@sn:Section III
@snt: Emerging Technologies
@cn:Chapter 8
@ct:Peer-to-Peer Computing
@ics:I@ic:n the public eye, peer-to-peer (P2P) computing has become almost synonymous with Napster, but it really goes beyond Napster and music sharing to include a broad sharing of resources and information through the direct linking of systems within networks. The underlying technology will change the way the growing number of participants in the wired world share information, content, and resources. The question that remains, however, is whether a viable business model exists underneath all the hype. Even in today's relatively unfriendly technology market, over a hundred companies occupy the peer-to-peer computing space. Although many companies and individuals have entered this space, only a few contenders appear to have a sustainable business model. By examining the existing models, we can develop a framework indicating who will be the ultimate winners.
@h1:History of Network Computing
@cq:The emergence of peer-to-peer computing signifies a revolution in connectivity that will be as profound to the Internet of the future as Mosaic was to the Web of the past.
@cqa:—Patrick Gelsinger, Vice President and Chief Technology Officer, Intel Corporation
@$:Back in the 1960s, when the Internet was conceived, it was a peer-to-peer system. The Advanced Research Projects Agency (ARPA), of the Department of Defense, created ARPANET, which was the precursor of the Internet. The agency's goal was to develop a host-to-host protocol that would improve and increase computer-research productivity through resource sharing over a common network. The University of California-Los Angeles and the Stanford Research Institute were the first two hosts, with independent computing sites and equal status on the network; and ARPANET connected these universities as computing peers. After ARPANET was up and running, the researchers and developers

realized that assisting human communication was the network’s most important contribution.

Early killer apps of the network provided a fundamental symmetry that made it radical in relation to other network structures. FTP and Telnet were client/server applications. While an FTP client sent and received files from a file server, a Telnet client logged into a computer server. The applications were client/server, but the usage patterns were symmetrical. The first e-mail program for a distributed network was released in 1971, and Telnet appeared in 1974. By 1983, ARPA research had developed the TCP/IP protocol (Transaction Control Protocol/Internet Protocol), to which it converted all the interconnected research networks of ARPANET, and the “Internet” became official. It had 500 hosts.

In 1984, the Domain Name System (DNS) was established as a solution to a file-sharing problem. DNS blends peer-to-peer networking with a hierarchical model of information ownership, and it was developed to distribute the data sharing across the peer-to-peer Internet. Also in 1984, Mosaic, the first graphical Web browser, was released and took the Internet by storm. The Internet now comprised over three million hosts.

The explosion of the World Wide Web caused the Internet to become more and more restricted and pushed development toward the client/server model. This structure has plagued the Internet with firewalls and thus with extreme partition. There is a need to return the Internet to its initial design to effectively serve its original purposes of sharing and cooperating. Figure 8.1 diagrams this evolution of networks.

@h1:Current Scene

@$:With the development of networks, information has multiplied and become a critical resource. Within this sea of information, we have some that others may want. There is no reason for us not to provide it to them, but we have no system that allows us to find one another effectively. Napster overcame the challenge for music. By allowing its users to find the content they were looking for anywhere in the world, Napster greatly increased the content available to all.

Most of the technical requirements to handle this type of file-sharing community are already available. Today's file-sharing software has only marginal use outside the music industry, but that will change. Peer-to-peer computing allows people and companies to significantly decrease the complexity and expense of networking. It is not difficult to create Napster- or Gnutella-like programs that are specific for other communities, such as business to business (B2B). XML will be a major enabler for interoperability among networks.

Peer-to-peer technology is already giving rise to new ways of networking and collaborating directly with peers and avoiding the complexities of centralization. It is not hard to imagine the value of a private collaborative network. The idea of sharing spaces with peers and working together without redundancy on a project is fundamental in any industry. Security, however, is a serious concern even in traditional networks that depend on centralized servers, and peer-to-peer networks exacerbate the problem.

“The old days [i.e., the current Internet] were all about centralization and control, almost Soviet style,” says Miko Matsumura, the CEO and co-founder of Kalepa Networks, a six-month-old [[Au: As of what date?]] start-up that plans to link peer-to-peer networks into a sort of alternative Internet. “In this new topology, everyone brings his or her own resources. The new network will be built on top of the old network, like Rome was built in different layers.”[[Au: Source note for quote?]]

One recognized application using aggregated resources is the SETI@home project. Launched May 17, 1999, at the University of California-Berkeley, the project set out to help find extraterrestrial life beyond our solar system. Radio-frequency signals received from outer space are distributed among the participating computers in the network to enable an unprecedented number-crunching application, as those computers use the minute gaps in their processing time to analyze chunks of the signals for a repeated pattern that would indicate an intelligent source. To date, over 2.8 million users in more than 226 countries have contributed nearly 583,000 years of computing to the cause. This computing rate of 25 teraflops more than doubles the speed of IBM's ASCII White, the fastest supercomputer in the world.

