Peer-to-Peer Computing and Business Networks: More Than Meets the Ear
Part 1 – What is P2P?

A CTOMentor White Paper

December, 2001
CTOMentor
A Service of Stratvantage Consulting, LLC

Published December, 2001

About This White Paper
CTOMentor has published this white paper in two parts. This first part offers an introduction to peer-to-peer computing and general information about the P2P trend. Part 2 features in-depth analysis of current P2P applications and vendors as well as recommendations for P2P use. The second part is available for a nominal fee at CTOMentor. See more at www.ctomentor.com/p2p/part2.htm.

About CTOMentor
CTOMentor is a subscription advisory service tailored to a technical leader's industry-specific and personal technology information needs. Four times a year CTOMentor will provide a four-hour briefing on the most important emerging technology trends that could affect your business. See more at www.ctomentor.com.

About Stratvantage Consulting
Stratvantage brings years of Internet business-to-business (B2B) electronic commerce experience to the corporate strategy planning process. Stratvantage helps develop a strategy and a plan of attack to exploit the opportunities fast moving emerging technologies like the Internet, wireless, and peer-to-peer computing can bring to business. Stratvantage has worked with companies of all sizes, from the Fortune 15 to startups. We specialize in Internet strategy and vision as well as permission marketing and strategy execution.

Thanks to Geneer Corporation, which funded a previous version of this white paper.
www.geneer.com

Table of Contents

Introduction

Right now your computer is wasting resources. Millions of CPU cycles are doing nothing more than warming the air. Megabytes of disk space sit unused. Your broadband connection idles as you read downloaded Web pages. On a daily basis, you and everyone you work with typically use only a small fraction of your computers' capacity.

What if you could put all these underemployed resources to work? And what if you could earn money doing so? That's the business proposition hundreds of peer-to-peer (P2P) startups are chasing today, trying to be the first to spin the estimated 10 billion MHz of processing power and 10,000 terabytes of storage[1] connected to the Web into Internet gold.

If you've heard of P2P at all, you probably associate it with Napster, the consumer music file-sharing service that's been in the news, and in trouble, over copyright infringement issues. Or perhaps you've heard of Morpheus or KaZaA, two of the services vying to replace Napster as the most popular music download service.

But P2P is much more than college kids ripping CDs and sharing them with the world. It is a movement that seeks to change the fundamental way the Internet is organized, returning it to its peer-to-peer roots.

This white paper is published in two parts. Both examine the business uses of P2P, as opposed to the consumer, file-sharing uses. In this first part, CTOMentor defines P2P, takes a look at the four prongs of the P2P effort, delineates who will benefit and who will be threatened by the technology, and presents some of the problems facing P2P.

In Part 2, CTOMentor looks at examples of the types of P2P applications currently in development as well as some of the possible future uses for the technology. We also take a look at what some of the major players in the computer industry have to say about P2P and make some predictions for how the technology will affect the evolution of the Internet and business networks. Part 2 is available for a nominal fee at the CTOMentor Web site.

But first, let's define what it is we're talking about when we say P2P.

[1] http://www.business2.com/content/magazine/breakthrough/2000/11/20/22119

What is P2P?

P2P is not all about music, or Napster, or file sharing, or intellectual property. It is not even really a class of applications, but rather a technology or, more properly, a computing concept. The concept of peer-to-peer computing has been with us since the dawn of the Internet, which was conceived as a network of computer peers – the large computers at universities and government installations. As the Internet has grown, concerns about security and a scarcity of IP (Internet Protocol) addresses led to the growth of firewalls, which isolate private networks from the Internet and conceal the identity of the computers behind them.

If you're behind a firewall, Network Address Translation (NAT) masks your true IP address by transforming it into a common address for all private network users. While this is efficient and secure, it means outsiders cannot locate your computer, and thus you cannot run a Web server that the world can see. Similarly, Internet Service Providers (ISPs) maximize their scarce IP resources by dynamically assigning your computer an IP address whenever you connect. Thus your address is different today than it will be tomorrow, and this fact also prevents you from running a Web server or otherwise sharing information from your computer.

The peers in the phrase peer-to-peer are users' computers. Rather than the currently prevalent client/server model, in which all communication passes through and is controlled by central servers (Web servers, FTP servers, application servers and the like), in P2P, the communication goes directly from your computer to another user's computer.

This quality of P2P is both its promise and its threat. Without central servers, there is no central authority, and without a central authority, many people believe there can be no central control. And lack of control makes IT managers very nervous. IT departments are used to being able to dictate what happens on their networks, and on their users' desktops. The ability to directly access the contents of an individual computer, despite the presence of a firewall, can also be a security problem. You may want people to be able to download your music files, but you certainly don't want them to see the rest of your computer, or your network.

P2P Defined

Clay Shirky, a principal in the acceleratorgroup, a venture capital firm, defines P2P this way:

P2P is a class of applications that takes advantage of resources – storage, cycles, content, human presence – available at the edges of the Internet. Because accessing these decentralized resources means operating in an environment of unstable connectivity and unpredictable IP addresses, P2P nodes must operate outside the [Domain Name System, which assigns names to computers] and have significant or total autonomy from central servers.

http://www.openp2p.com/pub/a/p2p/2000/11/24/shirky1-whatisp2p.html

The current P2P craze began as a grass roots movement to empower ordinary people to be content publishers, if only of their digital music files. Without some alternative way to find users with dynamic or NAT-translated IP addresses, Web publishing is a frustrating experience. You would need to obtain a permanent network address and domain name (which takes money), and install and maintain a Web server security system, and content publishing application (which takes significant software knowledge). Or you can just download Morpheus and be sharing music files with anonymous strangers in minutes.

File sharing applications such as Napster and Morpheus, and even Instant Messaging (IM) applications such as AOL Instant Messenger (AIM), solve the problem of locating peers by maintaining an address service outside of the normal Domain Name Service (DNS) system. When your browser wants to locate a Web site, it looks up the name you typed using DNS. When you want to locate a buddy on AIM or a song file on Morpheus, your computer contacts a non-DNS central directory for the information.

P2P Doesn't Mean Serverless

A common misconception about P2P is that the technology does not employ centralized servers of any kind. Perhaps in the pure definition of peer-to-peer this is true. But the applications being developed today are intended to solve problems, not be ideologically pure. This means, like AIM and Napster (and many of its successors), many of the most popular P2P applications do indeed use central servers to facilitate peer communication in an architecture called Brokered or Hybrid P2P.

Ironically, it is the use of a central server to house the music index that was the key to Napster's legal troubles. Although all Napster transfers happen between peers, searches are performed on the Napster index, which the company maintains in a central location. Maintaining this index exposed Napster to legal attack by recording industry copyright owners. The adjacent figures show the differences between client/server computing and the two dominant P2P configurations.

