highly trafficked sites. Google lines up the advertisers, provides the targeting technology, serves the ads, and handles advertiser payment collection. To participate, content providers just sign up online, put a bit of Google-supplied HTML code on their pages, and wait for Google to send them cash (Web sites typically get about seventy to eighty cents for every AdSense dollar that Google collects). [2]

Google originally developed AdSense to target ads based on keywords automatically detected inside the content of a Web site. A blog post on your favorite sports team, for example, might be accompanied by ads from ticket sellers or sports memorabilia vendors. AdSense and similar online ad networks provide advertisers with access to the long tail of niche Web sites by offering both increased opportunities for ad exposure as well as more-refined targeting opportunities.

Running ads on your Web site is by no means a guaranteed path to profits. The Internet graveyard is full of firms that thought they'd be able to sustain their businesses on ads alone. But for many Web sites, ad networks can be like oxygen, sustaining them with revenue opportunities they'd never be able to achieve on their own.

For example, AdSense provided early revenue for the popular social news site Digg, as well as the multimillion-dollar TechCrunch media empire. It supports Disaboom, a site run by physician and quadriplegic Dr. Glen House. And it continues to be the primary revenue generator for AskTheBuilder.com. That site's founder, former builder Tim Carter, had been writing a handyman's column syndicated to some thirty newspapers. The newspaper columns didn't bring in enough to pay the bills, but with AdSense he hit pay dirt, pulling in over \$350,000 in ad revenue in just his first year! [3]

BEWARE THE CONTENT ADJACENCY PROBLEM

Contextual advertising based on keywords is lucrative, but like all technology solutions it has its limitations. Vendors sometimes suffer from content adjacency problems when ads appear alongside text they'd prefer to avoid. In one particularly embarrassing example, a *New York Post* article detailed a gruesome murder where hacked up body parts were stowed in suitcases. The online version of the article included contextual advertising and was accompanied by...luggage ads. [4]

To combat embarrassment, ad networks provide opportunities for both advertisers and content providers to screen out potentially undesirable pairings based on factors like vendor, Web site, keyword, and category. Ad networks also refine ad-placement software based on feedback from prior incidents (for more on content adjacency problems, see [Chapter 7 "Facebook: Building a Business from the Social Graph"](#)).

Google launched AdSense in 2003, but Google is by no means the only company to run an ad network, nor was it the first to come up with the idea. Rivals include the Yahoo! Publisher Network, Microsoft's adCenter, and AOL's Platform-A. Others, like Quigo, don't even have a consumer Web site

yet manage to consolidate enough advertisers to attract high-traffic content providers such as ESPN, *Forbes*, Fox, and *USA Today*. Advertisers also aren't limited to choosing just one ad network. In fact, many content provider Web sites will serve ads from several ad networks (as well as exclusive space sold by their own sales force), oftentimes mixing several different offerings on the same page.

Ad Networks and Competitive Advantage

While advertisers can use multiple ad networks, there are several key strategic factors driving the industry. For Google, its ad network is a distribution play. The ability to reach more potential customers across more Web sites attracts more advertisers to Google. And content providers (the Web sites that distribute these ads) want there to be as many advertisers as possible in the ad networks that they join, since this should increase the price of advertising, the number of ads served, and the accuracy of user targeting. If advertisers attract content providers, which in turn attract more advertisers, then we've just described network effects! More participants bringing in more revenue also help the firm benefit from scale economies—offering a better return-on-investment from its ad technology and infrastructure. No wonder Google's been on such a tear—the firm's loaded with assets for competitive advantage!

Google's Ad Reach Gets Bigger

While Google has the largest network specializing in distributing text ads, it had been a laggard in graphical display ads (sometimes called image ads). That changed in 2008, with the firm's \$3.1 billion acquisition of display ad network and targeting company DoubleClick. Now in terms of the number of users reached, Google controls both the largest text ad network and the largest display ad network. [5]



KEY TAKEAWAYS

- Google also serves ads through non-Google partner sites that join its ad network. These partners distribute ads for Google in exchange for a percentage of the take.
- AdSense ads are targeted based on keywords that Google detects inside the content of a Web site.
- AdSense and similar online ad networks provide advertisers with access to the long tail of niche Web sites.
- Ad networks handle advertiser recruitment, ad serving, and revenue collection, opening up revenue earning possibilities to even the smallest publishers.

QUESTIONS AND EXERCISES

1. On a percentage basis, how important is AdSense to Google's revenues?
2. Why do ad networks appeal to advertisers? What do they appeal to content providers? What functions are assumed by the firm overseeing the ad network?
3. What factors determine the appeal of an ad network to advertisers and content providers? Which of these factors are potentially sources of competitive advantage?
4. Do dominant ad networks enjoy strong network effects? Are there also strong network effects that drive consumers to search? Why or why not?
5. How difficult is it for a Web site to join an ad network? What does this imply about ad network switching costs? Does it have to exclusively choose one network over another? Does ad network membership prevent a firm from selling its own online advertising, too?
6. What is the content adjacency problem? Why does it occur? What classifications of Web sites might be particularly susceptible to the content adjacency problem? What can ad networks do to minimize the



content adjacency problem for their advertisers? Are there tradeoffs to these approaches?

8.6 More Ad Formats and Payment Schemes

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Know the different formats and media types that Web ads can be displayed in.
2. Know the different ways ads are sold.
3. Know that games can be an ad channel under the correct conditions.

Online ads aren't just about text ads billed in CPC. Ads running through Google AdSense, through its DoubleClick subsidiary, or on most competitor networks can be displayed in several formats and media types, and can be billed in different ways. The specific ad formats supported depend on the ad network but can include the following: image (or display) ads (such as horizontally oriented banners, smaller rectangular buttons, and vertically oriented "skyscraper" ads); rich media ads (which can include animation or video); and interstitials (ads that run before a user arrives at a Web site's contents). The industry trade group, the Internet Advertising Bureau (IAB) sets common standards for display ads so that a single creative (the design and content of the advertisement) can run unmodified across multiple ad networks and Web sites. ^[1] And there are lots of other ways ads are sold besides cost-per-click. Most graphical display ads are sold according to the number of times the ad



appears (the *impression*). Ad rates are quoted in CPM, meaning cost per thousand impressions (the M representing the roman numerical for one thousand). Display ads sold on a CPM basis are often used as part of branding campaigns targeted more at creating awareness than generating click-throughs. Such techniques often work best for promoting products like soft drinks, toothpaste, or movies.

Cost-per-action (CPA) ads pay whenever a user performs a specified action such as signing up for a service, requesting material, or making a purchase. Affiliate programs are a form of cost-per-action, where Web sites share a percentage of revenue earned when a user click-throughs and buys something from their Web site. Amazon runs the world's largest affiliate program, and referring sites can earn 4 percent to 15 percent of sales generated from these click-throughs. Purists might not consider affiliate programs as advertising (rather than text or banner ads, Amazon's affiliates offer links and product descriptions that point back to Amazon's Web site), but these programs can be important tools in a firm's promotional arsenal. And rather than buying targeted ads, a firm might sometimes opt to become an exclusive advertiser on a site. For example, a firm could buy access to all ads served on a site's main page; it could secure exclusive access to a region of the page (such as the topmost banner ad); or it may pay to sponsor a particular portion or activity on a Web site (say a parenting forum, or a "click-to-print" button). Such deals can be billed based on a flat rate, CPM, CPC, or any combination of metrics.

Ads in Games?

As consumers spend more time in video games, it's only natural that these products become ad channels, too. Finding a sensitive mix that introduces ads without eroding the game experience can be a challenge. Advertising can work in racing or other sports games (in 2008 the Obama campaign



famously ran virtual billboards in EA's *Burnout Paradise*), but ads make less sense for games set in the past, future, or on other worlds. Branding ads often work best, since click-throughs are typically not something you want disrupting your gaming experience.

Advertisers have also explored sponsorships of Web-based and mobile games. Sponsorships often work best with casual games, such as those offered on Yahoo! Games or EA's *Pogo*. Firms have also created online mini games (so-called *advergames*) for longer term, immersive brand engagement (e.g., Mini Cooper's *Slide Parking* and Stride Gum's *Chew Challenge*). Others have tried a sort of virtual product placement integrated into experiences. A version of *The Sims*, for example, included virtual replicas of real-world inventory from IKEA and H&M.

In-game ad-serving technology also lacks the widely accepted standards of Web-based ads, so it's unlikely that ads designed for a Wii sports game could translate into a PS3 first-person shooter. Also, one of the largest in-game ad networks, Massive, is owned by Microsoft. That's good if you want to run ads on Xbox, but Microsoft isn't exactly a firm that Nintendo or Sony want to play nice with.

In-game advertising shows promise, but the medium is considerably more complicated than conventional Web site ads. That complexity lowers relative ROI and will likely continue to constrain growth.

KEY TAKEAWAYS

- Web ad formats include, but are not limited to, the following: *image* (or *display*) *ads* (such as horizontally oriented *banners*, smaller rectangular *buttons*, and vertically oriented *skyscraper ads*), *rich media ads* (which can include animation or video), and *interstitials* (ads that run before a user arrives at a Web site's contents).

- In addition to cost-per-click, ads can be sold based on the number of times the ad appears (impressions), whenever a user performs a specified action such as signing up for a service, requesting material, or making a purchase (cost-per-action), or on an exclusive basis which may be billed at a flat rate.
- In-game advertising shows promise, with successful branding campaigns run as part of sports games, through in-game product placement, or via sponsorship of casual games, or in brand-focused advergaming.
- A lack of standards, concerns regarding compatibility with gameplay, and the cost of developing and distributing games are all stifling the growth of in-game ads.

QUESTIONS AND EXERCISES

1. What is the IAB and why is it necessary?
2. What are the major ad format categories?
3. What's an interstitial? What's a rich media ad? Have you seen these? Do you think they are effective? Why or why not?
4. List four major methods for billing online advertising.
5. Which method is used to bill most graphical advertising? What's the term used for this method and what does it stand for?
6. How many impressions are recorded if a single user is served the same ad one thousand times? How many if one thousand users are served the same ad once?
7. Imagine the two scenarios below. Decide which type of campaign would be best for each: text-based CPC advertising or image ads paid for on a CPM basis). Explain your reasoning.
 - a. Netflix is looking to attract new customers by driving traffic to its Web site and increase online subscriptions.
 - b. Zara has just opened a new clothing store in major retailing area in your town. The company doesn't offer online sales; rather, the majority of its sales come from stores.



Which firm runs the world's largest affiliate program? Why is this form of advertising particularly advantageous to the firm (think about the ROI for this sort of effort)?

Given examples where in-game advertising might work and those where it might be less desirable. List key reasons why in-game advertising has not been as successful as other forms of Internet-distributed ads.

8.7 Customer Profiling and Behavioral Targeting

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Be familiar with various tracking technologies and how they are used for customer profiling and ad targeting.



2. Understand why customer profiling is both valuable and controversial.
3. Recognize steps that organizations can take to help ease consumer and governmental concerns.

Advertisers are willing to pay more for ads that have a greater chance of reaching their target audience, and online firms have a number of targeting tools at their disposal. Much of this targeting occurs whenever you visit a Web site, where a behind-the-scenes software dialog takes place between Web browser and Web server that can reveal a number of pieces of information, including IP address, the type of browser used, the computer type, its operating system, and unique identifiers, called cookies.

And remember, *any* server that serves you content can leverage these profiling technologies. You might be profiled not just by the Web site that you're visiting (e.g., nytimes.com), but also by any ad networks that serve ads on that site (e.g., Platform-A, DoubleClick, Google AdSense, Microsoft adCenter).

IP addresses are leveraged extensively in customer profiling. An IP address not only helps with geolocation, it can also indicate a browser's employer or university, which can be further matched with information such as firm size or industry. IBM has used IP targeting to tailor its college recruiting banner ads to specific schools, for example, "There Is Life After Boston College, Click Here to See Why." That campaign garnered click-through rates ranging from 5 to 30 percent ^[1] compared to average rates that are currently well below 1 percent for untargeted banner ads. DoubleClick once even served a banner that include a personal message for an executive at then-client Modem Media. The ad, reading "Congratulations on the twins, John Nardone," was served across hundreds of sites, but was only visible from computers on the Modem Media corporate network. ^[2]



The ability to identify a surfer's computer, browser, or operating system can also be used to target tech ads. For example, Google might pitch its Chrome browser to users detected running Internet Explorer, Firefox, or Safari; while Apple could target those "I'm a Mac" ads just to Windows users. But perhaps the greatest degree of personalization and targeting comes from cookies. Visit a Web site for the first time, and in most cases, a behind-the-scenes dialog takes place that goes something like this:

Server: *Have I seen you before?*

Browser: *No.*

Server: *Then take this unique string of numbers and letters (called a cookie). I'll use it to recognize you from now on.*

The cookie is just a line of identifying text assigned and retrieved by a given Web server and stored in your browser. Upon accepting this cookie your browser has been tagged, like an animal. As you surf around the firm's Web site, that cookie can be used to build a profile associated with your activities. If you're on a portal like Yahoo! you might type in your zip code, enter stocks that you'd like to track, and identify the sports teams you'd like to see scores for. The next time you return to the Web site, your browser responds to the server's "*Have I see you before?*" question with the equivalent of "*Yes, you know me,*" and it presents the cookie that the site gave you earlier. The site can then match this cookie against your browsing profile, showing you the weather, stock quotes, sports scores, and other info that it thinks you're interested in.

Cookies are used for lots of purposes. Retail Web sites like Amazon use cookies to pay attention to what you've shopped for and bought, tailoring Web sites to display products that the firm suspects you'll be most interested in. Sites also use cookies to keep track of what you put in an online "shopping cart," so if you quit browsing before making a purchase, these items will reappear the next time you visit. And many Web sites also



use cookies as part of a “remember me” feature, storing user IDs and passwords. Beware this last one! If you check the “remember me” box on a public Web browser, the next person who uses that browser is potentially using *your* cookie, and can log in as you!

An organization can’t read cookies that it did not give you. So *businessweek.com* can’t tell if you’ve also got cookies from *forbes.com*. But you can see all of the cookies in your browser. Take a look and you’ll almost certainly see cookies from dozens of Web sites that you’ve never visited before. These are third-party cookies (sometimes called *tracking cookies*), and they are usually served by ad networks or other customer profiling firms.

By serving cookies in ads shown on partner sites, ad networks can build detailed browsing profiles that include sites visited, specific pages viewed, duration of visit, and the types of ads you’ve seen and responded to. And that surfing might give an advertising network a better guess at demographics like gender, age, marital status, and more. Visit a new parent site and expect to see diaper ads in the future, even when you’re surfing for news or sports scores!

But What If I Don’t Want a Cookie!

If all of this creeps you out, remember that you’re in control. The most popular Web browsers allow you to block all cookies, block just third-party cookies, purge your cookie file, or even ask for your approval before accepting a cookie. Of course, if you block cookies, you block any benefits that come along with them, and some Web site features may require cookies to work properly. Also note that while deleting a cookie breaks a link between your browser and that Web site, if you supply identifying information in the future (say by logging into an old profile), the site might be able to assign your old profile data to the new cookie.



While the Internet offers targeting technologies that go way beyond traditional television, print, and radio offerings, none of these techniques is perfect. Since users are regularly assigned different IP addresses as they connect and disconnect from various physical and wi-fi networks, IP targeting can't reliably identify individual users. Cookies also have their weaknesses. They're assigned by browsers, so if several people use the same browser, all of their Web surfing activity may be mixed into the same cookie profile. Some users might also use different browsers on the same machine, or use different computers. Unless a firm has a way to match up these different cookies with user accounts or other user-identifying information, a site may be working with multiple, incomplete profiles.

KEY TAKEAWAYS

- The communication between Web browser and Web server can identify IP address, the type of browser used, the computer type, its operating system, and can read and assign unique identifiers, called cookies—all of which can be used in customer profiling and ad targeting.
- An IP address not only helps with geolocation; it can also be matched against other databases to identify the organization providing the user with Internet access (such as a firm or university), that organization's industry, size, and related statistics.
- A cookie is a unique line of identifying text, assigned and retrieved by a given Web server and stored in your browser, that can be used to build a profile associated with your Web activities.
- The most popular Web browsers allow you to block all cookies, block just third-party cookies, purge your cookie file, or even ask for your approval before accepting a cookie.



QUESTIONS AND EXERCISES

1. Give examples of how the ability to identify a surfer's computer, browser, or operating system can be used to target tech ads.
2. Describe how IBM targeted ad delivery for its college recruiting efforts. What technologies were used? What was the impact on click-through rates?
3. What is a cookie? How are cookies used? Is a cookie a computer program? Which firms can read the cookies in your Web browser?
4. Does a cookie accurately identify a user? Why or why not?
5. What is the danger of checking the "remember me" box on a public Web browser?
6. What's a tracking cookie? What kinds of firms might use these? How are they used?
7. How can users restrict cookie use on their Web browsers? What is the downside of blocking cookies?
8. Work with a faculty member and join the Google Online Marketing Challenge (held spring of every year— see <http://www.google.com/onlinechallenge>). Google offers ad credits for student teams to develop and run online ad campaigns for real clients and offers prizes for winning teams. Some of the experiences earned in the Google Challenge can translate to other ad networks as well; and first-hand client experience has helped many students secure jobs, internships, and even start their own businesses.



8.8 Profiling and Privacy

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand the privacy concerns that arise as a result of using tracking cookies to build user profiles.
2. Be aware of the negative consequences that could result from the misuse of tracking cookies.
3. Know the steps Google has taken to demonstrate its sensitivity to privacy issues.
4. Know the kinds of user information that Google stores, and the steps Google takes to protect the privacy of that information.

Until 2009, Google hadn't used tracking cookies on its AdSense network. While AdSense has been wildly successful, contextual advertising has its



limits. For example, what kind of useful targeting can firms really do based on the text of a news item on North Korean nuclear testing? [1] So in March 2009, the firm announced what it calls “interest-based ads.” Google AdSense would now issue a third-party cookie and would track browsing activity across AdSense partner sites, and Google-owned YouTube. AdSense would build a profile, initially identifying users within thirty broad categories and six hundred subcategories. Says one Google project manager, “We’re looking to make ads even more interesting.” [2] Of course, there’s a financial incentive to do this too. Ads deemed more interesting should garner more clicks, meaning more potential customer leads for advertisers, more revenue for Web sites that run AdSense, and more money for Google.

But while targeting can benefit Web surfers, users will resist if they feel that they are being mistreated, exploited, or put at risk. Negative backlash might also result in a change in legislation. The U.S. Federal Trade Commission has already called for more transparency and user control in online advertising and for requesting user consent (opt-in) when collecting sensitive data. [3] Mishandled user privacy could curtail targeting opportunities, limiting growth across the online advertising field. And with less ad support, many of the Internet’s free services could suffer.

Google’s roll-out of interest-based ads shows the firm’s sensitivity to these issues. First, while major rivals have all linked query history to ad targeting, Google steadfastly refuses to do this. Other sites often link registration data (including user-submitted demographics such as gender and age) with tracking cookies, but Google avoids this practice as well.

Google has also placed significant control in the hands of users, with options at program launch that were notably more robust than those of its competitors. [4] Each interest-based ad is accompanied by an “Ads by Google” link that will bring users to a page describing Google advertising



and which provides access to the company’s “Ads Preferences Manager.” This tool allows surfers to see any of the hundreds of potential categorizations that Google has assigned to that browser’s tracking cookie. Users can remove categorizations, and even add interests if they want to improve ad targeting. Some topics are too sensitive to track, and the technology avoids profiling race, religion, sexual orientation, health, political or trade union affiliation, and certain financial categories.^[5] Google also allows users to install a cookie that opts them out of interest-based tracking. And since browser cookies can expire or be deleted, the firm has gone a step further, offering a browser plug-in that will remain permanent, even if a user’s opt-out cookie is purged.

Google, Privacy Advocates, and the Law

Google’s moves are meant to demonstrate transparency in its ad targeting technology, and the firm’s policies may help raise the collective privacy bar for the industry. While privacy advocates have praised Google’s efforts to put more control in the hands of users, many continue to voice concern over what they see as the increasing amount of information that the firm houses. ^[6] For an avid user, Google could conceivably be holding e-mail (Gmail), photos (Picasa), a Web surfing profile (AdSense and DoubleClick), medical records (Google Health), location (Google Latitude), appointments (Google Calendar), transcripts of phone messages (Google Voice), work files (Google Docs), and more.

Google insists that reports portraying it as a data-hording Big Brother are inaccurate. The firm is adamant that user data exists in silos that aren’t federated in any way, nor are employees permitted access to multiple data archives without extensive clearance and monitoring. Data is not sold to third parties. Activities in Gmail, Docs, or most other services isn’t added to targeting profiles. And any targeting is fully disclosed, with users



empowered to opt-out at all levels. [7] But critics counter that corporate intentions and data use policies (articulated in a Web site's Terms of Service) can change over time, and that a firm's good behavior today is no guarantee of good behavior in the future. [8]

Google does enjoy a lot of user goodwill, and it is widely recognized for its unofficial motto "Don't Be Evil." However, some worry that even though Google might not be evil, it could still make a mistake, and that despite its best intentions, a security breach or employee error could leave data dangerously or embarrassingly exposed.

When AOL released search history on over six hundred and fifty thousand of its Web searchers, these log files included queries such as "How to tell your family you're a victim of incest," "Surgical help for depression," "Can you adopt after a suicide attempt," "Gynecology oncologists in New York City," "How long will the swelling last after my tummy tuck," and perhaps most damning, queries that included specific names, addresses, and phone numbers. While AOL offered the data in a way that disguised individual user accounts, in many cases aggregate query detail contained terms so specific, they provided a strong indication of who conducted the searches. [9]

While Google has never experienced a blunder of that magnitude, it has suffered minor incidents, including a March 2009 gaffe in which the firm inadvertently shared some Google Docs with contacts who were never granted access to them. [10]

Privacy advocates also worry that the amount of data stored by Google serves as one-stop shopping for litigators and government investigators. The counter argument points to the fact that Google has continually reflected an aggressive defense of data privacy in court cases. When Viacom sued Google over copyright violations in YouTube, the search giant successfully fought the original subpoena, which had requested user-identifying information. [11] And Google was the only one of the four largest

search engines to resist a 2006 Justice Department subpoena for search queries. [12]

Google is increasingly finding itself in precedent-setting cases where the law is vague. Google's Street View, for example, has been the target of legal action in the United States, Canada, Japan, Greece, and the United Kingdom. Varying legal environments create a challenge to the global rollout of any data-driven initiative. [13]

Ad targeting brings to a head issues of opportunity, privacy, security, risk, and legislation. Google is now taking a more active public relations and lobbying role to prevent misperceptions and to be sure its positions are understood. While the field continues to evolve, Google's experience will lay the groundwork for the future of personalized technology and provide a case study for other firms that need to strike the right balance between utility and privacy. Despite differences, it seems clear to Google, its advocates, and its detractors that with great power comes great responsibility.

KEY TAKEAWAYS

- Possible consequences resulting from the misuse of cookies include user resistance and legislation. Mishandled user privacy could curtail targeting opportunities and limit growth in online advertising. With less ad support, many of the Internet's free services could suffer.
- Google has taken several steps to protect user privacy and has thus far refused to link query history or registration data to ad targeting.
- Google's "Ads Preferences Manager" allows surfers to see, remove, and add to, any of the categorizations that Google has assigned to that browser's tracking cookie. The technology also avoids targeting certain sensitive topics.

- Google allows users to install a cookie or plug-in that opts them out of interest-based tracking.
- Some privacy advocates have voiced concern over what they see as the increasing amount of information that Google houses.
- Even the best-intentioned and most competent firms can have a security breach that compromises stored information.

QUESTIONS AND EXERCISES

1. Gmail uses contextual advertising. The service will scan the contents of e-mail messages and display ads off to the side. Test the “creep out” factor in Gmail—create an account (if you don’t already have one), and send messages to yourself with controversial terms in them. Which ones showed ads? Which ones didn’t?
2. Google adamantly about not building user profiles based on Gmail messages. Ads are served based on a real-time scanning of keywords. Is this enough to make you comfortable with Google’s protection of your own privacy? Why or why not?
3. List the negative consequences that could result from the misuse of tracking cookies.
4. What steps has Google taken to give users control over the ads they wish to see?
5. Which topics does “Ads Preferences Manager” *not* target?
6. Visit Google’s Ad Preferences page. Is Google tracking your interests? Do you think the list of interests is accurate? Browse the categories under the “Ad Interest” button. Would you add any of these categories to your profile? Why or why not? What do you gain or lose by taking advantage of Google’s “Opt Out” option? Visit rival ad networks. Do you have a similar degree of control? More or less?
7. List the types of information that Google might store for an individual. Do you feel that Google is a fair and reliable steward for this information?



Are there Google services or other online efforts that you won't use due to privacy concerns? Why?

8. What steps does Google take to protect the privacy of user information?
9. Google's "interest-based advertising" was launched as an opt-out effort. What are the pros and cons for Google, users, advertisers, and AdSense partner sites if Google were to switch to an opt-in system? How would these various constituencies be impacted if the government mandated that users explicitly opt in to third-party cookies and other behavior-tracking techniques?
10. What is Google's unofficial motto?
11. What is "Street View"? Where and on what grounds is it being challenged?
12. Cite two court cases where Google has mounted a vigorous defense of data privacy.
13. *Wired News* quoted a representative of privacy watchdog group, The Center for Digital Democracy, who offered a criticism of online advertising. The representative suggested that online firms were trying to learn "everything about individuals and manipulate their weaknesses" and that the federal government should "investigate the role [that online ads] played in convincing people to take out mortgages they should not have." ^[14] Do you think online advertising played a significant role in the mortgage crisis? What role do advertisers, ad networks, and content providers have in online advertising oversight? Should this responsibility be any different from oversight in traditional media (television, print, radio)? What guidelines would you suggest?



8.9 Search Engines, Ad Networks, and Fraud

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Be able to identify various types of online fraud, as well as the techniques and technologies used to perpetrate these crimes.
2. Understand how firms can detect, prevent, and prosecute fraudsters.

There's a lot of money to be made online, and this has drawn the attention of criminals and the nefarious. Online fraudsters may attempt to steal from advertisers, harm rivals, or otherwise dishonestly game the system. But bad guys beware—such attempts violate terms-of-service agreements and may lead to prosecution and jail time.

Studying ad-related fraud helps marketers, managers, and technologists understand potential vulnerabilities, as well as the methods used to combat them. This process also builds tech-centric critical thinking, valuation, and risk assessment skills.

Some of the more common types of fraud that are attempted in online advertising include the following:

- *Enriching click fraud*—when site operators generate bogus ad clicks to earn PPC income.
- *Enriching impression fraud*—when site operators generate false page views (and hence ad impressions) in order to boost their site's CPM earnings.
- *Depleting click fraud*—clicking a rival's ads to exhaust their PPC advertising budget.
- *Depleting impression fraud*—generating bogus impressions to exhaust a rival's CPM ad budget.
- *Rank-based impression fraud*—on sites where ad rank is based on click performance, fraudsters repeatedly search keywords linked to rival ads or



access pages where rival ads appear. The goal is to generate impressions without clicks. This process lowers the performance rank (quality score) of a rival's ads, possibly dropping ads from rank results, and allowing fraudsters to subsequently bid less for the advertising slots previously occupied by rivals.

- *Disbarring fraud*—attempting to frame a rival by generating bogus clicks or impressions that appear to be associated with the rival, in hopes that this rival will be banned from an ad network or punished in search engine listings.
- *Link fraud (also known as spamdexing or link farming)*—creating a series of bogus Web sites, all linking back to a page, in hopes of increasing that page's results in organic search.
- *Keyword stuffing*—packing a Web site with unrelated keywords (sometimes hidden in fonts that are the same color as a Web site's background) in hopes of either luring users who wouldn't normally visit a Web site, or attracting higher-value contextual ads.

Disturbing stuff, but firms are after the bad guys and they've put their best geeks on the case. Widespread fraud would tank advertiser ROI and crater the online advertising market, so Google and rivals are diligently working to uncover and prosecute the crooks.

Busting the Bad Guys

On the surface, enriching click fraud seems the easiest to exploit. Just set up a Web site, run PPC ads on the page, and click like crazy. Each click should ring the ad network cash register, and a portion of those funds will be passed on to the perpetrating site owner—*ka ching!* But remember, each visitor is identified by an IP address, so lots of clicks from a single IP make the bad guys easy to spot.

So organized crime tried to raise the bar, running so-called click farms to spread fraud across dozens of IP addresses. *The Times of India* uncovered one such effort where Indian housewives were receiving up to twenty-five cents for each ad click made on fraudster-run Web sites. [1] But an unusually large number of clicks from Indian IP addresses foiled these schemes as well.

Fraudsters then moved on to use zombie networks—hordes of surreptitiously infiltrated computers, linked and controlled by rogue software. [2] To create zombie networks (sometimes called *bot nets*), hackers exploit security holes, spread viruses, or trick users into installing software that will lie dormant, awaiting commands from a central location. The controlling machine then sends out tasks for each zombie, instructing them to visit Web sites and click on ads in a way that mimics real traffic. Zombie bot nets can be massive. Dutch authorities once took down a gang that controlled some 1.5 million machines. [3]

Scary, but this is where scale, expertise, and experience come in. The more activity an ad network can monitor, the greater the chance that it can uncover patterns that are anomalous. Higher click-through rates than comparable sites? Caught. Too many visits to a new or obscure site? Caught. Clicks that don't fit standard surfing patterns for geography, time, and day? Caught.

Sometimes the goal isn't theft, but sabotage. Google's Ad Traffic Quality Team backtracked through unusual patterns to uncover a protest effort targeted at Japanese credit card firms. Ad clicks were eventually traced to an incendiary blogger who incited readers to search for the Japanese word *kiyashinku* (meaning cashing credit, or credit cards), and to click the credit card firm ads that show up, depleting firm search marketing budgets. Sneaky, but uncovered and shut down, without harm to the advertisers. [4]

Search firm and ad network software can use data patterns and other signals to ferret out most other types of fraud, too, including rank-based impression fraud, spamdexing, and keyword stuffing. While many have tried to up the stakes with increasingly sophisticated attacks, large ad networks have worked to match them, increasing their anomaly detection capabilities across all types of fraud. [5]

Click Fraud: How Bad Is It?

Accounts on the actual rate of click fraud vary widely. Some third-party firms contend that nearly one in five clicks is fraudulent. [6] But Google adamantly disputes these headline-grabbing numbers, claiming that many such reports are based on site logs that reflect false data from conditions that Google doesn't charge for (e.g., double counting a double click, or adding up repeated use of the browser back button in a way that looks like multiple clicks have occurred). The firm also offers monitoring, analytics, and reporting tools that can uncover this kind of misperceived discrepancy. Google contends that all invalid clicks (mistakes and fraud) represent less than 10 percent of all clicks, that the vast majority of these clicks are filtered out, and that Google doesn't charge advertisers for clicks flagged as mistakes or suspicious. [7] In fact, Google says their screening bar is so high and so accurate that less than 0.02 percent of clicks are reactively classified as invalid and credited back to advertisers. [8]

So who's right? While it's impossible to identify the intention behind every click, the market ultimately pays for performance. And advertisers are continuing to flock to PPC ad networks (and to Google in particular). While that doesn't mean that firms can stop being vigilant, it does suggest that for most firms, Google seems to have the problem under control.

