December 2004

P2P Business Applications: Still Waiting for the Revolution?

C. Pankaj
Southern Illinois University at Carbondale

Micki Hyde
University of Michigan-Dearborn

P2P Business Applications: Still Waiting for the Revolution?

Pankaj
Southern Illinois University at Carbondale
Pankaj@siu.edu

Micki Hyde
University of Michigan-Dearborn
hydem@umich.umd.edu

ABSTRACT

Since the launch of Napster in June 1999, peer-to-peer (P2P) has become synonymous with file sharing applications that are the bane for the recording industry due to copyright infringements and revenue losses. P2P promised a revolution in business computing which has not arrived (DiSabatino 2000). Most P2P studies in the MIS[1] area have focused on the file sharing P2P applications like Kaaza, etc. There is a need for MIS researchers to examine and suggest ways to promote practical P2P solutions for businesses. This paper brings forward the distinguishing characteristics of P2P applications with respect to the Internet and ecommerce and argues that if the purist rhetoric in P2P application of complete decentralization is abandoned, P2P can be effectively used to develop practical business applications. It also points to certain issues that need to be considered when developing and implementing P2P applications and outlines three levels of progressing complexity at which P2P applications can exist.

KEYWORDS

Peer-to-peer; Business Applications; P2P; Characteristics; Issues; Business Models

INTRODUCTION

Over the last few years peer-to-peer (P2P) has been branded as the next killer technology. Chairman of Intel, Andy Grove (Rutheford, 2000) said, "P2P revolution will change the world as we know it." The real reason why P2P applications have become exciting in today's world is primarily due to the increase in the computing power of end-user computing devices like the PC, the cell-phone, handheld computers, and other appliances. These devices have been traditionally relegated to the status of a client due to their limited capabilities. But now they have the capabilities of performing tasks that have been previously performed on more powerful computers or servers. Whereas traditionally an end-user could not do much with his/her computing device, a current state of the art computing device places a more powerful computer in his/her hands. At the same time, the use of resources on these computing devices is neither extensive nor continuous. As a result, these computing devices have spare resources available. There is a potential for the resources to be employed for other tasks. These resources can be tapped through an application stack that provides the needed functionality of connecting these devices on the network for accessibility, running applications from third parties, routing, etc. This is the promise of the P2P technologies. But for most, P2P technologies have come to stand for anonymous file sharing and copyright infringements that are the bane of the music and movie industry. Various researchers have focused their attention to this topic (Alexander, 2002; Clark and Tsiaparas, 2002; Krishnan, Smith, et al., 2003)

This paper has two purposes. The first is to examine that exact nature of P2P technologies that are currently making rounds in the industry and clearly point out its distinguishing features. We demonstrate here that there are not many. In principle everything offered by P2P technologies has existed and has been used for several years. The only difference now is the potential large scale deployment using end-user computing devices. Secondly, based on the distinguishing features of the current P2P technologies, we examine if these distinguishing features provide any real value for business applications. Specifically we examine what are the issues that need to be addressed before the technology can be used. We argue that there

[1] A survey of the peer reviewed journals on PROQUEST, when the search was done for articles with "peer to peer" in the title, yielded 11 articles. 9 of these articles dealt with issues in P2P. 5 of these 9 articles dealt with file sharing applications. The other 4 dealt with technical aspects of P2P.

are models at various levels that can be employed to use P2P technologies in for-profit businesses. Given the space constraints this paper is intended to be a treatise on P2P. It is more intended to direct attention to certain issues that we feel are relevant.

WHAT IS DIFFERENT IN P2P?