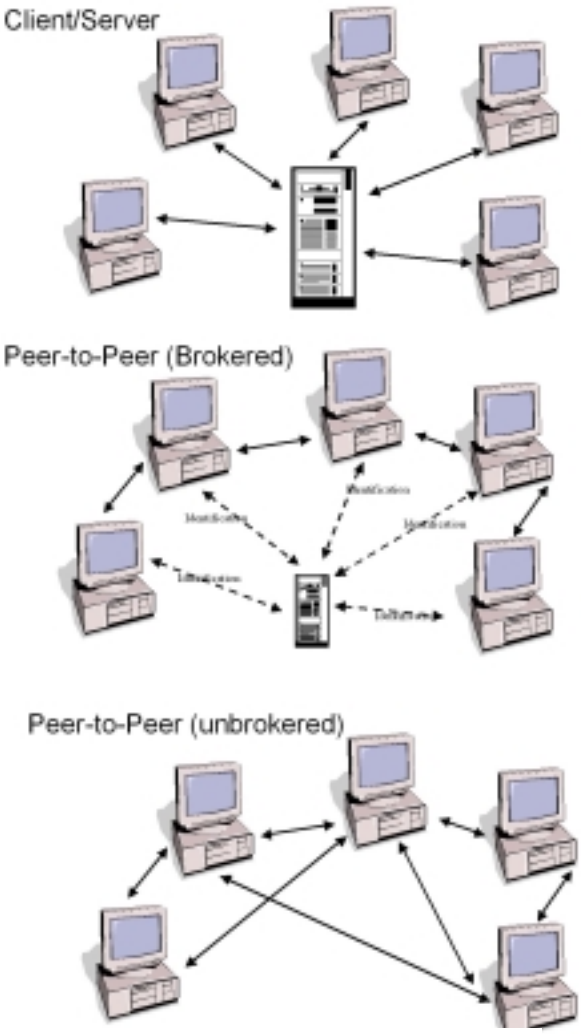

Source: Yankee Group

Collaboration P2P vendor Groove Networks[2] is a good example of the Hybrid approach. Groove, Lotus Notes developer Ray Ozzie's brainchild, uses a central server for some types of file distribution. Among its many services, Groove offers a collaborative document editing service. When a workgroup member is offline and other members edit a document, Groove employs a relay server that tracks the changes and transmits them to the absent member once he or she is connected again.

Although this latest P2P wave started with music file sharing, P2P has grown well beyond its content-oriented roots to embrace a whole range of network applications.

[2] http://www.groove.net/products/

The Four-Pronged P2P Effort

There are four major prongs in the P2P effort today. All four involve direct communication between peers as well as putting unused computing resources to use. All four also delegate significant authority to peered computers on the edges of the Web rather than emphasizing centralized servers.

Content Serving

Also known as file sharing, content serving refers to what Napster, Wrapster, Aimster, Gnutella, Morpheus, and other consumer-to-consumer (C2C) technologies do. These services enable people to serve content off their local hard drives and share files with other people. These applications represent a growing trend of decentralized content, and consequently, decentralized control.

Putting aside the legal and moral issues, in hindsight, one can see that content serving services were inevitable. Increases in PC capabilities combined with wider availability of always-on bandwidth made decentralizing content feasible. In the consumer arena that means sharing music and other entertainment files, but this is only one possible application. As this technology makes its way into business, it may threaten the business models of centralizing forces such as content aggregators, demand aggregators, and catalog aggregators.

An obvious place for P2P content serving is the corporate intranet. While intranets have been somewhat successful, it can be argued that they have been hampered by the necessity for central control. People have jobs to do, and have little patience for the bureaucratic and technical procedures often required to publish their content on the intranet. The Hybrid approach, a centralized index with distributed, redundant content, could be used to break the intranet logjam and enable increased communication of company best practices and competitive intelligence. One good example of this usage is Jibe, which enables employees to quickly and easily share information on a corporate network. There's more on Jibe in Part 2[3] of this white paper.

Bandwidth Sharing

Video producers and other multimedia event producers have traditionally struggled to serve high bandwidth content such as video events from central servers. There have been several spectacular flameouts when thousands of users converged on a single site to watch a video. The most famous was the 1999 Victoria's Secret debacle, when 1.5 million viewers tried to watch a catwalk video. Most viewers had an unsatisfactory experience, with 5 percent not able to see the video at all.

For subsequent shows, the company has employed companies known as Content Distribution Networks (CDNs). CDNs such as Akamai farm the video out to hundreds of machines they've located in geographically dispersed data centers. When a user requests the content, it is served from an Akamai machine close by rather than from a huge central server, thus saving producers lots of money. But there's a flaw in Akamai's model: It involves buying thousands of servers (13,000 currently) and locating them all over the world (in 54 countries). These servers will become obsolete eventually and incur huge capital costs to replace, and the logistics of operating in 54 countries will force Akamai's prices to remain high.

[3] http://www.CTOMentor.com/p2p/part2.htm

Contrast this with the P2P approach being taken today by vTrails, AllCast, and ChainCast. These companies enroll users worldwide whose PCs will relay content to other users nearby. It is likely that this scheme will not attract volunteer users, so event producers will incur costs to rent users' bandwidth and disk space to serve their content. While delivery costs using these networks are a currently a fraction of traditional CDNs' charges, the technology is so new, and the P2P CDNs customer lists so limited, it remains to be seen whether the P2P scheme can scale and be effective. There's more on all three companies in Part 2 of this white paper.

Collaboration

Although P2P collaboration involves processor, bandwidth, and disk sharing, it really is a different category of application because it emphasizes the real-time communication among a group of users. While services such as Instant Messaging (IM) and online presentations are examples of P2P collaboration, some of the current crop of P2P collaboration vendors bundle many services to enhance interactivity. For example, consumer services such as Aimster bundle Instant Messaging with content serving to enable friends to set up collaborative networks to chat and share files.

More-sophisticated uses include Placeware's[4] mediated business meetings, and Groove Networks. Groove wraps several P2P technologies in a business-friendly package. Using Groove, business users can co-edit documents, display shared whiteboards, share files, and use Instant Messaging, live voice, video and threaded discussions. Groove's value-add is to make these services available in a secure environment. There's more on Groove in Part 2 of this white paper.

Hive Computing

In hive computing, computers' unused processor power is harnessed to attack huge computing problems like virus modeling or Wall Street securities analysis. The problem is split up and assigned to cooperating computers over a network. These computers work on their bit of the problem and return the results to a central location. Typically, hive-computing applications run only when the user is not using the computer, and often take the form of a screen saver.

Non-profit effort SETI@Home is perhaps the best known of the hive computing efforts. Millions of people have donated unused computer processing cycles to the Search for Extraterrestrial Intelligence. When their computers are idle, a screen saver analyzes radio telescope information for evidence of intelligence. We discuss SETI@Home in more depth in Part 2 of this white paper.

This P2P prong may turn out to be one of the most important for businesses. Intel estimates that a typical large business has two-and-a-half times the computing power in individual computers than is available from its servers.[5] There are a host of startups focusing on this area, including DataSynapse, United Devices, and Parabon. For a more in-depth look at hive computing, see Geneer's white paper, *The Buzz About Hive Computing*[6]*.* There's more on DataSynapse and Parabon in Part 2 of this white paper.

[4] http://www.placeware.com/
[5] http://news.cnet.com/news/0-1003-200-2603611.html
[6] http://whopper.geneer.com/web/inforeq.nsf/whitepaper+download+2?OpenForm&whitepaper=hives

What is the P2P Market?

There's no denying there's a lot of furious activity in the P2P space, and the applications, both existing and planned, sure are cool. But is there a market for all this innovation? Or is P2P yet another technology looking for a problem to solve?

In the rapidly accelerating technology marketplace, new ideas are coming ever more quickly. The problem is, more-rapid hype is resulting in more-rapid deflation. A new trend like P2P, which only gained real momentum in Summer 2000, has already spawned a backlash, and some pundits are predicting its doom.[7] At the end of 2001, after 18 months of furious hype, company failures and consolidations, and despite a tremendous range of software development, P2P has definitely cooled. It may only take a few incidents like the Mandragore virus that struck music sharing service Gnutella in early 2001 to put the P2P trend into full retreat.[8]

Hype Cooling

At the end of 2001, after 18 months of furious hype, company failures and consolidations, and despite a tremendous range of software development, P2P has definitely cooled.