KEY TAKEAWAYS

- Fraud can undermine the revenue model behind search engines, ad networks, and the ad-based Internet. It also threatens honest competition among rivals that advertise online.
- There are many forms of online fraud, including enriching fraud (meant to line the pockets of the perpetrators), depleting fraud (meant to waste the ad budgets of rivals), disbaring fraud (meant to frame the innocent as fraudsters), and methods to lower rival ad rank performance, or gain search engine ranking algorithms.
- While fraudsters have devised ingenious ways to exploit the system (including click farms and zombie attacks), IP addresses and detailed usage pattern monitoring increasingly reveal bogus activity.
- Fraud rates are widely disputed. However, it is clear that if widespread fraud were allowed to occur, advertisers would see lower ROI from online ad efforts, and Internet business models would suffer. The continued strength of the online advertising market suggests that while fraud may be impossible to stop completely, most fraud is under control.

QUESTIONS AND EXERCISES

1. Why is it difficult for an unscrupulous individual to pull off enriching click fraud simply by setting up a Web site, running ad network ads, and clicking?
2. Why did hackers develop zombie networks? What advantage do they offer the criminals? How are they detected? Why do larger ad networks have an advantage in click fraud detection?
3. How can you prevent zombies from inhabiting your computers? Are you reasonably confident you are “zombie-free?” Why or why not?
4. What are spamdexing and keyword stuffing? What risks does a legitimate business run if it engages in these practices, and are they discovered by search engines?



5. Which types of fraud can be attempted against search advertising? Which are perpetrated over its ad network?
6. What are the consequences if click fraud were allowed to continue? Does this ultimately help or hurt firms that run ad networks? Why?

8.10 The Battle Unfolds

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand the challenges of maintaining growth as a business and industry mature.
2. Recognize how the businesses of many firms in a variety of industries are beginning to converge.
3. Critically evaluate the risks and challenges of businesses that Google, Microsoft, and other firms are entering.



4. Appreciate the magnitude of this impending competition, and recognize the competitive forces that will help distinguish winners from losers.

Google has been growing like gangbusters, but the firm's twin engines of revenue growth—ads served on search and through its ad networks—will inevitably mature. And it will likely be difficult for Google to find new growth markets that are as lucrative as these. Emerging advertising outlets such as social networks and mobile have lower click-through rates than conventional advertising, suggesting that Google will have to work harder for less money.

For a look at what can happen when maturity hits, check out Microsoft. The House that Gates Built is more profitable than Google, and continues to dominate the incredibly lucrative markets served by Windows and Office. But these markets haven't grown much for over a decade. In industrialized nations, most Windows and Office purchases come not from growth, but when existing users upgrade or buy new machines. And without substantial year-on-year growth, the stock price doesn't move.

For big firms like Microsoft and Google, pushing stock price north requires not just new markets, but *billion-dollar* ones. Adding even one hundred million dollars in new revenues doesn't do much for firms bringing in twenty-one billion dollars and fifty-one billion dollars a year, respectively. That's why you see Microsoft swinging for the fences, investing in the uncertain, but potentially gargantuan markets of video games, mobile phone software, cloud computing (see [Chapter 10 "Software in Flux: Partly Cloudy and Sometimes Free"](#)), music and video, and of course, search and everything else that fuels online ad revenue.



Search: Google Rules, But It Ain't Over

PageRank is by no means the last word in search, and offerings from Google and its rivals continue to evolve. Google supplements PageRank results with news, photos, video, and other categorizations (click the “Show options...” link above your next Google search). Yahoo! is continually refining its search algorithms and presentation (click the little “down” arrow at the top of the firm’s search results for additional categorizations and suggestions). And Microsoft’s third entry into the search market, the “decision engine” Bing, sports nifty tweaks for specific kinds of queries. Restaurant searches in Bing are bundled with ratings stars, product searches show up with reviews and price comparisons, and airline flight searches not only list flight schedules and fares, but also a projection on whether those fares are likely go up or down. Bing also comes with a one-hundred-million-dollar marketing budget, showing that Microsoft is serious about moving its search market share out of the single digits. And in the weeks following Bing’s mid-2009 introduction, the search engine did deliver Microsoft’s first substantive search engine market share gain in years.

New tools like the Wolfram Alpha “knowledge engine” (and to a lesser extent, Google Squared) move beyond Web page rankings and instead aggregate data for comparison, formatting findings in tables and graphs. Web sites are also starting to wrap data in invisible tags that can be recognized by search engines, analysis tools, and other services. If a search engine can tell that a number on a restaurant’s Web site is, for example, either a street address, an average entrée price, or the seating capacity, it will be much easier for computer programs to accurately categorize, compare, and present this information. This is what geeks are talking about when they refer to the semantic Web. All signs point to more innovation, more competition, and an increasingly more useful Internet!



Both Google and Microsoft are on a collision course. But there's also an impressive roster of additional firms circling this space, each with the potential to be competitors, collaborators, merger partners, or all of the above. While wounded and shrinking, Yahoo! is still a powerhouse, ranking ahead of Google in some overall traffic statistics. Add in eBay, Facebook, Twitter, Amazon, Salesforce.com, Apple, Netflix, the video game industry, telecom and mobile carriers, cable firms, and the major media companies, and the next few years have the makings of a big, brutal fight.

Strategic Issues

Google's scale advantages in search and its network effects advantages in advertising were outlined earlier. The firm also leads in search/ad experience and expertise and continues to offer a network reach that's unmatched. But the strength of Google's other competitive resources is less clear.

Within Google's ad network, there are *switching costs* for advertisers and for content providers. Google partners have set up accounts and are familiar with the firm's tools and analytics. Content providers would also need to modify Web sites to replace AdSense or DoubleClick ads with rivals. But choosing Google doesn't cut out the competition. Many advertisers and content providers participate in multiple ad networks, making it easier to shift business from one firm to another. That likely means that Google will have to retain its partners by offering superior value.

Another vulnerability may exist with search consumers. While Google's brand is strong, switching costs for search users are incredibly low. Move from Google.com to Ask.com and you actually save three letters of typing! Still, there are no signs that Google's search leadership is in jeopardy. So far users have been creatures of habit, returning to Google despite heavy marketing by rivals. And in Google's first decade, no rival has offered



technology compelling enough to woo away the googling masses—the firm’s share has only increased. Defeating Google with some sort of technical advantage will be difficult, since Web-based innovation can often be quickly imitated.

The Google Toolbar helps reinforce search habits among those who have it installed, and Google pays the Mozilla foundation (the folks behind the Firefox browser) upwards of sixty-six million dollars a year to serve as its default search option for the open source browser. [1] But Google’s track record in expanding reach through distribution deals is mixed. The firm spent nearly one billion dollars to have MySpace run AdSense ads, but Google has publicly stated that social network advertising has not been as lucrative as it had hoped (see [Chapter 7 "Facebook: Building a Business from the Social Graph"](#)). The firm has also spent nearly one billion dollars to have Dell preinstall its computers with the Google browser toolbar and Google desktop search products. But in 2009, Microsoft inked deals that displaced Google on Dell machines, and it also edged Google out in a five-year search contract with Verizon Wireless. [2]

How Big Is Too Big?

Microsoft could benefit from embedding its Bing search engine into its most popular products (imagine putting Bing in the right-mouseclick menu alongside *cut*, *copy*, and *paste*). But with Internet Explorer market share above 65 percent, Office above 80 percent, and Windows at roughly 90 percent, [3] this seems unlikely.

European antitrust officials have already taken action against Redmond’s bundling Windows Media Player and Internet Explorer with Windows. Add in a less favorable antitrust climate in the United States, and tying any of these products to Bing is almost certainly out of bounds. What’s not clear is whether or not regulators would allow Bing to be bundled with less

dominant Microsoft offerings, such as mobile phone software, Xbox, and MSN.

But increasingly, Google is also an antitrust target. Microsoft has itself raised antitrust concerns against Google, unsuccessfully lobbying both U.S. and European authorities to block the firm's acquisition of DoubleClick. [4] Google was forced to abandoned a fall, 2008, search advertising partnership with Yahoo! after the Justice Department indicated its intention to block the agreement. The Justice Department is also investigating a Google settlement with the Authors' Guild, a deal in which critics have suggested that Google scored a near monopoly on certain book scanning, searching, and data serving rights. [5] And yet another probe is investigating whether Google colluded with Apple, Yahoo! and other firms to limit efforts to hire away top talent. [6]

Of course, being big isn't enough to violate U.S. antitrust law. Harvard Law's Andrew Gavil says, "You've got to be big, and you have to be bad. You have to be both." [7] This may be a difficult case to make against a firm that has a history of being a relentless supporter of open computing standards. And as mentioned earlier, there is little forcing users to stick with Google—the firm must continue to win this market on its own merits. While Google may escape all of these investigations, increased antitrust scrutiny is a downside that comes along with the advantages of market-dominating scale.

More Ads, More Places, More Formats

Google has been a champion of increased Internet access. But altruism aside, more net access also means a greater likelihood of ad revenue. Google's effort to catalyze Internet use worldwide comes through on multiple fronts. In the United States, Google has supported (with varying degrees of success) efforts to offer free wi-fi in San Francisco and Mountain

View. But most ambitiously, Google is also a major backer (along with Liberty Global and HSBC) of the O3b satellite network. O3b stands for “the other three billion” of the world’s population who currently lack Internet access. O3b hopes to have sixteen satellites circling the globe by late 2010, blanketing underserved regions with low latency (low delay), high-speed Internet access. [8] With Moore’s Law dropping computing costs as world income levels rise, Google hopes to empower the currently disenfranchised masses to start surfing. Good for global economies, good for living standards, and good for Google.

Another way Google can lower the cost of surfing is by giving mobile phone software away for free. That’s the thinking behind the firm’s Android offering. With Android, Google provides mobile phone vendors with a Linux-based operating system, supporting tools, standards, and an application marketplace akin to Apple’s AppStore. Android itself isn’t ad-supported—there aren’t Google ads embedded in the OS. But the hope is that if handset manufacturers don’t have to write their own software, the cost of wireless mobile devices will go down. And cheaper devices mean that more users will have access to the mobile Internet, adding more ad-serving opportunities for Google and its partner sites.

While Android started as a mobile phone operating system and software stack, its use has been expanded by hardware manufacturers; it is now serving as a Windows alternative on low-cost, Internet-equipped laptops (netbooks) and eBook readers. And in 2009, Google announced a direct strike at the netbook—its Chrome OS. Powered by a combination of open source Linux and Google’s open source Chrome browser, the Chrome OS is specifically designed to provide a lightweight, but consistent user interface for applications that otherwise live “in the clouds,” preferably residing on Google’s server farms (see [Chapter 10 "Software in Flux: Partly Cloudy and Sometimes Free"](#)). And a tailored Android or Chrome OS could conceivably



power net access in a host of devices, including televisions, set top boxes, and automobiles. Google has dabbled in selling ads for television (as well as radio and print). There may be considerable potential in bringing variants of ad targeting technology, search, and a host of other services across these devices.

Google has also successfully lobbied the U.S. government to force wireless telecom carriers to be more open, dismantling what are known in the industry as walled gardens. Before Google's lobbying efforts, mobile carriers could act as gatekeepers, screening out hardware providers and software services from their networks. Now, paying customers of carriers that operate over the recently allocated U.S. wireless spectrum will have access to a choice of hardware and less restrictive access to Web sites and services. And Google hopes this expands its ability to compete without obstruction.

YouTube

Then there's Internet video, with Google in the lead here too. It's tough to imagine any peer-produced video site displacing YouTube. Users go to YouTube because there's more content, while amateur content providers go there seeking more users (classic two-sided network effects). This critical advantage was the main reason why, in 2006, Google paid \$1.65 billion for what was then just a twenty-month-old startup.

That popularity comes at a price. Even with falling bandwidth and storage costs, at twenty hours of video uploaded to YouTube *every minute*, the cost to store and serve this content is crippling large. ^[9] Credit Suisse estimates that in 2009, YouTube will bring in roughly \$240 million in ad revenue, pitted against \$711 million in operating expenses. That's a shortfall of more than \$470 million. Analysts estimate that for YouTube to break even, it would need to achieve an ad CPM of \$9.48 on each of the roughly seventy-five billion streams it'll serve up this year. A tough task.



For comparison, Hulu (a site that specializes in offering ad-supported streams of television shows and movies) earns CPM rates of thirty dollars and shares about 70 percent of this with copyright holders. Most user-generated content sports CPM rates south of a buck. ^[10] Some differ with the Credit Suisse report—RampRate pegs the losses at \$174 million. In fact, it may be in Google’s interest to allow others to think of YouTube as more of a money pit than it really is. That perception might keep rivals away longer, allowing the firm to solidify its dominant position while getting the revenue model right. Even as a public company, Google can keep mum about YouTube specifics. Says the firm’s CFO, “We know our cost position, but nobody else does.” ^[11]

The explosion of video uploading is also adding to costs as more cell phones become net-equipped video cameras. YouTube’s mobile uploads were up 400 percent in just the first week following the launch of the video-equipped iPhone 3GS. ^[12] Viewing will also skyrocket as mobile devices and television sets ship with YouTube access, adding to revenue potential. The firm is still experimenting with ad models—these include traditional banner and text ads, plus ads transparently layered across the bottom 20 percent of the screen, *pre-roll* commercials that appear before the selected video, and more. Google has both the money and time to invest in nurturing this market, and it continues to be hesitant in saturating the media with ads that may annoy users and constrain adoption.

Apps and Innovation

In 2007 the firm announced a tagline to sum up its intensions: “search, ads, and apps.” Google is king of the first two, but this last item hasn’t matured to the point where it impacts the firm’s financials.

Experimentation and innovation are deeply ingrained in Google’s tech-centric culture, and this has led to a flood of product offerings. Google

released more than 360 products in 2008, and another 120 in Q1 2009. ^[13] It's also cancelled several along the way, including Jaiku (which couldn't beat Twitter), Google Video (which was superseded by the YouTube acquisition), and a bunch more you've likely not heard of, like Dodgeball, Notebook, Catalog Search, and Mashup Editor. ^[14]

What's Google Up To?

With all this innovation, it's tough to stay current with Google's cutting edge product portfolio. But the company does offer "beta" releases of some projects, and invites the public to try out and comment on its many experiments. To see some of these efforts in action, visit Google Labs at <http://googlelabs.com>. And to see a current list of more mature offerings, check out <http://www.google.com/options>.

Google's "Apps" are mostly Web-based software-as-a-service offerings. Apps include an Office-style suite that sports a word processor, presentation tool, and spreadsheet, all served through a browser. While initially clunky, the products are constantly being refined. The spreadsheet product, for example, has been seeing new releases every two weeks, with features such as graphing and pivot tables inching it closer in capabilities to desktop alternatives. ^[15] And new browser standards, such as HTML 5, will make it even easier for what lives in the browser to mimic what you're currently using on your desktop.

Google also offers Gears, a product that allows compatible apps to be used offline when net access isn't available. That'll be critical as long as Internet access is less reliable than your hard drive, but online collaboration is where these products can really excel (no pun intended). Most Google apps allow not only group viewing, but also collaborative editing, common storage, and version control. Google's collaboration push also includes its



wiki-like Google Sites tool, and a new platform called Wave, billed as a sort of next-step evolving beyond e-mail and instant messaging.

Unknown is how much money Google will make off all of this. Consumers and small businesses have free access to these products, with usage for up to fifty users funded by in-app ads. But is there much of a market serving ads to people working on spreadsheets? Enterprises can gain additional, ad-free licenses for a fee. While users have been reluctant to give up Microsoft Office, many have individually migrated to Google's Web-based e-mail and calendar tools. Google's enterprise apps group will now do the same thing for organizations, acting as a sort of outsourcer by running e-mail, calendar, and other services for a firm; all while handling upgrades, spam screening, virus protection, backup, and other administrative burdens. Arizona State University, biotech giant Genentech, and auto parts firm Valeo are among the Google partners that have signed on to make the firm's app offerings available to thousands. ^[16]

And of course, Microsoft won't let Google take this market without a fight. Office 10 was announced along with a simplified, free, ad-supported, Web-based, online options for Word, Excel, PowerPoint, and OneNote; and Microsoft can also migrate applications like e-mail and calendaring off corporate computers and onto Microsoft's server farms.

It's not until considered in its entirety that one gets a sense for what Google has the potential to achieve. It's possible that increasing numbers of users worldwide will adopt light, cheap netbooks and other devices powered by free Google software (Android, Google's Chrome browser and Chrome OS). Productivity apps, e-mail, calendaring, and collaboration tools will all exist "in the cloud," accessible through any browser, with files stored on Google's servers in a way that minimizes hard drive needs. Google will entertain you, help you find the information you need, help you shop, handle payment (Google Checkout), and more. And the firms you engage online may



increasingly turn to Google to replace their existing hardware and software infrastructure with corporate computing platforms like Google Apps Engine (see [Chapter 10 "Software in Flux: Partly Cloudy and Sometimes Free"](#)). All of this would be based on open standards, but switching costs, scale, and increasing returns from expertise across these efforts could yield enormous advantages.

Studying Google allowed us to learn about search and the infrastructure that powers this critical technology. We've studied the business of ads, covering search advertising, ad networks, and ad targeting in a way that blends strategic and technology issues. And we've covered the ethical, legal, growth, and competitive challenges that Google and its rivals face. Studying Google in this context should not only help you understand what's happening today, it should also help you develop critical thinking skills for assessing the opportunities and threats that will emerge across industries as technologies continue to evolve.

KEY TAKEAWAYS

- For over a decade, Google's business has been growing rapidly, but that business is maturing.
- Slower growth will put pressure on the firm's stock price, so a firm Google's size will need to pursue very large, risky, new markets—markets that are also attractive to well-financed rivals, smaller partners, and entrepreneurs.
- Rivals continue to innovate in search. Competing with technology is extremely difficult, since it is often easy for a firm to mimic the innovations of a pioneer with a substitute offering. Microsoft, with profits to invest in infrastructure, advertising, and technology, may pose Google's most significant, conventional threat.

- Although Microsoft has many distribution channels (Windows, Internet Explorer, Office) for its search and other services, European and U.S. regulators will likely continue to be prevent the firm from aggressive product and service bundling.
- Google is investing heavily in methods that promote wider Internet access. These include offering free software to device manufacturers, and several telecommunications and lobbying initiatives meant to lower the cost of getting online. The firm hopes that more users spending more time online will allow it to generate more revenue through ads and perhaps other services.
- YouTube demonstrates how a firm can create a large and vastly influential business in a short period of time, but that businesses that host and serve large files of end-user content can be costly.
- Google, Microsoft, and smaller rivals are also migrating applications to the Web, allowing Office-style software to execute within a browser, with portions of this computing experience and storage happening off a user's computer, "in the cloud" of the Internet. Revenue models for this business are also uncertain.
- With scale and influence comes increased governmental scrutiny. Google has increasingly become a target of antitrust regulators. The extent of this threat is unclear. Google's extreme influence is clear. However, the firm's software is based on open standards; competitors have a choice in ad networks, search engines, and other services; switching costs are relatively low; users and advertisers aren't locked into exclusive contracts for the firm's key products and services; and there is little evidence of deliberate, predatory pricing or other "red-flag" activity that usually brings government regulation.



QUESTIONS AND EXERCISES

1. Perform identical queries on both Google and on rival search engines. Try different categories (research for school projects, health, business, sports, entertainment, local information). Which sites do you think give you the better results? Why? Would any of these results cause you to switch to one search engine versus the other?
2. Investigate new services that attempt to extend the possibilities for leveraging online content. Visit Bing, Google Squared, Wolfram Alpha, and any other such efforts that intrigue you. Assume the role of a manager and use these engines to uncover useful information. Assume your role as a student and see if these tools provide valuable information for this or other classes. Are you likely to use these tools in the future? Why or why not? Under what circumstances are they useful and when do they fall short?
3. Assume the role of an industry analyst: Consider the variety of firms mentioned in this section that may become competitors or partners. Create a chart listing your thoughts on which firms are likely to collaborate and work together, and which firms are likely to compete. What are the advantages or risks in these collaborations for the partners involved? Do you think any of these firms are “acquisition bait?” Defend your predictions and be prepared to discuss them with your class.
4. Assume the role of an IT manager: to the extent that you can, evaluate online application offerings by Google, Microsoft, and rivals. In your opinion, are these efforts ready for prime time? Why or why not? Would you recommend that a firm choose these applications? Are there particular firms or users that would find these alternatives particularly appealing? Would you ever completely replace desktop offerings with online ones? Why or why not?



5. Does it make sense for organizations to move their e-mail and calendaring services off their own machines and pay Google, Microsoft, or someone else to run them? Why or why not?
6. What are Chrome, the Chrome OS, and Android? Are these software products successful in their respective categories? Investigate the state of the market for products that leverage any of these software offerings. Would you say that they are successful? Why or why not? What do you think the outlook is for Chrome, the Chrome OS, and Android? As an IT manager, would you recommend products based on this software? As an investor, do you think it continues to make sense for Google to develop these efforts? Why or why not?
7. Google's unofficial motto is "Don't be evil." But sometimes it's challenging for managers to tell what path is "most right" or "least wrong." Google operates in countries that require the firm to screen and censor results. Short term, this is clearly a limitation on freedom of speech. But long-term, access to the Internet could catalyze economic development and spread information in a way that leads to more democratization. Investigate and consider both of these arguments and be prepared to argue the case either for limiting work in speech-limiting countries, or in working within them as a potential agent of change. What other pressures is a publicly traded firm under to choose one path or the other? Which path would you choose and why?

Chapter 9

Understanding Software: A Primer for Managers

9.1 Introduction



LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Recognize the importance of software and its implications for the firm and strategic decision making.
2. Understand that software is everywhere; not just in computers, but also cell phones, cars, cameras, and many other technologies.
3. Know what software is and be able to differentiate it from hardware.
4. List the major classifications of software and give examples of each.

We know computing hardware is getting faster and cheaper, creating all sorts of exciting and disruptive opportunities for the savvy manager. But what's really going on inside the box? It's software that makes the magic of computing happen. Without software, your PC would be a heap of silicon wrapped in wires encased in plastic and metal. But it's the instructions—the software code—that enable a computer to do something wonderful, driving the limitless possibilities of information technology.

Software is everywhere. An inexpensive cell phone has about one million lines of code, while the average car contains nearly one hundred million. ^[1] In this chapter we'll take a peek inside the chips to understand what software is. A lot of terms are associated with software: operating systems, applications, enterprise software, distributed systems, and more. We'll define these terms up front, and put them in a managerial context. A follow-up chapter, [Chapter 10 "Software in Flux: Partly Cloudy and Sometimes Free"](#), will focus on changes impacting the software business, including open source software, software as a service (SaaS), and cloud computing. These changes are creating an environment radically different from the software industry that existed in prior decades—confronting managers with a whole new set of opportunities and challenges.



Managers who understand software can better understand the possibilities and impact of technology. They can make better decisions regarding the strategic value of IT and the potential for technology-driven savings. They can appreciate the challenges, costs, security vulnerabilities, legal and compliance issues, and limitations involved in developing and deploying technology solutions. In the next two chapters we will closely examine the software industry and discuss trends, developments and economics—all of which influence decisions managers make about products to select, firms to partner with, and firms to invest in.

What Is Software?

When we refer to computer hardware (sometimes just hardware), we're talking about the physical components of information technology—the equipment that you can physically touch, including computers, storage devices, networking equipment, and other peripherals.

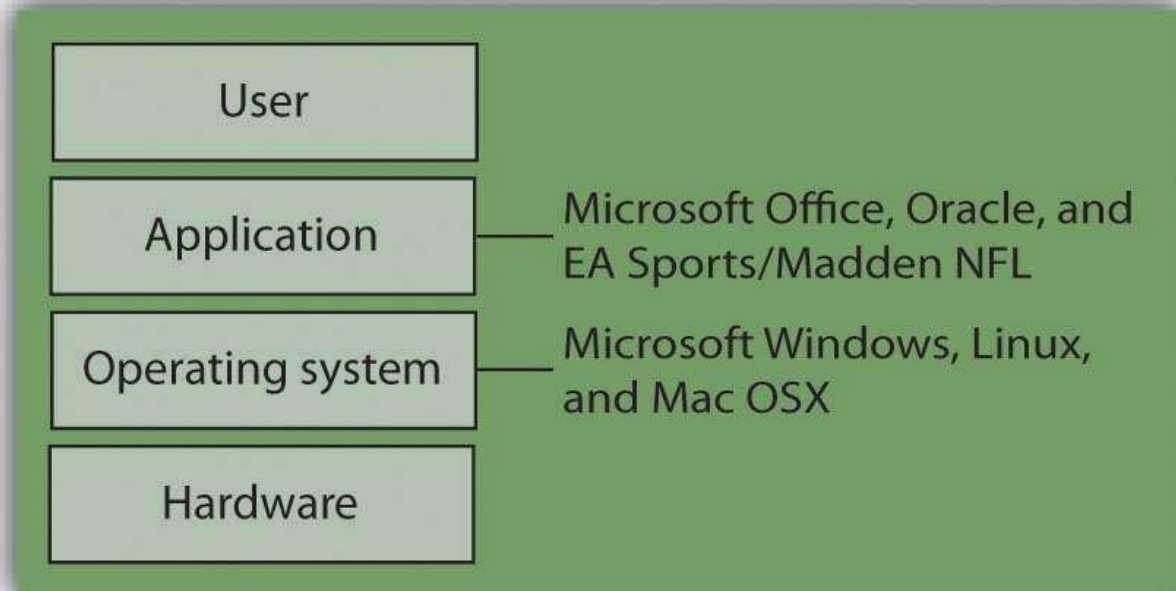
Software refers to a computer program or collection of programs—sets of instructions that tell the hardware what to do. Software gets your computer to behave like a Web browser or word processor, makes your iPod play music and video, and enables your bank's ATM to spit out cash.

It's when we start to talk about the categories of software that most people's eyes glaze over. To most folks, software is a big, incomprehensible alphabet soup of acronyms and geeky phrases: OS, VB, SAP, SQL, to name just a few. Don't be intimidated. The basics are actually pretty easy to understand. But it's not soup; it's more of a layer cake. Think about computer hardware as being at the bottom of the layer cake. The next layer is the operating system, the collection of programs that control the hardware. Windows, Mac OS X, and Linux are operating systems. On top of that layer are applications—these can range from end-user programs like those in



Office, to the complex set of programs that manage a business's inventory, payroll, and accounting. At the top of the cake are users.

Figure 9.1 The Hardware/Software Layer Cake



The flexibility of these layers gives computers the customization options that managers and businesses demand. Understanding how the layers relate to each other helps you make better decisions on what options are important to your unique business needs, can influence what you buy, and may have implications for everything from competitiveness to cost overruns to security breaches. What follows is a manager's guide to the main software categories with an emphasis on why each is important.

KEY TAKEAWAYS

- Software refers to a computer program or collection of programs. It enables computing devices to perform tasks.
- You can think of software as being part of a layer cake, with hardware at the bottom; the operating system controlling the hardware and

establishing standards, the applications executing one layer up, and the users at the top.

- How these layers relate to one another has managerial implications in many areas, including the flexibility in meeting business demand, costs, legal issues and security.
- Software is everywhere—not just in computers, but also in cell phones, cars, cameras, and many other technologies.

QUESTIONS AND EXERCISES

1. Explain the difference between hardware and software.
2. Why should a manager care about software and how software works? What critical organizational and competitive factors can software influence?
3. What role has software played in your decision to select certain products? Has this influenced why you favored one product or service over another?
4. Find the *Fortune* 500 list online. Which firm is the highest ranked software firm? While the *Fortune* 500 ranks firms according to revenue, what's this firm's profitability rank? What does this discrepancy tell you about the economics of software development? Why is the software business so attractive to entrepreneurs?
5. Refer to earlier chapters (and particularly to [Chapter 2 "Strategy and Technology"](#)): Which resources for competitive advantage might top software firms be able to leverage to ensure their continued dominance? Give examples of firms that have leveraged these assets, and why they are so strong.

9.2 Operating Systems

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand what an operating system is and why computing devices require operating systems.
2. Appreciate how embedded systems extend Moore's Law, allowing firms to create "smarter" products and services

Computing hardware needs to be controlled, and that's the role of the operating system. The operating system (sometimes called the "OS") provides a common set of controls for managing computer hardware, making it easier for users to interact with computers and for programmers to write application software. Just about every computing device has an operating system—desktops and laptops, enterprise-class server computers, your mobile phone. Even specialty devices like iPods, video game consoles, and television set top boxes run some form of OS.

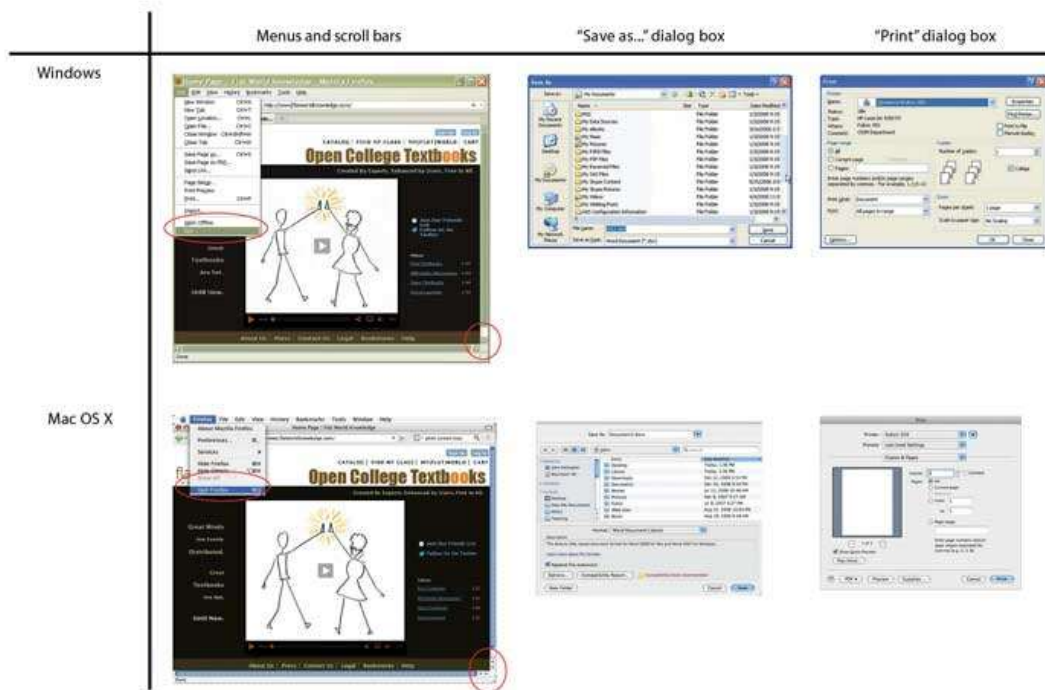


Some firms, like Apple and Nintendo, develop their own proprietary OS for their own hardware. Microsoft sells operating systems to everyone from Dell to the ATM manufacturer Diebold (listen for the familiar Windows error beep on some cash machines). And there are a host of specialty firms, such as Wind River (purchased by Intel), that help firms develop operating systems for all sorts of devices that don't necessarily look like a PC, including cars, video editing systems, and fighter jet control panels. Anyone who has used both a PC and a Mac and has noticed differences across these platforms can get a sense of the breadth of what an operating system does. Even for programs that are otherwise identical for these two systems (like the Firefox browser), subtle differences are visible. Screen elements like menus, scroll bars, and window borders look different on the Mac than they do in Windows. So do the dialog boxes that show up when you print or save.

