15th ICCRTS

"The Evolution of C2"

Privacy-Enabled Identity Management for Network-centric Command and Control (C2) Systems

Topic(s):

Networks and Networking, Information Sharing and Collaboration Processes

and Behaviors

Name of Author(s):

Gerald Beuchelt

Point of Contact:

Gerald Beuchelt
The MITRE Corporation
202 Burlington Rd.
Bedford, MA 01731
(781) 271-2000
gbeuchelt@mitre.org

Privacy-Enabled Identity Management for Network-centric Command and Control (C2) Systems

Gerald Beuchelt
MITRE Corporation
202 Burlington Road
Bedford, MA 01730

Abstract

In this paper we investigate existing identity management and privacy concepts and best practices and evaluate benefit for command and control (C2) systems. C2 systems are generally distributed among many semi-independent components, both logically and physically. Our basic proposition is that the application of privacy best practices such as minimization of data collection – even among components of the C2 system — significantly enhances the overall system security of distributed service-oriented C2 systems. At the same time we recognize that certain other privacy assumptions such as user consent are not necessarily relevant in such a setting[1]. Our intent is to demonstrate how the implementation of privacy principles can improve distributed system security, especially in the context of Network Centric system design.

[1] There can be situations, however, that might warrant user or organizational consent prior to releasing identity information, especially in classified mission scenarios or cross-organizational transactions.

1 Introduction

This paper is largely concerned with the security and privacy implication of this recent transformation from traditional, hierarchical, and monolithic C2 system to distributed and edge-oriented systems. For a more thorough discussion, see [1]. At the same time it should be noted that the ideas presented in this paper are not limited to military C2 systems: the theory of Command and Control has direct applicability to commercial applications in managing the value and supply chain of manufacturing processes.

In the first section we will review some of the existing principles that underlie modern privacy management frameworks and their relationship to information disclosure in existing and emerging identity management technologies. In the next section we review the relevant concepts of network centricity and service orientation, as outlined by the U.S. Department of Defense (DoD) Command and Control Research Program (CCRP)[2], and related initiatives. We then proceed to investigate the application of privacy-enabling technologies to the network-centric system Identity and Access Management (IdAM) design and spell out high-level patterns for using these within a service-oriented architecture.

In addition we look at passing of authority for action. Limitations on passing authority (i.e., passing minimum authority to get the job done, with tightest limits on further delegation) can improve the security of the entire system. This is the service chaining problem and depending on how authority is passed, it may not involve any identification information or privacy issues.

We conclude with an evaluation of the benefits, its applicability to network-centric design, and discuss open issues.

2 Concepts of Identity and Privacy

2.1 Identity Management

In recent years, driven especially by a fast growing Internet, identity management has established itself as a discipline in its own right, and not simply an aspect of security. It is possible to deploy existing identity management technologies without having to employ security technologies such as cryptography [2]. At the same time, practical applications and deployments obviously still require security technologies, at least for the majority of transactions and services. From this point-of-view, information security is an enabling technology for identity management. The identity management architecture has in turn a deep impact on operational security by defining the requirements for authentication and authorization, and establishing a trust framework between the actors.

[2] For more information, please visit http://www.dodccrp.org/

Most of our discussion should be accessible to the general technical reader. A few sections require familiarity with concepts and current technologies t for identity management. In particular, we will refer to the Security Assertion Markup Language (SAML) [3], WS-Trust [4], Liberty Alliance protocols [5], and other relevant technologies. Basic knowledge of network security related protocols and technologies such as LDAP, Kerberos v5, or X.509 based PKI, is also assumed. The reader should be familiar with basic deployment scenarios such as Federation or Circle or Trust or the InfoCard model.

It should be noted that the enterprise identity management community has seen strong influences from the "user-centric" IdM community[3] and other related emerging research fields such as Vendor Relationship Management[4]. While the underlying technologies are relevant for our discussion, the design philosophy of "user-centric" IdM has only limited compatibility with the organizational needs of network centric C2.

2.2 Common privacy principles

Privacy is a complex issue that has traditionally been largely defined by its legal and social dimensions. The concepts of privacy can be traced throughout history. With the advent of comprehensive data management technologies and—later—the internet, the concept of privacy and its maintenance has acquired a technical dimension as well.

