

Mobile Java objects

R.J.Hayton, M.H.Bursell, D.I.Donaldson, A.J.Herbert

APM Ltd.

Poseidon House, Castle Park,

Cambridge CB4 3JB

Tel: +44 1223 713111 Fax: +44 1223 359779

Email: apm@ansa.co.uk

Abstract

In this paper we discuss the engineering requirements for adding object mobility to the Java programming language, and give an overview of the design and implementation of our mobile object system. We show that it is helpful to cluster objects for mobility, and that if these clusters represent untrusted pieces of code (for example Agents) then they must be encapsulated both to control their access and to control access to them. We show that managing large numbers of mobile objects in an open environment is a difficult problem, but has its roots in the management of large distributed name spaces. We propose an architecture for re-locating moved objects that is both scaleable and tuneable.

The mobile object system we describe has been implemented, and is currently in use as part of an ESPRIT agent project. We are currently evolving the design and implementation to provide additional security and distribution facilities.

Keywords

Java, mobile objects, agents, security, middleware, binding

1 INTRODUCTION

Java is an ideal language for developing distributed applications. It provides both object and interface abstractions, which gives a useful distinction between an object's interface and implementation. It also provides language level introspection, and allows for dynamic code creation. These features make it particularly suited to the design of middleware systems.

With the design of our middleware platform, FlexiNet, we took the approach of extending Java language concepts by adding selectively transparent remote invocation to calls on any interface. We used Java's strong typing support and runtime introspection to allow us to build a strongly typed reflexive binding framework. We further used Java's support for runtime code loading to allow us to create stubs transparently on the fly during program execution. This gives a middleware system which is both extremely flexible and a natural extension of Java's features.

When we turned our attention to mobility, we took the view that this too should be a natural extension of the Java language. We designed a system in which we maintained the FlexiNet view of a "sea of objects" but which allowed objects to move between hosts. Clients of mobile objects need not be aware if a service object is mobile, nor if it has actually moved; they simply continue to use Java references to exported interfaces.

This approach has several advantages. As well as providing a straightforward and familiar programming paradigm, we remain within the well-understood domain of distributed systems. This allows us to leverage existing research when tackling issues of scalability, robustness and security.

2 RELATED WORK

There are an ever increasing number of "agent" systems available, either commercially or as academic projects. Although some provide high level abstractions that might be considered as support for intelligent behaviour, the majority of these systems are basically mobile code platforms. They allow applet-like pieces of code to move through a distributed system, and allow varying degrees of communication between different agents and between agents and hosts or other resources. Examples of these systems include Telescript (GenMagic 1995) and Aglets (Lange 1996).

In general, communications in these systems takes place using messages sent between the various parties. We believe that the advantages of object oriented computing, and in particular distributed object computing, are well accepted, and it is equally advantageous to build mobile systems using objects. Until the advent of the Java programming language, it was difficult to create "deployable" objects so most mobile code systems relied on scripting technology. It is only recently that the possibility of combining mobility and object orientation has emerged.

Existing Java systems that claim some degree of object mobility, do so either by abandoning Java's method invocation abstraction (and taking the message passing route) or by utilising Sun's remote method invocation (Sun 1996). Unfortunately when designing RMI, Sun dictated that each object must be explicitly tagged as a server object or as a data object. Only data objects may be passed by value on a remote invocation. This precludes the possibility of moving server objects, and is a show stopper for most mobile object systems.

Voyager (ObjectSpace 1997), is an exception to this. This is a Java middleware platform that supports mobile services by wrapping them with pre-processor generated classes. However, objects are moved individually and there is no notion of clustering. Although Voyager is closest system to ours, we believe we have a more flexible and scaleable approach. In FlexiNet, an introspection based serialization engine is used, which does not have the limitations of RMI's serializer, and is more portable as it is written using 100% Java.

FlexiNet embraces ODP principles (ODP 1995), and by adding mobile objects to FlexiNet we are providing the ODP relocation and migration transparencies that most other middleware platforms lack.

