2006 CCRTS
THE STATE OF THE ART AND THE STATE OF THE PRACTICE

Command Authority & Information Flows in Net-Centric Operations

C2 Concepts and Organizations
Policy
Social Domain Issues

Linsey O'Brien
Scott Renner
Arnie Rosenthal
Jay Scarano

POC: Linsey O'Brien
The MITRE Corporation
202 Burlington Road, MS M380
Bedford, MA 01730
781-271-6340 / 781-271-2101 (fax)
lobrien@mitre.org

Approved for Public Release; Distribution Unlimited. 06-0306

Command Authority & Information Flows in Net-Centric Operations

Linsey O'Brien
Scott Renner
Arnie Rosenthal
Jay Scarano
{lobrien,sar,arnie,jgs}@mitre.org

Abstract

The DoD net-centric transformation will bring extended reach & flexible capabilities through seamless information sharing. This requires breaking down the stovepiped systems that limit commanders from taking advantage of external information and assets. However, stovepiped systems are not *all* bad: one side-effect of those arbitrary and rigid barriers is to ensure access only to vetted information by authorized participants. As we take down these old barriers, many new information flows and decision procedures become possible. Some of those possibilities are wrong, and should not be permitted. The question for this paper is: Who decides, and what is required to enforce those decisions? We need new procedures and technical features to ensure that the right information gets to the right people, that all information is protected, and that overall information flow policy is preserved for the enterprise. This paper discusses the impact of net-centric operations on technical architectures and some of the options and capabilities that new technologies can provide.

Introduction

Network-centric warfare (NCW) is a theory of military operations that is at the heart of the force transformation process underway in the DoD. NCW begins with information technology supporting seamless connectivity and information sharing between the elements of the friendly force. This seamless information exchange leads to the shared situational awareness that in turn produces increased speed of command, synchronized effects in the battlespace, and dramatically increased mission effectiveness [1,2].

Information technology resources are initially scarce and expensive, but over time they become commodities available to our adversaries. In the long term, the advantage of NCW will go to those best able to rapidly change their doctrine, organization, and training [3,4]. Rigid information systems that take months and years to modify will act as a brake on that change. Flexible, easily-changed information exchanges are an essential part of the NCW advantage.

Seamless and flexible access to information was motivation for the DoD Net-Centric Data Strategy (NCDS), that describes a vision for a net-centric environment and the data goals for achieving that vision [5]. The NCDS goals for data – make it visible, accessible, understandable, interoperable, and trusted – will lead to the information exchange that is essential to net-centricity. When producers post data to shared spaces, they make it available to unanticipated users and applications, leading to improved

flexibility and increased warfighter agility. This is part of a deliberate shift from a "need to know" culture of information protection to a "need to share" culture of integration.

The NCDS goal of data accessibility does not mean an end to access control restrictions. Access to data will continue to be limited by applicable law, policy, and security classification [6]. The NCDS goal does mean that access restrictions must be based on deliberate policy decisions (which may be quickly changed), and not on accidents of incompatible implementation (which change very slowly). The strategy is to open up the system "stovepipes" through information services, and to instill a policy mindset of sharing information as widely as possible. The objective is to enable many new information exchanges that are not possible today, without creating an insecure information "free-for-all."

In this paper we will show why this objective cannot be supported with the existing technology and procedures used to enforce access control. We discuss requirements for Net-centric policy management and enforcement, and the procedures and technical features needed to support those requirements.

Limitations of Physically-Based Current Approaches

In current practice, information and its users are grouped in long-term organizations; each organization typically embodies its information flow policy in physical implementation. In other words, there is separate equipment – networks, servers, etc. – for each functional as well as each geographical group of people. When a sharing policy changes, it takes major effort to merge the separate physical infrastructures to implement the combined policy. Furthermore, the merger essentially opens everything to everyone within the combined organization.

Maintaining control of information flows in an increasingly blended environment requires increasingly overlapping and dynamic organizations, that need to define increasingly fine-grained and dynamic policies. This becomes prohibitively expensive to implement if policy boundaries require system boundaries – costs scale quadratically as each organization stands up a separate policy implementation to interact with every other organization. In fact, the cost of this approach creates an inherent limitation in how adaptable an organization can be. In stable, predictable circumstances, or for organizations that have extremely deep pockets, this limitation may not be particularly oppressive. However, if the organization is in rapid flux or constrained by resources, limitations on agility may be intolerable.

