Peer-To-Peer Technologies and Collaborative Work Management: The Implications of "Napster" for Document Management

B. John Masters, Jr., Information Management Consultants, Inc.

Can over eighty-four million people be wrong? Bertelsmann (BMG) says that's how many users have registered for Napster. This pioneering Peer-to-Peer (P2P) file-sharing platform stood the world of Internet collaboration on its head in a big way. Even in bankruptcy, Bertelsmann was willing to pay nearly a billion dollars for the service.

Peer-to-Peer (P2P) computing is a parallel or distributed system with at least one of its sub-systems (e.g. control, data, processing or interface) distributed over more than one computing device, and with each one of them capable of communicating directly with any other.[1] Its main aim is to optimize resource sharing on the network. As this model eliminates the distinction between publishing resources and consuming them, it's "goodbye" to client/server, and "hello" to peers.

During 2000, file-sharing peer-to-peer networks such as Napster were in vogue. Napster allowed users to exploit the worldwide reach of the Internet to share music files. However, many well known P2P systems have been doing yeoman's duty for many years. These include Domain Name System (DNS) servers, Internet mail servers, distributed caching systems or even database servers. Although the Napster program initially gained notoriety because of its music-pirating ability, the real news was the impact of its very powerful distributed search and delivery platform search. The combination of changes and improvements in PCs (including the availability of newer Internet-aware technologies) and the wildly popular Napster phenomenon, has given credibility to an entirely new type of business model based on P2P computing.

The concept of peer-to-peer opens the way for some pretty extreme possibilities since the resources that can be shared between peers can be very diverse:

- Data sharing: Imagine the richness of a universal file system (video, documents, music...) that could be shared by all the peers connected to the network.
- CPU resource sharing: Think about the SETI@Home project, in which small bits of collected data are served out to

[1] N. Drakos, "Gartner Research Note COM-12-4960," February 23,2001.

thousands of computers for processing on the individual CPUs.

- Service sharing: Each peer can offer services to other peers, such as a simulation calculation service.[2]

In addition, P2P is not restricted to computers. It can apply to devices of all different types and sizes. Expect to see the convergence of digital and wireless technologies further increase the dissemination and distribution of information processing activities.

COLLABORATION

According to Suneeth Nayak, Chief Technology Officer of IMC, Inc., "The challenge is one of getting the right information to the right people at the right time... so effective decisions can allow organizations to succeed. Workflow, in its many forms, is part of the solution to this challenge, and P2P technologies can facilitate workflow in the new 'boundless' organizations and amongst the ad hoc project teams."

Workflow generally pushes filtered (and sometimes augmented) information to the next participant in a defined process.

According to IMC's Nayak, "A primary issue for any workflow management system is resource location. The whole concept of 'users and roles' exists to map people to the operations they are capable of performing to advance a business objective. Thus, workflow clients construct elaborate 'inboxes' that deliver work to a 'qualified' user."

In contrast to workflow, work management is where all participants have equal access capabilities and the desired resource is available to any person with the skills to service the request. "Information is free and ubiquitous. This is why Road Runner organizations capitalize on openness, for they know it's not the information and data that is crucial but what smart, creative people do with it."[3]

Work management systems can better and more timely support changes and exceptions to planned business processes. Disintermediation, demassification, and disaggregation have become the watchwords of cyberspace. New technologies are apparently breaking collectives down into individual units. It is predicted that any form of coherence and coordination beyond the individual will be the result of self-organizing systems.[4]

Collaborative work management systems enable the bi-directional exchange of "value" by two or more parties working together over

---

[2] Nicolas Farges and Habib Guergachi, "P2P and its Impact on the Enterprise," Intranet Journal, September 26, 2001.

[3] Chip R. Bell & Oren Harari, Beep! Beep! *Competing In the Age Of The Road Runner*, Warner Books, p.16.

