ISSN: 2167-1435
v29 n1931

Common Access Control Terminology Used in Multilevel Security Systems

Robert L. Marchant
marchant@psu.edu
Technical Council
Sotera Defense Solutions
Herndon, VA 20171, USA

Abstract

Access to computer data can be controlled by many methods ranging from simply ensuring that the data is contained in a secure environment where only approved personnel have access to more complex access methods associated with public cloud infrastructures. Regardless of where a system resides, controlling access to data must start with fundamental understanding of the terms used in deciding who (or what) has access to the data. In Multilevel Security (MLS) systems where users (or services acting for users) may have disparate privilege to access the data and the data may have disparate sensitivity, access based on attributes (both data attributes and user attributes in combination) may be required. This paper is intended to describe some of the terms that are used when discussing classification systems and these types of systems. Its purpose is to provide common.

Keywords: Classification, compartment, dissemination, discretionary access control, mandatory access control, multilevel security.

1. CLASSIFICATIONS, COMPARTMENTS AND DISSEMINATION CODES

Governments and organizations world-wide have methods of classifying sensitive data based on some decision criteria usually associated with the potential damage the data may cause the nation or organization if that data were compromised. The Unites States has three levels (Director of National Intelligence (DNI) Special Security Center (SSC), Controlled Access Program Coordination Office (CAPCO)(12 May 2008)): TOP SECRET, SECRET, and CONFIDENTIAL. Other nations and organizations have similar systems, some with more or less categories. For example, the North Atlantic Treaty Organization (NATO) has four levels: COSMIC TOP SECRET, NATO SECRET, NATO CONFIDENTIAL, and NATO RESTRICTED. These classification systems are hierarchical; with a clear and undisputed dominance. In NATO, COSMIC TOP SECRET is more restrictive since release of the data is more damaging, than NATO SECRET, and is therefore a higher or more dominant classification. In all of these systems, access to lower levels of classification is inherited. An individual who is trusted enough to be granted authority to access COSMIC TOP SECRET, because COSMIC TOP SECRET dominates NATO SECRET, will automatically be granted access to NATO SECRET and below. Someone who is granted access to NATO SECRET, but not COSMIC TOP SECRET, will have access to NATO SECRET, NATO CONFIDENTIAL, and NATO RESTRICTED, but will not have access to COSMIC TOP SECRET.

If access to classified data were based simply on classification levels as described above, access control would be very simple. For example, anyone who has a COSMIC TOP SECRET clearance would be assigned the role on the computer systems that is allowed to access the areas containing the COSMIC TOP SECRET data as well as NATO SECRET, NATO CONFIDENTIAL, and NATO RESTRICTED. Someone with NATO SECRET can access NATO SECRET, NATO CONFIDENTIAL, and NATO RESTRICTED but not COSMIC TOP SECRET. The access control for this simple "dominance" model can be as simple as placing the data in a directory that has controlled access by either individual userid or

group membership. On systems that can control access through Role Based Access Control (some version of RBAC is available on all modern operating systems), the user would then be granted access to those areas (or in the case of MLS systems, those data) containing COSMIC TOP SECRET through the operating systems implementation of Role Based Access Control (RBAC). Similarly, a user granted clearance for data at NATO SECRET would be assigned the role on the computer systems that is allowed to access the areas containing the NATO SECRET data as well as NATO CONFIDENTIAL, and NATO RESTRICTED.

Unfortunately, classification alone often does not provide sufficient protection for data. Some sensitive data need special handling or special protection that will lead to its being given additional attributes that further restrict access to it. Special subcategories of data (often called compartments) are created to define the type of data and restrict access to the data to individuals who have the need to access the data and been specifically trained on how to handle that type of data. For example, salary data for most companies is sensitive, usually requires special handling, and is usually restricted to access by a select few workers in Finance, Human Relations and Management who have the need to see and analyze the data and have been trained in how to handle (protect) this data. Another example is drug test results. Users of these special data are therefore trained on the need for special protection of the data, how it is supposed to be used, and are provided procedures on how to ensure it does not get compromised. Special controls can also be assigned to data that restrict how the data may be transmitted or released to other individuals or organization (e.g. the press). For example, a large corporation may have data at its branch in Michigan that because of federal law (e.g. International Trade in Arms Regulations restrictions), must get special approval or special censorship before it can be released to its branch in Ireland.

Governments often must restrict how they share intelligent data with other nations. Because of this and similar types of problems, a third level of control is often added to sensitive data to define how that data may be shared among nations or organizations (this is often called release-ability or dissemination control). There are, of course, other types of controls that can be associated with sensitive data (e.g. special (eyes only type) distribution codes, declassification guidance, destruction deadlines, etcetera) but, in general, a classification schema will include a classification hierarchy, some method of defining compartments and some method for controlling distribution (release) are sufficient for most sensitive data description. Using the above as a guide then, you would expect to see sensitive data labeled with a classification level, you may see the data labeled with a compartment code, and you may see one or more dissemination or release-ability codes. For example data that is at the CONFIDENTIAL level, that is intended for users who need the data to track drug test results on cleared personal (let's call it DT for drug test) and is only to be viewed by U.S. citizens would be labeled with classification code CONFIDENTIAL, compartment code DT, and release-ability code US. To view this data, a person would have to be cleared (have access privilege) to at least Confidential, be identified as needing and be trained on how to handle DT data, and must be a US citizen.