Need for Information-Based Policy Enforcement

In the past information security was indivisible from securing the physical medium or container that held it. Sharing information involved a relatively expensive and unreliable copying process; if you controlled the container-copying process, you controlled the information sharing. With the rise of digital technology, information propagation is so much better, faster *and* cheaper that organizations can more easily afford to share

information on demand during those periods of rapid change when sharing is most valuable, enabling them to adapt without the massive capital investments previously required. However, enterprises that depend only on old policy methods based on container management will have many of the costs of the new technologies, without the cross-organizational benefits. Policy and procedures have to shift their emphasis from control of physical perimeters to enforcement of perimeters specified in terms of information constructs. Information constructs make strong multi-level damage control feasible because they enable policy enforcement wherever needed, not just at physical perimeters. Corruption or deletion of information is limited to only that which is accessible. Even if the aggregate exposure of all instances is large, it is still limited to the sum of many small chunks of information rather than many large ones. Furthermore, since policies can be enforced automatically, information sharing can be transparent to the authorized users, further easing coordinated actions among them.

The essence of information-based policy is to specify which aspects of a collection of information cause it to be considered sensitive, and which aspects of a consumer gives them the access rights to match. In classic chains of command we can continue to consider all information within a physical operational location to be implicitly labeled as appropriate for consumption by the people in that space. Once that organization is geographically distributed or shares *any* physical facilities with other organizations, or the consumers potentially have time- or mission-varying levels of authority, the information infrastructure must enable much greater control and flexibility in specifying access policy. The granularity may need to be different or the information may not be equivalently structured – one organization may combine certain sensitive information into a single virtual document or database view, while another may distribute it throughout several information service provider streams. As long as each piece has equivalent attributes (or a mapping between them), only a single policy need be written.

Attribute-based access control (ABAC) can enable organizations to express minimalist policies; information classification resolution can match the natural granularity of the information itself rather than that of its medium. ABAC generalizes role based access control (RBAC) in that it does not force all categorization into a single, one dimensional role structure and instead defines roles as collections of attributes. Multiple humans can support a single role over time, and multiple objects can be treated in the same way, without requiring policy updates for every new member or retraining the human clients and/or gatekeepers for each policy update – a cumulatively expensive process.

Challenges with ABAC

Controlling access using information-based mechanisms requires capabilities and CONOPS that extend beyond the classic physically-based ones. The difficulty becomes specifying the policy you want, in a vocabulary and structure that is meaningful to the administrator, and that can still be enforced on objects defined in other terms.

Networking organizations together often generates extremely large-scale and complex operations that make manual access control and policy management infeasible.

Technologies and tools that can help translate commander's intent into formal machine-compatible policies without getting the humans mired in the details or generating a mass of conflicting spaghetti-like rules will be a necessary capability. Improved resolution should not overwhelm the policy makers and managers with a mountain of constantly changing mini-policies. Constantly re-tagging each jot or twiddle with a list of access classification labels ten times its size is unworkable at enterprise-scales. It may cause more vulnerabilities due to user confusion than it resolves and could even make the overall system more brittle.

To mitigate the issues caused by the shift to attribute based access control, we should employ new concepts of operation based on the new ABAC capabilities themselves, and employ attribute-based information structural analysis to generate access attributes. Because neither manufacturers nor integration contractors can predict operational sensitivity aspects, authorization aspects or changes in policy beforehand, the best that they can provide are tools allowing the operational team members to define them and then make adjustments in terms that make operational sense to the humans and implementation 'sense' to the system equipment.

Fortunately, there is a straightforward initial approach to using the tools included with most information management technologies such that the resulting formal specifications can be presented in non-technical operational terms. Standard usages and labeling conventions already exist, and can be used to define attributes. For example, commander delegation in a joint operation, that *must* be based on doctrine, can be used as a source of labels.

