

## **Evaluating Inter-Organizational Information Systems**

**A book chapter submitted for the book titled, “Inter-Organizational Information Systems in the Internet Age,” Sean B. Eom, Editor**

Submitted by:

Jill Drury, Sc.D. (corresponding author)

The MITRE Corporation

202 Burlington Road, MS K320

Bedford, MA 01730

781-271-2034

[jldrury@mitre.org](mailto:jldrury@mitre.org)

Jean Scholtz, Ph.D.

National Institute for Standards and Technology (NIST)

1000 Bureau Drive, MS 8940

Gaithersburg, MD 20899

301-975-2520

[jean.scholtz@nist.gov](mailto:jean.scholtz@nist.gov)

Dr. Jill Drury is a usability engineer and researcher with an interest in evaluating collaborative computing systems; she is an Associate Department Head in the Collaboration and Multimedia Department of The MITRE Corporation and an Adjunct Assistant Professor in the Computer Science Department of the University of Massachusetts Lowell.

Dr. Jean Scholtz is a computer scientist with an interest in evaluation of interactive systems and human-robot interaction; she is a research scientist in the Information Access Division of the Information Technology Laboratory at the National Institute of Standards and Technology.

## **Evaluating Inter-Organizational Information Systems**

### **ABSTRACT**

This chapter describes different means of evaluating the usability and suitability of computer-based inter-organizational information systems (IOISs). It begins with describing why doing so is important yet difficult, and provides an assessment of the advantages and disadvantages of the major types of evaluation. It continues with a case study focusing on determining whether an application provides the necessary insight into other collaborators' identities, presence, and activities while keeping sensitive information private from a subset of the collaborators. The goal of this chapter is to provide practical guidance to organizations seeking IOISs to help them choose (or develop) an IOIS that best meets their needs.

### **KEYWORDS**

Groupware, heuristics, data sharing, human-computer interaction, Human-Machine Systems, person/machine interaction, Systems Evaluation, System Evaluation, Software Evaluation, IT investment evaluation, Evaluation Criteria, Evaluation Methods, IT Evaluation Methods, Electronic Collaboration, Collaborative Work Systems, Collaborative Technologies, User Characteristics, User Needs, User Requirements, Requirements Definition, Requirements Analysis, Requirements Specification, User interface requirements analysis, User Needs Assessment

## **INTRODUCTION**

Information technology can be used to support collaborations and partnerships among organizations for competitive purposes. Organizations have developed the notion of inter-organizational systems (IOSs), also known as inter-organizational information systems (IOISs), to support these collaborations and partnerships. An IOIS is defined as an automated information system shared by two or more organizations (Cash and Konsynski, 1985) in a collaborative fashion.

Compare the definition of an IOIS with the definition of Computer Supported Collaborative Work (CSCW) applications: applications that support coordinated activity carried out by groups of collaborating individuals (Greif, 1988). CSCW applications are also known as multi-user, groupware, or collaborative applications.

Collaborative applications normally provide capabilities beyond simple information access to facilitate communication and collaboration among the partners. Depending upon the collaborative application, both synchronous and asynchronous communications may be supported and documents can be shared. Some collaborative applications incorporate video to support communications and negotiations. These coordination mechanisms are essential to efficient collaboration among cooperating organizations. In fact, since IOISs are computer-based systems used to collaborate across organizations, they are a subset of collaborative applications.

Hong and Kim (1998) built on Cash and Konsynski's (1985) work by developing a framework for classifying the various types of IOISs. Their classification scheme is based on three categories: vertical linkage, horizontal linkage, and cross linkage.

Vertical systems connect suppliers with sellers with the goal of more efficient marketing. This type of system gives sellers, for example, the capability to place orders quickly but also gives suppliers sales data to help them plan production. Horizontal systems link homogeneous groups of businesses. Partnerships within an industry, often consisting of smaller businesses, benefit from improved access to information. Cross systems are an attempt to integrate both horizontal and vertical links into one complete system.

It is necessary to understand the roles of the participants or collaborators in IOISs in order to provide the necessary system capabilities to support a variety of tasks. For example, consider a vertical IOIS that links a manufacturer with a number of suppliers. A subset of those suppliers may be competitors who are negotiating terms with the manufacturer. Suppliers may want to use this system to share contractual information with the manufacturer but not with each other. Vertical and cross IOISs will need to support the most diverse set of users (e.g., suppliers, manufacturers, and retailers), though horizontal IOISs might also need to support differing groups of collaborators (e.g., manufacturers from the Eastern US versus manufacturers from the Western US). The roles of participants and their different information sharing needs should be taken into account when evaluating which IOIS is appropriate for a set of cooperating organizations.

To help organizations evaluate which IOIS they should adopt, and provide guidance for developing an IOIS, this chapter includes an assessment of the advantages and disadvantages of using different types of evaluation methods for determining the suitability of IOIS applications. The other contributions of this chapter are:

- An explanation of why IOISs are difficult, yet important, to evaluate
- A description of how the Synchronous Collaborative Awareness and Privacy (SCAPE) awareness framework could be used to evaluate an IOIS application
- A case study of evaluating the Groove<sup>TM1</sup> application's suitability for use by a collaborating team that includes members from organizations with different goals

## **Outline**

The rest of this section discusses the importance of evaluating IOISs, the difficulties of doing so, and the critical distinctions between evaluating single-user computing applications versus multi-user applications, such as IOISs. The second section describes evaluation methods for multi-user applications in general, and one method in particular, SCAPE, is singled out for description in the third section. The fourth section presents a case study of using SCAPE to evaluate Groove, a popular tool that aids inter-organizational information sharing. Finally, a discussion of what can be learned by evaluating IOISs completes the chapter.

## **The Importance of Evaluating IOISs**

Much research has centered on evaluating the usability of collaborative applications because it is extremely important to ensure that these applications can be effectively and

efficiently used by their intended audience. The success of a collaborative application normally depends on a “critical mass” of users accepting and making proper use of the application. For example, picture:

- An inventory control system so cumbersome to use that some of the staff receiving inventory neglect to log the inventory into the system or log it incorrectly.
- An instant messaging application that users did not find easy and rewarding to use, so very few of a user’s business or social contacts bother to remain accessible via instant messaging.
- An automated calendar management application that takes a lot of work to enter activities into, so many people within a workgroup don’t make an effort to keep their calendars up-to-date.

Clearly, each of these situations constitutes a recipe for failure. These cases illustrate the fact that a collaborative application is likely to fail if the work people need to put into the application exceeds the perceived value of their benefits from using the application (Grudin, 1988).

Besides an imbalance in work versus benefits, there are other reasons why adoption of an IOIS may fail. For example, we have seen adoption of a collaborative application fail when:

- Users did not perceive a need to collaborate. Such a finding is consistent with Rogers' (1995) work on diffusion of innovations, which notes that the rate of adoption of innovations is related to the extent to which the innovation (e.g., a new collaborative application) satisfies users' needs.
- The application did not provide functionality that users felt was relevant. Relevant functionality is dependent on the tasks the users need to perform and the conditions under which they normally perform those tasks. For example, an IOIS intended for use by personnel driving delivery trucks that requires a substantial amount of typing (instead of, say, using a bar code reader) would be likely to be unsuccessful.
- Users were not available to log in frequently. Users who often attend meetings or engage in other activities that preclude access to the IOIS, for example, would be unlikely to embrace use of the IOIS.
- Users did not develop a well-articulated communications strategy. An example of an incomplete communications strategy is one that does not define situations in which to use the collaborative application versus email.
- The application was not easy to learn or use. Rogers (1995) notes that adoption of innovations is related to the complexity or ease with which an innovation can be understood.