So what does the current wave of P2P technologies offer that is exciting and different from what has existed for last three decades? If one examines the history of computer networks, perhaps not much. The foundations have existed for a long time. Computer networks, by definition, are comprised of peer computers that communicate with each other. Chat programs like ICQ and AOL instant messenger are good examples. It should however be mentioned that in chat applications the detection of presence is still centralized and once the presence is registered then communication can happen in a peer to peer mode[2]. As per the seven layer OSI model (refer to Figure 1) and TCP/IP protocol, communication happens in a P2P mode in the first four layers of the OSI model (physical, data link, network, and transport). Routers and gateways are peer equipment and share the burden of tasks like routing, protocol conversion, and network management in a distributed fashion. Traditionally it is the top three layers of the OSI model (application, session, and presentation) that have functioned in a master-slave or non-peer mode. P2P technologies aim to make the behavior in the top three layers a peer behavior. For example algorithms for routing and searching in P2P networks like Gnutella are implemented at application level (Jovanovic, Annexstein, et al., 2001). The peer behavior at the application level is again nothing new and has existed for a long time in distributed computing applications, clustering applications, etc. But the current wave of P2P technologies does incorporate something new and that is the philosophy of use of end-user computing equipment and decentralization. While the use of end-user computing equipment opens up a host of new opportunities and issues, the decentralization opens up a host of issues primarily on account of several dimensions of what may be decentralized and how they may be decentralized (Boyle, 2001), the most important dimension being control. Possibility of performance degradation due to loss of control with high decentralization is a real issue (Miller, 2002).

A distinguishing feature of the current wave of P2P technologies is their aim to harness the computing power of the end-user computing devices (Shirky, 2000). This can be done firstly because end-user computing devices have become sufficiently powerful, that they have some spare processor time (Hayes, 1998), and they have storage that can be used for different tasks while simultaneously being used for routine tasks by end-users. Many PCs can perform functions that were traditionally performed by servers or more powerful computers about four to five years ago. Secondly, most of these end-user computing devices are connected to the Internet and hence can be accessed through the Internet. The P2P business models and applications as such then need to focus on the use of the end-user computing devices on a large scale to solve practical business problems. Scale using end-user computing devices not withstanding, the current P2P computing is neither new nor fancy when compared with what has existed for more than last two decades. Thus when one talks about P2P technologies in the current context, one needs to focus specifically on:

1. An application that uses the end-user computing equipment connected to the Internet (or a data network in a restrictive sense)
2. The end-user computing equipment that is being used by end-users for their normal day-to-day work and/or routine tasks.
3. Processing cycle and/or storage resources from the end-user equipment used for some business task.
4. The pool of computing equipment is dynamic and computing equipment enter and leave the pool at will or in a random fashion.
5. Each node can acts as a server and client at the same time.
6. There is high peer diversity.

[2] No doubt the one will agree that centralized presence detection or discovery is the only feasible way of discovery when Internet access with dynamic IP addresses. Static IP addresses can though enable direct conversation like with telephones

