

# Security for Network Places

Tim Kindberg

Dept. of Computer Science, Queen Mary & Westfield College, University of London  
London E1 4NS. timk@dcs.qmw.ac.uk

## 1. Introduction

The Mushroom project<sup>1</sup> is developing a framework to support network ‘places’ for wide-area collaboration and interaction [1-4]. A network place is a persistent environment in which a collection of users can share information objects and communicate with one another. We envisage a world consisting of many places, each supporting a group of users performing a specific task or closely related set of tasks. For example, a group of authors spread around the world can use a network place to manage drafts of a book on which they are collaborating; a software company can use a place to hold sources and design documents for each piece of software that its geographically distributed programmers are developing.

Much of the information that users store in places and their communications within them are confidential: that is, users require that they be kept secret. Users also require that the integrity of their information is maintained. We are currently investigating the use of network places to support diabetic patient care, where it is essential that both the confidentiality and integrity of patient records are maintained. Similarly, companies need their workers to collaborate without fear of eavesdropping by their competitors, or fear of working with corrupted or inconsistent information. To achieve these goals, users apply security and integrity policies which allow working only with users they can trust, and which take into account the degree of trust. Collaborators require the system to enforce their secrecy and integrity policies against attack, and also to provide an audit trail in case security is breached.

This paper discusses the issue of security for network places, and outlines our approach. Under the heading of security, we include both maintaining the secrecy of stored and transmitted information, and maintaining its integrity<sup>2</sup>. There are two key questions which our project is exploring:

**How to provide an access control model for naive computer users** such as doctors, who require discretionary access control [5]. Existing security models for group working are often so complex even their designers have trouble elaborating them [6]. Ellis et al [7] state: “Changing an object’s access permissions should, for example, be as easy as dragging the object from one container to another.” This paper describes how we have taken up that challenge. Our approach is based on *boundaries*. A network place has a boundary object, which implements a policy for the movement of objects and communications across itself, and for the set of operations that a user may perform in the place. When security is applied, no-one can interact with any of the place’s objects, or communicate or listen within the place, unless s/he has entered the place. Once inside, limits may be placed on what s/he can do.

**How to protect the past.** Records of past events and object states are often of interest in collaborative activity. First, users need to be able to undo actions. Second, they review object histories and derive alternative objects from previous versions. Third, users need to be able to work disconnectedly, storing the operations or intermediate versions, and merging their work on reconnection. Mushroom stores previous versions of some objects. Moreover, it stores updates, in the form of objects called *events*, which Mushroom uses to propagate changes. Protecting historic data is an important issue with some subtle ramifications. Consider a doctor who writes a comment into a patient’s record, and subsequently deletes it in order not to reveal it to a newly-joined clinician. The new clinician must not have access to the object’s history, if the comment is to remain secret, for it will contain the event that described the original comment.

---

<sup>1</sup>UK EPSRC grants GR/L14602 (1996-97), and GR/L64300 (1998-2000).

<sup>2</sup>Integrity is not only a security matter: it may be breached by friendly accesses, for example in the absence of adequate concurrency control. But we concentrate here on security.

In Section 2 we describe and motivate the features of the Mushroom system needed to understand the security requirements. Section 3 describes the access control model and discusses the issues raised, particularly with respect to securing events. We conclude with a summary and outline of the open issues.

## 2. The Mushroom system

Users perform their work by sharing objects in places. They either work synchronously (in real-time, often while communicating over an audio or video link), or asynchronously (at different times) and even disconnectedly. The shared objects are written in Java and we shall refer to them as applets<sup>3</sup>. A simple applet is either a wrapper for a legacy document such as a word processing file or a web page, or a group-aware tool such as an electronic whiteboard or patient record. Other applets provide links to places. In particular, a place that is reachable only via another place is called a nested place. At the system level, each place is an object which maintains a directory of the applets it contains.

For reasons of security or otherwise, each object (place or applet) may exist in several *flavours*. Objects of different flavours vary in their functionality or state but perform related functions. Developers implement flavours either using several classes or using a single class that may function as any particular flavour. An instance of a particular flavour is called an *alternate* of the object<sup>4</sup>. For example, clinician A may view and keep a local copy of a patient's record in its entirety, whereas clinician B may only use an alternate which does not contain certain sensitive information, or on which certain operations are not possible. The clinicians think of themselves as sharing a single object (although they may realise that it appears differently to each), but in reality each has his or her own alternate.

Updates are communicated in Mushroom as events, which are delivered wherever an object has subscribed to them. We say that an object *applies* an event that it handles. An event [8] is an object which describes a state of affairs, for example "Sid added the paragraph p at position x in document d", encoded as a set of values. Events differ from object invocation messages in that they do not necessarily designate a particular target object, and they are first-class objects which can be stored and queried.