There is normally a large financial difference to a collaborating group of organizations between failure to adopt an IOIS and adopting—and making good use of—an IOIS that is well-suited to those organizations. When adoption failure occurs, the organizations must

count as a loss the purchase price of the IOIS plus the loss in productivity represented by the hours spent installing, training, and experimenting with the IOIS. Collectively, these costs could be extremely substantial, especially for large organizations that may have asked hundreds of people to try the IOIS. In addition, there is an intangible cost: members of the organizations may be less open to use of IOISs in the future once they have had a negative experience with an IOIS.

Contrast this failure situation with the case in which organizations choose an IOIS that meets their needs and streamlines their business processes. Depending upon how an IOIS is used, an effective IOIS may result in a decreased need to travel (because IOIS technology can often mitigate the need for face-to-face meetings), shorter document production review cycles, decreased time-to-market, increased sales, and/or better customer support.

The difference between adoption failure and success hinges on defining collaboration requirements that take into account the work characteristics of users, the likely benefits to users, as well as ease of use from the point of view of the intended set of users. A further prerequisite for success is an evaluation program that examines how well an IOIS is likely to meet those requirements.

Evaluation goals differ depending on whether an IOIS is being chosen from among a set of existing products or a custom (bespoke) IOIS product is being developed. If an IOIS is being chosen, the candidate applications are each examined against a tailored set of

requirements, using one or more evaluation methods such as those discussed later in this chapter. Since the commercial IOISs cannot normally be modified substantially, the one that comes closest to meeting the requirements is chosen. If a custom IOIS is being developed, the goal of evaluation is to find problems that can be corrected as early as possible in the product design and development lifecycle. The later in the development process that interface problems are found, the more costly they are to correct. Mantei and Teorey (1988) found that changes made to the interface designs of systems after production coding had begun cost four times as much as changes to the designs made during prototyping phases.

### **The Challenges of Evaluating IOIS Applications**

There are several reasons why evaluating collaborative applications is more difficult than evaluating single-user computing applications. Malone (1985) cites the difficulties in assembling a group of people in a lab that reflect the social, motivational, economic, and political characteristics of typical users—yet these characteristics are likely to affect performance when using the collaborative system. If evaluation is attempted in the users' normal work environments (“in the field”), Grudin (1988) observes that it is extremely difficult to disperse evaluators to the various locations of the collaborators as well as take into account the wide variation in user group composition and work environments. Regardless of whether they occur in a lab or in the field, Grudin (1988) notes that evaluations of collaborative applications take much more time than evaluations of single user applications because the relevant group interactions “typically unfold over weeks.”

Adding to the difficulty of collaborative application evaluation is the fact that sophisticated applications allow users to take on a number of different roles. Users' expectations of an applications' behavior may change depending upon the roles users are playing at the time and the specific tasks they are performing. More generally, collaborative applications are challenging to evaluate due to the need to take into account how the application mediates users' interactions with each other.

Although difficult to perform, evaluations of collaborative applications are extremely important due to the cost implications described above. The purpose of this chapter is to provide insight into the state-of-the-art in evaluating collaborative applications in general, and into one evaluation method in particular (which will form the basis for the case study presented later in this chapter).

### **The Critical Difference Between Evaluating Single-User and Multi-User (e.g., IOIS) Computer Applications**

The preceding subsection touched upon the crucial distinction between evaluations of single-user systems, such as word processors, and multi-user systems, such as IOIS applications. An evaluation of multi-user (collaborative) systems needs to investigate whether the application adequately supports collaborators' awareness of each others' presence, identities, and activities. Awareness is important in collaborative applications because it aids coordination of tasks and resources, and it assists in transitions between individual and shared activities (Dourish and Bellotti, 1992). Dourish and Bellotti (1992) defined awareness as "an understanding of the activities of others, which provides a

context for your own activities.” “Workspace awareness” was defined by Gutwin, et al., (1995) as the up-to-the-minute knowledge of other participants’ interactions with the shared workspace, such as where other participants are working, what they are doing, and what they have already done in the workspace.

To understand the importance of awareness, picture trying to use a chat application without being able to read the contributions of the other participants; clearly the application is useless without an understanding of the other participants’ activities. Chat is an example of a synchronous collaborative application (those that are used by collaborators at the same time, although not necessarily at the same place). An “up to the moment” awareness of others’ activities is especially pertinent to the class of synchronous (as opposed to asynchronous) collaborative applications. An example of an asynchronous collaborative application is email.

Awareness and privacy are in tension with one another, as are awareness and information overload. Hudson and Smith (1996, p. 247) expressed the tradeoffs very well:

This dual tradeoff is between privacy and awareness, and between awareness and disturbance. Simply stated, the more information about oneself that leaves your work area, the more potential for awareness of you exists for your colleagues. Unfortunately, this also represents the greatest potential for intrusion on your privacy. Similarly, the more information that is received about the activities of colleagues, the more potential awareness we have of

them. However, at the same time, the more information we receive, the greater the chance that the information will become a disturbance to our normal work.

This dual tradeoff seems to be a fundamental one.

Any evaluation method that pertains to collaborative applications should be sensitive to issues of privacy and awareness. For example, a student using a distance learning application may want to make the instructor aware of all of his or her on-line activities, while the instructor would want to keep grading activity private (except to the student directly affected).

It is difficult to apply evaluation techniques developed for single-user applications to multi-user applications because they do not address the issues of awareness and privacy. Only recently has there been a push towards developing evaluation methods specifically for collaborative applications.

## **USABILITY/SUITABILITY EVALUATION METHODS FOR IOIS**

### **APPLICATIONS**

An important prerequisite for applying usability and/or suitability evaluation methods is a thorough understanding of the users' requirements and desires. Thus, the first part of this section provides a brief discussion of how to acquire such an understanding.

The remainder of this section provides an overview of the three broad categories of evaluation methods: formal methods (analytic methods such as dialog and task modeling techniques), empirical methods (experiments and user studies involving human test subjects), and inspection methods (expert examination of user interfaces). Despite the difficulties enumerated above, some usability evaluation methods have been developed for collaborative systems.

#### **Understanding Users' Needs**

There are many techniques for learning about users' collaboration needs. Some of them are:

- **Task analysis:** a family of different techniques that involve breaking apart users' tasks, from the standpoints of either cognitive or physical activities, at a high level of abstraction or in great detail, depending upon the particular task analysis technique chosen. For practical advice on performing task analyses, we recommend Mayhew (1999).

- **Ethnographic observation:** a broad-based approach originating in anthropology in which users are observed while they pursue their normal activities; observers become participants by immersing themselves in the users' environment. For examples of ethnographic observation applied to adoption of collaborative applications, see the work of Bonnie Nardi (e.g., (Nardi and O'Day, 1999)).
- **Contextual inquiry:** an ethnographic-based technique in which the observer becomes an apprentice of the person being observed; besides observation, contextual inquiry involves focused interviews, discussion, and reconstruction of past events (Holtzblatt and Jones, 1993).
- **Critical incident interviews:** a method in which users are interviewed about the events and activities surrounding an unusual or high-impact event. Klein (2000) describes the use of critical incident interviews for collaborating teams.

