Prepared by: Richard Taylor
under the supervision of Dr. Blake Ives

Peer-to-Peer Computing: Applications in the Business World

Introduction

"P2P is a revolution that will change computing as we know it."
Andy Grove, chairman of Intel (Rutherford 2000)

What's so special about P2P? It's simply another way of doing things we have always used the Internet and other networks for – to exchange and find data, or to collaborate remotely – making use of distributed, interconnected computers to do those things faster, more flexibly and more dynamically. The power of peer-to-peer computing, better known as P2P, was demonstrated in 1999 when a 19-year-old university student, Shawn Fanning created Napster. Napster permitted users across the world to share MP3 music files. An instant success, Napster, in just 18 months, attracted nearly 40 million users (Laudon & Laudon 2002).

(image from www.napster.com)

Napster opened up a wonderful universe of music to millions of people. But Napster was much more than a free source of popular music; it was an exploration of the unknown that widened cultural horizons (Oram 2001). At its peak, Napster boasted 70 million registered users with up to 1.57 simultaneously online (Shirky 2001). Napster was a challenge to intellectual property laws and eventually was shut down by court order; but P2P refused to die. As of June 2002, an estimated 19% of Americans over age 12 had downloaded music files from various P2P systems according to cyberatlas.internet.com. P2P file sharing applications accounted for five of the top 10 downloads for the download.com Web site in the last week of June 2002, together constituting 4.5 million downloads (Lee 2003).

Downloading music is not the only use of the P2P architecture. Forrester predicts that by 2005, 43% of broadband households will be using P2P to share things such as digital photos, digital video, calendars, and addresses (Forrester 2001). Online consumers are already armed with tools for creating personal rich media, and those with broadband access are actively exchanging personal rich media at least once a week (Forrester 2001). But some are already tiring of continually posting materials to Web sites or having their e-mail folders clogged with attachments. It would be far easier for them to set aside secure areas on their PCs' hard drives and permit monitored access to family and friends.

The initial success of Napster and other file sharing systems forcefully demonstrated that the P2P computing model can readily accommodate explosive growth and support millions of client computers. IT managers are beginning to explore ways of leveraging P2P's strengths within their organizations, while avoiding or overcoming its management challenges. They are joined by a number of major vendors, including Intel, IBM, Sun Microsystems, and Microsoft who have embraced P2P computing as demonstrated by new products and well-funded initiatives (Edwards 2002). This paper will examine the business applications for the P2P architecture. It begins by examining the P2P architecture and its roots. The paper then discusses business uses for P2P computing followed by an investigation of the concerns business have regarding P2P. Finally, the paper will look into the future of P2P.

What is P2P?

P2P has been around for over 25 years. Remember original computer modems? Their connections were P2P (Fox 2001). Another early, but less sophisticated, version of a P2P network was "sneaker net". Users had personal computers on their desks, but not connected. To transfer information, data was copied to a disk and carried to another computer. The most frequent endpoint of a typical sneaker net was a PC with a printer attached (Cope 2002). The telephone system, the discussion forums of Usernet, and the early form of the Internet can also be classified as P2P systems (Minar & Hedlund 2001). The earliest Internet programs communicated from computer to computer, with no server required in between to facilitate the connection. One of the breakthroughs of ARPANet was that is was based on the concept of connecting computers as equal peers. (Spangler 2001). The initial use of P2P networks in business followed the deployment in the early 1980s of free-standing PCs. One of the first large-scale users of P2P was in 1994 when two scientists at the Goddard Space Flight Center in Maryland networked 16 processors together and created a single cluster computer (Rutherford 2000).

But what exactly is P2P…a set of protocols, an IT architecture, a design philosophy stressing decentralization, a business model, or merely a fad? In its current form it encompasses all of these attributes. P2P refers to technology that enables two or more peers to collaborate spontaneously in a network of equals (peers) by using appropriate information and communication systems without the necessity for central coordination (Schoder & Fischbach 2003). P2P computing is a form of distributed processing that links computer via the Internet or private networks so that they can share processing tasks (Laudon & Laudon 2002). This means that users share files via a number of interconnected virtual private or public networks. They hardly ever crash because they are decentralized and they can handle huge numbers of users simultaneously. The architecture of the networks also makes them more scaleable and far less vulnerable to distributed denial of service attacks (Krebs 2001).