The enthusiasm for information technology in general has diminished this year, and so has the P2P hype. "Clearly there's not the same level of wild enthusiasm for P2P as there was in the early part of this year," said Kevin Werbach[9], editor of EDventure Holdings' Release 1.0 newsletter. "Still, there are a lot of things happening. More people are believing in the concepts, and more practical work is being done. Although things are moving slower, and fewer concepts are getting funded, this is still a sector that is drawing more than its fair share of interest and investment."

Werbach said that a year ago, there were no real products or customers, but today there are both, with many of the Fortune 500 doing P2P pilots. The most straightforward sell, he feels, is hive computing. Enterprises with huge processing needs will implement these solutions on their intranets. It was a mistake, in Werbach's estimation, for many hive computing companies to have started with the SETI@Home model of computing over the Internet (see Part 2 of this white paper for more information on this model). Concerns over security and reliability have made the computing as an Internet service model a tough sell for corporations. A market for Internet hive computing will likely come later as companies become more familiar with the concept and the security models.

Despite the hype and deflation, though, new technologies that really solve problems are likely to survive. In the following section we take a look at what some other industry analysts are saying about P2P.

[7] "Peer-to-Peer Party Comes to a Halt As Companies Morph or Disappear," Wall Street Journal, April 4, 2001.
[8] http://www.idg.net.nz/webhome.nsf/UNID/533B8A279D66D96DCC256A01000C57AD!opendocument
[9] Interview with Kevin Werbach, 11/13/01.

What Are the Analysts Saying?

Industry observers have been fairly uniformly positive about P2P's market potential, although some rosy early predictions have not quite panned out. Hurwitz Group predicted that 500 P2P companies would be funded in 2001. Stratvantage's *P2P4B2B* directory[10] counts 252 companies, not all of which were founded this year. Add to that the P2P companies created on the business-to-consumer (B2C) side, and the Hurwitz prediction may be close.

According to the Yankee Group early in the year, P2P was to be a big buzzword in 2001, but only certain segments would see real dollar growth. They predicted that specialized distributed computing applications in areas such as bioinformatics, online disk storage, and statistical analysis will provide real opportunity. While there are not a whole lot of companies making big money in the P2P space, certainly the hive computing firms have attracted significant interest. As far as buzz goes, P2P was eclipsed as the reigning hot buzzword at mid-year 2001 by Web services, a somewhat related decentralized technology.

> **Important Technology**
>
> "Any IT manager who fails to look at peer-to-peer should be fired."
>
> Cheryl Currid, president of Currid & Co

Analysts looking farther out still have time to be proven right. Currid & Co. estimated the distributed computing market alone will reach $9.5 billion in services, $7 billion in hardware, and $13 billion in software by 2006. Cheryl Currid, president of Currid & Co., said earlier in the year, "Any IT manager who fails to look at peer-to-peer should be fired."[11] More recently, Currid said, "It's not going to happen before IT managers are comfortable with security, stability, and reliability, and that could take some time." [12]

GartnerGroup predicts that by the end of 2002, more than half of all Internet users will regularly sign onto at least two peer-to-peer (P2P) applications. Since Gartner's model of P2P seems to include Instant Messaging (IM), the prediction hardly seems outrageous. The firm predicts that, "By 2005, instant messaging will surpass e-mail as the primary way in which consumers interact with each other electronically. By 2005, IM will be integrated into 50 percent of the applications that businesses use to directly interact with their customers."[13] On the other hand, Gartner predicts hive computing applications will be limited to specific heavy-processing and low-security applications through 2005.[14]

Gartner further believes that by 2003, 30 percent of corporations will have experimented with P2P applications for content distribution. "The way people get things done is by working together in small or large groups," said Rob Batchelder[15], Gartner's research director for Internet infrastructure. "With peer to peer, we move from personal to interpersonal computing." Overall, Gartner says most enterprises "can afford to take a wait-and-watch approach to P2P technology. For those that like to be on the leading edge, products are now emerging that provide a basis for experimentation and enable exploitation of the radical decentralization that P2P enables. Even

[10] http://www.Stratvantage.com/directories/p2pcos.htm
[11] http://www.intel.com/pressroom/archive/speeches/pg082400.htm
[12] http://www.informationweek.com/story/IWK20010802S0011
[13] http://www4.gartner.com/DisplayDocument?id=338943&acsFlg=accessBought
[14] http://www3.gartner.com/DisplayDocument?id=323309&acsFlg=accessBought
[15] http://www.nytimes.com/2001/06/13/technology/13BURT.html?searchpv=day05

those enterprises that choose not to engage in P2P at this stage must recognize the technology's longer-term impact and factor it into their strategic planning."

An Omni Consulting Group study found that there might be more companies using P2P technology today than you might think. These companies are already reaping benefits from P2P usage. "Overall, there's 18 percent to 19 percent greater efficiency in the use of system resources with peer-to-peer than with conventional methods," said Omni managing principal Frank Bernhard[16].

Bernhard predicts strong growth for P2P overall. "P2P computing increased 2.7 times from 1999 to 2000," he said. "However, we project a 3.3 times increase from 2000 to 2001." Although most P2P efforts appear to be skunk works at this point, Bernhard is seeing an increase in companies' projected spending on the technology. Interestingly, Bernhard said that those companies using P2P generally tended to be less Internet savvy. The major impetus behind their use of P2P was the grassroots desire to circumvent IT roadblocks and get things done.

Increase Efficiency

"Overall, there's 18 percent to 19 percent greater efficiency in the use of system resources with peer-to-peer than with conventional methods."

Omni Consulting Group managing principal Frank Bernhard

This ability to end-around IT contributes to a fear of losing control that could present a barrier to business acceptance of P2P computing. IT professionals are used to being able to control the computing environment from end to end. Yet this isn't possible when anonymous thousands of Internet and intranet computers are brought to bear on a problem. "CIOs should not feel threatened and should not relinquish the strategy of the organization when considering peer-to-peer," Bernhard said. "There's infancy in it right now, but peer-to-peer has real value for large corporations."

Cheryl Currid said, "Peer-to-peer computing will dramatically change the way enterprise-computing platforms are built, what they cost, and how they tackle computing needs. Using existing resources, organizations will be able to leverage the latent power of linked PCs across the Internet to deliver supercomputing-like processing power to solve complex problems, or even challenge tasks that were previously impossible to undertake."[17] Mark Eggleston, emerging technologies analyst at Currid & Co. said, "P2P is the inevitable evolution of the Internet economy, but that doesn't mean it's happening this year. This year people will pay a lot more attention to it. It will become the B2B of 2001."[18]

Yankee Group sees P2P as somewhat of an equalizer for smaller firms. "Where, for example, in the past only the largest catalog vendors could do massive data processing, P2P computing could enable smaller retailers and catalogs access to the same kind of revenue-lifting activities, such as analyzing cross-sell uptakes and product sales correlation. However, without help from the platform vendors, enterprises will have a hard time figuring it out by themselves."[19]

[16] Interview with Frank Bernhard, March, 2001.
[17] http://industry.java.sun.com/javanews/stories/story2/0,1072,30225,00.html
[18] http://www.crn.com/sections/special/etech/etech.asp?ArticleID=23956
[19] Distributed Computing Offers New Opportunities for Enterprises, www.yankeegroup.com

Steve Kleynhans, an analyst at Meta Group, is somewhat skeptical. "It's too good a technology and appealing of an idea for it not to have some place, but we have to be cautious and not think this is a panacea or the next big thing that will take over the world," he said. "This is another tool in the toolkit." [20]

The big online search portals have much to fear from P2P, according to Forrester researcher Simon Yates. "Vendors are developing technology to perform natural language searches of millions of peer systems that will crush today's centralised searching technology from Inktomi and Google, allowing users to discover other users, not just data and Web pages."[21] Nonetheless, Yates believes, "There's a long way to go before this technology is really something that companies are going to be implementing on a large scale."