Without understanding users' characteristics and work environment, it is impossible to determine whether an IOIS would be "natural" or "intuitive" for those users, or whether the IOIS would be compatible with the users' normal work practices. Consider an application targeted at scientists and mathematicians, such as Mathematica™.

Mathematicians expect to see terminology in the interface such as "factorial" and "cosine;" they don't need definitions of these terms. Factorial and cosine functions also exist in Excel™, which was designed for a general audience. In Excel, mathematical terms are defined and the definitions are readily visible (they are not buried in a "help" file, for example).

## **Formal Methods**

An example of a formal method that can be used to evaluate collaborative systems is Critical-Path-Method (CPM)-Goals, Operators, Methods, Selection rules (GOMS) (John and Kieras, 1994). CPM-GOMS is also known as Cognitive-Perceptual-Motor-GOMS because of its purpose to model the parallel, multi-stage processor nature of human information processing. CPM-GOMS is a task modeling technique that allows the analyst to break down a task at a very fine level of granularity, such as individual eye and hand movements. The method does not assume that each subtask happens serially; it takes into account the parallel nature of performing activities (e.g., both hands can be moving at the same time while eye movement is also occurring). The end results are predictions for task execution times. While CPM-GOMS was originally envisioned as a method for analyzing a task performed by an individual, its assumption of parallelism lends it to analyzing a task performed by a team of individuals.

An advantage of using a formal method such as CPM-GOMS is the fact that no user participation is required, which simplifies the problem of trying to recruit and schedule users and either replicate a realistic work environment in the lab or capture all facets of the users' environment in the field. A disadvantage of using a formal method is that evaluators normally require extensive training in the method because they are usually complex and can require grounding in specific theories. CPM-GOMS is useful for obtaining a very detailed understanding how quickly a particular task can be done using an interface but cannot be used to answer broader questions such as, "how satisfied will the users be with this interface?"

## **Empirical Methods**

We are not aware of any empirical methods that have been tailored or created specifically for collaborative systems. Some researchers have applied empirical methods developed for single-user systems to small-scale collaborative systems with success. For example, Gutwin and Greenberg (1998) performed usability testing to compare two different interface approaches for a collaborative computer-assisted welding application. In general, usability tests consist of typical users performing typical tasks under controlled, but realistic, conditions, either in a laboratory or in the field. In Gutwin and Greenberg's test, the subjects worked in pairs in two different locations and performed their tasks over the course of a few hours. The study goals were focused enough to involve only two people at a time performing a few tasks with minimal training. As a result, the study conductors could obtain a rich amount of data and insights within a manageable time period.

Usability tests are often considered to be the “gold standard” in terms of the amount of data and the subtlety of problems that may be uncovered; thus, they are advantageous to perform whenever it is not too difficult to do so. The difficulties normally arise when duplicating a realistic environment of use, recruiting appropriate users, and scheduling them in groups. The challenges only increase when “typical” user group sizes rise; testing with two to five people at a time is much more tractable than fifty, for example.

Although it is highly desirable to perform realistic usability tests, such tests have proven to be too difficult and expensive to perform on collaborative applications in many cases. Baker, Greenberg, and Gutwin (2001) states, “we have not yet developed techniques to make groupware [collaborative] evaluation cost-effective within typical software project constraints” (Baker, Greenberg, and Gutwin, 2001). Thus, collaborative applications are often developed or chosen without any evaluation whatsoever. Recent work in tailoring inspection methods for collaborative applications has taken place in an attempt to provide a reasonable means for collaborative application evaluation to take place.

### **Inspection Methods**

Inspection methods are promising since they can often be performed more quickly and inexpensively than the other usability evaluation methods; savings accrue because they do not involve scheduling users (as do empirical methods) or extensive training of evaluators (often needed for formal methods). They are often used when there is insufficient time or budget to perform usability testing or to analyze an interface using a formal method. Further, they are often used early in the development process on low-fidelity prototypes to gain early insight into whether the proposed design is consistent with general principles of human-computer interaction (even if empirical evaluations are scheduled for later versions of the application). The disadvantage with inspection methods is that they do not always result in finding the subtle problems that occur due to mismatches between the application design and the user’s mental model of how the application is working.

The classic inspection method is heuristic evaluation (Molich and Nielsen, 1990). It is useful to describe it because several methods have been developed for evaluating multi-user systems that are adaptations of heuristic evaluation. When performing a heuristic evaluation, inspectors (often, but not necessarily, usability specialists) judge whether each user interface element conforms to established usability principles known as heuristics. Examples of heuristics are, “The interface should be consistent” and “The interface should provide clearly marked exits.” To apply the heuristics, individual evaluators independently step through all parts of a user interface, noting cases where the interface violates the heuristics. After looking at the interface, each evaluator may assign a score to how well the interface meets each heuristic in general. Once each individual assessment is complete, evaluators normally discuss their findings and agree upon a joint set of problems and scores. The power of this method comes from combining the observations of several inspectors, because people normally find somewhat different subsets of the problems. Heuristic evaluation is straightforward enough so that people other than human-computer interaction experts or human factors engineers can successfully perform a heuristic evaluation with as little as an hour’s training.

Other inspection methods compare an application against a set of guidelines (either general or application-specific) or a “capabilities” (function) checklist tailored to the users’ needs. An example of a tailorable function checklist for collaborative applications can be found in Drury et al. (Drury, Damianos, Fanderclai, Hirschmann, Kurtz, and Linton, 1999).

Three inspection methods developed for collaborative systems employ heuristics-based inspection: benchmarks for workspace awareness (Villegas and Williams, 1997), the Locales Framework heuristics (Greenberg, Fitzpatrick, Gutwin and Kaplan, 2000), and the “Mechanics of Collaboration” (Baker, Greenberg, and Gutwin, 2001). An additional inspection method, Synchronous Collaborative Awareness and Privacy Evaluation (SCAPE) (Drury, 2001) provides both a means of specifying awareness and privacy requirements and evaluating whether the application satisfies the requirements via a heuristic approach. We describe two methods in more detail—the Mechanics of Collaboration and SCAPE—because they are more recent and more mature than the others.

Gutwin and Greenberg (2000) maintain that there are some basic collaboration activities that should be supported by any collaborative application: “These activities, which we call the mechanics of collaboration, are the small scale actions and interactions that group members must carry out in order to get a shared task done. Examples include communicating information, coordinating manipulations, or monitoring one another” (Gutwin and Greenberg, 2000, p. 98). Gutwin and Greenberg have proposed that the mechanics of collaboration framework can be used to construct heuristics. They have formed eight heuristics (Baker, Greenberg, and Gutwin, 2001, p. 125):

Provide the means for intentional and appropriate verbal communication

Provide the means for intentional and appropriate gestural communication

Provide consequential communication of an individual’s embodiment

Provide consequential communication of shared artifacts

Provide protection

Provide management of tightly and loosely coupled collaboration

Allow people to coordinate their actions

Facilitate finding collaborators and establishing contact

The idea behind the Mechanics of Collaboration method is that evaluators inspect the interface using the heuristics from Baker, Greenberg, and Gutwin (2001) instead of the ones developed by Molich and Nielsen (1990) or, more recently, Nielsen (1994).

Otherwise, the method is essentially the same as that developed by Molich and Nielsen (1990).