The power behind this direct two-way exchange of information and the vast pool of underused resources represent a wealth of opportunities waiting to be exploited. Pioneering individuals and firms have organized, planned, and are already off and running in search of the pot of gold that awaits them. Before declaring any winners, it is imperative to consider what lies beneath the peer-to-peer technology, understand the sources of value, consider the landscape as it exists today, identify the technology enablers, and address the barriers that still need to be overcome.

@h1:Peer-to-Peer Model

@$:Clay Shirky, from the Accelerator Group, indicates that two questions test whether we have a peer-to-peer network:

@nlfi: 1. Does it allow for variable connectivity and temporary network addresses?

@nls: 2. @nlfi:Does it give the nodes at the edges of the network significant autonomy?

@$:If the answer to both questions is yes, the application qualifies as peer-to-peer.

Peer-to-peer computing allows for the sharing of resources and information through the direct linking of systems within a network. Since the Internet is the most extensive computing network available, peer-to-peer computing can take advantage of that existing infrastructure to connect peers. In a typical Internet client/server interaction, the computer at the edge of the network acts as a client while central servers store information. In the peer-to-peer system, an additional software layer enables the computers at the edge of the network to act as clients and/or servers depending on the need. Now the computer is able to respond to a request for information or computing resources from another peer computer. This deviation from the classic client/server system has led to common network topologies: pure P2P and hybrid P2P. Each offers advantages and is coexisting with the current Internet topology.

@h2:Pure P2P Network Topology

@$: Figures 8.2 and 8.3 illustrate the change from a client/server topology to a pure peer-to-peer topology. Pure peer-to-peer models, as illustrated in Figure 8.3, do not use any central servers to direct or control interactions between peers. Individuals or companies in the network bypass central exchanges to relay information directly with one another. Without the central server, peers poll adjacent peers in search of the desired information or resource. Those peers will, in turn, poll other adjacent peers, quickly resulting in an exponentially growing tree. Gnutella-type applications have become one of the most common applications under this network structure.

@h2:Hybrid Peer-to-Peer Network Topology

@$:Unlike the pure peer-to-peer model, hybrid peer-to-peer models, such as Napster, incorporate some traces of the client/server relationship. Central servers within the network fulfill two primary functions. First, they act as central directories where either connected users or indexed content can be mapped to the current IP address. Second, the servers act as a traffic cop and direct traffic among the peers. The hybrid form of peer-to-peer computing is a compromise between the pure peer-to-peer networks of the ARPANET structure and the client-server structure. The alternative topology, shown in Figure 8.4, reflects the additional role of the server in this case.

Whether the model incorporates the pure peer-to-peer or the hybrid form, the key is that, unlike the client-server hierarchical approach, the peers share the resource or information. Clay Shirky has indicated that complete decentralization is not the goal; instead, it is that the application can provide just enough decentralization. Both topologies are the basis for several applications in the market today.

@h1:Today’s Peer-to-Peer Players

@$:Even though the lack of emerging standards for the peer platforms would seem to indicate that peer-to-peer technology is still in its infancy, a host of players have already suited up and taken the field. This section provides an assessment of the major players within the primary peer-to-peer categories of content indexing and file sharing, instant messaging, and distributed computing. Many of today’s players are based on a single functional capability.

@h2:Content Indexing and File Sharing

@$:Napster is the ultimate example of peer-to-peer file sharing. Users open up their hard drives to anyone who is registered on the system in exchange for the same access to others’ files. This concept is extremely powerful in that it opens truly millions of doors to valuable information.

Content indexing and file sharing go hand in hand, for content indexing is the essential first step in file sharing. Since data is scattered throughout a network, an efficient search engine and indexing system must be available to find the critical information and organize it. This is similar to the search method used by Napster or another search engine. Several companies are jumping into the general file-sharing business space, including Aimster and Mangosoft.

Many of the original players in the content-indexing and file-sharing space are related to sharing music since that is the killer app that brought peer-to-peer computing to the forefront. In addition to music, other players in this space allow for sharing programs, videos, images, and documents.

@h3s:Napster.@h3:|em|Napster is now well known as the grandfather of P2P computing in the media, even though it was not even the first P2P music-sharing application. The first was Scour, but Napster has become the most popular. Napster focuses exclusively on sharing music files and recently reached an agreement with Bertelsmann to establish a subscription-based system to compensate the record companies for copyrighted material. Its long-term business model is under question because of the great degree of change that it is facing. Napster uses a hybrid P2P architecture in which several of its own servers maintain the index of music files available on peers’ computers.

@unl:Business model: subscription based

@unl2i:Topology: hybrid P2P

@h3s:Gnutella.@h3:|em|Gnutella is software that was developed as a protocol by Nullsoft engineers. The program was placed on America Online (AOL) servers for a short time and soon removed because AOL-Time Warner saw its possibilities as a conflict of interest with the media industry. Gnutella can be used to share any type of file. As Gnutella’s code was released and is now open source, no single entity owns the rights to it, which has resulted in many Gnutella look-alikes. In addition, Gnutella’s architecture is pure peer-to-peer; rather than relying on central servers to index content, it uses the other clients within a user’s “horizon” to complete an extensive search. This decentralized architecture is challenging not only to shut down, but also to manage.

