Justification for "The Insider Problem Revisited"

Matt Bishop
Dept. of Computer Science
University of California at Davis
One Shields Ave.
Davis, CA 95616-8562
email: bishop@cs.ucdavis.edu

The insider problem, or threat, has long been considered the most difficult of all security problems because the insider already has access to sensitive information and resources. But the term "insider problem" covers a multitude of problems, ranging from an insider abusing their authoritty on the system to an insider selling proprietary secrets to a conmpetitor. This lack of clarity in the problem statement adds to the difficulty of solving the problem—or even studying it.

This paper argues for a definition of the insider based upon trust, and upon the acceptance of an insider's ability to violate a security policy precisely *because* an insider is trusted to do so appropriately. It rejects the notion of an insider as an employee or someone who knows something about the system, unless that employment or knowledge includes trust. It ties the notion of an insider to assumptions made by the policy.

This makes the insider problem not a single problem, but a continuum of problems. The paper proposes some general ideas for handling these problems.

Discussion points are:

1. Does the definition of "insider" include those definitions in common usage?
2. Will this definition provide a basis for formalizing the problem? If so, how, and what will we learn from such a formalization?
3. How does this definition improve our ability to capture the notion of insider actions with respect to policy?
4. The paper introduces th notion of a "shadow policy" for insiders; does this improve our ability to reason about insider attacks, and detect or remediate against them?
5. What elements of anomaly detection would be most useful, and how do we select them? Would feature selection help, and if so, how?
6. What other types of protection mechanisms would be useful here?