Note that the heuristics of Baker, Greenberg, and Gutwin (2001) are very broad and make the assumption that the role of the user is not a factor in the evaluation. The SCAPE method was developed to provide a finer-grained evaluation technique, acknowledging that users of an application may have different awareness and privacy needs.

## **THE SCAPE METHOD**

Later in this chapter, we will give an example of how SCAPE was used to evaluate the awareness and privacy support of a particular application being considered for use by a retailer and supplier. Since we will give examples of developing SCAPE analysis materials that were culled from our case study, we introduce the case study scenario in the next subsection. The fine-grained awareness framework and awareness relationships that underpin SCAPE are introduced in the following subsections, ending this section with a by a step-by-step description of the SCAPE method.

### **Case Study Scenario**

We developed a scenario based on recent events in the popular business press (e.g., (Bianco and Zellner, 2003)). The scenario centers on the partnership between a fictitious major retailer (“Wal-Store”) and an equally fictitious nation-wide distributor of juice drinks (“Sea Spray”).

Joy Brown is a beverages buyer at Wal-Store and Jenn Smith is a manufacturer’s representative from Sea Spray. Jenn works with two product managers, Rod Leeds and Sally Steele. Joy wants to get the Sea Spray juice drink products on the shelves at Wal-Store at the lowest possible price by probing Sea Spray’s manufacturing, inventorying, and distribution processes and suggesting ways to streamline these processes. Jenn would like to get as many different Sea Spray products on Wal-Store’s shelves as possible and so is willing to work with Joy—up to a point. She does not want to give Joy information about the recipes for the

newest juice drinks because Wal-Store has a history of using “inside” product information to manufacture similar products under the “Sal’s Choice” in-house label, thus eating into the market share of national brands.

Jenn would like to use the same IOIS with Joy as well as Rod and Sally, her product managers for the two new juice drinks. She would like to keep the proprietary information private from Joy, however.

By highlighting the natural tension between competitors, this Wal-Store scenario includes the notion of *adversarial collaboration* (Cohen, Cash, and Muller, 2000): situations in which collaborators have widely divergent goals yet must work together to perform specific tasks. Other situations that involve adversarial collaboration are merger/acquisition negotiations and document sharing among opposing legal teams. We expect that, as technology becomes more widely used to facilitate communications between organizations with divergent goals, computer-based adversarial collaboration support will become increasingly more important.

Both Cohen et al.’s study and the Wal-Store scenario illustrate a need for IOISs to support detailed awareness and privacy requirements, and to ensure that everything that is visible to others is revealed intentionally instead of being revealed accidentally or by default.

## **Awareness Framework**

Before describing the SCAPE method, we first need to refine the definitions of awareness cited earlier into a more fine-grained analysis framework. Although there are many definitions of awareness (a few of which were cited above), there is no standard definition. We define awareness as follows (Drury, 2001):

**Awareness:** Given two participants  $p_1$  and  $p_2$  who are collaborating via a synchronous collaborative application, awareness is the understanding that  $p_1$  has of the identity and activities of  $p_2$ .

To support  $p_1$ 's understanding of  $p_2$ , an application provides  $p_1$  with information about the identity and activities of  $p_2$  *without  $p_1$  having to request the information or  $p_2$  having to explicitly transmit it*. Awareness information is intended to emulate the kinds of non-verbal cues that people get about each other when they work face to face in the same physical environment.

The awareness information that  $p_1$  might have access to regarding  $p_2$  includes (but is not necessarily limited to)  $p_2$ 's presence in the shared workspace,  $p_2$ 's identity, the task that  $p_2$  is performing, the tools and artifacts that  $p_2$  is using, the changes  $p_2$  is making, the area of the workspace viewable by  $p_2$ , and  $p_2$ 's focus within that viewable area. We refer to the aggregation of all of the awareness information that all of the participants may have about each other as the *awareness information space*.

Although the awareness literature largely assumes that the more awareness information available, the better, there are times when awareness information should be withheld. For our purposes, we define *privacy* as follows:

**Privacy:** the intentional withholding of awareness information.

Evaluation of awareness support has two underlying principles: (1) ensure that awareness information that should be provided is provided; and (2) ensure that awareness information that should not be provided is kept private. A different type of awareness-related error is associated with each of these principles (Drury and Williams, 2002):

**Type 1. Awareness violation.** Awareness information that should be provided is not (a violation of the first principle).

**Type 2. Privacy violation.** Awareness information that should be kept private is revealed (a violation of the second principle).

A method for evaluating awareness support must provide for the identification of both awareness violations and privacy violations.

There is currently no theory *per se* of awareness support. We constructed a general-case default specification for the kinds of awareness information after examining eight theories and frameworks for collaborative work (see Drury and Williams (2002)). We developed a general, application-independent awareness specification, expressed as

heuristics (see Figure 1). These *awareness heuristics* are of two types: *activity heuristics* pertaining to activities only, regardless of who is performing them; and *identity heuristics* pertaining to participants.

The general awareness specification has two assumptions built into it: (1) that all participants using a collaborative application have the same awareness needs, and (2) that participants require access to all possible awareness information about each other. The general specification may be tailored for specific applications where these assumptions do not hold.

Type	Heuristic
<b>Activity</b>	Show the tasks being performed Show the tools being used Show the changes being made. Show the historical changes made. Show the time of each historical change.
<b>Identity</b>	Show participants' identities. Show the immediate intentions of each participant. Show the focus of each participant. Show area viewable by each participant. Show the participant(s) performing each task. Show the participant(s) using each tool. Show the participant(s) making changes. Show the participant(s) who made each historical change.

**Figure 1. General awareness specification, expressed as activity heuristics and identity heuristics.** This specification is application-independent.

We refer to the situation where a participant has access to all possible awareness about another participant as *complete awareness*:

**Complete awareness:** Given two participants  $p_1$  and  $p_2$  who are collaborating via a synchronous collaborative application, if all information regarding  $p_2$  in the awareness information space is available to  $p_1$ , we say that  $p_1$  has complete awareness with respect to  $p_2$ .

It is not uncommon for a participant to need (or to be permitted) only limited, rather than complete information about another participant, so that *partial awareness* should be provided:

**Partial awareness:** If some, but not all, information regarding  $p_2$  in the awareness information space is available to  $p_1$ , we say that  $p_1$  has partial awareness with respect to  $p_2$ .

Whether a participant  $p_1$  needs complete or partial awareness of another participant  $p_2$  depends on the *roles* played by  $p_1$  and  $p_2$ . For example, in the application supporting the retail scenario,  $p_1$  may be a buyer and  $p_2$  may be a sales representative. We define a *role* to be a participant's activities and responsibilities with respect to the other participants in a collaborative session.

Participants using a collaborative application can be partitioned into role-based equivalence classes (e.g., buyer, sales representative). All participants in the same role have the same awareness and privacy needs.

### **Awareness Relationships**

Each role is related to every other role by an *awareness relationship* that characterizes the awareness information participants in one role may have about participants in another.

An awareness relationship may be either complete or partial, depending whether participants have complete or partial awareness of each other. It is also possible for participants in one role to have no awareness information at all about participants in another role, in which case we describe the awareness relationship as *no awareness*.

Awareness relationships are unidirectional. Given two roles  $r_1$  and  $r_2$ , there is an awareness relationship for  $r_1$  with respect to  $r_2$  characterizing the information available to participants in  $r_1$  about participants in  $r_2$ . The relationship may be complete, partial, or no awareness. Similarly, there is an awareness relationship for  $r_2$  with respect to  $r_1$ . It also may be complete, partial or no awareness, independent of the type of the relationship for  $r_1$  with respect to  $r_2$ .