These items look and behave differently because each of these functions touches the hardware, and the team that developed Microsoft Windows created a system distinctly different from their Macintosh counterparts at Apple. Graphical user interface (GUI) items like scroll bars and menus are displayed on the hardware of the computer display. Files are saved to the hardware of a hard drive or other storage device. Most operating systems also include control panels, desktop file management, and other support programs to work directly with hardware elements like storage devices, displays, printers, and networking equipment. The Macintosh Finder and the Windows Explorer are examples of components of these operating systems. The consistent look, feel, and functionality that operating systems enforce across various programs help make it easier for users to learn new software, which reduces training costs and operator error. See [Figure 9.2](#) for similarities and differences.



Figure 9.2



Differences between the Windows and Mac operating systems are evident throughout the user interface, particularly when a program interacts with hardware.

Operating systems are also designed to give programmers a common set of commands to consistently interact with the hardware. These commands make a programmer's job easier by reducing program complexity and making it faster to write software while minimizing the possibility of errors in code. Consider what an OS does for the Wii game developer. Nintendo's Wii OS provides Wii programmers with a set of common standards to use to access the Wiimote, play sounds, draw graphics, save files, and more. Without this, games would be a lot more difficult to write, they'd likely look differently, be less reliable, would cost more, and there would be fewer titles available.

Similarly, when Apple provided developers with a common set of robust, easy-to-use standards for the iPhone and (via the App Store) an easy way for users to install these applications on top of the iPhone/iPod Touch OS, software development boomed, and Apple became hands-down the most versatile mobile computing device available. ^[1] In Apple's case, some *fifty thousand apps* became available through the App store in less than a year. A good OS and software development platform can catalyze network effects (see [Chapter 5 "Understanding Network Effects"](#)). While the OS seems geeky, its effective design has very strategic business implications!

Firmware and Embedded Systems

Most personal computers have an operating system installed on their hard drives. This system allows the OS to be replaced or upgraded easily. But many smaller, special-purpose computing devices have their operating systems installed on nonvolatile memory, often on read-only memory (ROM) chips. Control programs stored on chips are sometimes referred to as firmware. The OS in an iPod, mobile phone, or your TV's set-top box is most likely stored as firmware. Your PC also has a tiny bit of firmware that allows it to do very basic functions like startup (boot) and begin loading its operating system from disk.

Another term you might hear is embedded systems. As computing gets cheaper, special-purpose technology is increasingly becoming embedded into all sorts of devices like cars, picture frames, aircraft engines, photocopiers, and heating and air conditioning systems. The software programs that make up embedded systems are often stored as firmware too.

Moore's law (see [Chapter 4 "Moore's Law and More: Fast, Cheap Computing and What It Means for the Manager"](#)) enables embedded systems, and these systems can create real strategic value. The Otis

Elevator Company, a division of United Technologies, uses embedded systems in its products to warn its service centers when the firm's elevators, escalators, and moving walkways need maintenance or repair. This warning provides Otis with several key benefits:

1. Since products automatically contact Otis when they need attention, these systems generate a lucrative service business for the firm and make it more difficult for third parties to offer a competing business servicing Otis products.
2. Products contact service technicians to perform maintenance based on exact needs (e.g., lubricant is low, or a part has been used enough to be replaced) rather than guessed schedules, which makes service more cost effective, products less likely to break down, and customers happier.
3. Any product failures are immediately detected, with embedded systems typically dispatching technicians before a client's phone call.
4. The data is fed back to Otis's R&D group, providing information on reliability and failure so that engineers can use this info to design better products.

Collectively, software embedded on tiny chips yields very big benefits, for years helping Otis remain at the top of its industry.

KEY TAKEAWAYS

- The operating system (OS) controls a computer's hardware and provides a common set of commands for writing programs.
- Most computing devices (enterprise-class server computers, PCs, phones, set-top boxes, video games, cars, the Mars Rover) have an operating system.
- Some products use operating systems provided by commercial firms, while others develop their own operating system. Others may leverage

open-source alternatives (see [Chapter 10 "Software in Flux: Partly Cloudy and Sometimes Free"](#)).

- Embedded systems are special-purpose computer systems designed to perform one or a few dedicated functions, and are frequently built into conventional products like cars, air conditioners, and elevators.
- Embedded systems can make products and services more efficient, more reliable, more functional, and can enable entire new businesses and create or reinforce resources for competitive advantage.

QUESTIONS AND EXERCISES

1. What does an operating system do? Why do you need an operating system? How do operating systems make a programmer's job easier? How do operating systems make life easier for end-users?
2. How has the market for desktop, server, and mobile operating systems changed in recent years? Do certain products seem to be gaining traction? Why do you think this is the case?
3. What kinds of operating systems are used in the devices that you own? On your personal computer? Your mobile phone? The set-top box on top of your television? Are there other operating systems that you come into contact with? If you can't tell which operating system is in each of these devices, see if you can search the Internet to find out.
4. For your list in the prior question (and to the extent that you can), diagram the hardware/software "layer cake" for these devices.
5. For this same list, do you think each device's manufacturer wrote all of the software that you use on these devices? Can you add or modify software to all of these devices? Why or why not? What would the implications be for cost, security, complexity, reliability, updates and upgrades, and the appeal of each device?
6. Some ATM machines use Windows. Why would an ATM manufacturer choose to build its systems owing Windows? Why might it want to avoid

this? Are there other non-PC devices you've encountered that were running some form of Windows?

7. What are embedded systems? When might firms want to install software on chips instead of on a hard drive?
8. It's important to understand how technology impacts a firm's strategy and competitive environment. Consider the description of Otis elevator's use of embedded systems. Which parts of the value chain does this impact? How? Consider the "five forces": How does the system impact the firm's competitive environment? Are these systems a source of competitive advantage? If not, explain why not? If they are, what kinds of resources for competitive advantage can these kinds of embedded systems create?
9. Can you think of other firms that can or do leverage embedded systems? Provide examples and list the kinds of benefits these might offer firms and consumers.
10. Research the Americans with Disabilities Act of 1990 (or investigate if your nation has a similar law), and the implications of this legislation for software developers and Web site operators. Have firms been successfully sued when their software or Web sites could not be accessed by users with physical challenges? What sorts of issues should developers consider when making their products more accessible? What practices might they avoid?



9.3 Application Software

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Appreciate the difference between desktop and enterprise software.
2. List the categories of enterprise software.
3. Understand what an ERP (enterprise resource planning) software package is.
4. Recognize the relationship of the DBMS (database system) to the other enterprise software systems.
5. Recognize both the risks and rewards of installing packaged enterprise systems.



Operating systems are designed to create a platform so that programmers can write additional applications, allowing the computer to do even more useful things. While operating systems control the hardware, *application software* (sometimes referred to as *software applications*, *applications*, or even just *apps*) perform the work that users and firms are directly interested in accomplishing. Think of applications as the place where the users or organization's real work gets done. As we learned in [Chapter 5 "Understanding Network Effects"](#), the more application software that is available for a platform (the more games for a video game console, the more apps for your phone), the more valuable it potentially becomes. Desktop software refers to applications installed on a personal computer—your browser, your Office suite (e.g., word processor, spreadsheet, presentation software), photo editors, and computer games are all desktop software. Enterprise software refers to applications that address the needs of multiple, simultaneous users in an organization or work group. Most companies run various forms of enterprise software programs to keep track of their inventory, record sales, manage payments to suppliers, cut employee paychecks, and handle other functions. Some firms write their own enterprise software from scratch, but this can be time consuming and costly. Since many firms have similar procedures for accounting, finance, inventory management, and human resource functions, it often makes sense to buy a software package (a software product offered commercially by a third party) to support some of these functions. So-called enterprise resource planning (ERP) software packages serve precisely this purpose. In the way that Microsoft can sell you a suite of desktop software programs that work together, many companies sell ERP software that coordinates and integrates many of the functions of a business. The leading ERP vendors include the firm's SAP and Oracle,



although there are many firms that sell ERP software. A company doesn't have to install all of the modules of an ERP suite, but it might add functions over time—for example, to plug in an accounting program that is able to read data from the firm's previously installed inventory management system. And although a bit more of a challenge to integrate, a firm can also mix and match components, linking software the firm has written with modules purchased from different enterprise software vendors.

Figure 9.4 *ERP in Action* [1]



An ERP system with multiple modules installed can touch many functions of the business:

- *Sales*—A sales rep from Vermont-based SnowboardCo. takes an order for five thousand boards from a French sporting goods chain. The system can verify credit history, apply discounts, calculate price (in Euros), and print the order in French.
- *Inventory*—While the sales rep is on the phone with his French customer, the system immediately checks product availability, signaling that one

thousand boards are ready to be shipped from the firm's Burlington warehouse, the other four thousand need to be manufactured and can be delivered in two weeks from the firm's manufacturing facility in Guangzhou.

- *Manufacturing*—When the customer confirms the order, the system notifies the Guangzhou factory to ramp up production for the model ordered.
- *Human Resources*—High demand across this week's orders triggers a notice to the Guangzhou hiring manager, notifying her that the firm's products are a hit and that the flood of orders coming in globally mean her factory will have to hire five more workers to keep up.
- *Purchasing*—The system keeps track of raw material inventories, too. New orders trigger an automatic order with SnowboardCo's suppliers, so that raw materials are on hand to meet demand.
- *Order Tracking*—The French customer can log in to track her SnowboardCo order. The system shows her other products that are available, using this as an opportunity to cross-sell additional products.
- *Decision Support*—Management sees the firm's European business is booming and plans a marketing blitz for the continent, targeting board models and styles that seem to sell better for the Alps crowd than in the U.S. market.

Other categories of enterprise software that managers are likely to encounter include the following:

- customer relationship management (CRM) systems used to support customer-related sales and marketing activities;
- supply chain management (SCM) systems that can help a firm manage aspects of its value chain, from the flow of raw materials into the firm through delivery of finished products and services at the point-of-consumption; and



- business intelligence (BI) systems, which use data created by other systems to provide reporting and analysis for organizational decision making.

Major ERP vendors are now providing products that extend into these and other categories of enterprise application software, as well.

Most enterprise software works in conjunction with a database management system (DBMS), sometimes referred to as a “database system.” The database system stores and retrieves the data that an application creates and uses. Think of this as another additional layer in our cake analogy. Although the DBMS is itself considered an application, it’s often useful to think of a firm’s database systems as sitting above the operating system, but under the enterprise applications. Many ERP systems and enterprise software programs are configured to share the same database system so that an organization’s different programs can use a common, shared set of data. This system can be hugely valuable for a company’s efficiency. For example, this could allow a separate set of programs that manage an inventory and point-of-sale system to update a single set of data that tells how many products a firm has to sell and how many it has already sold—information that would also be used by the firm’s accounting and finance systems to create reports showing the firm’s sales and profits.

Firms that don’t have common database systems with consistent formats across their enterprise often struggle to efficiently manage their value chain. Common procedures and data formats created by packaged ERP systems and other categories of enterprise software also make it easier for firms to use software to coordinate programs between organizations. This coordination can lead to even more value chain efficiencies. Sell a product? Deduct it from your inventory. When inventory levels get too low, have your computer systems send a message to your supplier’s systems so that they can automatically build and ship replacement product to your firm. In



many cases these messages are sent without any human interaction, reducing time and errors. And common database systems also facilitate the use of BI systems that provide critical operational and competitive knowledge and empower decision making. For more on CRM and BI systems, and the empowering role of data, see [Chapter 11 "The Data Asset: Databases, Business Intelligence, and Competitive Advantage"](#).

Figure 9.5



An organization's database management system can be set up to work with several applications both within and outside the firm.

The Rewards and Risks of Packaged Enterprise Systems

When set up properly, enterprise systems can save millions of dollars and turbocharge organizations. For example, the CIO of office equipment maker Steelcase credited the firm's ERP with an eighty-million-dollar reduction in operating expenses saved from eliminating redundant processes and making data more usable. The CIO of Colgate Palmolive also praised their ERP, saying, "The day we turned the switch on, we dropped two days out of our order-to-delivery cycle." [2] Packaged enterprise systems can streamline processes, make data more usable, and ease the linking of systems with software across the firm and with key business partners. Plus, the software

that makes up these systems is often debugged, tested, and documented with an industrial rigor that may be difficult to match with proprietary software developed in-house.

But for all the promise of packaged solutions for standard business functions, enterprise software installations have proven difficult.

Standardizing business processes in software that others can buy means that those functions are easy for competitors to match, and the vision of a single monolithic system that delivers up wondrous efficiencies has been difficult for many to achieve. The average large company spends roughly fifteen million dollars on ERP software, with some installations running into the hundreds of millions of dollars. [3] And many of these efforts have failed disastrously.

FoxMeyer was once a six-billion-dollar drug distributor, but a failed ERP installation led to a series of losses that bankrupted the firm. The collapse was so rapid and so complete that just a year after launching the system, the carcass of what remained of the firm was sold to a rival for less than eighty million dollars. Hershey Foods blamed a \$466 million revenue shortfall on glitches in the firm's ERP rollout. Among the problems, the botched implementation prevented the candy maker from getting product to stores during the critical period before Halloween. Nike's first SCM and ERP implementation was labeled a "disaster"; their systems were blamed for over one hundred million dollars in lost sales. [4] Even tech firms aren't immune to software implementation blunders. HP once blamed a \$160 million loss on problems with its ERP systems. [5] Manager beware—there are no silver bullets. For insight on the causes of massive software failures, and methods to improve the likelihood of success, see [Section 9.6 "Total Cost of Ownership \(TCO\): Tech Costs Go Way beyond the Price Tag"](#) later in this chapter.

KEY TAKEAWAYS

- Application software focuses on the work of a user or an organization.
- Desktop applications are typically designed for a single user. Enterprise software supports multiple users in an organization or work group.
- Popular categories of enterprise software include ERP (enterprise resource planning), SCM (supply chain management), CRM (customer relationship management), and BI (business intelligence) software, among many others.
- These systems are used in conjunction with database management systems, programs that help firms organize, store, retrieve, and maintain data.
- ERP and other packaged enterprise systems can be challenging and costly to implement, but can help firms create a standard set of procedures and data that can ultimately lower costs and streamline operations.
- The more application software that is available for a platform, the more valuable that platform becomes.
- The DBMS stores and retrieves the data used by the other enterprise applications. Different enterprise systems can be configured to share the same database system in order share common data.
- Firms that don't have common database systems with consistent formats across their enterprise often struggle to efficiently manage their value chain, and often lack the flexibility to introduce new ways of doing business. Firms with common database systems and standards often benefit from increased organizational insight and decision-making capabilities.
- Enterprise systems can cost millions of dollars in software, hardware, development, and consulting fees, and many firms have failed when attempting large-scale enterprise system integration. Simply buying a system does not guarantee its effective deployment and use.



- When set up properly, enterprise systems can save millions of dollars and turbocharge organizations by streamlining processes, making data more usable, and easing the linking of systems with software across the firm and with key business partners.

QUESTIONS AND EXERCISES

1. What is the difference between desktop and enterprise software?
2. Who are the two leading ERP vendors?
3. List the functions of a business that might be impacted by an ERP.
4. What do the acronyms ERP, CRM, SCM, and BI stand for? Briefly describe what each of these enterprise systems does.
5. Where in the “layer cake” analogy does the DBMS lie.
6. Name two companies that have realized multimillion-dollar benefits as result of installing enterprise systems.
7. Name two companies that have suffered multimillion-dollar disasters as result of failed enterprise system installations.
8. How much does the average large company spend annually on ERP software?

9.4 Distributed Computing

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand the concept of distributed computing and its benefits.
2. Understand the client-server model of distributed computing.
3. Know what Web services are and the benefits that Web services bring to firms.

4. Appreciate the importance of messaging standards and understand how sending messages between machines can speed processes, cut costs, reduce errors, and enable new ways of doing business.

When computers in different locations can communicate with one another, this is often referred to as distributed computing. Distributed computing can yield enormous efficiencies in speed, error reduction, and cost savings and can create entirely new ways of doing business. Designing systems architecture for distributed systems involves many advanced technical topics. Rather than provide an exhaustive decomposition of distributed computing, the examples that follow are meant to help managers understand the bigger ideas behind some of the terms that they are likely to encounter.

Let's start with the term server. This is a tricky one because it's frequently used in two ways: (1) in a hardware context a server is a computer that has been configured to support requests from other computers (e.g., Dell sells servers) and (2) in a software context a server is a program that fulfills requests (e.g., the Apache open source Web server). Most of the time, server *software* resides on server-class *hardware*, but you can also set up a PC, laptop, or other small computer to run server software, albeit less powerfully. And you can use mainframe or super-computer-class machines as servers, too.

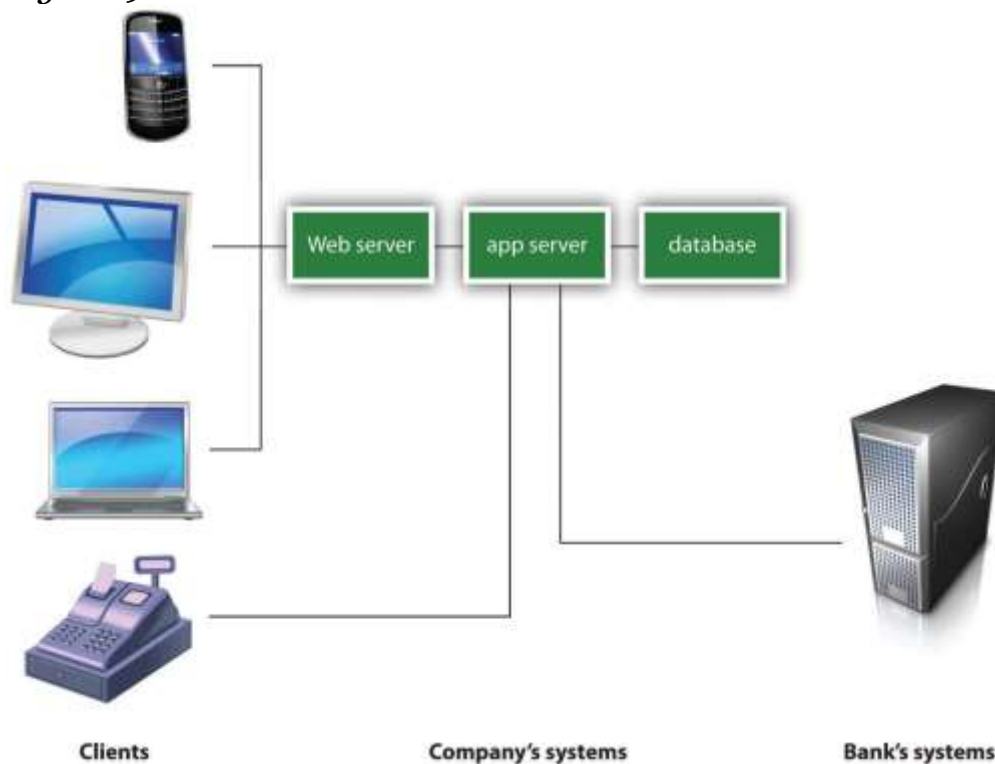
The World Wide Web, like many other distributed computing services, is what geeks call *client-server* system. Client-server refers to two pieces of software, a client that makes a request, and a server that receives and attempts to fulfill the request. In our WWW scenario, the client is the browser (e.g., Internet Explorer, Firefox, Safari). When you type a Web site's address into the location field of your browser, you're telling the client



to “go find the Web server software at the address provided, and tell the server to return the Web site requested.”

It is possible to link simple scripting languages to a Web server for performing calculations, accessing databases, or customizing Web sites. But more advanced distributed environments may use a category of software called an application server. The application server (or app server) houses business logic for a distributed system. Individual Web services served up by the app server are programmed to perform different tasks: returning a calculation (“sales tax for your order will be \$11.58”), accessing a database program (“here are the results you searched for”), or even making a request to another server in another organization (“Visa, please verify this customer’s credit card number for me”).

Figure 9.6



In this multitiered distributed system, client browsers on various machines (desktop, laptop, mobile) access the system through the Web

server. The cash register doesn't use a Web browser, so instead the cash register logic is programmed to directly access the services it needs from the app server. Web services accessed from the app server may be asked to do a variety of functions, including perform calculations, access corporate databases, or even make requests from servers at other firms (for example, to verify a customer's credit card).

Those little chunks of code that are accessed via the application server are sometimes referred to as Web services. The World Wide Web consortium defines *Web services* as software systems designed to support interoperable machine-to-machine interaction over a network. ^[1] And when computers can talk together (instead of people), this often results in fewer errors, time savings, cost reductions, and can even create whole new ways of doing business! Each Web service defines the standard method for other programs to request it to perform a task and defines the kind of response the calling client can expect back. These standards are referred to as application programming interfaces (APIs).

Look at the advantages that Web services bring a firm like Amazon. Using Web services, the firm can allow the same order entry logic to be used by Web browsers, mobile phone applications, or even by third parties who want to access Amazon product information and place orders with the firm (there's an incentive to funnel sales to Amazon—the firm will give you a cut of any sales that you send Amazon's way). Organizations that have created a robust set of Web services around their processes and procedures are said to have a service-oriented architecture (SOA). Organizing systems like this, with separate applications in charge of client presentation, business logic, and database, makes systems more flexible. Code can be reused, and each layer can be separately maintained, upgraded, or migrated to new hardware—all with little impact on the others.



Web services sound geeky, but here's a concrete example illustrating their power. Southwest Airlines had a Web site where customers could book flights, but many customers also wanted to rent a car or book a hotel, too. To keep customers on Southwest.com, the firm and its hotel and rental car partners created a set of Web services and shared the APIs. Now customers visiting Southwest.com can book a hotel stay and rental car on the same page where they make their flight reservation. This process transforms Southwest.com into a full service travel destination and allows the site to compete head-to-head with the likes of Expedia, Travelocity, and Orbitz. [2] Think about why Web services are important from a strategic perspective. By adding hotel and rental car services, Southwest is now able to eliminate the travel agent, along with any fees they might share with the agent. This shortcut allows the firm to capture more profits or pass on savings to customers, securing its position as the first place customers go for low-cost travel. And perhaps most importantly, Southwest can capture key data from visitor travel searches and bookings (something it likely couldn't do if customers went to a site like Expedia or Travelocity). Data is a hugely valuable asset, and this kind of customer data can be used by Southwest to send out custom e-mail messages and other marketing campaigns to bring customers back to the airline. As geeky as they might at first seem, Web services can be very strategic!

Messaging Standards

Two additional terms you might hear within the context of distributed computing are EDI and XML. EDI (electronic data interchange) is a set of standards for exchanging information between computer applications. EDI is most often used as a way to send the electronic equivalent of structured documents between different organizations. Using EDI, each element in the electronic document, like a firm name, address, or customer number, is

coded so that it can be recognized by the receiving computer program. Eliminating paper documents makes businesses faster and lowers data entry and error costs. One study showed that firms that used EDI decreased their error rates by 82 percent and their cost of producing each document fell by up to 96 percent. [3]

EDI is a very old standard, with roots stretching back to the 1948 Berlin Air Lift. While still in use, a new generation of more-flexible technologies for specifying data standards are taking its place. Chief among the technologies replacing EDI is extensible markup language (XML). XML has lots of uses, but in the context of distributed systems, it allows software developers to create a set of standards for common data elements that, like EDI messages, can be sent between different kinds of computers, different applications, and different organizations. XML is often thought of as easier to code than EDI, and it's more robust because it can be extended—organizations can create formats to represent any kind of data (e.g., a common part number, photos, the complaint field collected by customer support personal). In fact, most messages sent between Web services are coded in XML (the technology is a key enabler in mashups, discussed in [Chapter 6 "Peer Production, Social Media, and Web 2.0"](#)). Many computer programs also use XML as a way to export and import data in a common format that can be used regardless of the kind of computer hardware, operating system, or application program used. And if you design Web sites, you might encounter XML as part of the coding behind the cascading style sheets (CSS) that help maintain a consistent look and feel to the various Web sites in a given Web site.



Rearden Commerce: A Business Built on Web Services

Web services, APIs, and open standards not only transform businesses, they can create entire new firms that change how we get things done. For a look at the mashed-up, integrated, hyperautomated possibilities that Web services make possible, check out Rearden Commerce, a Foster City, California, firm that is using this technology to become what AMR's Chief Research Office referred to as "Travelocity on Steroids."

Using Rearden, firms can offer their busy employees a sort of Web-based concierge/personal assistant. Rearden offers firms a one-stop shop where employees can not only make the flight, car, and hotel bookings they might do from a travel agent, they can also book dinner reservations, sports and theatre tickets, and arrange for business services like conference calls and package shipping. Rearden doesn't supply the goods and services it sells. Instead it acts as the middleman between transactions. A set of open APIs to its Web services allows Rearden's one hundred and sixty thousand suppliers to send product and service data to Rearden, and to receive booking and sales data from the site.

In this ultimate business mashup, a mobile Rearden user could use her phone to book a flight into a client city, see restaurants within a certain distance of her client's office, have these locations pop up on a Google map, have listings accompanied by Zagat ratings and cuisine type, book restaurant reservations through Open Table, arrange for a car and driver to meet her at her client's office at a specific time, and sync up these reservations with her firm's corporate calendaring systems. If something unexpected comes up, like a flight delay, Rearden will be sure she gets the message. The system will keep track of any cancelled reservation credits, and also records travel reward programs, so Rearden can be used to spend those points in the future.



In order to pull off this effort, the Rearden maestros are not only skilled at technical orchestration, but also in coordinating customer and supplier requirements. *AsTechCrunch's* Erick Schonfeld put it, “The hard part is not only the technology—which is all about integrating an unruly mess of APIs and Web services—[it also involves] signing commercially binding service level agreements with [now over 160,000] merchants across the world.” For its efforts, Rearden gets to keep between 6 percent and 25 percent of every nontravel dollar spent, depending on the service. The firm also makes money from subscriptions, and distribution deals.

The firm’s first customers were large businesses and included ConAgra, GlaxoSmithKline, and Motorola. Rearden’s customers can configure the system around special parameters unique to each firm: to favor a specific airline, benefit from a corporate discount, or to restrict some offerings for approved employees only. Rearden investors include JPMorgan Chase and American Express—both of whom offer Rearden to their employees and customers. Even before the consumer version was available, Rearden had over four thousand corporate customers and two million total users, a user base larger than better-known firms like Salesforce.com. ^[4]

Connectivity has made our systems more productive and enables entire new strategies and business models. But these wonderful benefits come at the price of increased risk. When systems are more interconnected, opportunities for infiltration and abuse also increase. Think of it this way—each “connection” opportunity is like adding another door to a building. The more doors that have to be defended, the more difficult security becomes. It should be no surprise that the rise of the Internet and distributed computing has led to an explosion in security losses by organizations worldwide.



KEY TAKEAWAYS

- Client-server computing is a method of distributed computing where one program (a client) makes a request to be fulfilled by another program (a server).
- Server is a tricky term and is sometimes used to refer to hardware. While server-class hardware refers to more powerful computers designed to support multiple users, just about any PC or notebook can be configured to run server software.
- Web servers serve up Web sites and can perform some scripting.
- Most firms serve complex business logic from an application server.
- Isolating a system's logic in three or more layers (presentation or user interface, business logic, and database) can allow a firm flexibility in maintenance, reusability, and in handling upgrades.
- Web services allow different applications to communicate with one another. APIs define the method to call a Web service (e.g., to get it to do something), and the kind of response the calling program can expect back.
- Web services make it easier to link applications as distributed systems, and can make it easier for firms to link their systems across organizations.
- Popular messaging standards include EDI (older) and XML. Sending messages between machines instead of physical documents can speed processes, drastically cut the cost of transactions, and reduce errors.
- Distributed computing can yield enormous efficiencies in speed, error reduction, and cost savings and can create entirely new ways of doing business.
- When computers can communicate with each other (instead of people), this often results in fewer errors, time savings, cost reductions, and can even create whole new ways of doing business.



- Web services, APIs, and open standards not only transform businesses, they can create entire new firms that change how we get things done.

QUESTIONS AND EXERCISES

1. Differentiate the term “server” used in a hardware context, from “server” used in a software context.
2. Describe the “client-server” model of distributed computing. What products that you use would classify as leveraging client-server computing?
3. List the advantages that Web services have brought to Amazon.
4. How has Southwest Airlines utilized Web services to its competitive advantage?
5. What is Rearden Commerce and which technologies does it employ? Describe Rearden Technology’s revenue model. Who were Rearden Technology’s first customers? Who were among their first investors?
6. What are the security risks associated with connectivity, the Internet, and distributed processing?



9.5 Writing Software

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand, at a managerial level, what programming languages are and how software is developed.
2. Recognize that an operating system and microprocessor constrain the platform upon which most compiled application software will run.
3. Understand what Java is and why it is significant.
4. Know what scripting languages are.

So you've got a great idea that you want to express in software—how do you go about creating a program?

Programmers write software in a programming language. While each language has its strengths and weaknesses, most commercial software is written in C++ (pronounced “see plus plus”) or C# (pronounced “see sharp”). Visual Basic (from Microsoft) and Java (from Sun) are also among the more popular of the dozens of programming languages available. Web developers may favor specialty languages like Ruby and Python, while languages like SQL are used in databases.

Most professional programmers use an integrated development environment (IDE) to write their code. The IDE includes a text editor, a debugger for sleuthing out errors, and other useful programming tools. The most popular IDE for Windows is Visual Studio, while Apple offers the Xcode IDE. Most IDEs can support several different programming languages. The IDE will also compile a programmer's code,



turning the higher-level lines of instructions that are readable by humans into lower-level instructions expressed as the patterns of ones and zeros that are readable by a computer's microprocessor.

Look at the side of a box of commercial software and you're likely to see system requirements that specify the operating system and processor that the software is designed for (e.g., "this software works on computers with Windows 7 and Intel-compatible processors"). Wouldn't it be great if software could be written once and run everywhere? That's the idea behind Java—a programming language developed by Sun Microsystems. Java programmers don't write code with specific operating system commands (say for Windows, Mac OS X, or Linux), instead they use special Java commands to control their user interface or interact with the display and other hardware. Java programs can run on any computer that has a Java Virtual Machine (JVM), a software layer that interprets Java code so that it can be understood by the operating system and processor of a given computer. Java's platform independence—the ability for developers to "write once, run everywhere"—is its biggest selling point. Many Web sites execute Java applets to run the animation you might see in advertisements or game. Java has also been deployed on over six billion mobile phones worldwide, and is popular among enterprise programmers who want to be sure their programs can scale from smaller hardware up to high-end supercomputers. As long as the machine receiving the Java code has a JVM, then the Java application should run. However, Java has not been popular for desktop applications. Since Java isn't optimized to take advantage of interface elements specific to the Mac or Windows, most Java desktop applications look clunky and unnatural. Java code that runs through the JVM interpreter is also slower than code compiled for the native OS and processor that make up a platform. ^[1]



Scripting languages are the final category of programming tool that we'll cover. Scripting languages typically execute within an application. Microsoft offers a scripting language called VB Script (a derivative of Visual Basic) to automate functions in Office. And most browsers and Web servers support JavaScript, a language that helps make the Web more interactive (despite its name, JavaScript is unrelated to Java). Scripting languages are interpreted within their applications, rather than compiled to run directly by a microprocessor. This distinction makes them slower than the kinds of development efforts found in most commercial software. But most scripting languages are usually easy to use, and are often used both by professional programmers and power users.