For the purpose of this paper, we will identify a number of privacy principles that can be applied for enhancing the security of distributed systems. We recognize that this list does not cover all aspects of privacy; instead they have been selected based on their relevance for improving the security of the IdAM subsystem. For a more thorough discussion of privacy principles and their implementation in various legal instruments, see e.g., [6], [7].

Privacy Principle	Scope
Collection Limitation (CL)	Limit the creation, transmission, and collection of PII during the execution of mission threads.
Use Limitation (UL)	Use PII only for the purpose for which it was requested.
Access and Correction (ACC)	Enable access to PII and the ability to correct such data within the limits of policy.
Anonymity and Pseudonymity (P)	Use transient pseudonyms when possible and limit the exposure of identifiers.
Security and Safeguards (SECSAFE)	Provide a secure Information Assurance (IA) stance to protect the IdAM system.

Table 1: Privacy Principles Overview

[3] See e.g. http://www.burtongroup.com/Research/PublicDocument.aspx?cid=736 as in introduction to "user-centric" identity management

[4] See e.g. http://cyber.law.harvard.edu/research/projectvrm for more information on Vendor Relationship Management.

2.2.1 Collection Limitation (CL)

Any authorization system operating under the principle of collection limitation should only obtain the information necessary to make an informed authorization decision. No other information about the end-user or any other data subject should be requested or granted from one subsystem to another, unless it is necessary to meet needs. Note that accounting for later audit purposes might constitute a reason for data collection, as long as it is limited to only those attributes that are required for creating an audit trail through log aggregation.

Note that in traditional privacy frameworks, the principle of collection limitation includes the requirement for fair and lawful collection means. Obviously, this is assumed to be abided by in C2 systems.

2.2.2 Use Limitation (UL)

Use limitation typically restricts the use and retention of collected data for the purposes stated in the privacy statement and the system of records notice. In the absence of such instruments, the intent of this principle can still be applied in our context: any data collected for authentication, authorization, or accounting (AAA) will only be used for their respective use and logging. Any other data received will be discarded whenever this is possible.

2.2.3 Access and Correction (ACC)

From this generic principle we can adopt the notion that an end-user or a data subject might be authoritative for a subset of attributes and should thus be allowed to correct information pertaining to him. This is quite common in enterprise account management tools, where users can add information about them, such as cell phone numbers or alternative addresses. In the context of C2 systems, user-certified attributes must be carefully selected to not interfere with role or attribute based authorization systems.

2.2.4 Anonymity and Pseudonymity (P)

Anonymity or pseudonymity can significantly enhance the security of a system. This principle is—essentially—a corollary to the collection limitation principle which requires the limitation of data collection about the end-user or data subject. Completely anonymous authorization can be achieved by using trust-anchor signed attribute statements bound to ephemeral identifiers, but for the purposes outlined in this paper it is not useful.

2.2.5 Security and Safeguards (SecSafe)

While completely obvious, this privacy principle should still be mentioned: Without proper security[5] and safeguards, any system will quickly be compromised by intentional or unintentional interception of its normal operations. It is important to keep in mind that no security measure

[5] For at least the information and physical domain, but ideally also the other domains of C2, where applicable.

will guarantee that any attack can be avoided. As such, the architecture of the underlying security model must reflect the risk analysis for the system: what risks can be associated with the failure of the security and identity management infrastructure, and how much effort is justified or the prevention of security breaches.

2.3 Intersections

Traditionally, the security community has seen identity management as an aspect of information security. From such a point-of-view, security technologies such as encryption are at the heart of the problem, and the management of entities becomes an exercise left to the implementer.

2.3.1 Privacy and security

When comparing the goals of privacy management and use of personally identifiable information (PII) for information security, one might see some areas of agreement (e.g., encrypt identity data in transit) but other areas as conflicting. While security is concerned about restricting and auditing access to all sorts of information including PII, traditional privacy technology approaches have very often concerned themselves with the protection only of the data subject, including the end-user[6]. This point of view does not take into account the more complex inter-relationships between the technical aspects of privacy and information security. While it is fairly well established that security technologies such as ephemeral cryptographic keys, pseudonyms, etc., can be used to enable aspects of privacy through privacy enhancing technologies (PETs) [8], this paper demonstrates that privacy principles can also be used to improve the operational security of distributed systems. This is achieved by limiting the collection of, and protecting access to personally identfiable information about the end-user that is used during authentication or authorization. This applies to information about the user in both databases, as well as information provided by the user during the authentication and authorization processes. This way, privacy and security best practices and technologies are complementary, and not opposite to each other.