Figure 2 shows a matrix of awareness relationships for evaluating a hypothetical collaboration application used by buyers and sales representatives, assuming use by one participant in the buyer role, one in the sales representative role, and two in the product managers' roles. Some roles (e.g., product managers) have awareness relationships with

themselves, which characterize what participants in the same role may know about each other (what one product manager may know about another). Since there is only one buyer (and one sales representative), there is no buyer-buyer (or sales representative-sales representative) awareness relationship.

What these roles... ↓	...can know about the roles below		
	Buyer	Sales representative	Product manager
Buyer (1)	-----	Partial	No awareness
Sales representative (1)	Partial	-----	Partial
Product manager (2)	No awareness	Partial	Complete

**Figure 2. Awareness relationship matrix for the hypothetical application used by retailers and suppliers.** Numbers indicate the number of participants in each role.

### Steps for Performing a SCAPE Evaluation

SCAPE is based on the awareness framework and awareness relationships. It can be used to determine whether users' awareness and privacy needs would be met by an application. In particular, SCAPE is especially useful for evaluating a synchronous collaborative application because such systems require an up-to-the-moment awareness of fellow participants' identities and activities. Figure 3 illustrates the difference in awareness needs for synchronous versus asynchronous applications. A "yes" under the "Sync." column means the heuristic is applicable to synchronous applications; similarly, a "yes" in the "Asynch." column implies the heuristic pertains to an asynchronous

application. A “no” in either column means the heuristic is not relevant for that type of collaborative application.

<b>Heuristic</b>	<b>Sync.</b>	<b>Async.</b>
Show the tasks being performed	Yes	No
Show the tools being used	Yes	No
Show the changes being made.	Yes	No
Show the historical changes made.	Yes	Yes
Show the time of each historical change.	Yes	Yes
Show participants’ identities.	Yes	Yes
Show the immediate intentions of each participant.	Yes	No
Show the focus of each participant.	Yes	No
Show area viewable by each participant.	Yes	No
Show the participant(s) performing each task.	Yes	No
Show the participant(s) using each tool.	Yes	No
Show the participant(s) making changes.	Yes	No
Show the participant(s) who made each historical change.	Yes	Yes

**Figure 3. Awareness Needs in Synchronous Versus Asynchronous Applications.** “Yes” means the heuristic is applicable; “No” means the heuristic is not applicable.

SCAPE is an inspection method performed by evaluators, and is based on the awareness framework described earlier in this section. It is designed to help evaluators find both awareness violations and privacy violations. SCAPE takes into account the roles that participants play and the awareness relationships between the roles.

The SCAPE method has two parts: (1) development of a detailed, application-specific specification of awareness and privacy requirements, and (2) evaluation of the application for compliance with the specification. The SCAPE Handbook (Drury, 2001) contains

more detailed explanations than can be included here, as well as examples, advice, and worksheets.

There are three steps to developing an awareness specification for a collaborative application: define the awareness relationships, develop role-based awareness policies, and identify activity-based exceptions to the policies.

**Step 1: Define awareness relationships.** The goal of this step is to identify roles and the high-level awareness relationships between them. Knowledge of the application domain is needed to perform this step. Roles are identified and an awareness relationship matrix, such as the one shown in Figure 2, is created. The output of this step is the matrix.

**Step 2: Develop role-based awareness policies.** The goal of this step is to develop a set of awareness policies, based on the relationships between roles. The approach is to modify the general awareness specification from Figure 1 according to the awareness relationships identified in Step 1. While Awareness Policies can be created for each role pair, in practice usually only one Awareness Policy is needed; it indicates the superset of what participants can know about other participants. Exceptions to this policy are the identified in the next step. The Awareness Policy for the case study is shown in Figure 4.

For each awareness relationship, the evaluator begins with the general awareness specification and deletes any portions of the specification that are not applicable. (For example, if the application requires anonymity, then all identity heuristics are deleted.)

---

**Awareness Policy:  
What Participants CAN Know About Other Participants  
Worksheet for Step 2**

**Date:** 13 May 2003  
**Evaluator:** J. L. Drury  
**Application:** Groove

**Instructions:**

1. Add missing heuristics.
2. Cross out heuristics that are not applicable to participants in any role.

**Identity awareness heuristics:**

Show the identity of each participant.  
Show the immediate intentions of each participant.  
Show area viewable by each participant.  
Show the focus of each participant.  
Show the participant(s) performing each task.  
Show the participant(s) using each tool.  
Show the participant(s) making changes.  
Show the participant(s) who made each historical change.

*Other heuristics for identity awareness information:*

**None.**

**Activity awareness heuristics:**

Show the task(s) being performed.  
Show the tool(s) being used.  
Show the change(s) being made.  
Show the historical change(s) made.  
Show the time of each historical change.

*Other heuristics for activity awareness information:*

**None.**

---

**Figure 4. Awareness Policy for the case study.** None of the awareness heuristics are deleted in this example, and no extra heuristics have been added. Notations to the worksheet are indicated by this Comic Sans MS font.

We also leave open the possibility that the evaluator may need to add to the specification.

The evaluator performs this step as many times as necessary:

- once for all of the complete awareness relationships, since they all have the same requirements
- once for each partial awareness relationship, since they may all have different requirements

The output of this step is a set of role-based awareness policies expressed as heuristics.

**Step 3: Identify activity-based exceptions.** The goal of this step is to tailor the role-based awareness policies developed in Step 2 so that they include any activity-related exceptions that are necessary to ensure privacy. All of the policies for partial awareness relationships must be tailored in this way. (There is no need to tailor the policies for complete and “no awareness” relationships.)

The evaluator begins with the awareness relationship matrix from Step 1 and the set of awareness policies from Step 2, and uses domain knowledge about tasks that participants will perform. The evaluator identifies the activities during which parts of an awareness policy should be suspended. (For example, it may be OK for a product manager to be aware of what a buyer is doing *except* when the buyer is working on another purchase in which the product manager is not involved.)

The result can be seen in Figure 5, which is a portion of a privacy policy for the case study. The set of tailored awareness and privacy policies constitute an application-specific awareness requirements specification.

There are three steps to evaluating an application for compliance with an awareness specification: assess the level of effort a complete evaluation would entail, develop scenarios to use during evaluation, and perform the evaluation.

---

**Privacy Policy: What Participants  
CANNOT Know About Other Participants  
Worksheet for Step 3**

**Date:** 13 May 2003  
**Evaluator:** J. L. Drury  
**Application:** Groove

**Role pairs analyzed:** what Buyer CANNOT know about Sales rep.

**Instructions:**

- List privacy needs as exceptions to applying the heuristics.
- Note impacts to violating the privacy needs.

**Privacy Needs Table**

<b>Apply the <i>identity</i> awareness heuristics except for information related to the following activities:</b>	<b>Impact of privacy violation (low, medium, or high)</b>
Sales rep's bottom-line price (unless she chooses to reveal it deliberately)	High
Manufacturer's recipes for juice drinks	High
<b>Apply the <i>activity</i> awareness heuristics except for information related to the following activities:</b>	<b>Impact of privacy violation (low, medium, or high)</b>
Same as above	

---

**Figure 5.** Sample Privacy Policy for the case study, showing what the buyer cannot know about the sales representative's activities.

