

Roger Tagg

University of South Australia; Roger.Tagg@unisa.edu.au
AUSTRALIA

Within the life cycles of Virtual Enterprises, there are usually periods where representatives of different organisations collaborate in creative ways to set guidelines for the operational stages of the VE. Examples of what has to be agreed or negotiated include legal contracts, financial terms, human responsibilities, joint processes to be followed, common terminology, data formats and common software. This collaboration may be a mixture of synchronous (physical or virtual meetings) and asynchronous activities. This position paper assesses the gap between what support is required and what is available in currently available tools, and suggests a number of key developments to the state of the art.

1. INTRODUCTION

This position paper addresses the issue of IT support for those periods within the operation of a Virtual Enterprise when groups of people have to develop ideas, designs and solutions, which are acceptable to all participants and which further both their joint and individual objectives.

The motivation for the work arises from a range of application needs. Some colleagues of this author are focussing on support for groups working in emergency situations, often involving the participation of different authorities such as Fire, Ambulance, Police and the Military. Others, including this author, are targeting requirements engineering, systems design and general administrative cooperative work. An additional goal has been to propose ways to overcome the epidemic of information overload.

The paper is organized as follows. First, the place of creative work in Virtual Enterprises is introduced. Then, the patterns of creativity in general are discussed. Next, the limitations and shortcomings of current facilities for supporting group creative activities are highlighted. A possible way forward is offered in two parts: a) an overall architecture; and b) specific software developments to support that architecture. A brief summary highlights some of the most critical success factors for these developments to become a reality.

2. CREATIVITY IN VIRTUAL ENTERPRISES

2.1 Characteristics of a Virtual Enterprise

In this paper the term “Virtual Enterprise” (VE) includes a range of forms of cooperation that involve the participation of more than one legal entity or individual. The common feature of all VEs is that they present a single face to the outside world of customers, suppliers, the general public or the environment. They can vary from more formal arrangements like joint ventures, through loose confederations of small enterprises, to voluntary associations of individuals. In all cases, some level of “ground rules” has to be negotiated, as well as discussions of the goal and priorities.

2.2 Virtual Enterprise Life Cycles

(Tagg, 2001) proposed 4 stages in the life of a Virtual Enterprise, namely Establishment, Business Development, Business Execution and Winding Up. In each of these, some activities involve creative work, whereas others can follow processes agreed at earlier stages. Some examples are shown in Table 1.

Life Cycle Stage	Creative activities	Other activities
Establishment	<ul style="list-style-type: none"> - Identify broad market opportunities - Develop processes and protocols for the VE 	<ul style="list-style-type: none"> - Apply joining procedures
Business Development	<ul style="list-style-type: none"> - Joint market research - Decide participants' responsibilities - Develop template business processes 	<ul style="list-style-type: none"> - Bid in competitive tenders - Evaluate benefits, costs and risks - Project management
Business Execution	<ul style="list-style-type: none"> - Negotiate individual B2B business processes 	<ul style="list-style-type: none"> - Customer Relationship Management - Accounting
Winding Up	<ul style="list-style-type: none"> - Apportion assets and intellectual property 	<ul style="list-style-type: none"> - Settle legal and financial issues

Table 1 – Examples of Creative and other Activities in Virtual Enterprise Life Cycles

When a VE reaches the Business Execution stage, periods of creativity are less dominant. Processes that have been negotiated earlier on will often be followed. Indeed, the development of good processes, which participants can be expected to follow, is more important in a VE than a single organization, since it has to compensate for the fact that individuals primarily report to different managements. However as the table shows, creativity, in the form of negotiation of individual contracts and protocols, is still present; and it reappears more strongly when a VE is wound up.

3. PATTERNS OF CREATIVITY

3.1 The Nature of Creativity

The nature of creativity, and models to describe creative work, have been discussed recently by (Hoffmann, 2005), (Blackburn, 2005), and elsewhere. Creative work has been distinguished by its semi-structured or unstructured nature, in contrast to following pre-defined procedures.