[4] See, for example, George Gilder, Life After Television (New York, NY: W.W. Norton, 1994) for disintermediation; Alvin Toffler, The Third Wave (New York, NY: Morrow, 1980) for demassification; Nicholas Negroponte, Being Digital (New York, NY: Alfred A. Knopf, 1996) for disaggregation.

time. It usually transcends two or more organizations and accomplishes two primary goals: 1. facilitates agreement; and, 2. enacts the decision. Additionally, the systems must maintain a history of the process and facilitate post-agreement changes.

Peer-to-peer collaboration applications can be used for real-time meetings and communications and secure file sharing in ad hoc groups. Business groups can form and dissolve self-organized "webs" for collaboration on projects. Peer-to-peer collaboration can also be used to speed the development of new products and to decrease the cost and time involved in developing manufactured products by providing incremental increases in the speed of sharing information and ideas (collaboration), and organizing that collaboration into a rational and effective process.

ORGANIZATIONAL BENEFITS

There are few more dramatic impacts on today's businesses than the breakneck speed of technological change. Where decision cycles were measured in days or weeks, we now expect minutes, or even seconds; where people and businesses formed relationships for years or a lifetime, today the connections change constantly; where once we were content to make decisions based on the limited information available, technology's ability to process vast amounts of data have whetted our appetite for better and more information.

Today's world is not just about speed, it's also about anticipation, responsiveness, imagination, and most of all, agility. Speed is about pace; agility is about nimbleness.

In a world where linear speed is a commodity, any organization can use technologies that compress time. The key is this: What do you *do* with those technologies? What can you create that is different, unique, or special? In a world where speed is a commodity, being fast equals no more than surviving. The key to thriving is no longer just the ability to work fast, but *what* you are able to do fast.[5]

Unfortunately, all this collecting, indexing, and storing, has buried us under an avalanche of "data." In an average day, a typical worker must filter through 220 discrete messages. These may be e-mail, fax, voice mail, phone, letters, or even instant messages. Eighty percent more feature films are released today than 10 years ago. One estimate puts the number of publicly available Web pages at 2.1 billion, with 7.3 million new pages added every day, and we're expected to manage this information explosion in real time. According to the GartnerGroup, the average white collar professional spends 10 to 40 percent of every workday searching or waiting for information. The challenge that exists is to then discern what is valuable. Even our nomenclature recognizes this need. What was once a "data processing" (DP) department is more likely to be described today as an "Information Technology (IT) department" as new technologies

[5] Bell & Harari, P. 23.

are integrated into our organizations and lives. DP Managers begat IT Managers, who begat Chief Information Officers, and as we enter the 21st century, organizations have Chief Knowledge Officers and white collar workers are "knowledge workers."

Organizing knowledge across hybrid communities is the essential activity of organizational management. It is also difficult, though why is not often appreciated. Certainly, most managers will acknowledge that getting knowledge to move around organizations can be difficult. In general, however, discussions of such problems are reduced to issues of information flow. If, as the saying goes, organizations don't always know what they know, the problem lies with the organization: given the opportunity, information appears to flow readily. Hence the belief that technology, which can shift information efficiently, can render organizations, which shift it inefficiently, obsolete.[6]

The distribution of knowledge in an organization, or in society as a whole, reflects the social division of labor. As Adam Smith insightfully explained, such division of labor is a great source of dynamism and efficiency. The specialized tasks undertaken by communities of practice develop corresponding particular, local, and highly specialized knowledge within the community. [7]

In other words, loan-processing specialists know best how to process loans, while records management specialists know best how to manage records. Trying to move the knowledge without the practice involves moving the know-what without the know-how. Due to its social origins, knowledge moves differently within communities than it does between them. Within communities, knowledge is continuously embedded in practice and thus circulates easily. Members of a community implicitly share a sense of practice and standards for judgment, supporting the spread of knowledge. Without this shared sense of context, the community disintegrates.[8]

Business processes are made up of these embedded practices. Ideally, processes should allow groups, through negotiation, to align themselves with one another and with the organization as a whole to accomplish the business purposes of the organization. Business processes can enable productive cross-boundary relations by establishing a shared interpretation among different groups within an organization and across organizations. "In the right circumstances, the interlocking practices that result from such negotiations should cohere both with one another and with the overall strategy of the company. The processes provide some structure, the negotiations

---

[6] John Seely Brown and Paul Duguid, "Organizing Knowledge," University of California, 1998.