**Step 4: Assess level of effort.** The goal of this step is to assess the level of effort that a full-fledged evaluation would entail. A SCAPE evaluation involves a series of mini-evaluations, since each awareness relationship needs to be evaluated. The level of effort

can be substantial, because the worst-case number of awareness relationships can be roughly estimated as the square of the number of roles. Thus, it is prudent to calculate the level of effort needed for a complete SCAPE evaluation, compare the result to the resources available, and select the highest-priority awareness relationships to evaluate.

The evaluator begins with the role relationship matrix from Step 1, the specification from Step 3, and an understanding of resources (time, money, etc.) available. If the evaluator expects that evaluating one awareness relationship will take  $e$  effort (measured in evaluators' time, money, etc.), and if there are  $a$  awareness relationships, then the level of effort for a complete evaluation is proportional to  $e * a$ .

To identify the awareness relationships that are the highest priority for evaluation, the evaluator annotates the awareness policies, indicating whether violations would have high, medium, or low impact. The output is a list of prioritized awareness relationships.

**Step 5: Develop scenarios.** The goal of step 5 is to specify activities that will cause SCAPE evaluators to perform sequences of actions that will exercise the application's awareness support capabilities. Scenarios are a well-established technique for evaluation; Nielsen has pointed out the advantages of using scenarios for heuristic evaluation when it is important to examine participants' interactions (Nielsen, 1995). The evaluator uses the awareness specification from Step 3 and the prioritized list of awareness relationships from Step 4 to develop a master scenario and scenario worksheets to use during the

evaluation. The master scenario worksheet from the case study can be seen as an example in Figure 6.

**Step 6: Perform the evaluation.** The goal of this step is to identify awareness violations and privacy violations. Teams of evaluators perform the scenarios developed in Step 5 and examine the application for compliance with the specification from Steps 2 and 3. The output of this step is a set of problem reports.

Before performing a SCAPE evaluation, the evaluators need to understand the users' tasks and roles and the applications' functionality. To understand the users, we performed a brief task analysis and critical incident analysis based on structured interviews with a buyer for a large Eastern-US retailer. One of the evaluators already had significant prior experience with using the collaborative application and helped train the other (supplementing training provided by the developer's tutorial information).

An explanation of the evaluation performed for the case study can be found below, after a description of the application evaluated.

**Master Scenario  
Worksheet for Step 5**

**Date:** 26 May 2003  
**Evaluator(s):** Jill Drury, Jean Scholtz  
**Application:** Groove

**Instructions:**

1. Determine who will participate in what roles. Not everyone who participates needs to be an evaluator.
2. List sequences of actions or subtasks that cover all evaluated role relationships and include both typical situations and situations where violating the awareness and privacy policy would have a high impact.

<b>Participants</b>		
Name	Role	Evaluator?
Rod	Product Manager (PM)	No
Jill	Buyer (B)	Yes
Jean	Sales Rep (SR)	Yes

<b>Master Scenario</b>			
Time	Activity	Roles	Verify privacy/awareness needs:
8:00	Discuss the new juice flavors	PM, B, SR	B cannot see juice recipes
8:20	Share manufacturing process info	PM, B, SR	B cannot see juice recipes
8:40	Discuss pricing strategies	B, SR	B cannot see SR's bottom price; SR cannot see B's top price

**Figure 6.** Example portion of master scenario worksheet developed for the case study.

**CASE STUDY: EVALUATING GROOVE USING SCAPE**

We evaluated Groove™ to understand how well it can support the awareness and privacy needs of a retail buyer/supplier team. Groove was chosen because there is currently a lot of interest in using Groove (Groove's customer list, as published on its web site, shows approximately 100 organizations, including large and small businesses, non-profits, educational, and government entities). Also, Groove's functionality is similar to that of

several others in its class; it is a general-purpose collaboration application intended to be used by a wide variety of organizations.

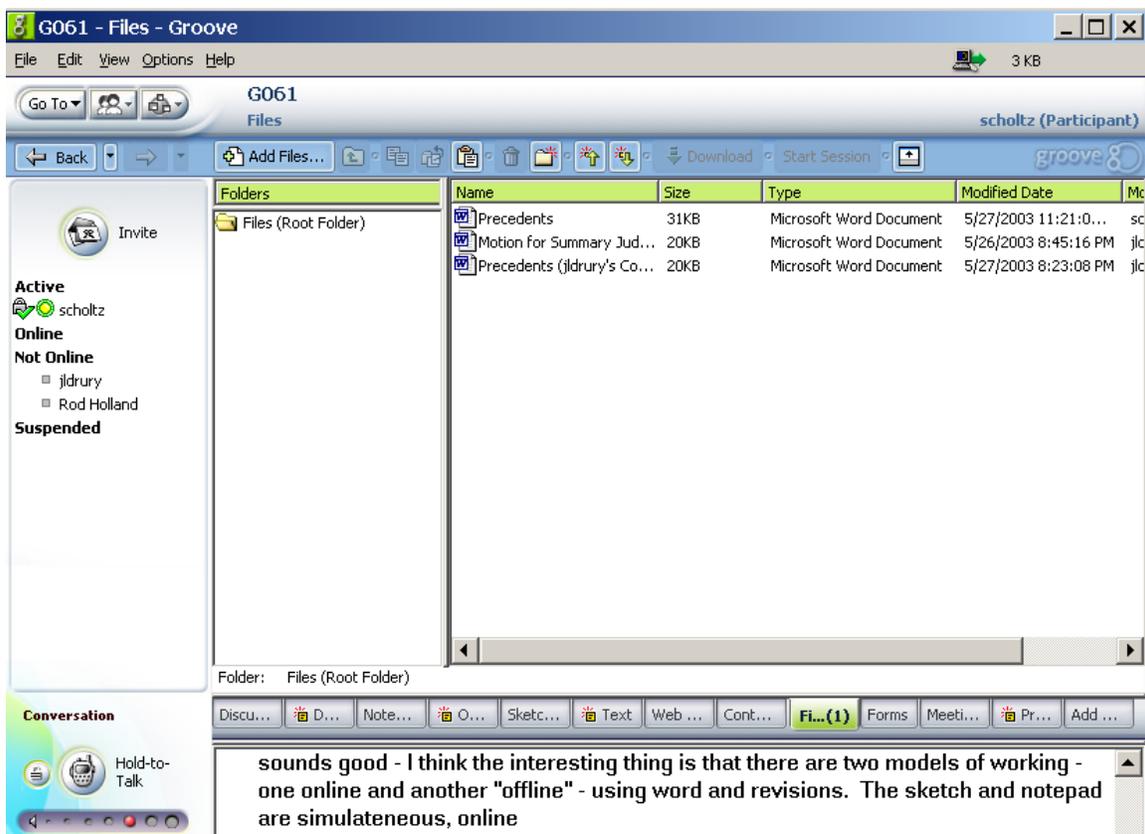
After the Groove overview, we discuss how we performed the case study and summarize its results.

### **Overview of Groove Functionality**

Groove peer-to-peer collaboration software provides users with synchronous as well as asynchronous collaboration. Groove uses the metaphor of shared spaces to create collaboration environments. Templates are provided for spaces and a number of different functionalities can be incorporated into a space. Functions provided by Groove include file storage, calendar facilities, discussion spaces, sketch pads, project planning, and meeting spaces with support for agendas, minutes, and action items. A built-in browser lets participants pull up web pages within Groove and browse the pages together.