Events allow us to manage heterogeneous alternates of the same applet. The alternates of the same applet can raise and respond (differently, if need be) to a common set of events in order to make their state consistent. Another advantage of events is that they provide a convenient way to schedule operations upon alternates. If we order event-delivery according to the consistency constraints that apply in the application, then workstations and servers need only apply concurrency control locally to ensure consistent updates [9]. Finally, since events are first-class objects, we can construct services which process them on behalf of applications. For example, services can log and replay them [12].

Places require the following system support:

**An active store.** Persistent and replicated storage is needed for the place, its applets' state, and events, to allow asynchronous and disconnected working. We may replicate alternates between several servers and client sites, in order to achieve high availability. Thus what appears to the users as a single object is represented by alternates of several flavours, each of which may be one of several replicas. When we speak of the 'alternates' of an object, we include all replicas. The store needs to be active, first, since it must supply appropriate alternate state on request, which is typically achieved by having a comprehensive alternate produce state of the appropriate flavour. For example, an alternate that contains the entire patient record can produce a restricted form appropriate for a nurse. The second reason that the store needs to be active is that object histories need to be *agreed*. Users can make incompatible changes to an object when working concurrently. To support sharing, we must define what is an alternate of the same shared object, as opposed to what must be regarded as a derived version of the object with a separate identity. Agreement defines the set of events that belong to any legitimate alternate of the object, and any ordering constraints that apply. Agreement is sometimes carried out at a single site (typically a server), but it need not be so. For example, contributions to a 'chat' applet could be agreed by default. And the agreed ordering is not necessarily total: some applets have more relaxed consistency constraints.

---

<sup>3</sup>Our applets are not the same as instances of the Java Applet class, but are similar in terms of being downloadable active content.

<sup>4</sup>An alternate is similar to the notion of view in the database world, but in CSCW 'view' typically means a user-interface representation of a shared object.

**Shared channels for audio and video communication and events.** Both media and events need to be pushed on a best-effort basis to users who are collaborating in real-time, in order that they can interact. These channels are manufactured and destroyed as they are required.

### 3. Securing places

Securing a place amounts to securing its active store and the event and media channels. Our design goals are:

- **Providing security** with respect not only to the state of objects in the place, but also with respect to what are the agreed events, and who may access them.
- **Supporting dynamic groups of users.** We allow both the set of users who can access objects in the place (the *members* of a place), and their roles within the place, to change. An important problem is how to maintain the consistency of the security attributes of a place.
- **User-friendly operation.** This is not just a convenience. A framework that makes security policies difficult to implement, and security attributes difficult to keep track of, is liable to lead to mistakes and therefore vulnerability to security attacks.

#### Access control model

We are investigating a container-based access control model for places [15], on the assumption that this is more tractable for users such as clinicians than per-object access control lists. A place may be made into a secure container by giving it a *boundary object*, with methods *import* to bring an object into the place and *export* to obtain an alternate of the object which those who may not enter the place may use. (Of course, a boundary may refuse an export or import altogether. But we cannot secure a place against leaking by one of the collaborators.) To give a user access to the place is to import an object representing the user into the place. Only those with a representation in the place may access the objects within it. Users can see at a glance who has access to the objects within a place. Recall that places may be nested. One must be a member of the enclosing place to access a nested place. Consider a place for keeping data pertaining to a patient. Data which may be read by all clinicians is kept at the top-level place. More secure data – for example, data which is intended only for the patient and the GP to see, and not nurses or other doctors – is kept in a secure nested place.

It is easy to see how read-only data can be protected using such a scheme. Standard classifications such as “top secret”, “confidential” and “unclassified” can be realised using nested places to store the respective objects. How can we give access rights to those who need to update objects? Protocols exist for users who have the same rights to shared objects [11], but users such as clinicians require different rights according to their role [14]. We are investigating a model in which a user has uniform rights to all objects within a place (but not a nested place), where these rights take the form of one of a small number of types, for example *edit* (read/overwrite), *annotate* (read/add), *read*, *append* (add), *manage* (read/synthesise events) and *own* (change place’s ACL). The rights pertaining to a principal (an individual or a user playing a role) need to be mapped onto the system’s right set of *read*, *append*, *edit* etc. For example, in a place containing certain items from a patient record, those with the role of a nurse who is only able to enter data may have at most append rights, whereas a GP may have edit rights and a consultant annotate rights.

