I-0451: When To Use IFF/IFC And ACF/ACC

TYPE:	Guidance
NUMBER:	I-0451
STATUS:	Ready for External Repost after Interpretations Board Rework/Review
TITLE:	When To Use IFF/IFC And ACF/ACC
FIRST POST:	[cc-cmt 00357]
MOST RECENT REPOST:	[cc-cmt 00715]
COMMENTS DUE BY:	Friday, February 6, 2004 to cc-cmt@nist.gov
SOURCE REFERENCE:	CC v2.1 Part 2 Subclause 6.1 FDP_ACC CC v2.1 Part 2 Subclause 6.2 FDP_ACF CC v2.1 Part 2 Subclause 6.5 FDP_IFC CC v2.1 Part 2 Subclause 6.6 FDP_IFF
RELATED TO:	<None>

ISSUE:

When is it appropriate to use IFF/IFC or ACF/ACC? In particular, for an environment where there is an externally defined policy and no attributes on objects, should IFF/IFC be used?

STATEMENT

ACF/ACC should be used when the intent is to control access to an object; IFF/IFC should be used when the intent is to control the flow of information.

Access to an object means that there is a set of rules that define whether some entity (a "subject") may have a particular form of access to a data container (an "object") for some particular type of operation. There are no controls based on the information itself; that is: if the subject is permitted to write the object, it may write any data into the object; similarly, if a subject is permitted to read an object, it can do whatever it wishes with the data it has read.

Information flow control is based on some fundamental characteristic associated with the information (not the container), and may not involve an active subject. Information flow policies dictate whether information with a particular characteristic can move from one controlled entity to another.

SUPPORT:

When the Common Criteria was first written, the goal on the functional side was to capture the policies that were present in the previous criteria that served as inspiration. In the case of the Trusted Computer System Evaluation Criteria, this was Discretionary Access Control (DAC) and Mandatory Access Control (MAC). However, at this time, there was an effort not to use terms that brought with them specific connotations--hence, many TCSEC terms were not used in the CC.

While the TCSEC separated controls based upon who was allowed to make changes (discretionary verses mandatory), the CC has separated controls based upon the type of control. Since IFF/IFC type controls were often non-discretionary and ACF/ACC were often discretionary the confusion arose that this was always the case.

Access controls (ACF/ACC) may be used to model what in the TCSEC-paradigm was called Discretionary Access Control, but this is not their sole use. They could also be used to implement Role-Based Access Controls, as well as a variety of other controls. The key characteristic is that the controls are based on the object containing the information: they address the fundamental question of whether a particular subject can access a particular object. Note that this is independent of the actual implementation: it could be an access control list on an object, but it could also be a permitted access list on a subject, or some fixed rule set stored completely independently of either subject or object. Access control policies typically have an element of active decision making.

Information flow controls (IFF/IFC) may be used to model what in the TCSEC-paradigm was called Mandatory Access Control (more properly, non-Discetionary Access Control, although even that is a misnomer), but that is not their sole use. The key characteristic is that the controls are based on a characteristic associated with the information, such as a sensitivity label, the quality of the data, or the source of the data. More importantly, the characteristic stays associated with the data as it moves through the TSF, and serves to provide the basis for the controls (although this characteristic may be discarded, or the association destroyed, after the decision has been made). Note that the characteristic need not be embedded in the information; for example, consider unlabelled information entering a multi-level system through a labelled port.

In Information Flow, subjects need not be involved, as might happen in a network device that connects two ports. In fact, there might not be a decision made during run-time at all; the decision may have been made ahead of time by the TOE design and captured in the TOE design (for example, an information diode).

In determining which policy to use in writing a security target or profile, it is extremely important not to let the actual or planned implementation affect the choice of policy. The type of policy should be chosen based on the fundamental type of control: is it subject access to an object, or is it based on characteristics of the information.