@unl:Business model: none

@unl2i:Topology: pure P2P

@h3s:Scour.@h3:|em| Scour was the first peer-to-peer music-sharing program. Under the threat of looming lawsuits, Scour closed shop during the last quarter of 2000. CenterSpan Communications bought the rights to Scour and plans to roll out a new version shortly. [[Au: Update?]] It will soon be determined if Scour is going to continue to be a player in the file-sharing space.

@unl:Business model: N/A

@unl2i:Topology: hybrid P2P

@h3s:Aimster.@h3:|em|Aimster is not all that different from Napster. Like Napster, the software is geared toward the consumer market. Rather than allowing sharing of information with everyone on a network, the available network is limited to a smaller set of “buddies.” These people have mutual access to files on a designated location on their hard drive. Aimster latches onto the user’s instant-messaging (IM) service and uses the buddies within that system as the community for file sharing. Aimster allows the

user to conduct natural-language searches of peer systems, permitting the exchange of files and information within a core group of a community.
@unl:Business model: portal partnership
@unl2i:Topology: hybrid P2P
@h3s:iMesh.@h3:|em|iMesh is a general file-sharing program that has attracted over 2 million users. The system is similar to Napster’s in that there is a central server, but it allows for file sharing beyond music files. In addition, iMesh developed a technology that allows it to identify and remove copyrighted files. This may prove to be a valuable functionality in the software as copyright owners push to preserve their intellectual property rights.
@unl:Business model: advertising
@unl2i:Topology: hybrid P2P
@h3s:Mangosoft.@h3:|em|Mangosoft is geared toward the commercial market and provides a means to share and modify files in a team environment. It is a multiuser, Internet-based, file-sharing system whose product, Mangomind, is a secure way for multiple users to access, share, and store files.
@unl:Business model: software licensing
@unl2i:Topology: hybrid P2P
@h2:Instant Messaging
@$:Instant messaging was one of the first peer-to-peer applications. The launch of ICQ in 1996 meant that for the first time individual users could directly address intermittently connected personal computers; any two users on the system could directly communicate with one another. The development enabled future peer-to-peer applications, since many users are still connected using intermittent Internet addresses (i.e., the IP address is different each time the user connects). As the ICQ Web site states: “You can chat; send messages, files, and URLs; play games; or just hang out with your fellow “netters” while still surfing the Net.” Instant-messaging applications for the consumer market have proliferated and use different standards, so to communicate, users must use the same application. The

applications are relatively small and run in the background of a user’s computer. Many of these systems are hybrid peer-to-peer systems in that servers act as critical brokers in facilitating communication.
@h3s:ICQ.@h3:|em|ICQ was the first instant-messaging system that addressed intermittently connected personal computers. The application allows for instant messaging and also includes other modes of communication such as voice, message board, data conferencing, and file transfer. The information regarding users is stored on central servers that maintain the database and check for users on the network. It is best for group chat and file transfers among instant-messaging applications. ICQ is attempting to develop a platform that it can use to build peer-to-peer applications, since it has developed the infrastructure.
@unl:Business model: free software
@unl2i:Topology: hybrid P2P
@h3s:America Online Instant Messenger.@h3:|em|AOL’s Instant Messenger, AIM, has jumped up as one of the leading instant-messaging services because of AOL’s presence on the Internet. It handles basic one-on-one chat effectively and adds additional capabilities. AIM targets the novice computer user, but because of the large number of America Online members, it controls a major portion of the instant-messaging market.
@unl:Business model: free software, subscriptions from America Online users
@unl2i:Topology: hybrid P2P
@h3s:Windows Messenger.@h3:|em|Microsoft’s new Windows Messenger, included with the new Windows XP operating system, is the company’s effort to bundle instant messaging with what it hopes will remain the world’s ubiquitous operating system. It allows users to page cell phones, send pictures and music, and complete free voice and video calls to another PC on the Internet. Windows Messenger also makes it possible to call from a computer to a telephone just as easily, using the PC-to-phone service provider that the user selects. It is a direct challenge to AOL’s Instant

Messenger, currently the market leader. Increasingly, AOL and Microsoft are going after the same customer.
@unl:Business model: free software, free subscription required
@unl2i:Topology: hybrid P2P
@h2:Distributed Computing
@$:The overwhelming response to SETI@home's distributed computing initiative has led both nonprofit organizations and profiteers alike to explore alternative applications. Other firms and institutions with processing-intensive applications have already made use of the supercomputer power embedded in the network. Cancer and AIDS research, earthquake simulation, and particle physics research are just a few of the fields that have benefited from this raw processing horsepower.
@h3s:Popular Power.@h3:|em|Popular Power is one such broker that has entered the market by offering users $5 per month to capitalize on their underused CPU cycles. This aggregated processing power is then auctioned off to bioengineers, mathematicians, and others in need of a computation community.
@unl:Business model: transaction fees
@unl2i:Topology: hybrid P2P
@h3s:Entropia.@h3:|em|Entropia has also entered this space, focusing primarily on nonprofit applications such as FightAIDS@home. By marketing with the slogan "Activate your PC's power to save the environment," Entropia has created the image of performing tasks for the good of humankind.
@unl:Business model: transaction fees
@unl2i:Topology: hybrid P2P
@h3s:Envive.@h3:|em|Envive is one of the leading providers of Web-application performance-management solutions. Peer-to-peer Web testing provides real-world Web site-performance testing. Instead of using servers that are located within ISPs to simulate users, peer-to-peer Web testing uses actual PCs that are connected to the Internet, which is a more accurate method. It incorporates all the inherent delays and upsets that are normally present on the network into the test.
@unl:Business model: transaction fees