We do not yet know whether this model will in fact suffice for the clinicians. Boundary-based security is inconvenient for changing the access to one of a group of objects, although by moving objects between places any combination of rights can be achieved. However, we believe that boundaries have certain advantages even if, eventually, per-object ACLs prove necessary:

- boundaries can be used to maintain the integrity of collections of objects, through their ‘import’ functionality. For example, a boundary could enforce the integrity constraint “data pertaining to no more than one patient may exist in place P”.
- they correspond intuitively to a model of paper security which is familiar to our users: papers are kept in locked filing cabinets inside locked rooms. Network places have the advantage over physical places of being able to represent to the user precisely who is able to enter each place, and what type of updates, if any, they can perform on objects in that place.
- Even if per-object ACLs were to be used, it is useful to give rights only to principals of the form (Tim **and** memberOfPlace) – that is, Tim must be a member of the place to access any object within it. One has only to remove membership to deny all access to a place and to places nested within it.

Note that this model is independent of the user-interface representation. A patient record could in principle appear in the form of a set of WWW pages which are in a collection of nested containers. When a user accesses the root page,

what she sees (the alternate she retrieves) depends upon who she is and/or what role she plays. She may not see a link at all to certain information. On following a link, she may or may not be able to update the data referred to. Each 'page' (sub-object) is in a secure container, but this fact is relevant only to an owner of the root place. An owner could be given a container-like view in order to set the protection settings.

## Securing alternates and events

The basic features of the security design are:

- trust only certain servers and user machines operated by certain principals to supply information; the degree of trust may vary – clinicians tend to be distrustful of servers which aggregate data [5]
- ensure that only those with read rights may read alternates (of a flavour appropriate to the requesting principal) and events
- provide secure event agreement, in particular with respect to places, so that the security attributes of a place are unambiguous for any point on its time line; and only events that do not conflict with the security attributes appropriate for the place and time will be applied by 'good' principals
- allow the owners of a place (who may set its security attributes) to bound a user's access rights so that he may obtain no alternate state or event prior to a given time
- prevent all principals from altering history – changing the set of past events or substituting saved alternate state.

The strategy to secure objects and events is:

- events and snapshots of alternate state are stored and transmitted in encrypted form
- snapshots of alternates are signed so that the recipient can check that they were manufactured by a trusted party
- conformant sites (i.e. those that are operated by trusted parties) securely check alternate snapshots, and events before applying them to alternates.

In order to consider what it means to check the legitimacy of an event before applying it, let us first distinguish between events of different types:

**update events** describe an update to an individual object, such as "Tim added paragraph P to D at time T"; on application to an alternate, they are equivalent to an operation upon the object

**relational events** make an assertion about one or more objects, but do not imply that the announcer made an update. For example, "Tim entered place P at time T".

In the first case we need to check that Tim had the right to add the paragraph at time T. An applet maps the update events it consumes into one of the categories *append*, *annotate* etc. Each alternate checks that Tim had the corresponding rights in the place – which means that it must have its own local copy of the place's security attributes. In the case of relational events, the question is whether the announcer has (had) *jurisdiction* over the assertion made by the event. An applet such as a log which applies this event is not concerned about any rights over itself: it is concerned about whether it is recording a valid state of affairs. To support this case, a principal may be granted *manage* rights within a container. This means that that principal may read objects within the container and announce relational events. These rights imply more privilege than read rights, since others are asked to believe the accuracy of what is described. They may apply non-trivial updates in response to a relational event. For example, if two objects in a virtual world are told that they have collided with one another, then they may each respond by updating their shapes.

We now turn to the questions of managing the evolution of security attributes over time, and of controlling access to historic data. To deal with the latter first, consider a history object which contains a version tree of snapshots of the development of one or more objects, and the events that map the development. For example, a place could have a history, or a document within a place. The problem of securing the past reduces to the problem of securing history objects, which are just another type of object in a place. We do not have space to consider this problem further here, but some obvious policies that users may require are:

- No user may alter any event or snapshot once it has been stored in a history.
- No user may insert an event before an existing event in the history.
- Only trusted users may delete parts of a history, and such parts must correspond to the history's tail (so that no gaps are produced).

- Users are given access to the history object only from a certain point in time – this deals with the problem of securing events (and saved snapshots) which record confidential information that has now been deleted.

We must solve the problem of managing the evolution of security attributes while bearing in mind that users may not want events to be totally ordered, and while allowing historic data to be retrieved from any of several sites that may have ordered events differently. We pointed out above that agreement may allow more relaxed ordering than total ordering – for example, the items in a shared bulletin board do not have to have the same ordering everywhere.