When a superior communicates intent it is often sub-divided into different functional areas, with a subordinate specialist responsible for each area such as intelligence, operations or logistics. In information terms, functional types of information assets, such as target type or surveillance video stream type, imply various offensive and defensive aspects of intent. Each of the commander's asset types is delegated to a particular type of subordinate; each instance to a particular instance (e.g., splitting geographically). Also in information terms, each asset should have a standard authority attribute that identifies the designated responsible agent; the system should also track the chain of delegation (usually the chain of command). To reduce labor, the default type and responsible agent are inherited from the parent asset. The collection of assets is a type of asset in its own right and has as its authoritative steward the commanding officer. The component assets have either subordinate stewards and type, or default to the superior.

Another example exploits the implicit authority contained in formal organizational entity labels. It is quite common to infer that an agent has authority based on its network name or protocol assignment. The Public Key Infrastructure (PKI) X.509 security certificates assume that their authority chains are defined by a formally structured list of organization names and that trust is a function of recognizing agents by their organization names and

matching organizational identity management policy identifiers[1] to trusted authentication procedures.

Managing Attribute Definitions

Managing large sets of automatically-generated attribute definitions requires some overall structure for their creation. (The alternative is a massive "laundry list" that would grow every time there was a new commander, a new type of information, or a new asset.) Without solid attribute definition management procedures, interoperability plummets, testing and validation become impossible and administrators are overwhelmed. Such procedures need to identify and eliminate overlaps, justify new additions beyond local usage, and otherwise ensure that the system is sustainable over the long term. They cannot be the responsibility of a single enterprise administration organization. At these scales, there must be one mechanism that serves as a federation council and another mechanism that methodically and formally sub-divides and delegates responsibility further down each chain. This combination of unity of effort at the top, and divide & conquer underneath, is the only way to ensure that the system is responsive to the needs of the people and organizations it serves while maintaining sufficient coherence throughout.

Attribute administration must cover more than delegation of administrative authority. It also must specify the formal methods that enable attributes and policies generated in one part of the administrative tree to interoperate with those from other branches. This requires a foundational understanding of what attribute and policy specifications *must* have in common, what they *might* have in common, and where they *must* and *can* agree to differ. It also requires a methodology that enables developers and administrators to tell which is which.

Information Perimeters

In order for a policy to control access to an asset, there must be a selective perimeter barrier 'surrounding' and 'containing' the asset. These perimeters are infrastructure constructs that implement a set of policy enforcement points. For information assurance (IA), access policy is usually based on combining the proposed asset utilization and sensitivity together with the information consumer's inherent and delegated capabilities. The essential nature of any selective mechanism is based on defining some distinguishable characteristics. In the case of information access control, it is helpful to support both attributes that represent asset utilization and sensitivity, and attributes that represent consumer authorization. The two types of attributes are called out here because the descriptive processes that tag assets have significant operational differences from those that describe and tag consumers. They are often implemented in different sub-systems for legal, financial and performance reasons.

[1] The Computer Security Objects Register is a National Institute of Standards and Technology organization one of whose purposes is to deconflict and standardize security labeling conventions within the federal government. http://csrc.nist.gov/csor/

Asset Access Description Points: Documents, Objects & Streams

Asset description for assurance purposes is essentially a risk-sensitivity evaluation process that adds the potential cost of service denial (due to overuse of newly published information, also known as the Slashdot or Digg effect) to the cost of damage if information is inappropriately revealed by publication. For better performance, a label or meta-data tag that signifies the result may be attached literally or figuratively to the asset where it can be assessed without access to the asset proper. As description procedures migrate from hardcopy information containers with physical labels to electronic documents to information objects, the perimeter type shifts from a physical one to a spectrum of electronic types that define much more flexible containers such as information objects or flows within streams. Description points may be implemented within asset creation, or separately, as part of a conventional classification process or as a part of policy creation. All three description points may coexist in order to accommodate the entrance of separately generated legacy and 'foreign' information as well as 'native' assets.

Authorization Points: Assigned versus Implicit Roles

Just as an asset description procedure is necessary for asset control, an authorization procedure requires identification and similar descriptions for information consumers. Like the asset description procedure, automation of the consumer description procedure requires formal specification of the authorization inputs, the sources and combinatorial methods. There is considerable previous work in this area, as legal sources of authority and methods of delegation have been the subject of millennia of analysis and are embodied in our legal institutions. Assigned roles are simply collections of explicitly delegated institutional authority attributes.