[7] Eric Von Hippel, "`Sticky Information' and the Locus of Problem Solving: Implications for Innovation," Management Science, 40 (1994): 429-439.

[8] John Seely Brown and Paul Duguid, "Organizing Knowledge," University of California, 1998.

provide room for improvisation and accommodation, and the two together can result in coordinated, loosely coupled, but systemic behavior."[9]

Traditional workflow systems have worked best in the organizational models of the recent past. The client/server model these organizations deployed lent itself to automated workflow based on precise business rules. Computing systems have largely reflected the strictly organized, rules-based hierarchical organizations that arose as a result of the need to perform decision-making in an environment with partially complete, partially correct and untimely information. This reflected the centralization of responsibility for risk management. An obvious extreme example was Soviet central economic planning. In this structure, all decision-making was done from a central committee with limited input from the "field" organizations that were impacted by the decisions. The controls were very rigid and discouraged entrepreneurial thinking from lower-level workers.

Today's organizations are more apt to flat and flexible in their organizational model. This has an impact on the types of information processing technologies that will be deployed. In their book, *Beep! Beep! Competing In The Age of The Road Runner,* Chip Bell and Oren Harari tell us:

Successful post-year 2000 organizations of two people or 200,000 are interlocked webs of alliances working anytime, anywhere to add new value. They are collaborative confederations of people with a common purpose: consolidating minds and energy to create something new. These alliances (often temporary) are confederations of equals, inside and outside—with permeability, and the ability to cross boundaries. A confederation means: 'being united in an alliance.' The connotation is: friend, companion, associate, accomplice, accessory, and ally.[10]

Knowledge-based repositories maintain information on three tiers. In going up each level from data to information to knowledge, facts and context are added to help understand the level below. Data forms the initial level in a digital archive, and includes the digital objects, which may themselves be very complex, plus the syntax of where records start and stop. The next level is information, which adds the tagged data, or metadata, to the data. Information includes

[9] For the notion of "loosely coupled" systems, see Karl B. Weick, "Organizational Culture as a Source of High Reliability," California Management Review, 29/2 (Winter 1987): 112-127; J. Douglas Orton and Karl E. Weick," Loosely Coupled Systems: A Reconceptualization," Academy of Management Review, 15/2 (April 1990): 203-223.

[10] Chip R. Bell & Oren Harari, Beep! Beep! *Competing In The Age Of The Road Runner*, Warner Books, P. 29.

the attributes or fields in the collection, plus an explanation of what they mean.[11]

Finally, the knowledge level includes relationships between metadata elements. This includes implied knowledge, which is information plus context and rules not necessarily contained or implicit in the data itself, which may come from sources external to the data.[12]

While a knowledge-based repository contextually linked to work management applications is not sufficient to deal with all exceptions, it can help facilitate re-use of best practices or identify the resources (e.g., colleague) that can help in a specific process situation.[13]

Without the rigid organizational structures and business rules of hierarchical workflows, advanced and flexible work allocation methods will use real-time data on the availability of resources to complete the process. A collaborative work management application could use P2P technologies to link to information dynamically and adapt work assignment decisions using context-sensitive allocation rules.

The dominant work allocation method in today's workflows is based on routing work items on predefined roles/groups instead of individuals, and workload balancing. This approach is not sufficient to address the requirement of a dynamic business environment where individuals play multiple roles. Ad hoc extension (or reduction) of tasks assigned to roles/groups, user-defined allocation rules, reservation or voting mechanisms for users/work items need to be added. Advanced work allocation methods will also be important as conflicting requirements on the usually scarce skilled resources surface in growing organizations.[14]

TECHNOLOGY IN THE P2P ENVIRONMENT

A number of new technologies enable P2P processes: Faster processors allow “clients” to acts as “servers;” the rollout of broadband allows for “always-on” connections and supports spontaneous messaging; cheaper storage allows for distributed caches; fungible naming provided by Virtual Name Spaces (AOL Instant Messaging, for example) allows an alias to be dynamically associated with an IP address.