Much IT support for creative work has up to now been targeted primarily at face-to-face or other synchronous meetings of participants, with the intention of “brainstorming”. (Blackburn, 2005) proposes the monitoring of “cognitive dust” – the lowest level of recordable events in a high-tech meeting room, in order to provide a knowledge base for an electronic facilitator. This can be extended to include a number of “different place” – and even “different time” environments.

However creative work does not all happen within meetings. Much creativity takes place between meetings, often by individuals working alone, e.g. preparing ideas for discussion. Individuals often develop their ideas when totally off-line, such as when walking in the park or lying in bed.

Meetings are often the forum for debating proposals and individual ideas, rather than the place where the new ideas emerge. They also enforce a discipline, including convening the meeting (usually with an agenda) before, and recording decisions, allocated actions and other outcomes afterwards. This discipline still applies to “single meeting” creative work. Much of the information to do with meetings is in fact of a procedural nature.

As opportunities for creativity, meetings – especially face-to-face ones - have well-known problems. Regular meetings may become enslaved by their agendas and the creep of administrative overheads. All meetings, but especially recurring ones, may become bear-pits for egos and political agendas, and contributions of less pushy participants may be overlooked.

3.2 Cycles of Creativity

Creativity, especially in business, is often cyclic in character. There is almost always a goal that is to be achieved, but it may take several cycles to achieve the best result – or even an acceptable one. Creative work may also need to be broken down into smaller tasks and re-combined later.

The creative cycle often involves putting ideas to the test, or evaluating their worth by calculation, simulation or other models. Predicting risk, acceptability and ease of use of a proposed solution are more difficult in creative work than in operational work, where simple financial models are usually adequate.

One period of creativity may lead on into a subsequent one; for example developing requirements for a system typically leads into designing a technical solution.

In information work today, individuals are typically involved in multiple creativity cycles simultaneously. Some of these are related to the same area of creativity, while some are totally unrelated. Being involved in too many parallel activities is known to contribute to information overload, which in turn threatens creativity. Addressing this overload is one of the motivations for this paper.

4. CURRENT SUPPORT AND ITS LIMITATIONS

4.1 Current Support

Tools to support individual creativity are relatively plentiful in the software market. These range from the MS Word “outline view” to a graphical mind-mapping tool such as Mind Manager (www.mindjet.com). In specialist fields there are dedicated tools, e.g. CAD (Computer Aided Design) and music composition software.

Commercial Groupware such as Lotus Notes or Microsoft Exchange supports creativity, but only in an unstructured (and usually asynchronous) sharing and sending of messages and artefacts. Apart from version control, they do not generally offer any specific support for creativity. In a VE environment, there are also the problems that file formats may not be exchangeable, and there is no shared version control or access control, unless provided by one participant on behalf of the group.

For synchronous meetings, a number of commercial products and research prototypes are now available. Typically, several participants can view or manipulate an electronic version of the artefacts being discussed (e.g. text or diagrams). However once again, use of this technology in a VE requires all participants to support the same software.

4.2 Problems with Current Support

Heterogeneity of data formats, semantics and presentation is the first problem with any IT support for VEs. Different organizations in a VE may use different vendors’ products, or different versions from the same vendor. IT tools that are currently used are usually the products of a competitive market place. Marketing considerations often result in the data formats for a vendor’s proprietary software being non-standard and obscure. The vendor may provide support for importing foreign data into their format, but rarely for exporting. A good example of this is the MS Outlook .pst format for mailboxes.

Proliferation of user interfaces (UI) is a somewhat related problem. Whenever a user has to adopt a new tool – as someone involved in VE work is very likely to do – he or she has to learn yet another UI. Many individuals already suffer from having to use too many interfaces, and resent the waste of time when yet another one has to be learnt and remembered. This can result in a good tool not getting adopted.

Intra-enterprise bias is a problem because the cost of most computer infrastructure is almost always paid by real organizations (i.e. companies, or large public enterprises). Therefore, software that supports group work often favours groups that are entirely within that organization, and does not take account of the needs of VEs. The common use of the adjective “Enterprise” to describe future integrated systems, reveals this bias.

As discussed earlier, there is also a **synchronous meetings bias**. Products and prototypes have been developed to support synchronous meetings, but not so many are designed specifically to support what happens in between meetings. These periods are left to be supported by “ad-hoc” use of groupware tools, e.g. email or chat with attachments.