"There is no single P2P market, but there is a market for P2P," Alex Veytsel, research associate and analyst at the Aberdeen Group said at the Peer-to-Peer Working Group meeting this summer. "Right now it is clear that there is a diverse portfolio of opportunities for P2P, the question is, which ones will gain the most traction in the months and years to come."

What Is the Industry Saying?

Pundits and prophets aside, the real test of the power of P2P is whether the industry heavyweights are willing to back it. The following section presents some of what the big computer powers are saying and doing about P2P.

Intel

Intel started all the P2P furor by organizing the P2P Working Group in August 2000. "Peer-to-peer is what we think is possibly the next computing frontier," said Pat Gelsinger, chief technology officer of the Intel Architecture Group[22]. The chipmaker has used P2P internally for years to do chip design, and saved millions by not buying mainframes for the task. "We were able to raise our internal compute efficiency from 30 percent to about 80 percent within the company," said Gelsinger. "We now want to take P2P resource sharing and generalize it for more applications across the corporate fabric."[23]

Intel is putting its might where its mouth is, working with P2P developers, contributing to the P2P Working Group, and, most importantly, incorporating a new feature into the new Pentium 4 called "NetBurst," which the company said will improve P2P computing tasks. Intel also developed the Peer-to-Peer Trusted Library, a freely available security toolset.

Microsoft

Microsoft is developing Farsite, a distributed file system, and recently announced HailStorm, a platform based on XML Web services that is part of their .NET effort. HailStorm offers an authentication service based on the company's existing Microsoft Passport service. Groove Networks is among the early users of HailStorm, and received a $54 million investment in a strategic partnership with Microsoft.

---

[20] op. cit www.crn.com.
[21] http://www.futurecompany.co.za/2001/01/19/covstoryd.htm
[22] http://news.cnet.com/news/0-1003-200-2603611.html
[23] http://www.websphereadvisor.com/Articles.nsf/AID/ASMUP01

One of Microsoft's key understandings of the P2P space involves two related concepts: user authentication and Presence Management (PM). In the new world of distributed P2P and Web services, Microsoft wants to be the authority on users' identities and personal details. By making Passport the central repository for all user details, and incorporating Passport as the authentication authority for distributed services, Microsoft can benefit from others' P2P efforts as well as their own.

Presence Management means letting users control who knows they are online. In Instant Messaging (IM) applications, users can designate buddies who can be notified when they are online and available to chat. The use of IM applications on wireless phones extends PM into a very powerful force for connecting people. Of course, users will want to be very careful about who knows they're online, via phone or computer. PM becomes even more of a privacy issue if wireless presence management gets integrated with global positioning, revealing not only online status, but location as well.

In the always-connected world we're heading for, Microsoft's strategy to be the information center of the digital universe makes tremendous sense. There's more on IM and PM in Part 2 of this white paper.

Sun

With their "The Network Is The Computer™" philosophy, Sun obviously is interested in the P2P space and is one of its largest proponents. The company devotes significant resources to its JXTA Open Source P2P development effort as well as their hive computing efforts such as the Sun Grid Engine. Sun even is backing, along with IBM, the Globus Project™, an Open Source hive computing platform effort. In fact, Sun's eagerness to make JXTA and Grid code freely available could be a major threat to hive computing companies that use proprietary distributed computing solutions.

Sun also purchased P2P search engine InfraSearch early in 2001 and rolled it into JXTA. It appears that Sun will make P2P technology a prominent part of their Sun ONE environment, which aims to compete with Microsoft's .NET initiative.

Sun is also making efforts to get JXTA technology incorporated into Independent Software Vendors' (ISVs) products. Improv Technologies, iMulet, Oculus Technologies, and Texar Corporation have all integrated JXTA technology into their P2P applications.

There's more on JXTA and Sun's Grid computing efforts in Part 2 of this white paper.

IBM

In addition to a $4 billion investment in 50 worldwide centers to do hive computing, IBM is sponsoring a research project at the University of California at Berkeley called OceanStore. It's a distributed file system that involves storing portions of a file on many different servers in a method similar to RAID (Redundant Array of Inexpensive Disks).

IBM has also created SashJab, an Instant Messaging client based on the Open Source Jabber client. SashJab is part of the Sash Weblications for Windows effort that aims to simplify application development.

IBM is a backer of the Globus Project™, an Open Source hive computing platform effort. Recently IBM partnered with Entropia, a hive computing startup, and announced a project called the

TeraGrid. "The TeraGrid will connect a heterogeneous collection of existing high-performance computers at the National Center for Supercomputing Applications (NSCA), the San Diego Supercomputing Center (SDSC), Argonne National Laboratory, and the California Institute of Technology, creating a giant virtual computer accessible from any point on the Grid," said Irving Wladawsky-Berger, Vice President, Technology and Strategy, IBM Server Group. "The Grid will also include high-resolution visualization environments, and data storage integrated into an information infrastructure called the 'TeraGrid.' It will store more than 600 terabytes of data, the equivalent of 146 million full-length novels, be connected via a superfast, 40 gigabits per second network, and amount to a single computing system capable of 13.6 trillion calculations per second."[24]

IBM's Vision

"The TeraGrid will store more than 600 terabytes of data, the equivalent of 146 million full-length novels, be connected via a superfast, 40 gigabits per second network, and amount to a single computing system capable of 13.6 trillion calculations per second."

Irving Wladawsky-Berger, Vice President, Technology and Strategy, IBM Server Group

IBM is a partner with Microsoft, Ariba and others in the Universal Description, Discovery, and Integration (UDDI) business registry. The partners maintain a repository that contains information on companies and the way they want to do eCommerce. IBM hopes UDDI will become the business yellow pages of eCommerce, and permit companies to initiate P2P connections with one another.

There's more on Jabber in Part 2 of this white paper.

Hewlett-Packard

HP is committed to P2P development as part of its effort to "turn the information super-highway into a business super-highway." The company aims to create an autonomous but extensible platform focused on providing dynamic application support, service provisioning, information transfer and synchronization, and workflow management. To this end, HP is creating Internet Business Servers[25], which the company claims turn ordinary Web servers into communicating and collaborating servers. HP claims to be the first to feature automated, agent-based process management, enabling decentralized workflow management.

What Is the Developer Community Saying?

The recent interest in P2P began as a grassroots response to the problem of sharing files on consumers' computers. Before the big industry players knew it, Napster hit 30 million users, and spawned dozens of imitators. The P2P movement thus has much in common with the Open Source movement in that it is individuals providing alternatives to the accepted way of doing things. Consequently, many grassroots developers in the P2P movement resent what they see as an attempt of the huge computer companies to dominate the nascent field.

This attitude became quite clear to Intel in October 2000 when it convened the first meeting of the P2P Working Group (P2Pwg) it started with six other industry heavyweights. The chipmaker was obviously stunned when it met with vocal opposition to its proposed plans for organizing the working group with a steering committee of seven powerful companies and fees ranging from

[24] http://www-1.ibm.com/servers/events/grid.html
[25] http://www.hpl.hp.com/techreports/2001/HPL-2001-14.pdf

$500 to $25,000 for participation. Many were critical of Intel, who had developed no P2P code at that point and had nothing concrete to contribute, for "leading the parade." "Why are we doing research and development for Intel?" asked one developer.[26] Intel subsequently released the Peer-to-Peer Trusted Library in February.