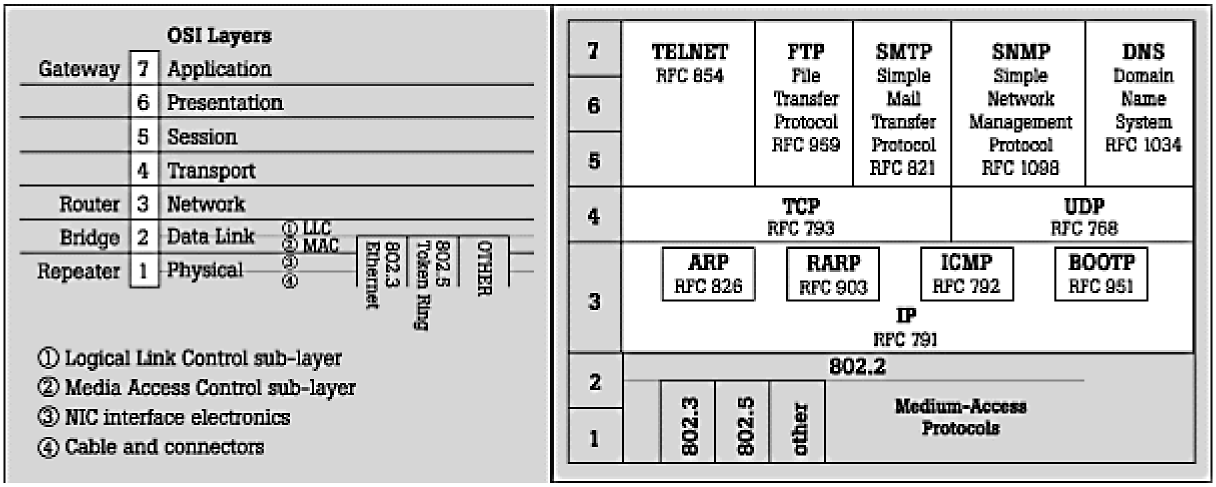

Figure 1. OSI Seven Layer Model (http://www.netc.org/network_guide/c.html)

All these characteristics are important. For example, many supercomputers have been built using PCs working in unison. The end-user computing device in this context is solely dedicated to the supercomputing architecture and not used by end-users and so may not be the focus of P2P computing that is focus of this paper.

The other distinguishing feature of the current P2P technologies is that of decentralized application performance without central control. If for no other reason, technical performance considerations preclude the use of highly decentralized operations. For example it was shown that pure P2P networks like Gnutella do not scale well and exhibit the small network effect (Jovanovic, Annexstein, et al., 2001). In the context of a for-profit business performance and security may be the main concerns which preclude the use of decentralized applications. These are discussed in the next section.

P2P, Utility Computing, and Grid Computing

A relevant question is how the current wave of P2P technology is related to concepts like utility computing and grid computing. We define utility computing as the business model and associated technologies for provisioning computing capacity to organizations based on their demand patterns, primarily to meet peak demands. This allows organizations to save money by not investing in equipment meant primarily to meet peak demands. P2P technologies may be used to provide additional capacity but this is not a primary aim. Utility computing uses server-class computing equipment in a data center serving several organizations simultaneously. Grid computing involves running a single task across several (maybe thousands) computing devices under some sort of centralized or distributed control. These computing devices may be end-user computing equipments or servers, though most commercial uses of grids employ servers (Anonymous, 2004) with the aim of increasing server utilization. The following table provides salient differences in the two modes of computing (Zhang, 2002):

Grid Computing	P2P Computing
Most grid nodes act primarily as a server	Each node acts both as a client and a server
Grid offers direct access to resources in remote and targeted sites	Access to resources is random without any specific target
Grid has pre-determined registered clients and servers	No distinction between clients and servers, peers and not registered and enter and exit at random
Services are reliable and guaranteed	Services are only partially reliable and guaranteed
Security is assured for each participating node	Security is not guaranteed to a participating node
Resource monitoring, allocation and scheduling is done centrally	No central control and relies primarily on self-organization

Table 1. P2P and Grid Computing

We would like to stress that the grid computing may be the only feasible model for the use of P2P technologies for business applications with deterministic performance considerations. Also the P2P technologies would benefit from the use of the specific term of grid computing rather than the use of a more generic term of P2P computing. Grid computing involves splitting up a task into several subtasks which may be run on various computers. The subtasks in most cases are all essentially similar and can be run independently of each other. Subtasks may be decomposed to a level such that can be run on machines with progressively smaller and smaller computing power. The results from the subtasks are combined to achieve the final results. Grid computing projects like Seti@Home and the American Diabetes Association have used end-user computers connected to the Internet to achieve significant milestones. The following table depicts the statistics on 02/15/2004 from Seti@Home. Table 2 below demonstrates the immense potential of grid computing for computationally intensive tasks. Tasks in areas like risk, assessment, simulating economic conditions etc. can benefit from grid computing applications (Anonymous, 2004).