An important step for this process is to leave the data subject centric model for privacy behind and consider the role of the data steward as well. Shapiro et al. have argued for considering the role of “Enterprise Privacy Enhancing Technologies” (ePETs) that focus on the data steward and the protection of personally identifiable information (PII) once it has been collected[7].

[6] Note in this context that at least 3 of the privacy principles are not about technology or about keeping secrets.

[7] Shapio recently presented on ePETs at the “10th Annual Privacy and Security Conference”, Victoria, BC, Canada. See http://www.rebootconference.com/privacy2009/agenda.php

2.3.2 Privacy and Identity Management

The identity management architecture determines if and how privacy of authentication and authorization information can be preserved during normal operation, and how the identity of a pseudonymous transaction can be resolved though log auditing. When designing identity management protocols such as SAML [3], [9], Liberty Web Services [5], or the InfoCard Identity Metasystem [2], enabling privacy was already part of the underlying design principle. The editors of SAML created a security assertion language that could not only be used to convey authentication information, but instead only assert specific attributes of an entity in an authentication or authorization transaction.

These privacy-focused design principles of the identity management community are documented in a Liberty Alliance Deployment Guide [10]. For an end-user/data-subject specific sub-set of guidelines, one might also refer to the "Laws of User-centric Identity" [11].

3 Network Centric C2 Systems

It is beyond the scope of this paper to review the theory and concepts of command and control systems. Throughout the history of warfare, there has been a constant effort to understand the underlying processes and variables of C2 and improve its efficiency. Crucial to the success of Network Centricity is the ability to adapt to the needs of the mission and the warfighter. Traditionally, the focus for C2 has been on successfully completing complex large scale missions that could involve tens of thousands of individual participants. With the change in the global security environment over the past 20 years, research in C2 has gone beyond the industrial age, top-down C2 approach, and started to explore Network Centric C2 and "Power to the Edge" [12] concepts.

We do not advocate the decentralization of decision processes at all cost. Despite the current trend to de-centralization, some problems can still be solved best at a global level. For example, the U.S. National Airspace System, which encompasses the management of the entire U.S. airspace, about half of the North Atlantic, and the better part of the Northern Pacific is a highly centralized system, where a lot of the effectiveness is achieved through the ability of the Air Traffic Control System Command Center (ATCSCC)[8] to set the operating parameters for the entire system. True Network Centricity requires the agility to change the operational approach from a centralized to a de-centralized approach.

3.1 Network centric system design concepts

To illustrate network-centric operations 'requirements, let us revisit the four basic tenets, as originally described in [13]:

[8] See e.g. http://www.fly.faa.gov/

- *A robustly networked force improves information sharing.*
- *Information sharing enhances the quality of information and shared situational awareness.*
- *Shared situational awareness enables collaboration and self-synchronization, and enhances sustainability and speed of command.*
- *These, in turn, dramatically increase mission effectiveness*

These tenets, combined with the governing principles [14] determine the needs for the organization of the Information Domain in the Conceptual Framework. This organization requires enabling end-users to quickly obtain situation awareness and to contribute to the overall information base. This transformation of command and control happens on all levels: from the enterprise to the individual, thus empowerment, self-sustainability, and enhanced collaboration become ubiquitous.

3.2 Creating an Information Infrastructure

The autonomy of individuals and organizational units is mirrored in how systems should be build for a network-centric approach: system should be service oriented, meaning that complex services and are built from autonomous sub-systems that are interconnected through a fabric of interoperable networks. In the vision for a future architecture of Global Information Grid (GIG) [15], the DoD has determined that loosely coupled[9] service-oriented approaches are fundamental to creating an information infrastructure suitable for Network Centric Operations (NCO).

This general approach closely mirrors the architecture found in modern enterprises, where service-orientation has now been practiced for some time. As such, the identity management issues found in enterprise SOAs such as reservation systems, inventory management or business activity monitoring [16], are also present in the information systems enabling network-centric operations.

4 Identity Management for Network Centric System Design

Privacy has not traditionally been a major concern in operational Command and Control (C2) systems. Privacy has a profound impact on how we operate within society and protect our individuality, but application developers are generally concerned with executing their tasks. In the C2 environment, the mission focus and fact that most participants are associated with the military further led to reduced emphasis on privacy. However, the operational security and

[9] For more information on loose coupling, especially in the context of organizational science, see [27].

integrity of each of the C2 system's components is recognized as critical – and can actually be improved by improved privacy.