Many of the tools available in Groove are compatible for use with Microsoft Office; these functions provide asynchronous collaboration. In addition, Groove provides text chat and audio chat for synchronous collaboration. Microsoft NetMeeting can be launched from within Groove as well.

Figure 7 shows a Groove shared space with the files tool open. The participant bar is to the left and the audio chat tool is at the bottom left. A portion of a text chat message is shown at the bottom of the window.



**Figure 7.** A shared Groove space.

Groove provides role-based access for three predefined, standardized roles. A manager sets up the shared space and can invite people as participants or guests. The privileges that one has in a shared space are based upon these roles. In general, guests are allowed read-only privileges while participants are able to post and edit as well as read.

Groove has two models of working. In tools such as the sketchpad multiple users can work together at the same time and they see each other's contributions in real time (as they are made). However, in tools such as the file space, participants open a file that creates a local copy on the user's machine. Groove contains synchronization detection that updates the document when a user stores a modified file back to the file manager.

Document creation and revision is supported via a document revision tool and an outliner. Document revision automatically creates a folder for the participants who have edited a document, separate from the folder for the original document. Documents in Groove are marked as to whether they have been read or not; this feature enables users to keep track of work they have done.

Groove also provides a “navigate together” feature. All participants who have selected the “navigate together” feature will be taken from one tool in the space to another when one of the participants changes location within the space. This feature can be used during a meeting to make sure that everyone is working in the same location.

Groove provides some mechanisms for being aware of the presence, identities, and activities of fellow Groove space users. For example, there is a persistent part of the display showing the participants who are logged in, those who are active, and those who belong to the space but are currently logged off or have been inactive for sometime (suspended). The tool bar shows the number of participants who are viewing any particular tool space. A notice is shown when someone enters a tool space. When users are engaged in text chat, Groove shows when one of the participants is typing.

### **Evaluating Groove Using the SCAPE Method**

Prior to evaluation, we became familiar with Groove functionality and prepared the SCAPE materials described in earlier in this chapter. We (the two evaluators) decided

upon the roles we would play and recruited a third person to assist us. While not an evaluator, this third person agreed to play a role and answer specific questions relating to the levels of privacy and awareness that he observed in the Groove shared spaces during the session.

During the evaluation session, the three of us logged into a shared Groove space from three different geographic locations. We assumed the roles of a buyer, sales representative, and product manager. The buyer created the space and thus was considered the Groove Manager. The other two people were given Groove Participant privileges at first. Over the course of the session, the manager changed their Groove roles to Guest and back to Participant to see how these changes affected their awareness and privacy.

We followed the Master Scenario excerpted in Figure 6, beginning with a discussion of the new juice flavors. After performing the actions indicated in the Master Scenario worksheet, we checked to see if the privacy and awareness requirements highlighted in the worksheet were supported by Groove. We made note of the situations and conditions under which the requirements were not supported.

We communicated with each other nearly continuously during the evaluation session using the Groove chat tool. This tool provides a persistent transcript of all messages; we noted our evaluation observations in chat messages to each other and saved the chat file for post-session analysis. The evaluation took approximately four hours.

## **Results of Evaluating Groove**

We combed through the chat file notes and compared our observations to the Awareness and Privacy Policies and the specific requirements noted in the Master Scenario worksheet. Highlights of the results are below.

Groove does a reasonably good job of providing awareness but does not support many of the privacy needs for the case in which proprietary information needs to be seen by some collaborators but not by others. For example, Groove provides at least read-only privileges to all users for all documents. It is impossible to mark documents as being readable by only a subset of Groove users. This means that the Sea Spray team members cannot work on a document in a shared Groove space and keep that document private from members of Wal-Store's team. This situation violates the privacy need specified as "it is OK to be able to read some documents but not others."

Of course, the problem of keeping some documents private from some collaborators can be resolved by creating new Groove "spaces" for subgroups of collaborators, such as the Sea Spray team members. This approach has the drawback of requiring collaborators to remember what documents and other information was stored in what spaces, and manage version control of documents that are stored in both spaces for the convenience of avoiding moving back and forth between spaces.

Similar to the document privacy issue, chat messages in Groove cannot be addressed to a subset of the shared space users; the messages are displayed to all users. Thus team members from the manufacturer cannot caucus privately among themselves using the shared chat tool. This means that the privacy need specified as “The buyer cannot see the sales representative’s bottom price; the sales representative cannot see the buyer’s top price” cannot be satisfied if the relevant pricing material is included in the Groove space.

Groove does not provide the means for two people to have management privileges in a shared space; thus either the buyer or sales representative would be the manager and the other could be a participant at best. While not violating an awareness or privacy requirement, such a situation would not be palatable to two people who are trying to achieve a status in which they are co-equals in their collaborations.

By default, Groove participants are allowed to invite others into the shared space. Unless the manager revokes this “invitation” privilege for everyone, it is possible for a participant to invite another to be a participant, who invites someone else, etc. Thus it may not be known to everyone who has been invited and may subsequently join a session. Once someone has joined, their user name will be visible to everyone in the space unless that person restricts what group can see their online identity; but by that time the damage may have been done if they were not supposed to have had any access to the materials in the shared Groove space. Imagine the consequences of the relatively uncontrolled invitation mechanism for a sensitive shared space containing pricing or budget documents.

While Groove provides awareness of who has revised documents, it does not provide notification of who has revised other shared artifacts. For example, a meeting agenda can be erased or revised by anyone—even a guest—and there is no indication that such a revision has taken place. In fact, most people would be likely to assume that the agenda was authored by the person who scheduled the meeting; this may not be true. Also, the original author does not have any notification that their work was changed. This situation violates the heuristics in the Awareness Policy: “Show the participant(s) making changes” and “Show the participant(s) who made each historical change.”

## **DISCUSSION**

Based on our knowledge of other collaborative applications, we do not believe that Groove is unique in its lack of privacy support. Privacy is enhanced in applications that provide access control on a document-by-document basis, and that provide the capability to address chat messages to subsets of participants (e.g., a private, “whispered” chat message). In general, however, most collaborative applications are akin to Groove in that they aim to provide support to participants who wish to share everything, all the time. Such an approach is not consistent with the move towards greater computer-aided collaboration across organizations with somewhat divergent goals.

Of all the evaluation methods discussed, SCAPE is the only method to specifically focus on whether an application supports the awareness and privacy needs of a particular group of users. SCAPE does not address basic usability issues such as whether the interface

was designed to be consistent or to have clearly marked exits; inspection techniques developed for single-user applications can provide this type of insight even for multi-user systems and should generally be used. Thus, SCAPE should be used in conjunction with techniques such as Nielsen's heuristic evaluation and function checklists. If an IOIS will be used in situations where awareness and privacy are important concerns, however, a SCAPE analysis would be appropriate and beneficial.

SCAPE does not involve users as evaluation subjects; this is both a strength and a weakness. Usability testing (structured testing with users) can be very expensive and time-consuming but yields rich data. Rather than forgoing evaluation completely, however, an inspection evaluation method constitutes a less expensive and less time-consuming approach. We recommend that usability testing be conducted if doing so involves small numbers of users, the users' environment can be easily recreated or accessed, and the user population is readily available. Even if usability testing is performed and a full SCAPE evaluation or other inspection evaluation is not attempted, we recommend performing the first three steps of SCAPE to understand the roles that users play and the need to provide information to, or conceal information from, each other. Once SCAPE Awareness and Privacy Policies have been developed, they can act as a specification against which the application's performance can be compared during a usability test. Awareness and Privacy Policies can also provide guidance for which tasks the users should be asked to perform to ensure that situations are exercised in which sensitive or critical information must be revealed or concealed.