The main concept behind P2P is decentralization. Without a centralized server to connect users, users are connected through other users. This allows users to communicate directly with one another. The PCs in this interconnected network are referred to as "peers". They offer their resources to other PCs on the network (Fox

2001). Resources can include both disk space and processing power. The term "peer" indicates that each node is treated as an equal. In reality, some peers are "more equal" than others. Computers have different CPUs, memory, storage capacity, and network connectivity. Some computers are professionally managed and highly available while others are not. Some computers reside at network hubs while others are at the edges. Some are locked in machine rooms and some are public (Kubiatowicz 2003). However, when accessing information they are all clients, when providing information to other peers, they are servers. P2P generally assumes that each peer is acquainted with a small number of other peers with which it can exchange information and services. Acquaintances change constantly, there is no central control, and peers remain autonomous throughout their participation in a P2P network (Penserini et al. 2003).

In a pure P2P environment, every client needs to know where all the others are, and must be able to qualify and connect to them. A variation of this computing model, known as hybrid P2P, introduces a server, which may be used for managing the peer devices or to store information, such as replicated data, for disconnected peers. In this situation, the server plays a supporting role, unlike the leading one it has in client/server or Web server applications. The hybrid model resembles a ring network topology, where pieces of the traditional application and directory, and the message store and files, reside on the hard drives or the networked PCs. However, unlike client/server applications, hybrid P2P applications would only draw upon server resources when some part of the ring is not available, such as when a remote user is no longer logged onto the network, when the network is experiencing heavy traffic, or when a connection is severed (Edwards 2001). In the hybrid architecture, the central server may store content replications for disconnected peers, and other information to manage peer devices and user profiles. With larger systems, multiple servers may store the directory information required to locate other systems and identify services of facilities available on them. Servers may also be used to connect to the system for content exchange, replication, or management purposes (Edwards 2001).

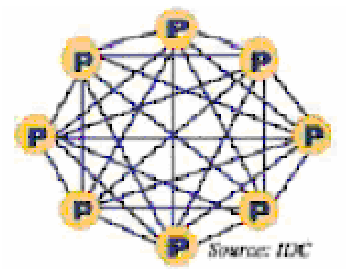

Exhibit 1: Pure P2P
(Image from Edwards 2001)

Hybrid P2P networks are appealing to businesses, because they still afford a level of control that is lost with pure P2P. Napster is an example of a hybrid P2P network, an architecture which proved to be its downfall. The Recording Industry Association of America was able to totally shut down Napster by a court order shutting down the primary server that was used to direct request for all of Napster's millions of users. Other file sharing sites, such as KaZaA, have been able to stay in operation because they are pure P2P networks. There is no central server that could be shut down

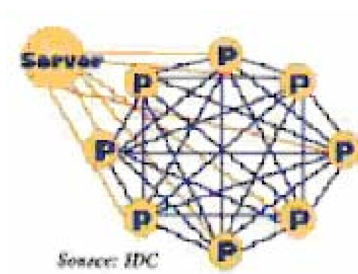

Exhibit 2: Hybrid P2P
(image from Edwards 2001)

to stop the transfer of files. Each peer is connected directly to a group of peers. Legal action must now be, and has recently been, directed at individual users.

Whether pure or hybrid, P2P is divided into three basic divisions: distributed computing, instant messaging, and affinity communities (Rutherford 2000). Distributed computing environments pool the processing power of many computers. One popular application involves utilizing spare CPU time on client machines across the network to complete tasks more quickly and make more efficient use of computing resources. This so-called "grid computing" model has been used primarily for heavy number crunching in scientific research, but is now being considered for corporate applications such as data mining and 3-D modeling (Edwards 2002). Grid computing promises to enable supercomputing capabilities at a fraction of current costs by breaking down large computations into much smaller ones that can be run on numerous laptops and PCs with spare CPU time. Also, whereas supercomputers require expensive, specialized software, grid computing works with a wide variety of windows and UNIX programs available off the shelf (Edwards 2002). Vendors such as Distributed.net, Entropia, SETI@home and United Devices recruit computer users from the general public to volunteer their idle processing power. SETI@home has created the world's largest supercomputer by using the spare resources of millions of individuals' PCs to process data collected from space in a search for extraterrestrial intelligence (Jones 2001). Grid computing is becoming popular with financial services, biotech, and science firms that need intense processing power (Rutherford 2000). P2P proponents claim that the untapped resources of personal computers owned by ordinary people can be combined to build something greater and more reliable than the sum or its parts (Kubiatowicz 2003).