Since at least the beginning of this decade, identity management practitioners have been addressing the public's concern for protecting user privacy in identity management systems. As such, identity management became an enabler for privacy protection. The central idea presented here is somewhat opposite of the traditional approach of relating federated identity management and privacy [17]: We are here proposing to use privacy best practices and privacy enabling technologies for improving the quality of federated identity management, in particular the operational security aspects.

4.1 The case for privacy-enabled Identity Management

In the C2 community, accountability and chain of command are important. It is thus usually assumed that the identity of all actors must be well established at all points and at all times within the 'shell' of a C2 system. Authorization models such as Attribute Based Authorization Control (ABAC) require the presence of reliable attributes for a given digital identity to compute the access policy decision for any given service invocation – but can then often substitute for identity when communicating with other components. Thus, ABAC can facilitate the usage of privacy-enabled identity management technology, as will be shown below.

On the other hand, limiting the amount of information—especially sensitive information such as identity data—that gets transported or permanently shared across a distributed system has a number of positive effects:

- Restricting the content of over-the-network security tokens limits the possibility of obtaining sensitive or classified data through token interception.
- Federating user accounts allows the creation of local user profiles, thus eliminating the need to manage complex user schemas in the central account database. Trusting another system's authentication and authorization results in less personally identifiable information about users being stored in and transmitted across networks.
- Limiting local user account databases to contain only service/capability module specific user attributes diminishes the value of these local databases to intruders.

It is public knowledge, that foreign agents have already gained access to military systems [18] and to critical U.S. infrastructure [19]. There is a concrete potential peril not only to national security or critical infrastructure, but also for individuals. For example, they might gain access to security tokens that, beyond attributes necessary for authorization, unnecessarily reveal PII such as a full name. If the token for a high-ranking official with multiple Top Secret/SCI compartment authorizations was intercepted by an unfriendly insider, the identity of that official could be exploited by external attackers on the system, and could therefore result in a national and personal safety issue.

In general terms, attribute information that is bound to an end user's or a device's identity is very valuable to protect. Capture and correlation of such data should be prevented to avoid negative impact. In commercial environments this need to protect PII is well understood in the context of Identity Theft. Through capture and exploitation of information that is specific to a person's identity (both digital and real) criminals may impersonate that person to obtain financial, medical, or other services. Any successful impersonation in a C2 environment can have devastating effects on the mission success. This is amplified by the highly decentralized character of NCO C2 processes.

By applying privacy protection principles—which originated as a legal concept—to the identity management architecture of C2 information systems, users can be protected from these emerging threats.

Beyond these immediate benefits, privacy-aware identity management architecture may provide other, subsequent benefits:

- In today's highly connected environment, information exchange between systems of different security levels (e.g., classified/unclassified) or compartments typically requires strict operational procedures, such as the enforcement of mandatory access controls. More modern, "service-friendly" approaches for Cross Domain Solutions (CDS) are complex and require strong meta-data tagging. Privacy-enabled identity management can help mitigating some of the CDS concerns such as revealing to one domain's users the identity of a source in the other system, As a result, it reduces the risk of a more connected architecture for systems of different security domains.
- The U.S. Department of Defense implements the requirements of OMB Circular A-130 and the 1974 Privacy Act through DoDI 5400-11R [20]. This program requires all DoD business systems to provide comprehensive privacy impact assessments and implement risk mitigation strategies for all DoD business systems that might store or access personally identifiable information. While there is currently no regulation for operational C2 systems, future changes might require advanced PII protection for all DoD systems. This may become especially important in the context of the required use of the PIV Federal Agent Smart Card Number (FASC-N)[10].

[10] See FIPS PUB 201-1 for more information on the use of the FASC-N.

4.2 Applying Privacy Principles to loosely-coupled C2 systems

With the basic privacy principles we identified in Section 2.2, we can now examine how these improve the overall system security. For each principle we will first examine the security ramifications, and offer a simple example, use-case, or deployment scenario for applying this principle.

