Risk-Adaptable Access Control (RAdAC)

This paper considers the impediments of traditional access control approaches to sharing of information. It describes a concept for an access control model that emulates real world decision-making, considering operational need and security risk as part of each access control decision, and recognizing that situational conditions will drive the relative weight of these two factors in determining access. Access control decisions can adapt to varying situational conditions in accordance with an access control policy. Thus the model can support extremely restrictive policies and those that provide for the widest sharing, with added risk, under specific conditions.

Robert W. McGraw, Author

Contents

Access Control and the Information Sharing Problem

Mechanisms for controlling access to information that are available in today's systems do not have the flexibility and basis for decision-making needed to support the goals of information sharing. In the real world, decisions are regularly made by commanders in the field to give access to, and share classified information under less than ideal security conditions. These decisions are driven by situational factors and operational need[1]. They are made with the belief that the operational benefits of sharing the information outweigh the potential security risk[2] of sharing it. The focus of such decisions is achieving operational success at the expense of added security risk, given any number of situational factors. The basis for making these decisions is an understanding of the operational need, the resultant security risk, the policies and operating procedures governing the situation, and the knowledge of the effects of similar decisions from the past. Note that while it is essential for commanders to have the latitude to make these decisions in order to execute their missions, it is unlikely they have a comprehensive understanding of the security risk associated with their decisions. The local ramifications may be understood but not the enterprise-wide effects.

Contrast this highly dynamic, operationally driven, security risk-based decision process with access control processes prescribed for information systems in Figure 1. Most of these approaches derive from the computer security activities of the 1980s, and strictly fulfill the laws, policies, and implementing directives for protecting classified information. They require the intended recipient of classified information, the subject, to hold a security clearance commensurate with the classification of the information object (Mandatory Access Control), and have a "need-to-know" for the information (Discretionary Access Control). The problem with this approach from a sharing perspective is that it assumes that it is too risky to share information if these criteria are not met. It does not recognize that in some situations, the consequences to national security of not sharing information might be graver than those of sharing it, even under the worst of security conditions. The risk-averse security policy of the enterprise is hard-coded into the access control logic, with no room to support the dynamic and situational conditions in the real world. Such policies and supporting access control logic assume a homogenous environment, where all people that could potentially require access to information have a clearance and are located in secure environments, and all the computers and associated networks that would process classified information have the pedigree needed to do so. This situation is not reality, particularly in an enterprise as diverse, complex, and situational as is the DoD.

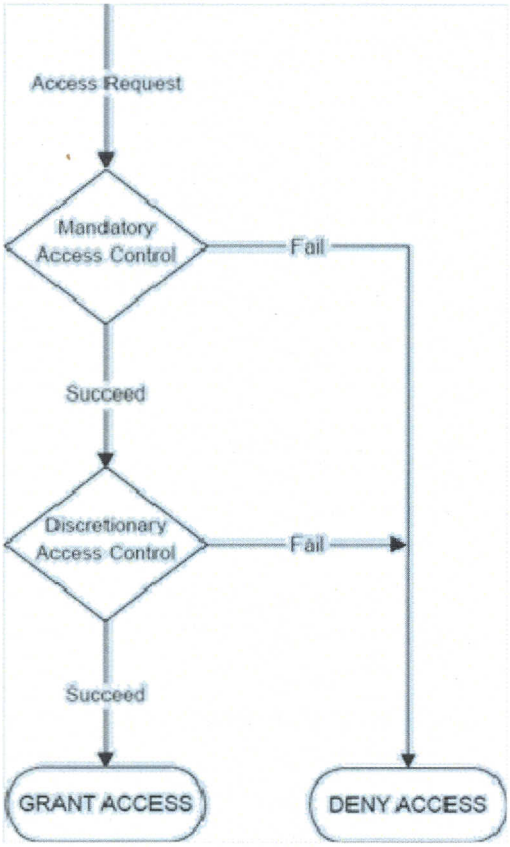

Figure 1: Traditional Access Control

Changing the Access Control Formula

A critical part of implementing effective information sharing then is to implement an object-level access control process that can deal with the realities of the information sharing environment. The proposed access control concept to achieve this environment has been named Risk-Adaptable Access Control (RAdAC, pronounced Raid-ack). What distinguishes RAdAC from traditional models is flexibility and adaptability - flexibility to adapt access control decisions to the situation at hand, much as the commander's decision process described earlier.