3 FLEXINET

The FlexiNet Platform is a Java middleware system built as part of a larger project to address some of the issues of configurable middleware and application deployment (APM 1996). Its key feature is a component based 'white-box' approach with strong emphasis placed on reflection and introspection. This allows programmers to tailor the platform for a particular application domain or deployment scenario.

Interfaces on remote objects are represented by **proxies**. Proxies enforce the typing of the remote interface, and perform remote access by utilising **binders**. Each binder is an object capable of creating a *generic* binding between a local proxy object and the remote object it represents. Different binders embody different application requirements or engineering strategies such as choice of protocols and the imposition of security policies or atomicity. Binders may make use of other binders in a recursive way. This keeps individual binders small, and allows application domain-specific binders to be easily created. To manage the flexibility, FlexiNet supports the notion of **multiple name spaces** for interfaces. Names are both strongly typed and structured. Names may be constructed out of other names, or arbitrary data, making the management of aggregate and indirected names straightforward.

3.1 Generic Communications

Rather than using stubs to convert an invocation directly into a byte array representation, we leverage Java's runtime typing support to represent the

invocation in a generic (but fully typed) form. This is a reflexive technique central to the design of FlexiNet. The layers of the FlexiNet communication stack may then be viewed as reflexive meta-objects that manipulate the invocation before it is ultimately invoked on the destination object using Java core reflection.

This approach allows middleware (or application) components to examine and modify the parameters and semantics of the invocation using the full Java language typing support. Third part meta-objects can be fully general and are fully type-safe. When designing the Mobile Object Workbench, we decided to implement it as a set of FlexiNet binders, binding protocols and services. In particular the mechanism for communication with a mobile object is essentially communication with a static object together with two reflective layers: a server side layer, which raises an exception if an object has moved, and a client side layer, which relocates and rebinds to the object.

4 REQUIREMENTS FOR MOBILE OBJECTS

4.1 Unbinding

The key feature of a mobile object is that it must be able to move. When we move an object, we effectively copy it to a new location, and then arrange that all references to the old object are replaced with references to the new object. This requires a mechanisms for *unbinding* a previous binding to an object. If objects are referenced directly using language level pointers, unbinding would require changes to the implementation of the Java virtual machine, reducing the advantages of Java as a portable language. Instead we arrange that mobile objects (or more correctly interfaces on mobile objects) are referred to indirectly, using FlexiNet stubs as proxies. This level of indirection allows us to rebind references dynamically whenever objects move.

In addition, to avoid the need to track distributed references, we only perform this rebinding when a reference to a moved object is first de-referenced. The use of stubs and the rebind is transparent to the application programmer, although they may reflect the process, for example to deal with errors.

4.2 Consistency and Threads

At any point in time, an object may be *active* or *passive*. An active object is one that has a thread of control currently executing in it, or passing through it. A passive object is one that is not currently being executed, and hence has no threads active in it. When we move an object we must ensure that the move is *atomic*. To do this the object must be *passivated* - i.e. all processing of the object must be suspended. One approach to this would be to pause any threads running in an object, move the object, and then restart the object and the threads at the new location. This, unfortunately, is impractical, as Java does not allow us to determine

the complete state of an active thread at an arbitrary point of execution.

We have chosen to *encapsulate* the object and enforce a locking strategy which ensures that no threads are executing the object at the point of movement. This does not prevent the existence of active mobile objects, or those that contain completely internal threads, but it does require a degree of co-operation with such objects, so that they can be 'shut down' prior to movement, and then 'restarted' at the new location.

The encapsulation mechanism is integrated with the mechanism for transparent re-binding, so that external threads that have blocked pending an object's movement restart and relocate the newly moved object.

4.3 Grouping

In the discussion so far, we have been describing the migration of single Java objects. However there is little utility in moving a single Java object as they are typically very small (for example a linked list will consist of many objects). A more useful unit for mobility is a set of related objects, and we need a mechanism for deciding which parts of a program should move together.

We introduce the notion of a *cluster* as both a grouping and encapsulating construct. References that pass across a cluster boundary are treated differently from those entirely internal or external to it. In particular, when resolving an external reference, the system may have to locate a cluster on a remote machine (possibly after it has moved). References entirely within a cluster can be ordinary Java language references, as no special action needs to be taken when they are de-referenced.