KEY TAKEAWAYS

- Programs are often written in a tool called an IDE, an application that includes an editor (a sort of programmer's word processor), debugger, and compiler, among other tools.
- Compiling takes code from the high-level language that humans can understand and converts them into the sets of ones and zeros in patterns representing instructions that microprocessors understand.
- Popular programming languages include C++, C#, Visual Basic, and Java.
- Most software is written for a platform—a combination of an operating system and microprocessor.
- Java is designed to be platform independent. Computers running Java have a separate layer called a Java Virtual Machine that translates (interprets) Java code so that it can be executed on an operating system/processor combination. In theory, Java is “write once, run everywhere,” as opposed to conventional applications that are written for an operating system and compiled for an OS/processor combination.



- Java is popular on mobile phones, enterprise computing, and to make Web sites more interactive. Java has never been a successful replacement for desktop applications, largely because user interface differences among the various operating systems are too great to be easily standardized.
- Scripting languages are interpreted languages, such as VB Script or Java Script. Many scripting languages execute within an application (like the Office programs, a Web browser, or to support the functions of a Web server). They are usually easier to program, but are less powerful and execute more slowly than compiled languages.

QUESTIONS AND EXERCISES

1. List popular programming languages.
2. What's an IDE? Why do programmers use IDEs? Name IDEs popular for Windows and Mac users.
3. What is the difference between a compiled programming language and an interpreted programming language?
4. Name one advantage and one disadvantage of scripting languages.
5. In addition to computers, on what other technology has Java been deployed? Why do you suppose Java is particularly attractive for these kinds of applications?
6. What's a JVM? Why do you need it?
7. What if a programmer wrote perfect Java code, but there was a bug on the JVM installed on a given computer? What might happen?
8. Why would developers choose to write applications in Java? Why might they skip Java and choose another programming language?
9. Why isn't Java popular for desktop applications?
10. Go to <http://www.java.com>. Click on "Do I have Java?" Is Java running on your computer? Which version?



9.6 Total Cost of Ownership (TCO): Tech Costs Go Way beyond the Price Tag

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. List the different cost categories that comprise total cost of ownership.



2. Understand that once a system is implemented, the costs of maintaining and supporting the system continue.
3. List the reasons that technology development projects fail and the measures that can be taken to increase the probability of success.

Managers should recognize that there are a whole host of costs that are associated with creating and supporting an organization's information systems. Of course, there are programming costs for custom software as well as purchase, configuration, and licensing costs for packaged software, but there's much, much more.

There are costs associated with design and documentation (both for programmers and for users). There are also testing costs. New programs should be tested thoroughly across the various types of hardware the firm uses, and in conjunction with existing software and systems, *before* being deployed throughout the organization. Any errors that aren't caught can slow down a business or lead to costly mistakes that could ripple throughout an organization and its partners. Studies have shown that errors not caught before deployment could be one hundred times more costly to correct than if they were detected and corrected beforehand. ^[1] Once a system is "turned on," the work doesn't end there. Firms need to constantly engage in a host of activities to support the system that may also include the following:

- providing training and end user support;
- collecting and relaying comments for system improvements;
- auditing systems to ensure compliance (i.e., that the system operates within the firm's legal constraints and industry obligations);
- providing regular backup of critical data;
- planning for redundancy and disaster recovery in case of an outage; and
- vigilantly managing the moving target of computer security issues.



With so much to do, it's no wonder that firms spend 70 to 80 percent of their information systems (IS) budgets just to keep their systems running. [2] The price tag and complexity of these tasks can push some managers to think of technology as being a cost sink rather than a strategic resource. These tasks are often collectively referred to as the total cost of ownership (TCO) of an information system. Understanding TCO is critical when making technology investment decisions. TCO is also a major driving force behind the massive tech industry changes discussed in [Chapter 10 "Software in Flux: Partly Cloudy and Sometimes Free"](#).

Why Do Technology Projects Fail?

Even though information systems represent the largest portion of capital spending at most firms, an astonishing one in three technology development projects fail to be successfully deployed. [3] Imagine if a firm lost its investment in one out of every three land purchases, or when building one in three factories. These statistics are dismal! Writing in *IEEE Spectrum*, risk consultant Robert Charette provides a sobering assessment of the cost of software failures, stating, “The yearly tab for failed and troubled software conservatively runs somewhere from sixty to seventy billion dollars in the United States alone. For that money, you could launch the space shuttle one hundred times, build and deploy the entire 24-satellite Global Positioning System, and develop the Boeing 777 from scratch—and still have a few billion left over.” [4]

Why such a bad track record? Sometimes technology itself is to blame, other times it's a failure to test systems adequately, and sometimes it's a breakdown of process and procedures used to set specifications and manage projects. In one example, a multimillion-dollar loss on the NASA Mars Observer was traced back to a laughably simple oversight—Lockheed Martin contractors using English measurements, while the folks at NASA

used the metric system. [5] Yes, a \$125 million taxpayer investment was lost because a bunch of rocket scientists failed to pay attention to third grade math. When it comes to the success or failure of technical projects, the devil really is in the details.

Projects rarely fail for just one reason. Project post-mortems often point to a combination of technical, project management, and business decision blunders. The most common factors include the following: [6]

- Unrealistic or unclear project goals
- Poor project leadership and weak executive commitment
- Inaccurate estimates of needed resources
- Badly defined system requirements and allowing “feature creep” during development
- Poor reporting of the project’s status
- Poor communication among customers, developers, and users
- Use of immature technology
- Unmanaged risks
- Inability to handle the project’s complexity
- Sloppy development and testing practices
- Poor project management
- Stakeholder politics
- Commercial pressures (e.g., leaving inadequate time or encouraging corner-cutting)

Managers need to understand the complexity involved in their technology investments, and that achieving success rarely lies with the strength of the technology alone.

But there is hope. Information Systems organizations can work to implement procedures to improve the overall quality of their development practices. Mechanisms for quality improvement include capability maturity model integration (CMMI), which gauge an



organization's process maturity and capability in areas critical to developing and deploying technology projects, and provides a carefully chosen set of best practices and guidelines to assist quality and process improvement. [7]

Firms are also well served to leverage established project planning and software development methodologies that outline critical businesses processes and stages when executing large-scale software development projects. The idea behind these methodologies is straightforward—why reinvent the wheel when there is an opportunity to learn from and follow blueprints used by those who have executed successful efforts. When methodologies are applied to projects that are framed with clear business goals and business metrics, and that engage committed executive leadership, success rates can improve dramatically. [8]

While software development methodologies are the topic of more advanced technology courses, the savvy manager knows enough to inquire about the development methodologies and quality programs used to support large scale development projects, and can use these investigations as further input when evaluating whether those overseeing large scale efforts have what it takes to get the job done.

KEY TAKEAWAYS

- The care and feeding of information systems can be complex and expensive. The total cost of ownership of systems can include software development and documentation, or the purchase price and ongoing license and support fees, plus configuration, testing, deployment, maintenance, support, training, compliance auditing, security, backup, and provisions for disaster recovery. These costs are collectively referred to as TCO, or a system's total cost of ownership.



- Information systems development projects fail at a startlingly high rate. Failure reasons can stem from any combination of technical, process, and managerial decisions.
- IS organizations can leverage software development methodologies to improve their systems development procedures, and firms can strive to improve the overall level of procedures used in the organization through models like CMMI. However, it's also critical to engage committed executive leadership in projects, and to frame projects using business metrics and outcomes to improve the chance of success.
- System errors that aren't caught before deployment can slow down a business or lead to costly mistakes that could ripple throughout an organization. Studies have shown that errors not caught before deployment could be 100 times more costly to correct than if they were detected and corrected beforehand.
- Firms spend 70 to 80 percent of their IS budgets just to keep their systems running.
- One in three technology development projects fail to be successfully deployed.
- IS organizations can employ project planning and software development methodologies to implement procedures to improve the overall quality of their development practices.

QUESTIONS AND EXERCISES

1. List the types of total ownership costs associated with creating and supporting an organization's information systems.
2. On average, what percent of firms' IS budgets is spent to keep their systems running?
3. What are the possible effects of not detecting and fixing major system errors before deployment?



4. List some of the reasons for the failure of technology development projects.
5. What is the estimated yearly cost of failed technology development projects?
6. What was the reason attributed to the failure of the NASA Mars Observer project?
7. What is capability maturity model integration (CMMI) and how is it used to improve the overall quality of a firm's development practices?
8. Perform an Internet search for "IBM Rational Portfolio Manager." How might IBM's Rational Portfolio Manager software help companies realize more benefit from their IT systems development project expenditures? What competing versions of this product offered by other organizations?



Chapter 10

Software in Flux: Partly Cloudy and Sometimes Free

10.1 Introduction

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand how low marginal costs, network effects, and switching costs have combined to help create a huge and important industry.
2. Recognize that the software industry is undergoing significant and broadly impactful change brought about by several increasingly adopted technologies including open source software, cloud computing, and software-as-a-service.

For many, software has been a magnificent business. It is the two-hundred-billion-dollar-per-year juggernaut ^[1] that placed Microsoft's Bill Gates and Oracle's Larry Ellison among the wealthiest people in the world. Once a successful software product has been written, the economics for a category-leading offering are among the best you'll find in any industry. Unlike physical products assembled from raw materials, the marginal cost to produce an additional copy of a software product is effectively zero. Just



duplicate, no additional input required. That quality leads to businesses that can gush cash. Microsoft generates a one and a half billion dollars a month from Windows and Office, alone. [2] Network effects and switching cost can also offer a leading software firm a degree of customer preference and lock in that can establish a firm as a standard, and in many cases creates winner-take-all (or at least winner-take-most) markets.

But as great as the business has been, the fundamental model powering the software industry is under assault. Open source software (OSS) offerings—free alternatives where anyone can look at and potentially modify a program’s code—pose a direct challenge to the assets and advantages cultivated by market leaders. Giants shudder—“How can we compete with free,” while others wonder, “How can we make money and fuel innovation on free?” And if free software wasn’t enough of a shock, the way firms and users think about software is also changing. A set of services referred to as cloud computing is making it more common for a firm to move software out of its own IS shop so that it is run on someone else’s hardware. In one variant of this approach known as software as a service (SaaS), users access a *vendor’s software* over the Internet, usually by simply starting up a Web browser. With SaaS, you don’t need to own the program or install it on your own computer. Hardware clouds can let firms take *their software* and run it on someone else’s hardware—freeing them from the burden of buying, managing, and maintaining the physical computing that programs need. Another software technology called virtualization can make a single computer behave like many separate machines. This function helps consolidate computing resources and creates additional savings and efficiencies.

These transitions are important. They mean that smaller firms have access to the kinds of burly, sophisticated computing power than only giants had access to in the past. Startups can scale quickly and get up and running with



less investment capital. Existing firms can leverage these technologies to reduce costs. Got tech firms in your investment portfolio? Understanding what's at work here can inform decisions you make on which stocks to buy or sell. If you make tech decisions for your firm or make recommendations for others, these trends may point to which firms have strong growth and sustainability ahead, or which may be facing troubled times.

KEY TAKEAWAYS

- The software business is attractive due to near-zero marginal costs and an opportunity to establish a standard—creating the competitive advantages of network effects and switching costs.
- New trends in the software industry, including open source software (OSS), hardware clouds, software as a service (SaaS), and virtualization are creating challenges and opportunity across tech markets. Understanding the impact of these developments can help a manager make better technology choices and investment decisions.

QUESTIONS AND EXERCISES

1. What major trends, outlined in the section above, are reshaping how we think about software? What industries and firms are potentially impacted by these changes? Why do managers, investors, and technology buyers care about these changes?
2. Which organizations might benefit from these trends? Which might be threatened? Why?
3. What are marginal costs? Are there other industries that have cost economics similar to the software industry?
4. Investigate the revenues and net income of major software players: Microsoft, Google, Oracle, Red Hat, and Salesforce.com. Which firms have higher revenues? Net income? Which have better margins? What

do the trends in OSS, SaaS, and cloud computing suggest for these and similar firms?

5. How might the rise of OSS, SaaS, and cloud computing impact hardware sales? How might it impact entrepreneurship and smaller businesses?



10.2 Open Source

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Define open source software and understand how it differs from conventional offerings.
2. Provide examples of open source software and how firms might leverage this technology.

Who would have thought a twenty-one-year-old from Finland could start a revolution that continues to threaten the Microsoft Windows empire? But Linus Torvalds did just that. During a marathon six-month coding session, Torvalds created the first version of Linux^[1] marshalling open-source revolutionaries like no one before him. Instead of selling his operating system, Torvalds gave it away. Now morphed and modified into scores of versions by hundreds of programmers, Linux can be found just about everywhere, and most folks credit Linux as being the most significant product in the OSS arsenal. Today Linux powers everything from cell phones to stock exchanges, set top boxes to supercomputers. You'll find the OS on 30 percent of the servers in corporate America, ^[2]and supporting most Web servers (including those at Google, Amazon, and Facebook). Linux forms the core of the TiVo operating system, it underpins Google's Android and Chrome OS offerings, and it has even gone interplanetary. Linux has been used to power the Phoenix Lander and to control the Spirit and Opportunity Mars rovers. ^[3] Yes, Linux is even on Mars!

How Do You Pronounce Linux?

Most English speakers in the know pronounce Linux in a way that rhymes with “cynics.” You can easily search online to hear video and audio clips of



Linus (whose name is actually pronounced “Lean-us” in Finish) pronouncing the name of his OS. In deference to Linux, some geeks prefer something that sounds more like “*lean-ooks*.” [4] Just don’t call it “*line-uks*,” or the tech-savvy will think you’re an open source noob! Oh yeah, and while we’re on the topic of operating system pronunciation, the Macintosh operating system OS X is pronounced “oh es ten.”

Open source software (OSS) is often described as free. While most OSS can be downloaded for free over the Internet, it’s also “free” as in liberated. The source code for OSS products is openly shared. Anyone can look at the source code, change it, and even redistribute it, provided the modified software continues to remain open and free. [5] This openness is in stark contrast to the practice of conventional software firms, who treat their intellectual property as closely guarded secrets, and who almost never provide the source code for their commercial software products. At times, many software industry execs have been downright hostile toward OSS. The former President of SAP once referred to the open source movement as “socialism,” while Microsoft’s Steve Balmer has called Linux a “cancer.” [6] But while execs at some firms see OSS as a threat undermining the lifeblood of their economic model, other big-name technology companies are now solidly behind the open source movement. The old notion of open source being fueled on the contributions of loners tooling away for the glory of contributing to better code is now largely inaccurate. The vast majority of people who work on efforts like Linux are now paid to do so by commercially motivated employers. [7] Nearly every major hardware firm has paid staff contributing to open source projects, and most firms also work together to fund foundations that set standards and coordinate the release of product revisions and improvements. Such coordination is critical—helping, for example, to ensure that various versions of Linux work alike. Sun Microsystems claims to have eleven thousand engineers



contributing to OSS. [8] Guido van Rossum, the inventor of the open source Python programming language, works for Google where he continues to coordinate development. IBM programmers work on several open source projects, including Linux. The firm has even deeded a commercially developed programming tool (including an IDE) to the Eclipse foundation, where it's now embraced and supported by dozens of firms.

Turn on the LAMP—It's Free!

Open source is big on the Web. In fact, you'll often hear Web programmers and open source advocates refer to the LAMP stack. LAMP is an acronym that stands for the Linux operating system, the Apache Web server software, the MySQL database, and any of several programming languages that start with the letter "P"—Perl, Python, and PHP. From Facebook to YouTube, you'll find LAMP software powering many of the sites you visit each day.

KEY TAKEAWAYS

- OSS is not only available for free, but also makes source code available for review and modification (for the Open Source Initiatives list of the criteria that define an open source software product, see <http://opensource.org/docs/osd>).
- While open source alternatives are threatening to conventional software firms, some of the largest technology companies now support OSS initiatives and work to coordinate standards, product improvements, and official releases.
- The flagship OSS product is the Linux operating system, now available on all scales of computing devices from cell phones to supercomputers.
- The LAMP stack of open source products is used to power many of the Internet's most popular Web sites. Linux can be found on 30 percent of



corporate servers, supports most Web servers, and is integral to TiVo and Android-based cell phones.

- The majority of persons who work on open source projects are paid by commercially motivated employers.

QUESTIONS AND EXERCISES

1. Who developed Linux?
2. Who develops it today?
3. List the components of the LAMP stack. Which commercial products do these components compete with (investigate online, if necessary)?
4. Why do commercial firms contribute to open source consortia and foundations?
5. Free doesn't always win. Why might a firm turn down free software in favor of a commercial alternative?

10.3 Why Open Source?

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Know the primary reasons firms choose to use OSS.
2. Understand how OSS can beneficially impact industry and government.

There are many reasons why firms choose open source products over commercial alternatives:

Cost—Free alternatives to costly commercial code can be a tremendous motivator, particularly since conventional software often requires customers to pay for every copy used and to pay more for software that runs on increasingly powerful hardware. Big Lots stores lowered costs by as much as ten million dollars by finding viable OSS [1] to serve their system needs. Online broker E*TRADE estimates that its switch to open source helped save over thirteen million dollars a year. [2] And Amazon claimed in SEC filings that the switch to open source was a key contributor to nearly twenty million dollars in tech savings. [3] Firms like TiVo, which use OSS in their own products, eliminate a cost spent either developing their own operating system or licensing similar software from a vendor like Microsoft.

Reliability—There’s a saying in the open source community, “Given enough eyeballs, all bugs are shallow.” [4] What this means is that the more people who look at a program’s code, the greater the likelihood that an error will be caught and corrected. The open source community harnesses the power of legions of geeks who are constantly trawling OSS products, looking to squash bugs and improve product quality. And studies have shown that the quality of popular OSS products outperforms proprietary commercial competitors.[5] In one study, Carnegie Mellon University’s Cylab estimated the quality of Linux code to be less buggy than commercial alternatives by a factor of two hundred! [6]

Security—OSS advocates also argue that by allowing “many eyes” to examine the code, the security vulnerabilities of open source products come to light more quickly and can be addressed with greater speed and reliability. [7] High profile hacking contests have frequently demonstrated the strength of OSS products. In one well-publicized 2008 event, laptops running Windows and Macintosh were both hacked (the latter in just two minutes), while a laptop running Linux remained uncompromised. [8] Government agencies and the military often appreciate

the opportunity to scrutinize open source efforts to verify system integrity (a particularly sensitive issue among foreign governments leery of legislation like the USA PATRIOT Act of 2001). [9] Many OSS vendors offer “security focused” (sometimes called “hardened”) versions of their products. These can include systems that monitor the integrity of an OSS distribution, checking file size and other indicators to be sure that code has not been modified and redistributed by bad guys who’ve added a back door, malicious routines, or other vulnerabilities.

Scalability—Many major OSS efforts can run on everything from cheap commodity hardware to high-end supercomputing. Scalability allows a firm to scale from startup to blue chip without having to significantly rewrite their code, potentially saving big on software development costs. Not only can many forms of OSS be migrated to more powerful hardware, packages like Linux have also been optimized to balance a server’s workload among a large number of machines working in tandem. Brokerage firm E*TRADE claims that usage spikes following 2008 U.S. Federal Reserve moves flooded the firm’s systems, creating the highest utilization levels in five years. But E*TRADE credits its scalable open source systems for maintaining performance while competitors’ systems struggled. [10]

Agility and Time to Market—Vendors who use OSS as part of product offerings may be able to skip whole segments of the software development process, allowing new products to reach the market faster than if the entire software system had to be developed from scratch, in-house. Motorola has claimed that customizing products built on OSS has helped speed time-to-market for the firm’s mobile phones, while the team behind the Zimbra e-mail and calendar effort built their first product in just a few months by using some forty blocks of free code. [11]



KEY TAKEAWAYS

- The most widely cited benefits of using OSS include low cost; increased reliability; improved security and auditing; system scalability; and helping a firm improve its time to market.
- Free OSS has resulted in cost savings for many large companies in several industries.
- OSS often has fewer bugs than its commercial counterparts due to the large number of persons who have looked at the code.
- The huge exposure to scrutiny by developers and other people helps to strengthen the security of OSS.
- “Hardened” versions of OSS products often include systems that monitor the integrity of an OSS distribution, checking file size and other indicators to be sure that code has not been modified and redistributed by bad guys who have added a back door, malicious routines, or other vulnerabilities.
- OSS can be easily migrated to more powerful computers as circumstances dictate, and also can balance workload by distributing work over a number of machines.
- Vendors who use OSS as part of product offerings may be able to skip whole segments of the software development process, allowing new products to reach the market faster.

QUESTIONS AND EXERCISES

1. What advantages does OSS offer TiVo? What alternatives to OSS might the firm consider and why do you suppose the firm decided on OSS?
2. What’s meant by the phrase, “Given enough eyeballs, all bugs are shallow”? Provide evidence that the insight behind this phrase is an accurate one.
3. How has OSS benefited E*TRADE? Amazon? Motorola? Zimbra? What benefits were achieved in each of these examples?



4. Describe how OSS provides a firm with scalability. What does this mean, and why does this appeal to a firm? What issues might a firm face if chosen systems aren't scalable?
5. The Web site NetCraft (<http://www.netcraft.com>) is one of many that provide a tool to see the kind of operating system and Web server software that a given site is running. Visit NetCraft or a similar site and enter the address of some of your favorite Web sites. How many run open source products (e.g., the Linux OS or Apache Web server)? Do some sites show their software as "unknown"? Why might a site be reluctant to broadcast the kind of software that it uses?

10.4 Examples of Open Source Software

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Recognize that just about every type of commercial product has an open source equivalent.
2. Be able to list commercial products and their open source competitors.

Just about every type of commercial product has an open source equivalent. SourceForge.net lists over two hundred and thirty thousand such products! ^[1] Many of these products come with the installation tools, support utilities, and full documentation that make them difficult to distinguish from traditional commercial efforts. ^[2] In addition to the LAMP products, some major examples include the following:

- Firefox—a Web browser that competes with Internet Explorer;
- OpenOffice—a competitor to Microsoft Office;
- Gimp—a graphic tool with features found in Photoshop;



- Alfresco—collaboration software that competes with Microsoft Sharepoint and EMC’s Documentum;
- Marketcetera—an enterprise trading platform for hedge fund managers that competes with FlexTrade and Portware;
- Zimbra—open source e-mail software that competes with Outlook server;
- MySQL, Ingres, and EnterpriseDB—open source database software packages that each go head-to-head with commercial products from Oracle, Microsoft, Sybase, and IBM;
- SugarCRM—customer relationship management software that competes with Salesforce.com and Siebel;
- Asterix—an open source implementation for running a PBX corporate telephony system that competes with offerings from Nortel and Cisco, among others; and
- Free BSD and Sun’s OpenSolaris—open source versions of the Unix operating system.

KEY TAKEAWAYS

- There are thousands of open source products available, covering nearly every software category. Many have a sophistication that rivals commercial software products.
- Not all open source products are contenders. Less popular open source products are not likely to attract the community of users and contributors necessary to help these products improve over time (again we see network effects are a key to success—this time in determining the quality of an OSS effort).
- Just about every type of commercial product has an open source equivalent.



QUESTIONS AND EXERCISES

1. Visit <http://www.SourceForge.net>. Make a brief list of commercial product categories that an individual or enterprise might use. Are there open source alternatives for these categories? Are well-known firms leveraging these OSS offerings? Which commercial firms do they compete with?
2. Are the OSS efforts you identified above provided by commercial firms, nonprofit organizations, or private individuals? Does this make a difference in your willingness to adopt a particular product? Why or why not? What other factors influence your adoption decision?
3. Download a popular, end-user version of an OSS tool that competes with a desktop application that you own, or that you've used (hint: choose something that's a smaller file or easy to install). What do you think of the OSS offering compared to the commercial product? Will you continue to use the OSS product? Why or why not?



10.5 Why Give It Away? The Business of Open Source

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand the disproportional impact OSS has on the IT market.
2. Understand how vendors make money on open source.
3. Know what SQL and MySQL are.

Open source is a sixty-billion-dollar industry, ^[1] but it has a disproportionate impact on the trillion-dollar IT market. By lowering the cost of computing, open source efforts make more computing options accessible to smaller firms. More reliable, secure computing also lowers costs for all users. OSS also diverts funds that firms would otherwise spend on fixed costs, like operating systems and databases, so that these funds can be spent on innovation or other more competitive initiatives. Think about Google, a firm that some estimate has over 1.4 million servers. Imagine the costs if it had to license software for each of those boxes!

Commercial interest in OSS has sparked an acquisition binge. Red Hat bought open source application server firm JBoss for three hundred fifty million dollars. Novell snapped up SUSE Linux for two hundred ten million dollars. And Sun plunked down over one billion dollars for open source database provider MySQL. ^[2] And with Oracle's bid for Sun, one of the



world's largest commercial software firms has zeroed in on one of the deepest portfolios of open source products.

But how do *vendors* make money on open source? One way is by selling support and consulting services. While not exactly Microsoft money, Red Hat, the largest purely OSS firm, reported half a billion dollars in revenue in 2008. The firm had two and a half million *paid* subscriptions offering access to software updates and support services. [3] Oracle, a firm that sells commercial ERP and database products, provides Linux for free, selling high-margin Linux support contracts for as much as five hundred thousand dollars. [4] The added benefit for Oracle? Weaning customers away from Microsoft—a firm that sells many products that compete head-to-head with Oracle's offerings. Service also represents the most important part of IBM's business. The firm now makes more from services than from selling hardware and software. [5] And every dollar saved on buying someone else's software product means more money IBM customers can spend on IBM computers and services. Sun Microsystems was a leader in OSS, even before the Oracle acquisition bid. The firm has used OSS to drive advanced hardware sales, but the firm also sells proprietary products that augment its open source efforts. These products include special optimization, configuration management, and performance tools that can tweak OSS code to work its best. [6]

Here's where we also can relate the industry's evolution to what we've learned about standards competition in our earlier chapters. In the pre-Linux days, nearly every major hardware manufacturer made its own, incompatible version of the Unix operating system. These fractured, incompatible markets were each so small that they had difficulty attracting third-party vendors to write application software. Now, much to Microsoft's dismay, all major hardware firms run Linux. That means there's a large,



unified market that attracts software developers who might otherwise write for Windows.

To keep standards unified, several Linux-supporting hardware and software firms also back the Linux Foundation, the nonprofit effort where Linus Torvalds serves as a fellow, helping to oversee Linux's evolution. Sharing development expenses in OSS has been likened to going in on a pizza together. Everyone wants a pizza with the same ingredients. The pizza doesn't make you smarter or better. So why not share the cost of a bigger pie instead of buying by the slice? [7] With OSS, hardware firms spend less money than they would in the brutal, head-to-head competition where each once offered a "me too" operating system that was incompatible with rivals but offered little differentiation. Hardware firms now find their technical talent can be deployed in other value-added services mentioned above: developing commercial software add-ons, offering consulting services, and enhancing hardware offerings.

Linux on the Desktop?

While Linux is a major player in enterprise software, mobile phones, and consumer electronics, the Linux OS can only be found on a tiny fraction of desktop computers. There are several reasons for this. Some suggest Linux simply isn't as easy to install and use as Windows or the Mac OS. This complexity can raise the total cost of ownership (TCO) of Linux desktops, with additional end-user support offsetting any gains from free software. The small number of desktop users also dissuades third party firms from porting popular desktop applications over to Linux. For consumers in most industrialized nations, the added complexity and limited desktop application availability of desktop Linux just isn't worth the one to two hundred dollars saved by giving up Windows.

But in developing nations where incomes are lower, the cost of Windows can be daunting. Consider the OLPC, Nicholas Negroponte’s “one hundred dollar” laptop. An additional one hundred dollars for Windows would double the target cost for the nonprofit’s machines. It is not surprising that the first OLPC laptops ran Linux. Microsoft recognizes that if a whole generation of first-time computer users grows up without Windows, they may favor open source alternatives years later when starting their own businesses. As a result, Microsoft has begun offering low-cost versions of Windows (in some cases for as little as seven dollars) in nations where populations have much lower incomes. Microsoft has even offered a version of Windows to the backers of the OLPC. While Microsoft won’t make much money on these efforts, the low cost versions will serve to entrench Microsoft products as standards in emerging markets, staving off open source rivals and positioning the firm to raise prices years later when income levels rise.

MySQL: Turning a Ten-Billion-Dollars-a-Year Business into a One-Billion-Dollar One

Finland is not the only Scandinavian country to spawn an open source powerhouse. Uppsala Sweden’s MySQL (pronounced “my sequel”) is the “M” in the LAMP stack, and is used by organizations as diverse as FedEx, Lufthansa, NASA, Sony, UPS, and YouTube.

The “SQL” in name stands for the structured query language, a standard method for organizing and accessing data. SQL is also employed by commercial database products from Oracle, Microsoft, and Sybase. Even Linux-loving IBM uses SQL in its own lucrative DB2 commercial database product. Since all of these databases are based on the same standard, switching costs are lower, so migrating from a commercial product to MySQL’s open source alternative is relatively easy. And that spells trouble



for commercial firms. Granted, the commercial efforts offer some bells and whistles that MySQL doesn't yet have, but those extras aren't necessary in a lot of standard database use. Some organizations, impressed with MySQL's capabilities, are mandating its use on all new development efforts, attempting to cordon off proprietary products in legacy code that is maintained but not expanded.

