EXECUTIVE BRIEFING SERIES

WHITE PAPER

Volume 5

The Buzz About Hive Computing:

Putting Peer-to-Peer Computing to Work

Geneer®

Building Software. Building Business.SM

VOLUME 5

The Buzz About Hive Computing:

Putting Peer-to-Peer Computing to Work

Contents:

INTRODUCTION

"Fifty years ago, Thomas Watson estimated there was a worldwide need for maybe five computers. We now know that that number was wrong. He overestimated by four."

Author and peer-to-peer (P2P) pundit Clay Shirky, a partner with venture capital firm The Accelerator Group, recently uttered this provocative statement at O'Reilly's Peer-to-Peer conference. He was referring to former IBM Chairman Thomas Watson and his famous 1943 quote, which is often held up as an example of gross misunderstanding of market potential. Watson was looking at the new device from an old frame of reference: office calculators. Shirky is looking at computing today from a new frame of reference: hive computing.

What Shirky was talking about is an innovation, P2P, that promises to make Sun Microsystems' mantra, "The Network is the Computer," come true. P2P is a technology that involves linking together the resources of multiple computers to solve a computing, collaboration, or communication problem. The most well known example of P2P computing is the music file sharing service, Napster, but there are dozens of other uses of the technology. One use is to combine the idle computing power of Internet connected computers to solve big computational problems.

Just as the many cells of a honeycomb make up a hive, the millions of Internet-connected computers make up a huge potential supercomputer. If these tremendous resources, an estimated 10 billion MHz of processing power and 10,000 terabytes of storage[1], could act as one, the network would really be the computer. While Shirky was being a bit facetious in turning Watson's quote around, the reality is, there's real work being done today using hive computing.

Hive computing has gone by many names over the years, distributed computing being the most popular. In general, the concept involves bringing more than one computer together to solve a problem. There have been numerous attempts to do this in the context of a single system. The current fastest supercomputers are, after all, federations of thousands of independent processors.

Where hive computing gets difficult is in coordinating the effort of standalone computers that are physically distant and connected via a network. Making these federations work like a coherent whole on a single task has proved to be extremely challenging. For example, if the solution to a problem requires processors from three different computers, memory from five others and hard disk space from another dozen, the coordination effort and the communications delays caused by the network can slow the process considerably.

For this reason, the current best approach to hive computing involves splitting up a very large problem into discrete parts that can be independently solved by a loose federation of computers. Thus, the best problems for hive computing schemes are those that require a low degree of communication and coordination to solve. This "divide and conquer" method, however, implies some sort of central control. Therefore, is it really a peer-to-peer solution?

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 would be true. However, the applications being developed today are intended to solve problems, not be ideologically pure. This means several of the most popular P2P applications do indeed use central servers to facilitate peer communication in an architecture called Brokered P2P. The Napster music file sharing service is the best example of this hybrid approach. In fact, the use of a central server to house the music index is key to Napster's continuing legal troubles. Although all transfers happen between peers, searches are performed on the Napster index, which the company maintains in a central location.

Peer-to-Peer Models

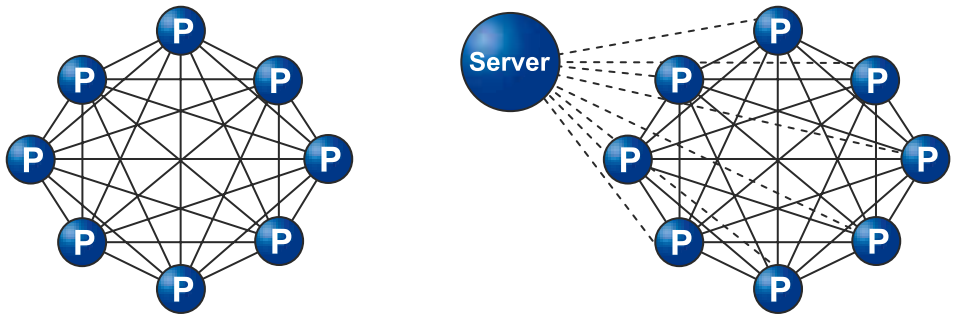

Pure P2P

Hybrid/Brokered P2P
The dotted lines between peers and the server reflect online peers connecting to the server to store information deltas (changes) for later access by offline peers.

Source: IDC #23576, December 2000

Collaboration P2P vendor Groove Networks[2] also employs a central server for some types of file distribution. Among other 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.

So what does P2P mean? Once again, Clay Shirky[3]:

> 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 DNS system and have significant or total autonomy from central servers.

Basically, P2P began as a grass roots movement to empower ordinary people to be content publishers. Without these easy-to-use applications, Web publishing is a frustrating experience, involving obtaining a permanent network address and domain name, and installing and maintaining a Web server, security system, and content publishing application.

An obvious place for P2P content serving is the corporate intranet. Intranets have been somewhat successful, but, it can be argued, have been hampered by the necessity for central control.

P2P has grown well beyond its content-oriented roots, however, to embrace a whole range of network applications.

The Four-Pronged P2P Effort

There are four major prongs in the P2P effort today. All four involve putting unused computing resources to use, at least to some extent. 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, and other consumer-to-consumer (C2C) technologies do. These services enable people to serve content off their local hard drives and share files with anonymous 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 on the consumer scale. In the consumer arena that means sharing music and other entertainment files, but this is only one possible application. If the resulting technology makes its way into business, one promise of distributed computing will be fulfilled, and centralizing forces such as content aggregators, demand aggregators, and catalog aggregators may be threatened.

An obvious place for P2P content serving is the corporate intranet. Intranets have been somewhat successful, but, it can be argued, they have been hampered by the necessity for central control. People have jobs to do, and have little patience for

the bureaucratic procedures often required to publish their content on the intranet. The Napster 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 WorldStreet, which enables Wall Street traders to quickly and easily share information

Bandwidth Sharing

Video producers and other multimedia event producers can serve high bandwidth content by farming it out to hundreds of machines. When a user requests the content, it is served from a machine close by rather than from a huge central server, thus saving producers a lot of money. 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. Akamai is a non-P2P company that seeks to solve this problem by providing regional content distribution centers for clients.