To a programmer, clusters are a straightforward concept. A special function is used to create the initial object populating a cluster, and after this any new object is created in the same cluster as its creator. Generally, clusters are completely transparent to the programmer.

4.4 Failure Modes

When designing distributed systems, there is always the possibility of host or network failure. In particular network partition can result in hosts incorrectly assuming that other hosts have failed. When designing a mobile object system, a key decision is the semantics in the worst case scenario of a network partition during object migration. There are three possibilities. We could allow the possibility of the object existing on both sides of the partition - this was rejected as it introduces an unwanted degree of complexity. The second possibility is to ensure that an object is destroyed if it cannot be uniquely determined which side of a partition it exists in. This is the default semantics chosen in the Mobile Object Workbench. The third possibility is to suspend use of the object until the network is restored. This is being considered in the context of the integration of transactional

mechanisms into FlexiNet. Section 6.4 gives a fuller description of the state transitions required to ensure consistent movement.

4.5 Scalability

There are two issues relating to the scalability of a mobile object system. Firstly, some design choices would require the registration either of all objects, or worse, of all references to all objects. We rejected these approaches as we wish to use the mobile object workbench in an Internet environment, which is both open and has no central administration. The second issue relates to the rebinding to interfaces on an object once it has moved. We would like this to be possible, even if the original host has since failed. This issue is discussed in detail in section 8.

5 REQUIREMENTS FOR MOBILE AGENTS

The mobile object workbench is not a mobile agent system, however it was developed as part of an ESPRIT agent project called FollowMe (FAST 1997), and one of the other partners is developing an agent system on top of it. As mobile agents are an obvious application of mobile objects, it is worthwhile to consider their specific requirements.

5.1 Autonomy

Mobile agents are generally considered to be 'autonomous'. That is to say that it is the agent itself that determines the actions it takes, and in particular controls its movement. In terms of mobile objects and clusters, this gives a requirement for cluster mobility to be initiated only by the cluster itself. The encapsulation mechanism gives provision for this; only threads within a cluster have access to the objects within it, and by giving one of these objects a handle to the infrastructure which controls mobility, this is effectively hidden from the outside world. There are two exceptions to this. Firstly, a malicious implementation of the infrastructure can overcome the FlexiNet encapsulation mechanisms; this is a necessary evil of distributed computing - you have to trust the host. The second exception is that a host may 'legally' destroy an object and reclaim the resources it is using. This is necessary to allow hosts to be manage their own resources. The normal procedure is for a host to inform a cluster that destruction is imminent, in order to allow it to move or shut down cleanly, but a host must always be able to perform a 'dirty' shutdown in order to protect itself from malicious or erroneous agents.

5.2 Security

We have found that the approach of designing and engineering from the point of view of a mobile object system, allowed us build on established security principles,

in particular we identify six basic areas of security concern within the Mobile Object Workbench:

1. **Host integrity** - protecting the integrity of a hosting machine and data it contains from possible malicious acts by visiting objects.
2. **Cluster integrity** - it should be possible to determine if a cluster has been tampered with, either in transit or by a host at which it was previously located. We may wish to allow hosts to modify parts of a cluster (e.g. data) but not others (e.g. code).
3. **Cluster confidentiality** - a cluster may wish to carry with it information that should not be readable by other clusters, or by (some) of the hosts which it visits.
4. **Cluster authority** - a cluster should be able to carry authority with it, for example a user's privileges, or credit card details. To provide this we need both cluster integrity and cluster confidentiality.
5. **Access control** - a host should be able to impose different access privileges on different clusters that move to it. Clusters and hosts should also be able to enforce access control on exported methods.
6. **Secure communications** - clusters and hosts should be able to communicate using confidential and/or authenticated communication. Some applications may also require other security communication features, such as non-repudiation.

We believe that unless all of these aspects of security are addressed, any mobile object system will not prove secure enough for real world applications, and we have therefore adopted the principle of including security issues from the outset, rather than as an "add-on", bolted on at a later date. Section 7 discusses our approach to these issues.