4.2.1 Collection Limitation

Limiting the amount of end-user or data subject specific information that is collected by the C2 System and it subsystems, but also non-authorized parties has a significant impact on the overall system security. Limiting the amount each service provider or other component is allowed to collect helps still further. In applying this concept we propose the following:

> **CL1**: Any service provider will only request and collect the amount of information about a data subject or end-user that it needs.
>
> **CL2**: Information about a data subject or end-user that gets sent unsolicited must be discarded.
>
> **CL3**: A record of the relevant data should be created, as long as data retention and data-at-rest policies are applied.

Achieving these guidelines is fairly easy with existing technology: SAML allows the creation of signed and encrypted attribute statements that only contain attributes that pertain to the subject. Limiting the content of these attributes allows restricting the amount of shared information, especially when including derived attributes[11]. Alternatively, authorization protocols such as XACML [21] may be used to restrict the networked shared information to simple decisions.

A concrete example of using privacy principles CL2 and CL3 in the context of distributed systems was published by the U.S. National Security Agency (NSA) in a SAML attribute profile for the Net-Centric Enterprise Service (NCES) architecture [22]. This profile employs XACML for creating a distributed attribute provisioning system, using a X.509 PKI for authentication. The intended use of this profile is for creating an attribute based access control (ABAC) capability.

4.2.2 Use Limitation

When information about the end-user or data subject gets collected, this principle specifies that it should only be used for the intended purpose, like e.g., authorization. There is a conflict with net-centricity, which attempts to make data usable for unanticipated purposes.

[11] Such as a Boolean value for the attribute “Over 21” instead of an attribute statement on the birth date.

When reusing any data for non-intended purposes, there is always the risk of misinterpreting the semantics of the transmitted data. Also, as in commercial applications, reuse of information can damage the trust between parties, especially if the reuse is not communicated back to the individual.

Thus, the following principles are stated as absolutes for PII. In reality, they should be considered, and traded against other organizational and personal goals, such as the value of analyzing after-action reports.

> **UL1**: End-user or data subject data received for authentication, authorization, or any other IdAM purpose must not be re-used for any other purpose. In particular, such data must only be used for authentication or authorization steps for which it was released.
>
> **UL2**: Data retention policies should be established and implemented, so data is retained only as long as needed.

For the scope where these apply (i.e. PII) these guidance statements conflict with the principles of NCW, which prescribe the reuse of information whenever possible. However, Use Limitation applies to the use of PII only and does not cover data in general. Since PII is critical for ABAC-based authorization, UL1 and UL2 apply specifically to data exchanged for AAA.

4.2.3 Access and Correction

Access and Correction to data within the components of a distributed C2 system can be defined as the ability to obtain and modify or annotate information for which the participant is responsible. The extent to which this is permissible within any C2 systems depends on how this instance operates: C2 theory identifies three fundamental dimensions for Command and Control approaches [23]: allocation of decision rights, distribution of information, and patterns of interactions. The right to change security-enabling information about themselves will probably only be found in highly "edge oriented" organizations, with peer-to-peer hierarchies, highly distributed interaction patterns, and a fairly broad dissemination of data.

> **ACC1**: Components, data subjects, or end-users should be enabled to access all information for which they are authoritative and be allowed to make corrections or amends, as long as these are permissible under the access control policy.
>
> **ACC2**: The access control policies for Access and Correction should be adjustable over purpose and time.

User managed corrections are particularly useful for access regimes that allow broad access, but audit afterwards. In such cases, the user might supply information (e.g., location, purpose, person one is to meet) that can later be audited.

The parameters determining the level of Access and Correction can easily change over time or purpose, so any distributed C2 systems that aims at supporting a highly agile organization must be able to support changes to the permissions.

The critical problem for access and correction is the determination of authority for a given data source. There are two major categories of data sources: one the one hand those, for which an authoritative source exists[12]. For example, the FAA is the authoritative source for who is a licensed pilot in the U.S. or an employer is authoritative on the employee identification number. Many other data describe a real world situation rather than an assertion. For these, the information may come from different sources. There may be no way to determine which is correct, but the consumer's organization may designate a particular source as the one to be used. One example is the location of an aircraft: there might be sensor data available from different sources (including the aircraft crew itself). There is no final authority that can determine the "correct" location, but a consumer may choose to say, for our purposes, we use the following source.

The goal of applying this privacy principle is to enable the IdAM subsystem to determine dynamically the role of the data subject and allow self-assertion in cases where appropriate. For C2 systems, this can becomes critical for self-synchronizing[13] elements in the force.