While services such as Instant Messaging and online presentations are examples of P2P collaboration, the current crop of P2P collaboration vendors bundle many services to enhance interactivity.

Bandwidth sharing obviously involves sharing other computer resources, particularly processor power. Consider an application like InfraSearch, recently acquired by Sun. InfraSearch's OpenSearch aims to be a distributed search engine. Rather than relying on Web spiders – automated processes that visit sites and build a huge central database of URLs – InfraSearch will be a real time search engine. When a query is made, it is farmed out to participating users' computers that then search for relevant Web sites in real time. InfraSearch uses the participants' computing power as well as their bandwidth. Bandwidth sharing is really the key to this scheme, however, since it would not be feasible to perform such searches using hundreds of computers at a central location due to the bandwidth required.

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 and online presentations are examples of P2P collaboration, 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 Lotus Notes developer Ray Ozzie's brainchild, Groove Networks. Groove attempts to wrap 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.

Hive Computing

In hive computing, computers' unused processor power is harnessed to attack huge computing problems like virus modeling or Wall Street securities analysis. In hive computing, a huge problem requiring significant computing power is split up and assigned to cooperating computers over the Internet. 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.

This P2P prong may turn out to be 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]

Red Herring magazine picked hive computing as the number one trend for 2001. Noted industry consultant Cheryl Currid said, "Any IT manager who fails to look at peer-to-peer should be fired."

P2P – The Next Big Thing?

Red Herring magazine picked hive computing as the number one trend for 2001. Noted industry consultant Cheryl Currid said, "Any IT manager who fails to look at peer-to-peer should be fired." More than 200 P2P companies have sprung up, hoping to monetize the Net's idle capacity. A listing of many of these companies appears in Appendix A.

In the sections that follow, we analyze several current hive computing projects and project the growth and the usefulness of this technique for business. First, let's take a look at what is arguably the first significant public P2P application: SETI@Home.

IN THE BEGINNING – SETI@HOME

The SETI@Home project is popularly regarded as the granddaddy of hive computing, although the idea has been around for 30 years or more. In fact, researchers at Xerox's Palo Alto Research Center first explored distributed computing in the 1970s. Along the way, there have been some pretty high profile hive computing projects, including the rendering of the Toy Story movies via server farm and the cracking, via a hive network in 1997, of Netscape's 56-bit encryption key[6]. Computer chip giant Intel has used hive computing to design chips since 1990, saving half a billion dollars using spare computing resources within its organization.

SETI, however, deserves prominence as the largest scale voluntary hive-computing project to date. SETI (Search for Extraterrestrial Intelligence) farms out radio telescope data to volunteer computers that analyze it, looking for signs of extraterrestrial intelligence. It has 2.8 million users in 226 countries, has accumulated 585,000 years of CPU time, and has analyzed 45 terabytes of data in less than two years of operation. The consolidated volunteer computing power is 25 teraflops, which is twice the speed of IBM's $110 million ASCI White, the fastest supercomputer in the world. A staggering 740 million teraflops of computing power have been donated. SETI has accomplished all this with a staff of three to five people at a cost of $500,000.[7]

Hive Best Practices

There are several keys to the success of this effort and they form a good set of rules for other hive computing ventures. These are the requirements of successful hive computing:

- A problem that **needs huge computing** resources, but that is **not time sensitive** – If one of the participating SETI PCs goes down before finishing its analysis, the work can be parceled out to another PC. The overwhelming amount of data SETI has to analyze (and the small likelihood of success) makes such delays inconsequential.
- Server software that can **break a large problem** into **smaller pieces** – A monolithic and serial problem, where each step must be calculated before the next is begun, is not a good candidate for hive computing. Since the SETI data could be sliced up by time, it was ideal.

- **Minimal data communications requirements** – Many SETI computers are connected via modem. Any problem that requires a lot of data traffic would limit the number of machines that could participate. SETI sends a 250KB packet of data, called a work unit, to the worker computer. Once the worker computer starts analyzing the problem, it can be disconnected from the Internet until it is finished with the work unit. There are a variety[8] of hive computing schemes in development that take a different approach: linking discrete computing resources together into a single machine. In this scheme, a calculation could use networked resources such as memory and disk space, thus requiring a significant amount of coordination between hive members. This results in a huge volume of messages on the network as the various resources talk to one another and signal their availability. It seems unlikely that such an approach will work effectively on today's Internet.
- **Minimal data storage requirements** – In dealing with volunteer computers, SETI needed to make sure that it didn't take over the users' hard disks. Participation requires the installation of a small program that acts as a screen saver. The data for the work units is also quite small. Although one of the prongs of P2P involves using donated disk space to distribute content, using significant disk space in a hive computing setting requires a higher level of user commitment than just agreeing to run a screen saver.
- **Minimal inconvenience to the user** – Because SETI runs as a screen saver, it uses volunteered resources only when their owners are not using them. There are some hive computing schemes that attempt to use the spare computing power available while users perform tasks on the computer. For example, when you are using a word processing program, even if you type 100 words per minute, your computer is not even breathing hard; there are plenty of spare cycles. Because software to exploit this spare computer power is more likely to interfere with the owner's use of the computer, it is likely to be successful only in for-pay business models. The SETI application uses 16MB of memory when running, and thus runs typically only as a screen saver, although it can run as a background task.
- **Give incentives for greater participation** – Although SETI does not pay for use of their volunteers' computers, they do encourage them to maximize their utilization by posting a leader board on the SETI site. SETI breaks down users by school type, company size, clubs, and government agencies (all classifications are user-determined; currently The Ministry of Silly Walks is leading the Air Force in the Government division). Volunteers can band together in groups to try to place more highly in the standings, and a real competition has developed among regional or company-affiliated groups. Some users dig out old 386 and even older computers, slap a network card in them, and stick them on their broadband connection in hopes of getting an edge.