5.3 Thread Encapsulation

As clusters representing agents are mutually distrustful pieces of code, it is important that one cluster cannot adversely affect another. In particular one cluster must not be able to invoke a method on a second cluster, and then destroy the thread performing the call, so as to leave the second cluster in an inconsistent state. Equally, if a cluster crashes or intentionally blocks whilst servicing a request, the client must be able to recover, and must not also fail or block indefinitely.

In order to achieve this degree of strong encapsulation, we de-couple all threads that enter or leave a cluster, by spawning and rendezvous, so that the failure of the caller and callee are independent. This thread de-coupling is integrated with the binding system and is transparent to the application programmer.

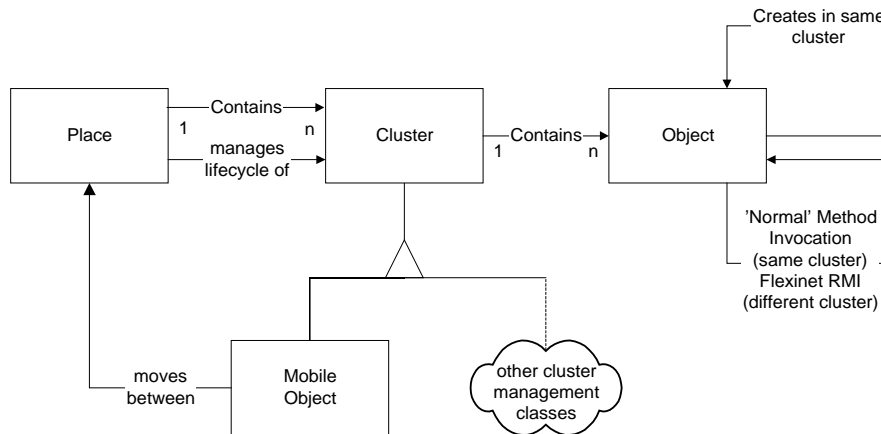


Figure 1 Objects, Clusters and Mobile Objects.

6 THE MOBILE OBJECT WORKBENCH

6.1 Clusters

Figure 1 shows the relationship between (Java) objects, Clusters and Mobile Objects. A Cluster is a Java object containing a grouping of objects which are managed together. A Mobile Object is a specialisation of this which is able to move between Places. Places are themselves objects which abstract execution environments, typically with one Place per JVM. Protection, movement, destruction, charging and other management functions are considered in terms of the lifecycle of Clusters and the interaction between them. It is sometimes useful to consider a Cluster and its contents as a virtual process, and the encapsulation and security concerns around Clusters encourage this view.

An object is the basic building block out of which applications may be built. Objects may contain references to interfaces on other objects anywhere in the system. Objects may directly create other objects, but only within the same Cluster. They may be able to arrange the creation of objects in other Clusters via communication with a Place. Within a Cluster, access to methods/data on objects is determined by standard Java language protection means and takes place using standard Java method invocation. Between Clusters, encapsulation is enforced so that object in one Cluster may only access methods on objects in other Clusters if these methods form part of an interface passed between the clusters.

As communication between mobile (and non mobile) objects takes place using application-level exported interfaces, the Mobile Object Workbench API is entirely concerned with the lifecycle of mobile objects, and is analogous to the Applet or Bean API. The key classes and methods are illustrated in Figure 2.

public class Cluster

```
public synchronized void lock()  
public synchronized void unlock()
```

Change the number of locks held on the object. Whilst a lock is held, new calls made on objects in this cluster from other clusters will block.

```
public void init()
```

Called upon object instantiation. A subclass that requires initialisation arguments, or wishes to return an interface to its creator, should provide an alternative `init(...)` method.

```
public void stop()
```

A call made by the place if it wishes the cluster to cease processing. The cluster is expected to clean up and then return. When the call returns, the place will destroy the object.

```
public void restart(Exception e)
```

Called after the cluster is restarted. A subclass which wishes to take action after a restart should override this method.

public class MobileObject extends Cluster

```
public void syncMove(Place dest)
```

Request a move to the identified place. The current thread will attempt to perform the move. If successful it will exit. The move will not take place until there are no other threads within the cluster.

public interface Place

```
public Tagged newCluster(Class cls, Object[] args)
```

Create a new cluster at this place. Once created, `init(arg0, arg1, ...)` will be called on the new object. The `init` method may return an interface, which is passed to the creator

Figure 2 Mobile Object Workbench API.

