



UNIVERSITY OF
CALGARY

Security Concerns for Mobile Agents

**Tutorial Report for SENG 609.22
David A. Baker**

Course Instructor: Dr. Behrouz H. Far

Introduction

Mobile agents are agents that can travel to different hosts through a network and perform a variety of tasks based on their goals. The fact that a mobile agent is not dependant on executing on particular system makes it very useful in distributed systems. A few of the advantages of mobile agents are listed below:

- Mobile agents can execute autonomously and asynchronously.
- Mobile agents can adapt dynamically to changes in the overall system.
- Mobile agents are ideal for collecting and processing information over distributed networks.
- Mobile agents are robust and fault tolerant.
- Mobile agents encapsulate protocols.

Security Issues for Mobile Agents

The use of mobile agents for e-commerce, network management, and information gathering is becoming more and more popular in industry today. One of the potential stumbling blocks for mobile agents is the issue of security. An agent running without any security measures could have its data or functionality compromised. The threats to a mobile agent can be grouped into the following four categories [Milagres & Moreia, 2002]:

- Disclosure of private information.
- Denial of service.
- Corruption of code or information.

- Interference or nuisance.

These attacks on a mobile agent can come from the following sources:

- *Malicious Host* – The host server that an agent that is running on has access to the code and data encapsulated within the agent. In the case of a malicious host, the privacy and state of the code and data in the agent could be compromised.
- *Malicious Agent* – The exchange of services and information is vital to the functionality of a multi-agent system. A malicious agent can purposely corrupt services and information that is requested by another agent. A malicious agent could also corrupt the data and services of the host system. Perhaps the most common type of malicious agent is the software virus. A virus has the capability to create duplicates of itself and infect other agents/hosts by modifying their code or data.

Strategies to Mitigate Security Risks

Risks against host system

The security threats to a host system from a malicious agent can be partially mitigated by limiting the access rights for incoming agents, and setting limits on the resources available to each agent. The use of firewalls and anti-virus software can be used to reduce the risk of viral agents infecting a host system.

Risks against agents

Mobile agents can execute a variety of hosts. The diagram below shows the typical lifecycle of a mobile agent with a finite lifespan.




Figure 1: Mobile Agent Execution Path [Fritz Hohl]

The threat to the security of a mobile agent is higher than in traditional software, due to the fact that mobile agents do typically execute a variety of hosts. When running a conventional program, it is very likely that the owner of the software and the owner of the computer system are the same person (i.e. running Microsoft Word on your home PC). However in mobile agent systems, it is very likely that the owner of the agent and the owner of the host system may be different. This creates opportunities for a malicious host to interfere with the normal operation of agents on its system.

Even in cases where the host itself is not malicious, an agent's security could still be compromised. If a malicious agent is running on a host with weak control mechanisms, it is possible that the malicious agent could attempt to modify the code or data of other agents present on the host.

One method of protecting an agent against malicious (or even poorly secured hosts) is to only allow an agent to run on hosts that the agent owner knows are trustworthy. However this approach severely limits the mobility of agents, thereby eliminating much of the advantages of the mobile agent approach. In effect, the agents are running on a closed network.

A less restrictive approach would add the functionality to the software to detect and/or prevent an agent from being tampered with. Verifying the entire code and data base of an agent (as well as corresponding communications with a suspect host) would likely require very large data transfers. A common technique that gets around this problem is to use 'reference states'. Fritz Hohl supplies the following definition of a reference state in his paper 'A Framework to Protect Mobile Agents by Using Reference States':

"A reference state consists of the variable parts (i.e. the state) of a mobile agent executed by a host showing reference behavior"

Basically, using a reference state for the agent allows for checking the difference in the variable parts of an agent that is executed on a suspect host against a reference host, ideally given the complete set of input during the execution. The corruption of an agent by a malicious host would result in a measurable difference in state when compared with the results from the reference host. Although this method will protect the agent against corruption of its code or information, it will not protect the agent from a malicious host reading and disclosing the agent's private information.

There are several existing approaches to mobile security that are based on the idea of reference states. Listed below are the descriptions for three of these approaches:

- *State Appraisal* – A "state appraisal" mobile agent security mechanism was suggested by Farmer, Guttman and Swarup. Following this method, the validity of the state of an agent is verified as the first step when it arrives at a host. The check is performed by the host that receives the agent, as it probably does not want to execute the agent if it has been tampered with (as the agent could be infected by a virus). The set of conditions and rules for how the agent reference

- check is to be performed are set by the programmer of the agent, and presumably (speculation on my part) encrypted within the agent. This approach does have its limitations though, as there are attacks on the agent that would not be detected due to the fact that the host inputs to the agent are not being recorded.
- *Server Replication* – Another mobile security agent mechanism suggested by Minsky et al, was to assume that for every execution session on a host, there are a set of independent hosts that offer the same services with the exact same data. Every execution session on a particular host is mirrored on this set of independent hosts. When the agent execution is complete on all the hosts, they all vote on the results of that execution session. The execution with the most votes wins. Of course this approach fails when the number of malicious hosts in the set outnumbers the number of sincere hosts.
 - *Execution Traces* – Vigna suggested an agent security approach in which the agent owner is allowed to check the execution section of the agent at different hosts when a malicious host is suspected. Every host records a trace of the data passed into the agent. After the agent has finished executing, the host generates a hash of the trace and resulting agent state. The hash is signed by the host, and then continues with the agent on to the next host. After the agent has completed its assigned tasks and returned to its home host, the agent owner can examine all the host traces if there is a suspicion of tampering. If the hash of the resulting state on the home host is equal to the one on the suspect host (using the same input values), then the suspect host did not perform a malicious operation on the agent. This approach only works if each host does not lie about input data that was sent to the agent.

Summary

The value of mobile agent societies for distributed systems and e-commerce applications is becoming more and more apparent. Correspondingly, the need to secure agents and hosts against attack by malicious agents/hosts has become very important. Although there have been several mechanisms suggested for protecting an agent from having its data or code compromised, there does not appear to be a definite solution that will ensure that an agent is never going to be compromised. However new research into agent security continues, and hopefully in the future there will be security mechanisms that are flexible, but still more robust than the current alternatives.

References

- “Securing Mobile Agents for Electronic Commerce: An Experiment” by A. H. W. Chan, K. M. Wong, T. Y. Wong, and M. R. Lyu, The Chinese University of Hong Cong, 2002.
- “A Framework to Protect Mobile Agents by Using Reference States”, by Fritz Hohl, University of Stuttgart.

- “A Protocol to Detect Malicious Hosts Attacks Using Reference States” by Fritz Hohl, University of Stuttgart.
- “MASSA: Mobile Agents Security through Static/Dynamic Analysis” by Alessandro Orso, Giovanni Vigna, and Mary Jean Harrold, Georgia Institute of Technology and University of California.
- “Security Analysis of a Multi-Agent System in EU’s DEEPSIA Project” by Francisco Gomes Milagres and Joao Paulo Pimentao, University de Sao Paulo and UNINOVA.