@unl2i:Topology: hybrid P2P

@$: As of February 2001, over 110 companies were relying on the peer-to-peer structure as the heart of their business. About 35 of these companies are competing in the file-sharing arena; 15 companies are in cycle sharing; and the remainder are spread out as combinations of these categories or are in smaller niche areas.

Up to this point, we have considered the functionality behind peer-to-peer computing and some of the participants in the field, but we have not looked into the value proposition that it brings.

@h1:Value Drivers: Computing at the Edge of the Network

@$:The direct link to peers enabled by bypassing larger networks produces value from several sources. Shifting from the client/server architecture to the peer connections makes better use of the available resources. The change in network architecture takes computing power, which in the current Internet architecture resides on servers, and moves it to the edge of the network. This is not a new concept, as networks were originally built on a peer-to-peer architecture. With the growth of the World Wide Web, the network architecture evolved into a hub-and-spoke model that placed servers at the core of the network and individual PCs at the edges. The model developed into a system in which knowledge and information transfer was mostly a one-way street. When someone surfing the World Wide Web requests information, a server retrieves it. However, the PC user is not sending back much information that increases the knowledge base of the network. The PCs connected to the Internet are restricted simply to being clients that receive information. This model works well when users need content and information stored in a central location but is not as efficient when the information is needed at the edges of the network.

The value drivers that individuals and corporations are capitalizing on with peer-to-peer systems are described in the following section.

@h2:Distributed Parallel Processing

@$:Move over supercomputer, here comes grid computing. Creating a network of directly connected peers now enables processing resources to be aggregated and used at a large scale. While processor speeds on even home PCs are surging, only a fraction of the processing power is actually used. Peer computing now allows complex and time-consuming processing applications to be distributed and processed at the end nodes. Math-intensive operations that were once impossible without a supercomputer are now feasible at a more economical level.

@h2:Distributed Content Storage

@$:Like distributed parallel processing where resources are underused, distributed content storage takes advantage of underused storage resources among all participants in a peer network. Since even within a large corporation, the use of hard-drive storage is only about 50 percent, terabytes of storage capacity go underused. Although dropping prices per unit of storage negates this value proposition, the real value lies in the aggregation of the available capacity on a larger scale within a corporation.

In addition, edge services act as caching mechanisms as they move data closer to the point of consumption. A global company that gives online training to its employees through the Web can store the video on local clients, which act essentially as local database servers. Because the streaming goes over the local area network (LAN) instead of the wide area network (WAN), the company saves time; it also saves money by eliminating the need for local storage on servers.

@h2:Direct Communication/ Information Distribution

@$:The majority of Internet applications are structured such that information is pulled from the business to the consumer (B2C). Though this arrangement obviously produces valuable applications, it doesn’t enable the end user to share anything meaningful with others on the network. Content sharing is at the heart of what has brought recognition to peer computing. Previously, making content available to the public was a tedious process that most users were unwilling to undertake. They needed to obtain a fixed

IP address from an ISP and wade through the technicalities of setting up a Web site. Now users can download a free interface, install it, and within minutes share their content with peers on the network.

Peer-to-peer networks need to operate outside the DNS system with nearly complete autonomy. Therefore, collaboration through peer-to-peer technology takes advantage of resources such as storage, human presence, and content at the edge of the Internet. These new platforms will fully embrace existing systems. Current collaboration platforms through client/server arrangements are not flexible in that a collaborative team needs to adapt to it and its applications. Through peer-to-peer computing, a collaboration environment will leverage the strength of existing assets of applications and computing. Companies will be able to adapt their databases stored through their Enterprise Resource Planning (ERP) systems and share information with partners in a seamless and secure environment without relying on a third party.

@h2:Dynamic, Distributed Search

@$:Unlike Napster, in which a directory of each node’s content is stored on a central server, Gnutella applications have been based on searching the content on the peer’s storage device. Typically, a handful of adjacent peers are polled in search of the desired content. Each of these peers in turn polls a handful of other adjacent peers. The exponential growth of this search quickly returns the results from potentially thousands of peers.

To be a viable source of value, a software interface is essential to match the desired content with the available content residing among peers on the network. The Napster model and other search engines using the server-based directory are very effective in mapping the content available on the network and are limited only in the time between mapping updates. To be advantageous, the pure peer-to-peer application would benefit most in cases in which the real-time updates are needed when searching the network.

@h1:Secret of Success

@$:Though peer-to-peer computing is not completely revolutionary, technological advances have enabled it to progress to where it is today and will further propel it into the future. Increasing bandwidth, hybrid centralized/decentralized models, scalability, and the sheer growth in the number of participants have enhanced the peer-to-peer offering.

