



## ***Introduction to Groove***

## The Internet Goes Peer-to-Peer

A new sense of excitement runs through the Internet these days. Napster, Gnutella, Freenet – these technologies have re-introduced a style of computing – peer-to-peer – that has captured the imagination of consumers, developers, businesspeople, entrepreneurs, copyright holders, and technology companies.

In fact, each week the computer and business press seems to uncover another half dozen or so peer computing companies with a new twist on Internet search, file swapping, micropayments, personal communication and other applications.

So far, peer computing innovations have come in three different flavors.

**Direct access to information.** Peer-to-peer search and file transfer tools, such as Napster and gnutella, have burst onto the Internet, disrupting not only the business models of copyright holders, but reorienting users' notions of what content is available and desirable. While the Web is full of interesting and useful content, often times the most urgent and personally relevant content – music, photos, work files, drafts and works-in-progress, calendars and schedules – are found on the desktops of individuals, beyond the reach of any centralized or Web-based search engine. In this light, there is no mystery to the popularity among end users of tools that easily discover and retrieve content at the “grass roots” or “edge” of the network.

**Direct access to computing power.** There is a class of computations involving such massive amounts of data that they require supercomputers to perform them. Examples of these include pattern detection algorithms to help discover trends for weather forecasting, credit card fraud detection, stock market tracking, economic analysis, corporate data mining. Peer-to-peer computing can greatly economize these computations by distributing the number crunching to peer computers found across a network. SETI@Home, which distributes the computations that listen for clues to the existence of extraterrestrial intelligence, is the most popular example of such a peer computing application deployed over the Internet.

**Direct access to people.** When one looks at the most popular uses of the Internet in general, it is readily apparent that email far outdistances Web browsing and electronic commerce. In

*The Internet transformed the world by linking computer networks.*

*The World Wide Web transformed the Internet by making it easy to link files on those computers.*

*Now another major transformation is occurring.*

*By enabling millions of computer users to search for files and transfer them from one desktop computer to another, the balance of power shifts from the commercial interests that now dominate the Internet to the individual.*

*New York Times  
June 29, 2000*

survey after survey, email remains the primary driver of Internet use and adoption. It is no wonder then, that the most widespread use of Internet peer computing is not music- or file-swapping, but rather instant messaging. Interestingly, like Napster and SETI@Home, this communications application is actually a blend of computing architectures: centralized, server-based “awareness” capabilities combined with direct, text-driven, peer-to-peer interaction.

Like peer-to-peer file sharing, the adoption of instant messaging has been driven by *individuals*. In the absence of a single organization imposing a technology, this peer technology has naturally lent itself to ad hoc conversations within relatively small circles of “buddies” in business and non-business settings alike.

Clearly, peer computing has a bottom-up, grassroots appeal to it. This appeal is mirrored by a business-driven dimension that attracts innovative developers, systems integrators and business managers. So, what is it about peer computing that has immediate appeal to end users and technologists and business people alike?

**Personal Control.** There is an immediacy to peer computing that is absent from Web computing. As mentioned earlier, peer computing provides direct access to content. A person not only can *view* the content, but also *move* the content from one peer device directly to his own. There is no greater degree of control over content than when the user has it stored locally, on his or her own machine. The user now “owns” that content; it is always available, whether or not the user is connected to the network. Anyone who has saved the contents of email in a local directory, or detached a file, has participated in and appreciated the difference between “having” the content and “viewing” the content on a Web site.

**Context.** This immediacy is even more apparent in direct communications, where the peers are not just devices but people. Personal peer interaction takes place in real time. That is, peers are able to go beyond merely sharing or swapping content, to talking with each other about that content, to jointly interacting with or modifying that content. In this way, an intuitive and shared context emerges from the marriage of content and activity.