Instant Messaging allows interconnected users to exchange text messages and files synchronously. AOL and Microsoft both offer free messaging tools. Lands' End has capitalized on this tool by allowing on-line shoppers to utilize "Lands End Live" and "Shop with a Friend" (Ives & Piccoli 2003). "Lands' End Live" allows customers to chat online directly with a customer service representative. "Shop with a Friend" allows two shoppers to browse the site together, communicate with each other and add items to a single shopping cart. Title Data, Inc. recently released an IM product called Sonork, which will enable customers to link directly to needed files, thereby reducing support-related calls by as much as $10,000 per month (Kontzer 2003).

Affinity communities are direct file sharing groups, such as Napster and Kazaa. P2P holds considerable potential for improved collaboration both within and outside the enterprise, and for software interaction, allowing programs to send data inputs and outputs from one application to another. The affinity communities emphasize knowledge integration over acquisition and learning (Tiwana 2003). Vendors who utilize these communities include Groove Networks, GoneSilent, Pointera, Roku, uRoam, Hilgraeve, FlyCode, Hotline Communications, Kalepa Networks, Centrata, and the infamous Napster (Rutherford 2000). Groove Networks was founded in 1997 with the mission of allowing business teams to collaborate on a broad range of activities within secure, shared virtual spaces, in real time or offline. These systems combine Napster's file-sharing abilities with instant messaging and other capabilities to allow work groups to communicate and share information in a secure environment (Edwards 2002). In 2001, Microsoft invested $51 million in Groove (Groove; Microsoft). Since then they have been working together to unleash the power of collaboration in Windows XP and Microsoft SharePoint Technology. This type of P2P advances the key premise that new value comes from sharing information and building on it (Oram 2001).

P2P systems are attractive because (1) the barriers to starting and growing such systems are low, since they usually do not require any special administrative or financial arrangements, unlike centralized facilities; (2) P2P systems offer a way to aggregate and make use of the tremendous computation and storage resources on computers across the Internet; and (3) the decentralized and distributed nature of P2P systems gives them the potential to be robust to faults or intentional attacks, making them ideal for long-term storage as well as for lengthy computations (Balakrishnan et al. 2003). P2P can also be a way to solve bandwidth issues. By decentralizing data and therefore redirecting users so they download data directly from other users' computers, Napster reduced the load on its servers to the point where it could cheaply support tens of millions of users. How much bandwidth does a simple P2P system like Napster save? Rough estimates made by a company called CenterSpan, which makes a P2P content-sharing system called C-Star, claim that if you put together Napster and the various Gnutella systems and all the knock-offs, you would see about three billion songs traded every month. If you delivered all of those songs from a central server, you would need 25,000 T1 lines costing approximately $25 per month each (Oram 2001). This same philosophy can be successfully used in other P2P systems. In short, P2P cannot only distribute files; it can also distribute the burden of supporting network connections, thus eliminating bandwidth issues.

Bandwidth may be seen as an advantage in some P2P instances, however the bandwidth issue is a reason many organizations avoid P2P all together. While universities are often supportive of new technologies and therefore help them reach critical mass, this was not the case with P2P. University administrators tried to stop it because of the strain it was putting on bandwidth at universities and more recently because of legal risk associated with illegal downloading of both music and video. (Oram 2001) It was estimated that if everyone on campus turned off the outbound KaZaA traffic, approximately 50% more bandwidth could be freed for other Internet traffic (Lee 2003).

From a technical perspective, P2P computing suggests the possibility of providing such system attributes as fault tolerance, performance, and security. From a social perspective, P2P suggests the possibility of powerful communication technologies in distributed form, leading to more robust person-to-person interaction structures (Lethin 2003). The benefits of P2P are strengthened by the increasing availability of powerful communication networks, a growing number of agreed upon technical standards for interfaces and protocols, and more user-friendly clients that make P2P architectures transparent to the user (Schoder & Fischbach 2003).

P2P returns the Internet to its original version, in which everyone creates as well as consumes. However there is no reason to assume that P2P architecture will replace client/server systems. P2P technologies must achieve a comfort level that involves developing hybrid solutions that take the best of P2P technologies and marries them with client/server, resulting in "centralized P2P" (Scannell 2001). Therefore, in the future expect to see hybrid systems that combine the advantages of both approaches (Schoder & Fischbach 2003).