If the buyer and sales representative in the case study scenario were truly looking for an IOIS to manage information sharing in an adversarial collaboration environment, they would be well advised to perform SCAPE evaluations on several systems as part of a comparative analysis. At a minimum, the most sensitive situations should be included in scenarios developed for the SCAPE evaluation. For example, SCAPE quickly highlighted the fact that Groove does not provide for documents to be read by only a subset of the application's users; this fact alone means that Groove may not be a suitable choice for this particular group of users.

## REFERENCES

- Baker, K, Greenberg, S. and Gutwin, C. (2001). Heuristic Evaluation of Groupware based on the Mechanics of Collaboration. In M.R. Little and L. Nigay (eds) Engineering for Human-Computer Interaction (8th IFIP International Conference, EHCI 2001, Toronto, Canada, May 2001), Lecture Notes in Computer Science Vol 2254, 123-139, Springer-Verlag.
- Bianco, A. and Zellner, W. (2003). Is Wal-Mart Too Powerful? McGraw-Hill: *Business Week*, 6 October 2003, 100 - 110.
- Cash, J. I., Jr., and Konsynski, B. R. (1985). "IS Redraws Competitive Boundaries," *Harvard Business Review* (63:2), pp. 134-142.
- Cohen, A. L., Cash, D. and Muller, M. J. (2000). Designing to support adversarial collaboration. ACM: Proceedings of the Computer Supported Cooperative Work (CSCW) 2000 conference , Philadelphia, PA, 31 - 39.
- Dourish, P. and Bellotti, V. (1992). Awareness and Coordination in Shared Workspaces. ACM: Proceedings of the Computer Supported Cooperative Work (CSCW) '92 conference , Toronto, Canada, 107 - 114.
- Drury, J., Damianos, L., Fanderclai, T., Hirschmann, L., Kurtz, J., and Linton, F. (1999). Methodology for Evaluation of Collaborative Systems, v. 4.0. DARPA Intelligent Collaboration and Visualization Project: published at <http://zing.ncsl.nist.gov/nist-icv/documents/methodv4.htm>.

Drury, J. (2001). Extending usability inspection evaluation techniques for synchronous collaborative computing applications. Sc.D. thesis, Department of Computer Science, University of Massachusetts Lowell.

Drury, J. and Williams, M. G. (2002). A framework for role-based specification and evaluation of awareness support in synchronous collaborative applications. IEEE: Workshops on Enabling Technologies, Infrastructure for Collaborative Enterprises (WETICE), Pittsburg, PA, 12 - 17.

Greenberg, S., Fitzpatrick, G., Gutwin, C. and Kaplan, S. (2000). Adapting the Locales Framework for Heuristic Evaluation of Groupware. *Australian Journal of Information Systems (AJIS)* 7(2), 102-108.

Greif, I. (1988). Computer-Supported Cooperative Work: A Book of Readings, San Mateo, CA: Morgan Kaufmann.

Groove web site (2003). Groove Networks: [www.groove.net](http://www.groove.net).

Grudin, J. (1988). Why CSCW Applications Fail. ACM: Proceedings of the Computer Supported Cooperative Work (CSCW) 88 conference, Portland, Oregon, 85 - 93.

Gutwin, C. and Greenberg, S. (1998). Effects of awareness support on groupware usability. ACM: Proceedings of the CHI 98 Conference on Human Factors in Computing Systems, Los Angeles, CA, 511 - 518.

Gutwin, C. and Greenberg, S. (2000). The Mechanics of Collaboration: Developing Low Cost Usability Evaluation Methods for Shared Workspaces. IEEE: WETICE Workshop on Collaborative Enterprises, Gaithersburg, MD, 98 - 103.

Gutwin, C., Stark, G., and Greenberg, S. (1995). Support for Workspace Awareness in Educational Groupware. Lawrence Erlbaum Associates: Proceedings of the First CSCL Conference on Computer Supported Collaborative Learning, Bloomington, Indiana, 147-156.

Holtzblatt, K. and Jones, S. (1993). Contextual Inquiry: a participatory technique for systems design. In D. Schuler, and A. Namioka (eds.) *Participatory Design: Principles and Practice*. Hillsdale, NJ: Lawrence Elbaum Associates, 177 - 210.

Hong, I. B., and Kim, C. (1998). Toward a new framework for interorganizational systems: a network configuration perspective. IEEE: Proceedings of the Thirty-First Hawaii International Conference on Systems Science, Volume 4, 92 – 101.

Hudson, S. and Smith, I. (1996). Techniques for Addressing Fundamental Privacy and Disruption Tradeoffs in Awareness Support Systems. ACM: Proceedings of the CSCW 96 Conference on Computer Supported Cooperative Work, 248 – 257.

- John, B.E. and Kieras, D. E. (1994). The GOMS Family of analysis techniques: Tools for design and evaluation. Pittsburgh, PA: Carnegie Mellon University.
- Klein, G. (2000). Cognitive task analysis of teams. In Schraagen, J. M, Chipman, S. F., and Shalin, V. L. (eds), *Cognitive Task Analysis*. Mahwah, N.J.: Lawrence Erlbaum Associates, Inc., 417 – 429.
- Malone, T. W. (1985). Designing organizational interfaces. ACM: Proceedings of the CHI 85 Conference on Human Factors in Computing Systems, San Francisco, CA, 66 - 71.
- Mantei, M. M. and Teorey, T. J. (1988). Cost/Benefit Analysis for Incorporating Human Factors in the Software Lifecycle. *Communications of the ACM* **31**(4): 428 - 439.
- Mayhew, D. (1999). *The Usability Engineering Lifecycle: A Practitioner's Handbook for User Interface Design*. San Francisco, CA: Morgan Kaufmann.
- Molich, R. and Nielsen, J. (1990). Improving a human-computer dialog. *Communications of the ACM* **33**(3): 338-348.
- Nardi, B. A. and O'Day, V. L. (1999). *Information Ecologies: Using Technology with Heart*. Cambridge, MA: MIT Press.
- Nielsen, J. (1994). Enhancing the Explanatory Power of Usability Heuristics. ACM: Proceedings of the CHI 94 Conference on Human Factors in Computing Systems, Boston, MA, 152 - 158.
- Nielsen, J. (1995). Scenarios in Discount Usability. In *Scenario-Based Design: Envisioning Work and Technology in System Development*. J. M. Carroll, ed. New York, N.Y.: John Wiley and Sons, 59 - 83.
- Rogers, E. M. (1995). *Diffusion of Innovations*, 4<sup>th</sup> edition. NY: Free Press.
- Villegas, H. and Williams, M. G. (1997). Benchmarks for Workspace Awareness in Collaborative Environments. International Institute of Informatics and Systemics: Proceedings of the 1997 World Multiconference on Systemics, Cybernetics and Informatics, Caracas, Venezuela, 480 - 486.

---

<sup>1</sup> The identification of any commercial product or trade name does not imply endorsement or recommendation. Groove is a trademark of Groove Networks.