

Organisational impact of Peer-to-Peer technology

Introducing Peer-to-Peer technology (P2P) as a justifiable business and educational technology usually stirs memories of Napster, Kazaa, Gnutella etc and the impact that these programs had, and continue to have, on the music recording industry. Information Technology managers are justifiably cautious about P2P because of the disruptive nature of these applications to security, bandwidth, intellectual property and content management. Additionally, some people understand the value and advantage of P2P technology, but have based a whole architecture on centralised or web-based models and are reluctant to change the status quo. Nevertheless, P2P systems can address these issues while providing more efficient work processes for workers and more educationally friendly online environments for teachers and students. It is just a case of understanding P2P and deciding whether this is a worthwhile technology to use.

The business need

Drivers of change

Consider two types of change influencing educational institutions¹ today. Firstly, there are those changes that affect all organisations regardless of their industry. For example, organisations have increasingly developed into complex and hectic work environments that operate with less clearly delineated lines of authority and responsibility and with less rigorous management structures. Such organisations are being asked to manage more and more issues, deliverables, commitments, and relationships at both an organisational and an individual employee level.

Secondly, there are those changes that alter the very nature of education and learning.

- **The convergence of Internet Communication Technologies (ICT)** and the application of these technologies in an educational – learning context.
- **The changing nature of students and learning.** Learners are wanting to learn without disrupting work and family commitments. Learning must be relevant to and integrated within the work context. Students are no longer just an 18-25 year-old cohort. Rather, the student population is ageing, taking courses part-time, studying off-campus, technologically enabled and Internet friendly. It expects high quality customer service, not just student administration. Learning is becoming a lifelong activity with each student having individual learning wants and needs that they are seeking to fulfil.

¹ The term 'institutions' is used in this report as a generic term that refers to vocational educational & training colleges or organisations, and universities.

- **The changing nature of work and increasing volume of knowledge** now requires that workers be able to learn and adapt as they work. It is no longer sufficient to complete a course, pass an exam, and then be expected to apply knowledge and skills on the job, as organisations are already seeking technical systems and business processes that ‘embed’ learning with knowledge management and performance support techniques. The completion of a ‘course of study’ out-of-context no longer satisfies the needs of a knowledge society.
- **Changing nature of teaching and learning.** Although the teacher as lecturer and ‘distributor of knowledge’ worked in the industrial era of mass production, limited knowledge and defined occupations and skills, the industrial model of education is no longer an appropriate context in which to teach and learn. Internet Communication Technologies (ICT) have provided a wealth of opportunities for teachers and learners to redefine what they do in order to better meet the present needs of the modern workplace.

Why collaborative software is the way to go².

Many organisations have chosen to embrace collaboration and business process automation technologies. At the macro level, where such business processes can be understood and mapped out in a top-down fashion, these technologies (e.g. workflow, notes, web, databases) have been effective. However, in more complex organisations where hierarchy has diminished, it is clear that much - if not most - of the process of ‘how the organisation operates’ defies formalisation, and thus defies automation attempts through any means other than perhaps electronic mail.

Why does this occur? Because in the end it all comes down to **people**. In order to solve problems, serve customers and address challenges, employees need to interact with one another. Within an organisation people ‘network’ with one another, gaining value and effectiveness through their relationships, leveraging each other’s knowledge, ability and experience. Unfortunately this networking often occurs in an ad hoc manner.

The breadth and effectiveness of an organisation's culture and ‘interpersonal value network’ can, in many ways, shape its real operational effectiveness. In fact, by focusing on improving the effectiveness of interpersonal relationships and collaboration, a business one can often gain as much – if not more – long-term improvement than can be obtained through macro business process reengineering. For while there are limits to the amount of efficiency that can be milked out of a business process, the incubation of ideas and the management of knowledge and relationships directed toward corporate advancement can yield continued vitality and potentially sustainable growth. A vital key to achieving a dynamic organisation is **the individual** and the nature of their collaboration with other **individuals**. While continual improvement of business processes, and large-scale communication and collaboration systems are necessary for organisational efficiency and survival, sustainable growth and

² Based on the postings contained in the weblog of Ray Ozzie technologies. Accessed 10th of June 9 2003 at <http://www.ozzie.net/blog/stories/2002/10/01/marketOpportunity.html>

overall operational effectiveness are best achieved through focusing on the enablement of the individuals to network within an organisation.