The basis and formula for access control decisions must change in order to achieve this flexibility and adaptability. Thus, RAdAC decisions are made based on a number of factors as depicted in Figure 2 and discussed here:

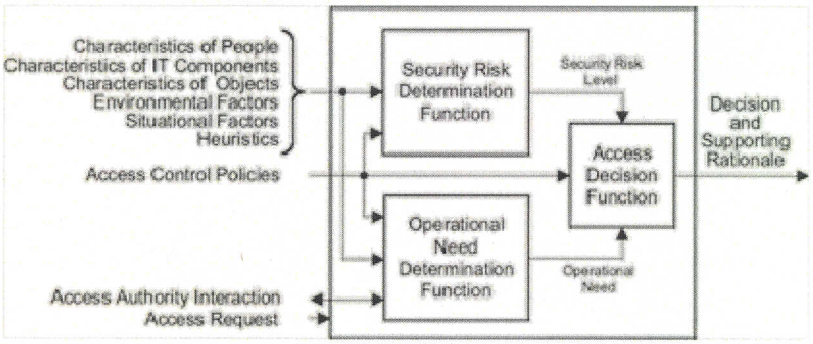

Figure 2: RAdAC Functional Depiction

Operational Need

The proposed model allows operational need to enable access, and under specified conditions, to outweigh security risk in determining access. Operational need was considered in traditional models under the guise of 'need to know', but was used to restrict access rather than to enable it. It can manifest itself in many ways such as a person's membership in some community of interest or organization, or their location. A supervisor or other approving authority might have to attest to a person's need to have specific information. This interaction is shown in the functional model as "Access Authority Interaction." Given its emphasis here, Operational Need must be characterized and parameterized in such a way that the RAdAC process can use it. It must convey some quantifiable measure, not a binary indication.

Security Risk

RAdAC incorporates a real time, probabilistic determination of security risk into the access control decision rather than just using a hard comparison of the attributes of the subject and object as in traditional models. Further, the determination of security risk takes into account a number of factors as shown in Figure 3. It shows that people use information technology (IT) components, which connect and use other IT components, and control the information objects that are to be shared. All of these items exist in environments that include a physical location and an adversarial threat element. The trustworthiness of people, the protection capabilities and robustness of IT components, and the threat level of their environments, in conjunction with the value and access history of the information object being accessed, all contribute to the security risk. The process will determine the risk associated with each of these, as well as a composite risk. Using this risk-based approach enables RAdAC to be responsive to the broad range of operational situations.

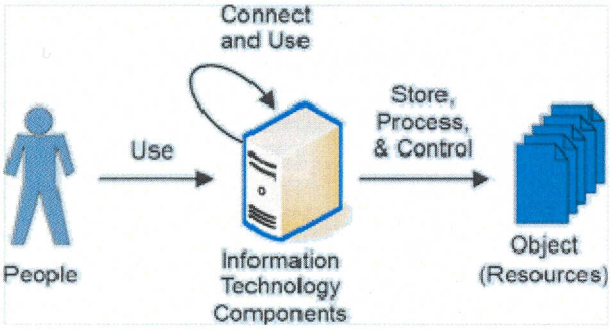

Figure 3: Security Risk Considerations

Situational Factors

The conditions under which the access decision is being made are factored into the process. National, enterprise or local situations may determine these conditions. The national terrorist threat level, or an indication that the enterprise is under cyber attack could tighten or loosen access rules. A situation where troops are under severe enemy fire may drive the weighting between operational need and security risk in making an access control decision. It may be critical that those troops be given some piece of information regardless of the risk that the information might be compromised. Thus, such conditions may dictate whether operational need can outweigh security risk, regardless of the severity of the security risk.