@h2:Bandwidth Availability

@$:Widespread use of broadband enables peer-to-peer computing because it facilitates communication between the nodes at the edge of the network. “Last mile” bottlenecks limit the speed of data transfer on the network. This becomes even more critical in peer-to-peer networking, since a majority of the communication occurs over the last mile. The rollout of broadband makes peer-to-peer computing much more powerful. Although the increase of broadband has not lived up to expectations, its steady growth improves the prospects for the success of peer-to-peer computing.

Peer-to-peer applications are changing the assumption that end users want only to download from the Internet and never upload. The Web was the killer application of the Internet, and it is made up mostly of clients, not servers. New peer-to-peer players face great challenges sparked by the rise of asymmetrical network connections such as cable modems and Asymmetric Digital Subscriber Line (ADSL). Until this asymmetry is resolved, digital subscriber line (DSL) providers have an advantage in the race to enable peer-to-peer networking because of their relatively symmetrical technologies.

@h2:Hybrid Centralized/Decentralized Models

@$:Napster struck oil by integrating the benefits of the direct peer-to-peer connection with the centralized directory search to efficiently pair up the seeker and provider of content. Peer-to-peer success stories will properly balance this blend of centralization and autonomy to facilitate people’s tasks in a more efficient manner or allow them to complete new functions that are not even out there today. As noted, the goal is to be just decentralized enough.

@h2:Network Externalities

@$:As the number of participants grows, the available shared resources (content, storage, or processing) increase as well. The increase in available resources then attracts new participants, and the cycle continues. Napster managed to build a database of over 50 million addresses. Virtually any song from any genre was available because of the enormous number of participants.

Napster changed the economics of the music industry because the marginal cost of downloading a song is virtually zero. By creating one more copy of a song in the network, a user increased the chances that the next user who searched for that song would find it. Therefore, accidentally, value was created in the network. Designing systems that create value automatically will be key for succeeding in the peer-to-peer space.

@h2:Scalability

@$:By using the resources of the end user, peer networks benefit from the scalability inherent in the network. In the corporate environment, instead of fighting the bureaucratic channels to obtain additional storage and processing capacity for the servers as the users grow in numbers, each new user brings the incremental resources to grow the network.

@h1:Challenges

@$:Even as technological advances are ushering in peer computing, critics maintain that issues of security and data integrity pose threats to its usefulness and growth. Also, bandwidth not only contributes to peer-to-peer penetration, but potentially poses a barrier by increasing traffic flow across the network. Finally, the rush to establish industrywide de facto standards is resulting in applications flooding the market and a lack of interoperability.

@h2:Security

@$:Security and authentication are two crucial factors that will determine the success or failure of peer-to-peer networks. It is predictable that peer-to-peer networks will become the killer application for virtual private networks. It is hard to imagine those networks without state-of-the-art security systems. Some players in the peer-to-peer space have already created maverick forms of security. These include one pass phrase per account, one asymmetrical key per account, and one asymmetrical key pair per identity for signature/verification, and another asymmetrical key for encrypting/decrypting symmetrical keys. There is little doubt that the latest technological advances that allow for increasing security will be key for these applications.

@h2:Bandwidth Limitations

@$:Transferring client-server traffic to peer-to-peer connections can be likened to moving traffic from highways to residential streets. The move may alleviate the traffic on the highways and provide potentially more efficient routes, but the residential network is not designed to handle the traffic. As the growth and penetration of broadband continues, it will alleviate the congestion in the last mile of the network and facilitate efficient peer-to-peer computing.

@h2:Interoperability

@$:Because of the rapid proliferation of peer-to-peer computing technologies, several such technologies are available. Although that situation provides many choices for the consumer, the lack of standardization also prevents computers with either different computing systems or software platforms from communicating with one another.

@h1:Current Applications and Value Drivers

@$:The main applications for peer-to-peer computing include file sharing, instant messaging, and distributed computing. These individual functions are evaluated along with the crucial value drivers. Table 8.1 contains a summary.

@h2:File Sharing

@$:The current killer app for peer-to-peer computing has certain key value drivers. The success of file sharing depends on both distributed content storage and dynamic search capabilities. Initially, the information must be indexed in a way that makes it accessible and understandable. Then a powerful search capability is required to find the data. This can be implemented through a system in which the information either resides on central servers or is created dynamically by each user who logs on. The ability

to distribute content storage prevents overburdening any individual node and optimizes storage usage.
@h2:Instant Messaging
@$:Instant messaging is another popular function in peer-to-peer computing. The critical value driver in instant messaging is direct communication/information distribution. To find another user on the network, the instant-messaging system must allow users to directly access intermittently connected computers that have changing IP addresses. This capability allows the program to connect any two or more users and allow them to communicate directly.
@h2:Distributed Computing
@$: The key value driver of distributed computing is parallel processing. The ability to take a large task, break it into manageable pieces, and then send it out to individuals is at the heart of distributed computing. The value of direct communication further enhances this proposition by enabling direct connectivity with the nodes.
@h1:Viable Business Models
@$:The question that still remains, however, is how value will be captured and who will capture it. Even Forrester Research has stated that although millions of consumers and dozens of companies are jumping on the bandwagon, profitable business models remain elusive. Is the end result a networked world in which consumers are simply sharing and receiving content for free?