Another problem in supporting creativity is the **difficulty of capturing ideas when the human with the idea is not on-line**. Technology to support voice dictation and handwriting is improving, but to capture ideas expressed in hand-drawn diagrams is less simple.

Perhaps the hardest problem is **Context Recognition**, here defined as recognizing what a document or message is about and how it should be categorized within the structure of the group's (or individual's) area of concern. So far, even the best software has not yet matched the ability of a good secretary or filing clerk. Two approaches have been tried. One is **statistical text mining**, in which the frequency of words and phrases is mapped and candidate topic clusters are identified. An example of a tool using this approach is Leximancer (Smith, 2005). However this approach does not make use of any knowledge about the topic area already held by the groups or individuals. **Ontologies** (Noy, 2001) offer a complementary approach, by formalizing the structure of knowledge about a domain of discourse. However this also has drawbacks. In creative work, the focus of the knowledge often evolves rapidly. Also, significant effort is needed to build and maintain an ontology. In the VE environment there is the additional problem that no central ontology may exist unless all participants agree to follow the same one.

5. PROPOSED ARCHITECTURAL MODEL

5.1 Collaboration Architecture

Central to the work described in this paper is the observation that an increasing amount of collaborative work these days is "many-to-many". One human can participate in many groups, while collaborating groups can each involve several organizations or legal entities. Figure 1 illustrates this many-to-many concept.

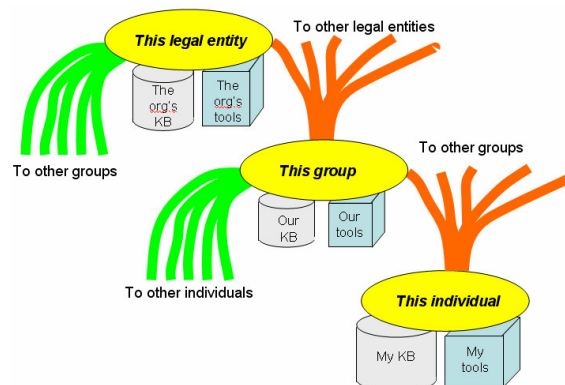


Figure 1 – Many-to-many Group Participation

Each group - and each individual - has a characteristic world view, represented in the diagram by the various levels of Knowledge Base (KB). There is also a set of tools that are used at each level. In a VE, there may be few group-level tools and

everything has to be done by using individual tools and passing versions of artefacts. No distinction has been made between creative and non-creative work here; as discussed earlier, the two are interwoven within the cycles of creativity.

5.2 Software Architecture

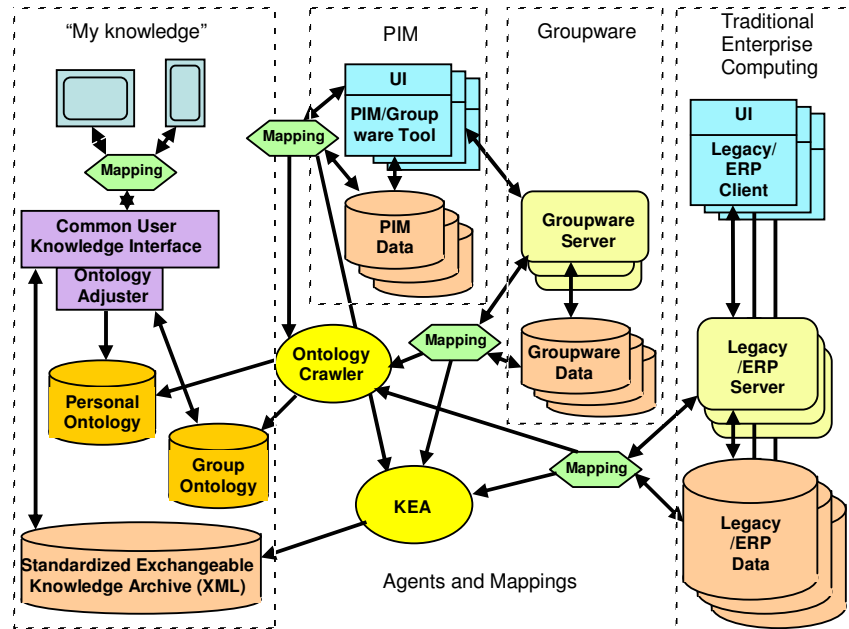


Figure 2 – Proposed Architecture for VE Creativity Support

Figure 2 shows one way in which the combined needs of a many-to-many model (Figure 1), a virtual enterprise and support for creative work can be addressed in terms of a set of software components. Two specific additions to normal software environments should be noted: **my knowledge** and **agents and mappings**.