- **Ensure security** – Both the data being crunched and the worker computers must be secure. SETI claims their application is more secure than the standard Web browser and we don't doubt it. By hard coding the address of the computers that parcel out the work and receive the results, SETI ensures that their volunteers are working only for them. The application also uses the standard Web traffic port, 80, making it easier to site volunteer machines behind firewalls.
- **Ensure reliability** – With so much spare power at its disposal, SETI now sends the same work unit to more than one computer, increasing the chance it will get finished. SETI also resends a work unit if it isn't completed within set time limits. When high reliability is required, hive projects can bring multiple computers together to work on a single part of the problem.
- **Make participation voluntary** – With SETI, the only way to participate is to volunteer. With other hive computing projects, particularly those that might be undertaken within a business, this may not be the case. Noted tech writer Steve Steinberg puts the problem this way:

  > 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.[9]

  Any hive computing scheme that enrolls a user's computer without the support and assent of that user could be doomed to failure. Workers regard their workstations as their personal computers. It makes no difference if management considers the computers the property of the business to do with as it wishes. A workforce not on board with the goals of a hive project can sabotage it. It is wise to remember that the "P" in P2P also means People.

With these key learnings from the largest hive computing project to date in mind, many startups have sprung up to try not only to adapt this concept to business, but to make some money on it as well.

PAY FOR PROCESSING – DATASYNAPSE

DataSynapse is a startup targeting Wall Street's huge data crunching needs. Every day securities brokers need to make quick decisions about pricing, risk assessment and market conditions. There is plenty of data to support these decisions, but the typical brokerage lacks the computing muscle to deliver these complex analyses quickly. Some analyses can take up to 50 hours of processing time on a single computer. As a result, brokers often make critical decisions using limited data and seat-of-the-pants know how.

DataSynapse's WebProc™ software links computers together to solve financial problems, reducing exponentially the time it takes to do complex tasks. By harnessing the idle power of JP Morgan Chase's workstations, DataSynapse completed in a matter of minutes a test calculation that normally took eight hours overnight. According to the company, 12 high-performance financial applications are being optimized over two farms of 400 workstations and 250 PCs.

The application is minimally invasive: As soon as a user activates his or her computer, WebProc takes the work that computer was doing and moves to another idle client. WebProc is available for installation on a company's intranet for a monthly subscription fee. For companies with computing needs larger than their available internal resources, DataSynapse offers "power by the hour" from two external resource pools:

- **Extranet** – DataSynapse has signed up extranet partners including Web hosting firms and other providers with large reserves of idle, underutilized and/or dedicated computing power. These partners maintain physically secure hardware installations which DataSynapse claims the security comparable to or greater than the Virtual Private Network (VPN) solutions that are already in place at most firms.
- **Internet** – DataSynapse aggregates power from a 100 percent broadband subscriber network of home PCs. It incents these users using Flooz, an online currency redeemable at many online merchants. The first 10,000 customers will be paid $5 in Flooz for joining and credited with $1 for every referral they make. Once the business is established, customers will be paid based on their PC's processing power, available idle PC power and other variables. While that will vary, the company says it will probably be in the $5 per month range.

Businesses might be wary of using Internet resources for potentially sensitive applications due to concerns about security. DataSynapse asserts the built-in encryption and authentication in WebProc solves this problem. The company is also partnering with Zone Labs Inc., maker of the ZoneAlarm firewall, to offer WebProc users a personal firewall as well as digital encryption and digital signature technology. Nonetheless, it remains to be seen if businesses will accept these assurances and use Internet-based hive computing services.

DataSynapse is initially focusing on financial services companies but plans to branch out into the insurance and energy sectors. It is also partnering with the National Center for Supercomputing Applications to create a distributed search engine for the Web.

VOLUNTEER COMPUTING – PARABON

Parabon's motto is *Put your computer to good use—all the time*℠. Their hive computing application, Pioneer™, communicates with their Frontier™ server to receive work assignments. The Pioneer client application can work on the work task while disconnected from the Internet and then upload the results when connected. Parabon's users, called providers, currently donate their time to one of several cancer-related research efforts at the Colorectal Cancer Network, Cancer Treatment Research Foundation, National Cancer Institute, West Virginia University, or the University of Maryland. The company plans to offer computing resources to commercial clients in the first part of 2001.

Providers are incented to contribute resources by a feeling of making a difference in cancer research through Parabon's $100 daily sweepstakes drawing and monthly $1,000 grand prize drawing.

Parabon claims its application is safer than surfing the Net. They use Java, with its robust security model, to isolate their application from a provider's PC. The Pioneer application verifies the integrity of a task's code before it can be executed. In addition, a security mechanism prohibits tasks from making any network connections except to the Frontier server.

It is not at all certain that Parabon or other volunteer hive computing efforts targeted at non-profits will be able to make the transition to for-pay service. However, hive computing provider Distributed Science has made this transition, recently signing one of the first commercial hive computing accounts for its Process Tree™ network. Performance testing and monitoring services provider Envive Corporation will use Distributed Science's network of over 145,000 computers from 72,000 unique suppliers. Envive will be able to test applications using actual, real-world computer configurations.[10] Process Tree plans to offer users $100 to $1,000 per year for their computer time.[11]

Regardless of the company's business model viability, attendees at a recent CIO conference selected Parabon as the winner among 100 companies chosen for having new technologies that CIOs should be aware of.[12]

SHARPENING SEARCHES – PANDANGO

Pandango, a project currently in stealth mode at intellectual property developer i5 Digital, aims to be a distributed search engine, much like InfraSearch, the startup recently acquired by Sun. Pandango (Hawaiian for "wedding money dance") plans a different approach to the problem of sifting the valuable needles from the digital haystack of the Web. Unlike most search engines, which use traditional academic citation analysis (the more sites that link to a page, the higher the relevance), Pandango will determine relevance by examining a radiating network of "referrers."

A referrer downloads the Pandango application and joins the P2P network. When he or she does a keyword search, Pandango examines the Web histories and bookmarks of an initial network of 100 referrers, and then those 100 referrers' combined 10,000 referrers, and finally all of those people's referrers for a total of 1 million people's bookmarks. The reasoning behind this approach is similar to the standard search engine approach: The most valuable resources will have the most hits in the referrers' history and bookmark files. Over time, frequently selected referrers (the ones with the most in common with a user) would move to the top of the users' referrer heap, while others would drop off and new ones would join. Thus, the more a user uses Pandango, the more relevant his or her searches will become due to the assembly of a group of like-minded Web users.