**Security.** The Web is often an appropriate medium for sharing information among a group. On the other hand, some content is sensitive enough that users may not be comfortable using a

*Now combine the server qualities of a Napster-like application with the ability of the central computer to note that you're online.*

*Suddenly, you can be a Web site yourself — a publisher — creating content on your own computer, dishing it out to anyone you want.*

*Keep in mind that this world doesn't exist yet, but it could.*

*San Jose Mercury News  
August 5, 2000*

Web server outside the home or corporate firewall to share it with others. Malicious or accidental access or deletion of Web-based content is not an altogether uncommon occurrence. Because peer computing does not rely on a third-party trusted intermediary, it lends itself nicely to highly secure communications and content sharing.

**Flexibility.** When a user owns the content and the collaboration environment, he or she can control how it is used. If the user wishes to do something with the data that the collaboration application doesn't readily support, he can add or use a different tool. For example, if a user does not like to use the built-in player that comes with a music swapping application, he can use an alternative music player. Because the content belongs to the end user, and not to the Web hosting service, the end user has broad discretion about how to view, manipulate and use that content.

**Cost Effectiveness.** Peer computing can have a significant positive impact on the system and human resources required for group interaction. Let's say a team of coworkers and contractors needs to meet regularly, investigate options, make decisions, and share documents, updates and ideas. Without much expense, they could use a shared space on a Web site. However, if the Web site needed to accommodate hundreds of such teams, any number of which may be actively online at any given time, the Web site would need to expand in capacity. Because of the resources consumed, administrators have to police who can and who cannot create spaces; they have to do backups of these spaces; and, since most of the time no single member of a team assumes the responsibility of declaring the collaboration to be "over," administrators have to prod people to delete their shared spaces.

In an edge-based peer-to-peer environment, each member of a team makes use of his or her own local computing resources, without any worst-case capacity planning necessary. There is no need to have bandwidth, processing and storage enough to support the entire universe of users, only enough to support his or her own usage. And, when an individual user finds that stored assets no longer merit storage on a local machine, he or she can delete them without having to ask permission of the rest of the team. This eliminates the administrative burden of managing access to services, backing up resources that may be obsolete, and tracking down abusers of the system.

Consider an even simpler person-to-person interaction: e-mail file attachments. Everyone is aware of the network inefficiency

*Boosters say the technology is one of those once-in-a-decade ideas that will change the face of computing...*

*Skeptics, and there are many, say the accumulating froth over peer-to-peer reminds them of other big ideas that came and went without ever living up to overheated predictions.*

*Wall Street Journal  
July 5, 2000*

of sending a single email with a file attachment to ten recipients, only to have some recipients reply to all with the file still attached. Not only does this unnecessarily tax the network, but most client/server messaging systems will store a replica of the message and its attachments on both the client and the server. Peer-to-peer file transfer can minimize network traffic while eliminating redundant storage. This is not to say that peer file transfer will displace email, but rather that the appropriate and judicious use of peer computing will naturally lead to a reduction in server-based storage requirements.

**Speed.** Some of the advantages of peer computing mentioned above – creating a shared context among a group, having the flexibility to add function on an as needed basis, and making connections without having to go through a centralized resource – naturally combine to speed purposeful action. In a peer computing environment, users can more easily gather people and content together to make more informed decisions more quickly.

## **Peer Computing: Tools, Platform, Infrastructure**

The pattern of peer-to-peer adoption is reminiscent of other significant shifts or inflection points in the history of computing. This pattern has distinctive phases: early adoption of innovative tools by consumers and by ‘change agents’ in businesses; the broad adoption of a more complete platform for general purpose business use; and the full-fledged integration of a robust platform as part of a strategic business technology infrastructure.

This pattern played itself out with the advent of the personal computer, which was first used mostly for games, then used by early adopters in business to automate individual or departmental functions, and then, eventually, embraced by centralized information technology departments as an essential part of a computing infrastructure. The Web went through a similar (and swifter) evolution, from the informal province of academics, to unsanctioned use as business intranets, to critical computing infrastructure supporting the full spectrum of an eBusiness.

