Connecting Multi-National Command Centers
Technical Report

FOR THE

PORTABLE, REUSABLE, INTEGRATED SOFTWARE MODULES (PRISM) PROGRAM

CONTRACT NO. F19628-92-C-0006

Prepared for:

Electronic Systems Division
Air Force Systems Commands, USAF
Hanscom AFB, MA 01731-5000

Prepared Jointly by:

G. G. Lonardo
PRISM Technical Director
RAYTHEON COMPANY
Equipment Division
528 Boston Post Road
Sudbury, MA 01776

J. D. Wallin
PRISM Technical Director
HUGHES TECHNICAL SERVICES COMPANY
Ground Systems Operations
1901 West Malvern Avenue
Fullerton, CA 92634

STATEMENTS:

DISCLAIMER - The United States Government is absolved from any responsibility and/or liability for the contents of this document.

SUBMITTAL - This submittal is made in accordance with the Statement of Work (SOW) section 4.1.2.2.3

DEVELOPMENT - The primary authors of this document were Jeff Williams and Dave Wichers of Arca Systems.

APPROVAL:

_____ _____

J. D. Wallin Date
PRISM Technical Director
Hughes Aircraft Company

_____ _____

G. G. Lonardo Date
PRISM Technical Director
Raytheon Company

1.0 Problem Statement

Several GCC mission objectives have a critical need for multilevel secure (MLS) connectivity with multi-national command centers. These objectives address command center operations in a diversity of situations, including peacetime, crises, and wartime. Some examples of operations where MLS connectivity is required include:

- Situation Monitoring
- Force and Resource Monitoring
- Situation Assessment
- Hostilities Termination and Negotiation
- Strength and Readiness
- Data Communications

This paper examines some of the fundamental security difficulties in establishing MLS connections between multinational command centers. There are several different tiers of issues:

Need	The need for information sharing comes from the mission to be accomplished. For example, the U.S. might need to share information with Canada in order to accomplish a particular mission.
Agreement	An agreement between the parties who need to share information should be reached. The agreement should cover what information will be shared and the protection requirements for the information.
Policy	Each party must set a policy on how the agreement will be upheld. Policies can address how information will be separated, shared, or isolated from other information.
Implementation	Implementation refers to the mechanisms needed to enforce the policies necessary to meet the agreement. While some mechanisms are directly related to a particular policy, others are more general.

The remainder of this paper examines the policies and implementations involved in addressing these issues.

2.0 Background

2.1 Security Domains

Security domains are environments where there is a common enforcement of security. Domains are defined by:

- Policies to be enforced
- Levels/sensitivities of information handled
- Connections allowed
- Range of information allowed on each connection
- Sensitivity label form/label meaning

Once a piece of information has left a security domain, there is *no* further control over that information. Therefore, passing information to another domain requires a great deal of trust. First, the information must be protected while in transit between domains. Second, because there is no way for the automated system to retain control over released information, the security domain must trust that the receiver will protect the information appropriately. .

When information is shared between domains, each domain must uphold their agreement by protecting that information appropriately. Therefore, both domains are likely to have similar policies for protecting that information. In this sense, security domains *overlap* to protect shared information.

This discussion assumes that *within* a domain, all components are interoperable and support the same policies. This may be a broad assumption, but the issues identified in this paper apply to domains of all sizes, from standalone computers to large networks.

2.2 Policies for Exchanging Information Between Domains

In general, a security policy is the set of laws, rules, or practices that regulate how a domain protects and distributes sensitive information. Clearly understanding each domain's policy is critical for identifying common ground which will allow domain's to share information, and protect this shared information in a manner consistent with the policy of each participating domain. Countries are used as examples of domains throughout this paper for illustrative purposes, but these could just as easily be armed services, Government agencies, or companies.