This approach resembles another computing trend, collaborative filtering, which analyzes the behavior of large numbers of people to find affinities and makes recommendations based on the preferences of people with similar behavior. For example, collaborative filtering is used when Amazon makes book recommendations based on the buying trends of other book buyers.

Since Pandango is pre-release, it is not known how efficient the analysis of a million people's bookmarks will be. Nor is it clear what the resource impact of servicing searches will be on the participating computers. It is likely that a certain amount of processor power, memory, and disk access will be required on each peer to handle search requests. What will happen, for example, if you become a popular peer and hundreds of searches a day access your Web history?

When you consider that Pandango will be competing against search engines that return results in seconds, performance of the network will be critical to acceptance. On the other hand, many users might be willing to wait for more relevant results. Also unknown is whether users will be willing to open their browsing habits, even anonymously, for use in the Pandango network. Pandango will require that any user of the service also be part of the network, with their bookmarks available for use.

The biggest challenge Pandango faces, however, is developing a business model that makes sense. How many people will be willing to pay for a commodity (searching the Web) that they can get for free? The company will need to develop a compelling ROI model based on wasted searching time in order justify any charges. Such a model could undoubtedly be built, based on the pitiful performance of free search engines. But gaining visibility for the costs of inefficient Web use could be a challenge.

i5 Digital is investigating licensing the patent-pending Pandango technology to companies, portals and other search engine companies. There is precedent for this, as search engine Google licenses its engine to Yahoo, for example. However, there is likely to be considerable hive computing search engine competition. Sun's InfraSearch, which uses participating peers to actively canvas Web sites in real time, is a likely strong player.

Whichever technology wins out, you can expect some sort of distributed search engine to become dominant within the next 12 to 18 months. Searching may well become the most visible of the hive computing applications now in development.

THE BUSINESS OF HIVE COMPUTING

Right now, hive computing is somewhat of a solution looking for a problem. Various companies have found that they can get people to volunteer to help solve important problems like medical research or the search for extraterrestrial intelligence. They are just starting to find out if the compute-for-pay model will work with Internet users. The emerging hive computing industry is also just beginning to discover appropriate applications for hive computing. Forrester Research analyst Eric Scheirer sums up the problem this way: "I'd say peer-to-peer is an infrastructure, not a business model. But that doesn't mean there aren't a lot of interesting concepts around peer-to-peer applications."[13]

The Gartner Group has predicted 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.

Market Potential

Despite the uncertainties, there appears to be huge potential for hive computing. Industry analysts Currid and Company estimates the distributed computing market will reach $9.5 billion in services, $7 billion in hardware, and $13 billion in software by 2006.

The Gartner Group has predicted 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, the prediction hardly seems outrageous. Gartner further predicts hive computing applications will be limited to specific heavy-processing and low-security applications through 2005.[14]

There may be more companies using P2P technology today than you might think. Omni Consulting Group ran a study of 1,590 companies that are using some version of peer-to-peer technology. "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 analyst Frank Bernhard.

The Omni study found that although 38 percent of the companies were using P2P technology for information sharing, a miniscule number were using hive computing applications. Nonetheless, 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.[15]

Fear of losing control could present a barrier to business acceptance of hive computing. IT professionals are used to being able to control the computing environment from end to end. Yet this isn't possible when thousands of anonymous Internet 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."

Hive Computing Applications

In this paper we have just scratched the surface of the tremendous possibilities for hive computing. The applications of this technology are by no means limited to the examples we've given. Below are a few more examples of business uses of this emerging computing force.

Industry	Application Area
Insurance firms and Re-Insurers	• Assess insurance risk and probabilities • Investment management • Enterprise risk • Derivatives trading • Options trading • Variable annuities • A/L Management • C-3a Risk Based Capital analysis • Extreme Value Theory analysis • Worst case scenarios • Market assessment • Liquidity assessment • Credit risk • Monte Carlo simulation-based models
Securities Broker/Dealers	• Derivatives trading • Options trading • MBS • CDOs • Swap books • Multi-tiered securitization packaging • Hedge and risk books • VaR calculations • Market assessment • Liquidity assessment • Credit risk assessment • Monte Carlo simulation-based models • Demographic behavior patterns • Predictive modeling
Energy	• Distribution grid modeling • Energy trading • Risk management • Credit books • VaR calculations • Market assessment • Liquidity assessment • Credit risk assessment • Monte Carlo simulation-based models

Industry	Application Area
Asset Managers	• Trading • Risk management • Derivatives trading • Convertible arbitrage • Options trading • Hedged positions • Position analysis • VaR calculations • Extreme Value Theory analysis • Worst case scenarios • Market assessment • Liquidity assessment • Credit risk assessment • Monte Carlo simulation-based models
Financial ASPs	• Derivative and structured products • Swap books • Hedge and risk books • VaR calculations • Market assessment • Liquidity assessment • Credit Risk assessment • Monte Carlo simulation-based models
Manufacturing	• New factory site design and placement • CAD model rendering • Product simulations • Wind tunnel and other testing simulations • Interactive visual exploration of large datasets • Electrochemical simulations
Science/Medical	• Virus modeling • Cancer investigation • Prime number computation • Weather/climate simulations • Earthquake simulations • Astrophysical simulations • Particle physics simulations • Protein modeling • Genome-related research
Entertainment/Film	• Animation

The chances are good that your industry has a large computational problem that may fit well into the hive computing scheme. Geneer can help you evaluate whether hive computing is an appropriate solution.

THE DOWNSIDE OF HIVE COMPUTING

All this sharing of computing cycles is wonderful and the idea of making a buck off the idle resources of your computer sounds great as well. But there are some very real impediments to the adoption of hive computing.

There are significant questions regarding users' acceptance of hive computing. As Steve Steinberg pointed out in the excerpt quoted earlier, people regard their computers as their property, regardless of the legal ownership.

Finally, perhaps the biggest question about hive computing concerns the financial viability of the emerging business models.

Unfortunately, at least one Internet Service Provider (ISP), Juno, believes it has the right to use its users' processors as it sees fit. The ISP's user agreement[16] was recently amended 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 further asserts that they may amend the agreement at any time, and logging on to use their service implies assent to the current version of the agreement. Although their press release[17] announcing the Juno Virtual Supercomputer Project mentions that they will ask for volunteers, the company states that users of its free ISP service may be required to contribute idle resources. Press release notwithstanding, the Juno user agreement makes it clear that no permission is required.