Savings from using MySQL can be huge. The Web site PriceGrabber pays less than ten thousand dollars in support for MySQL compared to one hundred thousand to two hundred thousand dollars for a comparable Oracle effort. Lycos Europe switched from Oracle to MySQL and slashed costs from one hundred twenty thousand dollars a year to seven thousand dollars. And the travel reservation firm Sabre used open source products such as MySQL to slash ticket purchase processing costs by 80 percent. [8] MySQL does make money, just not as much as its commercial rivals. While you can download a version of MySQL over the net, the flagship product also sells for four hundred ninety-five dollars per server computer compared to a list price for Oracle that can climb as high as one hundred sixty thousand dollars. Of the roughly eleven million copies of MySQL in use, the company only gets paid for about one in a thousand. [9] Firms pay for what's free for one of two reasons: (1) for MySQL service, and (2) for the right to incorporate MySQL's code into their own products. [10] Amazon, Facebook, Gap, NBC, and Sabre pay MySQL for support; Cisco, Ericsson, HP, and Symantec pay for the rights to the code. [11] Top-level round-the-clock support for MySQL for up to fifty servers is fifty thousand dollars a year, still a fraction of the cost for commercial alternatives. Founder Marten Mickos has stated an explicit goal of the firm is "turning the \$10-billion-a-year database business into a \$1 billion one." [12]

When Sun Microsystems spent over one billion dollars to buy Mickos' MySQL in 2008, Sun CEO Jonathan Schwartz called the purchase the

“most important acquisition in the company’s history.” [13] Sun hoped the cheap database software could make the firm’s hardware offerings seem more attractive. And it looked like Sun was good for MySQL, with the product’s revenues growing 55 percent in the year after the acquisition. [14] But here’s where it gets complicated. Sun also had a lucrative business selling hardware to support commercial ERP and database software from Oracle. That put Sun and partner Oracle in a relationship where they were both competitors and collaborators (the “coopetition” or “frenemies” phenomenon mentioned in [Chapter 5 "Understanding Network Effects"](#)). Then in Spring 2009, Oracle announced it was buying Sun. Oracle CEO Larry Ellison mentioned acquiring the Java language was the crown jewel of the purchase, but as of this writing, many issues remain unknown. Will Oracle keep all of Sun or sell off some of the assets? Will international regulators allow the merger to go through without conditions? Will the firm continue to nurture MySQL and other open source products, even as this software poses a threat to its bread-and-butter database products? Will the development community continue to back MySQL as the de facto standard for open source SQL databases, or will they migrate to an alternative? Or will Oracle find the right mix of free and fee-based products and services that allow MySQL to thrive while Oracle continues to grow? The implications are serious for investors, as well as firms that have made commitments to Sun, Oracle, and MySQL products. The complexity of this environment further demonstrates why technologists need business savvy and market monitoring skills and why business folks need to understand the implications of technology and tech-industry developments.

Legal Risks and Open Source Software: A Hidden and Complex Challenge

Open source software isn’t without its risks. Competing reports cite certain open source products as being difficult to install and maintain (suggesting



potentially higher total cost of ownership, or TCO). Adopters of OSS without support contracts may lament having to rely on an uncertain community of volunteers to support their problems and provide innovative upgrades. Another major concern is legal exposure. Firms adopting OSS may be at risk if they distribute code and aren't aware of the licensing implications. Some commercial software firms have pressed legal action against the users of open source products when there is a perceived violation of software patents or other unauthorized use of their proprietary code.

For example, in 2007 Microsoft suggested that Linux and other open source software efforts violated some two hundred thirty-five of its patents. ^[15] The firm then began collecting payments and gaining access to the patent portfolios of companies that use the open source Linux operating system in their products, including Fuji, Samsung, and Xerox. Microsoft also cut a deal with Linux vendor Novell in which both firms pledged not to sue each other's customers for potential patent infringements.

Also complicating issues are the varying open source license agreements (these go by various names, such as GPL and the Apache License), each with slightly different legal provisions—many of which have evolved over time. Keeping legal with so many licensing standards can be a challenge, especially for firms that want to bundle open source code into their own products. ^[16] An entire industry has sprouted up to help firms navigate the minefield of open source legal licenses. Chief among these are products, such as those offered by the firm Black Duck, which analyze the composition of software source code and report on any areas of concern so that firms can honor any legal obligations associated with their offerings. Keeping legal requires effort and attention, even in an environment where products are allegedly “free.” This also shows that even corporate lawyers

had best geek-up if they want to prove they're capable of navigating a twenty-first-century legal environment.

KEY TAKEAWAYS

- Business models for firms in the open source industry are varied, and can include selling services, licensing OSS for incorporation into commercial products, and using OSS to fuel hardware sales.
- Many firms are trying to use OSS markets to drive a wedge between competitors and their customers.
- Linux has been very successful on mobile devices and consumer electronics, as well as on high-end server class and above computers. But it has not been as successful on the desktop. The small user base for desktop Linux makes the platform less attractive for desktop software developers. Incompatibility with Windows applications, switching costs, and other network effects-related issues all suggest that Desktop Linux has an uphill climb in more mature markets.
- MySQL is the dominant open source database software product. Adoption of the SQL standard eases some issues with migrating from commercial products to MySQL.
- OSS also has several drawbacks and challenges that limit its appeal. These include complexity of some products and a higher total cost of ownership for some products, concern about the ability of a product's development community to provide support or product improvement, and legal and licensing concerns.

QUESTIONS AND EXERCISES

1. Describe the impact of OSS on the IT market.
2. Show your understanding of the commercial OSS market. How do Red Hat, Oracle, Oracle's Sun division, and IBM make money via open source?



3. Visit Mozilla.org. Which open source products does this organization develop? Investigate how development of these efforts is financed. How does this organization differ from the ones mentioned above?
4. What is the Linux Foundation? Why is it necessary? Which firms are members, underwriting foundation efforts?
5. List the reasons why Linux is installed on only a very small fraction of desktop computers. Are there particular categories of products or users who might see Linux as more appealing than conventional operating systems? Do you think Linux's share of the desktop market will increase? Why or why not?
6. How is Microsoft combating the threat of open source software and other free tools that compete with its commercial products?
7. What is the dominant open source database software product? Which firms use this product? Why?
8. Which firm developed the leading OSS database product? Do you think it's more or less likely that a firm would switch to an OSS database instead of an OSS office suite or desktop alternative? Why or why not?
9. How has stewardship of the leading OSS database effort changed in recent years? Who oversees the effort today? What questions does this raise for the product's future? Although this book is updated regularly, current events continue to change after publication of this chapter. Investigate the current status of this effort—reaction of the developer community, continued reception of the product—and be prepared to share your findings with class.
10. List some of the risks associated with using OSS. Give examples of firms that might pass on OSS software, and explain why.



10.6 Cloud Computing: Hype or Hope?

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand the concept of cloud computing.
2. Identify the two major categories of cloud computing.

Oracle Chairman Larry Ellison, lamenting the buzzword-chasing character of the tech sector, once complained that the computer industry is more fashion-focused than even the women's clothing business. ^[1] Ellison has a point: when a technology term becomes fashionable, the industry hype machine shifts into overdrive. The technology attracts press attention,



customer interest, and vendor marketing teams scramble to label their products and services as part of that innovation. Recently, few tech trends have been more fashionable than *cloud computing*.

Like Web 2.0, trying to nail down an exact definition for cloud computing is tough. In fact, it's been quite a spectacle watching industry execs struggle to clarify the concept. HP's Chief Strategy Office "politely refused" when asked by *BusinessWeek* to define the term cloud computing. [2] Richard Stallman, founder of the Free Software Foundation said about cloud computing, "It's worse than stupidity. It's a marketing hype campaign." [3] And Larry Ellison, always ready with a sound bite, offered up this priceless quip, "Maybe I'm an idiot, but I have no idea what anyone is talking about. What is it? It's complete gibberish. It's insane." [4] Insane, maybe, but also big bucks. By year-end 2008, the various businesses that fall under the rubric of cloud computing had already accounted for an estimated thirty-six-billion-dollar market. That represents a whopping 13 percent of global software sales! [5] When folks talk about cloud computing they're really talking about replacing computing resources—either an organization's or an individual's hardware or software—with *services* provided over the Internet. The name actually comes from the popular industry convention of drawing the Internet or other computer network as a big cloud.

Cloud computing encompasses a bunch of different efforts. We'll concentrate on describing, providing examples, and analyzing the managerial implications of two separate categories of cloud computing: (1) *software as a service (SaaS)*, where a firm subscribes to a third-party software-replacing service that is delivered online, and (2) models often referred to as utility computing, *platform as a service*, or *infrastructure as a service*. Using these latter techniques, an organization develops its own systems, but runs them over the Internet on someone else's hardware. A later section on virtualization will discuss how some organizations are



developing their own private clouds, pools of computing resources that reside inside an organization and that can be served up for specific tasks as need arrives.

The benefits and risks of SaaS and the utility computing-style efforts are very similar, but understanding the nuances of each effort can help you figure out if and when the cloud makes sense for your organization. The evolution of cloud computing also has huge implications across the industry: from the financial future of hardware and software firms, to cost structure and innovativeness of adopting organizations, to the skill sets likely to be most valued by employers.

KEY TAKEAWAYS

- Cloud computing is difficult to define. Managers and techies use the term cloud computing to describe computing services provided over a network, most often commercial services provided over the Internet by a third party that can replace or offload tasks that would otherwise run on a user or organization's existing hardware or software.
- Software as a service (SaaS) refers to a third-party software-replacing service that is delivered online.
- Hardware cloud computing services replace hardware that a firm might otherwise purchase.
- Estimated to be a thirty-six-billion-dollar industry, cloud computing is reshaping software, hardware, and service markets, and is impacting competitive dynamics across industries.

QUESTIONS AND EXERCISES

1. Identify and contrast the two categories of cloud computing.
2. Define cloud computing.

10.7 The Software Cloud: Why Buy When You Can Rent?

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Know how firms using SaaS products can dramatically lower several costs associated with their information systems.
2. Know how SaaS vendors earn their money.
3. Be able to list the benefits to users that accrue from using SaaS.
4. Be able to list the benefits to vendors from deploying SaaS.

If open source isn't enough of a threat to firms that sell packaged software, a new generation of products, collectively known as SaaS, claims that you can now get the bulk of your computing done through your Web browser. Don't install software—let someone else run it for you and deliver the results over the Internet.

Software as a service (SaaS) refers to software that is made available by a third party online. You might also see the terms ASP (application service provider) or HSV (hosted software vendor) used to identify this type of offering. SaaS is potentially a very big deal. Firms using SaaS products can dramatically lower several costs associated with the care and feeding of their information systems, including software licenses, server hardware, system maintenance, and IT staff. Most SaaS firms earn money via a usage-



based pricing model akin to a monthly subscription. Others offer free services that are supported by advertising, while others promote the sale of upgraded or premium versions for additional fees.

Make no mistake, SaaS is yet another direct assault on traditional software firms. The most iconic SaaS firm is Salesforce.com, an enterprise customer relationship management (CRM) provider. This “un-software” company even sports a logo featuring the word “software” crossed out, *Ghostbusters*-style. [1]

Other enterprise-focused SaaS firms compete directly with the biggest names in software. Some of these upstarts are even backed by leading enterprise software executives. Examples include NetSuite (funded in part by Oracle’s Larry Ellison—the guy’s all over this chapter), which offers a comprehensive SaaS ERP suite; and Workday (launched by founders of Peoplesoft), which has SaaS offerings for managing human resources. Several traditional software firms have countered startups by offering SaaS efforts of their own. IBM offers a SaaS version of its Cognos business intelligence products, Oracle offers CRM On Demand, and SAP’s Business ByDesign includes a full suite of enterprise SaaS offerings. Even Microsoft has gone SaaS, with a variety of Web-based services that include CRM, Web meeting tools, collaboration, e-mail, and calendaring.

SaaS is also taking on desktop applications. Intuit has online versions of its QuickBooks, TurboTax, and Quicken finance software. Adobe has an online version of Photoshop. Google and Zoho offer office suites that compete with desktop alternatives, prompting Microsoft’s own introduction of an online version of Office. And if you store photos on Flickr or Picassa instead of your PC’s hard drive, then you’re using SaaS, too.



The Benefits of SaaS

Firms can potentially save big using SaaS. Organizations that adopt SaaS forgo the large upfront costs of buying and installing software packages. For large enterprises, the cost to license, install, and configure products like ERP and CRM systems can easily run into the hundreds of thousands or even millions of dollars. And these costs are rarely a one time fee.

Additional costs like annual maintenance contracts have also been rising as rivals fail or get bought up. Less competition among traditional firms recently allowed Oracle and SAP to raise maintenance fees to as much as 20 percent. [2]

Firms that adopt SaaS don't just save on software and hardware, either. There's also the added cost for the IT staff needed to run these systems. Forrester Research estimates that SaaS can bring cost savings of 25 to 60 percent if all these costs are factored in. [3]

There are also accounting and corporate finance implications for SaaS. Firms that adopt software as a service never actually buy a system's software and hardware, so these systems become a variable operating expense. This flexibility helps mitigate the financial risks associated with making a large capital investment in information systems. For example, if a firm pays Salesforce.com sixty-five dollars per month per user for its CRM software, it can reduce payments during a slow season with a smaller staff, or pay more during heavy months when a firm might employ temporary workers. At these rates, SaaS not only looks good to large firms, it makes very sophisticated technology available to smaller firms that otherwise wouldn't be able to afford expensive systems, let alone the IT staff and hardware required to run them.

In addition to cost benefits, SaaS offerings also provide the advantage of being highly scalable. This feature is important because many organizations



operate in environments prone to wide variance in usage. Some firms might expect systems to be particularly busy during tax time or the period around quarterly financial reporting deadlines, while others might have their heaviest system loads around a holiday season. A music label might see spikes when an artist drops a new album. Using conventional software, an organization would have to buy enough computing capacity to ensure that it could handle its heaviest anticipated workload. But sometimes these loads are difficult to predict, and if the difference between high workloads and average use is great, a lot of that expensive computer hardware will spend most of its time doing nothing. In SaaS, however, the vendor is responsible for ensuring that systems meet demand fluctuation. Vendors frequently sign a service level agreement (SLA) with their customers to ensure a guaranteed uptime and define their ability to meet demand spikes. When looking at the benefits of SaaS, also consider the potential for higher quality and service levels. SaaS firms benefit from economies of scale that not only lower software and hardware costs, but also potentially boost quality. The volume of customers and diversity of their experiences means that an established SaaS vendor is most likely an expert in dealing with all sorts of critical computing issues. SaaS firms handle backups, instantly deploy upgrades and bug fixes, and deal with the continual burden of security maintenance—all costly tasks that must be performed regularly and with care, although each offers little strategic value to firms that perform these functions themselves in-house. The breadth of a SaaS vendor's customer base typically pushes the firm to evaluate and address new technologies as they emerge, like quickly offering accessibility from mobile platforms like the Blackberry and iPhone. For all but the savviest of IT shops, an established SaaS vendor can likely leverage its scale and experience to provide better, cheaper, more reliable standard information systems than individual companies typically can.



Software developers who choose to operate as SaaS providers also realize benefits. While a packaged software company like SAP must support multiple versions of its software to accommodate operating systems like Windows, Linux, and various flavors of Unix, an SaaS provider develops, tests, deploys, and supports just one version of the software executing on its own servers.

An argument might also be made that SaaS vendors are more attuned to customer needs. Since SaaS firms run a customer's systems on their own hardware, they have a tighter feedback loop in understanding how products are used (and why they fail)—potentially accelerating their ability to enhance their offerings. And once made, enhancements or fixes are immediately available to customers the next time they log in.

SaaS applications also impact distribution costs and capacity. As much as 30 percent of the price of traditional desktop software is tied to the cost of distribution—pressing CD-ROMs, packaging them in boxes, and shipping them to retail outlets. [4] Going direct to consumers can cut out the middleman, so vendors can charge less or capture profits that they might otherwise share with a store or other distributor. Going direct also means that SaaS applications are available anywhere someone has an Internet connection, making them truly global applications. This feature has allowed many SaaS firms to address highly specialized markets (sometimes called vertical niches). For example, the Internet allows a company writing specialized legal software, for example, or a custom package for the pharmaceutical industry, to have a national deployment footprint from day one. Vendors of desktop applications that go SaaS benefit from this kind of distribution, too.

Finally, SaaS allows a vendor to counter the vexing and costly problem of software piracy. It's just about impossible to make an executable, illegal copy of a subscription service that runs on a SaaS provider's hardware.



Gaming in Flux: Is There a Future in Free?

PC game makers are in a particularly tough spot. Development costs are growing as games become more sophisticated. But profits are plummeting as firms face rampant piracy, a growing market for used game sales, and lower sales from rental options from firms like Blockbuster and GameFly. To combat these trends, Electronic Arts (EA) has begun to experiment with a radical alternative to PC game sales—give the base version of the product away for free and make money by selling additional features.

The firm started with the Korean version of its popular FIFA soccer game. Koreans are crazy for the world's most popular sport; their nation even cohosted the World Cup in 2002. But piracy was killing EA's sales in Korea. To combat the problem, EA created a free, online version of FIFA that let fans pay for additional features and upgrades, such as new uniforms for their virtual teams, or performance-enhancing add-ons. Each enhancement only costs about one dollar and fifty cents, but the move to a model based on these so-called “microtransactions” has brought in big earnings. During the first two years that the microtransaction-based Korean FIFA game was available, EA raked in roughly one million dollars a month. The two-year, twenty-four-million-dollar take was twice the sales record for EA's original FIFA game.

Asian markets have been particularly receptive to microtransactions—this revenue model makes up a whopping 50 percent of the region's gaming revenues. But whether this model can spread to other parts of the world remains to be seen. The firm's first free, microtransaction offering outside of Korea leverages EA's popular *Battlefield* franchise. *Battlefield Heroes* sports lower quality, more cartoon-like graphics than EA's conventional *Battlefield* offerings, but it will be offered free online. Lest someone think they can rise to the top of player rankings by buying the best military hardware for their virtual armies, EA offers a sophisticated



matching engine, pitting players with similar abilities and add-ons against one another. [5]

Players of the first versions of *Battlefield Heroes* and *FIFA Online* needed to download software to their PC. But the startup *World Golf Tour* shows how increasingly sophisticated games can execute within a browser, SaaS-style. WGT doesn't have quite the graphics sophistication of the dominant desktop golf game (EA's *Tiger Woods PGA Golf*), but the free, ad-supported offering is surprisingly detailed. Buddies can meet up online for a virtual foursome, played on high-resolution representations of the world's elite courses stitched together from fly-over photographs taken as part of game development. *World Golf Tour* is ad-supported. The firm hopes that advertisers will covet access to the high-income office workers likely to favor a quick virtual golf game to break up their workday. Zynga's *Farmville*, an app game for Facebook, combines both models. Free online, but offering added features purchased in micropayment-sized chunks, *Farmville* made half a million dollars in three days, just by selling five-dollar virtual sweet potatoes. [6] *FIFA Online*, *Battlefield Heroes*, *World Golf Tour*, and *Farmville* all show that the conventional models of gaming software are just as much in flux as those facing business and productivity packages.

KEY TAKEAWAYS

- SaaS firms may offer their clients several benefits including the following:
 - *lower costs* by eliminating or reducing software, hardware, maintenance, and staff expenses.
 - *financial risk mitigation* since startup costs are so low.
 - potentially *faster deployment times* compared with installed packaged software or systems developed in-house.
 - costs are a *variable operating expense* rather than a large, fixed capital expense.



- *scalable systems* make it easier for firms to ramp up during periods of unexpectedly high system use.
- *higher quality and service levels* through instantly available upgrades, vendor scale economies, and expertise gained across its entire client base.
- *remote access and availability*—most SaaS offerings are accessed through any Web browser, and often even by phone or other mobile device.
- Vendors of SaaS products benefit by:
 - *limiting development to a single platform*, instead of having to create versions for different operating systems.
 - *tighter feedback loop* with clients, helping fuel innovation and responsiveness.
 - ability to *instantly deploy bug fixes and product enhancements* to all users.
 - *lower distribution costs*.
 - *accessibility* to anyone with an Internet connection.
 - *greatly reduces software piracy*.
- Microtransactions and ad-supported gaming present alternatives to conventional purchased video games. Firms leveraging these models potentially benefit from a host of SaaS advantages, including direct-to-consumer distribution, instant upgrades, continued revenue streams rather than one-time purchase payments, and a method for combating piracy.

QUESTIONS AND EXERCISES

1. Firms that buy conventional enterprise software spend money buying software and hardware. What additional and ongoing expenses are required as part of the “care and feeding” of enterprise applications?
2. In what ways can firms using SaaS products dramatically lower costs associated with their information systems?
3. How do SaaS vendors earn their money?
4. Give examples of enterprise-focused SaaS vendors and their products. Visit the Web sites of the firms that offer these services. Which firms are

- listed as clients? Does there appear to be a particular type of firm that uses its services, or are client firms broadly represented?
5. Give examples of desktop-focused SaaS vendors and their products. If some of these are free, try them out and compare them to desktop alternatives you may have used. Be prepared to share your experiences with your class.
 6. List the cost-related benefits to users that accrue from using SaaS.
 7. List the benefits other than cost-related that accrue to users from using SaaS.
 8. List the benefits realized by vendors that offer SaaS services instead of conventional software.
 9. Microtransactions have been tried in many contexts, but have often failed. Can you think of contexts where microtransactions don't work well? Are there contexts where you have paid (or would be willing to pay) for products and services via microtransactions? What do you suppose are the major barriers to the broader acceptance of microtransactions? Do struggles have more to do with technology, consumer attitudes, or both?
 10. Search online to find free and microtransaction-based games. What do you think of these efforts? What kind of gamers do these efforts appeal to? See if you can investigate whether there are examples of particularly successful offerings, or efforts that have failed. What's the reason behind the success or failure of the efforts that you've investigated?



10.8 SaaS: Not without Risks

LEARNING OBJECTIVE

After studying this section you should be able to do the following:

1. Be able to list and appreciate the risks associated with SaaS.

Like any technology, we also recognize there is rarely a silver bullet that solves all problems. A successful manager is able to see through industry hype and weigh the benefits of a technology against its weaknesses and limitations. And there are still several major concerns surrounding SaaS. The largest concerns involve the tremendous dependence a firm develops with its SaaS vendor. Having all of your eggs in one basket can leave a firm particularly vulnerable. If a traditional software company goes out of business, in most cases its customers can still go on using its products. But



if your SaaS vendor goes under, you're hosed. They've got all of your data, and even if firms could get their data out, most organizations don't have the hardware, software, staff, or expertise to quickly absorb an abandoned function.

Beware with whom you partner. Any hot technology is likely to attract a lot of startups, and most of these startups are unlikely to survive. In just a single year, the leading trade association found the number of SaaS vendors dropped from seven hundred members to four hundred fifty. ^[1] One of the early efforts to collapse was Pandesic, a joint venture between SAP and Intel—two large firms that might have otherwise instilled confidence among prospective customers. In another example, Danish SaaS firm “IT Factory” was declared “Denmark’s Best IT Company 2008” by *Computerworld*, only to follow the award one week later with a bankruptcy declaration. ^[2] Indeed, despite the benefits, the costs of operating as a SaaS vendor can be daunting. NetSuite’s founder claimed it “takes ten years and one hundred million dollars to do right” ^[3]—maybe that’s why the firm still wasn’t profitable, even a year and a half after going public.

Firms that buy and install packaged software usually have the option of sticking with the old stuff as long as it works, but organizations adopting SaaS may find they are forced into adopting new versions. This fact is important because any radical changes in a SaaS system’s user interface or system functionality might result in unforeseen training costs, or increase the chance that a user might make an error.

Keep in mind that SaaS systems are also reliant on a network connection. If a firm’s link to the Internet goes down, its link to its SaaS vendor is also severed. Relying on an Internet connection also means that data is transferred to and from a SaaS firm at Internet speeds, rather the potentially higher speeds of a firm’s internal network. Solutions to many of



these issues are evolving as Internet speeds become faster and Internet service providers become more reliable. There are also several programs that allow for offline use of data that typically stored in SaaS systems, including Google Gears and Adobe AIR. With these products a user can download a subset of data to be offline (say on a plane flight or other inaccessible location), and then sync the data when the connection is restored. Ultimately, though, SaaS users have a much higher level of dependence on their Internet connections.

And although a SaaS firm may have more security expertise than your organization, that doesn't mean that security issues can be ignored. Any time a firm allows employees to access a corporation's systems and data assets from a remote location, a firm is potentially vulnerable to abuse and infiltration. Some firms may simply be unacceptably uncomfortable with critical data assets existing outside their own network. There may also be contractual or legal issues preventing data from being housed remotely, especially if a SaaS vendor's systems are in another country operating under different laws and regulations. "We're very bound by regulators in terms of client data and country-of-origin issues, so it's very difficult to use the cloud," says Rupert Brown, a chief architect at Merrill Lynch. [4]

SaaS systems are often accused of being less flexible than their installed software counterparts—mostly due to the more robust configuration and programming options available in traditional software packages. It is true that many SaaS vendors have improved system customization options and integration with standard software packages. And at times a lack of complexity can be a blessing—fewer choices can mean less training, faster startup time, and lower costs associated with system use. But firms with unique needs may find SaaS restrictive.

SaaS offerings usually work well when the bulk of computing happens at the server end of a distributed system because the kind of user interface you



can create in a browser isn't as sophisticated as what you can do with a separate, custom-developed desktop program. A comparison of the first few iterations of the Web-based Google office suite, which offers word processing, presentation software, and a spreadsheet, reveals a much more limited feature set than Microsoft's Office desktop software. The bonus, of course, is that an online office suite is accessible anywhere and makes sharing documents a snap. Again, an understanding of tradeoffs is key. Here's another challenge for a firm and its IT staff: SaaS means a greater *consumerization* of technology. Employees, at their own initiative, can go to SocialText or Google Sites and set up a wiki, WordPress to start blogging, or subscribe to a SaaS offering like Salesforce.com, all without corporate oversight and approval. This work can result in employees operating outside established firm guidelines and procedures, potentially introducing operational inconsistencies or even legal and security concerns. The consumerization of corporate technology isn't all bad. Employee creativity can blossom with increased access to new technologies, costs might be lower than home grown solutions, and staff could introduce the firm to new tools that might not otherwise be on the radar of the firm's IS Department. But all this creates an environment that requires a level of engagement between a firm's technical staff and the groups that it serves that is deeper than that employed by any prior generation of technology workers. Those working in an organization's information systems group must be sure to conduct regular meetings with representative groups of employees across the firm to understand their pain points and assess their changing technology needs. Non-IT managers should regularly reach out to IT to ensure that their needs are on the tech staff's agenda. Organizations with internal IT-staff R&D functions that scan new technologies and critically examine their relevance and potential impact on the firm can help guide an organization through the promise and peril of new technologies.



Now more than ever, IT managers must be deeply knowledgeable about business areas, broadly aware of new technologies, and able to bridge the tech and business worlds. Similarly, any manager looking to advance his or her organization has to regularly consider the impact of new technologies.

KEY TAKEAWAYS

The risks associated with SaaS include the following:

- *dependence on a single vendor.*
- concern about the long-term *viability of partner firms.*
- *users may be forced to migrate to new versions*—possibly incurring unforeseen training costs and shifts in operating procedures.
- *reliance on a network connection*—which may be slower, less stable, and less secure.
- *data asset stored off-site*—with the potential for security and legal concerns.
- *limited configuration, customization, and system integration options* compared to packaged software or alternatives developed in house.
- *the user interface of Web-based software is often less sophisticated and lacks the richness of most desktop alternatives.*
- ease of adoption *may lead to pockets of unauthorized IT* being used throughout an organization.

QUESTIONS AND EXERCISES

1. Consider the following two firms: a consulting startup, and a defense contractor. Leverage what you know about SaaS and advise whether each might consider SaaS efforts for CRM or other enterprise functions? Why or why not?

2. Think of firms you've worked for, or firms you would like to work for. Do SaaS offerings make sense for these firms? Make a case for or against using certain categories of SaaS.
3. What factors would you consider when evaluating a SaaS vendor? Which firms are more appealing to you and why?
4. Discuss problems that may arise because SaaS solutions rely on Internet connections. Discuss the advantages of through-the-browser access.
5. Evaluate trial versions of desktop SaaS offerings (offered by Adobe, Google, Microsoft, Zoho, or others). Do you agree that the interfaces of Web-based versions are not as robust as desktop rivals? Are they good enough for you? For most users?



10.9 The Hardware Cloud: Utility Computing and Its Cousins

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Distinguish between SaaS and hardware clouds.
2. Provide examples of firms and uses of hardware clouds.
3. Understand the concepts of cloud computing, cloudbursting, and black swan events.
4. Understand the challenges and economics involved in shifting computing hardware to the cloud.

While SaaS provides the software *and* hardware to replace an internal information system, sometimes a firm develops its own custom software but wants to pay someone else to run it for them. That's where hardware clouds, utility computing, and related technologies come in. In this model, a firm replaces computing hardware that it might otherwise run on-site with a service provided by a third party online. While the term utility computing was fashionable a few years back (and old timers claim it shares a lineage with terms like hosted computing or even time sharing), now most in the industry have begun referring to this as an aspect of cloud computing, often referred to as hardware clouds. Computing hardware used in this scenario exists in the cloud, meaning somewhere on the Internet. The costs of systems operated in this manner look more like a utility bill—you only pay for the amount of processing, storage, and telecommunications used. Tech research firm Gartner has estimated that 80 percent of corporate tech



spending goes toward data center maintenance. [1] Hardware-focused cloud computing provides a way for firms to chip away at these costs.

Major players are spending billions building out huge data centers to take all kinds of computing out of the corporate data center and place it “in the cloud.” Efforts include Sun’s Network.com grid, IBM’s Cloud Labs, Amazon’s EC2 (Elastic Computing Cloud), Google’s App Engine, Microsoft’s Azure, and Salesforce.com’s Force.com. While cloud vendors typically host your software on their systems, many of these vendors also offer additional tools to help in creating and hosting apps in the cloud. Salesforce.com offers Force.com which includes not only a hardware cloud, but also several cloud-supporting tools, including a programming environment (IDE) to write applications specifically tailored for Web-based delivery. Google’s App Engine offers developers a database product called Big Table, while Amazon’s offers one called Amazon DB. Traditional software firms like Oracle are also making their products available to developers through various cloud initiatives.

Still other cloud computing efforts focus on providing a virtual replacement for operational hardware like storage and backup solutions. These include the cloud-based backup efforts like EMC’s Mozy, and corporate storage services like Amazon’s Simple Storage Solution (S3). Even efforts like Apple’s MobileMe and Microsoft’s Live Mesh that sync user data across devices (phone, multiple desktops) are considered part of the cloud craze. The common theme in all of this is leveraging computing delivered over the Internet to satisfy the computing needs of both users and organizations.

Clouds in Action: A Snapshot of Diverse Efforts

Large, established organizations, small firms and startups are all embracing the cloud. The examples below illustrate the wide range of these efforts.

Journalists refer to the *New York Times* as, “The Old Gray Lady,” but it turns out that the venerable paper is a cloud-pioneering whippersnapper. When the *Times* decided to make roughly one hundred fifty years of newspaper archives (over fifteen million articles) available over the Internet, it realized that the process of converting scans into searchable PDFs would require more computing power than the firm had available. [2] To solve the challenge, a *Times* IT staffer simply broke out a credit card and signed up for Amazon’s EC2 cloud computing and S3 cloud storage services. The *Times* then started uploading terabytes of information to Amazon, along with a chunk of code to execute the conversion. While anyone can sign up for services online without speaking to a rep, someone from Amazon eventually contacted the *Times* to check in after noticing the massive volume of data coming into its systems. Using one hundred of Amazon’s Linux servers, the *Times* job took just twenty-four hours to complete. In fact, a coding error in the initial batch forced the paper to rerun the job. Even the blunder was cheap—just two hundred forty dollars in extra processing costs. Says a member of the *Times* IT group: “It would have taken a month at our facilities, since we only had a few spare PCs....It was cheap experimentation, and the learning curve isn’t steep.” [3]

NASDAQ also uses Amazon’s cloud as part of its Market Replay system. The exchange uses Amazon to make terabytes of data available on demand, and uploads an additional thirty to eighty gigabytes every day. Market Replay allows access through an Adobe AIR interface to pull together historical market conditions in the ten-minute period surrounding a trade’s execution. This allows NASDAQ to produce a snapshot of information for regulators or customers who question a trade. Says the exchange’s VP of Product Development, “The fact that we’re able to keep so much data online indefinitely means the brokers can quickly answer a question without having to pull data out of old tapes and CD backups.” [4] NASDAQ isn’t the



only major financial organization leveraging someone else's cloud. Others include Merrill Lynch, which uses IBM's Blue Cloud servers to build and evaluate risk analysis programs; and Morgan Stanley, which relies on Force.com for recruiting applications.