Within the U.S. DoD/Intelligence community, the policy for manual handling of classified documents includes numerous classification, dissemination, and handling restrictions. These restrictions are represented by easily identifiable markings applied to each document. Their primary policy is:

- Basic Classification Policy: Information is kept isolated from any person or organization not authorized for the information based on two characteristics. The sensitivity of the document (in the form of its sensitivity level (e.g., Secret, Top Secret) and any associated categories (e.g., Compartment A, Category X)) and the clearance of the intended recipient (in the form of a clearance level and any compartments they have a need-to-know for (i.e., has been read into)).

This basic policy is sometimes augmented with additional restrictions or handling caveats. Some examples include:

- Releasability Policy (REL): Information is authorized for release to a country or international organization.
- NOFORN Policy: Information is Not Releasable to Foreign Nationals
- ORCON Policy: Dissemination and Extraction of Information Controlled by Originator

The INFOSEC community has started to develop automated technologies to enforce some aspects of these policies. NSA's Trusted Computer System Evaluation Criteria (TCSEC) and DIA's Requirements for Compartmented Mode Workstations (CMW) lists requirements for computer systems which can label the sensitivity of objects and the restrict access to those objects based on the clearance of the users. Vendors have started to develop systems that can protect files based on these requirements.

Systems built to the TCSEC or CMW can enforce the basic classification policy described above. The CMW requirements include the ability to enforce releasabilities while TCSEC systems do not. The general nature of TCSEC and CMW labeling requirements supports the ability to do NOFORN type labeling and enforcement as well. However, these types of systems do not support the ability to enforce an ORCON policy. Even if they did, they would be limited to enforcing the ORCON policy internally. Once information was released outside the system, the ORCON policy (as well as any other policy for that matter) could not be enforced by the system, except through the cooperation of the outside systems receiving the information.

These capabilities illustrate that, although specific policies may be agreeable to different parties, it may be difficult to implement them with current technology. Therefore, one must be careful when developing a policy for protecting exchanged information to ensure that the technology can support such a policy.

2.3 Treaties and Agreements

The first step in establishing connections between domains is to understand what information is to be communicated, particularly the sensitivity levels of the information. The nature of the information is not important (e.g., mission data, positioning data, troop movement data) for security enforcement, the protection requirements based on its

classification level are. This negotiation is based on the need for sharing defined by a domain's mission. The understanding is documented in treaties and agreements between the two countries in question.

The question becomes very difficult to resolve if the two parties cannot arrive at an agreement about how to protect transferred information. Both sides must agree to sufficiently protect the other's information from compromise. This agreement can be complicated if the meanings of terms are not well defined. For example, the U.S. protects SECRET data according to its set of rules. Canada also has data called SECRET, but protects it according to a different set of rules. Therefore, the identifier of the information is the same, but the meaning is different.

3.0 Policy Coordination

This section deals with establishing system policies that define how an agreement will be enforced. Both the compartmented and releasability policies are relevant for interconnected multi-national command centers. Information can be shared only as long as each domain can keep it in the proper compartment and only release it to authorized parties.

The simplest case of policy coordination is when all the domains are enforcing the same set of policies. In this case, the only possible problems will be in implementing the policy in the different systems. This section examines some possible differences in policies and their effects on multilevel interdomain communication.

3.1 Sharing Between Domains with Different Compartment Policies

Connecting domains that have the same compartment policy is purely an implementation problem. It is possible to connect domains that do not have the same policy for compartmenting information, but several issues arise.

If one domain supports compartments and another does not, the domains can only share uncompartmented information. Therefore, the connection is largely one-way. Figure 1 illustrates the problems that can occur when domains with different comparment policies are connected. In this case, England can share Sensitive data with the U.S., but is not allowed access to Sensitive A information.