Access Control Policy

This specifies the rules for access control for various classes of information objects under different conditions. It allows the enterprise to describe the degree of operational need required to "override" acceptable or normal security risk, and to set acceptable levels of risk. It must be capable of specifying the policy for each step of the RAdAC process. The policy might specify the relative weighting of personnel risk, IT component risk, and environmental risk in computing a composite risk. Effectively implementing and managing digital access control policies is a critical element to making the RAdAC model successful.

Heuristics

Knowledge of past access control decisions will be used in making each subsequent decision. Such knowledge can be used to develop better algorithms for determining risk and operational need, and help to fine-tune the access control policy to improve the rate of positive access control decisions. Knowledge of compromises that have resulted under various access conditions, for example, might help the system more accurately determine risk and make better decisions. Policy must specify the degree to which heuristics should be considered in each access decision, as well as how each decision should be incorporated into the learning process.

RAdAC Notional Process Model

A notional representation of the RAdAC process flow is shown in Figure 4. Note that there are other process flows that could accomplish the same or similar results. A request to access an object initiates the RAdAC process, which includes these steps:

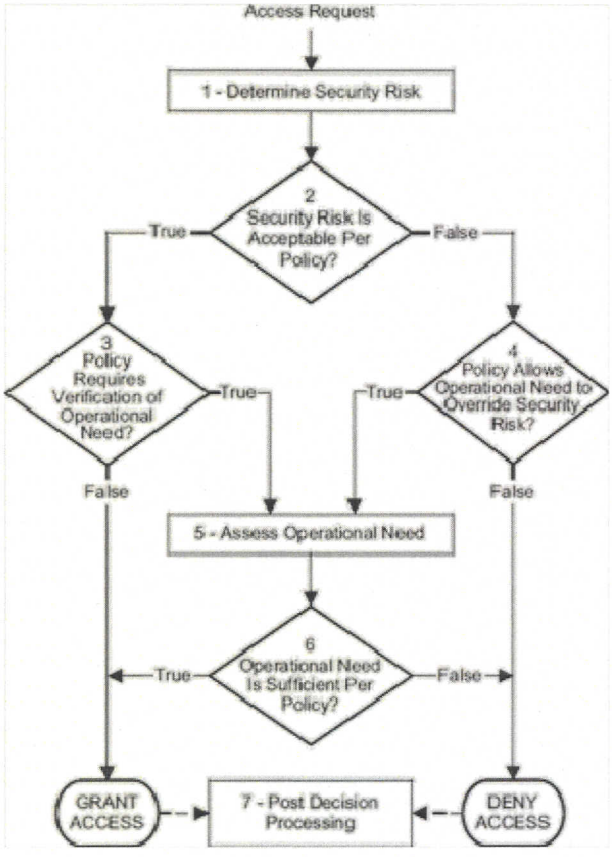

Figure 4: RAdAC Notional Process Model

Step 1 – Determine Security Risk

A real-time, probabilistic determination of the security risk associated with granting the requested access is made based on examining several external factors. The level of risk will be determined in several different areas such as the risk associated with the people, IT components, and environment involved in the access. The result of this process is some quantitative indication of the level of risk for each area, as well as a composite risk.

Step 2 – Comparison of Security Risk Against Policy

In this step, the result of the measured security risk is compared with the access control policy that identifies the acceptable level of risk for the object being accessed. The level of risk will be compared

for several different areas such as people, IT components, and a composite risk. The policy will have to specify an acceptable risk level for each area, or a risk range (e.g. low, normal, and high).

Step 3 – Policy for Verifying Operational Need

At this point the security risk of granting access has been determined to be acceptable, but the requestor might not have an operational need to access the information. The policy will specify whether verification of operational need is required for access, and if required, the criteria for determining it. If the policy requires operational need to be verified further processing is required, otherwise, access is granted.

Step 4 – Policy for Operational Need Overriding Security Risk

This step occurs if the security risk was determined to be unacceptable in one or more areas, but the requestor might have an operational need to access the information regardless of that risk. The model and the policy must be capable of specifying whether operational need may outweigh security risk, specifically for which areas of risk operational need may take precedence (e.g. person's trustworthiness, their location, weak IT components, etc.), and under what conditions (e.g. situational factors). If override is allowed then further processing is required to determine whether the requestor's operational need is critical enough to outweigh the security risk, the criteria for which must be specified in the policy. If override is not allowed in any area where the risk was not acceptable, then access is denied.