	Total	**Last 24 Hours**
Users	4882229	1300
Results received	1249512223	1468874
Total CPU time	1814799.447 years	1265.581 years
Floating Point Operations	4.457283e+21	5.728609e+18 (66.30 TeraFLOPs/sec)
Average CPU time per work unit	12 hr 43 min 23.0 sec	7 hr 32 min 51.3 sec

Table 2. Seti@Home Statistics (2/15/04)

ISSUES IN BUSINESS P2P APPLICATIONS

Given the immense potential of the P2P technologies, what is holding business applications back? Scale using the end-user computing equipment opens up a host of opportunities and issues. There are both technical and organizational issues. What is probably lacking is a sound business model. If a successful business model can be put in place then business applications may be rolled out.

The first issue one needs to examine is the issue of control. P2P computing using a scale that is employed by applications like Seti@Home takes an organization far beyond its boundaries. A business has certain performance expectations that it needs to meet (Galbraith, 1977). To meet its performance requirements, every business needs to control the resources so that they can be employed to use as per the business needs. In information systems terminology one may talk about service level agreements (SLAs) for the resources. To ensure that the employed resources meet the SLAs implies that a business should be able to control the resources through some means. In a decentralized collection of computing devices with inherent heterogeneity between the nodes and the dynamic constitution of the pool the co-ordination and control poses immense burdens. The only way to establish and maintain control is to communicate the service level requirements to the end-users and motivate them to meet these by means of a rewards system. Applications like Seti@Home induce participation for philanthropic reasons and do not have any control over the end-user computing devices. The organization using Seti@Home is non-profit and working for the betterment of the human race. Issues related to rewarding or paying the end-users range from assigning a fair value to the resources used on an end-user computing device, payment mechanisms, cost of acquiring and running a comparable big computer in house, etc. Additional cost items like the cost of dealing with thousand of end-users whose computing devices are being used also needs to be considered. These end-users will need to be serviced with problems ranging from application faults on their computers, unpaid credits for work done, etc.

An associated question with ensuring service level agreements and inclination to control the end-user equipment is the question of security and liability. A business may not want to be held liable for issues relating to damage of the end-user equipment, violation of privacy due to bugs in application, etc. The other side of the security and liability coin relates to the end-user liability in case data is stolen from the end-user computer through a hack attack, backdoor, Trojan, or corrupted results are deliberately sent back to the business, etc.

This paper can not possibly discuss all the security and liability issues associated with the use of end-user computing equipment on a large scale. Broadly when we talk about use of P2P technologies in a for-profit organization we may look at four primary issues:

- meeting service levels on the end-user equipment,
- compensation for the end-user computing device usage,
- liability of the organization, and
- liability of the customer.
- security[3]

In addition, for use of P2P to become mainstream, off-the-shelf technologies that address these issues need to be developed and made available to organizations. Accounting models which can measure the overall contribution of a computing resource to a task are also needed. The models need to be built into the P2P technologies. Available P2P development toolkits like Globus (Unknown, 2004) address some of these issues, but not all. Globus complies with the Open Grid Services Architecture (OGSA) and provides grid security, remote job submission and control, data transfer and other facilities which may ensure some level of service (though not necessarily high performance) and does address some liability issues but does not have any mechanism for compensating end-users.

These are not issues in the most popular P2P file sharing applications. Though the issue of liability of the business and customers are on the forefront, it is primarily a byproduct of the nature of tasks being carried out rather a primary consideration. File sharing networks do not guarantee any service levels. End-users are participating purely for what one may designate as self actualization, rebelling against the big corporations, and monetary benefit (without any explicit compensation mechanism). The end-users also participate for reciprocation or in a spirit to give back when they have gained something. Liability questions are complex and subject of various lawsuits across the globe. Security has been a concern with spyware that come bundled with most popular file sharing application and viruses spreading through innocuous looking files. But these are no different than what people come across while surfing on the Internet and downloading email attachments.