How to achieve secure evolution is an open question, which we are currently addressing. One approach worthy of investigation is to order the ‘security’ events for each place totally with respect to other events. Security events are ones which announce membership changes, rights changes and users joining and leaving active work sessions. We require that the order of such events and the set of events that occur between any two of them is well-defined, but the ordering of the intermediate events is not constrained for security purposes. A principal’s events will only be agreed with logical timestamps sanctioned by a session certificate. The certificate is obtained when applying to start a work session, or when submitting events after a disconnected work session. Sanctioned timestamps are labelled with an ‘epoch’ number drawn from a non-repeating series. Each epoch (work session) is securely terminated.

In order to apply an event securely, a site first verifies that the event has been agreed (signed) by an appropriate principal, extracts its timestamp, and verifies that the announcing principal held the corresponding rights at that time. The site needs to verify the existence of a certificate issuing the rights – one which is not contradicted by a subsequent certificate withdrawing the rights before the time of the event. These verifications, of course, raise several implementation issues which we are exploring.

## 4. Conclusions

We have outlined the Mushroom model of persistent shared places, and the security requirements attaching to places for collaborative activities such as patient care. We are pursuing a boundary-based security model, which we believe to be widely applicable. Boundaries enforce the secure import and export of objects and events between places. Even if per-object access control lists prove desirable (as opposed to container-based security), boundary objects still have some desirable properties. These include being able to constrain access to the objects within a place globally, and being able to manage the place’s integrity.

We have shown that events raise some interesting security issues. First, they are not necessarily the same as ‘operations’. Second, it is often desirable for collaborating users to have access to the past, but care must be taken over what a new member of a collaboration is allowed to access.

Finally, we have raised the issue of how to manage the evolution of security attributes. Data consistency is anyway a thorny problem, when considering how to enable distributed users to work on shared objects, sometimes disconnectedly. How can we ensure that distributed sites apply security policies consistently, while allowing relaxed consistency between (some) shared objects? The answer seems to involve ordering of security-related events, without forcing synchronisation of the system (which would be an unreasonable constraint for most users). We are currently investigating practical solutions.

## References

- [1] T. Kindberg (1996). Mushroom: a framework for collaboration and interaction across the Internet. Proc. CSCW & the Web, 5th ERCIM workshop, Busbach, U., Kerr, D., and Sikkel, K. (eds), GMD, pp. 43-53.
- [2] T. Kindberg (1996). A stake in Cyberspace. Proc. 7th. ACM SIGOPS European Workshop, Connemara, Eire, Sept., pp. 83-88.
- [3] T. Kindberg, G. Coulouris, J. Dollimore and Jyrki Heikkinen (1996). Sharing objects over the Internet: the Mushroom approach. Proc. IEEE Global Internet 1996, London, Nov., pp. 67-71.
- [4] Mushroom project home page: <http://www.dcs.qmw.ac.uk/research/distrib/Mushroom>.
- [5] R. Anderson (1996). A security policy model for clinical information systems. Proc. IEEE Symposium on Security and Privacy 1996.
- [6] K. Sikkel (1997). A group-based authorization model for cooperative systems. Proc. ECSCW97, pp. 345-360.
- [7] C. Ellis, S. Gibbs and Rein, G (1991). Groupware - some issues and experiences. Communications of the ACM, vol. 34, no. 1, pp. 38-58.
- [8] J. Bacon, J. Bates, R. Hayton, and K. Moody (1996), Using Events to Build Distributed Applications. Proc. 7th. ACM SIGOPS European Workshop, Connemara, Eire, Sept., pp. 9-16.

- [9] A. Schiper and M. Raynal (1996), From group communication to transactions in distributed systems. *Comm. ACM.*, vol. 39, no. 4, pp. 84-87.
- [10] A. Rowley and J. Dollimore (1996). Replicated secure shared objects for groupware applications. Tech. rep. 716, Dept. of Comp. Sci., Queen Mary & Westfield College, U. of London, Proc. ERSADS 97.
- [11] L. Gong (1996). Enclaves: enabling secure collaboration over the Internet. Proc. USENIX Unix and network security symposium, San Jose, July.
- [12] W. Edwards and E. Mynatt (1997), Timewarp: techniques for autonomous collaboration. Proc. CHI 97, Mar., pp. 218-225.
- [13] D. Dean, E. Felten, and D. Wallach (1996), Java security: from HotJava to Netscape and beyond. Proc. IEEE Symp. on Security and Privacy, May.
- [14] G. Coulouris, J. Dollimore and M. Roberts (1998), Secure communication in non-uniform trust environments. Submission to ECOOP Workshop on Distributed Object Security, Brussels, 1998.
- [15] S. Greenwald and R. Newman-Wolfe (1994), The distributed compartment model for resource management and access control. Tech report TR94-035, dept. of Computer and Information Sciences, Univ. of Florida.