For the application to represent a viable business model, it must accomplish one or more of the following three propositions for the end user:

@nls: 1. @nl:<I>Revenues must be generated above and beyond the costs to produce them.<I>|en|These revenues can fall into several models--

@nlsl:—Transaction—fees collected on the basis of an occurrence.
—Subscription—a fixed fee for the content or service over a predetermined period.

—Information selling—includes the selling of consumer-behavior information.
—Advertising—banner advertisements or other kinds of advertising placement.
—Software licensing—license fees for applications.
—Solutions providing—bundling of service and support with applications to suit the particular needs of an organization.
@nl2i: <I>Increasing revenues cannot be confused with market share.<I>|en|It is imperative that these revenues not be obtained “at any cost,” but that the costs be low enough to net a profit.
@nls: 2. @nl2i:<I>A cost savings must be realized that outweighs the implementation costs.<I>|en|Companies will look to peer-to-peer computing not only to enhance their revenues, but to reduce costs. Large corporations may be attracted to enterprise solutions that will actually enable them to operate at a lower cost. Once again, the gain must exceed the cost of implementing the solution.
@nls: 3. @nlfi:Better information must be received to enable better decision making. Though better decisions will ultimately lead to increased revenues or reduced costs, it is often difficult to justify a project when the financial gains are not easily quantifiable up front. For that reason, this third option may not target immediate cost savings or revenue gains, but rather may improve the overall decision-making capabilities within the corporation. Expedient, reliable information previously unavailable may now be delivered right to the desktop.
@h1:Future Applications
@$:The current network requires and allows minimal interaction between clients. Figure 8.5 illustrates the lack of communication between peers. The focus, as discussed, is on the interaction between the client and the server. Information has to be manually transferred between the clients with no direct means of communication.

Peer-to-peer applications will allow clients to interact in ways that do not erect barriers between them. Clients will be able to peer through the firewall and share information and resources. These applications will bring together in a unique way the functional capabilities already offered in the market. The appropriate set of functionalities from Figure 8.6 will come together to create a potential killer application.

@h2:Collaboration in Project Management
@h3s:Application.@h3:|em|ERP providers such as SAP, JD Edwards, and Peoplesoft are adding online collaboration networks through which different parties can participate in a given project with different levels of access. Other firms, such as Bidcom and Cephren, provide the same service mainly for the construction industry. They have been providing these services mostly through hosting, not through peer-to-peer applications. These application service providers rent the service and keep the project information on their servers, usually in huge data centers.
@$: Collaboration will evolve in an unimaginable way through peer-to-peer computing because of the symmetry of the participants in the network, as well as the economics of not using an administrator. Imagine a group of engineers working on the same drawing. Each engineer does not need to have his or her own drawing, so the drawing can be stored in pieces on computers on the network. Solution providers allow for files to be in bits and pieces, so instead of a 3 gigabyte file on somebody’s hard drive, there can be a (3 gigabyte)/n file stored on each computer on the network, with n being the number of computers on the network.

Alternatively, imagine an engineer checking an architect’s drawings. The file needs to be in only one space in the network. It does not matter where on the network the file is stored. The engineer can have access to it from his computer, and his changes or comments will be automatically saved on the other team member’s computer. Since the engineer does not have the file, all he is sending is information about the changes that need to be made. His computer is not processing that information; the computer that has the file is doing it.

Peer-to-peer computing fits perfectly with the collaborative feature of the Internet, as shown in the latest project from Ray Ozzie, the inventor of Lotus Notes. Ozzie is working on a product that will allow small teams to share information and collaborate on tasks and projects. His company, Groove Networks, is being watched with great interest.

For collaboration to succeed as a peer-to-peer application, several users on a team must have access to the same documents. Beyond that, the ability to modify the documents and update the shared versions at all locations is necessary. Finally, since project groups must work closely together to implement a project, direct communication between users is a key capability. On the basis of these requirements, the critical value drivers for collaboration are distributed content storage and direct communication. Distributed content storage will allow the essential information to be stored at the location where it is used the most. A requirement for partial content storage is that all users must be online at all times to assure access to the necessary information.

@h3s:Business Model.@h3:|em|We can expect vendors to charge subscription and monthly fees for their platform service. As in the traditional software industries, there are significant development costs. Whoever comes out with a flexible, easy-to-use collaboration platform will enjoy a first-mover advantage in a new market willing to pay dearly for significant cost reductions in day-to-day team-project management.

@h2:Knowledge Management

@h3s:Application.@h3:|em|As the economy develops toward a more information-driven system, the ability to manage knowledge within the enterprise is going to become even more critical. The power of using peer-to-peer computing for knowledge management is enormous, since it makes vast amounts of knowledge available to everyone. For reasons of security and intellectual-property protection, however, users are highly unlikely to open up and share knowledge with any stranger on the network.