If enablement and empowerment of the collaborating individual is the focus, then it is useful to examine how people collaborate. Exactly how do individuals utilise objects, tools, environments, and technology to collaborate effectively with one another when in supposedly 'optimal' contexts - that is, in the same place, at the same time, working toward a common goal? In meetings, and during brainstorming, individuals communicate with one another using whatever tools and environments are at their disposal to most effectively get their point across. Telephones, voice mail, whiteboards, pads of paper, napkins, electronic mail and faxes are just a few such examples. People on the path to innovation and discovery need to sketch, prototype and perform; they need to convey insights that only a visual, tactile or aural representation of the problem can evoke.

It appears clear that if an environment is to foster effective collaboration, it needs to be flexible enough to support a rich variety of interaction methodologies. Furthermore, collaborative tools and environments need to empower rather than handicap human expression. Take for example, the Web, an environment used by many organisations as a means of communication. The Web is inherently asymmetrical in nature, that is, nobody expects to write nearly as much as they read. The basic browser stresses navigation as opposed to editing or disposition, yet developers with an 'all I have is a hammer, so everything looks like a nail' mentality attempt to do everything - including collaboration and mail - within the otherwise clean and simple browser user interface. Electronic mail is principally a one-to-one (or at most one-to-many) communication mechanism, yet people awkwardly carry on many-to-many group discussions through the abuse of distribution lists and through custom mailbox views sorted by conversation, by person, and so on.

Many technologies are available for effective group information sharing, yet the management of such technologies is almost always centrally controlled. While most ad hoc sharing occurs through email due to it being easy to use and understand, the ability to communicate in ways other than words is inhibited: For example, email does not allow people to do a simple sketch, to annotate in the margins or accurately convey an emotion. Because of these constraints, people resort to use of smileys, angle brackets and underscores. Given these limitations, it is perhaps not surprising that fax and voicemail continue to be used even within an increasingly email-based business environment.

While it is acknowledged that the mechanism of email enables some degree of collaboration, if one believes that our increasingly chaotic business environment will require increasingly effective means of collaboration, it is surely arguable that one must think beyond the constraints of the current personal tools such as email. And if this is the case, it is similarly arguable that one must also think beyond the collaboration constraints of the current tools and environments. Rather, one must envision tools enabling individuals to collaborate in new and creative ways by defining a new environment with new languages for interaction - languages appropriate to the type of collaboration being undertaken. The goal must therefore be to embrace tools and environments that enhance and augment the productive relationship between collaborators.

When a small group is brought together, their most effective collaboration tool is often a whiteboard, a pad of paper, a flipchart, or a napkin. Are there not environments or products that can be developed to enhance same-time same-place collaboration? And when individuals are apart, should the flood of collaborative ideas generated by these individuals be forced through the funnel of constraining technologies? The answer to those questions lie within the domain of peer-to-peer technology.

Using Peer-to-Peer as an educational technology

Knowledge now relies on high levels of collaboration, drawing together a range of disciplines, insights and experiences. The most powerful unit in our society is a mutually supportive team of learners.

– Mark Latham³

The reasons for which P2P is used as a work collaboration technology are synergistic with its use in teaching and learning. In teaching and learning there is a need for relevant, credible and up-to-date information. Therefore, the role of web based systems, content management systems with reusable learning objects and other resource repositories are valid. However, the interaction of teaching and learning is more than just content and information dissemination. It is a dialogue, a discourse. It is story-telling and swapping experiences. It is questioning and answering, exploring and investigating. If the interaction of teaching and learning are all these things, then the relevance of P2P systems that enable such activities to be undertaken by teachers and learners alike is clear.