Juno may be the first ISP to try this ploy, but they aren't likely to be the last. ISPs' networks will be additionally burdened by people using hive computing applications, and many of them might be tempted to recoup their losses through this type of arrangement.

Seizing users' computing resources against their will highlights one of the other major stumbling blocks for hive computing success: security. Customers of hive computing companies may not trust sensitive data and applications to anonymous, hard to control computers. In turn, users may not trust companies not to snoop around on their systems. It may be quite some time before these issues of trust are resolved.

Finally, perhaps the biggest question about hive computing concerns the financial viability of the emerging business models. The Internet has taught us that, while everything devolves to free, advertising and other indirect revenue sources may not be enough to support a business. 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.

Despite these problems, hive computing represents such tremendous potential that issues such as these are very likely to get worked out in time.

THE FUTURE OF COMPUTING NETWORKS

Hive computing and the larger P2P initiative are 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. If this is the beginning, what will the future look like?

Once the initial obstacles are surmounted, business hive computing networks will go beyond the current artificial division of labor between file sharing and idle resource maximization to provide a rich, collaborative and 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. 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.

Once the initial obstacles are surmounted, business hive computing networks will go beyond the current artificial division of labor between file sharing and idle resource maximization to provide a rich, collaborative and integrated marketplace of business capabilities both within the enterprise and externally.

As they evolve, business computing networks will develop the following characteristics:

- **Dynamic** – Communications are interactive, extensible, flexible, and easily reconfigured in real-time. Systems 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, support any number of participants and enable the discovery of new networked resources. This feature is perhaps the most important in setting workers free of location-based business interaction. We've already seen the beginning of this with call forwarding and cell phones.
- **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.
- **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 on-line 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 can be easily personalized or otherwise adapted to the business purpose.

CONCLUSION

As novelist William Gibson once observed, “The future’s already arrived; it’s just not evenly distributed yet.”[18] Your business will be affected by hive computing, even if only because the Web search engine you use employs it. Once a few of the details are ironed out, computing “power by the hour” will become a standard way of solving business problems.

John Gantz, noted columnist and VP at industry analyst firm IDC, puts it this way: “Jump on P2P, and stay ahead of the curve! Learn everything you can about it, try a few applications in pilots, get your arms around it before your internal customers do, and become your company’s expert on it. You’ll either catch the wave, if there is one, or you’ll know how to protect yourself if there isn’t one. Just don’t get caught flat-footed.”[19]

APPENDIX A

Peer-to-Peer Companies

Much of the material in the following table was derived from the lists maintained at PeerProfit (http://www.peerprofit.com) and Peertal (http://www.peertal.com). These sites are a tremendous resource for anyone wanting to keep track of the rapidly changing P2P marketplace.

All links were visited and supplementary material added in Q1, 2001.

Company or Effort	Web Site	Description
.Net	www.microsoft.com/	Microsoft's .NET effort is based on networked components, which are inherently P2P. In addition, its FarSite project is developing a distributed file system.
100X	www.100x.com/	A startup accelerator and software development company for P2P businesses.
24Link	www.24link.com/	P2P secure collaboration software.
2AM	www.2am.com/newweb/speer.asp	StreamingPeer™ delivers rich media content safely, quickly and cost-effectively. 2nd Order www.2ndorder.com/ A startup in stealth mode.
3Path	www.3path.com/	A priority, permission-based, straightforward, and fully outsourced peer-to-peer content delivery service for all digital content. Basically a variation on push technology
acceleratorgroup	www.acceleratorgroup.com/	A P2P investment firm.
Actedge	www.actedge.com/	A startup in stealth mode. Apparently seeks to be a P2P resource portal.
Affiniti	www.affinitigroup.com/	LiveCache forms the platform for LiveMessenger, a peer-to-peer messaging product .
AgentWare	www.agentware.net/	Syndicator ™ is a software development platform that integrates applications.
Aimster	www.aimster.com/news.html	Combines Napster-like file sharing with the AOL Instant Messaging service.
Amaya Web Editor/Browser	www.w3.org/Amaya/Amaya.html	A complete web browsing and authoring environment that implements collaborative annotation. From W3C.
Applied Meta Computing	www.appliedmeta.com/	Legion is a computing environment that can find needed resources, coordinate and execute required processes, and return results employing scheduling, data management, fault tolerance, site autonomy, and a wide range of security options. Sounds similar to Microsoft's .NET effort. Boeing, Harvard Medical School, the Naval Research Laboratory and NASA are clients.
Bad Blue	www.badblue.com/	Bad Blue is a web-based file sharing and Web publishing service that includes a tiny version of the Microsoft IIS Web server.
Beowulf	www.beowulf-underground.org/ www.beowulf.org/	Distributed computing project started in 1994.
Bitcollider	www.bitcollider.com/	A startup in stealth mode.
Biz2peer	www.biz2peer.com/	Developing a P2P platform.
Blogger	www.blogger.com/	A Weblog application that allows you to easily publish content to the Web. Interestingly, Blogger got its users to chip in for more servers, voluntarily.
Bluetooth	www.bluetooth.com	A peer-to-peer wireless networking protocol originally developed by Ericsson. Bluetooth devices are starting to come into the market now. For more information on Bluetooth, see the Geneer Wireless White Paper.
Bodetella	www.felmlee.com/bodetella/	File sharing Gnutella client for corporate intranets. Features download resume.
Brazil Project	www.sun.com/research/brazil/	This Web application framework effort from Sun links together a variety of their technologies, from Java to Jini, as well as wireless networking.
BXXP		Blocks eXtensible eXchange Protocol is an XML-based messaging framework for building application protocols.