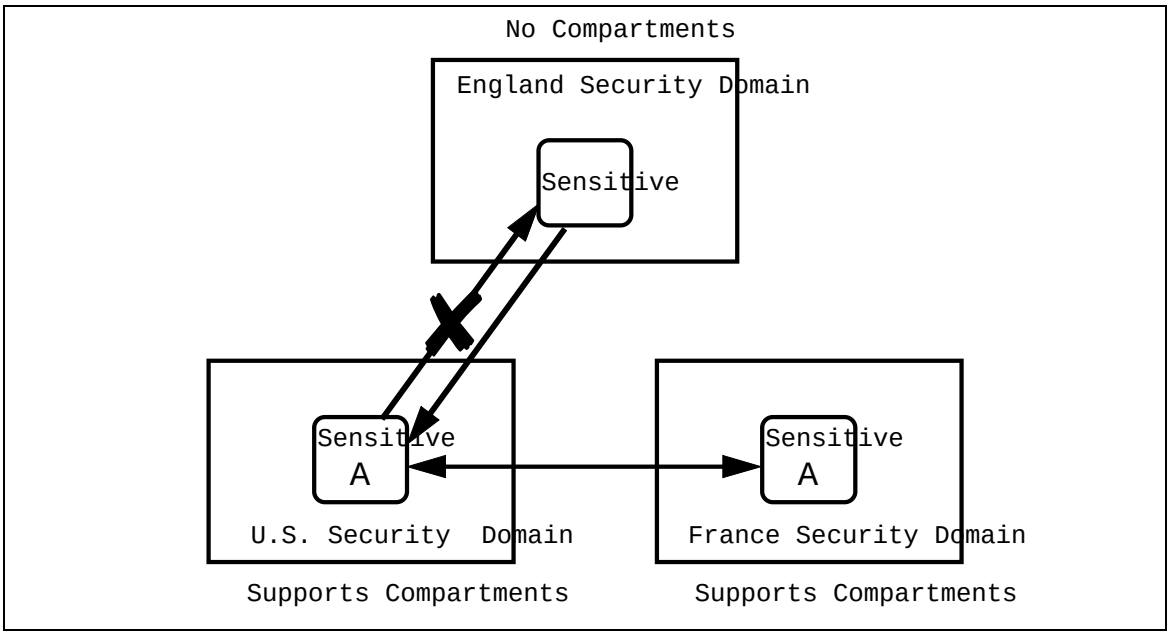

Figure 1 – Connecting domains with different compartment policies can result in limited information flows.

Even if both domains support compartments, they must also agree on exactly how the information is to be protected. Information is only as safe as the weakest protection it is given in any of the domains.

Therefore, in order to share compartmented information, both domains must support compartments and must agree to protect common compartments identically.

3.2 Sharing Between Domains with Different Releasability Policies

When information marked with releasabilities is transferred to a domain that does not support releasabilities, the information must “float” to a higher sensitivity level. This occurs because when information loses releasability, it is not as widely available.

In a large interconnected set of domains with a lot of information transfer, there is a real danger that a lot of information will “float.” This can result in data being overclassified and will therefore needlessly prevent otherwise authorized use. Several copies of identical information may also have to be maintained at several levels.