Innovative, standalone peer computing tools such as Napster and Gnutella may indeed represent the ‘thin edge of the wedge’ that will open up into an entirely new dimension of computing. If so, we can expect the next phase of this trend to unveil a general purpose platform – like the desktop operating system

*Whatever the outcome, the simple idea at its core – easy-to-use peer-to-peer computing – already has unleashed an intellectual storm that stands to change how digital bits, the lifeblood of the New Economy, are delivered, stored, and valued.*

*Newsweek  
August 14, 2000*

or the Web server – upon which a wide variety of applications can be built and deployed. And, beyond that stage, we can expect IT management and systems integrators to begin to tie this new breed of peer-to-peer applications together with existing systems.

Nevertheless, the appearance of a promising technology on its own is not enough to bring about a major inflection point in computing. Some technologies, such as push, have appeared and faded. Other technologies, such as email, have seen broad adoption as infrastructure, but have failed to drive the business value and usage of technology to a new order of magnitude. Indeed, file sharing has already had a less than dramatic impact on computing since it first appeared on corporate networks over a decade ago. For peer computing to be adopted like the Web and personal computing before it, it needs to generate demand for commercial business solutions.

Only when there is a financial stake in its success on the part of those who deploy it (and not just on the part of suppliers of peer technologies), will peer computing impel a true inflection point. Demand for business solutions will drive innovation among tool and platform vendors, solution developers and systems integrators.

It is at just this intersection – where a peer applications platform intersects the business demand for direct, personal communications – where Groove sits.

## **Introducing Groove**

Groove is a combination of software and services that transform the Internet into a personal medium for direct communication and interaction.

- For individuals, Groove is immediately useful for direct peer-to-peer multimedia communications related to projects, meetings, events, and relationships. Groove is equally appropriate for business work, school projects, social interactions, and games and entertainment. Groove's peer-to-peer design gives casual and business users a sense of individual control and creates a context that is immediate and personally relevant.
- As a peer-to-peer solutions platform, Groove naturally extends and enhances Web-based and other centralized solutions. Groove enables close, secure, "anytime" on-line working relationships with key customers and partners, enabling rapid response to critical incidents and opportunities, and reducing time-to-decision and time-to-resolution.

### **Groove: A Peer Communications Network for Individuals**

In Groove, users make immediate and direct connections with other users for a wide variety of activities – from working on a joint project with coworkers, brainstorming, planning an event, discussing issues, sharing drafts and proposals, coordinating schedules, and just getting stuff done: to hanging out online with friends or family, sharing music and photos, talking and chatting, browsing around the Web together, and just having fun.

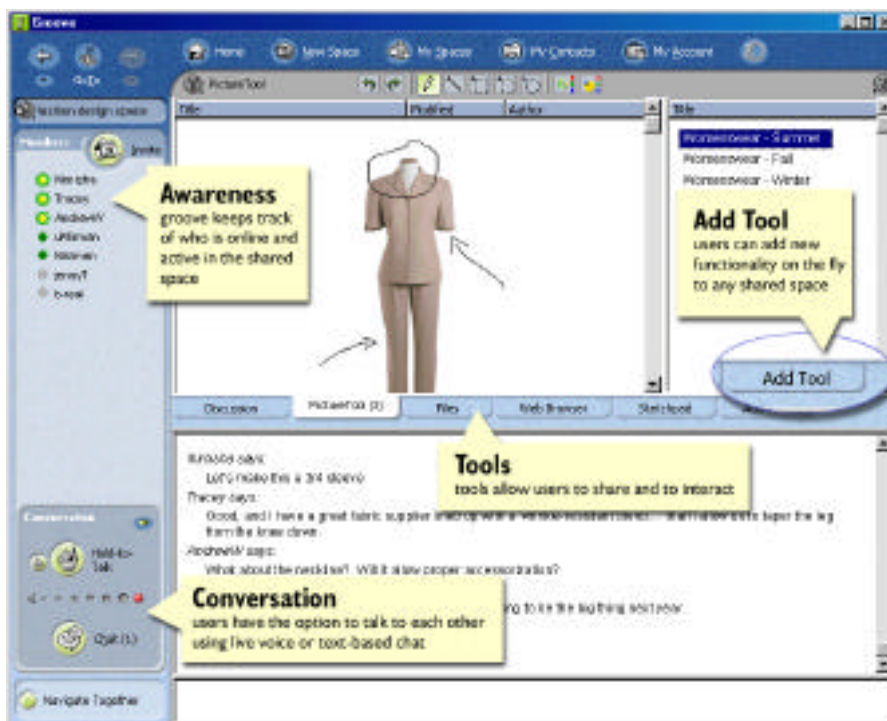
Groove is a virtual space for small group interaction. The Groove "transceiver" – the "space" in which you work – includes capabilities that lend themselves to natural and intuitive group interactivity. These include communication media (live voice over the Internet, instant messaging, text-based chat, and threaded discussion), tools for sharing (shared files, shared pictures, shared contacts), and activities (drawing, web browsing, rich text notepad). Starting with this basic environment, the user directly invites other members, adds new tools, and easily keeps track of activity and changes.