6.2 Binding architecture

Communication between clusters takes place by remote method invocation using a 'standard' FlexiNet binder, together with two additional reflexive layers. On the client side of an invocation is a "cluster location" layer. This examines the internal name used to represent the interface being accessed, and determines the host on which it resides. The procedure adopted is to try the last known location, and only contact the relocation service upon failure.

On the server side of the communication, there is a reflexive "encapsulation layer". This processes incoming calls, checks that they refer to clusters that are (still) located on the host, and performs the synchronisation required to ensure that the cluster is not in the process of moving. Part of the encapsulation process is to de-couple the calling thread, so that client and sever clusters cannot affect each other by killing or otherwise manipulating the thread. This is illustrated in Figure 3.




Figure 3 Implementation of inter-cluster calls.

6.3 Orchestrating Mobility

Figure 4 illustrates the states that an instance of a mobile object may be in. The object is initially created in state A_1 . This state represents an *active* object that has *one* thread in it (the thread that calls the constructor). When active, the object may create other threads, and methods on its interface may be invoked by objects in other clusters. It will therefore move between active states.

In order to move, a mobile object invokes a `pendMove` or `syncMove` call. Both of these request a move ‘as soon as possible’, the difference being whether the calling thread returns immediately (`pendMove`) or never returns (`syncMove`). When a move call is invoked, the object enters a pending state. These are identical to active states except that the object will be moved as soon as all executing threads exit (i.e. when it enters state P_0). As a side effect of executing a `pendMove` or `syncMove`, the cluster becomes locked. When locked, calls from other clusters block until the cluster is unlocked.

When a mobile object enters the state P_0 it will undergo a series of transitions that may result in the creation of a new mobile object at a different place. The original mobile object will then be discarded (it enters state X). If an error occurs during this process and it can be safely inferred that the new object has not been created, then this object is returned to state A_1 . If the move was initialised by a call

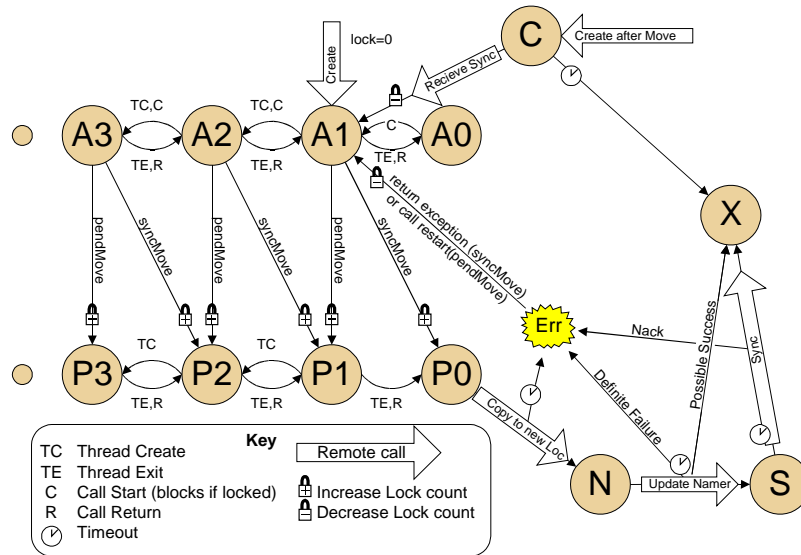


Figure 4 State transitions of a mobile object instance.

to `syncMove`, then the error status is returned as an exception. If the move was initiated by a call to `pendMove`, the object is restarted by calling the `restart` method, and the exception is passed as a parameter.