4.2.4 Anonymity or Pseudonymity

Anonymity or pseudonymity allow hiding the user's "real-world" identity from service providers and other relying parties. While true anonymity implies that the identity of the user cannot be dereferenced through log linking, a pseudonymous transaction can be reconstructed by linking and cross referencing [24]. In C2 applications, true anonymity is often undesirable, especially since this make the determination of accountability for command extremely difficult, if not impossible. This is also expressed in the non-repudiation security principle. As such, we will restrict the application of this principle to pseudonyms only:

P1: Data in transit should only contain references to pseudonyms.

P2: Pseudonyms can be persistent or transient; transient pseudonyms are preferable where permissible by operational requirements.

P3: Pseudonyms for a given user should be valid only between two systems. For different pairs, different pseudonyms should be used.

[12] In some cases, the assertion is of the form "designatedEntity asserts X" and that entity is by definition authoritative about what it asserts. Note that a detailed discussion on authority of attribute sources is outside the scope of this paper.

[13] Self-synchronization is one of the fundamental principles of network-centric operations. See e.g. [13].

Application of these principles reduces the amount of data that has to be transported over the network. It also ensures that logged identity information is valuable only to users that can correlate them across different systems. Requiring the use of limited validity pseudonyms limits the damage that can result from the exposure of any account database, for example through an unfriendly insider.

At the same time it should be noted that there are also applications where there is no operational need to identify the end-user and maintain a record about the transaction. For example, end-users will have to access commonly available databases to develop a comprehensive situational awareness. Since many of these databases contain commonly available, unclassified data, it is not necessary to log successful access to this data[14].

4.2.5 Security and Safeguards

We have included security and safeguard mostly to raise the awareness for their importance for privacy and identity management. There is a long history of excellent analysis on this topic, and there are many works on best practices in information security available[15].

4.3 Sample Patterns

In this section we will look at some existing high-level patterns that implement the principles outlined above. For a more comprehensive treatment of lower-level security patterns, especially for JEE, see e.g., [25]

4.3.1 Central authentication

This pattern is the traditional central authentication pattern, as implemented in an enterprise directory, Kerberos realm, or NIS domain: a central system (here denoted Identity Provider – IdP) maintains all information about the user. Relying party will either redirect the user to the IdP for authentication, or validate their credential by passing cryptographic artifacts (salted hashes, including nonces, etc.) to the IdP for verification. As a variation, the IdP may use multiple sources for aggregating a central authentication database from various sources.

Authorization is not addressed in this pattern.

Implemented principles: CL1, CL3, UL1, ACC2, P1

[14] In fact, when accessing an unclassified source from a secure network through a Cross Domain Solution, the identity of the end-user must not be revealed. This paper will not go into detail. Cross Domain problems will be further discussed (briefly) in section 5.2.

[15] See for example [28], [29]. As a minimal baseline for operational systems the DoD 8500 instructions series provides a set of Information Assurance controls.

4.3.2 Attribute-based Authorization

Verifiable subject attribute statements can be used to authorize access to resources. When using this pattern, the relying party requires the user to present an attribute statement that must satisfy certain conditions. These conditions typically include:

(i) Attribute type
(ii) Attribute source
(iii) Value

Other conditions may include the quality or assurance of correctness, the transport used for obtaining the statement, etc.

Authentication and authorization to see attribute information is not addressed in this pattern, but can reuse the basic pattern. The attribute provider is often collocated with the IdP, in which case a standard central authentication pattern might be sufficient.

This pattern is very valuable for authorization in typical Network Centric systems: since these systems may be highly partitioned, and access must be granted to very different classes of users[16], it is necessary to develop a system that is more scalable than traditional access control lists. The aforementioned NSA/DISA SAML profile [22] is an example for this pattern.

Implemented principles: CL1, CL2, UL1, ACC2, P1, P2

4.3.3 Federated authentication and authorization

In a federation, authentication can flow from one authentication provider (such as the IdP) to other members in the federation. Note that this means that there can be more than one IdP, and successful authentication at e.g., an SP may result in a full single-sign one.

A federation is more complex to setup than any of the earlier patterns, but with correct policies applied it is the most powerful. Any federation protocol should support single-sign off as well.