However, some authority derives implicitly from the organic capability and assets found within each agent's perimeter or can be inferred on the spot from other attributes. Consider nightclub admission – slipping the doorman cash, dressing well, and driving an expensive car all have been accepted as adequate authority for access. However, like delegated authority, derived authority requires an explicit definition that can be validated: implicit roles are collections of derived attributes. Without implicit, as well as assigned role attributes, formulating a complete policy is impossible because authority is not completely defined. 'Gaming the system' and 'social engineering' becomes much more likely without a comprehensive definition of authority. [7]

Policy Enforcement Points: Commander's Intent, Guards and Encryption Sandwiches

Given classification attributes and authority attributes, a *mapping and tradeoff specification* is the last part necessary for information-based policy. The *mapping* part of the specification define which classification attribute[s] are to be compared to which authority attribute[s]. It also defines any normalization or translation necessary, for example between combatant commander specification and area of responsibility GPS coordinates. Commander's intent and guidance define the *tradeoff* semantics part of the

specification – the acceptable ranges for each attribute and how much discretion there is in the match between classification and authority. By using risk-structured classification and capability-structured authority attributes, together with structured policies that capture intent and tradeoff guidance, instead of blindly matching opaque classification and authority attributes, it is much easier to develop and automate policy management tools to help with policy deconfliction.

Formalized policy permits automated enforcement. Inspection by either human guard cells or automated guard servers is a form of policy enforcement point (PEP) that presumes information containers are composed of smaller containers down to some minimum granularity. Each component has its own perimeter to which is assigned the access control meta-data:

- the storage array, the file system, the document, the sentence
- the InfoBase, the server, the view, the field
- the ISR stream, the sensor array, the sensor

Information-based access control perimeters that exist at multiple scales allow for redaction – partial access to a selection of smaller scale components based on a policy evaluation at each perimeter's enforcement point – as well as access to the assembly as a whole. Such information sieves can filter on size of risk.

Finally, by providing a way to characterize information relative to performance only – size, delay sensitivity, etc. – infrastructure administrator agents can be authorized to monitor and manage information flows without having to grant them unnecessary access to all the information.

Trusted Infrastructure in a Net-centric Enterprise

All of the above discussion about information perimeters presumes a trusted infrastructure that implements them, and that the classification, authorization and policy enforcement procedures can operate. This last section reviews the impact of ongoing technology advances driving the development of a trusted infrastructure to serve as the foundation for information-based access control.

Modular architectures, virtualization & scaling

The biggest challenge to a trusted net-centric infrastructure is the very thing that makes it both robust and flexible – a modular architecture. Each module is essentially another machine agent, that requires authorization as it is given access to the information asset in order to process, store or transfer it. In order to avoid drowning in policy administration, two simplifying approaches are available: trusted platforms, and automated policy-based configuration management

When systems were implemented as a whole, instead of being composed from separable components, the system designer, creator and factory could all be vetted and the resulting system could not be changed without breaking it. This led to brittleness, manufacturing expense, inability to scale, and shortened lifetimes with consequent high lifetime

capability cost. Some markets have spurned this approach and accepted greater vulnerability because the other costs are too high to remain competitive.

An alternative is to leverage the work done on automated assembly & configuration and add access policy to on-the-fly assembly or plug-and-play. This has had two problems. First, there has been a major struggle over both who defines authority (and how) and who defines configuration policy (and how). Second, adding access control attributes and policy to hardware interfaces will likely require major changes in existing standards, which takes time.

Commercial attempts to integrate access control in automated configuration management (i.e., via trusted computing platforms) have not been very successful because vendors tried to use them for access control monopolies – you don't get access unless the platform is trusted but there are limited sources for trusted platforms. A standard based on a more general definition of access control requires a mechanism in which a policy can permit access based on multiple authorities including implicit authority, not just that delegated from a single vendor. This is a necessary pre-requisite for a truly open interoperability standard. Interoperability is the essence of a net-centric enterprise running on composite systems because it ensures the ability to adjust capacity and capability on the demand of the organization, not at the pleasure of its vendors.