The Network.com offering from Sun Microsystems is essentially a grid computer in the clouds (see [Chapter 4 "Moore's Law and More: Fast, Cheap Computing and What It Means for the Manager"](#)). Since grid computers break a task up to spread across multiple processors, the Sun service is best for problems that can be easily divided into smaller mini jobs that can be processed simultaneously by the army of processors in Sun's grid. The firm's cloud is particularly useful for performing large-scale image and data tasks. Infosolve, a data management firm, uses the Sun cloud to scrub massive data sets, at times harnessing thousands of processors to comb through client records and correct inconsistent entries.

IBM Cloud Labs, which counts Elizabeth Arden and the U.S. Golf Association among its customers, offers several services, including so-called cloudbursting. In a cloudbursting scenario a firm's data center running at maximum capacity can seamlessly shift part of the workload to IBM's cloud, with any spikes in system use metered, utility style.

Cloudbursting is appealing because forecasting demand is difficult and can't account for the ultrarare, high-impact events, sometimes called black swans. Planning to account for usage spikes explains why the servers at many conventional corporate IS shops run at only 10 to 20 percent capacity. [5] While Cloud Labs cloudbursting service is particularly appealing for firms that already have a heavy reliance on IBM hardware in-house, it is possible to build these systems using the hardware clouds of other vendors, too.

Salesforce.com's Force.com cloud is especially tuned to help firms create and deploy custom Web applications. The firm makes it possible to piece

together projects using premade Web services that provide software building blocks for features like calendaring and scheduling. The integration with the firm's SaaS CRM effort, and with third-party products like Google Maps allows enterprise mash-ups that can combine services from different vendors into a single application that's run on Force.com hardware. The platform even includes tools to help deploy Facebook applications. Intuitive Surgical used Force.com to create and host a custom application to gather clinical trial data for the firm's surgical robots. An IS manager at Intuitive noted, "We could build it using just their tools, so in essence, there was no programming." [6] Other users include Jobscience, which used Force.com to launch its online recruiting site; and Harrah's Entertainment, which uses Force.com applications to manage room reservations, air travel programs, and player relations. These efforts compete with a host of other initiatives, including Google's App Engine and Microsoft's Azure Services Platform, hosting firms like Rackspace, and cloud-specific upstarts like GoGrid.

Challenges Remain

Hardware clouds and SaaS share similar benefits and risk, and as our discussion of SaaS showed, cloud efforts aren't for everyone. Some additional examples illustrate the challenges in shifting computing hardware to the cloud.

For all the hype about cloud computing, it doesn't work in all situations. From an architectural standpoint, most large organizations run a hodgepodge of systems that include both package applications and custom code written in house. Installing a complex set of systems on someone else's hardware can be a brutal challenge and in many cases is just about impossible. For that reason we can expect most cloud computing efforts to focus on new software development projects rather than options for old

software. Even for efforts that can be custom-built and cloud-deployed, other roadblocks remain. For example, some firms face stringent regulatory compliance issues. To quote one tech industry executive, “How do you demonstrate what you are doing is in compliance when it is done outside?” [7]

Firms considering cloud computing need to do a thorough financial analysis, comparing the capital and other costs of owning and operating their own systems over time against the variable costs over the same period for moving portions to the cloud. For high-volume, low-maintenance systems, the numbers may show that it makes sense to buy rather than rent. Cloud costs can seem super cheap at first. Sun’s early cloud effort offered a flat fee of one dollar per CPU per hour. Amazon’s cloud storage rates were twenty-five cents per gigabyte per month. But users often also pay for the number of accesses and the number of data transfers. [8] A quarter a gigabyte a month may seem like a small amount, but system maintenance costs often include the need to clean up old files or put them on tape. If unlimited data is stored in the cloud, these costs can add up. Firms should enter the cloud cautiously, particularly where mission-critical systems are concerned. When one of the three centers supporting Amazon’s cloud briefly went dark in 2008, startups relying on the service, including Twitter and SmugMug, reported outages. Apple’s MobileMe cloud-based product for synchronizing data across computers and mobile devices, struggled for months after its introduction when the cloud repeatedly went down. Vendors with multiple data centers that are able to operate with fault-tolerant provisioning, keeping a firm’s efforts at more than one location to account for any operating interruptions, will appeal to firms with stricter uptime requirements.



KEY TAKEAWAYS

- It's estimated that 80 percent of corporate tech spending goes toward data center maintenance. Hardware-focused cloud computing initiatives from third party firms help tackle this cost by allowing firms to run their own software on the hardware of the provider.
- Amazon, EMC, Google, IBM, Microsoft, Oracle/Sun, Rackspace, and Salesforce.com are among firms offering platforms to run custom software projects. Some offer additional tools and services, including additional support for cloud-based software development, hosting, application integration, and backup.
- Users of cloud computing run the gamut of industries, including publishing (the *New York Times*), finance (NASDAQ), and cosmetics and skin care (Elizabeth Arden).
- Benefits and risks are similar to those discussed in SaaS efforts. Benefits include the use of the cloud for handling large batch jobs or limited-time tasks, offloading expensive computing tasks, and cloudbursting efforts that handle system overflow when an organization needs more capacity.
- Most legacy systems can't be easily migrated to the cloud, meaning most efforts will be new efforts or those launched by younger firms.
- Cloud (utility) computing doesn't work in situations where complex legacy systems have to be ported, or where there may be regulatory compliance issues.
- Some firms may still find TCO and pricing economics favor buying over renting—scale sometimes suggests an organization is better off keeping efforts in-house.

QUESTIONS AND EXERCISES

1. What are hardware clouds? What kinds of services are described by this terms? What are other names for this phenomenon? How does this differ from SaaS?

2. Which firms are the leading providers of hardware clouds? How are clients using these efforts?
3. List the circumstances where hardware clouds work best and where it works poorly.
4. Research cloud-based alternatives for backing up your hard drive. Which are among the best reviewed product or services? Why? Do you or would you use such a service? Why or why not?
5. Can you think of “black swan” events that have caused computing services to become less reliable? Describe the events and its consequences for computing services. Suggest a method and vendor for helping firms overcome the sorts of events that you encountered.

10.10 Clouds and Tech Industry Impact

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand how cloud computing’s impact across industries is proving to be broad and significant.
2. Know the effects of cloud computing on high-end server sales and the influence on the trend shifting from hardware sales to service.
3. Know the effects of cloud computing on innovation and the influence on the changes in the desired skills mix and job outlook for IS workers.
4. Know that by lowering the cost to access powerful systems and software, cloud computing can decrease barriers to entry.



5. Understand the importance, size, and metrics of server farms.

Although still a relatively recent phenomenon, cloud computing's impact across industries is already proving to be broad and significant.

Cloud computing is affecting the competitive dynamics of the hardware, software, and consulting industries. In the past, firms seeking to increase computing capacity invested heavily in expensive, high margin server hardware, creating a huge market for computer manufacturers. But now hardware firms find these markets may be threatened by the cloud. The trend shifting from hardware to services is evident in IBM's quarterly numbers. The firm recently reported its overall earnings were up 12 percent, even though hardware sales were off by 20 percent. ^[1] What made up the difference? The growth of Big Blue's services business. IBM is particularly well positioned to take advantage of the shift to services because it employs more technology consultants than any other firm in the world, while most of its competitors are forced to partner to offer something comparable. Consulting firm Capgemini's partnership to offer cloud services through Amazon is one such example.

The shift to cloud computing also alters the margin structure for many in the computing industry. While Moore's Law has made servers cheap, deploying SaaS and operating a commercial cloud is still very expensive—much more so than simply making additional copies of conventional, packaged software. Microsoft surprised Wall Street when it announced it would need to pour at least two billion dollars more than analysts expected into the year's server farm capital spending. The firm's stock—among the world's most widely held—sank 11 percent in a day. ^[2] As a result, many portfolio managers started paying closer attention to the business implications of the cloud.

Cloud computing can accelerate innovation and therefore changes the desired skills mix and job outlook for IS workers. If cloud computing customers spend less on expensive infrastructure investments, they potentially have more money to reinvest in strategic efforts and innovation. IT careers may change, too. Demand for nonstrategic skills like hardware operations and maintenance are likely to decrease. Organizations will need more business-focused technologists who intimately understand a firm's competitive environment, and can create systems that add value and differentiate the firm from its competition. [3] While these tech jobs require more business training, they're also likely to be more durable and less likely to be outsourced to a third party with a limited understanding of the firm. By lowering the cost to access powerful systems and software, barriers to entry also decrease. Firms need to think about the strategic advantages they can create, even as technology is easily duplicated. This trend means the potential for more new entrants across industries, and since startups can do more with less, it's also influencing entrepreneurship and venture capital. The CTO of SlideShare, a startup that launched using Amazon's S3 storage cloud, offers a presentation on his firm's site labeled "Using S3 to Avoid VC." Similarly, the CEO of online payments startup Zuora claims to have saved between half a million and one million dollars by using cloud computing: "We have no servers, we run the entire business in the cloud." [4] And the sophistication of these tools lowers development time. Enterprise firm Apttus claims it was able to perform the equivalent of six months of development in a couple of weekends by using cloud services. The firm scored its first million-dollar deal in three months, and was break-even in nine months, a ramp-up time that would have been unheard of, had they needed to plan, purchase, and deploy their own data center, and create from scratch the Web services that were provided by its cloud vendor. [5]



So What's It Take to Run This Thing?

In the countryside surrounding the Columbia River in the Pacific Northwest, potato farms are yielding to server farms. Turns out the area is tailor made for creating the kinds of massive data installations that form the building blocks of cloud computing. The land is cheap, the region's hydroelectric power costs a fraction of Silicon Valley rates, and the area is served by ultrafast fiber-optic connections. Even the area's mild temperatures cut cooling costs.

Most major players in cloud computing have server farms in the region, each with thousands of processors humming away simultaneously.

Microsoft's Quincy, Washington, facility is as big as ten American football fields and has nearly six hundred miles of wiring, 1.5 metric tons of battery backup, and three miles of chiller piping to keep things cool. Storage is big enough to store 6.75 trillion photos. Just a short drive away, Yahoo has two facilities on fifty acres, including one that runs at a zero carbon footprint. Google has a thirty-acre site sprawled across former farmland in The Dalles, Oregon. The Google site includes two massive buildings, with a third on the way. And in Boardman, Oregon, Amazon has a three building petabyte palace that sports its own ten-megawatt electrical substation. [6]

While U.S. activity has been particularly intense in the Pacific Northwest, server farms that support cloud computing are popping up from Shanghai to São Paulo. Not only does a diverse infrastructure offer a degree of fault tolerance and disaster recovery (Oregon down? Shift to North Carolina), the myriad of national laws and industry-specific regulatory environments may require some firms to keep data within a specific country or region. To meet the challenge, cloud vendors are racing to deploy infrastructure worldwide and allowing customers to select regional availability zones for their cloud computing needs.

The build-out race has become so intense that many firms have developed rapid-deployment server farm modules that are preconfigured and packed inside shipping containers. Some of these units contain as many as three thousand servers each. Just drop the containers on site, link to power, water, and telecom, and presto—you've got yourself a data center. More than two hundred containers can be used on a single site. One Microsoft VP claimed the configuration has cut the time to open a data center to just a few days, claiming Microsoft's San Antonio facility was operational in less time than it took a local western wear firm to deliver her custom-made cowboy boots! [7] Microsoft's Dublin-based fourth generation data center will be built entirely of containers—no walls or roof—using the outside air for much of the cooling. [8]

While firms are buying less hardware, cloud vendors have turned out to be the computing industry's best customers. Amazon has spent well over two billion dollars on its cloud infrastructure. Google reportedly has 1.4 million servers operating across three dozen data centers. [9] Demonstrating it won't be outdone, Microsoft plans to build as many as twenty server farms, at costs of up to one billion dollars each. [10] Look for the clouds to pop up in unexpected places. Microsoft has scouted locations in Siberia, while Google has applied to patent a method for floating data centers on an offshore platform powered by wave motions. [11]

KEY TAKEAWAYS

- Cloud computing's impact across industries is proving to be broad and significant.
- Clouds can lower barriers to entry in an industry, making it easier for startups to launch and smaller firms to leverage the backing of powerful technology.

- Clouds may also lower the amount of capital a firm needs to launch a business, shifting power away from venture firms in those industries that had previously needed more VC money.
- Clouds can shift resources out of capital spending and into profitability and innovation.
- Hardware and software sales may drop as cloud use increases, while service revenues will increase.
- Cloud computing can accelerate innovation and therefore changes the desired skills mix and job outlook for IS workers. Tech skills in data center operations, support, and maintenance may shrink as a smaller number of vendors consolidate these functions.
- Demand continues to spike for business-savvy technologists. Tech managers will need even stronger business skills and will focus an increasing percentage of their time on strategic efforts. These latter jobs are tougher to outsource, since they involve an intimate knowledge of the firm, its industry, and its operations.
- The market for expensive, high margin, sever hardware is threatened by companies moving applications to the cloud instead of investing in hardware.
- Server farms require plenty of cheap land, low cost power, ultrafast fiber-optic connections, and benefit from mild climates.
- Sun, Microsoft, IBM, and HP have all developed rapid-deployment server farm modules that are pre configured and packed inside shipping containers.

QUESTIONS AND EXERCISES

1. Describe the change in IBM's revenue stream resulting from the shift to the cloud.
2. Why is IBM particularly well positioned to take advantage of the shift to services?

3. Describe the shift in skill sets required for IT workers that is likely to result from the widespread adoption of cloud computing.
4. Why do certain entry barriers decrease as a result of cloud computing? What is the effect of lower entry barriers on new entrants, entrepreneurship, and venture capital? On existing competitors?
5. What factors make the Columbia River region of the Pacific Northwest an ideal location for server farms?
6. What is the estimated number of computers operated by Google?
7. Why did Microsoft's shift to cloud computing create an unexpected shock among stock analysts who cover the firm? What does this tell you about the importance of technology understanding among finance and investment professionals?
8. Why do cloud computing vendors build regional server farms instead of one mega site?
9. Why would a firm build a container-based data center?



10.11 Virtualization: Software That Makes One Computer Act Like Many

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Know what virtualization software is and its impact on cloud computing.
2. Be able to list the benefits to a firm from using virtualization.

The reduced costs and increased power of commodity hardware are not the only contributors to the explosion of cloud computing. The availability of increasingly sophisticated software tools has also had an impact. Perhaps the most important software tool in the cloud computing toolbox is virtualization. Think of virtualization as being a kind of operating system for operating systems. A server running virtualization software can create smaller compartments in memory that each behave as a separate computer with its own operating system and resources. The most sophisticated of these tools also allow firms to combine servers into a huge pool of computing resources that can be allocated as needed. [1]

Virtualization can generate huge savings. Some studies have shown that on average, conventional data centers run at 15 percent or less of their maximum capacity. Data centers using virtualization software have increased utilization to 80 percent or more. [2] This increased efficiency means cost savings in hardware, staff, and real estate. Plus it reduces a firm's IT-based energy consumption, cutting costs, lowering its carbon footprint, and boosting "green cred." [3] Using virtualization, firms can buy and maintain fewer servers, each running at a greater capacity. It can also power down servers until demand increases require them to come online.



While virtualization is a key software building block that makes public cloud computing happen, it can also be used in-house to reduce an organization's hardware needs, and even to create a firm's own private cloud of scalable assets. Bechtel, BT, Merrill Lynch, and Morgan Stanley are among the firms with large private clouds enabled by virtualization. [4] Another kind of virtualization, virtual desktops allow a server to run what amounts to a copy of a PC—OS, applications, and all—and simply deliver an image of what's executing to a PC or other connected device. This allows firms to scale, back up, secure, and upgrade systems far easier than if they had to maintain each individual PC. One game start-up hopes to remove the high-powered game console hardware attached to your television and instead put the console “in the cloud,” delivering games to your TV as they execute remotely on superfast server hardware. Virtualization can even live on your desktop. Anyone who's ever run Windows in a window on Mac OS X is using virtualization software; these tools inhabit a chunk of your Mac's memory for running Windows and actually fool this foreign OS into thinking that it's on a PC. Interest in virtualization has exploded in recent years. VMware, the virtualization software division of storage firm EMC, was the biggest IPO of 2007. But its niche is getting crowded. Microsoft has entered the market, building virtualization into its server offerings. Dell bought a virtualization software firm for \$1.54 billion. And there's even an open source virtualization product called Xen. [5]

KEY TAKEAWAYS

- Virtualization software allows one computing device to function as many. The most sophisticated products also make it easy for organizations to scale computing requirements across several servers.

- Virtualization software can lower a firm's hardware needs, save energy, and boost scalability.
- Data center virtualization software is at the heart of many so called private clouds, scalable corporate data centers, as well as the sorts of public efforts described earlier.
- Virtualization also works on the desktop, allowing multiple operating systems (Mac OS X, Linux, Windows) to run simultaneously on the same platform.
- Virtualization software can increase data center utilization to 80 percent or more.
- While virtualization is used to make public cloud computing happen, it can also be used in-house to create a firm's own private cloud.
- A number of companies, including Microsoft and Dell, have entered the growing virtualization market.

QUESTIONS AND EXERCISES

1. List the benefits to a firm from using virtualization.
2. What is the average utilization rate for conventional data centers?
3. List companies that have virtualization-enabled private clouds.
4. Give an example of desktop virtualization.
5. Name three companies that are players in the virtualization software industry.

10.12 Make, Buy, or Rent

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Know the options managers have when determining how to satisfy the software needs of their companies.

2. Know the factors that must be considered when making the make, buy, or rent decision.

So now you realize managers have a whole host of options when seeking to fulfill the software needs of their firms. An organization can purchase packaged software from a vendor, use open source offerings, leverage SaaS or other type of cloud computing, outsource development or other IT functions to another firm either domestically or abroad, or a firm can develop all or part of the effort themselves. When presented with all of these options, making decisions about technologies and systems can seem pretty daunting.

First, realize that that for most firms, technology decisions are not binary options for the whole organization in all situations. Few businesses will opt for an IT configuration that is 100 percent in house, packaged, or SaaS. Being aware of the parameters to consider can help a firm make better, more informed decisions. It's also important to keep in mind that these decisions need to be continuously reevaluated as markets and business needs change. What follows is a summary of some of the key variables to consider.

Competitive Advantage—*Do we rely on unique processes, procedures, or technologies that create vital, differentiating competitive advantage?* If so, then these functions aren't a good candidate to outsource or replace with a package software offering. Amazon.com had originally used recommendation software provided by a third party, and Netflix and Dell both considered third-party software to manage inventory fulfillment. But in all three cases, these firms felt that mastery of these functions was too critical to competitive advantage, so each firm developed proprietary systems unique to the circumstances of each firm.



Security—*Are there unacceptable risks associated with using the packaged software, OSS, cloud solution, or an outsourcing vendor? Are we convinced that the prospective solution is sufficiently secure and reliable? Can we trust the prospective vendor with our code, our data, our procedures and our way of doing business? Are there noncompete provisions for vendor staff that may be privy to our secrets? For off-site work, are there sufficient policies in place for on-site auditing?* If the answers to any of these questions is no, outsourcing might not be a viable option.

Legal and Compliance—*Is our firm prohibited outright from using technologies? Are there specific legal and compliance requirements related to deploying our products or services?* Even a technology as innocuous as instant messaging may need to be deployed in such a way that it complies with laws requiring firms to record and reproduce the electronic equivalent of a paper trail. For example, SEC Rule 17a-4 requires broker dealers to retain client communications for a minimum of three years. HIPAA laws governing healthcare providers state that electronic communications must also be captured and stored. ^[1] While tech has gained a seat in the board room, legal also deserves a seat in systems planning meetings.

Skill, Expertise, and Available Labor—*Can we build it?* The firm may have skilled technologists, but they may not be sufficiently experienced with a new technology. Even if they are skilled, managers much consider the costs of allocating staff away from existing projects for this effort.

Cost—*Is this a cost effective choice for our firm?* A host of factors must be considered when evaluating the cost of an IT decision. The costs to build, host, maintain, and support an ongoing effort involve labor (software development, quality assurance, ongoing support, training, and maintenance), consulting, security, operations, licensing, energy, and real

estate. Any analysis of costs should consider not only the aggregate spending required over the lifetime of the effort but also whether or not these factors might vary over time.

Time—*Do we have time to build, test, and deploy the system?*

Vendor Issues—*Is the vendor reputable and in a sound financial position? Can the vendor guarantee the service levels and reliability we need? What provisions are in place in case the vendor fails or is acquired? Is the vendor certified via the Carnegie Mellon Software Institute or other standards organizations in a way that conveys quality, trust, and reliability?*

The list above is a starter. It should also be clear that these metrics are sometimes quite tough to estimate. Welcome to the challenges of being a manager! At times an environment in flux can make an executive feel like he or she is working on a surfboard, constantly being buffeted about by unexpected currents and waves. Hopefully the issues outlined in this chapter will give you the surfing skills you need for a safe ride that avoids the organizational equivalent of a wipeout.

KEY TAKEAWAYS

- The make, buy, or rent decision may apply on a case-by-case basis that might be evaluated by firm, division, project or project component. Firm and industry dynamics may change in a way that causes firms to reassess earlier decisions, or to alter the direction of new initiatives.
- Factors that managers should considered when making a make, buy, or rent decision include the following: competitive advantage, security, legal and compliance issues, the organization's skill and available labor, cost, time, and vendor issues.
- Factors must be evaluated over the lifetime of a project, not at a single point in time.



- Managers have numerous options available when determining how to satisfy the software needs of their companies: purchase packaged software from a vendor, use OSS, use SaaS or utility computing, outsourcing development, or developing all or part of the effort themselves.
- If a company relies on unique processes, procedures, or technologies that create vital, differentiating, competitive advantages, the functions probably aren't a good candidate to outsource.

QUESTIONS AND EXERCISES

1. What are the options available to managers when seeking to meet the software needs of their companies?
2. What are the factors that must be considered when making the make, buy, or rent decision?
3. What are some security-related questions that must be asked when making the make, buy, or rent decision?
4. What are some vendor-related questions that must be asked when making the make, buy, or rent decision?
5. What are some of the factors that must be considered when evaluating the cost of an IT decision?
6. Why must factors be evaluated over the lifetime of a project, not at a single point in time?

Chapter 11

The Data Asset: Databases, Business Intelligence, and Competitive Advantage

11.1 Introduction

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand how increasingly standardized data, access to third-party datasets, cheap, fast computing and easier-to-use software are collectively enabling a new age of decision making.
2. Be familiar with some of the enterprises that have benefited from data-driven, fact-based decision making.

The planet is awash in data. Cash registers ring up transactions worldwide. Web browsers leave a trail of cookie crumbs nearly everywhere they go. And with RFID, inventory can literally announce its presence so that firms can



precisely journal every hop their products make along the value chain: “I’m arriving in the warehouse,” “I’m on the store shelf,” “I’m leaving out the front door.”

A study by Gartner Research claims that the amount of data on corporate hard drives doubles every six months, [1] while IDC states that the collective number of those bits already exceeds the number of stars in the universe. [2] Wal-Mart alone boasts a data volume well over *125 times* as large as the *entire* print collection of the U.S. Library of Congress. [3]

And with this flood of data comes a tidal wave of opportunity. Increasingly standardized corporate data, and access to rich, third-party datasets—all leveraged by cheap, fast computing and easier-to-use software—are collectively enabling a new age of data-driven, fact-based decision making. You’re less likely to hear old-school terms like “decision support systems” used to describe what’s going on here. The phrase of the day is business intelligence (BI), a catchall term combining aspects of reporting, data exploration and ad hoc queries, and sophisticated data modeling and analysis. Alongside business intelligence in the new managerial lexicon is the phrase analytics, a term describing the extensive use of data, statistical and quantitative analysis, explanatory and predictive models, and fact-based management to drive decisions and actions. [4]

The benefits of all this data and number crunching are very real, indeed. Data leverage lies at the center of competitive advantage we’ve studied in the Zara, Netflix, and Google cases. Data mastery has helped vault Wal-Mart to the top of the *Fortune* 500 list. It helped Harrah’s Casino Hotels grow to be twice as profitable as similarly sized Caesars, and rich enough to acquire this rival. And data helped Capital One find valuable customers that competitors were ignoring, delivering ten-year financial performance a full ten times greater than the S&P 500. Data-driven decision making is even credited with helping the Red Sox win their first World Series in eighty-

three years and with helping the New England Patriots win three Super Bowls in four years. To quote from a *BusinessWeek* cover story on analytics, “Math Will Rock Your World!” [5]

Sounds great, but it can be a tough slog getting an organization to the point where it has a leveragable data asset. In many organizations data lies dormant, spread across inconsistent formats and incompatible systems, unable to be turned into anything of value. Many firms have been shocked at the amount of work and complexity required to pull together an infrastructure that empowers its managers. But not only can this be done; it must be done. Firms that are basing decisions on hunches aren’t managing; they’re gambling. And the days of uninformed managerial dice rolling are over.

While we’ll study technology in this chapter, our focus isn’t as much on the technology itself as it is on what you can do with that technology. Consumer products giant P&G believes in this distinction so thoroughly that the firm renamed its IT function as “Information and Decision Solutions.” [6] Solutions drive technology decisions, not the other way around.

In this chapter we’ll study the data asset, how it’s created, how it’s stored, and how it’s accessed and leveraged. We’ll also study many of the firms mentioned above, and more; providing a context for understanding how managers are leveraging data to create winning models, and how those that have failed to realize the power of data have been left in the dust.

Data, Analytics, and Competitive Advantage

Anyone can acquire technology—but data is oftentimes considered a defensible source of competitive advantage. The data a firm can leverage is a true strategic asset when it’s rare, valuable, imperfectly imitable, and lacking in substitutes (see [Chapter 2 "Strategy and Technology"](#)).



If more data brings more accurate modeling, moving early to capture this rare asset can be the difference between a dominating firm and an also-ran. But be forewarned, there's no monopoly on math. Advantages based on capabilities and data that others can acquire will be short-lived. Those advances leveraged by the Red Sox were originally pioneered by the Oakland A's and are now used by nearly every team in the major leagues. This doesn't mean that firms can ignore the importance data can play in lowering costs, increasing customer service, and other ways that boost performance. But differentiation will be key in distinguishing operationally effective data use from those efforts that can yield true strategic positioning.

KEY TAKEAWAYS

- The amount of data on corporate hard drives doubles every six months.
- In many organizations, available data is not exploited to advantage.
- Data is oftentimes considered a defensible source of competitive advantage; however, advantages based on capabilities and data that others can acquire will be short-lived.

QUESTIONS AND EXERCISES

1. Name and define the terms that are supplanting discussions of decision support systems in the modern IS lexicon.
2. Is data a source of competitive advantage? Describe situations in which data might be a source for sustainable competitive advantage. When might data not yield sustainable advantage?
3. Are advantages based on analytics and modeling potentially sustainable? Why or why not?
4. What role do technology and timing play in realizing advantages from the data asset?



11.2 Data, Information, and Knowledge

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand the difference between data and information.
2. Know the key terms and technologies associated with data organization and management.

Data is simply raw facts and figures. Alone it tells you nothing. The real goal is to turn data into information. Data becomes information when it's presented in a context so that it can answer a question or support decision making. And it's when this information can be combined with a manager's knowledge—their insight from experience and expertise—that stronger decisions can be made.

Trusting Your Data

The ability look critically at data and assess its validity are vital managerial skills. When decision makers are presented with wrong data, the results can be disastrous. And these problems can get amplified if bad data is fed to automated systems. As an example, look at the series of man-made and computer-triggered events that brought about a billion-dollar collapse in United Airlines stock.

In the wee hours one Sunday morning in September 2008, a single reader browsing back stories on the *Orlando Sentinel's* Web site viewed a 2002 article on the bankruptcy of United Airlines (UAL went bankrupt in 2002, but emerged from bankruptcy four years later). That lone Web surfer's



access of this story during such a low-traffic time was enough for the *Sentinel's* Web server to briefly list the article as one of the paper's "most popular." Google crawled the site and picked up this "popular" news item, feeding it into Google News.

Early that morning, a worker in a Florida investment firm came across the Google-fed story, assumed United had yet again filed for bankruptcy, then posted a summary on Bloomberg. Investors scanning Bloomberg jumped on what looked like a reputable early warning of another United bankruptcy, dumping UAL stock. Blame the computers again—the rapid plunge from these early trades caused automatic sell systems to kick in (event-triggered, computer-automated trading is responsible for about 30 percent of all stock trades). Once the machines took over, UAL dropped like a rock, falling from twelve to three dollars. That drop represented the vanishing of one billion dollars in wealth, and all this because no one checked the date on a news story. Welcome to the new world of paying attention! ^[1]

Understanding How Is Data Organized: Key Terms and Technologies

A database is simply a list (or more likely, several related lists) of data. Most organizations have several databases—perhaps even hundreds or thousands. And these various databases might be focused on any combination of functional areas (sales, product returns, inventory, payroll), geographical regions, or business units. Firms often create specialized databases for recording transactions, as well as databases that aggregate data from multiple sources in order to support reporting and analysis. Databases are created, maintained, and manipulated using programs called database management systems (DBMS), sometimes referred to as *database software*. DBMS products vary widely in scale and capabilities.