---

[11] Regan Moore & Arcot Rajasekar, Persistent Digital Archives: A Knowledge Based Approach, NPACI & SDSC Online, Vol. 4, Issue 25, Dec. 2000.

[12] A. Gupta, B. Ludäscher, M. E. Martone, Knowledge-Based Integration of Neuroscience Data Sources , 12th Intl. Conference on Scientific and Statistical Database Management (SSDBM), Berlin, Germany, IEEE Computer Society, July, 2000.

[13] Regine Casonato, “Introducing Adaptive Work Management Systems,” Gartner Research Note, February 25, 1999.

[14] Casonato.

Extensible Markup Language (XML) will become a key enabling technology of P2P-based work process management systems. Besides being a major search enabler, XML makes content portable. It can permit the encoding of "rights" into the content, and be used to "carry" the required process metadata. .As Kahlil Gibran wrote in *The Prophet*, "A little knowledge that acts is worth infinitely more than much knowledge that is idle."

Available technologies can:

- Optimize human involvement
- Automate recurrent tasks
- Minimize interaction related delays
- Optimize concurrency

All these capabilities act to compress time and increase agility.

Workflow has been an important part of many information management strategies for some years now. The workflow systems that have developed during this period have been reflective of the hierarchical, rules-based organizations and the technology infrastructures that support them.

Peer-to-peer tools can lessen the need for corporate IT to expand some of their services, such as Web servers, backup storage, and replacing outdated documents. Peer-to-peer computing also has the potential to allow a certain amount of network traffic to move from the corporate processing resources to less expensive infrastructure, such as switches, hubs, and routers.

Many organizations have stretched resources for the servers to the limits, while a network of underutilized client PCs sit idly by. Peer-to-peer computing can extend computing throughout the Internet, allowing every computer to serve as a server to those computers around it. This represents an enormous untapped resource. Large companies might be able to utilize their client layer in order to offer as much as 10 terabytes of spare storage (2,000 clients x 5GB/client) with trillions of operations per second of spare processing power available for intense calculations performed over the network without placing additional strain on the backbone.[15]

Cycle sharing allows workstations on the network to access the computing resources of underutilized machines. Design teams that require massive computing resources can leverage the machines from other groups who are not currently using their machines for heavy computations. This could reduce the time to market for the development and manufacturing of new products.

P2P technologies integrated into work management systems will provide knowledge-based repositories containing information on similar cases, participants, best practices, and representative user skill profiles for potential contributors.

---

[15] Serena Lambiase, "Peer-To-Peer Computing Technologies: An Introduction," GartnerGroup Research Note No. DPRO-97205April 5, 2001.

REQUIREMENTS

What is the required services infrastructure for peer-to-peer architectures?

A peer must *publish* a resource (allowing it to be located and accessed by other peers) with sufficient precision for the other peers to be able to broadcast their needs and receive meaningful responses.

Before a resource can be "used" it must located on the network. This can be difficult due to the ubiquity of resources and the unpredictability of their connectivity. Some real-time directory mechanism, such as Virtual Name Spaces, is essential.

Once located, the resource must be invoked. This will require some level of standardization since a peer is not necessarily familiar with the workings of the resource that it is invoking. One protocol that will aid this standardization is Simple Object Access Protocol (SOAP). SOAP lets one application invoke a remote procedure call (RPC) on another application or pass an object to a remote location using an XML message and the Internet.

SOAP satisfies the growing need for business partners to exchange structured data over the Web independently of each other's underlying application platform. It is designed to let organizations publish data and services over the Web as easily as they can publish HTML pages. As such, it functions as a wire protocol to connect multiple Web portals, each of which might use an information server, object broker, or other facilities to integrate and process the information.[16] Services such as security management (authentication, non-repudiation, confidentiality), transactional unity, availability (by failure recovery), or performance (by balancing loads across different peers) may be required. Like AOL and MSN Instant Messaging, many of these will be hosted or facilitated services.