Even at a follow-up conference organized by book publisher and Open Source standard-bearer O'Reilly, the feeling that the big companies were trying to co-opt the P2P movement was evident. Then the second P2Pwg meeting resulted in a coup that denied Intel chairmanship of the steering committee and an attempt to transform the working group along the lines of the Internet Engineering Task Force (IETF), which oversees most of the standards that make the Internet work.

It is obvious that the question of P2P standards and huge corporations' role in them will foster a political battle that will continue for some time. Whether the P2Pwg or some new group emerges as the arbiter depends a lot on how comfortable the grassroots developers feel with the process.

[26] http://www.zdnet.com/zdnn/stories/news/0,4586,2640166,00.html

Who is Likely to Benefit from P2P?

Some of the largest companies in the software industry are backing the P2P trend. Obviously they think there's something in it for them. In most cases, the something is increased sales of hardware or software systems. There are others who stand to benefit as well.

Internet Service Providers

There is a potential upside for this group: Tapping their users' idle resources. At least one Internet Service Provider (ISP), Juno, believes it has the right to use its users' processors for its own gain, and without their permission. Juno's user agreement[27] was amended this year to include language that permits Juno to download applications onto the user's computer, run them, and retrieve the results. All this is without any compensation to the user. The agreement even stipulates that the user is responsible for not only the cost of the call to return the data, but also for all electricity and other operating expenses.

Juno may be the first ISP to try this ploy, but they aren't likely to be the last. In fact, one is tempted to believe there may be an arms race as various ISPs large and small (Juno, which combined recently with NetZero to form United Online, is the 2nd largest US ISP) try to offer the largest hive computing pool. However, the Yankee Group believes it's not all about size. Security, application domain knowledge, and the ability to scale applications for use with millions of processors will be much more important for business use. Yankee points out that Juno has no particular expertise in advising their potential clients on how best to adapt their applications for hive computing usage.

Here is one area, perhaps one of the very few, in which free ISPs can have an advantage. Other no-cost ISPs such Address.com may decide to follow Juno's lead and require their free users to donate cycles. For-pay ISPs will need to compensate their users in some way for the use of their resources, and the market will set the price.

Data Centers

Web hosting company Exodus is collaborating on P2P Web site testing with partner United Devices. Their current model is for Exodus to basically resell United Devices' network. The model of many data centers, which is to own the environment but not the actual hosting machines, may preclude them from offering P2P services. However, Managed Service Providers (MSPs) and others who own the hosting hardware may be able to benefit, although it is likely to be only through partnership with the clients who are paying for the use of the computers. There will likely be many sticky points to negotiate, and these hosters may end up selling only off-peak capacity.

Companies With Big Computing Problems

Any business that requires huge amounts of computing power, disk space, or bandwidth is likely to benefit from P2P, even if it never uses a P2P application. The availability of P2P will put price pressure on equivalent services (supercomputers, network storage, content distribution networks), forcing these competitive providers into commodity pricing.

[27] http://help.juno.com/privacy/agreement.html

Computer Owners

Computer owners, especially consumers, will benefit from the new revenue stream involved in selling their excess capacity. In addition, P2P turns every computer user into a publisher, offering a potential worldwide content marketplace.

Who Would Be Threatened by P2P?

P2P doesn't offer an upside for everyone. The emergence of a new way of doing things is rarely a win-win for all involved. While some will benefit from new revenue streams or the availability of new resources, other will have their business models threatened.

ISPs

As we mentioned, ISPs can benefit from the P2P wave. But it is far more likely they will be threatened by it. ISPs and even Internet backbone providers have built their networks based on their clients using only a fraction of the bandwidth/hours available. The average consumer only spends an hour or two per day on their 56Kb line; broadband users average more hours. But no matter how many elapsed hours you use, you still don't usually saturate your connection; you stop to read the Web pages. P2P applications, on the other hand, could use a significantly higher percentage of the available bandwidth, all day long. With the typical ISP oversubscribing their networks by 40-to-1, something's got to give. It's likely to be flat-rate pricing. "The ISPs are in a tough situation," said Clay Shirky. "All of their pricing models were based on everyone at the edge of the network tolerating second-class citizenship. They're totally unprepared for the idea that end users will distribute content to each other."[28]

If consumers can download MP3 files in the background, there's going to be precious little unused bandwidth while they're using their computers. And if users sell their computer resources when they're not using them, their computers could stay active 24/7. ISPs will need to deal with this, probably by revising their pricing and usage agreements. Already many broadband providers prohibit users from hosting servers. Look to see some tough enforcement actions against violators who run Napster's successors in the near future.

ASPs' Role

"This is the start of the next big wave in Internet computing. ASPs need to be watching it now, adjusting to this new mindset, and starting to build it into their offerings."

Phil Wainewright, founder of ASP industry site ASPNews

ASPs

Clay Shirky said in an article from April 2000, "Just as the ASP space is taking off, Napster's success represents the revenge of the PC. By removing the need to upload data (the single biggest bottleneck to using the ASP model for everything), the content-at-the-edges model points the way to a re-invention of the desktop as the center of a user's data, only this time the user will no longer need physical access to the PC itself . . . Napster's rise shows us that the versatility of the PC as a hardware platform will give the millions of desktop machines currently in use a new lease on life. This in turn means that the ASP revolution will be not be as swift nor will the death of the PC be as total as the technology press would have us believe."[29]

[28] http://www.zdnet.com/intweek/stories/news/0,4164,2689521,00.html
[29] http://www.business2.com/magazine/2000/04/20762.htm

However, hive-computing startup Parabon believes there's a place for ASPs. It has developed an ASP development kit for their Frontier P2P Platform to enable ASPs to offer distributed computing services to their users. Phil Wainewright, founder of ASP industry site ASPNews, believes there's both opportunity and threat in P2P: "This is the start of the next big wave in Internet computing, probably the most important next step since the advent of the web-based application server. It's new and it will take maybe five or more years to start to become mature, but ASPs need to be watching it now, adjusting to this new mindset, and starting to build it into their offerings."[30] Wainewright argues that ASPs are really about managing applications, no matter where they may reside. He asserts that the P2P movement provides opportunity for ASPs to manage applications not only on PCs, but also on the variety of wireless devices that are emerging.

There's more on Parabon in Part 2 of this white paper.

Content Providers and Copyright Holders

Napster certainly demonstrated P2P's threat to content owners. If it becomes ridiculously easy to share digital media, a ridiculous amount of digital media will be shared, mostly illegally. Yet the threat involves more than just the recording and motion picture industries. Anyone who wants to trade content for money (Lexis/Nexis, Dow Jones) could get Napstered, and this will put pressure on existing pricing schemes. The P2P phenomenon has already spurred the development of various encryption and digital watermarking schemes as well as efforts to defeat these schemes.

Unfortunately, the best that content owners can hope for is to reduce the amount of illegal trading, since there's a huge pool of talented crackers poised to reverse engineer any protection system. In fact, noted security expert Bruce Schneier said in a November, 2001 speech[31], "Bits are bits. By their nature, they're copyable. You might as well ask water not to be wet." Schneier feels that every Digital Rights Management scheme will eventually be cracked. Being able to access supercomputer-class computing power through P2P will only make the cracking that much more likely.

Supercomputer Makers

A tightly integrated, multiprocessor supercomputer is always going to be more efficient than a hive-computing network. Many computing problems, such as real-time virtual reality simulations, are insoluble without high-speed data transfer among cooperating parts – processor, memory, disk, and peripherals. But for certain types of problems, P2P will really make inroads into the supercomputer and workstation cluster market.

