Post Napster: Peer-to-Peer Revisited

Sean Mays
February 20, 2000

Peer-to-peer (hereafter referred to as P2P) is the accepted term for the revolution in file sharing and related technologies that has cropped up within the past year. The popularity of programs like Napster and Gnutella, regardless of their legality, has brought new attention to P2P technologies and their potential benefit within the workplace to harness the dormant power of desktop PCs. The benefits are so great that major vendors such as IBM, HP, and Intel want to standardize and commercialize the technology. With the adoption of this technology into the mainstream, I will attempt to define P2P, examine the reasons for P2P adoption, and examine security models contained within Groove.

P2P Defined

peer-to-peer architecture
A type of network in which each workstation has equivalent capabilities and responsibilities. This differs from client/server architectures, in which some computers are dedicated to serving the others. Peer-to-peer networks are generally simpler, but they usually do not offer the same performance under heavy loads.
http://webopedia.internet.com/Computer_Science/Client_Server_Computing/peer_to_peer_architecture.html

The architecture underpinning the P2P definition above is not new. The concept of computers acting as peers on the Internet has long been established with the usage of utilities such as telnet and ftp. Microsoft operating systems since Windows for Workgroups have long permitted users to share resources amongst peers in their workgroup.

Taken literally, this definition of P2P could easily apply to phones, email and even peer-to-peer games such as Doom or Quake. Napster, the most of infamous of P2P applications, could be removed by the definition since it utilizes a centralized server to store pointers and resolve addresses. Are we to deny Napster its P2P status and herald Quake as the perfect peer application? Does this definition really seem to fit the recent changes in Internet usage? Or do we need a new label to define these technologies?

Actually, it’s a mixture of both.

The architectural definition is useful for recognizing the transformation in roles of PCs that underline this model. PCs are transformed from passive clients dependent upon other servers for resources to active participants in the new server roles they perform. Even Napster users who never add to the collection become active participants in offering the music they have downloaded to other members within the network. They raise the once anonymous PC to the level of contributor to a larger effort.

Typically servers have fixed IP addresses for offering their services so that users could readily access their resources without too much trouble. P2P applications allow users to interact and share resources over a variety of network connections and many without fixed IP addresses. Even without fixed IP's, P2P applications have built-in mechanisms for finding resources for their users regardless of their location and status on the Internet.

The best definition that combines these qualities within the literature states:

> 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. (1)

The author further proposes two criteria for determining whether an application is truly P2P. If an application fails to meet one of the criteria, then it is not P2P.

1. Does it treat variable connectivity and temporary network addresses as the norm?
2. Does it give the nodes at the edges of the network significant autonomy?

P2P applications

Generally, P2P applications can be categorized into three models of usage: instant messaging applications, workgroups, and distributed computing.

Instant messaging includes applications such as Jabber and Aimster as well as some file-sharing applications such as Napster and Gnutella. As Tim O'Reilley has noted on many occasions, "Napster is really just instant messaging where the question isn't 'Are you there?' but 'Do you have this file?'" (2)

Workgroups permit individuals to collaborate over the Net on a joint project. Groove Networks provides a groupware "LAN on demand" for ad-hoc groups of peers to share not only their files and chat, but for a wide variety of shared applications as well. Distributed computing permits an organization to take advantage of the available underutilized cycles on thousands of PCs and collect the data. Recent studies estimate that most companies utilize less than 25% of their PCs computing and storage capacity. A well know example of this is SETI@Home that utilizes the spare cycles of more than one million PCs to analyze radio telescope data in search of extra terrestrial intelligence. Even businesses are utilizing this technology to reduce operating costs. Intel has been using the technology since 1990 to slash the cost of its chip-design process. The company uses a homegrown system called NetBatch to link 10,000 computers, giving its engineers access to globally distributed processing power. (3)

Appeal of P2P

With their popularity and ease of use, Napster and ICQ have set the course for what lays ahead for us. P2P applications will find their way onto user's machines and appeal to users for a number of reasons: ease of installation, maintenance and use; minimal requirements of IT resources or dependencies; and productivity enhancement.

Simplicity is the driving factor for the adoption of the majority of these P2P applications. Users want "one stop, one button, quick download, install this" type of applications. Before Napster, if a user wanted to serve files from their PC, you needed a permanent IP address, a domain name, and registration with domain name servers and properly configured Web server software on the PC. Within minutes, not hours or days, users have all of the functions of a Web server with none of the hassle. Software so simple that it is easy for non-technical people to feel comfortable setting it up themselves.

There is minimal dependency upon IT resources for implementation. Traditionally, messaging technologies like ICQ would have required months of IT resources to procure, develop, test and implement. Today, these services are readily available for free and require no resources in the way of money, time or services from IT departments.

Workers themselves are apt to realize the need for better products to enhance productivity in the workplace. The appeal of P2P applications such as Groove is driven by a general need for better communication mediums than email and email attachments. If users see these products as fulfilling needs missed by currently deployed IT applications, then P2P applications will find their way onto user's machines whether we want them to or not. When users are offered "better" solutions to their needs, users will take control of the matter into their own hands.