My knowledge is essentially a single portal to a variety of information that is extracted or mapped from its original sources. The information is a combination of individual information and information from all the groups in which the individual participates.

Agents and mappings are a set of translators and transformation rules respectively, that deliver information in a consistent form. This might involve any combination of reformatting, indexing or linking. There may be agents operating for groups as well as individuals. The mappings depend on the meta-level definitions of the formats of the various source information types.

The diagram shows different user interfaces for the knowledge portal, traditional groupware and enterprise applications. The term PIM (Personal Information Management) has been introduced in this context by (Jones, 2005). These latter interfaces are typically access-restricted for each member organization in the VE.

6. PROPOSED DEVELOPMENTS

The four main new components in the architecture are the Knowledge Extraction Agent (KEA), the Ontology Crawler, The Ontology Adjuster and the Common User Knowledge Interface.

6.1 Knowledge Extraction Agent

Analogous to the work of a filing clerk, this architecture includes an autonomous agent known as the Knowledge Extraction Agent (KEA). It would operate in the background in a similar fashion to Google Desktop (Sullivan, 2004). The group would specify, through a Wizard, the sources of data and knowledge to be processed. This could include material reachable through networks as well as that stored on a specific computer.

The goal of a KEA is to make information available in a form that enables it to be read consistently by any collaborating participant. The output of KEA is an archive of marked up (in XML) records that either contain, or give links to, all information that the agent's "principal" (i.e. group or individual) deems relevant for further reference. It does not necessarily involve complete duplication – a set of keyword tags providing an indexing mechanism may be enough. Links would then point to the original content or service, assuming that a translator can be generated to at least render the source information readable.

In the VE scenario, multiple KEAs may operate on behalf of specific groups and individuals, each with their own objectives. The group coordinators or individuals can then determine who else can access the information by sharing or sending.

In order to categorize source knowledge, a KEA would need to be aware of the group's "knowledge structure", which could include ontologies, lists of relevant entities and rules derived from statistical text mining. For non-text sources such as pictures and multimedia, additional categorization rules would be required.

It is envisaged that the archive produced by a KEA would include standard formats for entities that appear frequently in knowledge bases. Examples are events, text messages, tasks, activities (i.e. task instances), rules, tables, diagrams, audio/video clips, annotations, processes, data structures and persistent business objects (e.g. Customers, Products).

For example, all data representing *tasks* would be accessible (either physically stored or realizable through mappings from data held in the original format) in a format such as PMXML (Curran, 2004). This would apply wherever those tasks originally come from, e.g. a Project Management System, a Workflow Management System that creates task instances from templates, or from a Groupware tool (e.g. MS Outlook) that supports tasks or "To Do" lists.

6.2 Ontology Crawler

As discussed earlier, ontologies have the disadvantage that they need to be created and maintained, especially where the creativity element is high. Without most of this work being automated however, the approach proposed in this paper will probably fail to gain acceptance.

At least an approximation to every human user's (or group's) knowledge structure can be deduced from an analysis of the files and applications he or she has

access to. Valuable sources include:

- directory structures of local or networked data storage drives or places
- schemas of relational databases and fixed format files
- IDL (interface definition language) for remotely accessed programs (WSDL for web services)
- address books in an email client
- folder structures of web browser bookmark files

as well as registries, imported ontologies and the results of text mining.