At the very least, P2P will bring supercomputer power within reach of businesses that never could have afforded it before. Rather than leasing time on a supercomputer, these businesses can engage the services of United Devices, Parabon, Data Synapse, or any of the other hive computing vendors. Even large companies such as Intel, which has not bought a mainframe in 10 years, can replace supercomputer capacity with internal idle workstation resources. There's more on Data Synapse in Part 2 of this white paper.

[30] http://www.aspnews.com/analysis/analyst_cols/article/0,2350,4431_568701,00.html
[31] http://www.faircopyright.org/press/schneier.html

Net Appliance Makers

Net Appliance vendors have sought for several years to replace full-featured personal computers with low-cost, often hard-diskless terminals that rely on more-powerful servers to deliver applications and information. Many companies have lost money trying to realize this dream. P2P could put the final nail in the Net Appliance coffin by placing emphasis on having computing resources to share on P2P networks. If your Net Appliance doesn't have local storage, significant memory or computer power, or broadband Internet access, you'll be missing out on a significant part of the new Internet, not to mention a possible revenue stream.

On the other hand, P2P can provide content and services to limited computers, enhancing their usability. Thus more for-pay services will be available to users of Net Appliances. This ability to service limited computing devices will also hasten the development of Net Devices, from Net-connected refrigerators to pocket wireless handsets. These devices will be able to consume the content from P2P networks even if they can't share resources.

What Are Some of the Problems Facing P2P?

The future of P2P is by no means certain. There are huge questions facing the technology, not the least of which involves its appropriateness for use in anything but small workgroups. We take a brief look at some of the challenges in this section.

Reliability

P2P applications must be reliable despite the unreliability of the connections and the intermittent presence of the peers that provide services. This is no big deal for users of the large music sharing networks like Morpheus, KaZaA, and Gnutella because of the tremendous redundancy in those networks. If the PC from which you were downloading the latest Britney Spears single crashes, is shut down, or loses connectivity, you can find the song on another peer. All you've lost is some download time. But if a node crashes while computing a critical part of an automobile crash simulation, for example, the network needs to handle the interruption gracefully.

Reliable Networks

SETI@Home, Parabon and other hive computing companies deal with reliability problems by sending the same work packet to multiple computers, thus leveraging the breadth of their networks.

P2P applications typically ensure reliability through redundancy. SETI@Home, Parabon and other hive computing companies deal with reliability problems by sending the same work packet to multiple computers, thus leveraging the breadth of their networks. Having multiple computers working on the problem increases the likelihood that a result will be returned. Bandwidth sharing vendors AllCast, vTrails, ChainCast and other P2P content delivery networks have developed software that allows another peer to almost immediately pick up a dropped stream so that end users rarely will see interruptions in service. Some file sharing schemes enable users to download the packets that make up a file from several different computers, maximizing bandwidth and increasing reliability.

There's more on SETI, Parabon, AllCast, vTrails, and ChainCast in Part 2 of this white paper.

Scalability

Old-timer network administrators warn that peer-to-peer networking has been tried before. Systems like LANtastic and AppleTalk attempted to create serverless networks years ago. This network model didn't survive for several reasons, chief among them: high network traffic and lack of central control. When every peer is constantly casting about the network for other peers, files, or services, it becomes very easy to saturate the network's bandwidth with messages.

A real world contemporary example of this occurred in the days following a judge's ruling against Napster in October 2000. Afraid that the free music was going away, thousands of Gnutella users clogged the network downloading their favorites, making the network virtually unusable for two days. Gnutella, which uses the Pure P2P model, differs from Napster, which uses the Brokered P2P model, in that its search method creates cascading messages throughout the network.

In fact Napster developer Jordan Ritter thinks the problem of scalability in Pure P2P networks like Gnutella is hopeless. "Scaling Gnutella will require more than just better resource management

tools – in its current incarnation Gnutella is mathematically and technologically unable to scale to a network of any reasonably large size."[32] Ritter calculated that a single search for a music file on a network of 4,000 hosts would result in 800MB of traffic. The Gnutella network was paralyzed by a mere 10 queries per second. The exponential proliferation of messages throughout a Pure P2P network will thus prevent it from scaling much beyond a few thousand users, given the bandwidth available today.

This tendency of P2P networks to proliferate messages casts doubt on the scalability of any message-based P2P search facility such as JXTA Search, Pandango, and OpenCola Folders. There's more on each in Part 2 of this white paper.

It seems obvious, then, that the Brokered or Hybrid P2P model is the only hope for networks that want to scale. That means there needs to be central servers, and perhaps replicators if you want to support many millions of users. More-intelligent, hierarchical searches such as that employed by OpenCola Folders may also help.

Requirement for Robust Peers

P2P seems diametrically opposed to the Net Appliance trend, which has sought to slim down computers. To participate in P2P networks, people will want powerful computers, with resources to share, to avoid paying subscription fees for services. It's likely that a model will emerge whereby you get compensated based on the amount of CPU cycles, disk space, or bandwidth you can contribute.

This is one reason why Intel has sought to lead the P2P charge. With demand for PCs declining this year, Intel's banking on consumers and businesses wanting more power so they can share. Substantial computers, no matter how cheap they get, will still command a premium over Net Appliances or more-average computers, and that means more money for Intel and other computer component makers.

P2P also seems at odds with aspects of the Application Service Provider (ASP) movement, with its emphasis on servicing applications from a central server. When information and applications can be provided by a peer network, having a centralized ASP seems like a bottleneck and a single point of failure. P2P networks may put price pressure on some ASPs as well, assuming they achieve sufficient quality of service and reliability to compete.

Increased Energy Costs

Even though computer makers have made great strides in decreasing the power consumption of personal computers, if the computer is left on 24 hours a day to participate in a P2P application, it is going to use more power.

A typical consumer PC and peripherals consume between 200 and 400 watts per hour of operation. This means that if SETI@Home's 3.3 million users were to all participate at once and leave their computers on constantly, assuming neither the monitor nor the PC are in sleep mode, they'd consume more than half a gigawatt of power an hour, or 13 gigawatt hours per day. That's enough power to run more than half a million non-P2P-PC-using houses. Turning off the monitor or using an Energy Star monitor would cut these numbers about in half.

[32] http://www.darkridge.com/~jpr5/doc/gnutella.html

Contrast that with the typical data center, which uses between 10 and 30-megawatt hours per day, although usage in large centers can peak at more than 700 megawatts per second.[33] While the increase in P2P usage may result in an equivalent decrease in the amount of data center acreage, a key difference will be that the owner of the PC pays the electric bill directly, rather than part of a hosting fee.

Possible Legal Liability

In their February, 2001 decision on the Napster case, the US 9th Circuit Court of Appeals held that the so-called Betamax Supreme Court decision was not available as a defense. That decision found that copyright holders could not outlaw a device (in this case, the Betamax VCR) that, although capable of infringing uses, was also capable of substantial non-infringing uses. In this case, the Court found that Sony's knowledge that the machine was used for copyright infringement was irrelevant, since copyright infringement was not the primary intended use. The 9th Circuit said that Napster was liable precisely because it knew copyright infringement was taking place, and it did nothing to stop it. This argument, if applied to handguns, argues noted law expert Lawrence Lessig, would hold gun makers responsible for the murders committed with their products[34].

This decision could cast a chill over any P2P network that enables file sharing. Even if you set up the network for a legitimate business purpose – to enable workgroup collaboration, for example – if people start trading copyrighted information over it, you can become liable.

A further legal liability could result from the US ratification of the international Cybercrimes Convention[35] treaty. This treaty enables any signatory nation to enforce their cybercrime laws against people in any other signatory nation. This means if it is illegal to transmit a particular document in Bulgaria, a peer receiving that document in the US could be legally and criminally liable. The US would cooperate with any Bulgarian warrant to search and seize assets of US companies in order to investigate the case.