2. ACCESS CONTROL

Access control systems can be categorized as either Discretionary Access Control (DAC) or Mandatory Access Control (MAC) (DoD 5200.28-STD). This is not a question involving how to define group access or even Role Based Access Control (RBAC – in RBAC, access is based on the role of the user. For example, users may have escalated access privileges when assigned as the shift supervisor that are not available when the user is not filling this role); this is a question of how the data is contained or labeled and what decides who or what has access to the data. DAC systems allow the user to define the sharing of objects (e.g. data). The user controls where an object is placed (e.g. the user can create a folder and determine what objects are placed within that folder) and the user controls what users (or groups of users) have access to that folder. For example, the user puts the data in a directory for TS data and shares that directly (by name or by group) with whoever the users wants to have access.

A MAC system uses systems security policy, set by and maintained by an administrator, to determine the access rights of a user and to assign the access attributes to the data. In a MAC system, the user cannot control how the data is labeled, where the data is placed, and who has access to the data. For example, in a

MAC system, the source of the data might determine its attributes; data coming from a specific source is always tagged with the attributes that identify the data at predetermined level (e.g. it is always SECRET, always SD, and always only releasable to US, UK, AU, NZ, or CA). The key distinction here is that DAC is a user controlled access decision and MAC is a non-user, policy based decision. Multilevel Security (MLS) systems implement a MAC policy to permit simultaneous access by users with different security attributes to shared resources at the same time preventing users from obtaining access to data (and resources) for which they lack authorization. For our example above, regardless of access control category (DAC or MAC), the access control mechanism must answer TRUE to three types of logical questions:

First: Does the user's clearance level dominate the data's classification level? If the user has a clearance level equal to or higher than the data's classification level, then access can be allowed. Logically this is the simple greater than or equal to relationship (User Clearance) >= (Data Classification).

Second: Is the user authorized access to all of the compartments the data has associated to it? The system must check to ensure that for every compartment code associated with the data there is a matching access privilege associated with the user. The data may require from none to many compartment codes and the user may have access privilege to more compartments than are associated with the data. It is not acceptable for the user to not have privilege for any compartment associated with the data. The algorithm for this check must first compare the compartment codes of the data with the matching privilege of the user (a logical AND) then ensure that all the resulting binary values are TRUE (another logical AND). The defining characteristic in this comparison is that the set of data compartment codes must be a subset of the matching compartment codes associated with the user.

Third: Is the user authorized to receive data defined by the data's dissemination code(s)? The system must check to determine if any of the data's distribution codes match the user's dissemination code. The defining characteristic of this comparison is that at least one member of the set of the data's dissemination codes must match the user dissemination code.

As an example, let us assume we have a user who is cleared for SECRET (S), must work with Special Data (SD) and Sensitive Research Data(SRD), and is a US citizen (Nation code is very often used as a dissemination code). This user would be allowed access to CONFIDENTIAL (C) that is SD and is releasable to US, UK, FR, AU, CA, NZ because:

- NS dominates NC,
- SD is a subset of the set {SD, SRD},
- US is a member of the set {US, UK, FR, AU, CA, NZ}.

This same user would not be allowed access to data that is CONFIDENTIAL (C) that is SD and SRD and is releasable to UK, FR, AU, CA, NZ because:

- Even though NS dominates NC and,
- The set {ASI, SSS} is a subset of the set {SSS, ASI};
- US is not a member of the set {UK, FR, AU, CA, NZ}.

3. TAGS, LABELS AND PROVENANCE

For a system to provide MLS and to perform access control using the attributes discussed above (i.e. Attribute Based Access Control), that system must have some trusted mechanism for safeguarding and asserting the attributes of and for a user (e.g. a trusted directory, a reliable method of authentication, some form of security assertion). The system must also have some method of permanently and reliably associating (e.g. tagging) the attributes of the data to the data, or securely defining the endpoint attributes of the data's communications channel. For example, for digital data, one method that can be used is attaching the appropriate labels through header data or xml tags then digitally signing the data. Essentially, to provide MLS, access control based on trusted attributes is needed. Trusted attributes implies some form of digital signature and some form of trusted directory.

Provenance means history of ownership. Data provenance means the origin and history of computerized data. Secure Provenance means providing integrity and confidentiality guarantees to provenance information (not to be confused with maintaining the integrity and confidentiality of the object computerized data). Provenance is critical in any environment where both the origin and history of any modification or replication (cloning) of data is required (e.g. forensics). Many provenance methods involve

ISSN: 2167-1435
v29 n1931

using some form of identifier associated with the computerized data that uniquely associates the data with its provenance record. Although not essential to MLS systems, maintaining provenance is aided by ABAC.

4. CONCLUSION

Conclusion: MLS, RBAC, and ABAC systems rely on metadata that describe the sensitivity of the data and the privilege of the user. Understanding the terminology and the methods used to evaluate this metadata is foundational to working with MLS, RBAC and ABAC. The algorithms developed to handle data in MLS systems must correctly handle three types of relationships. This paper provided a high level reference on the meanings of the terms: Classification, compartment, dissemination, discretionary access control, mandatory access control, multilevel security, attribute based access control, role based access control, provenance, DAC, MAC, MLS, ABAC, RBAC.

5. REFERENCES

Director of National Intelligence (DNI) Special Security Center (SSC), Controlled Access Program Coordination Office (CAPCO)(12 May 2008), Authorized Classification and Control Markings Register, Volume 1, Edition 2 Retrieved June 1, 2012 from http://www.dni.gov/electronic_reading_room.htm).

DoD 5200.28-STD, "Department of Defense Trusted Computer System Evaluation Criteria. Retrieved June 1, 2012 from http://csrc.nist.gov/publications/history/dod85.pdf).