Once Groove is launched, the user quickly and easily creates a "shared space" – a secure space in which he and the people he

invites carry on business or personal conversations that last a minute, a week or a year. Each Groove shared space is stored locally, on the computers of each of the members of the shared space. When one member adds something new to the space, that change is reflected on everyone's machine – every member of the space remains completely in synch with all other members.

With Groove, all members are in complete control over the content, membership and duration of shared spaces. As in email, the only people who can see the content and participate in a Groove shared space are those that the members decide to invite. As with instant messaging, all interaction among members in a Groove shared space is immediate. And like a Web site, the content and functionality of a Groove shared space is rich. Beyond that, Groove has its own unique characteristics:

- **Flexibility and adaptability.** As a member of a shared space, whenever any user needs or wants additional functionality, he can add it from a more extensive selection of tools, developed by Groove and its partners (and by in-house corporate developers) available on groove.net. There is no waiting for a Web site manager or hosting service to add a tool for its entire user base.





- **Offline use.** If a member has disconnected from the Internet, Groove keeps that member in synch with all the other members of the shared space. Groove saves all the member's content and changes locally. When the member reconnects, all of his changes and additions are sent to all other members. In turn, all changes other members have made (from new files to new comments to a new functionality added to the space) are immediately sent to the reconnected member's space as well.
- **Never have to save.** Every change made in Groove – every keystroke, drawn line, activity, gesture – is stored right away in Groove's XML storage system. That means even if the system crashes, all content is saved. There is no "save" command in Groove.
- **Self-updating.** Groove automatically checks its own software settings and versions with the latest versions available on groove.net. When a more up to date Groove object is available, Groove automatically fetches the object and installs it without any direct user intervention. Also, Groove will check with other members with whom a user interacts, to ensure that all members are using compatible tools and functionality. If a single user has two different shared spaces that need two different versions of the same tool, Groove can manage multiple versions of the tool.
- **Multiple identities.** Many people will use Groove with different business and social groups. Groove promotes the use of different "personas." A user can have a "work" identity, an "online game" identity and a "family" identity all in one place.
- **Multiple computers.** Many people will use Groove with a desktop, a laptop, a home machine, a work machine, etc. Users do not need to manually keep each of them synchronized: Groove automatically keeps the content of all spaces on all machines in synch. In addition, if a person has spaces shared among several computers and one is lost, damaged or stolen, that person will not have lost the content and history of his or her spaces.

Business interaction in Groove includes many of the elements of ad hoc business interactions. In Groove the user finds – in a single context and in close proximity – voice conversations, shared files, joint planning documents, marked up edits and revisions, ongoing text-based questions and answers, schedules, decision histories and ad hoc problem solving – all the assets that are otherwise found in a somewhat random series of emails, telephone and conference calls, fax transmissions and overnight package deliveries.

Groove is ideally suited to the types of just-in-time, decentralized, cross-organizational business activities that characterize 21<sup>st</sup> century commerce. Over the past decade, most businesses have invested heavily in defining and automating core *business processes* – product development, supply chain management, customer relationship management, regulatory reporting, etc. In Groove, colleagues, customers and partners conduct the essential conversations, the *business practices* – negotiation, clarification, review and selection, exception handling – that surround and inform each of those business processes.