Implemented principles: CL1, CL2, UL1, ACC2, P1, P2, P3

[16] It is not uncommon that user from different services (Army, Navy, Air Force, Marines) operate within a single command and control center during joint operations. Access to systems may be different for members of different services.
In addition, there are many operations that are conducted with international coalition partners which require a high level of coordination. In such scenarios, a command center can requires members of foreign militaries to have access to some C2 systems which requires centrally managed and universally enforced authorization policy.

4.3.4 Delegation

This pattern is commonly used in service-to-service interactions where a service (Service Consumer [SC]) needs to act on behalf of the user and interact with another service (Service Provider [SP]). Through delegation, the SC should be able to obtain a user-specific token for use with the SP. This is fairly straightforward in a centralized or federated environment, where an access token can be created by the IdP or attribute provider.

With an appropriate authorization and policy framework in place, one can extend this pattern to a constrained delegation pattern can be constructed where the token that the SC obtains is limited to a subset of uses, single use with the SP, or even only a single operation.

Network Centric operations require delegation: a commander must be able to delegate authority for an operation to a number of different combat units. With the commander's intent as a background, these units self-synchronize their operations and must get access through delegation to resources, as authorized by their commander.

Implemented principles: CL1, CL2, UL1, ACC2

4.4 Auditing requirements

For this paper we will not address the requirements for a secure log architecture. Privacy-enablement of the identity management architecture outlined in this paper does not conflict with the requirements of accountability – which is a legitimate use. Even completely pseudonymous transactions can be logged, the ephemeral identifiers recorded and if there is a need, resolved to the user accounts that performed the transaction. It is important however, that the entirety of all logs is available—therefore, logging information should be aggregated and archived at a higher organizational level.

5 Conclusions

5.1 Benefits

5.1.1 Immediate security benefits

The most immediate benefit and the original motivator for this paper is an increase in system security for the overall system. We have already discussed the impact of the application of each privacy principle in its respective section. In general, applying the privacy principles to the development and operation of the identity management component, limits the creation, distribution, and retention of sensitive data, i.e., the sensitive or classified data-at-rest and data-in-transit.

To our knowledge, no existing C2 (or enterprise) systems were developed with comprehensive identity management architecture in mind. In fact, many systems have evolved from small function specific entities to complex distributed and inter-connected systems of systems. As a result, provisioning and de-provisioning users is complex and prone to mistakes that can result in orphaned accounts. Combined with an insufficient authentication strategy, incompletely de-provisioned accounts can result in significant security issues [26].

By applying federated principles one limits the usefulness of user accounts of individual system components or capability modules: de-provisioning of the central account databases is limited to a small number of well-managed identity provider systems. Local accounts are—effectively—rendered useless, by terminating the federation. This assumes that authentication function have been delegated to the IdPs.

5.1.2 Decoupling of security and identity

As indicated in the beginning of the paper, identity management and security are related, but not the same: the former requires some form of the latter for secure operation, but they are decoupled enough so that the security system can be replaced without affecting the identity management layer. Given the rapid development of computing systems, especially in the context of security and cryptography, a clean separation between the identity management layer[17] and the security layer allows a decoupled evolution of the two: once the deployed security technology becomes too weak (e.g., though the availability of cheap high-speed AES decryption tools, or the potential vulnerability of ECC algorithms), an improved security layer can be deployed without having to modify the existing user management infrastructure.

It should be noted that this clear separation will be hard to achieve in practice, since in most systems user identity is still tightly coupled to the security subsystem. As identity management technologies evolve, this goal should be easier to achieve[18].

5.2 Open Issues

5.2.1 CDS

As mentioned before, systems that operate in different security domains are currently facing complex issues when trying to interoperate. While CDS is fairly well understood, its implementation in a highly dynamic service-centric environment remains problematic. As such, any identity architecture that builds on top of web protocols faces significant problems, especially when complex meta-data tagging is required for connecting high and low systems.

[17] The abstract digital identity of the end-user which includes his attributes.

[18] For example, many SAML-related products allow today already the integration of different authentication and security subsystems, while maintaining a high-level notion of the user's identity across different systems.

The application of privacy principles such as collection limitation may help in downgrading information from higher security levels to lower ones by providing a finer data granularity with respect to the attributes needed for authorization. For example, if authorization for access to a higher level resource was granted through a security token, it would be much easier to clear this token if it contained only authorization-relevant attributes (e.g. the role of the authorizer and not the full name, FASC-N, or affiliation).