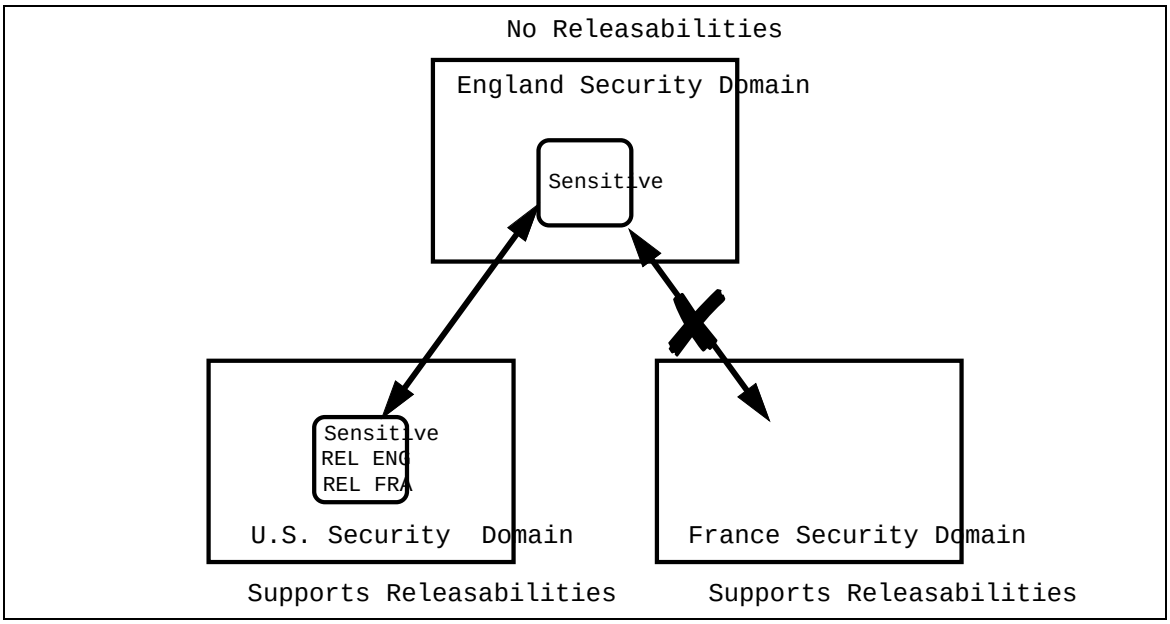

Figure 2 – Information may "float" when transferred between domains with different releasability policies.

Figure 2 shows how a piece of information marked "Sensitive Releasable to England and France" floats to being marked "Sensitive" when it is transferred to England which does not support releasabilities. This results in France not being able to receive the information from England.

Another aspect of sharing releasable information is that all domains involved must agree to protect releasable information consistently. If the domains do not enforce a common policy, then information may get propagated much farther than originally intended. Figure 3 depicts three homogeneous domains which each support compartment and releasability policies, using the same label format. There is a danger that France will get unauthorized U.S. data if England fails to adequately protect it.

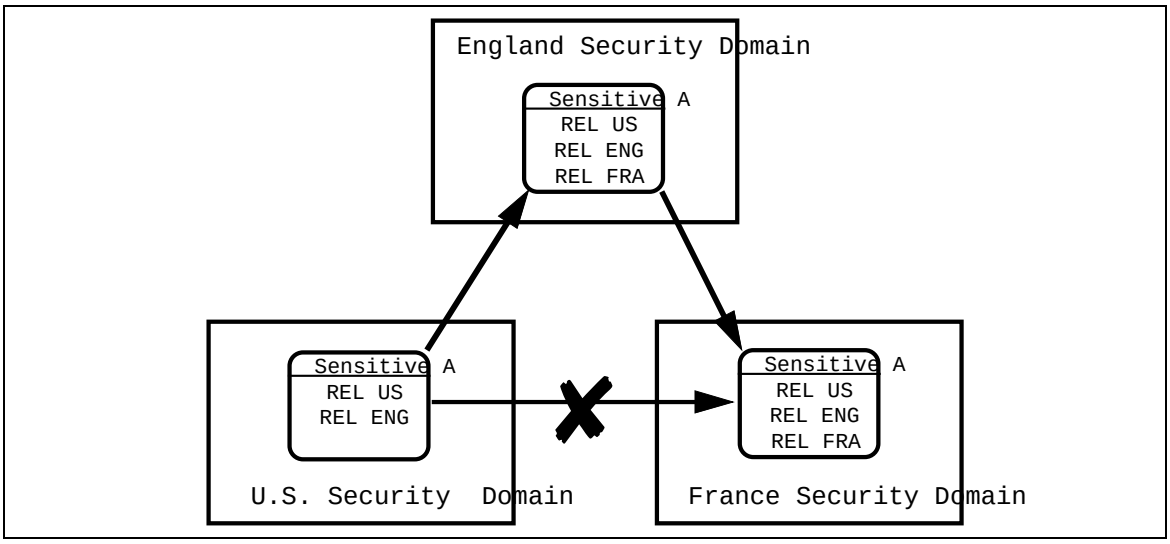

Figure 3 – Domains should not increase the releasability of information received from other domains.