A side product of the liability issues has been endowment of anonymity on the end-users through various technological solutions, most of which mainly stress total decentralization and hence lack of control on how these applications are being used and who is using them. The technology provided by the P2P application Overnet was the main step in this direction. This decentralization and anonymity has branched off a thought process with respect to the P2P applications where it seems like the real challenge is to design applications which are decentralized in an extreme sense (Schoder and Fischbach, 2003). But we would like to emphatically stress that, due to issues outlined earlier, decentralization is not a criteria for business P2P applications with the possible exception of some specialized applications. Centralized control appears to be a necessary to have reasonable performance expectations. Work using Dynamic Hash Tables to overlay a stable set of nodes on a dynamic set of P2P nodes to create good scaleable networks is based on the same requirement (Li, 2004).

What is feasible in a P2P business application scenario is applications with a strict central control so that the issues of performance, end-user compensation, business liability, end-user liability, and security are in a deterministic state at all times. Totally or highly decentralized models are infeasible. We should stress here though that there is research underway aimed at building up truly distributed P2P applications. The application/technology that will perhaps lead to true distributed P2P application is the ad-hoc networks standards developed by the Internet Engineering Task Force (IETF). These standards will enable the formation of networks on the fly without the need for central routing with nodes entering and leaving the network at will. However, these standards may not be able to meet the performance requirements for business applications which typically have large transaction volume, while taking care of the liability issues. Ad-hoc network standards are only being targeted for specialized applications like mobile networks for the army on the battlefield.

BUSINESS MODELS FOR BUSINESS P2P APPLICATIONS

We would like to stress that the way forward is not with a pure distributed model spread across thousands of end-users since this model would be infeasible for most practical business applications. The element of control can not be eliminated for

[3] We feel that security is an important issue but it is not a show stopper. Existing technologies in the area of encryption, firewalls, authentication etc. in our opinion are sufficiently strong to handle most of the concerns.

performance and liability (security) reasons. Businesses can jump start developing and implementing P2P applications by removing this consideration in their P2P application development and deployment.

We propose that feasible business models for P2P applications exist at three distinct levels. We compare these levels as akin to intranet, extranet, and Internet for information sharing. P2P applications may exist in ecommerce terminology as business-to-employee, business-to-business, and/or business-to-customer modes. As the organization moves from one level to another it expands its scope from internal to external leading to an increase in complexity of issues related to control (liability), pricing, and technology. These levels are depicted in Table 3.

Level	Scope	Complexity of Control	Pricing Models	Technologies
B2E (Business-to-Employee)	Completely internal (e.g., Intranet)	Low complexity since resources can be controlled	Allocating credits to end-users can be done through the existing accounting standards for allocating overhead, etc.	Existing desktop management technologies can be employed
B2B (Business-to-Business)	Extended to trusted business partners (e.g., Extranet)	Moderate complexity since resources are shared between businesses	Existing business models from utility computing and application service providers may be used	New applications needed that can be built on existing technologies for secure communications and distributed computing
B2C (Business-to-Consumer)	Extended to end-users and customers (e.g., Internet)	High Level of Complexity since several thousand end-user machines are involved that may be difficult to control. High resource coordination and control costs.	Microtransaction and aggregation-based models	New applications using the emerging technologies and standards like web services and existing technologies from distributed computing

Table 3. P2P Applications at Increasing Level of Complexity

A business can jump start the use of P2P applications that are implemented within the boundaries of the organization. There are some instances of it happening in various financial institutions as mentioned earlier in this paper. The aim is to use the servers with spare capacity to execute computationally intensive tasks through the use of grids (Anonymous, 2004). An instance of a server less file system has been implemented on an experimental basis within using desktops within Microsoft (Bolosky, Douceur, et al., 2001). The scale in this case will not be large as applications like Seti@Home, but it is the first step in experimenting with the technology. There are several applications that are in use in organizations and can be implemented without much pain. Oil companies use the spare capacities in the desktops to process oil exploration data. In case of professional services like law firms document storage can be distributed across the desktops in the organization and new documents can be added to the knowledge base as documents are generated at end-user machines. Routing, searching and other resource co-ordination can be done much more efficiently since the pool of computing devices is relatively static and deterministic. Technologies that exist for desktop management primarily to reduce the total cost of ownership will allow an organization to implement such solutions while addressing all issues outlines earlier. Off the shelf solutions for collaboration, oil exploration, and other areas are sold by vendors like SUN, UNISYS, etc.