Company or Effort	Web Site	Description
CenterSpan	www.getsocket.com/	Socket, a "Group Activity Hub" combines instant messaging with the ability to detect and launch multi-player games and other collaborative applications between users. It enables digital rights management.
Centrata	www.centrata.com/	A startup, currently in stealth mode, developing a distributed computing platform. It plans to convince PC makers to ship their software with new PCs.
Climate Dynamics Consilient	www.climate-dynamics.rl.ac.uk/ www.consilient.com/	Casino-21 is a climate simulation distributed computing project. Distributed computing infrastructure that aggregates any content and automates any business process.
CriticalPath	www.cp.net	Acquired PeerLogic. InJoin™ BATCH (formerly LiveContent BATCH) provides batch job execution environment on UNIX, featuring job step management and recovery, priority classes, conditioned execution, job monitoring and more.
Crowdshare (AKA Circlebox)	www.crowdshare.com/	Gnarly! Is a P2P file sharing application akin to Napster.
CuteMX.Com (GlobalScape, Inc.)	www.cutemx.com/	CuteMX is a music file sharing service that combines a powerful search engine, real-time chat rooms, instant messaging, friends (and enemies) lists, and a built-in media player.
Cytaq	www.cytaq.com/	Creator of Resource Router™, Universal Sharing Environment (USE™), Universal Query Language (UQL™), and other distributed platform-enabling technologies.
Datasynapse	www.datasynapse.com/	WebProc is a distributed, CPU-sharing system that offers CPU cycles for pay. Participants are paid in "flooz" credits they can use to buy on the Web.
Dcypher.net	www.dcypher.net/	Former name of ProcessTree.
Distributed.net	distributed.net/	At first a non-profit, Distributed.net has allied with United Devices to provide distributed computing resources.
Distributed Science	www.distributedscience.com/	ProcessTree enables peer-to-peer access to a commercial super-computing network of systems and other resources from across a registered heterogeneous network of 156,468 computers on the ProcessTree network.
Dotcast	www.dotcast.com	Creating a national high-speed digital network for the distribution of digital entertainment, interactive services, and multimedia communications.
Eazel	www.eazel.com/	Eazel Online Storage is a distributed file sharing application.
Ejacent	www.ejasent.com/	Ejasent UpScale™, supplies eBusinesses with on-demand Web application processing capacity instantly and transparently.
Elepar	www.elepar.com/	Creators of a P2P programming technology called Software Cabling.
eLiberation	www.eliberation.com/	ePilot™, is a free search engine and desktop portal with 953,000 members.
EMC	www.emc.com/	Storage vendor developed SymmAPI™-Sockets application programming interface for P2P file systems.
eMikolo	www.emikolo.com/	A startup currently in stealth mode, eMikolo has developed its PeerCasting™ and eMikolo Plus Platform™ technology to enable content owners, distributors and network operators to peer their content networks while dictating financial and copyright policies.
Endeavors Technology, Inc.	www.endtech.com/	Developed Magi suite of applications for P2P devices, remote control, and workflow.
Engenia Software, Inc.	www.engenia.com/	Developed EngeniaUnity™, an XML-based web-enabled collaborative workgroup solution and virtual distributed file system.
Entropia	www.entropia.com/	Entropia 2000, distributed computing for philanthropic organizations.
Evolvesoft	www.evolvesoft.com/	A startup in stealth mode.
eZ	www.ezmeeting.com/	EZmeeting software for virtual meetings.
File Navigator		A search engine client for Open-Nap networks supported by banner advertising.
FLIPR	www.flipr.com/	FLIPR (Future License of Intellectual Property Registry) is a P2P digital media distribution system.
Flycode	www.flycode.com/	Formerly AppleSoup, Flycode was started by a Napster founder, and is developing a network that lets content owners distribute anything digital.
Fracta Networks	www.fracta.com/	A startup in stealth mode.

Company or Effort	Web Site	Description
Freenet	freenet.sourceforge.net/	Freenet aims to create an information publication system similar to the World Wide Web based on the protocol in which information can be inserted into the system associated with a "key". Later anyone else can retrieve the information using the appropriate key.
Fuse Systems	www.fusesystems.com/	A startup in stealth mode.
Globus	www.globus.org/	The Globus Project aims to create computational grids, persistent environments that enable software applications to integrate instruments, displays, computational and information resources that are managed by diverse organizations in widespread locations.
Gnutella	gnutella.wego.com/	Gnutella is a fully distributed information-sharing technology incorporating client and server in the same application. Gnutella allows users to share files of all types.
gonesilent.com (aka InfraSearch)	www.gonesilent.com/	Recently bought by Sun and folded into the JXTA project, the company developed InfraSearch for real-time information sharing.
GPulp	gnutellang.wego.com/	A working group to develop the next version of the Gnutella protocol, the general Purpose Location protocol.
Groove Networks	www.groove.net/	Created by Lotus Notes developer Ray Ozzie, Groove is evolving groupware into peerware. Information sharing and real-time collaboration software are combined in this offering. Signed up 50 partners in its first 6 weeks of existence.
grub.org	www.grub.org/	Creates a distributed computing client, the first application of which is a Web indexer.
Hotline Communications Ltd	www2.bigredh.com/hotline3/	Hotline is a P2P community application featuring real time chat, conferencing, messaging, data warehousing, file transfer and streaming capabilities.
I5digital	www.i5digital.com/	Project Pandango allows users to conduct P2P searches on the web.
IBM	www-4.ibm.com/	IBM's SanFrancisco is a Java-based suite of business process components and services that can be used for P2P applications.
Ikimbo	www.ikimbo.com/	Ikimbo's Omniprise technology creates global, group-to-group communication and file sharing networks inside and outside the firewall.
iMaestro	www.imaestro.com	iMaestro Interact apparently is targeting ecommerce. Difficult to say from their site, though.
Imesh	www.imesh.com/	Distributed file sharing and collaboration.
IMXP	ietf.org/	A draft standard for instant messaging.
Infobot	www.infobot.org/	A Perl module that connects to an Internet Relay Chat (IRC) server, joins channels, and begins accumulating "factoids."
Inoize	www.inoize.com/	Distributed file sharing.
Intel	www.intel.com/	Intel formed a P2P working group that is meeting with some resistance from developers.
Interbind	www.interbind.com	Interbind's Exobind is a Java-based lightweight platform for mobile applications.
Invisible Worlds	www.invisibleworlds.com/	Developed Blocks eXtensible eXchange Protocol (BXXP) for P2P applications.
Jabber	www.jabber.org/	A cross-platform, XML-based instant messaging system.
Jini	www.jini.org/	P2P networking for devices, from Sun. Each device provides services that other devices in the community may use.
Jnutella	www.jnutella.org	Japanese Gnutella community.
Jungle Monkey	www.junglemonkey.net/	A Unix-based distributed file-sharing service.
Kalepa Networks, Inc	www.kalepa.com/	A startup currently in stealth mode, Kalepa will provide distributed content networks.
Killdara	www.killdara.com	Makes hardware and systems for P2P messaging.
KnowNow	www.knownow.com/	A startup in stealth mode, KnowNow has a technology that holds open the connection between a Web browser and a Web server. Has demonstrated a voting application.
Lightshare	www.lightshare.com	Lightshare enables e-commerce on peer-to-peer networks.
MangoSoft	www.mangosoft.com/	MangoSoft produces Cachelink, a software-based web caching product, and Mangomind, a business user file-sharing application.
Manila	manila.userland.com/	A content management system from Userland that can be used to create Weblogs.