Therefore, in order to share releasable information, both domains must support releasabilities and should enforce a common releasability policy.

3.3 Maintaining Originator Control

There is no guarantee that data will remain protected even between domains that enforce the same compartment and releasability policies. The problem is that once information has left a domain, the originator has no control over that information.

This is an example of need for the ORCON policy. Unfortunately, it is very difficult to implement ORCON across interconnected domains. Once a piece of information is outside the control of a domain, that information is impossible to protect with existing technology. The only approach to dealing with this problem is to negotiate strong agreements between domains that need to share information.

4.0 Implementation of Interdomain Sharing

This section discusses the coordination of policy implementation necessary to establish multilevel communication with other domains.

Specifically, this section concentrates on the use of labels, which are the most common means for implementing multilevel policies. In fact, there is very little difference in the way that the releasability and compartment policies are usually implemented. Usually, one bit in the label is assigned to each compartment or releasability. For example, data for a special program A to be shared with Canada would have one bit in its label for program A and one indicating releasability to Canada.

In the GCC, labels seem to be the most likely mechanism to implement the compartment and releasability policies. Two parts of the label are of particular interest, compartment bits and releasability bits.

A **compartment bit** is part of the label used to indicate that the information is part of a compartment. A compartment indicates a grouping of information to which only authorized personnel are allowed access. Compartment bits are used to *raise* the sensitivity of information.

A **releasability bit** is part of the label used to indicate that the information is authorized for transfer to another domain. Releasability bits are distinguished from compartment bits in that they are used to *lower* the sensitivity of information (i.e., make the information available to more people).

All implementations for interdomain sharing require assurance that sharing is performed correctly and effectively. This assurance should come from a rigorous evaluation of the sharing mechanisms and risk analysis of the system that needs it.

4.1 Use a Common Labeling Format

Selecting a common labeling format is the conceptually simple approach to achieving multilevel communication between domains. If all command centers could just adopt a particular standard, then the implementation problems would be minimal. However, there are many difficulties with achieving this goal.

First, there is no agreement in the security engineering community as to exactly what the right format for interdomain communication is. This format can be thought of as a language for communicating security information. There are many candidates for this language in the form of network protocols, but all have weaknesses.

Secondly, there are many systems that either do not or cannot support any label formats. Many systems today are not being built with the capability to label information. The GCC

environment requires the use of many systems that do not support labels. To integrate these types of systems, special considerations, such as guards, routers, and translators, must be used.

4.2 Label Translation

When domains support the same labeling policy, but do not use the same labeling format (syntax) or label meaning (semantics), translation is required. Label translation can be performed either by the sender, the receiver, or both. There are several ways to implement translation of labels. The specific technique selected depends on the differences between the labels and the desired approach.

First of all, the levels and compartments/categories of primary concern are those that are to be shared. By definition, those that are not shared would not require translation. For the shared levels/compartments, if the format of the labels are different then the label will have to translated from one format to the other. If the format is the same, but the meaning of the particular bits in the label are different, then translation will be required, but it should be easier than converting from one format to another. If most of the bits in the label are the same (e.g., the same meaning for sensitivity levels, and many identical compartments/categories) then the translation would only have to deal with the specific bits that are different, rather than translating the entire label. If the number of levels/categories to be shared is small, or the number of differences in meaning are small, translation should be fairly straightforward.

Translation unnecessary involves agreeing to the same format and meaning for all the levels and categories/compartments to be shared. Thus, when a label is sent from one system to the next, the label is automatically understood by the receiving system without the need to translate. This may not work for releasabilities, since they are relative the owner of the information (e.g., for U.S. data marked REL Canada, what does it mean when that data is sent to a Canadian system.)