The newly moved cluster is an exact replica of the original, and in addition all references to interfaces exported by the original cluster are re-mapped to the new one (effectively the original cluster has moved). It is started in state A_1 by a call to `restart`. The new cluster (or original after failure) will have the same lock status as the original - apart from the lock automatically taken when `pendMove` or `syncMove` was called, which is released. If a cluster wishes to restart in a locked state, then it may obtain an additional lock prior to calling `pendMove` or `syncMove`. This allows newly moved clusters to perform start-up cleanly, before allowing external access.

Copying, rather than moving, a cluster follows exactly the same procedure as `syncMove`. The copy operation blocks until there are no other threads and the new cluster has been created, or a failure is detected. After successful synchronisation, or failure, the original object enters state A_1 and the copy operation terminates. The newly created copy commences operation with a call to `copied` in state A_1 .

6.4 Method Invocation

When an object in one cluster attempts to invoke a method on an object in another cluster, it must block if the callee cluster is in the process of moving. The state diagram in Figure 5 indicates the process followed by the infrastructure. It

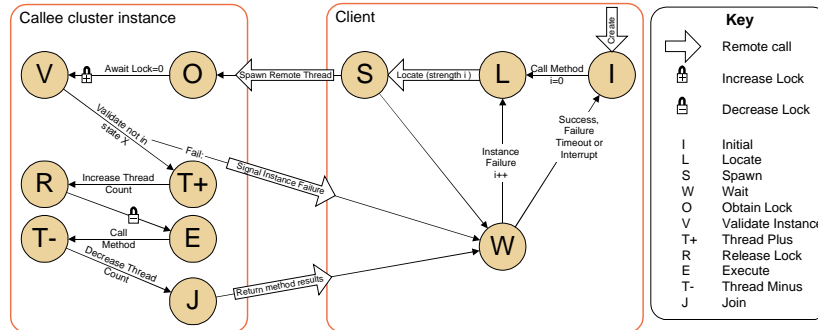


Figure 5 State transitions for an invocation on a remote cluster.

should be noted that the callee is able to interrupt a thread making a call, but that this will not affect the caller. This prevents the caller from blocking the callee's progress.

7 SECURITY

Within a mobile object context, issues of trust take on a different slant to non-mobile systems. In both mobile and non-mobile systems, questions of how much trust is placed in an object must be based on the provenance of that object - where it originated, and its history. In a non-mobile object oriented system, such as the base Java implementation, objects are typically instantiated from class files, having no other state. In Java, these classes may be signed, and a JVM may assign policies to their instantiations based on this signing. In a mobile object context, an object arriving from a remote host has a history that is not captured by the signature on the class file. Class signatures alone do not provide sufficient information on which to assign security policy.

For this reason, we have designed and implemented a Security Manager which extends Java's model by allowing policies to be assigned to instances of objects, rather than just their class. This has been possible because of the strong thread encapsulation we have employed within the Mobile Object Workbench, which gives each cluster its own thread group. As Java allows checking of the thread performing a particular operation, we may determine the cluster from which an invocation originated, and hence enforce the appropriate policies.

This security policy allows hosts to restrict operations allowed by particular clusters, thereby protecting their own integrity. It also provides a good base from which to extend cluster-to-cluster access restrictions.

Cluster integrity and confidentiality are enforced by encrypting and/or signing certain objects within a cluster. This prevents a host without sufficient access privileges from examining a cluster's state, and allows one host to detect if a cluster has been modified by a host that it visited earlier. In addition to this, we

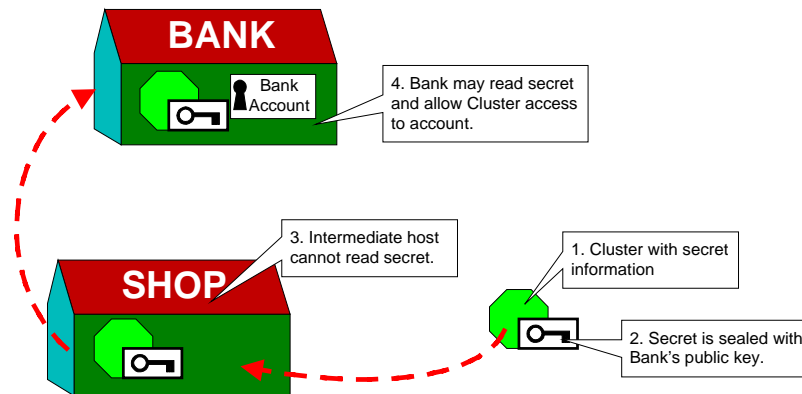


Figure 6 Clusters with Secrets.