Step 5 – Assess Operational Need

During this step, several factors are examined to determine if the requestor has the operational need required to access the object. The policy would specify different requirements for determining operational need, depending on whether the security risk was acceptable, or depending on why it was unacceptable. The requestor's membership in some community of interest or organization, their location, their rank or some other discretionary factor might be used to determine operational need. Another person or an automated service might have to attest to the requestor's operational need to access the information and thus, an external workflow process might be engaged to seek such approval.

Step 6 – Comparison of Operational Need Against Policy

The final step is to determine if all of the requirements for operational need, as specified in the policy, were met. If all requirements were met then access is granted. Otherwise it is denied. The policy must be capable of identifying the criteria for determining sufficient operational need under both stressed and unstressed security conditions.

Step 7 - Post Decision Processing

The actual decision, the rationale for the decision, and any other pertinent information are analyzed and stored by post RAdAC decision processing. The analysis of the results would be done automatically, in real time, and made available to aid and improve the RAdAC decision engine. Results of access control decisions would be made available to information owners/authorities to help them assess and adjust access control policies. The degree of information sharing occurring in the enterprise could also be measured and compared with performance metrics.

Challenges Abound In The Road Ahead

RAdAC reflects a significant departure from existing access control models and supporting technologies. There are numerous challenges that must be solved before it can be fully implemented in the 2016 and beyond timeframe. These challenges are in addition to the increasing demands for higher assurance in access control implementations. Despite the many challenges there is reason for cautious optimism that the technology needed can be developed in as much as the rudiments already exist, and there is some overlap between the needs and the business opportunities, which have already been identified. Undoubtedly the final solution will contain elements of many of the present capabilities and a variety of other new technologies beyond comprehension today. Here are few of the challenge areas that will need to be addressed:

Infrastructure Support

A supporting infrastructure that will supply various information and services must be in place in order for RAdAC to be successful. Actually, this problem is true for any object level access control mechanism, but RAdAC offers some added challenges. Figure 5 shows RAdAC in context of the supporting infrastructures that must be available to support it. Some of the infrastructure challenges are:

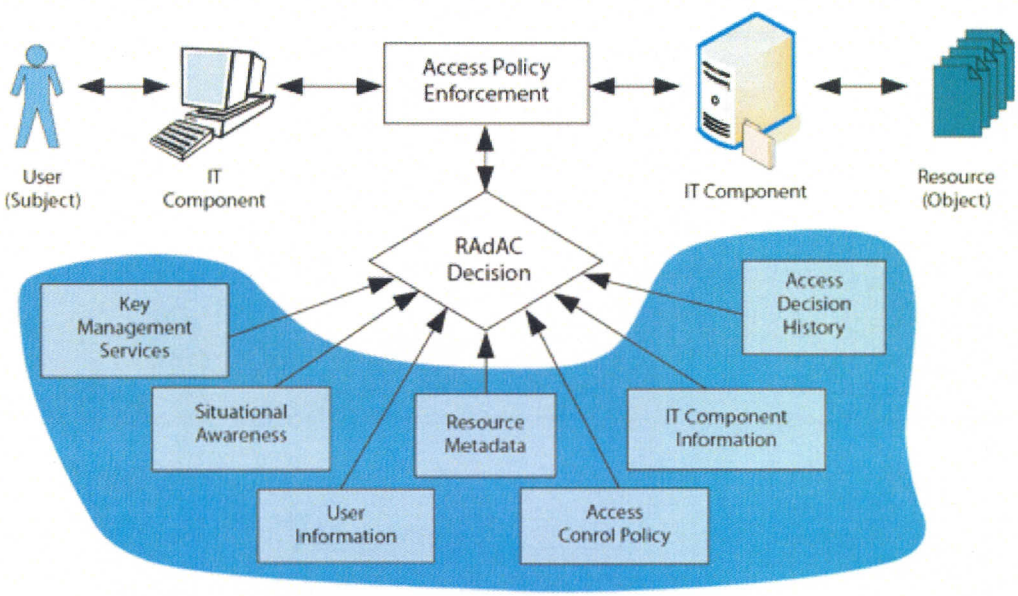

Figure 5: Access Control Context

User Information – This is the source of any information RAdAC would need to assess the trustworthiness of the people involved in the access decision, such as identification and authentication information, and authorizations such as their security clearance. Since RAdAC will have to render access decisions for people that do not hold security clearances, other information will need to be available to use in the risk determination process to determine a level of risk associated with granting them access. What sort of information might be valuable to determining their trustworthiness? Could a mini background investigation be done online?

IT Component Information – Sufficient knowledge of the information assurance capabilities and security robustness of a computing platform, as well as the risk associated with the environment in which it resides, will be required to determine the security risk of allowing access from that computing platform. While this may seem daunting, initiatives by the Trusted Computing Group (TCG)[3], an industry and government consortium, will help to meet this challenge. The TCG is developing a scheme as part of its Trusted Network Connect specification that enables a platform to "prove" its security goodness to another in a secure, verifiable manner. This approach could be used to determine the configuration of the platform, for which a security risk could then be assigned. It may also be possible to use information available from security risk assessments done as part of system certification and accreditation processes. Under this approach, risk calculations would be available online to support RAdAC risk determination. When information is to be sent to a user of a particular system, the certification and accreditation risk assessment values for that system could be used in the RAdAC risk calculation.

Access Control Policy – A robust infrastructure must exist to provide the access control policies needed to support RAdAC based decisions. At a high level, this infrastructure element must provide a repository from which machine-readable access control policies can be served. Far more challenging though is that it must provide a policy conversion function. This function must be able to capture the policymakers' intent for how information should be shared under various situations, and translate it into machine-readable policy statements. The language used for machine-readable policies must be capable of expressing the broad range of policy considerations associated with RAdAC. It must be extensible, provide rules for allowable and disallowable policy constructs, and for policy negotiation and deconfliction. The later item will be particularly important since policies could originate from the various hierarchies of the government, as well as from other governments whose information might be controlled by RAdAC. Inconsistencies are likely to abound and deconfliction will be essential. While much of this capability is beyond the limits of technology today, there is hope that positive improvements can be made in this area given the fact that the rudiments of such a capability already exist.

Security Risk and Making the Decision

At the core of RAdAC is the notion of determining risk associated with people, IT components, and their environment. Capabilities will need to be developed to produce meaningful and consistent risk calculations. Bayesian probability methods may hold promise, as well as use of fuzzy set theory for making decisions under the less than precise conditions that will often exist. Additionally, algorithms or heuristics developed to support decision making must be able to be dynamically tuned to support adaptive adjustments in information sharing policies. Because of the security criticality of all of this machinery, proof of correctness will be a critical factor in the implementation of these processes. Further, despite the complexity of its decision process, low latency performance will be critical to making RAdAC feasible. The access control decision must, with few exceptions, be instantaneous.

Cultural and Legal

"Giving uncleared people access to classified information??!!! That's illegal! How could you ever consider doing that? My information is far too sensitive to ever allow that to happen!" Culture will be a difficult barrier to overcome. RAdAC doesn't itself give unfettered access. Rather, its objective is to employ technology to allow policy makers to drive information sharing to the extent required for the situation at hand. By changing the basis for the access control decision as described earlier, RAdAC can support extremely restrictive policies and those that provide for the widest sharing, with added risk, under specific conditions. Gaining trust that this model can satisfy this claim will take time, and an effective transition strategy will be required that allows cultural and legal obstacles to be addressed along the way.

End Notes

1. [^] Operational need as used here means the degree to which success of an operation or mission is dependent on sharing the information.
2. [^] Notionally, security risk is an indication of the probability that a particular sharing decision will result in unintended consequences. Examples of unintended consequences are disclosure of information other than intended by the sharing decision, loss of integrity of the shared information, or loss of integrity of a system.
3. [^] For further information on the TCG see their website at https://www.trustedcomputinggroup.org/home