Company or Effort	Web Site	Description
Meerkat: An Open Wire Service	www.oreillynet.com/meerkat/	Meerkat is an XML-based syndicated content reader from publisher O'Reilly.
Microsoft	www.microsoft.com	Microsoft's .NET initiative enables creation of P2P systems.
Mithral Communications on projects.	www.mithral.com/	Mithral's Cosm Phase 1 is a set of cross-platform open protocols and applications designed to allow distributed computers to work together
Mojonation	www.mojonation.net/	Mojo Nation enables publishing and sharing of any kind of data and is creating a digital marketplace for the exchange of idle disk space, bandwidth, and CPU cycles. Users must contribute "mojo" in order to download.
MyCIO	www.mycio.com/	A distributed antivirus network.
MyFileShare	www.myfileshare.com	PeerGenius™ is a Digital Content Distribution Management ™ (DCDM™) file sharing solution.
Napster	www.napster.com/	The Granddaddy of the current P2P wave. It is primarily a music file sharing service.
NeurOK	www.neurok.com/	Developing Knowledge Network that uses P2P for knowledge management.
New Productivity Initiative	www.newproductivity.org/	Distributed Resource Management (DRM) ensures that disparate computing resources are united to optimize the use of the network as a computing platform.
Newtella	www.newtella.com/	P2P file sharing.
Nextpage	www.nextpage.com/	The NXT 3 e-Content Platform enable distributed content serving.
Novient	www.novient.com/	iServerNet collaboration platform enables the sharing of people, projects, and knowledge. Accenture is a client.
Nullsoft	www.shoutcast.com/	Audio publishing system SHOUTcast is Nullsoft's Free Winamp-based distributed streaming audio system.
OFSI	sourceforge.net/projects/ofsi	Open File Sharing Initiative is developing P2P file sharing.
Ohaha	www.ohaha.com/	An open source network of P2P networks, originating from the Ukraine.
OnSystems, Inc.	www.onsystems.com/	Formerly InterFriendly, OnSystems created Tijit, which develops Virtual Internet Networks (VINs) that create secure virtual networks.
OpenCOLA	www.opencola.com/	openCOLA (Open Collaborative Object Lookup Architecture) is an open-source autonomous and collaborative agent that collects, analyzes, and delivers dynamic content.
OpenDesign (FKA What U Want)	www.opendesign.com/	Software tools for P2P applications.
OpenNap	opennap.sourceforge.net/	An open source version of the proprietary Napster server.
Parabon Computation	www.parabon.com/	Parabon's products are Pioneer, a distributed computing application, and Frontier, a platform for secure distributed computing for "compute-intensive" projects. Currently powering some cancer research.
PeerClub (Korea)	www.peerclub.com/	Developing P2P based file-sharing program called Solomon.
Platform Computing	www.platform.com/	LSF optimally manages job processing by spreading the workload across the network, orchestrating all resources, regardless of differences in hardware architecture or operating system.
Plebio	www.plebio.com/	A search engine that searches a peer-to-peer search network.
Pointera	www.pointera.com/	The Pointera Sharing Engine allows users to find, download, and share any file format as well as doing a meta-search which automatically checks other search engines as well.
Popular Power	www.popularpower.com/	Along with SETI@home, Popular Power is the poster child for distributed computing. During testing, they are donating their users' cycles to model flu virus, but will eventually charge for them and compensate users.
Porivo Technologies, Inc.	www.porivo.com/	Porivo PEER(sm) is distributed computing technology. Establishing a computing exchange.
Proksim Software	www.proksim.com/	Developing P2P infrastructure.
Publius	www.cs.nyu.edu/~waldman/publius/	An anonymous Web publishing system that inserts a layer of encryption that protects the identity of a publisher. A project at NYU.