@$: Although it does not appear likely that we will be able to peer into the computers of total strangers on the Internet, the possibility of implementation on a smaller scale is nonetheless attractive. Even within a corporation, imagine the time that is spent reworking problems that were solved previously. In many cases, issues come up repeatedly; but because of personnel turnover and lack of appropriate social networks, those same problems are worked on again from scratch.

Another part of knowledge management that peer-to-peer computing can facilitate is automatic backups. Since peer-to-peer computing can use unused resources on computers such as storage, a knowledge-management system can include the capability of taking information that resides on a user’s hard drive and backing it up to another user’s hard drive. That eliminates the need for expensive central storage while maintaining enterprise-critical data.

@h3s:Business Model.@h3:|em|Two approaches potentially offer viable business models in the knowledge-management space. These are software licensing and solutions providing. The first, licensing, is straightforward in that a vendor would provide a software program allowing a company to index, store, and retrieve information. This is the first area of focus for companies interested in the market. As additional competitors enter the market, however, this space will become commoditized. Those vendors who understand the functions of the company will be able to add more value as solutions providers and generate higher revenues. In addition, a vendor might be able to use the transaction model on a megabyte-per-megabyte-of-data basis to generate revenue.

@h2:Enterprise Processing

@h3s:Application.@h3:|em|Distributed computing has tremendous potential because the price of supercomputers is high and a great number of unused computing assets are available. Many of the entrants in the market are developing applications that target general computer users on the Internet. The business model generally involves enticing users by paying them some nominal amount or asking them to donate computing power to nonprofits. The emphasis of these companies on the general public presents some challenges. These include the reliability of computing resources, payment of individual users, and the security issues that companies may have with sending their information outside the firewall.

@$: Using the computing resources within the enterprise offers tremendous opportunity. Corporate IT managers will likely be loathe to give up control and allow computing resources outside the company to become critical, but those same resources are available within the enterprise. Great computing power remains unused inside corporations. Peer-to-peer distributed-computing companies that focus within the enterprise have incredible resources.

In addition to providing the processing power to handle major computing, these software providers must provide solutions. Combining desktop computers to form a supercomputer will provide only part of the solution needed to gain traction in corporate environments. The providers must combine a killer app with the computing power. That means gaining a deep understanding of the end users' market and needs. Rather than just being computing experts, they will need to develop a vertical focus and be industry experts as well. The opportunities for enterprise processing lie in financial services, biotechnology, and general research and development within the semiconductor, aerospace, and defense fields.

@h3s:Business Model.@h3:|em|The aggregators of computing power who seek to serve the enterprise market must focus on securing a transaction-based revenue stream or becoming a solutions provider. The transaction-based system is the first step and is attractive for customers whose tasks are not critical. These customers would be willing to let their processing be done outside their firewalls and pay a per-usage charge. However, the majority of major corporations would like to maintain control over their processing. In these cases, peer-to-peer computing vendors will need to propose an enticing value proposition that will allow them to get their foot in the door. That entails becoming a solutions provider and understanding the inner workings of a customer's business and the intricacies of peer-to-peer computing.

@h2:Games

@h3s:Application.@h3:|em|Games are just not that much fun when played alone. Developers initially resolved this problem by building intelligence into the game, thereby setting up the "man versus machine" situation. Although that provided some form of competition, the human element and varied opponents were still missing. Peer-to-peer computing brings an abundance of opportunities into the games world. To be effective, two criteria must be established.

@$: To engage in multiperson game playing through the Internet, it is first necessary to find potential partners. For some time, online game players have used host servers to find opponents for Monopoly, chess, and other games. Although using a directory service through a server is an effective way of pairing these aficionados, the Gnutella technology could also be used. Instead of the players' registering in a directory, the game software could have an interface that would poll for peers out on the Internet to find others with the same interface.

Additionally, larger software developers could create a suite of game possibilities so that, depending on mood, a user could play any one of several possible games online. Consider a sports-games developer such as EASports. A Gnutella-type interface could be included on all software applications such that an enthusiast could connect and find other players thirsting for some competition in basketball, football, baseball, golf, and other sports. Whether the competition is matched up through a server directory or through the Gnutella search, the key is to link those individuals seeking to engage in competition in the same game.

The second essential ingredient for success is where peer-to-peer computing really delivers something new. Even though the client-server relationship enables locating other game players, it still requires the information to flow through the server. That may be fine for games in which response time is not critical, but for arcade-style applications, in which the actions of one player must be immediately fed back to the other, the client-server doesn’t suffice. Peer-to-peer computing provides the direct connection to engage in the action-packed games that thrill seekers are looking for.

@h3s:Business Model.@h3:|em|The transaction-revenue model in which game players pay a fixed fee per occurrence is a possible revenue source but probably not the best. The subscription plan would increase overall revenues more because the smorgasbord mentality of game players would lead these enthusiasts to devour more than they would have otherwise. The increase in revenue would outweigh the increase in cost because of the greater demand.

@$: Another area of opportunity is on the software-development side. Competition for the killer-app game has been stiff for a while, but now the focus will change to game applications that make the most of the peer-to-peer connections. Games that involve several individuals competing simultaneously with near real-time interactions will ultimately succeed. [[Au: Provide cite for Table 8.2.]]