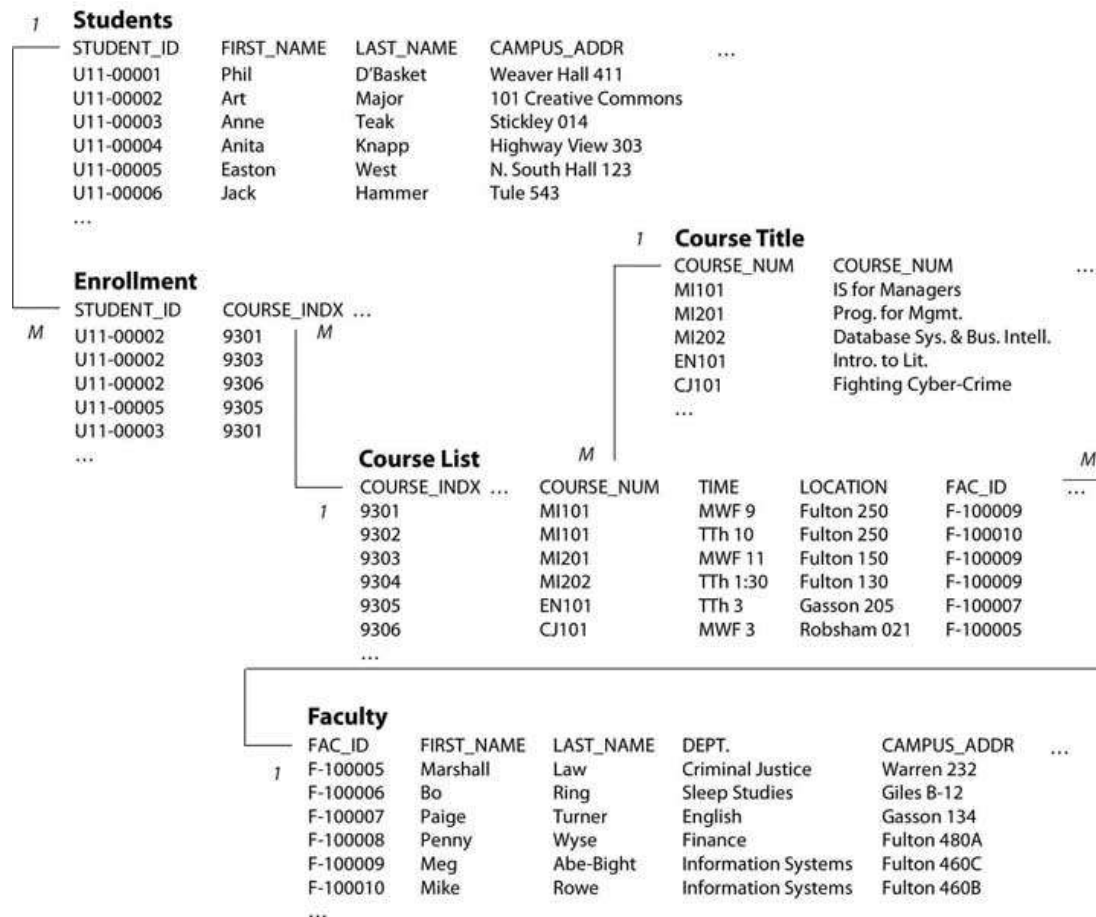
They include the single-user, desktop versions of Microsoft Access or Filemaker Pro, Web-based offerings like Intuit QuickBase, and industrial strength products from Oracle, IBM (DB2), Sybase, Microsoft (SQL Server), and others. Oracle is the world's largest database software vendor, and database software has meant big bucks for Oracle cofounder and CEO Larry Ellison. Ellison perennially ranks in the Top 10 of the *Forbes* 400 list of wealthiest Americans.

The acronym SQL (often pronounced *sequel*) also shows up a lot when talking about databases. Structured query language (SQL) is by far the most common language for creating and manipulating databases. You'll find variants of SQL inhabiting everything from lowly desktop software, to high-powered enterprise products. Microsoft's high-end database is even called SQL Server. And of course there's also the open source MySQL (whose stewardship now sits with Oracle as part of the firm's purchase of Sun Microsystems). Given this popularity, if you're going to learn one language for database use, SQL's a pretty good choice. And for a little inspiration, visit Monster.com or another job site and search for jobs mentioning SQL. You'll find page after page of listings, suggesting that while database systems have been good for Ellison, learning more about them might be pretty good for you, too.

Even if you don't become a database programmer or database administrator (DBA), you're almost surely going to be called upon to dive in and use a database. You may even be asked to help identify your firm's data requirements. It's quite common for nontech employees to work on development teams with technical staff, defining business problems, outlining processes, setting requirements, and determining the kinds of data the firm will need to leverage. Database systems are powerful stuff, and can't be avoided, so a bit of understanding will serve you well.



Figure 11.1 A Simplified Relational Database for a University Course Registration System



A complete discourse on technical concepts associated with database systems is beyond the scope of our managerial introduction, but here are some key concepts to help get you oriented, and that all managers should know.

- A table or file refers to a list of data.
- A *database* is either a single table or a collection of related tables. The course registration database above depicts five tables.
- A column or field defines the data that a table can hold. The “Students” table above shows columns for STUDENT_ID, FIRST_NAME, LAST_NAME, CAMPUS_ADDR (the “...” symbols above are meant to



indicate that in practice there may be more columns or rows than are shown in this simplified diagram).

- A row or record represents a single instance of whatever the table keeps track of. In the example above, each row of the “Students” table represents a student, each row of the “Enrollment” table represents the enrollment of a student in a particular course, and each row of the “Course List” represents a given section of each course offered by the University.
- A key is the field used to relate tables in a database. Look at how the STUDENT_ID key is used above. There is *one* unique STUDENT_ID for each student, but the STUDENT_ID may appear *many* times in the “Enrollment” table, indicating that each student may be enrolled in many classes. The “1” and “M” in the diagram above indicate the one to many relationships among the keys in these tables.

Databases organized like the one above, where multiple tables are related based on common keys, are referred to as relational databases. There are many other database formats (sporting names like *hierarchical*, and *object-oriented*), but relational databases are far and away the most popular. And all SQL databases are relational databases.

We’ve just scratched the surface for a very basic introduction. Expect that a formal class in database systems will offer you far more detail and better design principles than are conveyed in the elementary example above. But you’re already well on your way!

KEY TAKEAWAYS

- Data are raw facts that must be turned into information in order to be useful and valuable.
- Databases are created, maintained, and manipulated using programs called database management systems (DBMS), sometimes referred to as database software.



- All data fields in the same database have unique names, several data fields make up a data record, multiple data records make up a table or data file, and one or more tables or data files make up a database.
- Relational databases are the most common database format.

QUESTIONS AND EXERCISES

1. Define the following terms: table, record, field. Provide another name for each term along with your definition.
2. Answer the following questions using the course registration database system, diagramed above:
 - Imagine you also want to keep track of student majors. How would you do this? Would you modify an existing table? Would you add new tables? Why or why not?
 - Why do you suppose the system needs a “Course Title” table?
 - This database is simplified for our brief introduction. What additional data would you need to keep track of if this were a real course registration system? What changes would you make in the database above to account for these needs?
3. Research to find additional examples of organizations that made bad decisions based on bad data. Report your examples to your class. What was the end result of the examples you’re citing (e.g., loss, damage, or other outcome)? What could managers have done to prevent problems in the cases that you cited? What role did technology play in the examples that you cite? What role did people or procedural issues play?
4. Why is an understanding of database terms and technologies important, even for nontechnical managers and staff? Consider factors associated with both system use and system development. What other skills, beyond technology, may be important when engaged in data-driven decision making?



11.3 Where Does Data Come From?

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand various internal and external sources for enterprise data.
2. Recognize the function and role of data aggregators, the potential for leveraging third-party data, the strategic implications of relying on externally purchased data, and key issues associated with aggregators and firms that leverage externally sourced data.

Organizations can pull together data from a variety of sources. While the examples that follow aren't meant to be an encyclopedic listing of



possibilities, they will give you a sense of the diversity of options available for data gathering.

Transaction Processing Systems

For most organizations that sell directly to their customers, transaction processing systems (TPS) represent a fountain of potentially insightful data. Every time a consumer uses a point-of-sale system, an ATM, or a service desk, there's a transaction (some kind of business exchange) occurring, representing an event that's likely worth tracking. The cash register is the data generation workhorse of most physical retailers, and the primary source that feeds data to the TPS. But while TPS can generate a lot of bits, it's sometimes tough to match this data with a specific customer. For example, if you pay a retailer in cash, you're likely to remain a mystery to your merchant because your name isn't attached to your money. Grocers and retailers can tie you to cash transactions if they can convince you to use a loyalty card. Use one of these cards and you're in effect giving up information about yourself in exchange for some kind of financial incentive. The explosion in retailer cards is directly related to each firm's desire to learn more about you and to turn you into a more loyal and satisfied customer.

Some cards provide an instant discount (e.g., the CVS Pharmacy ExtraCare card), while others allow you to build up points over time (Best Buy's Reward Zone). The latter has the additional benefit of acting as a switching cost. A customer may think "I could get the same thing at Target, but at Best Buy, it'll increase my existing points balance and soon I'll get a cash back coupon."



Tesco: Tracked Transactions, Increased Insights, and Surging Sales

UK grocery giant Tesco, the planet's third-largest retailer, is envied worldwide for what analysts say is the firm's unrivaled ability to collect vast amounts of retail data and translate this into sales. [1]

Tesco's data collection relies heavily on its ClubCard loyalty program, an effort pioneered back in 1995. But Tesco isn't just a physical retailer. As the world's largest Internet grocer, the firm gains additional data from Web site visits, too. Remove products from your virtual shopping cart? Tesco can track this. Visited a product comparison page? Tesco watches which product you've chosen to go with and which you've passed over. Done your research online, then traveled to a store to make a purchase? Tesco sees this, too.

Tesco then mines all this data to understand how consumers respond to factors such as product mix, pricing, marketing campaigns, store layout, and Web design. Consumer-level targeting allows the firm to tailor its marketing messages to specific subgroups, promoting the right offer through the right channel at the right time and the right price. To get a sense of Tesco's laser-focused targeting possibilities, consider that the firm sends out close to ten million different, targeted offers each quarter. [2] Offer redemption rates are the best in the industry, with some coupons scoring an astronomical 90 percent usage! [3]

The firm's data driven management is clearly delivering results. In April 2009, while operating in the teeth of a global recession, Tesco posted record corporate profits and the highest earnings ever for a British retailer. [4]



Enterprise Software (CRM, SCM, and ERP)

Firms increasingly set up systems to gather additional data beyond conventional purchase transactions or Web site monitoring. CRM or customer relationship management systems are often used to empower employees to track and record data at nearly every point of customer contact. Someone calls for a quote? Brings a return back to a store? Writes a complaint e-mail? A well-designed CRM system can capture all these events for subsequent analysis or for triggering follow-up events.

Enterprise software includes not just CRM systems, as but also categories that touch every aspect of the value chain, including supply chain management (SCM) and enterprise resource planning (ERP) systems. More importantly, enterprise software tends to be more integrated and standardized than the prior era of proprietary systems that many firms developed themselves. This integration helps in combining data across business units and functions, and in getting that data into a form where it can be turned into information (for more on enterprise systems, see [Chapter 9 "Understanding Software: A Primer for Managers"](#)).

Surveys

Sometimes firms supplement operational data with additional input from surveys and focus groups. Oftentimes, direct surveys can tell you what your cash register can't. Zara store managers informally survey customers in order to help shape designs and product mix. Online grocer FreshDirect (see [Chapter 2 "Strategy and Technology"](#)) surveys customers weekly and has used this feedback to drive initiatives from reducing packaging size to including star ratings on produce. ^[5] Many CRM products also have survey capabilities that allow for additional data gathering at all points of customer contact.



Can Technology “Cure” U.S. Healthcare?

The U.S. healthcare system is broken. It’s costly, inefficient, and problems seem to be getting worse. Estimates suggest that healthcare spending makes up a whopping 18 percent of U.S. gross domestic product. [6] U.S. automakers spend more on healthcare than they do on steel. [7] Even more disturbing, it’s believed that medical errors cause as many as ninety-eight thousand unnecessary deaths in the United States each year, more than motor vehicle accidents, breast cancer, or AIDS. [8]

For years it’s been claimed that technology has the potential to reduce errors, improve health care quality, and save costs. Now pioneering hospital networks and technology companies are partnering to help tackle cost and quality issues. For a look at possibilities for leveraging data throughout the doctor-patient value chain, consider the “event-driven medicine” system built by Dr. John Halamka and his team at Boston’s Beth Israel Deaconess Medical Center (part of the Harvard Medical School network).

When docs using Halamka’s system encounter a patient with a chronic disease, they generate a decision support “screening sheet.” Each event in the system: an office visit, a lab results report (think the medical equivalent of transactions and customer interactions), updates the patient database. Combine that electronic medical record information with artificial intelligence on best practice, and the system can offer recommendations for care, such as, “Patient is past due for an eye exam” or, “Patient should receive pneumovax [a vaccine against infection] this season.” [9] The systems don’t replace decision making by doctors and nurses, but they do help to ensure that key issues are on a provider’s radar. More efficiencies and error checks show up when prescribing drugs. Docs are presented with a list of medications covered by that patient’s insurance, allowing them to choose quality options while controlling costs. Safety



issues, guidelines, and best practices are also displayed. When correct, safe medication in the right dose is selected, the electronic prescription is routed to the patients' pharmacy of choice. As Halamka puts it, going from "doctor's brain to patients vein" without any of that messy physician handwriting, all while squeezing out layers where errors from human interpretation or data entry might occur.

President Obama believes technology initiatives can save healthcare as much as one hundred twenty billion dollars a year, or roughly two thousand five hundred dollars per family. ^[10] An aggressive number, to be sure. But with such a large target to aim at, it's no wonder that nearly every major technology company now has a health solutions group. Microsoft and Google even offer competing systems for electronically storing and managing patient health records. If systems like Halamka's and others realize their promise, big benefits may be just around the corner.

External Sources

Sometimes it makes sense to combine a firm's data with bits brought in from the outside. Many firms, for example, don't sell directly to consumers (this includes most drug companies and packaged goods firms). If your firm has partners that sell products for you, then you'll likely rely heavily on data collected by others.

Data bought from sources available to all might not yield competitive advantage on its own, but it can provide key operational insight for increased efficiency and cost savings. And when combined with a firm's unique data assets, it may give firms a high-impact edge.

Consider restaurant chain Brinker, a firm that runs seventeen hundred eateries in twenty-seven countries under the Chili's, On The Border, and Maggiano's brands. Brinker (whose ticker symbol is EAT), supplements their own data with external feeds on weather, employment statistics, gas



prices, and other factors, and uses this in predictive models that help the firm in everything from determining staffing levels to switching around menu items. ^[11]

In another example, Carnival Cruise Lines combines its own customer data with third-party information tracking household income and other key measures. These data play a key role in a recession, since they help the firm target limited marketing dollars on those past customers that are more likely to be able to afford to go on a cruise. So far it's been a winning approach. For three years in a row, the firm has experienced double-digit increases in bookings by repeat customers. ^[12]

Who's Collecting Data about You?

There's a thriving industry collecting data about you. Buy from a catalog, fill out a warranty card, or have a baby, and there's a very good chance that this event will be recorded in a database somewhere, added to a growing digital dossier that's made available for sale to others. If you've ever gotten catalogs, coupons, or special offers from firms you've never dealt with before, this was almost certainly a direct result of a behind-the-scenes trafficking in the "digital you."

Firms that trawl for data and package them up for resale are known as data aggregators. They include Acxiom, a \$1.3 billion a year business that combines public source data on real estate, criminal records, and census reports, with private information from credit card applications, warranty card surveys, and magazine subscriptions. The firm holds data profiling some two hundred million Americans. ^[13]

Or maybe you've heard of Lexis-Nexis. Many large universities subscribe to the firm's electronic newspaper, journal, and magazine databases. But the firm's parent, Reed Elsevier, is a data sales giant, with divisions packaging criminal records, housing information, and additional data used to uncover corporate fraud and other risks. In February, 2008, the firm got even more



data rich, acquiring Acxiom competitor ChoicePoint for \$4.1 billion. With that kind of money involved, it's clear that data aggregation is very big business. [14]

The Internet also allows for easy access to data that had been public but otherwise difficult to access. For one example, consider home sale prices and home value assessments. While technically in the public record, someone wanting this information previously had to traipse down to their Town Hall and speak to a clerk, who would hand over a printed log book. Not exactly a Google-speed query. Contrast this with a visit to Zillow.com. The free site lets you pull up a map of your town and instantly peek at how much your neighbors paid for their homes. And it lets them see how much you paid for yours, too.

Computerworld's Robert Mitchell uncovered a more disturbing issue when public record information is made available online. His New Hampshire municipality had digitized and made available some of his old public documents without obscuring that holy grail for identity thieves, his social security number. [15]

Then there are accuracy concerns. A record incorrectly identifying you as a cat lover is one thing, but being incorrectly named to the terrorist watch list is quite another. During a five-week period airline agents tried to block a particularly high profile U.S. citizen from boarding airplanes on five separate occasions because his name resembled an alias used by a suspected terrorist. That citizen? The late Ted Kennedy, who at the time was the senior U.S. senator from Massachusetts. [16]

For the data trade to continue, firms will have to treat customer data as the sacred asset it is. Step over that "creep-out" line, and customers will push back, increasingly pressing for tighter privacy laws. Data aggregator Intellius used to track cell phone customers, but backed off in the face of customer outrage and threatened legislation.

Another concern—sometimes data aggregators are just plain sloppy, committing errors that can be costly for the firm and potentially devastating for victimized users. For example, from 2002 to 2003, a hacker stole 1.6 billion records from Acxiom; while in 2005, ChoicePoint accidentally sold records on one hundred and forty five thousand individuals to a cybercrime identity theft ring. The ChoicePoint case resulted in a fifteen-million-dollar fine from the Federal Trade Commission. ^[17] Just because you can gather data and traffic in bits doesn't mean that you should. Any data-centric effort should involve input not only from business and technical staff, but from the firm's legal team, as well (for more, see the box "Privacy Regulation: A Moving Target").

Privacy Regulation: A Moving Target

New methods for tracking and gathering user information appear daily, testing user comfort levels. For example, the firm Umbria uses software to analyze millions of blog and forum posts every day, using sentence structure, word choice, and quirks in punctuation to determine a blogger's gender, age, interests, and opinions. In 2009, Apple introduced facial recognition software while integrating iPhoto into Facebook. It's quite possible that in the future, someone will be able to upload a photo to a service and direct it to find all the accessible photos and video on the Internet that match that person's features. And while targeting is getting easier, a Carnegie Mellon study showed that it doesn't take much to find someone with a minimum of data. Simply by knowing gender, birth date, and postal zip code, 87 percent of people in the United States could be pinpointed by name. ^[18] Another study showed that publicly available data on state and date of birth could be used to predict U.S. Social Security numbers—a potential gateway to identity theft. ^[19]



Some feel that Moore's Law, the falling cost of storage, and the increasing reach of the Internet have us on the cusp of a privacy train wreck. And that may inevitably lead to more legislation that restricts data-use possibilities. Noting this, strategists and technologists need to be fully aware of the legal environment their systems face (see [Chapter 8 "Google: Search, Online Advertising, and Beyond..."](#) for examples and discussion) and consider how such environments may change in the future. Many industries have strict guidelines on what kind of information can be collected and shared. For example, HIPAA (the U.S. Health Insurance Portability and Accountability Act) includes provisions governing data use and privacy among healthcare providers, insurers, and employers. The financial industry has strict requirements for recording and sharing communications between firm and client (among many other restrictions). There are laws limiting the kinds of information that can be gathered on younger Web surfers. And there are several laws operating at the state level as well. International laws also differ from those in the United States. Europe, in particular, has a strict European Privacy Directive. The directive includes governing provisions that limit data collection, require notice and approval of many types of data collection, and require firms to make data available to customers with mechanisms for stopping collection efforts and correcting inaccuracies at customer request. Data-dependent efforts plotted for one region may not fully translate in another effort if the law limits key components of technology use. The constantly changing legal landscape also means that what works today might not be allowed in the future. Firms beware—the public will almost certainly demand tighter controls if the industry is perceived as behaving recklessly or inappropriately with customer data.



KEY TAKEAWAYS

- For organizations that sell directly to their customers, transaction processing systems (TPS) represent a source of potentially useful data.
- Grocers and retailers can link you to cash transactions if they can convince you to use a loyalty card which, in turn, requires you to give up information about yourself in exchange for some kind of financial incentive such as points or discounts.
- Enterprise software (CRM, SCM, and ERP) is a source for customer, supply chain, and enterprise data.
- Survey data can be used to supplement a firm's operational data.
- Data obtained from outside sources, when combined with a firm's internal data assets, can give the firm a competitive edge.
- Data aggregators are part of a multibillion-dollar industry that provides genuinely helpful data to a wide variety of organizations.
- Data that can be purchased from aggregators may not in and of itself yield sustainable competitive advantage since others may have access to this data, too. However, when combined with a firm's proprietary data or integrated with a firm's proprietary procedures or other assets, third-party data can be a key tool for enhancing organizational performance.
- Data aggregators can also be quite controversial. Among other things, they represent a big target for identity thieves, are method for spreading potentially incorrect data, and raise privacy concerns.
- Firms that mismanage their customer data assets risk lawsuits, brand damage, lower sales, fleeing customers, and can prompt more restrictive legislation.
- Further raising privacy issues and identity theft concerns, recent studies have shown that in many cases it is possible to pinpoint users through allegedly anonymous data, and to guess social security numbers from public data.

- New methods for tracking and gathering user information are raising privacy issues which possibly will be addressed through legislation that restricts data use.

QUESTIONS AND EXERCISES

1. Why would a firm use a loyalty card? What is the incentive for the firm? What is the incentive for consumers to opt-in and use loyalty cards? What kinds of strategic assets can these systems create?
2. In what ways does Tesco gather data? Can other firms match this effort? What other assets does Tesco leverage that helps the firm remain among top performing retailers worldwide?
3. Make a list of the kind of data you might give up when using a cash register, a Web site, a loyalty card, when calling a firm's customer support line. How might firms leverage this data to better serve you and improve their performance?
4. Are you concerned by any of the data use possibilities that you outlined in prior questions, discussed in this chapter, or that you've otherwise read about or encountered? If you are concerned, why? If not, why not? What might firms, governments, and consumers do to better protect consumers?
5. What are some of the sources data aggregators tap to collect information?
6. Privacy laws are in a near constant state of flux. Conduct research to identify the current state of privacy law. Has major legislation recently been proposed or approved? What are the implications for firms operating in effected industries? What are the potential benefits to consumers? Do consumers lose anything from this legislation?
7. Self-regulation is often proposed as an alternative to legislative efforts. What kinds of efforts would provide "teeth" to self-regulation. Are there



steps firms could do to make you believe in their ability to self-regulate?
Why or why not?

8. What is HIPPA? What industry does it impact?

9. How do international privacy laws differ from U.S. privacy laws?

11.4 Data Rich, Information Poor



LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Know and be able to list the reasons why many organizations have data that can't be converted to actionable information.
2. Understand why transactional databases can't always be queried and what needs to be done to facilitate effective data use for analytics and business intelligence.
3. Recognize key issues surrounding data and privacy legislation.

Despite being awash in data, many organizations are data rich but information poor. A survey by consulting firm Accenture found 57 percent of companies reporting that they didn't have a beneficial, consistently updated, companywide analytical capability. Among major decisions, only 60 percent were backed by analytics—40 percent were made by intuition and gut instinct. [1] The big culprit limiting BI initiatives is getting data into a form where it can be used, analyzed, and turned into information. Here's a look at some factors holding back information advantages.

Incompatible Systems

Just because data is collected doesn't mean it can be used. This limit is a big problem for large firms that have legacy systems, outdated information systems that were not designed to share data, aren't compatible with newer technologies, and aren't aligned with the firm's current business needs. The problem can be made worse by mergers and acquisitions, especially if a firm depends on operational systems that are incompatible with its partner. And the elimination of incompatible systems isn't just a technical issue. Firms might be under extended agreement with different vendors or outsourcers, and breaking a contract or invoking an escape clause may be costly. Folks working in M&A (the area of investment banking focused on



valuing and facilitating mergers and acquisitions) beware—it's critical to uncover these hidden costs of technology integration before deciding if a deal makes financial sense.

Legacy Systems: A Prison for Strategic Assets

The experience of one *Fortune* 100 firm that your author has worked with illustrates how incompatible information systems can actually hold back strategy. This firm was the largest in its category, and sold identical commodity products sourced from its many plants worldwide. Being the biggest should have given the firm scale advantages. But many of the firm's manufacturing facilities and international locations developed or purchased separate, incompatible systems. Still more plants were acquired through acquisition, each coming with its own legacy systems.

The plants with different information systems used *different* part numbers and naming conventions even though they sold *identical* products. As a result, the firm had no timely information on how much of a particular item was sold to which worldwide customers. The company was essentially operating as a collection of smaller, regional businesses, rather than as the worldwide behemoth that it was.

After the firm developed an information system that standardized data across these plants, it was, for the first time, able to get a single view of worldwide sales. The firm then used this data to approach their biggest customers, negotiating lower prices in exchange for increased commitments in worldwide purchasing. This trade let the firm take share from regional rivals. It also gave the firm the ability to shift manufacturing capacity globally, as currency prices, labor conditions, disaster, and other factors impacted sourcing. The new information system in effect liberated the latent strategic asset of scale, increasing sales by well over a billion and a half dollars in the four years following implementation.



Operational Data Can't Always Be Queried

Another problem when turning data into information is that most transactional databases aren't set up to be simultaneously accessed for reporting and analysis. When a customer buys something from a cash register, that action may post a sales record and deduct an item from the firm's inventory. In most TPS systems, requests made to the database can usually be performed pretty quickly—the system adds or modifies the few records involved and it's done—in and out in a flash.

But if a manager asks a database to analyze historic sales trends showing the most and least profitable products over time, they may be asking a computer to look at thousands of transaction records, comparing results, and neatly ordering findings. That's not a quick in-and-out task, and it may very well require significant processing to come up with the request. Do this against the very databases you're using to record your transactions, and you might grind your computers to a halt.

Getting data into systems that can support analytics is where data warehouses and data marts come in, the topic of our next section.

KEY TAKEAWAYS

- A major factor limiting business intelligence initiatives is getting data into a form where it can be used (i.e., analyzed and turned into information).
- Legacy systems often limit data utilization because they were not designed to share data, aren't compatible with newer technologies, and aren't aligned with the firm's current business needs.
- Most transactional databases aren't set up to be simultaneously accessed for reporting and analysis. In order to run analytics the data must first be ported to a data warehouse or data mart.

QUESTIONS AND EXERCISES

1. How might information systems impact mergers and acquisitions? What are the key issues to consider?
2. Discuss the possible consequences of a company having multiple plants, each with a different information system using different part numbers and naming conventions for identical products.
3. Why does it take longer, and require more processing power, to analyze sales trends by region and product, as opposed to posting a sales transaction?

11.5 Data Warehouses and Data Marts

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand what data warehouses and data marts are, and the purpose they serve.
2. Know the issues that need to be addressed in order to design, develop, deploy, and maintain data warehouses and data marts.

Since running analytics against transactional data can bog down a system, and since most organizations need to combine and reformat data from multiple sources, firms typically need to create separate data repositories



for their reporting and analytics work—a kind of staging area from which to turn that data into information.

Two terms you'll hear for these kinds of repositories are data warehouse and data mart. A data warehouse is a set of databases designed to support decision making in an organization. It is structured for fast online queries and exploration. Data warehouses may aggregate enormous amounts of data from many different operational systems. A data mart is a database focused on addressing the concerns of a specific problem (e.g., increasing customer retention, improving product quality) or business unit (e.g., marketing, engineering).

Marts and warehouses may contain huge volumes of data. For example, a firm may not need to keep large amounts of historical point-of-sale or transaction data in its operational systems, but it might want past data in its data mart so that managers can hunt for patterns and trends that occur over time.

It's easy for firms to get seduced by a software vendor's demonstration showing data at your fingertips, presented in pretty graphs. But as mentioned earlier, getting data in a format that can be used for analytics is hard, complex, and challenging work. Large data warehouses can cost millions and take years to build. Every dollar spent on technology may lead to five to seven more dollars on consulting and other services. ^[1]

Most firms will face a tradeoff—do we attempt a large-scale integration of the whole firm, or more targeted efforts with quicker payoffs? Firms in fast-moving industries or with particularly complex businesses may struggle to get sweeping projects completed in enough time to reap benefits before business conditions change. Most consultants now advise smaller projects with narrow scope driven by specific business goals. ^[2]

Firms can eventually get to a unified data warehouse but it may take time. Even analytics king Wal-Mart is just getting to that point. In 2007, it was



reported that Wal-Mart had seven hundred different data marts and hired Hewlett Packard for help in bringing the systems together to form a more integrated data warehouse. [3]

The old saying from the movie *Field of Dreams*, “If you build it, they will come,” doesn’t hold up well for large-scale data analytics projects. This work should start with a clear vision with business-focused objectives. When senior executives can see objectives illustrated in potential payoff, they’ll be able to champion the effort, and experts agree, having an executive champion is a key success factor. Focusing on business issues will also drive technology choice, with the firm better able to focus on products that best fit its needs.

Once a firm has business goals and hoped-for payoffs clearly defined, it can address the broader issues needed to design, develop, deploy, and maintain its system: [4]

- *Data relevance.* What data is needed to compete on analytics and to meet our current and future goals?
- *Data sourcing.* Can we even get the data we’ll need? Where can this data be obtained from? Is it available via our internal systems? Via third-party data aggregators? Via suppliers or sales partners? Do we need to set up new systems, surveys, and other collection efforts to acquire the data we need?
- *Data quantity.* How much data is needed?
- *Data quality.* Can our data be trusted as accurate? Is it clean, complete, and reasonably free of errors? How can the data be made more accurate and valuable for analysis? Will we need to ‘scrub,’ calculate, and consolidate data so that it can be used?
- *Data hosting.* Where will the systems be housed? What are the hardware and networking requirements for the effort?
- *Data governance.* What rules and processes are needed to manage data from its creation through its retirement? Are there operational issues



(backup, disaster recovery)? Legal issues? Privacy issues? How should the firm handle security and access?

For some perspective on how difficult this can be, consider that an executive from one of the largest U.S. banks once lamented at how difficult it was to get his systems to do something as simple as properly distinguishing between men and women. The company's customer-focused data warehouse drew data from thirty-six separate operational systems—bank teller systems, ATMs, student loan reporting systems, car loan systems, mortgage loan systems, and more. Collectively these legacy systems expressed gender in *seventeen* different ways: “M” or “F”; “m” or “f”; “Male” or “Female”; “MALE” or “FEMALE”; “1” for man, “o” for woman; “o” for man, “1” for woman and more, plus various codes for “unknown.” The best math in the world is of no help if the values used aren't any good. There's a saying in the industry, “garbage in, garbage out.”

E-discovery: Supporting Legal Inquiries

Data archiving isn't just for analytics. Sometimes the law requires organizations to dive into their electronic records. E-discovery refers to identifying and retrieving relevant electronic information to support litigation efforts. E-discovery is something a firm should account for in its archiving and data storage plans. Unlike analytics that promise a boost to the bottom line, there's no profit in complying with a judge's order—it's just a sunk cost. But organizations can be compelled by court order to scavenge their bits, and the cost to uncover difficult to access data can be significant, if not planned for in advance.

In one recent example, the Office of Federal Housing Enterprise Oversight (OFHEO) was subpoenaed for documents in litigation involving mortgage firms Fannie Mae and Freddie Mac. Even though the OFHEO wasn't a



party in the lawsuit, the agency had to comply with the search—an effort that cost six million dollars, a full 9 percent of its total yearly budget. [5]

KEY TAKEAWAYS

- Data warehouses and data marts are repositories for large amounts of transactional data awaiting analytics and reporting.
- Large data warehouses are complex, can cost millions, and take years to build.