Business uses of P2P

The Gartner Group predicts that P2P will "radically change business models" (Rutherford 2000). P2P proponents forecast that businesses can save billions by using distributed computing setups that take advantage of unused bandwidth and resources. Messaging tools and affinity communities can open up intellectual property and data that are otherwise hidden in departmental offices and servers (Rutherford 2000). Many also see P2P computing as a solution that will relieve network bottlenecks, unleash vast amounts of computing power from underutilized processors throughout an enterprise, and enhance collaboration within workgroups, both inside and outside the organization (Edwards 2001). However, with the image of P2P technology tarnished from legal actions taken against high-profile companies such as Napster many IT executives' feelings about it range from mild skepticism to pure paranoia (Scannell 2001). This paranoia, along with IT executives reliance on the centralized server as a way to control and secure their companies' mission critical data, carry out important functions such as backups, and host their e-commerce applications has led to their reluctance to deploy P2P technologies.

Intel is taking the lead in trying to improve management's perspective on P2P. They are spearheading a number of efforts to ensure that P2P networking is widely adopted by service providers and consumers. Pat Gelsinger, CTO of Intel, is one of the strongest proponents:

> *"We suggest that it could be ushering in the next computing revolution. Our job now as an industry is to change our perspective and now build the infrastructure that allows P2P computing to emerge in a broad way"* (Spooner 2000).

Intel's P2P visionaries see consumers and companies creating "self organizing webs", which would consist of employees at a company, family members or any group with common interests or goals. Users in these private networks would be able to share spare systems resources such as storage, or use the system to exchange files. Intel outlined possible applications for P2P networking for business; collaboration, distributed computing, file sharing, and edge services[1]. (Spooner 2000). Intel also organized the

[1] Edge services puts important content on distributed client PCs so it is more readily available to other nearby clients.

formation of a P2P working group of 18 companies including IBM, HP, and Applied MetaComputing LLC to help develop standards for P2P. Aside from driving P2P standards and developing P2P products, Intel intends to invest in P2P companies through its Intel Capital division, and lend some of its own technology to the field.

Collaboration is one of the applications that Intel identifies as being conducive to P2P networking. P2P collaboration will enable companies to tap into their own intellectual property locked up in boundaries such as separate offices and servers with restricted access. This can help corporations unleash the knowledge locked away on workers' PCs (Cortese 2001). P2P connections allow knowledge workers to communicate, collaborate and create. This collaboration now involves connecting people to people, people to machines, and enterprises to enterprises, increasing the overall knowledge management within organizations (Krill 2001). P2P even allows for the leveraging of individual expertise at the periphery of the network, extending organizational reach (Tiwana 2003).

The P2P model is naturally extensible to knowledge management applications because of its ability to spontaneously facilitate the rapid integration of previously unconnected expertise (Tiwana 2003). This was not the case with client/server networks, where knowledge transfer is only possible among individuals who previously agreed to collaborate. Resources exchanged via P2P networks include human expertise, including tacit knowledge, insight, rules-of-thumb, and lessons learned, in addition to files, processing cycles, and disk storage (Tiwana 2003). Each additional member of the network increases the network's potential value. Thus, P2P more closely adopts the conventions of face-to-face human communications.

Business objection to this type of P2P is that workers are not used to letting colleagues access their desktop PCs for documents. But young consumers have grown up with instant messaging, playing games online and downloading music; primarily peer-based activities. As these young people move up into the workforce, analysts believe they will bring with them a positive attitude to – and even a demand for- simpler P2P approaches for sharing documents and other files (Forrester 2001). It is even predicted that in a few years, corporations are going to be paying big money for IT professionals who can fill a hot new job title: *collaborative development manager*. This manager will marshal development teams dispersed over the globe and tie them together with P2P tools to create great software. The result will be cost savings that exceed the offshore model. The overall IT mega trend toward collaboration and P2P technologies is accelerating worldwide. One expert predicts that corporations that experiment now by finding the right people to make the new collaborative model work will enjoy cost savings not available by any other method (Mezick 2003).