These infrastructure services must operate on a network in which the connectivity of each peer is unpredictable. It is clear that these infrastructure services will need to be standardized in order to minimize problems of deployment for peers. GartnerGroup calls this the "crunchy edge" of peer-to-peer processing.

Several companies are currently developing P2P collaboration applications focused on the investment management industry. They claim their more simplified and automated workflow process will allow participants in the financial services community to connect directly with investment managers without a B2B exchange; such disintermediation could have significant impact on B2B exchanges.

Virtual Name Spaces allow almost any computing device to act as a server. This can enable more ad hoc workflows for project-specific activities. The Zaplet Appmail Suite is an enterprise technology plat-

[16] Deborah Hess, "SOAP, UDDI, and WSDL: An Introduction," GartnerGroup Research Note No. DPRO-92621, May 9, 2001.

form that claims to improve collaboration within corporations and among customers, partners, and suppliers.

The platform enables companies to create and host a kind of collaborative productivity application called “Zaplet Appmail.” Zaplet Appmail brings together the people and information needed to contribute, make decisions, and take action on collaborative business processes or projects. Zaplet Appmail can be created by anyone within the enterprise—not only technical staff—and deployed as easily as sending an email. It can then be accessed via the inbox or Web browser.

RISKS

Peer-to-peer applications enable networked access to resources previously available to those with physical access; this can lead to security problems. Steve Gibson, a longtime high-level PC programmer and a former antivirus code programmer, has spent more time than most of us studying and worrying about computer security issues. Gibson says, "Today's typical user PC is not generally serving any data. Most users are just sitting on the Net, with no ports open, not accepting connections from random outside machines. But this changes in a big way with P2P! Now you're in a situation where your machine will accept any incoming TCP connection—and that changes a lot! There are going to be many ports open that weren't before, [so hackers] will get in."

Of course he's right. We all remember our alarm on seeing the intrusion attempts on our home computers after installing our first firewall software. The potential openness of P2P may make it even easier to slip a Trojan horse into our systems, or tools to launch a denial-of-service attack later—or just to run unauthorized commands on someone else's machine.

Intel has developed security software code that other companies can use when developing P2P applications. The Peer-to-Peer Trusted Library (PtPTL) includes full API documentation and provides support for digital certificates, peer authentication, secure storage, public key encryption, digital signatures, and symmetric key encryption. The library also provides simple support for networking and some operating system primitives. The API is freely available, and allows developers to add the element of “trust” to their peer-to-peer applications.

Sun Microsystems Inc. is developing a Web-based programming language called Juxtapose for use by companies looking to build distributed peer-to-peer computing applications. There are four mechanisms planned for JXTA: the ability to connect peers, logically group them together, monitor and control what they do, and add a security layer.

Since P2P computing is not centralized, managing P2P within a corporation can be problematic. While federated search (where each node reports local results and forwards the request to all known

peers) is very powerful, the lack of centralization can pose a problem for search engines. Reliability is also a concern of many IT managers. What happens in a distributed computing application if certain PCs that are supposed to be cooperating (sharing storage or processing capacity) are unavailable?

Most enterprises do not have enough Internet bandwidth to sustain the flood of large file transfers that could result from use of P2P programs by their employees.

The still nascent peer-to-peer industry needs to develop common protocols and improve scalability, security, and management, and standards for interoperability. The lack of standards in P2P technology is not unusual for an embryonic process, but it does create concerns for IT managers. This is particularly true for companies that are decentralized. What happens if two different units invest in two different P2P platforms that are proprietary and not interoperable? What about a single division that uses several different P2P programs that require separate agents because they are not compatible?[17] Problems similar to those are already faced by corporate information management professionals, but with an added degree of complexity.