@h1:Future Landscape

@$:In the past several years, there has been talk about the death of the personal computer. The popular prediction was that the intelligence would lie at the heart of the network, within servers, and the users would connect to the network using simple terminals. Peer-to-peer computing will be a strong opposing force to this trend. Instead of using a “dumb” machine to connect into a powerful network supported by servers, the terminals of the future will be smart and provide the computing horsepower.

Servers within the system will not be eliminated either, for two reasons. First, companies will not walk away from their investments in servers. Second, servers will play an integral role in the peer-to-peer topology. PCs and servers will work together to take advantage of their appropriate strengths.

Rather than users’ receiving information from a network, there will be true two-way communication, and the computers at the edge of the network will play a critical role. This will require more computing resources in the networked computers. The resources will be used for many applications, including collaboration, knowledge management, games, and enterprise processing. As peer-to-peer computing continues to develop, we will see further integration of the different functions into more sophisticated applications. This includes the integration of instant messaging and file sharing into the collaboration application, which is already occurring. This trend will continue as more of the functions come together to create killer apps.

The ability to build a network around individual PCs dramatically changes the network architecture. Instead of relying on central server farms that are vulnerable to going down, the network is more stable by being spread throughout the system. In addition, the incremental investment to expand a network is dramatically reduced. Rather than buying a new server to expand the network, firms can simply add another PC to the system.

Building relationships will drive the future development of peer-to-peer computing; the technology will be an enabler. To build a peer-to-peer network, traditional walls of separation must be torn down. This applies in both the corporate setting and the home-user environment.

@h3s:Corporate Setting.@h3:|em|Even within corporate divisions, the traditional mode of operation is to divide the business into departments and maintain autonomy and information within those silos. The influence of peer-to-peer computing will help tear down these divisions. The information that is available in one part of the corporation will be readily accessible to all. That presents some management challenges in implementing peer-to-peer systems. Managers unfamiliar with the technology behind peer-to-peer systems may be reluctant to embrace what may appear to be radical changes. Although adoption may initially be difficult, tremendous returns await companies taking advantage of the benefits of peer-to-peer computing.

@h2:Home-User Setting

@$:Only a small percentage of home users have taken on the challenge of building their own Web sites to share their treasured images with family and friends. The increasing number of digital cameras, coupled with higher-resolution images, has created a strong demand for ways of sharing these memories. Forrester Research has projected that peer-to-peer computing offers a much easier solution to this once formidable and intimidating task. Users can now place their digital images and other content in a directory and make it available to those to whom they want to give access.

By 2005, 35 percent of all online households will be using peer-to-peer technologies for personal sharing (Figure 8.7). Though a potential market of nearly 24 million users is certainly worth writing home about, shared utilities do not present credible business models. Essentially, the functionality could be built directly into the operating system and simply provide one more compelling reason to purchase the OS upgrade. The real winners in the home-user domain are the consumers themselves. Although the development in the games frontier may create coveted revenues, most other applications will enhance the consumer’s computing experience without making billionaires out of the solutions providers.

@h1:Conclusion

@$:Though it is not certain who will discover the pot of gold, peer-to-peer computing will change the way we communicate and share information today. Whether it becomes a revolution, as Patrick Gelsinger, of Intel, predicted, or simply merges the centralized model of today with the decentralized model of yesterday, we all stand to gain from the functionality it brings.

@h1:References

@r:Boyd, Jade. “Load Testing Uses Distributed PCs.” <I>InternetWeek<I> (November 13, 2000).

Karel, Bruce, et al. “P2P’s Pervasive Future.” <I>Forrester Report<I> (January 2001).
Kutler, Jeffrey. “Rocking the Buy Side.” <I>Institutional Investor International<I> (November 2000).
Mahowald, Robert. “Peer-to-Peer Computing: A Primer.” <I>IDC Report<I> (January 2001).
Mahowald, Robert, and Mark Levitt. “Deja-Vu All Over Again: Collaborative Applications Meet Peer-to-Peer Computing.” <I>IDC Report<I> (December 2000).
McAfee, Andrew. “The Napsterization of B2B.” <I>Harvard Business Review<I> (November-December 2000).
McDougall, Paul. “Envive Taps Peer-to-Peer to Test Web Sites.” <I>Information Week<I> (December 4, (2000).
Minkoff, Jerry. “Napster-style Programs May Revolutionize Financial Services.” <I>Web Finance<I> (January 2001).
Oram, Andy. <I>Peer-to-Peer: Harnessing the Power of Disruptive Technology.<I> Sebastopol, CA.: O’Reilly & Associates, 2001.
Scannell, Ed. “P2P Takes on B2B.” <I>InfoWorld<I> (November 27, 2000).
Shirky, Clay. “Lessons from Napster.” Presentation at O’Reilly Peer-to-Peer, San Francisco. (2001).
Walsh, Mark. “Napster Wannabes See Big Bucks on Wall Street.” <I>Crain’s New York Business<I> (October 2000).