Distributed computing is another effective business use of P2P computing. The basic premise is that there are PCs around the organization doing nothing or very little. University of Wisconsin researchers estimate that companies use less than 25% of the computing and storage capacities that are already paid for (Fox 2001). What's in it for business? Distributed computing is a way to harness the dormant process power in desktop PCs. Companies like defense and aerospace giant Boeing, oil company Amerada Hess, and Intel say they have been able to reduce the need to buy high-end computer systems, including mainframes, by using P2P networking to tap into the processing power that is already available on their desktop PCs. P2P systems may also lessen bandwidth requirements, an important benefit for companies with networks jammed to capacity (McDougall 2000). Distributed computing can help transform the bottom line by defraying the cost of high performance hardware, reducing the administrative overhead of managing and maintaining long-running applications, and increasing opportunity by accelerating the decision-making process (Lee).

P2P computing can also increase an organization's storage capacity, thus reducing the need to buy higher priced server disk drives, or even implement Storage Area Networks (SANs)[2]. As previously mentioned, it is estimated that organizations use less than 25% of their storage capacity. Most of this unused capacity resides on individual users' desktops. Desktop computers are being purchased with larger and larger hard drives. Today's desktops typically arrive with hard drives of 40 to 80 gigabytes. For every 100 desktops that can translate into six untapped terabytes of storage. Companies can use P2P computing to gain efficiencies in their distributed storage networks to utilize this unused space. Software is available for determining the optimal distribution patterns and paths for storage loads, and for leaving a digital blueprint on each client identifying where to find the rest of the data. P2P can also help smaller businesses use their limited bandwidth more efficiently. By distributing storage loads, limiting downloads from Web servers, and sharing processing tasks at the LAN level, P2P can move activity away from the Internet to the corporate LAN, where bandwidth is more plentiful and easily managed (Edwards 2001).

Instant messaging is a tool that is already being utilized by many organizations. Instant messaging is the best example of a P2P application that has been rapidly adopted by businesses as a central communications tool. IDC expects the number of corporate users of IM programs to grow to over 180 million users by 2004 (Spangler 2001). Businesses can benefit from these P2P messaging tools. The messaging tools allow synchronous communication that is not possible with email, so remote workers collaborating on a project can instantaneously chat and complete a task, providing a richer media than email (Rutherford 2000).

The extent to which P2P is implemented for both internal and external business purposes will greatly depend on the technological, economic, and legal challenges posed by the technology and how they become resolved (Schoder et al. 2002). For business, server based P2P will probably win out, since P2P requires you to open new ports on your firewall. IT executives are more willing to downgrade to a port they feel comfortable with, which is usually HTTP, which requires you to have a server (Scannell 2001). However, the ultimate deciding factor on implementing a P2P solution is simple:

[2] For more information regarding Storage Area Networks (SAN) see an earlier ISRC Tech Briefing at www.uhisrc.com/FTB/SAN/Storage_Area_Networks.pdf.

to succeed, a P2P solution must be easily implemented and ultimately profitable (Alwang 2001).

Concerns with P2P

While the P2P architecture offers many promises, practitioners have raised a number of concerns. Will decentralized control be able to cope with challenges such as network control, security, interoperability, metadata, and cost sharing (Schoder & Fischbach 2003)? Unlike client/server architecture which favors a top-down design and planning approach, P2P encompasses a bottom-up approach. As a result, it is difficult to plan or predict the development, size and connections, as well as throughput and stability of P2P networks (Schoder & Fischbach 2003).

Security is still the primary concern with P2P technologies. It has not evolved to a point that allows groups to transparently communicate on both sides of the firewall. Many P2P technologies demand special ports to be opened in a firewall, leaving businesses vulnerable to hackers (Scannell 2001). The implementation of P2P technologies create additional security challenges as their use may require that third parties be allowed to access the resources of an organization's internal systems. As a result, conventional security mechanisms, such as firewalls, are frequently circumvented during communication in P2P networks (Schoder & Fischbach 2003). However, there are organizations that are declaring P2P to be secure. McAfee trusts P2P to distribute updates to their virus-detection software. McAfee ASaP is a service provided to large companies to let them distribute updates quickly throughout their organizations. Instead of making numerous individuals contact the McAfee Web site (a sure recipe for network overloads), a few initial systems contact the McAfee site, and then pass on the software to other systems in a chain. This form of P2P is called rumor technology and operates like a beneficial virus. McAfee's rumor technology is not only more efficient than routine Web downloads, but more secure. Employees of each company have to go outside their corporate network only a few times to get the software. Most of the networking takes place inside the corporate network, presumably protected by a firewall and general LAN architecture (Oram 2001).