must ensure that clusters are not dissected - or a malicious host could 'steal' parts of the cluster that represented encrypted passwords and use them to build its own clusters. To do this, we require a mechanism for specifying, and validating, integrity statements. For example we may annotate a cluster's definition to indicate that a particular field may only be modified by certain hosts. We may then use digital signature techniques to ensure that whenever the field is modified it obtains a signature from the current host, and when other hosts attempt to read this field we can throw an exception if the signature is incorrect.

We are currently developing a system to allow integrity policy statements to be specified. Once specified, the use of secure fields or objects can be made almost transparent to the programmer. All that is required is that they use accessor functions to access the protected fields.

Cluster authority can be implemented using cluster integrity and confidentiality. Together these allow a cluster to carry with it a password or other secret information, without the concern that this secret can be read at any host which is visited. Clearly, once the secret *is* revealed to a host, there is nothing that can be done to prevent the host from misusing it. For this reason we have a model that the mobile object moves into a secure environment before revealing a secret. Figure 6 gives an example; a cluster may move between several hosts before eventually arriving at a 'Bank' host. At this host, it may reveal a password to allow it access to a bank account. However, as the Bank host already knew the password, revealing it has not given the bank any additional privileges, and the security of the password has not been weakened.

Access control and secure communications may be implemented using standard techniques. We use FlexiNet's reflective binding system to allow a cluster to receive notification of an invocation immediately prior to its execution, so that it may implement its own security policy, and throw an access control exception if appropriate. Secure communication between places may take place using a FlexiNet binder that support SSL (Netscape 1996). Secure communication between

mobile clusters may also take place using SSL, but requires that clusters reveal the information use to prove their identity to the host from which they are communicating. This is reasonable in some circumstances, but should be used with caution.

8 NAME RELOCATION SERVICE

When a call is made on a mobile cluster, the encapsulation layer at the called host will determine whether the cluster (still) exists at this host. If it does not, then the host will raise an exception which is passed back to the callee and caught in the callee's locate layer. This then contacts the name relocation service to determine the new location of the object. The relocation service is a federation of a number of directories. Each directory contains a mapping from old to new cluster addresses. Our naming service was developed with five key properties:

1. We control what entities are able to update the directories - only hosts from which a cluster is moving may update the record for the cluster. This is possible as Cluster names (transparently to the applications programmer) contain information about their current network host. This prevents fraudulent changing of naming records by "spoof" hosts or clusters.
2. We provide a hierarchy of directories, for scale and robustness. This means that an instance of the relocation service may decide to copy the naming record for a cluster up the hierarchy to increase its stability, or to reduce the load placed upon it.
3. We allow naming records to be moved between directories so that an optimal directory location can be chosen for the record (e.g. following the movement of a cluster around the network).
4. We allow caching for performance. A naming record can be kept at a previous host, as well as being passed up the hierarchy, to reduce look-up time.
5. We arrange that a client can locate the appropriate directory for a cluster rapidly, without having to search other directories.

9 IMPLEMENTATION STATUS

The Mobile Object Workbench version 1.1 is currently in use by the members of the ESPRIT FollowMe project. The current implementation supports clusters, mobility and transparent communications as described in this paper. Current work is enhancing the workbench in three ways:

1. We have started work on the construction of a federated network class loader, to allow support for mobile objects that have different views on the Java class name space. Currently a Place cannot support two Clusters that have different interpretations on the mapping between the class named 'A', and the code

implementing this class. These enhancements are important in an Internet environment where there is no global consensus on class names - and to support evolution of classes.