Any automatic categorization also depends on recognizing (e.g. by name) instances of key entity types. Examples are persons (e.g. collaborators or correspondents), business partners, projects/subprojects, products (e.g. physical items and assemblies) and services. These must either be included in an ontology or linked as reference tables

This leads to the observation that any application using an ontology also needs access to a “lexicon” whereby strings (e.g. as they appear in documents) can be related to concepts in the knowledge structure. However this correspondence is non-deterministic; the same string could indicate more than one concept. For example, “Winter” could be a season or the name of a customer.

The ontology crawler represents the meta-level counterpart of the KEA. However changes to the structure of knowledge are not as frequent as changes to the body of knowledge itself. The biggest job would be the initial crawl; subsequent crawls might only need to take place occasionally (e.g. monthly).

6.3 Ontology Adjuster

The Ontology Crawler, however sophisticated, may not produce the knowledge structure that reflects the group’s (or individual’s) natural classification structure. There needs to be some means by which users can correct the structures detected.

(Einig, 2005) describes a prototype, built under the current author’s supervision, for editing an ontology. This tool is targeted for use by non-specialist human users, rather than experts.

6.4 Common User Knowledge Interface

The final proposal is for a common interface - “My Knowledge” in Figure 2 (by analogy with the traditional “My Documents” of MS Windows).

To try and overcome the UI proliferation problem, it is desirable to have a standard set of logical interface patterns that supports users interacting with the standardized knowledge archive described in 5.2 above.

It is envisaged that this interface will also bridge to existing software and their associated file formats. One approach would be to repackage the existing software as “services” in the SOA (Service Oriented Architecture) sense. Also, mappings will be required that enable translations to be generated to convert the native formats to the KEA standard.

A further hurdle to be overcome here is the variety of user devices, such as PDAs, web-enabled mobile phones and wearable computers. (Javahery, 2003) proposed a system of patterns which can be modelled at a logical level between the server application and the different devices.

Figure 3 shows the structure of such a user interface architecture.

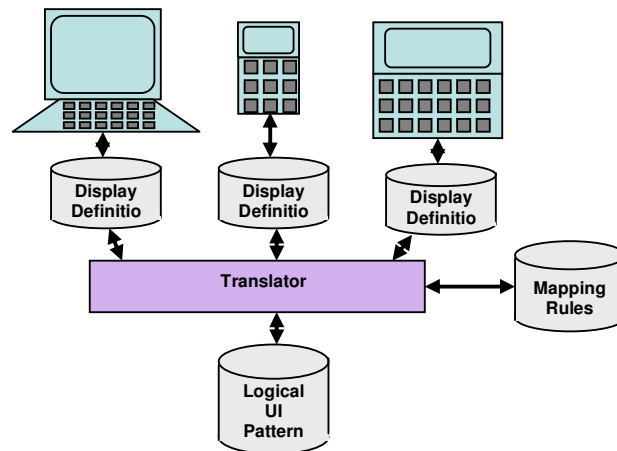


Figure 3 – Structure of a User Interface with a Device Independence Layer

7. SUMMARY

7.1 Work Done and Planned

In research projects supervised by the author, a number of test beds have been built and explored. We have built a simple ontology-based pre-processor for categorizing incoming email. Incoming email is one of the biggest irritants for many users. We searched message sender, subject and body text for strings that we could match with ontology concepts, then added keyword-style tags to create XML task records that represent the job of dealing the incoming request, enabling the tasks to be grouped by topic. We have also built a (hopefully!) end-user friendly tool for editing the knowledge structure stored in the RDF format commonly used for ontologies. However this prototype needs more serious testing with real users.

Our current work includes a combination of the above two techniques, development of a uniform model for tasks that derive from different applications such as workflow, enterprise applications, project management and PIM.

7.2 Related Work

The main vendors of groupware, including Microsoft and IBM, all have projects addressing some of these problems, but have not yet released much detail – see for example (Geyer, 2003) and the press release for IBM Lotus “Hannover” (IBM, 2005). University research seems more focussed on individual techniques such as text mining and categorization; one attempt to combine the statistical and ontology approaches is (Dittenbach, 2004). For details of a current project addressing this area, see (Ananiadou, 2005). The topic of Model Driven Architecture (Object Management Group, 2006) is also relevant, since MDA addresses the generation of mappings between different representations of information – something that is needed to provide the mappings in the architecture proposed.