A major problem with the implementation of P2P technology involves locating resources that may be significantly more difficult to identify than MP3 files. In order to convert raw data into usable information and make this available for efficient searches, acceptable metadata concepts will be required (Schoder & Fischbach 2003). One approach is to maintain a central database that maps a file name to the locations of servers that store the file as was used by Napster. However this solution has an inherent drawback; the database is a central point of failure. Another solution being used by P2P software such as Gnutella is a symmetric lookup algorithm. In this case, all nodes are treated equally and share only a small role in the lookup process. It begins when a user broadcasts a message to all its neighbors looking for X. When a node receives the request, it checks its local database. If it contains X then it responds to the user with the item. Otherwise it forwards the requests to its neighbors which execute the same process (Balakrishnan et al. 2003).

Reciprocation is another concern with business P2P systems. Many users of P2P networks are "freeloaders" who take advantage of the available resources but do not contribute in return. This undermines the characteristics of P2P and has a negative effect on the willingness of users to make resources available (Schoder & Fischbach 2003). It is estimated that 70% of P2P system users only download files without reciprocating by uploading files (Adar & Huberman 1998). Trust will be a big factor in creating a sharing environment within organizations. As the affinity community is smaller and more intimate, as would be the case in most organizations, reciprocation should become less of an issue. Reward systems can also be tweaked to encourage sharing.

Finally, the lack of standards is causing business to take a cautious approach to implementing P2P. Standards will promote ubiquity that is essential for P2P computing success. To assist in this endeavor, the Peer-to-Peer Working Group[3] was formed in October 2000 to aid in the advancement of infrastructure standards for P2P computing. The P2P Working Group is a consortium of corporations dedicated to developing a host of industry standards for P2P computing technology. The P2P Working Group plan to address P2P issues such as interoperability and performance of computing devices, security, management, privacy of data stored in web devices, common protocols for the way that information flows between, and is shared by, users of P2P devices. The overall goal is to develop infrastructure standards to enable P2P computing everywhere (Merkow 2000). If P2P is to become a standard business application, techniques and methods for authentication, authorization, availability, data integrity, and trust have to be integrated (Schoder & Fischbach 2003).

Conclusion

P2P computing has already been demonstrated to be a useful business architecture. For some current successful business uses see the Appendix. Yet to become a standard business architecture, many of the concerns addressed in this briefing will have to be addressed. P2P will continue to grow with the general public. P2P excites people because they can participate and make a difference. Even something as impersonal as SETI@home, where users downloaded software that perform calculations in the background, attracted millions of volunteers (Oram 2001). However, expect to see more business uses emerge in the near future. It is very likely that the design philosophy underlying P2P networks will gain importance in the development of mobile business and ubiquitous computing, especially when the goal is to establish communication between mobile network peers, including PDAs, laptops, and mobile telephones (Schoder & Fischbach 2003). P2P is probably the only workable architecture in the world where

[3] The Peer-to-Peer Working Group can help keep you updated on the latest P2P information. To become a member go to http://peer-to-peerwg.org/members.

billions of these devices can be connected to networks. Today's client-server architecture simply will not be able to handle such a large universe. (Gillmore 2001). We are not even beginning to use all the power of the machines and people at the edge of our networks, and the intelligence and creativity at the edges will provide some of the greatest value as we move forward (Gillmore 2001). P2P provides an architecture to allow businesses to tap this wealth of knowledge and power. In fact, both Microsoft Chairman Bill Gates and Intel Corp. Chairman Andrew Grove have both said they believe P2P is going to be very important (Disabatino 2000). Yet, today P2P remains an architecture still waiting for a killer application (Agree 2003).

References

Adar, E. and Huberman, B.A., Free riding on Gnutella, First Monday 5, 10, 1998.

Agre, P.E., P2P and the Promise of Internet Equality, *Communications of the ACM*, February 2003, Vol. 46, No.2.

Alwang, G., Peer-to-Peer for Grown-Ups, *PC Magazine*, January 16, 2001.

Balakrishnan, H., Kaashoek, M., Karger, D., Morris, R., Stoica, I., Looking Up Data in P2P Systems, *Communications of the ACM*, February 2003, Vol. 46, No. 2.