This does not mean that P2P is a 'cure-all' or the next 'big thing' in educational technology. Online educators know the value of using a blended learning approach, that is choosing the optimum mix of delivery strategies and teaching techniques to facilitate learning within a given learning context. Online educators also know that creating that blend is the use of appropriate technological choices, that is, the technology blend. What P2P does represent is one of a continuum of technologies available to educators to be employed when the context is appropriate.

Implementing P2P technology

Industry reviews (eg www.nwffusion.com/reviews/2003/0728Collab-BG-Excel.xls) reveal that there are a range of collaborative software platforms and applications available on the market that are both server-based and P2P. Although the web-based systems and P2P systems share many functions, there are some fundamental differences. To illustrate some of these differences, a comparison between the P2P system *Groove* (Groove Networks www.groove.net), and the server based collaborative system *Share Point* (Microsoft) is outlined below.

Conceptually, *Share Point* is a document-centric system, enhanced with real-time services for presence awareness and real-time communication. In contrast, *Groove* focuses more on the process of collaboration rather than just on the provision of collaborative documents. It assumes a decentralised shared workspace context and delivers collaboration and powerful off-line/resynchronisation capabilities by asynchronously disseminating actions

³ Mark Latham, *What did you learn today* (2001). Allen & Irwin.

within the shared context instead of replicating documents. This is certainly not a right-wrong, either-or scenario with *Share Point* and *Groove*, since they have been designed to address differing customer needs. It is perhaps more accurate to view their architectural differences as complementary rather than competitive.

There appear to be four major issues related to implementing and using P2P collaboration software. These issues are both technical and organisational in nature and are discussed here under four category headings: Security, Bandwidth, Content and People. The issues discussed within these categories represent both challenges and solutions to educational Institutes.

1. Security

Here are some common security issues related to the implementation and use of P2P:

The need to protect the confidentiality of content and the activity of users

Taking Groove as an example of P2P, we see that Groove encrypts all content on disk and over the network. The user's account and all shared space contents are automatically encrypted locally on the user's device. Users create a passphrase to encrypt their account and their shared spaces. A passphrase (a collection of words) provides for greater length and variety than a simple password, thereby making it more difficult for someone with malicious intent to attempt an educated guess at a user's secret code. In this way, should a user lose their laptop or allow someone else access to their desktop, shared space content is protected. Furthermore, even if an intruder gained access to the file system in which the spaces are listed, the contents of those spaces would remain encrypted and impossible to compromise without the user's passphrase. Not only is the user's device protected, but all content and activity within a shared space that is sent across the network to other shared space recipients is also encrypted and can therefore only be decrypted by other members of the shared space.

Authentication of users

Users establish trusted relationships via the exchange of contact information in the form of vCards. vCards include a person's public key. There are a number of convenient ways that this exchange can occur between two or more people. In the likeliest inter-company, cross-firewall scenario, one member will send an email invitation to another. For example, say I email to you an invitation to a shared space. When you accept the invitation, you automatically have accepted my vCard. Your reply sends to me your vCard in return. This peer distribution of keys ensures security without requiring centralized certificate and key management.

After you have accepted my invitation and we have authenticated each other, Groove then sends to you a set of keys and the shared space itself as well as a set of keys that allow you to decrypt the content of the space. This exchange of keys for the space itself provides an easy mechanism for managing membership for the duration of a space. Throughout the normal lifecycle of a shared space, it is likely that new members will be invited and that other members will be 'uninvited.' This presents a security challenge, since the 'uninvited' member still retains the keys to the shared space. However, whenever a member is uninvited from a shared space, Groove automatically

issues a new shared space key to all members so that all subsequent data is protected and kept private from all past members. As far as the 'uninvited' member is concerned, the 'old' shared space is still on their local device(s), and they can still review the content and history of that space. However, the past member can no longer access new content and activity.