QUESTIONS AND EXERCISES

1. List the issues that need to be addressed in order to design, develop, deploy, and maintain data warehouses and data marts.
2. What is meant by “data relevance”?
3. What is meant by “data governance”?
4. What is the difference between a data mart and a data warehouse?
5. Why are data marts and data warehouses necessary? Why can’t an organization simply query its transactional database?
6. How can something as simple as customer gender be difficult to for a large organization to establish in a data warehouse?



11.6 The Business Intelligence Toolkit

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Know the tools that are available to turn data into information.
2. Identify the key areas where businesses leverage data mining.
3. Understand some of the conditions under which analytical models can fail.
4. Recognize major categories of artificial intelligence and understand how organizations are leveraging this technology.

So far we've discussed where data can come from, and how we can get data into a form where we can use it. But how, exactly, do firms turn that data into information? That's where the various software tools of business intelligence (BI) and analytics come in. Potential products in the business intelligence toolkit range from simple spreadsheets to ultrasophisticated data mining packages leveraged by teams employing "rocket-science" mathematics.

Query and Reporting Tools

The idea behind query and reporting tools is to present users with a subset of requested data, selected, sorted, ordered, calculated, and compared, as



needed. Managers use these tools to see and explore what's happening inside their organizations.

Canned reports provide regular summaries of information in a predetermined format. They're often developed by information systems staff and formats can be difficult to alter. By

contrast, ad hoc reporting tools allow users to dive in and create their own reports, selecting fields, ranges, and other parameters to build their own reports on the fly. Dashboards provide a sort of heads-up display of critical indicators, letting managers get a graphical glance at key performance metrics. Some tools may allow data to be exported into spreadsheets. Yes, even the lowly spreadsheet can be a powerful tool for modeling "what if" scenarios and creating additional reports (of course be careful, if data can be easily exported then it can potentially leave the firm dangerously exposed, raising privacy, security, legal, and competitive concerns).

A subcategory of reporting tools is referred to as online analytical processing (OLAP)(pronounced "oh-lap"). Data used in OLAP reporting is usually sourced from standard relational databases, but it's calculated and summarized in advance, across multiple dimensions, with the data stored in a special database called a data cube. This extra setup step makes OLAP fast (sometimes one thousand times faster than performing comparable queries against conventional relational databases). Given this kind of speed boost, it's not surprising that data cubes for OLAP access are often part of a firm's data mart and data warehouse efforts.

A manager using an OLAP tool can quickly explore and compare data across multiple factors such as time, geography, product lines, and so on. In fact, OLAP users often talk about how they can "slice and dice" their data, "drilling down" inside the data to uncover new insights. And while conventional reports are usually presented as summarized list of information, OLAP results look more like a spreadsheet, with the various



dimensions of analysis in rows and columns, with summary values at the intersection.

Public Sector Reporting Tools in Action: Fighting Crime and Fighting Waste

Access to ad hoc query and reporting tools can empower all sorts of workers. Consider what analytics tools have done for the police force in Richmond, Virginia. The city provides department investigators with access to data from internal sources such as 911 logs and police reports, and combines this with outside data including neighborhood demographics, payday schedules, weather reports, traffic patterns, sports events, and more.

Experienced officers dive into this data, exploring when and where crimes occur. These insights help the department decide how to allocate its limited policing assets to achieve the biggest impact. While IT staffers put the system together, the tools are actually used by officers with expertise in fighting street crime—the kinds of users with the knowledge to hunt down trends and interpret the causes behind the data. And it seems this data helps make smart cops even smarter—the system is credited with delivering a single-year crime-rate reduction of 20 percent. ^[1]

As it turns out, what works for cops also works for bureaucrats. When administrators for Albuquerque were given access to ad hoc reporting systems, they uncovered all sorts of anomalies, prompting excess spending cuts on everything from cell phone usage to unnecessarily scheduled overtime. And once again, BI performed for the public sector. The Albuquerque system delivered the equivalent of two million dollars in savings in just the first three weeks it was used. ^[2]



Data Mining

While reporting tools can help users explore data, modern datasets can be so large that it might be impossible for humans to spot underlying trends. That's where data mining can help. Data mining is the process of using computers to identify hidden patterns and to build models from large data sets.

Some of the key areas where businesses are leveraging data mining include the following:

- *Customer segmentation*—figuring out which customers are likely to be the most valuable to a firm.
- *Marketing and promotion targeting*—identifying which customers will respond to which offers at which price at what time.
- *Market basket analysis*—determining which products customers buy together, and how an organization can use this information to cross-sell more products or services.
- *Collaborative filtering*—personalizing an individual customer's experience based on the trends and preferences identified across similar customers.
- *Customer churn*—determining which customers are likely to leave, and what tactics can help the firm avoid unwanted defections.
- *Fraud detection*—uncovering patterns consistent with criminal activity.
- *Financial modeling*—building trading systems to capitalize on historical trends.
- *Hiring and promotion*—identifying characteristics consistent with employee success in the firm's various roles.

For data mining to work, two critical conditions need to be present: (1) the organization must have clean, consistent data, and (2) the events in that data should reflect current and future trends. The recent financial crisis



provides lessons on what can happen when either of these conditions isn't met.

First let's look at problems with using bad data. A report in the *New York Times* has suggested that in the period leading up to the 2008 financial crisis, some banking executives deliberately deceived risk management systems in order to skew capital-on-hand requirements. This deception let firms load up on risky debt, while carrying less cash for covering losses. [3] Deceive your systems with bad data and your models are worthless. In this case, wrong estimates from bad data left firms grossly overexposed to risk. When debt defaults occurred; several banks failed, and we entered the worst financial crisis since the Great Depression.

Now consider the problem of historical consistency: Computer-driven investment models can be very effective when the market behaves as it has in the past. But models are blind when faced with the equivalent of the "hundred-year flood" (sometimes called "black swans"); events so extreme and unusual that they never showed up in the data used to build the model. We saw this in the late 1990s with the collapse of the investment firm Long Term Capital Management. LTCM was started by Nobel Prize-winning economists, but when an unexpected Russian debt crisis caused the markets to move in ways not anticipated by its models, the firm lost 90 percent of its value in less than two months. The problem was so bad that the Fed had to step in to supervise the firm's multibillion-dollar bailout. Fast forward a decade to the banking collapse of 2008, and we again see computer-driven trading funds plummet in the face of another unexpected event—the burst of the housing bubble. [4]

Data mining presents a host of other perils, as well. It's possible to over-engineer a model, building it with so many variables that the solution arrived at might only work on the subset of data you've used to create it. You might also be looking at a random but meaningless statistical fluke. In

demonstrating how flukes occur, one quantitative investment manager uncovered a correlation that at first glance appeared statistically to be a particularly strong predictor for historical prices in the S&P 500 stock index. That predictor? Butter production in Bangladesh. ^[5] Sometimes durable and useful patterns just aren't in your data.

One way to test to see if you're looking at a random occurrence in the numbers is to divide your data, building your model with one portion of the data, and using another portion to verify your results. This is the approach Netflix has used to test results achieved by teams in the Netflix Prize, the firm's million-dollar contest for improving the predictive accuracy of its movie recommendation engine (see [Chapter 3 "Netflix: David Becomes Goliath"](#)).

Finally, sometimes a pattern is uncovered but determining the best choice for a response is less clear. As an example, let's return to the data-mining wizards at Tesco. An analysis of product sales data showed several money-losing products, including a type of bread known as "milk loaf." Drop those products, right? Not so fast. Further analysis showed milk loaf was a "destination product" for a loyal group of high-value customers, and that these customers would shop elsewhere if milk loaf disappeared from Tesco shelves. The firm kept the bread as a loss-leader and retained those valuable milk loaf fans. ^[6] Data miner beware—first findings don't always reveal an optimal course of action.

This last example underscores the importance of recruiting a data mining and business analytics team that possesses three critical skills: information technology (for understanding how to pull together data, and for selecting analysis tools), statistics (for building models and interpreting the strength and validity of results), and business knowledge (for helping set system goals, requirements, and offering deeper insight into what the data really

says about the firm's operating environment). Miss one of these key functions and your team could make some major mistakes.

While we've focused on tools in our discussion above, many experts suggest that business intelligence is really an organizational processes as much as it is a set of technologies. Having the right team is critical in moving the firm from goal setting through execution and results.

Artificial Intelligence

Data Mining has its roots in a branch of computer science known as artificial intelligence (or AI). The goal of AI is create computer programs that are able to mimic or improve upon functions of the human brain. Data mining can leverageneural networks or other advanced algorithms and statistical techniques to hunt down and expose patterns, and build models to exploit findings.

Expert systems are AI systems that leverage rules or examples to perform a task in a way that mimics applied human expertise. Expert systems are used in tasks ranging from medical diagnoses to product configuration.

Genetic algorithms are model building techniques where computers examine many potential solutions to a problem, iteratively modifying (mutating) various mathematical models, and comparing the mutated models to search for a best alternative. Genetic algorithms have been used to build everything from financial trading models to handling complex airport scheduling, to designing parts for the international space station. [7]

While AI is not a single technology, and not directly related to data creation, various forms of AI can show up as part of analytics products, CRM tools, transaction processing systems, and other information systems.

KEY TAKEAWAYS

- Canned and ad hoc reports, digital dashboards, and OLAP are all used to transform data into information.
- OLAP reporting leverage data cubes, which take data from standard relational databases, calculating and summarizing data for super-fast reporting access. OLAP tools can present results through multidimensional graphs, or via spreadsheet-style cross-tab reports.
- Modern datasets can be so large that it might be impossible for humans to spot underlying trends without the use of data mining tools.
- Businesses are using data mining to address issues in several key areas including customer segmentation, marketing and promotion targeting, collaborative filtering, and so on.
- Models influenced by bad data, missing or incomplete historical data, and over-engineering are prone to yield bad results.
- One way to test to see if you're looking at a random occurrence in your data is to divide your data, building your model with one portion of the data, and using another portion to verify your results.
- Analytics may not always provide the total solution for a problem. Sometimes a pattern is uncovered, but determining the best choice for a response is less clear.
- A competent business analytics team should possess three critical skills: information technology, statistics, and business knowledge.

QUESTIONS AND EXERCISES

1. What are some of the tools used to convert data into information?
2. What is the difference between a canned reports and an ad hoc reporting?
3. How do reports created by OLAP differ from most conventional reports?
4. List the key areas where businesses are leveraging data mining.
5. What is market basket analysis?



6. What is customer churn?
7. For data mining to work, what two critical data-related conditions must be present?
8. Discuss occurrences of model failure caused by missing or incomplete historical data.
9. Discuss Tesco's response to their discovery that "milk loaf" was a money-losing product.
10. List the three critical skills a competent business analytics team should possess.
11. Do any of the products that you use leverage artificial intelligence? What kinds of AI might be used in Netflix's movie recommendation system, Apple's iTunes Genius playlist builder, or Amazon's Web site personalization? What kind of AI might help a physician make a diagnosis or help an engineer configure a complicated product in the field?

11.7 Data Asset in Action: Technology and the Rise of Wal-Mart

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand how Wal-Mart has leveraged information technology to become the world's largest retailer.
2. Be aware of the challenges that face Wal-Mart in the years ahead.

Wal-Mart demonstrates how a physical product retailer can create and leverage a data asset to achieve world-class supply chain efficiencies targeted primarily at driving down costs.



Wal-Mart isn't just the largest retailer in the world, over the past several years it has popped in and out of the top spot on the *Fortune 500* list—meaning that the firm has had revenues greater than *any* firm in the United States. Wal-Mart is so big that in three months it sells more than a whole year's worth of sales at number two U.S. retailer, Home Depot. [1]

At that size, it's clear that Wal-Mart's key source of competitive advantage is scale. But firms don't turn into giants overnight. Wal-Mart grew in large part by leveraging information systems to an extent never before seen in the retail industry. Technology tightly coordinates the Wal-Mart value chain from tip to tail, while these systems also deliver a mineable data asset that's unmatched in U.S. retail. To get a sense of the firm's overall efficiencies, at the end of the prior decade a McKinsey study found that Wal-Mart was responsible for some 12 percent of the productivity gains in the *entire* U.S. economy.[2] The firm's capacity as a systems innovator is so respected that many senior Wal-Mart IT executives have been snatched up for top roles at Dell, HP, Amazon, and Microsoft. And lest one think that innovation is the province of only those located in the technology hubs of Silicon Valley, Boston, and Seattle, remember that Wal-Mart is headquartered in Bentonville, Arkansas.

A Data-Driven Value Chain

The Wal-Mart efficiency dance starts with a proprietary system called Retail Link, a system originally developed in 1991 and continually refined ever since. Each time an item is scanned by a Wal-Mart cash register, Retail Link not only records the sale, it also automatically triggers inventory reordering, scheduling, and delivery. This process keeps shelves stocked, while keeping inventories at a minimum. An AMR report ranked Wal-Mart as having the seventh best supply chain in the country (the only other retailer in the top twenty was Tesco, at number fifteen). [3] The firm's



annual inventory turnover ratio of 8.5 means that Wal-Mart sells the equivalent of its entire inventory roughly every six weeks (by comparison, Target's turnover ratio is 6.4, Sears' is 3.4, and the average for U.S. retail is less than 2). [4]

Back-office scanners keep track of inventory as supplier shipments come in. Suppliers are rated based on timeliness of deliveries, and you've got to be quick to work with Wal-Mart. In order to avoid a tractor-trailer traffic jam in store parking lots, deliveries are choreographed to arrive at intervals less than ten minutes apart. When Levi's joined Wal-Mart, the firm had to guarantee it could replenish shelves every two days—no prior retailer had required a shorter than five day window from Levi's. [5]

Wal-Mart has been a catalyst for technology adoption among its suppliers. The firm is currently leading an adoption effort that requires partners to leverage RFID technology to track and coordinate inventories. While the rollout has been slow, a recent P&G trial showed RFID boosted sales nearly 20 percent by ensuring that inventory was on shelves and located where it should be. [6]

Data Mining Prowess

Wal-Mart also mines its mother lode of data to get its product mix right under all sorts of varying environmental conditions, protecting the firm from “a retailer's twin nightmares: too much inventory, or not enough.” [7] For example, the firm's data mining efforts informed buyers that customers stock up on certain products in the days leading up to predicted hurricanes. Bumping up prestorm supplies of batteries and bottled water was a no brainer, but the firm also learned that Pop-Tarts sales spike seven fold before storms hit, and that beer is the top prestorm seller. This insight has led to truckloads full of six packs and toaster

pastries streaming into gulf states whenever word of a big storm surfaces. [8]

Data mining also helps the firm tighten operational forecasts, helping to predict things like how many cashiers are needed at a given store at various times of day throughout the year. Data drives the organization, with mined reports forming the basis of weekly sales meetings, as well as executive strategy sessions.

Sharing Data, Keeping Secrets

While Wal-Mart is demanding of its suppliers, it also shares data with them, too. Data can help firms become more efficient so that Wal-Mart can keep dropping prices, and data can help firms uncover patterns that help suppliers sell more. P&G's Gillette unit, for example, claims to have mined Wal-Mart data to develop promotions that increased sales as much as 19 percent. More than seventeen thousand suppliers are given access to their products' Wal-Mart performance across metrics that include daily sales, shipments, returns, purchase orders, invoices, claims and forecasts. And these suppliers collectively interrogate Wal-Mart data warehouses to the tune of twenty-one million queries a year. [9]

While Wal-Mart shares sales data with relevant suppliers, the firm otherwise fiercely guards this asset. Many retailers pool their data by sharing it with information brokers like Information Resources and ACNielsen. This sharing allows smaller firms to pool their data to provide more comprehensive insight on market behavior. But Wal-Mart stopped data sharing data with these agencies years ago. The firm's scale is so big, the additional data provided by brokers wasn't adding much value, and it no longer made sense to allow competitors access to what was happening in its own huge chunk of retail sales.



Other aspects of the firm's technology remain under wraps, too. Wal-Mart custom builds large portions of its information systems to keep competitors off its trail. As for infrastructure secrets, the Wal-Mart Data Center in McDonald County, Missouri, was considered so off limits that the county assessor was required to sign a nondisclosure statement before being allowed on-site to estimate property value. ^[10]

Challenges Abound

But despite success, challenges continue. While Wal-Mart grew dramatically throughout the 1990s, the firm's U.S. business has largely matured. And as a mature business it faces a problem not unlike the example of Microsoft discussed at the end of [Chapter 8 "Google: Search, Online Advertising, and Beyond..."](#); Wal-Mart needs to find huge markets or dramatic cost savings in order to boost profits and continue to move its stock price higher.

The firm's aggressiveness and sheer size also increasingly make Wal-Mart a target for criticism. Those low prices come at a price, and the firm has faced accusations of subpar wages and remains a magnet for union activists. Others had identified poor labor conditions at some of the firm's contract manufacturers. Suppliers that compete for Wal-Mart's business are often faced with a catch-22. If they bypass Wal-Mart they miss out on the largest single chunk of world retail sales. But if they sell to Wal-Mart, the firm may demand prices so aggressively low that suppliers end up cannibalizing their own sales at other retailers. Still more criticism comes from local citizen groups that have accused Wal-Mart of ruining the market for mom-and-pop stores. ^[11]

While some might see Wal-Mart as invincibly standing at the summit of world retail, it's important to note that other megaretailers have fallen from grace. In the 1920s and 1930s, the A&P grocery chain once controlled 80



percent of U.S. grocery sales, at its peak operating five times the number of stores that Wal-Mart has today. But market conditions changed, and the government stepped in to draft antipredatory pricing laws when it felt A&P's parent was too aggressive.

For all of Wal-Mart's data brilliance, historical data offers little insight on how to adapt to more radical changes in the retail landscape. The firm's data warehouse wasn't able to foretell the rise of Target and other up-market discounters. And yet another major battle is brewing, as Tesco methodically attempts to take its globally honed expertise to U.S. shores. Savvy managers recognize that data use is a vital tool, but not the only tool in management's strategic arsenal.

KEY TAKEAWAYS

- Wal-Mart demonstrates how a physical product retailer can create and leverage a data asset to achieve world-class value chain efficiencies.
- Wal-Mart uses data mining in numerous ways, from demand forecasting to predicting the number of cashiers needed at a store at a particular time.
- To help suppliers become more efficient, and as a result lower prices, Wal-Mart shares data with them.
- Despite its success, Wal-Mart is a mature business that needs to find huge markets or dramatic cost savings in order to boost profits and continue to move its stock price higher. The firm's success also makes it a high impact target for criticism and activism. And the firm's data assets could not predict impactful industry trends such as the rise of Target and other upscale discounters.

QUESTIONS AND EXERCISES

1. List the functions performed by Retail Link. What is its benefit to Wal-Mart?

2. Which supplier metrics does Retail Link gather and report? How is this valuable to Wal-Mart and suppliers?
3. Name the technology does Wal-Mart require partners to use to track and coordinate inventory. Do you know of other uses for this technology?
4. What steps has Wal-Mart taken to protect its data from competitors?
5. List the criticisms leveled at Wal-Mart. Do you think these critiques are valid or not? What can Wal-Mart do to counteract this criticism? Should it take these steps? Why or why not?



11.8 Data Asset in Action: Harrah's Solid Gold CRM for the Service Sector

LEARNING OBJECTIVES

After studying this section you should be able to do the following:

1. Understand how Harrah's has used IT to move from an also-ran chain of casinos to become the largest gaming company based on revenue.
2. Name some of the technology innovations that Harrah's is using to help it gather more data, and help push service quality and marketing program success.

Harrah's Entertainment provides an example of exceptional data asset leverage in the service sector, focusing on how this technology enables world-class service through customer relationship management.

Gary Loveman is a sort of management major trifecta. The CEO of Harrah's Entertainment is a former *operations* professor who has leveraged *information technology* to create what may be the most effective *marketing* organization in the service industry. If you ever needed an incentive to motivate you for cross-disciplinary thinking, Loveman provides it.

Harrah's has leveraged its data-powered prowess to move from an also-ran chain of casinos to become the largest gaming company by revenue. The firm operates some fifty-three casinos, employing more than eighty-five thousand workers on five continents. Brands include Harrah's, Caesars Palace, Bally's, Horseshoe, and Paris Las Vegas. Under Loveman, Harrah's has aggressively swallowed competitors, the firm's \$9.4 billion buyout of Caesars Entertainment being its largest deal to date.



Collecting Data

Data drives the firm. Harrah's collects customer data on just about everything you might do at their properties—gamble, eat, grab a drink, attend a show, stay in a room. The data's then used to track your preferences and to size up whether you're the kind of customer that's worth pursuing. Prove your worth, and the firm will surround you with top-tier service and develop a targeted marketing campaign to keep wooing you back. [1]

The ace in the firm's data collection hole is its Total Rewards loyalty card system. Launched over a decade ago, the system is constantly being enhanced by an IT staff of seven hundred, with an annual budget in excess of one hundred million dollars. [2] Total Rewards is an opt-in loyalty program, but customers consider the incentives to be so good that the card is used by some 80 percent of Harrah's patrons, collecting data on over forty-four million customers. [3]

Customers signing up for the card provide Harrah's with demographic information such as gender, age, and address. Visitors then present the card for various transactions. Slide it into a slot machine, show it to the restaurant hostess, present it to the parking valet, share your account number with a telephone reservation specialist—every contact point is an opportunity to collect data. Between three hundred thousand and one million customers come through Harrah's doors daily, adding to the firm's data stash and keeping that asset fresh. [4]

Who Are the Most Valuable Customers?

All that data is heavily and relentlessly mined. Customer relationship management should include an assessment to determine which customers are worth having a relationship with. And because Harrah's has so much detailed historical data, the firm can make fairly accurate projections

of customer lifetime value (CLV). CLV represents the present value of the likely future income stream generated by an individual purchaser. [5] Once you know this, you can get a sense of how much you should spend to keep that customer coming back. You can size them up next to their peer group and if they fall below expectations you can develop strategies to improve their spending.

The firm tracks over ninety demographic segments, and each responds differently to different marketing approaches. Identifying segments and figuring out how to deal with each involves an iterative model of mining the data to identify patterns, creating a hypothesis (customers in group X will respond to a free steak dinner; group Y will want ten dollars in casino chips), then testing that hypothesis against a control group, turning again to analytics to statistically verify the outcome.

The firm runs hundreds of these small, controlled experiments each year. Loveman says that when marketers suggest new initiatives, “I ask, did we test it first? And if I find out that we just whole-hogged, went after something without testing it, I’ll kill ’em. No matter how clever they think it is, we test it.” [6] The former ops professor is known to often quote quality guru W. Edwards Deming, saying, “In God we trust; all others must bring data.”

When Harrah’s began diving into the data, they uncovered patterns that defied the conventional wisdom in the gaming industry. Big money didn’t come from European princes, Hong Kong shipping heirs, or the *Ocean’s 11* crowd—it came from locals. The less than 30 percent of customers who spent between one hundred and five hundred dollars per visit accounted for over 80 percent of revenues and nearly 100 percent of profits. [7]

The data also showed that the firm’s most important customers weren’t the families that many Vegas competitors were trying to woo with Disneyland-style theme casinos—it was Grandma! Harrah’s focuses on customers forty-

five years and older: twenty-somethings have no money, while thirty-somethings have kids and are too busy. To the premiddle-aged crowd, Loveman says, “God bless you, but we don’t need you.” [8]

Data Driven Service: Get Close (But Not Too Close) to Your Customers

The names for reward levels on the Total Rewards card convey increasing customer value—Gold, Diamond, and Platinum. Spend more money at Harrah’s and you’ll enjoy shorter lines, discounts, free items, and more. And if Harrah’s systems determine you’re a high-value customer, expect white-glove treatment. The firm will lavish you with attention, using technology to try to anticipate your every need. Customers notice the extra treatment that top-tier Total Rewards members receive and actively work to improve their status.

To illustrate this, Loveman points to the obituary of an Asheville, North Carolina, woman who frequented a casino Harrah’s operates on a nearby Cherokee reservation. “Her obituary was published in the Asheville paper and indicated that at the time of her death, she had several grandchildren, she sang in the Baptist choir and she was *a holder of the Harrah’s Diamond Total Rewards card.*” Quipped Loveman, “When your loyalty card is listed in someone’s obituary, I would maintain *you have traction.*” [9]

The degree of customer service pushed through the system is astonishing. Upon check in, a Harrah’s customer who enjoys fine dining may find his or her table is reserved, along with tickets for a show afterward. Others may get suggestions or special offers throughout their stay, pushed via text message to their mobile device. [10] The firm even tracks gamblers to see if they’re suffering unusual losses, and Harrah’s will dispatch service people

to intervene with a feel-good offer: “Having a bad day? Here’s a free buffet coupon.”^[11]

The firm’s CRM effort monitors any customer behavior changes. If a customer who usually spends a few hundred a month hasn’t shown up in a while, the firm’s systems trigger follow-up contact methods such as sending a letter with a promotion offer, or having a rep make a phone call inviting them back. ^[12]

Customers come back to Harrah’s because they feel that those casinos treat them better than the competition. And Harrah’s laser-like focus on service quality and customer satisfaction are embedded into its information systems and operational procedures. Employees are measured on metrics that include speed and friendliness and are compensated based on guest satisfaction ratings. Hourly workers are notoriously difficult to motivate: they tend to be high-turnover, low-wage earners. But at Harrah’s, incentive bonuses depend on an entire location’s ratings. That encourages strong performers to share tips to bring the new guy up to speed. The process effectively changed the corporate culture at Harrah’s from an every-property-for-itself mentality to a collaborative, customer-focused enterprise. ^[13]

While Harrah’s is committed to learning how to make your customer experience better, the firm is also keenly sensitive to respecting consumer data. The firm has never sold or given away any of its bits to third parties. And the firm admits that some of its efforts to track customers have misfired, requiring special attention to find the sometimes subtle line between helpful and “too helpful.” For example, the firm’s CIO has mentioned that customers found it “creepy and Big Brother-ish” when employees tried to greet them by name and talk with them about their past business history at Harrah’s, so the firm backed off. ^[14]

Innovation

Harrah's is constantly tinkering with new innovations that help it gather more data and help push service quality and marketing program success. When the introduction of gaming in Pennsylvania threatened to divert lucrative New York City gamblers from Harrah's Atlantic City properties, the firm launched an interactive billboard in New York's Times Square, allowing passers-by to operate a virtual slot machine using text messages from their cell phones. Players dialing into the video billboard not only control the display, they receive text message offers promoting Harrah's sites in Atlantic City. [15]

At Harrah's, tech experiments abound. RFID-enabled poker chips and under-table RFID readers allow pit bosses to track and rate game play far better than they could before. The firm is experimenting with using RFID-embedded bracelets for poolside purchases and Total Rewards tracking for when customers aren't carrying their wallets. The firm has also incorporated drink ordering into gaming machines—why make customers get up to quench their thirst? A break in gambling is a halt in revenue. The firm was also one of the first to sign on to use Microsoft's Surface technology—a sort of touch-screen and sensor-equipped tabletop. Customers at these tables can play bowling and group pinball games and even pay for drinks using cards that the tables will automatically identify. Tech even helps Harrah's fight card counters and crooks, with facial recognition software scanning casino patrons to spot the bad guys. [16]

Strategy

A walk around Vegas during Harrah's ascendancy would find rivals with bigger, fancier casinos. Says Loveman, "We had to compete with the kind of place that God would build if he had the money....The only thing we had was data." [17]



That data advantage creates intelligence for a high-quality and highly personal customer experience. Data gives the firm a service differentiation edge. The loyalty program also represents a switching cost. And these assets combined to be leveraged across a firm that has gained so much scale that it's now the largest player in its industry, gaining the ability to cross-sell customers on a variety of properties—Vegas vacations, riverboat gambling, locally focused reservation properties, and more.

Harrah's chief marketing officer, David Norton, points out that when Total Rewards started, Harrah's was earning about thirty-six cents on every dollar customers spent gaming—the rest went to competitors. A climb to forty cents would be considered monstrous. By 2005 that number had climbed to forty-five cents, making Harrah's the biggest monster in the industry. ^[18] Some of the firm's technology investments have paid back tenfold in just two years—bringing in hundreds of millions of dollars. ^[19] The firm's technology has been pretty tough for others to match, too. Harrah's holds several patents covering key business methods and technologies used in its systems. After being acquired by Harrah's, employees of Caesars lamented that they had, for years, unsuccessfully attempted to replicate Harrah's systems without violating the firm's intellectual property. ^[20]

Challenges

Harrah's efforts to gather data, extract information, and turn this into real profits is unparalleled, but it's not a cure-all. Broader events can often derail even the best strategy. Gaming is a discretionary spending item, and when the economy tanks, gambling is one of the first things consumers will cut. Harrah's has not been immune to the world financial crisis and experienced a loss in 2008.

Also note that if you look up Harrah's stock symbol you won't find it. The firm was taken private in January 2008, when buyout firms Apollo Management and TPG Capital paid \$30.7 billion for all of the firm's shares. At that time Loveman signed a five-year deal to remain on as CEO, and he's spoken positively about the benefits of being private—primarily that with the distraction of quarterly earnings off the table, he's been able to focus on the long-term viability and health of the business. [21]

But the firm also holds twenty-four billion dollars in debt from expansion projects and the buyout, all at a time when economic conditions have not been favorable to leveraged firms. [22] A brilliantly successful firm that developed best-in-class customer relationship management is now in a position many consider risky due to debt assumed as part of an overly optimistic buyout occurring at precisely the time when the economy went into a terrible funk. Harrah's awesome risk-reducing, profit-pushing analytics failed to offer any insight on the wisdom (or risk) in the debt and private equity deals.

KEY TAKEAWAYS

- Harrah's Entertainment provides an example of exceptional data asset leverage in the service sector, focusing on how this technology enables world-class service through customer relationship management.
- Harrah's uses its Total Rewards loyalty card system to collect customer data on just about everything you might do at their properties—gamble, eat, drink, see a show, stay in a room, and so on.
- Individual customers signing up for the Total Rewards loyalty card provide Harrah's with demographic information such as gender, age, and address, which is combined with transactional data as the card is used.
- Data mining also provides information about ninety-plus customer demographic segments, each of which responds differently to different marketing approaches.



- If Harrah's systems determine you're a high-value customer, you can expect a higher level of perks and service.
- Harrah's CRM effort monitors any customer behavior changes.
- Harrah's uses its information systems and operating procedures to measure employees based on metrics that include speed and friendliness, and compensates them based on guest satisfaction ratings.

QUESTIONS AND EXERCISES

1. What types of customer data does Harrah's gather?
2. How is the data that Harrah's collects used?
3. Describe Harrah's most valuable customers? Approximately what percentage of profits does this broad group deliver to the firm?
4. List the services a Rewards Card cardholder might expect.
5. What happens when a good, regular customer stops showing up?
6. Describe how Harrah's treats customer data.
7. List some of the technology innovations that Harrah's is using to help it gather more data, and help push service quality and marketing program success.
8. How does Harrah's Total Rewards loyalty card system represent a switching cost?
9. What is customer lifetime value? Do you think this is an easier metric to calculate at Harrah's or Wal-Mart? Why?
10. How did intellectual property protection benefit Harrah's?
11. Discuss the challenges Harrah's may have to confront in the near future.
12. Describe the role that testing plays in initiatives? What advantage does testing provide the firm? What's the CEO's attitude to testing? Do you agree with this level of commitment? Why or why not?