Cope, J., P2P Over The Internet, *Computerworld*, April 08, 2002.

Cope, J., Peer-to-Peer Network, *Computerworld*, April 08, 2002.

Cortese, A., Peer to Peer, *Business Week*, Spring 2001.

Disabatino, J., What's So New About Peer-to-Peer?, *Computerworld*, Nov. 20, 2000.

Edwards, M., P2P: Next computing wave or more vendor hype?, *Communications News* October 2001.

Edwards, M., Startups vie for P2P application niches, *Communication News*, Jan. 2002.

Forrester, P2P Gets Popular with the Public First, *Forrester Research*, 2001

Fox, P., Potential Uses Help Brighten Future of P2P, *Computerworld*, Feb. 26, 2001.

Gillmore, D., A Note to IT: Why You Need to Know P2P, *Computerworld*, Feb. 26, 2001.

Groove Networks and Microsoft, *Groove Networks*, www.groove.net/about/microsoft/.

Hall, M., Intel taps P2P technology for cancer research, *Computerworld*, April 03, 2001.

Ives, B. & Piccoli, G., Custom Made Apparel and Individualized Services at Lands End, *Communications of the Association for Information Systems*, Vol. 11, 2003.

Jones, M., Ray Ozzie, others spell out evolution of P2P, *InfoWorld*, June 20, 2001.

Krebs, B., Plug.In: Peer-To-Peer Networks Show Mainstream Potential, Newsbytes, July 24, 2001.

Krill, P., Intel preaches peer-to-peer, *InfoWorld*, June 13, 2001.

Kontzer, T., The Word Is Out on Instant Messaging, *Information Week*, May 14, 2003.

Kubiatowicz, J., Extracting Guarantees from Chaos, *Communications of the ACM* February 2003, Vol 46., No.2.

Lee, J., An End-User Perspective on File-Sharing Systems, *Communications of the ACM*, February 2003, Vol. 46, No.2.

Lethin, R., Technical and Social Components of Peer-To-Peer Computing, *Communications of the ACM*, February 2003, Vol. 46, No. 2.

Laudon, K, & Laudon J., Management Information Systems: Managing the Digital Firm, Prentice Hall: New Jersey, 7th Ed, 2002.

McDougall, P., The Power of Peer-to-Peer, *Information Week*, August 28, 2000.

Merkow, M., Peer-to-Peer Working Group Forms, *EC Outlook*, Dec. 28, 2000.

Mezick, D., Outsourcing 2.0: Collaborative Development, *Computerworld*, Aug. 18, 2003.

Minar, N. and Hedlund, M., A network for peers. Peer-to-peer models through the history of the Internet. In *Peer-to-Peer: Harnessing the Benefits of a Disruptive Technology*. A Oram, Ed. O'Reilly, Sebastopol, CA. 2001; 3-37.

Oram, A., Peer-to-Peer for Academia, *The O'Reilly Network*, Oct. 29, 2001.

Penserini, L., Liu, L., Mylopoulos, J., Panti, M., Spalazzi, L., Cooperation strategies for agent-based P2P systems, *Web Intelligence and Agent Systems: An international journal* 1(2003) 3-21.

Rutheford, E., The P2P Report, www.cio.com/knowledge/edit/p2p.html, December 01, 2000.

Scannell, E., Peering into the future, *InfoWorld*, April 27, 2001.

Schoder, D. and Fischbach, K., Peer-to-Peer Prospects, *Communications of the ACM*, Feb. 2003, Vol. 46, No. 2.

Schoder, D., Fishbach, K. and Teichmann, R. , Eds *Peer-to Peer*, Springer Verlag, Heidleberg, Germany, 2002; www.whu.edu/ebusiness/p2p-buch/.

Seymour, J., P2P Goes to Work, *PC Magazine*, May 8, 2001.

Shirky, C., Turelove, K., Dornfest, R. and Gonze, L. P2P Networking Overview. *The Emergent P2P Platform of Presence, Identity, and Edge Resources*. O'Reilly & Associates, 2001.

Spangler, T., Harnessing the Power of Corporate P2P, *ZDNet US*, June 13, 2991.

Spooner, J. and Popovich, K., Intel: The future is peer, *ZDNet News*, August 23, 2000.

Tiwana, A., Afinity to Infinity in Peer-to-Peer Knowledge Platforms, *Communications of the ACM*, May 2003, Vol. 46, No. 5.