Research attempts to integrate access control in automated configuration management tried to solve the most general authority problem by assuming that only implicit authority exists. In this environment, every agent only speaks for their organic resources and consequently must build a reputation of trustworthiness on the fly, based on their behavior. Much of the grid computing work operates on this basis, but it takes time and resource expenditures.

Note that while such efforts are necessary to establish (or re-establish) any authority delegation system, they are unnecessary where there is trust based on unity of command, and they are inefficient where there is trust based on unity of effort.

Service oriented architectures and shared infrastructure

Another aspect of a net-centric enterprise that has an impact on access control is the rise of service oriented architectures that expect to share an underlying infrastructure. One premise of net-centricity is that a significant part of an enterprise's information generation and processing will be non-local, i.e., not organically integrated. Any information produced and consumed must flow between enterprise participants; in order to avoid expensive duplication of the underlying infrastructure, resources must be shared, which means information with different classifications and agents with different role authorities must coexist. This adds another set of trust requirements to the infrastructure list. Not only must they have a configuration process that implements general attribute based access control – multiple policy enforcement points with different policies must also independently and peacefully coexist. Initial efforts to define 'independently and peacefully coexist' are captured in [8].

Conclusion

Information assurance requires both an enterprise view covering all the components, and individual component-oriented views. Key IA components needed for dynamic network-connected enterprises are: attribute-tagging for information assets, attribute-based access policy specification & management for consumers, and a trusted infrastructure implementing both. For large enterprises, the large number of both assets and their attributes will require a combination of information assurance conventions and automation to minimize confusion among users and administrators and maximize system performance. Some common legacy conventions such as classification categories for assets (and the labels that symbolize them), and authorized roles for consumers will simplify migration but ultimately impose limitations. Policy conventions and default policies such as delegation of authority or need to know based on an organizational identity or location attribute are more robust because they are embedded in our legal institutions such as the formal military chain of command. Understanding of access fundamentals and recognition of the limitations of current CONOPS will reduce the pain of future development and deployment.

Furthermore, we can't expect the developers or even integrators to do more than provide tools based on the fundamentals because they can't anticipate all relevant operational asset and consumer characteristics let alone operational policy requirements. We should attempt to influence commercial development such that they don't impose unacceptable limitations in our trusted infrastructure implementations. We should also attempt to influence open source research and development such that they don't impose unacceptable inefficiencies.

Acknowledgements

The authors would like to thank Mr. Bert Hopkins of the Air Force Electronic Systems Center (ESC), Global Information Grid (GIG) Systems Group for his support to this effort.

References

[1] D. Alberts, J. Gartstka, F. Stein, *Network Centric Warfare, 2nd Edition,* August 1999. http://www.dodccrp.org/publications/pdf/Alberts_NCW.pdf

[2] Department of Defense, Chief Information Officer, *Network-Centric Warfare: DoD Report to Congress,* July 2001. http://www.dod.mil/nii/NCW

[3] S. Renner, "Building Information Systems For Network-Centric Warfare", in *Proc. 8th Int. C2 Research and Technology Symposium*, Washington, DC, June 2003. http://www.dodccrp.org/events/2003/8th_ICCRTS/pdf/078.pdf

[4] Department of Defense, Office of Force Transformation, *The Implementation of Network-Centric Warfare*, January, 2005.
http://www.oft.osd.mil/library/library_files/document_387_NCW_Book_LowRes.pdf

[5] Department of Defense, Chief Information Officer, *DoD Net-Centric Data Strategy*, March 2003.
http://www.defenselink.mil/nii/org/cio/doc/Net-Centric-Data-Strategy-2003-05-092.pdf

[6] Department of Defense, Directive 8320.2, *Data Sharing in a Net-Centric Department of Defense*, December 2004.
http://www.dtic.mil/whs/directives/corres/pdf/d83202_120204/d83202p.pdf

[7] S. Chen, J. Dunagan, C. Verbowski, Y. Wang, "A Black-Box Tracing Technique to Identify Causes of Least-Privilege Incompatibilities", *Proceedings of Network and Distributed System Security Symposium*, February 2005.

[8] National Institute of Standards, National Security Agency, DRAFT *U.S. Government Protection Profile for Separation Kernels in Environments Requiring High Robustness*", July 2004, http://niap.nist.gov/pp/draft_pps/pp_draft_skpp_hr_v0.621.html