Protection from the spread of viruses

Strict guidelines often forbid, for good reason, the random download of executable code, some of which may have bugs or viruses, or which may violate standard desktop configurations. For this reason, Groove includes administrative controls that allow Information Technology (IT) staff to 'lock down' Groove desktops so that they may only download components that have been digitally signed (eg by Groove Networks or others) and/or that are found on approved server sites. This continues to allow end users a great degree of freedom while balancing necessary IT protections against viruses.

2. Bandwidth

Information Technology managers and end users often begin their evaluation of P2P by looking at how well it reduces infrastructure costs and delivers productivity benefits. One interesting aspect of P2P is the manner in which it optimizes network bandwidth. It is possible for P2P to generate less network traffic than traditional, server-based alternatives. Groove, for example, adapts to local network constraints and provides capabilities to make the most of dial-up and other low-bandwidth connections.

In order to understand how the Groove platform measurably reduces network bandwidth utilization, it is useful to first discuss the different types of collaborative activities people engage in, the content those activities produce, and the impact each has on bandwidth consumption. On the one hand, there is shared content that is *static* in nature, that is, information that, once posted tends to remain unchanged. Policies and procedures, best practices, reports, analyses, documentation, templates and news are examples of static content. As static content is often shared with a relatively large number of people (eg, a department or an entire enterprise), the center-based or web model of content distribution serves this type of content well. On the other hand, there is shared content that is *dynamic* in nature, that is content that is shared among participants who are actually composing or editing the documents in question. Drafts or revisions of documents such as budgets, customer profiles, re-engineered business processes, analysis of new competitive threats or emerging technologies and internal and external reports are common examples of dynamic content. Because dynamic content tends to be shared among a relatively small number of participants, for the vast majority of business people email has been the medium of choice for sharing these draft documents.

Regardless of whether collaboration data is static or dynamic, users typically gain access to content through a center-based, client-server system. This includes both web-based collaboration solutions, as well as their email solution. In a centralized model, users create network traffic every time they send a file initially, as well as each time they subsequently modify and share that file. While email hardly represents the most efficient means of distributing revised documents over and over again to a team of, say, three to ten people, it is probably the most convenient means available to from an end user's

perspective. Specifically, as email does not require any set up, there is no need to ask permission of IT, or anyone else to establish access or use. Also, email fairly seamlessly crosses enterprise boundaries. However, while there appear to be these advantages from an end user perspective, from an IT perspective a flurry of 'reply-to-all' emails with large file attachments is a regular and nettlesome source of network traffic spikes.

Web-based collaboration closely follows the network usage patterns of email. Similar to email, each time a member of the team 'reads' a Web file located on a central server, there is a network utilization cost associated with it. For example, when a user opens a 2 megabyte file from the server, that amount traverses the network. When all members of the team open that file, the amount of traffic multiplies. No matter what the user does with the file, whether it's simply reading it or making small edits to the file, the cost of the file transfer is there. The more dynamic and/or useful the content, the more traffic it creates.

In contrast to these problems inherent in using web-based collaboration, P2P's architecture is specifically designed to speed and improve the interaction within small groups or teams in which content is created and/or revised and in which business decisions are made. For example, dynamic content in a Groove shared space typically goes through multiple iterations. In a shared space, users create network traffic every time they initially add a file to that space. In contrast to web-based collaboration where files need to be re-downloaded each time they are re-opened, with P2P the content is, after the initial download, resident on the end users' machines. Consequently there is no network penalty to be paid to re-open it or re-read it. Thus, from a network efficiency perspective, a rule of thumb is that the more collaborative the sharing in a project, the better it is to distribute content to the end points. The more that shared content is of a static, broadcast mode, the better it is to centralise it (with a few exceptions in either case).

In many ways then, P2P is a complete reversal of conventional wisdom. In the past, anecdotal information and common logic suggested that decentralised collaboration was a danger to network bandwidth. Today, a simple analysis shows that the opposite is true: decentralised collaboration is a boon to network bandwidth utilisation.