5.2.2 Broader applicability, beyond PII

For identity management, much of the data is PII. Access management can include resource metadata or environmental factors such as unit status (e.g., currently on patrol at a particular location).

But the potential benefits of privacy principles are larger, applicable to any data that is sensitive. In particular, collection limitation, use of pseudonyms (e.g., for target locations) are of obvious value, and deserve to be incorporated into more processes.???

6 Acknowledgements

Special thanks for insightful comments and discussions: Steve Foote, Ray Modeen, Cliff Baker, and Arnie Rosenthal.

7 Bibliography

[1] D. S. Alberts and R. E. Hayes, *Understanding Command and Control*. DoD Command and Control Research Program, 2006.

[2] M. Jones and M. McIntosh, "Identity Metasystem Interoperability Version 1.0," 2009.

[3] OASIS Standard, "Assertions and Protocols for the OASIS Security Assertion Markup Language (SAML) V2.0," 2005.

[4] OASIS Standard. (2007, Mar.) WS-Trust 1.3. [Online]. http://docs.oasis-open.org/ws-sx/ws-trust/200512/ws-trust-1.3-os.html

[5] Project Liberty Alliance. (2002-2009) Project Liberty Alliance. [Online]. http://www.projectliberty.org/liberty/specifications__1

[6] International Security, Trust and Privacy Alliance, "Analysis of Privacy Principles: Making

Privacy Operational," 2007.

[7] American Institute of Certified Public Accountants, Inc., "Generally Accepted Privacy Principles: A Global Privacy Framework," 2006.

[8] IPCR (Ontario, Canada); Registratiekamer (The Netherlands), *'Privacy-Enhancing Technologies: The Path to Anonymity*. 1995.

[9] OASIS Standard, "Assertions and Protocol for the OASIS Security Assertion Markup Language (SAML) V1.1," 2003.

[10] Project Liberty Alliance, "Deployment Guidelines for Policy Decision Makers v2.9," 2008.

[11] K. Cameron. (2006, Jan.) IdentityBlog. [Online]. http://www.identityblog.com/?p=352

[12] D. S. Albert and R. Hayes, *Power to the Edge: Command and Control in the Information Age*. CCRP, 2003.

[13] Department of Defense, "Network Centric Warfare: Report to Congress," 2001.

[14] Office of Force Transformation, "The Implementation of Network Centric Warfare," 2005.

[15] Department of Defense CIO, "Global Information Grid Architectural Vision," Department of Defense Version 1.0, 2007.

[16] J. McKendrick. (2006) Service Oriented. [Online]. http://blogs.zdnet.com/service-oriented/?p=508

[17] S. Landau, H. Le van Gong, and R. Wilton, "Achieving Privacy in a Federated Identity Management System," 2009.

[18] A. Cole, Y. Dreazen, and S. Gorman. (2009, Apr.) Wall Street Journal Online. [Online]. http://online.wsj.com/article/SB124027491029837401.html

[19] S. Srinivasan. (2009, Apr.) Epoch Times. [Online]. http://www.theepochtimes.com/n2/content/view/15058/

[20] Department of Defense, Office of the Director Administration and Management, "Department of Defense Privacy Program ," DoDI 5400.11R, 2007.

[21] OASIS Standard, "eXtensible Access Control Markup Language (XACML) Version 2.0," 2005.

[22] National Security Agency, "Net-Centric Enterprise Services (NCES) Security Assertion Markup Language (SAML) Attribute Profile," 2008.

[23] SAS-050, "Exploring New Command and Control Concepts and Capabilities," 2006.

[24] L. F. Cranor and S. Spiekermann, "Engineering Privacy," *IEEE Transactions on Software Engineering*, vol. 35, no. 1, 2009.

[25] C. Steel, R. Nagappan, and R. Lai, *Core security patterns*. Prentice-Hall, 2006.

[26] M. Wheatley. (2003, Sep.) CSO Online. [Online]. http://www.csoonline.com/article/218425/Deprovisioning_Firing_Line

[27] K. E. Weick, *Sensemaking in organizations*. SAGE, 1995.

[28] NIST, *An Introduction to Computer Security: The NIST Handbook*. NIST, 1995.

[29] M. Swanson and B. Guttman, *Generally Accepted Principles and Practices for Securing Information Technology Systems*. NIST, 1996 .