7.3 Critical Success Factors

Apart from the immaturity in text categorization techniques, the most critical factor affecting whether or not the approach described here could ever be widely adopted is probably the user interface. User reaction to “yet another interface” would be understandably negative. To accept the proposed “My Knowledge” interface, users would need to be able to drop at least one, and preferably several, of their existing interfaces. However, history suggests that product timing, as much as technical design, determines whether a UI is widely adopted or not. The precedent of Visicalc, the first spreadsheet product, is perhaps a good example.

Although the challenge has been to provide an approach that addresses the need of creative work, analysis suggests that from an information point of view, creative work is usually closely interleaved with more routine administration. The following of processes may not be so important, but there seems no justification for a totally separate approach.

This paper has suggested a line of development that might lift IT support for groups, for both creative and other work, to a new level of effectiveness. Realization is perhaps a few years off, but research and development is starting to happen.

8. REFERENCES

1. Ananiadou, S., Chruszcz, J., Keane, J., McNaught, J., Watry, P., (2005): The National Centre for Text Mining: Aims and Objectives, <http://www.ariadne.ac.uk/issue42/ananiadou/>, 2005
2. Blackburn, T., Swatman, P.A., and Vernik, R. Extending CSCW theories to Model and Support Creative Group Processes, submitted for publication, <http://www.insyl.unisa.edu.au/publications/working-papers/2005-23.pdf>, 2005
3. Curran, K., Flanagan, L. and Callan, M. PMXML: An XML vocabulary intended for the exchange of task planning and tracking information, <http://www.xmlmania.com/documents.php?x=month&m=04&y=2004>, 2004
4. Dittenbach, M., Berger, H. and Merkl, D., Improving Domain Ontologies by Mining Semantics from Text, Proc. 1st Asia-Pacific Conference on Conceptual Modelling, Dunedin, New Zealand, <http://www.crpit.com/confpapers/CRPITV31Dittenbach.pdf> 2004
5. Einig, M., Tagg, R. and Peters, G. “Managing the knowledge needed to support an electronic personal assistant”, In Proc Int’l Conf on Enterprise Information Systems (ICEIS), 2006
6. Geyer, M., Vogel, J., Cheng, L., and Muller, M. Supporting Activity-centric Collaboration Through Peer-to-Peer Shared Objects, [http://domino.research.ibm.com/cambridge/research.nsf/0/ec360aeb7a39b10485256da9006506e2/\\$FILE/TR2003-09.pdf](http://domino.research.ibm.com/cambridge/research.nsf/0/ec360aeb7a39b10485256da9006506e2/$FILE/TR2003-09.pdf), 2003
7. Hoffmann, O., Cropley, D., Cropley, A., Nguyen, L. and Swatman, P.A. Creativity, Requirements, and World Views, Aust J of Info Systems, Vol 13, 2005; 1: 159-175
8. IBM, With the unveiling of the next release of Lotus Notes, code-named Hannover, press release, <http://www-306.ibm.com/software/swnews/swnews.nsf/n/nhan6dbjwg>, 2005
9. Javahery, H., Seffah, A., Engelberg, D. and Sinnig, D. Multiple User Interfaces: Multiple-Devices, Cross-Platform and Context-Awareness, chapter 12 “Migrating User Interfaces between Platforms Using HCI Patterns”. Wiley, 2003
10. Jones, W. et al. Proceedings of the PIM (Personal Information Management) Workshop, Seattle, WA, <http://pim.ischool.washington.edu/tools.htm>, 2005
11. Noy, N. and McGuinness, D.L. Ontology Development 101. Technical Report SMI-2001-0880, Stanford Medical Informatics, 2001
12. Object Management Group, MDA – the architecture of choice for a changing world, <http://www.omg.org/mda>
13. Smith, A. Leximancer home page, <http://www.leximancer.com/>, 2005
14. Sullivan, D. Google Desktop Search Launched, <http://searchenginewatch.com/searchday/article.php/3421651>, 2004
15. Tagg, R. Workflow in Different Styles of Virtual Enterprise, Proc of Workshop on IT in Virtual Enterprises, Aust. Computer Science Week, 2001