3. Content

Content management, intellectual property and copyright represent the main issues in regard to content. Management of 'administrative' and 'educational' content in Institutions is a topical issue. At one end of the spectrum, teachers' files and resources are mostly in a non-digital format, while at the other end, enterprise content management systems, learning object repositories and electronic documentation systems are employed. Using P2P need not compromise the control that Institutes have over their content. Using Groove for example, enterprise servers are able to link the shared spaces of controlled domain users with the Institute's legacy content systems.

Those with Learning Management Systems (LMS) need not convert their existing materials or extract the content for P2P. The web browser facility in Groove will still enable students to login to their web based course for resources, although it would be inefficient to then use the embedded communication features in the LMS as well as the P2P space. Early adopters of

P2P have used a blended P2P-LMS approach for course design and delivery, but have then migrated from the LMS entirely to the P2P environment, electing to put web based content on a simple website with password access.

P2P platforms, such as Groove, could provide the missing link in the current debate around Reusable Learning Objects (RLO). The concept of creating, meta-tagging and storing 'educational' digital data for sharing and reusability has always been appealing. Although the creating-tagging-storing process is now quite sophisticated, the actual application in an e-learning context is problematic. Some researchers have built engines for both students and teachers to 'assemble' RLOs in the context of their learning, mostly within some type of LMS environment. However, most of these initiatives seem to result in a clumsy linear type of information/object display that is difficult to put into the specific contextual need of the user.

In a P2P environment like Groove, it is clearly possible to link to existing RLO repositories. From there, teachers could use a search function within the P2P environment or the RLO Repository to locate and then insert the RLOs in their educational 'space'. The contextualisation of the objects is then more easily managed by the teacher and by the students because of the dynamic nature of the P2P learning environment. Also, there is no reason why a more sophisticated RLO 'builder' could not be developed in the P2P platform, that could act like a 'wizard' to assist teachers to manipulate the RLOs for context and/or easier use, or provide a more education-friendly interface than the current 'file explorer' presentation.

4. People and organisation

The successful implementation of any new technology is probably more dependent on people issues than anything else (eg staff training, organisational readiness, change management, culture of innovation, risk taking etc). Implementing P2P is no different, though perhaps the challenge here is greater with P2P being seen by some as 'disruptive technology' because of its challenge to the way people traditionally collaborate, communicate and work through enabling such to be undertaken online rather than face-to-face, the task of gaining converts to P2P will not necessarily be easy.

Illustrative scenario: small training organisation

The following scenario looks at a typical installation of P2P technology (in this case, Groove) in a small training organisation.

Hospitality Training Company (HTC) is a small training business that delivers training to the hospitality industry. HTC is comprised of a manager, administration officer and two workplace trainers. There are approximately 60 students enrolled at any one time.

Students complete hospitality traineeships by engaging in learning activities at work where they currently follow paper based workbooks. Employers are required to provide some training and supervision. Students submit assignments and record their tasks in their workbook, which is monitored by the employer and also a HTC workplace trainer. The workplace trainer visits the student at least six times during the traineeship to provide learning support and conduct workplace assessments. Although the minimum legal requirement is

three visits, the workplace trainers are unable to provide an adequate level of service in just three visits.

The HTC workplace trainer responds to student requests and queries mostly through email and telephone contact. The HTC administration officer despatches the workbooks to students, answers student questions (mostly via email) and collates all evidence and documentation from the workplace trainers either physically or via email.

HTC considered using online technologies in a blended delivery model to address the following issues:

- (a) reduce the number of face-to-face trainee visits to the minimum requirement as a cost cutting measure - yet still maintain the level of support required
- (b) use digital resources to decrease printing costs, distribution and version control
- (c) provide a means of better communication and learning support directly to the student as well as create a community of learners amongst the trainees; and
- (d) create better communications and document/report submission between the trainee, employer, HTC workplace trainer and the HTC administration officer.