Appendix
Examples of P2P Business Applications

- The Intel Philanthropic Peer-to-Peer Program enables computer users with Internet access to donate unused computer time to help solve medical research problems. By using the collective idle time of 6 million PCs connected to the Internet, the program could provide up to 50 teraflops of computing power. The National Foundation for Cancer Research has estimated that by using supercomputers, it can take 24 million computer hours to analyze the data on a single protein problem. Users can download software that will install on their PC. Craig Barrett, CEO of Intel said the software's security is "robust end to end", and is allowing it to be run within his company (Hall 2001).

- Intel is using a type of P2P, called Netbatch, for distributed processing. Netbatch allows engineers to take advantage of unused processing power on some 10,000 workstations across 25 worldwide Intel locations to run computer simulations for chip design. The process has shortened the time for developing new processors and reduced outlay for new mainframes. Intel claims to be saving up to $500 million a year with NetBatch (Spooner 2000).

- Intel also uses P2P to streamline the distribution of computer-based training materials to employees. The firm's IT department did not want employees to download huge multimedia files from a central server, so it developed an application on every desktop to reduce the network burden. When a user requests a course, the application searches for it on local desktops, gradually widening the search until it finds the closest source. (Edwards 2001)

- Ford Motor Co. is using P2P technology developed by Oculus Technologies of Boston to design and produce more fuel-efficient cars. P2P allows Ford to connect its far-flung design team members to help them evaluate more design iterations over a shorter period of time. "We think the result for Ford should be optimal product development, saving us from $5 million to $15 million per design program. This could have even greater impact on our final product by integrating the entire automobile design process," says John Goodman, fuel economy implementation manager for Ford Motor Co (Scannell 2001).

- Mojo Nation of Mountain View, CA, also cuts the cost of content distribution by breaking the task into lots of smaller data-delivery subtasks, and letting each peer device contribute as much as it can to the delivery effort. In this way, the company can efficiently aggregate numerous low-bandwidth peers (Edwards 2002).

- J.P. Morgan Chase & Co. is using P2P to distribute big processing tasks such as risk management calculations, across employees' computers at night (Cortese 2001).

- New York based Amerada Hess has been experimenting with P2P networking and resource sharing. Its Beowulf Project strings together 200 Dell desktop PCs with Ethernet and Linux to handle complex seismic data interpretation. This allowed them to replace a pair of IBM supercomputers. "We're running seven times the throughput at a fraction of the cost," say CIO Richard Ross. The company has two other projects in the works. One lets computers borrow processing from other workstations on the network. The other takes what Ross calls a Napster approach to knowledge management, where each desktop's storage joins with others to become a huge data repository (McDougall 2000).

- Law firm Baker & McKenzie is implementing P2P software to capture and share the knowledge of its 3,000 attorneys in 60 offices (Cortese 2001).

- Boeing engineers are using distributed computing to harness the MIPS required to run complex performance tests. Boeing, like Napster, is using a P2P model in which servers route traffic to the intended destination. "No single machine has been able to keep up with our requirements," says Ken Neves, director of Boeing's computer-science research organizations (McDougall 2000).

- Boston based WorldStreet has tailored its P2P system to the financial community. The issue WorldStreet's prospective customers face was the fact that analysts were bombarded with hundreds of email messages, without any indication of how critical the information actually was. WorldStreet's plug-in for Microsoft Outlook lets an investment bank's customers decide what kind of research they want to receive, for which companies and which financial analysts. "What's different about our P2P product is that it's a completely balanced relationship. You can set up profiles to accept only the information you care about. It's information per your specifications" says Rod Hodgman, WorldStreet's COO. Bear, Stearns & Co. signed on to use the product. "A typical portfolio manager gets 400 to 500 emails a day," says Stanley Salellson senior managing director of institutional equities, "what WorldStreet does is filter the information to determine whether something is really pertinent" (Spangler 2001).

- One P2P project, called the Free World Dialup project, aims to share phone lines. It takes Net phone calling a step further by allowing a person to "borrow" the phone line of someone else on the network and turn a long-distance call into a local call. This may sound like typical Internet telephony, but these are between telephones – not PCs. If you are in Houston and want to place a call to London, you just dial the number and the system routes the call through the Net to another PC that is physically located in London. The London PC then places the local call (Cortese 2001).