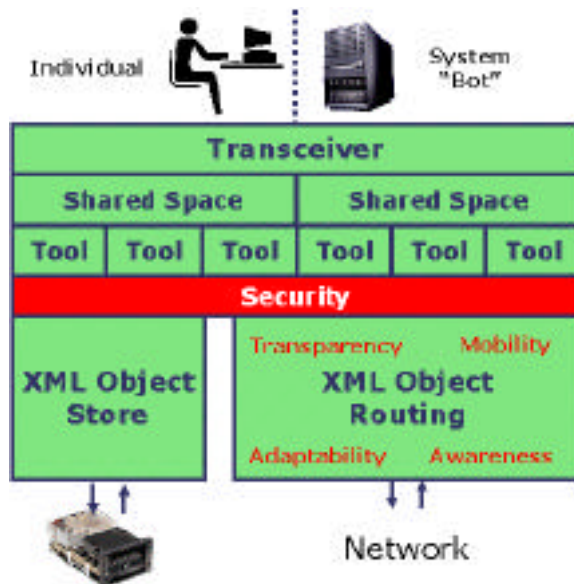
Without doubt, most business activities have benefited from process definition. Yet, in practice, there are always the unpredictable and unique circumstances that are addressed and resolved in an ad hoc fashion by individuals. For most businesses, how quickly and how well they adapt to such changing conditions, challenges and opportunities determines how well they perform. From its inception, Groove was specifically designed to make these ad hoc interactions effective and efficient.

## Groove: A Peer-to-Peer Platform for Business Solutions

So far, we have talked about Groove's suitability as peer computing tool for interactivity among business (and non-business) individuals. From the perspective of an eBusiness or a solution developer, Groove represents more than a tool – it is a platform upon which a wide variety of business solutions can be developed and deployed.

Groove's peer-to-peer architecture extends the benefits discussed earlier – personal control, flexibility, context, security, cost effectiveness, and speed – to all Groove solutions. As a platform, Groove includes the ability to integrate "edge-based" functionality with centralized systems and business processes. It easily extends the reach of those processes outside the corporate firewall to external partners and customers. And it provides users themselves with the ability to add functionality and other members on the fly.

Web developers and systems integrators will find Groove not only useful for its application functionality, but also for its full set of peer platform services. These peer services allow for full personal control and manage all traffic that flows among peers in a shared space. Below, we briefly describe the Groove platform architecture.



The portion of Groove that is visible to the end user includes

the Groove transceiver, the user's variety of shared spaces with differing sets of members, and the collection of tools available within each shared space. The underlying peer services include security, local XML object storage, and XML object routing services.

**Security.** All shared spaces in Groove are fully private: only those members specifically invited by other members can see or create content. Groove's peer-based authentication and end-to-end encryption ensures the confidentiality and integrity of all content and activity – both on the disk and over the network. What's more, Groove's security is flexible: it supports "whispering" between two (or more) members so that private conversations can take place within the context of a larger group. This is particularly useful when a shared space includes members from more than one company.

**XML Object Store.** All content and activities are stored locally on the user's device in Groove's XML object store. By using an object store instead of an XML file system, Groove improves system performance. In addition, Groove compresses the XML content for efficiency.

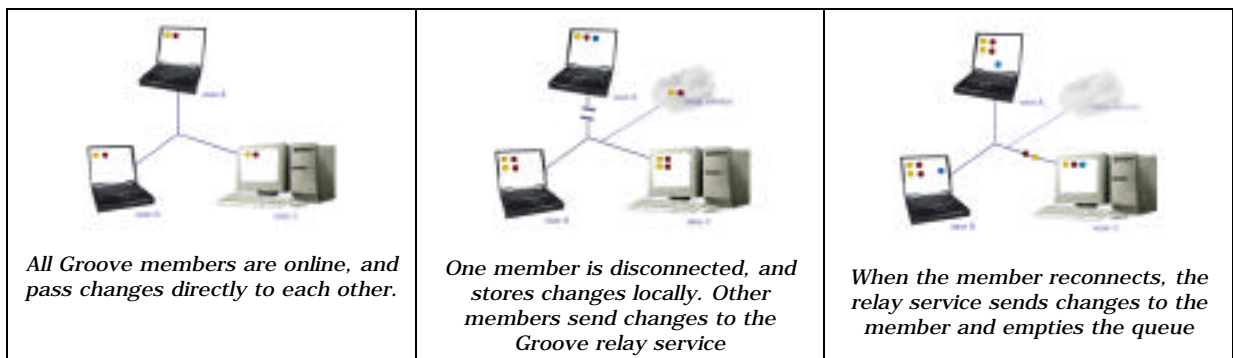
**XML Object Routing.** At the same time that Groove stores content and activities locally, it also passes all changes to the XML object routing service. The XML object routing service passes the encrypted content to the appropriate end points on the peer network. This service supports the following communications capabilities.