IT Charge

Given that some of these applications will sneak into the workplace, IT professionals should be actively evaluating and looking at the security models these applications utilize and select ones that mesh with our security practices and policies. After selection of an application(s), we should define an acceptable usage policy for and perform regular security seminars for these applications to our users. To counter usage of these P2P applications for the transmission and relay of viruses and Trojans, businesses should have AV software in place and automate the deployment of new virus definitions as they are released.

Key Technologies and Security, Inc., has charted the best practice and recommendations for businesses to deal with P2P software: (4)

--Establish security policy
--Define acceptable usage policy
--Perform regular security seminars for users
--Perform regular audits of security policies and procedures
--Install and perform AV software updates

The paper goes on to recommend blocking at the firewall to known P2P servers and clients. However, this task is daunting given that most of these technologies can easily bypass most Firewalls by abusing port 80. The key to maintaining our security objectives must be met by recommending applications that have sound security practices in place under the hood and optimize and work with our available resources.

Groove Security Model

Groove makes an impressive set of security guarantees to users and IT professionals alike in offering mechanisms for ensuring confidentiality, authentication, integrity, and fault-tolerant availability. (5)

Confidentiality, authentication, and integrity of the workgroup and its space are ensured at all times by strong security and encryption. The application automatically encrypts all materials on the user's disk and across the network as it travels between peers. Groove uses a 192-bit encryption passphrase to encrypt these shared spaces. Even if a user loses their machine or someone else gains access to their desktop, shared space is still protected by the user's passphrase.

Group members automatically exchange public keys via vCards when joining their group. Users can utilize these vCards to ensure integrity and authentication throughout their communications. Because invitations to join groups are dependent upon email, Groove provides additional security mechanisms within the system to verify the sender's identity to a prospective group member: voice annotation and a digital fingerprint via the vCard containing the user's public key. This peer distribution ensures security without requiring further centralized certificate and key management. (6)

Once a member is uninvited from a shared space, Groove automatically issues new shared space keys to current members so that all subsequent data is protected and kept private from past members. Uninvited members still retain access to previous content of the space, but they can no longer look at new content and activity.

Availability of services is facilitated in the Groove software by using relay servers. Relay servers provide flexibility for users whether they are firewalled, offline, or connected by a very slow link. Fault-tolerance is built in for the recovery of group data. In the event a peer's machine crashes, the data can be recovered from one's peers once Groove is reloaded. Groove minimizes network traffic by only relaying changes that occur rather than retransmitting the entire shared space.

All in all, Groove Networks delivers a very solid foundation for their product. There are some issues I discovered in playing with their product that need consideration in future developments. The option to remember a user's passphrase seems mute if the application is running on insecure products such as Windows 98. If passphrases are the key to the whole shared workgroup community, why offer to save it in the first place. My last complaint stems from the fact that uninvited members still can access previous contents of the space. In some situations, I can easily envision the need for removing the materials especially if they are of a highly sensitive nature.

Recommended Reading:

Cave, Damien. “Come together, right now, over P2P”. December 14, 2000. URL: http://www.salon.com/tech/feature/2000/12/14/popular_power/index.html (February 20, 2001).

Gillmor, Dan. “ Peer-to-peer computing: The next IT tsunami?” October 2, 2000. URL: http://www.computerworld.com/cwi/story/0,1199,NAV47-68-86-101_STO51550,00.html (February 20, 2001).

(6) Groove Networks. “Security Services Tech Brief”. URL: http://www.groove.net/feature/security/securitybrief.pdf (February 20, 2001).

Hamblin, Matt. “Rewards May Outweigh Risks of Peer Networking”. December 4, 2000. URL: http://www.computerworld.com/cwi/story/0,1199,NAV47_STO54614,00.html (February 20, 2001).

(4) Key Technologies and Security, Inc. “Security Concerns for Peer-to-Peer Software”. July 18, 2000. URL: http://www.ktsi.net/pdf_files/Security_Concerns_Peer-to-Peer_KTSI.pdf (February 20, 2001).

(3) McDougall, Paul. “The Power Of Peer-To-Peer”. August 28, 2000. URL: http://www.informationweek.com/801/peer.htm (February 20, 2001).

Minar, Nelson and Marc Hedlund. “A Network of Peers Peer-to-Peer Models Through the History of the Internet”. URL: http://www.oreilly.com/catalog/peertopeer/chapter/ch01.html (February 20, 2001).

O’Reilly Open P2P site. URL: http://openp2p.com/ (February 20, 2001).

O'Reilley, Tim. "Remaking the Peer-to-Peer Meme". December 15, 2000. URL: http://www.openp2p.com/pub/a/p2p/2000/12/05/book_ch01_meme.html (February 20, 2001).

Shirky, Clay. "P2P Smuggled In Under Cover of Darkness". February 14, 2001. URL: http://www.openp2p.com/pub/a/p2p/2001/02/14/clay_darkness.html (February 20, 2001).

(1) Shirky, Clay. "What is P2P… And What Isn't!" November 24, 2000. URL: http://www.openp2p.com/lpt/a/472/ (February 20, 2001).

(2) Sims, David. "P2P Directory". October 20, 2000. URL: http://www.openp2p.com/pub/a/p2p/2000/10/20/directory.html (February 20, 2001).

(5) Udell, John, Nimisha Asthagiri and Walter Tuvell. "Peer-To-Peer: Harnessing the Power of Disruptive Technologies". URL: http://www.groove.net/feature/security/ch18.gtml (February 20, 2001).