Increased Liability

Legal liability could result from the US ratification of the international Cybercrimes Convention treaty. This means if it is illegal to transmit a particular document in Bulgaria, a peer receiving that document in the US could be legally and criminally liable. The treaty also makes corporations liable for any cybercrimes their employees commit.

This treaty could dramatically curtail file sharing P2P networks, and could affect other P2P networks as well. For example, let's say a user in Malaysia wants to watch a Victoria's Secret runway video while at work. Using a P2P bandwidth sharing service, the video stream might pass through a user's computer in Brunei, where the video is illegal. Imagine the mess if authorities wish to prosecute one or both of the users. And imagine the liability of the business whose computers the first user was using. The treaty makes corporations liable for any cybercrimes their employees commit.

[33] http://www.idg.net/go.cgi?id=432715
[34] http://www.theregister.co.uk/content/4/22802.html
[35] http://conventions.coe.int/Treaty/EN/WhatYouWant.asp?NT=185

Physical and Logical Security

There are many aspects to the security puzzle for P2P applications. Security of data transmission, protection of host computer from attacks or data compromise, protection of the application or content from compromise by the peer, prevention of eavesdropping, and lack of peer privacy are just some of them. Centralized servers can be locked inside hardened sites and employ the latest software security practices. Yet anybody can walk up to a desktop computer and observe P2P processes, disrupt them, or hack into the computer and compromise it.

P2P networks may also be more vulnerable to crackers who take over one peer and use it to launch further attacks. P2P technology is actually at the core of a favorite cracker trick in which hundreds of computers are compromised and used to launch denial of service attacks against popular Web sites. The recent Code Red and Nimda worms exploited vulnerabilities on unpatched Windows computers to turn them into zombies. The zombies were then given instructions to bombard a particular site with bogus Web requests, causing a Denial of Service (DoS) outage. Many security professionals believe that these types of attacks will become the cybercriminals' preferred method of harassment in the future.

Lack of Corporate Network Control

How the question of control is handled will determine whether businesses will use P2P technology widely or just for niche applications. Network managers almost universally subscribe to the notion that what is not expressly permitted is forbidden. This is the safest way to approach the daunting task of managing a network and the computers attached to it. The key tool in managers' arsenals to enforce their rules is the centralized authentication and security service. Most P2P applications do not feature centralized user control, and this rightly scares network managers. The Groove approach of offering corporate users central points of control is likely to be the solution to this problem, at least in the short term. There's more on Groove in Part 2 of this white paper.

User Resistance

If enterprises treat the resources on their employees' desks as company property to do with as company leaders see fit, they might find significant employee resistance to P2P schemes. Workers regard their workstations as their *personal* computers. A workforce not on board with the goals of a P2P project can sabotage it, for example, by constantly pressing a key to prevent the computer from starting the P2P screensaver. Noted tech writer Steve Steinberg puts the problem this way, in the context of hive computing:

> The social barriers facing hive computing are the same that have long bedeviled socialism. Hive computing, after all, asks people to give up ownership of their property for the greater good. A user might come back from a coffee break to find his or her computer running someone else's program. That sort of infidelity can drive people nuts, and early attempts at hive computing were often sabotaged by users who periodically tapped their keyboards so their computers would always appear to be hard at work.[36]

Using employees' equipment without their buy-in is one thing. Using it without the users' knowledge or informed consent is another. Unfortunately, this issue is likely to become a problem, as we have seen from the previous discussion about ISP Juno.

It is wise to remember that the "P" in P2P also means People.

[36] http://www.wired.com/wired/archive/3.11/geek.html

Lack of P2P Standards

Right now, the infant P2P industry has not developed standards that will make various P2P efforts interoperate. Sun has made a play to create a standard with its JXTA stack, and there will be many more attempts. Some P2P leaders suggest that interoperability is more important than standards, and JXTA does make it easier for P2P applications to interoperate. Nonetheless, in an environment as fast evolving as P2P, there are likely to be winners and losers. Thus it's possible to back the wrong horse and end up with an isolated implementation, and nothing could be worse than an isolated peer-to-peer application.

Conversion of Volunteer Business Model

Right now, especially in the hive-computing segment, P2P depends a lot on volunteer or poorly paid labor. Users donate computer cycles or bandwidth or disk space for little or nothing in return. It is not at all certain that companies dependent on volunteer efforts will be able to make the transition to for-pay service. However, hive computing provider Parabon has made this transition, finding a niche in the bioinformatics industry with P2P versions of popular gene analyzing program, according to Jim Gannon[37], Parabon Chief Technology Officer. There's more on Parabon in Part 2 of this white paper.

The Internet has taught us that, while everything devolves to free, advertising and other indirect revenue sources may not be enough to build a business upon. Once users see the benefits hive computing companies are reaping from their resources, they are likely to demand more payment, and this could ruin many companies' business models.

[37] Interview with Jim Gannon, 11/14/01.

What is the Future of Business Computing Networks?

Despite the hype and the broad range of current P2P offerings, P2P technology is in the earliest stages of development. After all, it's only been over the last two years that broadband connections have achieved the kind of penetration necessary for such a scheme to work. So if this is the beginning, what will the future look like?

Once the initial obstacles are surmounted, the networks will evolve beyond the current artificial division of labor between file sharing and idle resource maximization to provide a rich, collaborative, integrated marketplace of business capabilities both within the enterprise and externally. The current four P2P prongs will blur and merge to produce networks that provide ubiquitous computing power, content delivery, bandwidth, and collaboration support.

Domain Expertise

"Because enterprises need help in understanding how to utilize this new low-cost computing infrastructure, P2P ventures must provide market-specific consulting or application knowledge to their prospective customer base."

Yankee Group

Justin Chapweske[38], one of the creators of OpenCola's Swarmcast content distribution networking software, believes bandwidth and file sharing solutions are the most likely to become integrated into the Internet. "The more proprietary the solution, the more vulnerable the multicast business will be in future. Certainly we'll see specialized Content Distribution Networks in niches like the current music sharing sites, but eventually, multicasting or P2P bandwidth sharing will migrate so that it's just a property of the network. Then the businesses that have been built up around solving this problem will be at risk."

IBM, with its multi-billion dollar investment in hive computing, is betting it can turn distributed computing into a network service, albeit a proprietary one. But with distributed computing services already being offered free in Sun's Grid Engine software and in open source efforts like Globus (backed by Sun and IBM), the Global Grid Forum and Beowulf, it's possible that Sun's "The Network Is The Computer" slogan will in fact be true in the future. When anyone can easily set up computer processor sharing networks, today's hive computing companies will be hard pressed to provide a competitive edge.

In any event, Yankee Group believes these companies need to develop vertical market expertise. "Because enterprises need help in understanding how to utilize this new low-cost computing infrastructure, P2P ventures must provide market-specific consulting or application knowledge to their perspective customer base. IBM, being one of the largest supercomputing vendors, is in the best shape to continue this practice in the peer-to-peer domain."[39]

It's possible that only the P2P collaboration companies will avoid commoditization. Although features such as secure Instant Messaging could be absorbed into the network, more-complex applications such as online meetings, shared workspaces, and offline document collaboration could retain their value.

[38] Interview with Justin Chapweske, 11/16/01.
[39] op. cit. Yankee Group.

The Future Network

We are moving toward a future in which your job is not defined as a place you go, but as an activity you undertake regardless of your location. Peer-to-peer computing technology, along with wireless networking standards such as Bluetooth and 802.11b, will help enable you to access content and services wherever you go.