- **Transparency.** Groove will be used by many business people to make direct connections with customers, partners, suppliers and others – some who may sit on the other side of a corporate or personal firewall. Groove's innovative peer connection services automatically and transparently cross firewalls and network address translators, so that users never need to go through special steps to set up a shared space.
- **Awareness.** Groove's awareness capability keeps track of what other devices and/or users are online at any given moment, and where that person is "located" in Groove. That is, users can tell at a glance who else is online at the moment, which shared spaces they are in, and what each person is doing in each space. For example, if there are three active members in a project planning space, all members can see that one member is in the shared document library, that another is browsing a competitive site on the public Web, and that the third is updating the

project timeline. They could also see that there are two members who are online at the moment, but not active in the space, and that the other three members are not online at this time.

- **Adaptability.** Groove will be used over a variety of networks (LANs, intranet, Internet) and connections (T1, cable modem, dial up modem, wireless). Groove's transport service checks for all members' online status, their position inside or outside a firewall or behind a network address translator, and how to most efficiently move content across the network.
- **Mobility.** Members of a shared business space will frequently be disconnected from the local network or the Internet. Groove allows disconnected members to continue to work in a shared space, propagating all changes made by all members when he or she reconnects. This keeps all members up to date even when they are not online at the same time.

Groove manages this using the Groove relay service. When a member(s) is offline, the online "sending" Groove members automatically forward their changes to a designated Groove relay service for the offline "receivers." The relay service queues the changes for later forwarding. When the offline user reconnects to the network and establishes his/her presence, the Groove relay service transmits all changes. All activity and content is delivered directly to the receiving Groove user's computer, and cleared from the Groove relay service queue.



Groove is designed for solutions at the intersection of peer-to-peer interaction and back-end business systems. Groove supports such hybrid "peer-to-Web" solutions by encapsulating content and function from server-based systems, so that

business process and business practice can be interwoven in a single shared space. Groove's security can flexibly manage who has access to back-end resources, both within and outside corporate firewalls, lending itself to reliable business-to-customer and business-to-business interaction. Groove achieves this integration using system "bots," or agents, that act as virtual members of a shared space. In this way, systems have a direct connection with the content and activities of a shared space, and can pull from or push information back into centralized systems, such as knowledge management systems, so that content, lessons learned, skills identification and other knowledge assets created or discovered in Groove can be shared across an enterprise.

It is possible to develop Groove peer-to-peer and "peer-to-Web" solutions because Groove itself is a peer-to-peer applications platform. Groove is a highly modular COM (Component Object Model) component framework specifically designed for peer collaborative applications. This architecture allows developers to create new collaborative tools and shared spaces using any COM-compliant language, and drawing upon the vast base of COM components available today from thousands of sources. As a platform, Groove extends its system-level services – local storage, security, firewall transparency, awareness, offline support, transport – to any component, tool or solution developed on top of it.

In fact, as a company, Groove Networks' strategy has been to develop and foster a complete ecosystem — individual and corporate developers, Web and eCommerce system designers, independent software vendors, systems integrators, business consultants and training services providers. Like the business environment that has grown up around the Web, this ecosystem, working in its own interest, will develop, implement and support a rich array of personal and business solutions. It is our goal that through its own collective innovative capacity, it will extend the functionality, applicability and reach of peer-to-peer technology and solutions.