The chosen blended delivery model utilised by HTC consisted of online learning, employer workplace training and supervision and three face-to-face workplace visits by workplace trainers. Because the trainees all had online access with suitable technology in their workplace, HTC selected P2P technology (Groove) for the online learning and administrative communication. HTC purchased 200 software licences to cover HTC staff, employers and students. Because this was a once-off licence fee, and the employer/student licences could be re-issued, the total cost of the licence was amortised over a 3 year period, reducing the cost of P2P per user.

Groove was installed on the three HTC office computers. The administration officer and manager had their own accounts and it was decided to allot an account to the office server, which was always online. This office server was used to distribute changes to all staff and students and was used as the backup computer (with the space being regularly saved as an archive). HTC staff completed the two hour online training with the Groove vendor to familiarise themselves with the P2P platform and completed a one-day workshop to develop their blended delivery and business model for the Hospitality Traineeship. Because HTC ran continuous enrolments in their program, it was decided to create four identical Hospitality Traineeship spaces. Each HTC workplace trainer would manage two spaces with each space having had a maximum of 15 trainees at any one time. It was deemed that this number would be sufficient to cope with a transient enrolment, whilst also providing enough critical mass to create a learning community.

The 'Hospitality Traineeship Space' was based on the ANTA 'Coach MCLE', which was downloaded off the net. HTC then modified the MCLE to include an outline of the program, a calendar of events (into which specific dates were

entered when a trainee enrolled) and the learning resources required (Word, PDF and PowerPoint files).

During the first visit by the workplace trainer to the workplace, Groove was installed on a designated workplace machine and the trainee and employer accounts were established. The space was loaded from CD to avoid any initial large space download. HTC felt that this initial face-to-face briefing was warranted to overcome any uncertainties and queries about the online component that may arise. Even though the trainee and the employer could synchronise their space with multiple computers (eg at home or in the office), it was felt that this could be suggested later after they had become more familiar with the P2P environment.

During the early stages of the program, some students and employers used telephone and emails for communication. When this happened the workplace trainer would politely re-direct them to the space for a response as well as via instant messaging. Quickly, the space became the preferred facility.

On the administrative side, the HTC administrator was invited to each space. The administrator found that communication was easier using the space as everyone had up-to-date information about schedules, assignments etc. Employers too felt that they had a reference point where they could see the status of their trainees and all the activities associated in the instance that they needed to know such information. A hospitality administrative space was created where workplace trainers were able to complete their administrative forms online. This enabled forms to be instantly viewed and retrieved by the office administration staff.

HTC workplace trainers found that the level of interaction and motivation amongst the trainees increased significantly. The trainers were pleased to see that a 'cross-pollination' of ideas between students who would normally never know of or even contact fellow students had developed. Each workplace trainer was a guest participant in the other colleague's space. This way the trainers were able to moderate assessment and share ideas more easily. Using their laptop both online and offline, the workplace trainers were also able to work from home and whilst conducting workplace visits

The HTC manager found that the P2P platform also doubled as an effective work tool for his staff. HTC staff used the Hospitality Administrative space to conduct staff meetings and to update procedures and forms. The HTC manager also found that the ability to interact with the client employers improved via use of the instant messaging facility – a facility that was used regularly.

Beyond completion of the training course, the vast majority of employers and trainees elected to acquire their own licence. This provided a convenient method for employers and trainees to maintain relationships with HTC and to undertake more training and learning.

A review of the program revealed that most of the business and educational objectives had been reached and that the Hospitality Training Company (HTC) is looking forward to expanding their services in the future.

Summary

The P2P discussion in this report has considered educational administration, issues relevant to teachers and learners alike. Some of the challenges and solutions to better implementing P2P within a collaborative learning environment have been explored.

This work has been produced with the assistance of funding provided by the Commonwealth Government through the Australian National Training Authority. Copyright for this document vests in ANTA. ANTA will allow free use of the material so long as ANTA's interest is acknowledged and the use is not for profit

The views and opinions expressed in this document are those of the authors and do not necessarily reflect the views of ANTA. ANTA does not give any warranty or accept any liability in relation to the content of this document.



© 2003 Australian National Training Authority