For releasabilities, if the U.S. uses a bit as REL Canada and the Canadians use the same bit as REL U.S., then when such information is sent back and forth between the two countries, the meaning of the bit will automatically change to the correct value, based on which system it is contained in. However, this simple trick will not work for more than two countries (consider the passage of REL Canada and Rel U.K. data from the U.S. to the U.K. to Canada and what the meaning of each bit would be without translation as it moved from system to system).

A solution to this releasability problem is to have each system use a self releasable bit (e.g., in the U.S. mark everything REL U.S.). Then went it is sent to another country, they can use the same bit to mean REL U.S. as well. (See Section 4.3 for more info) However, this may be confusing since people are not used to seeing information marked self releasable.

Specific translation involves translating labels according to their originating domain (either by sender, receiver, or a guard in-between). The translation required can be for any number of reasons. If the format or meaning of any bits are different, translation would be required. Translation could also be used to add on and strip off releasabilities between domains (e.g., for data going from Canada to U.S. the REL U.S. could be stripped off, and a REL Canada put on). The difficulty of the translation depends on the exact differences between the labels (both format and meaning).

The difficulty in this approach is developing the translators and deciding where the fit within an architecture and the interconnections between domains. If is it feasible to have only a small number of connections between two domains, then the translators can be concentrated at the connection points (e.g., routers), minimizing the affect on the rest of the domain and centralizing the management of the effort. If the two domains are heavily interconnected then possibly each machine would have to be able to perform translations, which might be difficult to implement and manage.

Specific translation can also get complicated if the number of domains one is connected to is large. Such a domain would have to be able to translate between each and every domain it connects to.

Universal translation involves translating each label to a standard format and meaning before transmission. This approach requires that all systems understand the universal format and meaning. This approach is much more feasible if the number of distinct connections between domains is small (ideally only one). The benefit of this approach is that each domain only needs to understand one other format. One translates to and from this universal format and meaning when sending and receiving any information.

The difficulty of this approach is coming up with an agreed upon format that everyone can translate to, and dealing with the security issues surrounding the sensitivity of compartment/category names that may be defined into this universal format and meaning. This clearly depends on the number and nature of the domains involved.

4.3 Self-Releasability Bits

Information is generally assumed to be releasable to the country or international organization where the information is located. In many computer systems, however, there is no bit to indicate this “self-releasability.” In other systems, the bit exists, but its use and meaning are not clear.

If a self-releasability bit is not used, a translation must be performed to create the appropriate markings for the destination. In this case, the marking must be translated from REL US to REL FRA based on its location.

If a self-releasability bit is used, all domains can share the same set of bits without translation. For example, information to be shared between the U.S. and France should be marked REL US/FRA regardless of its current location, if sharing is desired without label translation.