For implementing applications at the other two levels, we stress that a specialized intermediary is needed who can take care of the technological and business issues, keep the costs down, and ensure P2P applications affordability. Even with the technological issues ironed out interorganizational contractual issues may need to be sorted out. Most challenges though will come at the B2C levels where payment mechanisms like microtransactions and aggregation may be needed for the end-users.

CONCLUSION

It is important for the MIS academicians to take a holistic and practical approach to the P2P applications and look past the hype. Understanding what is feasible will allow us to channel our energies into the study of issues that will bring in both immediate and practical benefits to the business organizations. We should study issues related to applications running on end-user computers to benefit the organizations by better uses of slack resources.

There are many areas for future potential research, the most important of which is a payment scheme of microtransactions that will allow for-profit businesses to make a transition to the B2C model. We are currently working in this area as well as an empirical study of the characteristics and issues put forth in this paper. Other potential interesting areas of research are examining issues of return on investment on P2P computing applications. For instance, running applications on desktops inside the organization may increase the total cost of ownership for the desktop but lead to savings in server related costs. We hope that this paper will shift interest in the MIS community from simple file-sharing to a broader view of issues that will bring practical benefits to the organization.

REFERENCES

1. Alexander, P. J. (2002). Peer-to-Peer File Sharing - The Case of Music Recording Industry. *Review of Industrial Organization* 20: 151-161.
2. Anonymous (2004). Gearing up for Grid. *IT Utility Pipeline*. http://www.itutilitypipeline.com.
   Bolosky, W. J., J. R. Douceur, et al. (2001). Feasibility of a Serverless Distributed File System. http://research.microsoft.com/sn/Farsite/Sigmetrics2000.pdf.
3. Deployed on an Existing Set of Desktop PCs. *M. Research*. Redmond, Microsoft.
4. Boyle, T. (2001). Independent Axes for Centralization /Decentralization of Transaction Systems. http://www.ledgerism.net/P2Paxes.htm.
5. Clark, J. A. and A. Tsiaparas (2002). Bandwidth-on-Demand Networks - A Solution to Peer-to-Peer File Sharing. *BT Technology Journal* 20(1): 5-16.
6. Galbraith, J. R. (1977). *Organization Design*, Addison Wesley.
7. Hayes, B. (1998). Collective Wisdom, American Scientist. http://www.amsci.org/amsci/issues/Comsci98/copmsci1998-03.html
8. Jovanovic, M. A., F. S. Annexstein, et al. (2001). Scalability Issues in Large Peer-to-Peer Networks - A Case Study of Gnutella. *University of Cincinnati Technical Report*.
9. Krishnan, R., M. D. Smith, et al. (2003). The Economics of Peer-Peer Networks. *JITTA: Journal of Information Technology Theory and Application* 5(3): 31-45.
10. Li, S. (2004). JXTA 2: A High-Performance, Massively Scalable P2P Network, IBM. http://www-106.ibm.com/developerworks/library/j-jxta2/
11. Miller, E. (2002). Decentralization, Gnutella, and Bad Actors. Lawmeme, Yale. http://research.yale.edu/lawmeme/modules.php?name=News&file=article&sid=291
    Schoder, D. and K. Fischbach (2003). Peer-to-Peer Prospects. *Communications of the ACM* 46(2): 27-29.
12. Shirky, C. (2000). What is P2P... And What Isn't?, O'Riley Network. http://www.openp2p.com/lpt/a/1431
13. Unknown (2004). Status and Plans for Globus Toolkit 3.0, Globus Alliance. http://www-unix.globus.org/toolkit/
14. Zhang, X. (2002). Research Issues for Building and Integrating Peer-based and Grid Systems. Beijeng. http://www.nesc.ac.uk/talks/china_meet/zhang_beijing_talk.pdf