Finally, P2P implements processes "bottom-up," resulting in loss of managerial control. It creates a form of "functional anarchy." In today's business environment, this can be both frightening and enabling. Thomas Jefferson wrote, "I know of no safe repository of the ultimate powers of society but the people themselves; and if we think them not enlightened enough to exercise their control with a wholesome discretion, the remedy is not to take it from them, but to inform their discretion."

THE FUTURE OF P2P

New services being developed by several companies will allow users to synchronize messages from all their devices and retrieve them from a centralized Web application (e.g., http://www.fusionone.com). This new way of delivering data utilizes the Internet itself as an "operating system," thereby transparently managing problems created by conflicting formats and allowing users to store data safely on the Internet and access it from any personal computer, cell phone, or Web appliance.

Microsoft is expected to aid the creation of such online applications through its .NET initiative, which they claim will enable every developer, business, and consumer to benefit from the new Internet devices and programmable Web services that characterize the Next Generation Internet.

[17] Stan Schatt, "Market Overview: Peer-to-Peer Technology Landscape Complex, but Some Winning Niches Emerging", Giga Information Group, September 20, 2001.

Peer-to-peer technology has evolved away from the centralized directory of Napster, but there is reason to believe that the B2B applications will return toward more centralized control, as seen in the approach taken by Groove Networks. The reason for this is to ensure a higher level of security and a higher level of control. There are serious issues right now that need to be resolved to achieve true peer-to-peer distributed computing.[18]

Groove Networks, Inc. software and services enable members of the enterprise to connect quickly with customers, partners, and suppliers in a secure interactive environment to share information and get things done. A unique peer computing platform, Groove erases technical and organizational boundaries, bringing together the people, information, and tools needed to speed decision-making, solve problems, and reduce time-to-market for new goods and services.

Groove runs on each user's computer and handles all communication functions. Most activity in Groove takes place in shared spaces—virtual interactive workspaces open only to invited members. Shared space members can work together, online and in real-time, or independently, offline, at times of their own choosing. All shared space content is automatically saved and synchronized on users' machines, making ongoing meetings and project management a breeze. According to Groove, any business process that requires dynamic communication and collaboration can be enhanced and extended.

In Groove, all data is automatically encrypted, both on-disk and as it moves over the network. Shared spaces—accessible only by authenticated members using pass-phrase protected accounts—are protected from rogue components by IT controls included in the platform.

Another type of peer-to-peer model is Web services. A Web service is a process published by an enterprise that may be connected to the firm's existing IT system. The notion of clients and servers disappears, since a consumer of a Web service may also be the publisher of another Web service. What makes Web services special is the fact that they are based on a triptych of standards that makes interoperability between diverse systems a possibility. As examples, the Web uses two protocols to exchange information: HTTP, a real-time TCP protocol used by browsers and servers, and SMTP, a store-and-forward protocol for e-mail. XML documents can pass through firewalls, are easily translated to internal formats, and allow applications and servers to process information without exposing their internal workings to the Web.

It is not uncommon during the early stages of each wave of technology adoption for architectural discipline to be abandoned in a rush

[18] Stan Schatt and Colin Rankine, "Peer-to-Peer Networking: Lots of Turmoil, but Are There Practical Applications for Enterprise Clients?" Giga Information Group, February 2, 2001.

to implement. During two earlier shifts in particular—client/server and Internet applications—many "instant legacy" applications were created, i.e., applications without abstraction of business services, layering of user interface, business logic and data logic, integration with organization wide infrastructure for security and other services, or even, in some cases, Y2K compliance.

Such applications became maintenance headaches for the remainder of their useful lives. It will not be surprising to see the same phenomenon in the early adoption of Web Services, but the fact that some people are asking whether Web Services actually eliminate the need for enterprise architecture (at least, in the sense of centralized enterprise architecture) reflects how large a potential change Web Services represent. However, although Web Services may require enterprise architects to radically rethink their role within the organization, using Web Services as an excuse to abandon enterprise architecture altogether would significantly weaken an enterprise's ability to leverage and respond to future waves. Indeed, the lessons learned from those that have been most agile in leveraging past waves suggests a high correlation of this agility to those enterprises with the most effective architectural strategies.[19]

FUTURE PEER-TO-PEER WORKFLOWS

In future P2P workflow models, the work item and associated data will be packaged into a virtual envelope indicating the type of resource required. When a resource receives such an envelope, it can check its own profile for a skills match and/or forward the item to the next set of resources (perhaps based on the expected time to completion for the item; if the local resource can service it within some threshold, it will queue but not forward it; if the service time exceeds the threshold, it will queue and forward it; and if it lacks the appropriate skills, it will simply forward it).