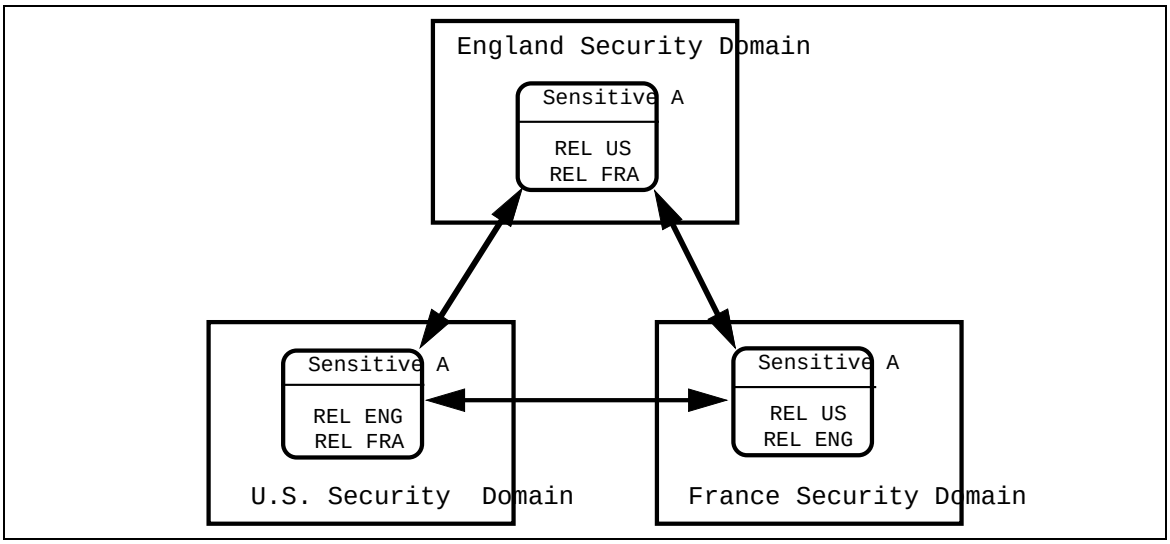

Figure 4 – Releasabilities must be translated if a "self-releasability" bit is not used.

Figure 4 shows three domains that all support compartments and releasabilities. In this example, all three systems share the same label format but do not use self-releasability bits. Therefore, as information gets transferred between systems, the labels must be translated. This translation can occur when data is sent, when it is received, or some combination.

5.0 Recommendations

The following are recommendations for establishing connectivity between multi-national command centers:

- Use physical, personnel, and procedural security measures to implement policies that cannot be implemented with today's technology.
- Plan the system with a multi-national concept in mind. Future systems are very likely to be interconnected and planning now will help to mitigate security problems in the future. The rules for information labeling must be defined. Try to avoid "hard-coding" the security policy into the implementation, as policies change frequently, especially in a multi-national environment.
- Adopt a risk-based approach to determining what sorts of policies are acceptable. This approach involves assessing the assets, threats, and vulnerabilities in a domain and determining the likelihood of a compromise. Select policies that address the highest priority risks. Determine the rules for exchanging information in terms of levels, compartments, and releasability.
- At an implementation level, choose a protocol that supports both releasability and compartment bits. Then use whatever translation components (routers, bridges, guards) necessary to connect domains. Standard protocols can be a very effective approach to establishing connectivity across multiple domains.
- Also at the implementation level, try to use the same label encodings across all domains if possible. At a minimum, share the encodings for the portions of the compartmenting and releasability policies that overlap.

There are many additional aspects of this issue that deserve additional research, especially those dealing with assurance. Interesting work is being done in the area of *composition*. Several areas of this work are applicable to the GCC, including composition of:

- identical stand-alone components
- different components at the same assurance level
- different components at different assurance levels
- components intended for different modes of operation
- untrusted components with trusted components

6.0 The GCC Demonstration Configuration

Recommendations:

- Establish exactly what sorts of policies are needed in a GCC.
- Examine the GCC domain to determine if consistent policies are enforced.
- Continue to seek out the latest protocols.
- Universal translation could potentially implemented with half-bridges UNO proposal (from a reviewer comment)
- Use common data dictionary for translation (from a reviewer comment).
- Use firewalls/guards/routers to implement translation.

7.0 References

Abrams, M., Heaney, J., King, O., LaPadula, L., Lazear, M., Olson, I., "Generalized Framework for Access Control: Towards Prototyping the ORGCON Policy," in Proceedings of the Fourteenth National Computer Security Conference. October 1991.

Director of Central Intelligence Directive No. 1/7, Control of Dissemination of Intelligence Information, 4 May 1981.

Hosmer, H., "The Multipolicy Paradigm," in Proceedings of the Fifteenth National Computer Security Conference, October 1992.

Neugent, B., "General Issues to be Resolved in Achieving MultiLevel Security (MLS)," in Proceedings of the Fifteenth National Computer Security Conference, October 1992.

Williams, J. and Day, M.L., "Sensitivity Labels and Security Profiles," Proceedings of the Eleventh National Computer Security Conference, October 1988.