2. We are currently implementing several of the security abstractions outlined in section 7. We have a Security Manager implementation which can enforce different security policies based on the identity of the cluster invoking a call.
3. The current implementation of the Relocation Service does not support the migration of directory entries. We intend to rectify this in the near future.

In addition to specific Mobile Object Workbench issues, work on FlexiNet is continuing. We have recently created added support for secure bindings using SSL and support for CORBA interoperability using IIOP (OMG 1997). Transactional support is underway.

10 SUMMARY

The Mobile Object Workbench was designed primarily to add *mobility transparency* to the distribution transparencies provided by FlexiNet. In this paper we have shown that in order to do this we must add clustering and re-binding mechanisms. If mobile objects are to be used to support autonomous agent systems, then security requirements lead to the need for *encapsulation* mechanisms so that hosts and agents may communicate and co-operate without the need for complete trust. We have approach the design of the Mobile Object Workbench as a distributed system problem, as it has all the traditional issues related to distributed systems; scale, robustness, independent failure modes, distrust, decentralised administration, multiple name spaces and diverging code bases. This approach has lead us to design an architecture, and implementation, that can evolve to meet future needs, and we believe this gives it clear advantages over the ad-hoc approaches of existing mobile agent systems. In addition we have designed the system as a natural extension of the Java language. This makes it straightforward to use, and allows programmers of mobile objects to use the full language facilities.

REFERENCES

- APM (1996) FlexiNet - Automating application deployment and Evolution
<http://www.ansa.co.uk/Research/FlexiNet.htm>
- FAST (1997) FollowMe project overview
<http://hyperwav.fast.de/generalprojectinformation>
- General Magic (1995) Telescript Language Reference
General Magic Inc.
- Hayton, R.J. and Herbert, A.J. (1998) A flexible component oriented middleware system. *SIGOPS European Workshop '98*

Lange D.B and Change D.T (1996) Aglets Workbench
<http://www.ibm.co.jp/tr/aglets>

Netscape Inc. (1996) The SSL Protocol
<http://home.netscape.com/newsref/std/SSL.html>

ObjectSpace (1997) Voyager Core Technology
<http://www.objectspace.com/Voyager/>

Object Management Group (1997) CORBA/IIOP 2.1 Specification
<http://www.omg.org/corba/corbiiop.htm>

ODP (1995) ODP Referenec Model
http://www.dstc.edu.au/AU/research_news/odp/ref_model/

Sun Microsystems (1996) Java Remote Method Invocation (RMI)" Specification
<http://www.sun.com/products/jdk/1.1/docs/guide/rmi/>

BIOGRAPHY

Richard Hayton graduated from Cambridge University in 1996 and joined the department as a research associate investigating LAN multimedia. He was awarded a PhD in 1996 for work on an Open Architecture for Secure Interworking Services (OASIS). He joined APM Ltd in October 1996 and has since been leading work on a flexible Java middleware platform (FlexiNet). His current research interests include support for complex trust relationships in distributed systems.

Douglas Donaldson graduated in Engineering and Computing from Oxford University in 1991. He worked in the Distributed Software Engineering section of the Department of Computing, Imperial College, London from 1992 to 1996, gaining a PhD in object modelling and design of distributed systems. Since 1996, he as worked for the ANSA research group of APM, investigating flexible, object-oriented middleware to support distribution transparencies.

Michael Bursell graduated from Cambridge University in 1994 and spent two years with Cambridge University Press engaged in electronic publishing projects and Internet development. Since joining APM in 1996 he has worked on Java middleware, database integration and FollowMe. Much of his work has focused on issues of mobile code, from Java applets and RMI to push technologies and mobile agents, and he has a particular interest in issues of security around mobile object systems.

Andrew Herbert, Chief Technical Officer of Digitivity Inc and its parent company, APM Ltd. With Digitivity Andrew is developing products for the secure enterprise deployment of Java. With APM, he leads the ANSA industry sponsored programme of research and advanced development into the use of distributed systems technology. The current focus of the ANSA programme includes mobile agents and component-oriented approaches to middleware deployment.