New security and authentication schemes will allow the line between intranet and Internet to blur, enabling the growth of software services that are unbundled and uncoupled from specific applications. Sun's JXTA project, with its emphasis on application pipes, is a first step toward the creation of a computing environment in which applications can be assembled as needed, on the fly. Developers will benefit from this freedom as well as from the larger Sun ONE effort and Microsoft's .NET initiative. Other Web services initiatives will improve users' access to applications wherever they are.

As they evolve, business computing networks, combining P2P, Web services, and client/server technologies, will develop the following ideal characteristics:

Dynamic – Communications and services must be interactive, extensible, flexible, and easily reconfigured in real-time. Systems will link people, applications, computer systems and devices such as cell phones, cameras, and printers. Computing becomes a fabric that responds to users' needs rather than being located within specific devices such as PCs or supercomputers.

Real-time – In many cases, sub-second response time is required to deliver rich, timely, personalized, on-demand information. Store and forward methodologies may have a place in the P2P network, but most communication is immediate.

Collaborative – Both people and applications need to work together to deliver value. These collaborations must be secure, must support any number of participants, and must enable the discovery of new networked resources. This feature is perhaps the most important in setting workers free of place-based business interaction. We've already seen the beginning of this with call forwarding, cell phones, and unified messaging.

Structured – Network services must support the vocabulary used by the business. This may extend to support for local languages as well. Applications must support business activities in a manner appropriate to the industry. In addition, searching for information, services, or people, will be supported in a structured manner through standard metadata. A main stumbling block to efficient business use of the Internet is the lack of effective search techniques. This problem becomes many times more complex as billions of devices and untold numbers of services are added to the network.

Relevant – The other side of the search problem is relevance. Information and services must be timely and focused on the participants' current business needs. Users need filters they can use to personalize information delivery. Today's collaborative filtering schemes will be refined and extended throughout the network.

Service-based – Network capabilities can be obtained and configured at a moment's notice. New business applications can be assembled on the fly by integrating new capabilities into existing workflows, systems, devices and applications.

Cost effective – The network reduces the costs of solving business problems as well as of establishing and maintaining online business relationships. Services are provided by low cost specialists and are easily integrated into the core business of a company.

Client focused – Services and capabilities all can be easily personalized or otherwise adapted to the business purpose.

Offer reputation evaluation and management – In a world where you may never meet your trading or computing partners, trusted intermediaries will need to provide introductions and reputation assessment. Finding potential partners is one thing, and this need may be filled in part by efforts evolving out of the Universal Description, Discovery, and Integration (UDDI) effort. Developing trust with a partner, however, is more difficult. Reputation evaluation services from Dun & Bradstreet, Société Générale de Surveillance, and Open Ratings will lay the groundwork for reputation management frameworks.

There is much work to do before this future network is achieved. P2P computing will be a significant technology to help enable its creation.

What Should You Do About P2P?

P2P is in its infancy. Risk-averse organizations may wish to wait until it becomes a little more defined before incorporating P2P technology into their systems. On the other hand, organizations in fast moving markets, particularly wireless and B2C eCommerce markets, need to evaluate P2P technology. Regardless of your stance on risky, cutting edge technology, remember Gartner's advice: "Even those enterprises that choose not to engage in P2P at this stage must recognize the technology's longer-term impact and factor it into their strategic planning."[40]

Businesses with business problems that have huge computational needs should evaluate hive computing as a possible cost-effective solution. Businesses of all types need to be aware of the various B2C P2P services on the Internet that could clog their bandwidth or potentially compromise their security. Any business with geographically dispersed workgroups can benefit from P2P secure collaboration services. Businesses with large amounts of information dispersed throughout the organization should take a look at P2P file sharing and distributed content management services. Businesses with massive multimedia content publishing challenges should consider using one of the P2P content distribution networks.

Continued in Part 2

This white paper continues in *Part 2 – How Are Businesses Using P2P?* In the second part, we examine real world examples of P2P applications used by businesses today and make recommendations about P2P applications use. Included are P2P vendors providing applications in the following categories:

- Collaboration
- Content Distribution
- Bandwidth Sharing
- Hive Computing
- Testing Web Applications
- Remote Access
- P2P eCommerce
- Collaborative Filtering and Searching

Also included are a consideration of future P2P applications for wireless devices and the evolving P2P platform, JXTA. See the full table of contents in Appendix B or at www.CTOMentor.com/p2p/part2.htm. Part 2 is available for a nominal charge.

Appendix A

Peer-to-Peer Resources

Web Sites

Besides Stratvantage's P2P4B2B directory (http://www.stratvantage.com/directories/p2pcos.htm) There are several very good information portals on the P2P computing movement, including:

- OpenP2P www.openp2p.com/
- P2PAnalyst www.p2panalyst.com/
- P2Ptracker www.p2ptracker.com/
- Peertal www.peertal.com/
- PeerToPeerCentral www.peertopeercentral.com/

These sites are a tremendous resource for anyone wanting to keep track of the rapidly changing P2P marketplace. The following additional resources also can help:

- Global Grid Forum www.gridforum.org/
- New Productivity Initiative www.newproductivity.org/
- P2P Terminology Glossary www.exocortex.org/p2p/terminology.html
- Peer-to-Peer Working Group www.peer-to-peerwg.org/
- Scientific American www.sciam.com/2000/1100issue/1100cyber.html
- TechRepublic www.techrepublic.com/article.jhtml?id=r00520001227gcn01.htm

Books

2001 P2P Networking Overview: The Emergent P2P Platform of Presence, Identity, and Edge Resources (O'Reilly Research) Clay Shirky, Kelly Truelov, Rael Dornfest, Lucas Gonze, Dale Dougherty, O'Reilly & Associates; ISBN: 0596001851

Building a P2P Supercomputer for the Enterprise, Harold Cabrera (to be published February 2002), Syngress Media Inc; ISBN: 1928994733

Discovering P2P Michael Miller, Sybex; ISBN: 0782140181

Early Adopter JXTA, Wrox (Editor), Wrox Author Team, Wrox Press Inc; ISBN: 1861006357

High Performance Cluster Computing: Architectures and Systems, Vol. 1 Rajkumar Buyya (Editor), Prentice Hall PTR; ISBN: 0130137847

High Performance Cluster Computing: Programming and Applications, Volume 2, Rajkumar Buyya (Editor), Prentice Hall PTR; ISBN: 0130137855

How to Build a Beowulf (Scientific and Engineering Computation) Thomas L. Sterling, John Salmon, Donald J. Becker, Savarese, Daniel F. Savarese, MIT Press; ISBN: 026269218X

JXTA: Java P2P Programming Daniel Brookshier, Nitin Borwankar, Darren Govoni, Navaneeth Krishnan, Juan Carlos Soto, Sams; ISBN: 0672323664

P2P: Getting Down to Business: Can Peer-to-Peer Processes Produce Profits? ManyWorlds (Downloadable), ManyWorlds, Inc.; ASIN: B00005R4KU

P2P: How Peer-to-Peer Technology Is Revolutionizing the Way We Do Business Hassan M. Fattah, Hassan M. Fattan, (to be published January 2002), Dearborn Trade; ISBN: 0793148782

Peer-to-Peer: Harnessing the Power of Disruptive Technologies Andy Oram (Editor), O'Reilly & Associates; ISBN: 059600110X

Peer-To-Peer Application Development: Cracking the Code Dreamtech (Editor), Hungry Minds, Inc; ISBN: 0764549049

Appendix B

Table of Contents for Part 2

The second part of this white paper, *Peer-to-Peer Computing and Business Networks: More Than Meets the Ear, Part 2 – How Are Businesses Using P2P?* is available for a nominal fee at http://www.CTOMentor.com/p2p/part2.htm. The following is the table of contents of the second part.