In work management systems, "who does what" is not necessarily predefined, and may be more dynamic in collaborative environments. General George S. Patton once said, "Don't tell people how to do something... Tell them what to do, and let them surprise you with their results." Work management systems of the future will drastically improve the flexibility of work allocation methods, replacing the traditional client/server model with which we are familiar.

Design the workflow to readily allow for changes to the business process. Foresight and planning are needed. Maximizing the business rules (or flow control logic) in the work management system may be limited by the actual capabilities of the chosen work management (i.e., Is there a dedicated workflow engine?) or may negatively impact system performance. The importance of allowing the

[19] Carl Zetie and Mike Gilpin, "Web Services vs. Enterprise Architecture: Paradigms Shift, Agile Architecture Is Forever, "Giga Information Group, September 25, 2001.

user organization to change the workflow independently of the IS organization will be one of the primary drivers of this approach. Maximizing business logic in the application system can address performance issues, but will then require programming changes to modify the business process. An organization will have to determine the trade-offs appropriate to its own business priorities and operational environment.[20]

Think of your company as a fire department. It cannot predict where the next fire will take place, so it has to shape a flexible and efficient team that responds well to unanticipated events, no matter how extreme—Andy Grove, Former Chairman, Intel

By extending business processes through work management, collaboration and peer-to-peer technologies, organizations can better manage and automate critical business processes across an extended enterprise. This should lower the cost of delivery for more complex processes.

CONCLUSIONS

Eighty-four million can't be wrong. Peer to peer is definitely here, it has a future, and it does not encroach on traditional heavy or light client models. Although the idea of a generic peer-to-peer architecture might seem to be a theoretician's ideal that cannot be applied in the real world, it can in fact be entirely feasible for some types of applications. Success stories like Napster or Gnutella are proof of the viability of the peer-to-peer model for certain problems.[21]

The approach to adopt for peer-to-peer architectures should be utterly pragmatic; a good example of this is the Web services standard that will enable enterprise applications to be extended by letting them communicate using the P2P model.

The value of a virtual company does not come from its products or services, but from its knowledge, its web of alliances, and its capacity to transform know-how into results.[22]

In his book, *The Minding Organization*, UCLA Professor Moshe Rubinstein writes about the leadership required in this new chaotic age: "To create an adaptive, innovative, problem-solving organization requires a new form of leadership, the ability to provide guidance in the search for purpose and goals. Classical leadership—implementing stipulated goals—is committed to the status quo."[23]

Tomorrow's organizations will be loose confederations of people working anytime and anywhere. This will require supporting tech-

---

[20] Regina Casonato, "Six Steps To Deploying Work Management Successfully," GartnerGroup Research Note TU-IDOM-384, July 7, 1997

[21] Farges and Guergachi.

[22] Bell & Harari, P. 122.

[23] Moshe F. Rubinstein, PhD and Iris R. Firstenberg, PhD, The Minding Organization: Bring The Future To The Present and Turn Creative Ideas Into Business Solutions, 1999, John Wiley & Sons, Inc. P. 147.

nologies that are as flexible and agile as the organizations themselves. These will be technologies capable of working outside the traditional hierarchical organizational structure of centralized command and control.

While there are significant challenges to implementing workflow in a P2P environment, the results can have significant value to an organization. It will be necessary to once again rethink our business processes and how we manage them, and to accept the fact that they will always remain in a state of flux. Workflow systems that exploit the strengths of P2P technologies will do much to empower these new organizations and provide a competitive edge.