Company or Effort	Web Site	Description
Quiq	www.quiq.com/	Collaborative customer service. Used in AskJeeves portal.
QuickCom	www.quickcom.com/	e-Courier™ uses IP Multicasting and JMS to support Guaranteed Quality of Service (GQoS) compliant, peer-to-peer networking over IP networks.
Radio Userland	radio.userland.com/	A "personal radio station," music organizer and player as well as a collaborative outliner, a writing tool that organizes stories, directories, presentations.
RDF	www.w3.org/Metadata/Activity.html	The World Wide Web Consortium's, Resource Description Framework, is a "declarative language and provides a standard way for using XML to represent metadata in the form of statements about properties and relationships of items on the Web."
ROKU	www.roku.com/	Roku Platform connects information together in "context" to view, use, and share information across all devices and networks. HP is reselling.
RSS 0.91	backend.userland.com/rss091	Rich Site Summary 0.91 is a lightweight syndication format for distributing news headlines that originated at Netscape.
RSS 1.0	www.egroups.com/	RDF Site Summary is an XML-based lightweight modular multipurpose extensible metadata description and syndication format. RSS 1.0 is a new version of RSS 0.91 that supports extensions.
Sandia National Laboratories	www.sandia.gov/	Developed a cyberagent intrusion protection application, still in the laboratory stage, which functions as a multiagent collective -- a distributed program.
Science Communications	www.sciencecommunications.com	Search engine for email.
SETI@home: The Search for Extraterrestrial Intelligence	setiathome.ssl.berkeley.edu/	Along with Popular Power, SETI@home is the poster child of distributed computing. It allows anyone with a computer and an Internet connection to take part in the search for extraterrestrial signals. It runs as a screensaver.
Simple Object Access Protocol (SOAP)	www.w3.org/TR/SOAP/	Originally developed by Microsoft and now a W3C effort, Simple Object Access Protocol is a lightweight, XML-based, protocol for exchange of information in a decentralized, distributed environment.
Spinfrenzy.com	www.spinfrenzy.com/	A membership-based media-sharing site for MP3s and video clips. Offers chat, instant messages, discussion boards and email.
Static	www.static.com/	Static Streamer™ reduces cutting bandwidth and server costs for streaming media or large file downloads using a distributed network and reflector technology that creates virtual bandwidth.
Surface Layer	www.surfacelayer.nu	An information portal for Gnutella users.
Texar	www.texar.com/	SecureRealms secure P2P file sharing.
The Free Haven Project	www.freehaven.net/	Enables anonymous publication.
theSupplyChain.com	www.thesupplychain.com/	ScEngine under development for supply chain P2P messaging.
Thinkstream	www.thinkstream.com/	A startup in stealth mode.
Toadnode	www.toadnode.com/	A P2P file-sharing program with multi-language support.
TurboLinux	www.turbolinux.com/products/enf/	EnFuzion cross-platform distributed computing application.
Ubero	www.ubero.com/	Ubero stands for "universal binding and execution of redundant objects," a distributed object computing platform.
United Devices, Inc.	www.ud.com/	Distributed computing. Currently helping with cancer research.
Universal Description, Discovery and Integration (UDDI)	uddi.org/	Universal Description, Discovery and Integration is a registry initiative lead by Ariba, IBM, and Microsoft to create a platform-independent, open framework for describing services, discovering businesses, and integrating business services using the Internet.
Uprizer	www.uprizer.com/	A startup in stealth mode.
Ventrada	www.ventrada.com/	Ventrada™ is an application that gives mobile professionals access to all the files in all their storage locations, no matter where they are.
ViralSounds	www.viralsounds.com	A startup in stealth mode.
Vtel	www.turbocast.vtel.com/	TurboCast Web is Webcast streaming technology for live events over an internal network or the Internet, and also can store any multimedia content for convenient, on-demand playback.
vTrails	www.vtrails.com/	Uses Full Duplex Packet Cascading (FDPC), a multicasting technique that relieves Web site congestion.

Company or Effort	Web Site	Description
WebDAV	www.webdav.org/	Web-based Distributed Authoring and Versioning is a set of extensions to the HTTP protocol which allows users to collaboratively edit and manage files on remote web servers. Supported by Adobe and Macromedia.
WebV2	www.webv2.com/	A startup in stealth mode that is developing an application platform and network infrastructure that extends to B2B collaboration between peers in the supply chain. WebV2 architecture is based on networked intelligent agents.
WorldOS Corporation	www.worldos.com/	An application server for decentralized applications. Apparently a one-man show.
WorldStreet	www.worldstreet.com/	WorldStreet Net offers distributed information distribution, workflow, and contact management for securities traders. Customers include Merrill Lynch, JP Morgan, and Deutsche Bank.
XDegrees	www.xdegrees.com/	A startup in stealth mode that claims to be working on a P2P metadata product.
XMethods	www.xmethods.com/	A portal listing P2P projects using SOAP.
XML-RPC	www.xml-rpc.org/	A precursor to SOAP, it uses XML and remote procedure calls to communicate information from one server to another.
XNS (eXtensible Name Service)	www.xns.org/	An open protocol and open-source platform for universal addressing promising a permanent identifier for a data container. Also XNS is a platform for "web agents" to negotiate the exchange, linking, and synchronization of information among different parties.
Yaga	www.yaga.com/	Secure file sharing.
Zion	www.zion.com/	P2P application platform. Jobster is a distributed job searching application.

FOOTNOTES

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

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

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

4 http://www.placeware.com/

5 http://news.cnet.com/news/0-1003-200-2603611.html

6 http://news.cnet.com/news/0,10000,0-1005-200-319851,00.html

7 http://openp2p.com/pub/a/p2p/2001/02/15/anderson.html

8 University of California at Berkeley (http://now.cs.berkeley.edu/) and Princeton University (http://www.cs.princeton.edu/shrimp/)

9 http://www.wired.com/wired/archive/3.11/geek.html

10 http://www.distributedscience.com/press/DSI_Envive.pdf

11 http://www.salon.com/tech/feature/2000/12/14/popular_power/index1.html

12 http://www.computerworld.com/cwi/story/0,1199,NAV47_STO54153,00.html

13 http://news.cnet.com/news/0-1005-201-3248711-3.html

14 http://www3.gartner.com/DisplayDocument?id=323309&acsFlg=accessBought

15 Interview with Frank Bernhard of Omni Consulting Group.

16 http://help.juno.com/privacy/agreement.html

17 http://www.juno.com/corp/news/supercomputer.html

18 http://www.saffo.org/sensors.html

19 http://www.computerworld.com/cwi/story/0,1199,NAV47_STO54153,00.html

ABOUT GENEER

Established in 1984, Geneer is a leading professional services firm specializing in the design and development of custom software applications. With a strong track record in using emerging software technologies to make its clients competitive in their market place, Geneer's longevity, dependability, and thorough understanding of technology and software development have earned it the honor of working for an outstanding group of clients, including many of the Fortune 100. Geneer was recently named as the first Microsoft Gold Certified Partner for eCommerce Solutions in the U.S.

GENEER EXECUTIVE BRIEFING SERIES

Geneer's Executive Briefing Series delivers essential insights on the future of e-business, the Internet, and the technology that can have a significant impact on your product and service offerings. If you are interested in receiving future Geneer white papers, please contact Geneer Business Development: 800-4-Geneer or 847-294-0300. For more information about Geneer, visit our web site at www.geneer.com or e-mail us at info@geneer.com.

Building Software. Building Business.℠

Geneer®

1400 East Touhy Avenue, Des Plaines, IL 60018-3340
phone: 847.294.0300 • toll free: 800.443.6337 • fax: 847.294.0358 • www.geneer.com