



NRL/MR/5542-97-7951

The JMCIS Information Flow Improvement (JIFI) Assurance Strategy

ANDREW P. MOORE

*Center for High Assurance Computer Systems
Information Technology Division*

19970710 067

May 30, 1997

DTIC QUALITY INSPECTED 4

Approved for public release; distribution unlimited.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE May 30, 1997	3. REPORT TYPE AND DATES COVERED 1 Oct. 1995 - 30 Sept. 1996		
4. TITLE AND SUBTITLE The JMCIS Information Flow Improvement (JIFI) Assurance Strategy			5. FUNDING NUMBERS PE - 33401G	
6. AUTHOR(S) Andrew P. Moore				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory Washington, DC 20375-5320			8. PERFORMING ORGANIZATION REPORT NUMBER NRL/MR/5542--97-7951	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Director National Security Agency 9800 Savage Rd. Ft. Meade, MD 20755-6000			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) The Joint Maritime Command Information System (JMCIS) provides a common operating environment for Naval tactical decision aids that currently operates two distinct system high enclaves, one at SECRET/GENSER and one at TOP SECRET/SCI. NRL Code 5540 is developing an extension of JMCIS, called JIFI (JMCIS Information Flow Improvement), to improve the timeliness and accuracy of GENSER information available to SCI JMCIS analysts while maintaining the security posture of the system. This document describes the strategy for developing the evidence that JIFI satisfies its critical security requirements. The strategy views databases in more classified enclaves as potential replica sites for data from less classified enclaves. Replicated data flows from lower enclaves to higher ones via simple one-way connections, yielding a high assurance MLS distributed system. The system high enclaves ensure discretionary security. The one-way connections are the only trusted component with respect to mandatory security. The JIFI architecture incorporates a one-way communications device, called the Pump, and existing COTS database replication technology to provide the extended JMCIS function. The JIFI assurance strategy described here complements and exploits modern system design methods, which separate data management from data processing, and enables effective low-cost MLS operation within the paradigm.				
14. SUBJECT TERMS JMCIS Concept of operation DoDIIS Assurance Copernicus Certification Information Security			15. NUMBER OF PAGES 48	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

Contents

Chapter 1	Introduction	1
1.1	Purpose	1
1.2	Certification and Accreditation.....	1
1.3	Document Structure.....	2
1.4	Prerequisites	3
Chapter 2	Background	4
2.1	Copernicus.....	4
2.2	JMCIS	6
2.3	The Problem	7
2.4	Approach	8
Chapter 3	Problem Definition	11
3.1	Current Operations	11
3.1.1	Tactical Command Organization	11
3.1.2	Structure and Operation of the Battle Group	12
3.2	Required Additional Capability	17
3.3	Abstract Model of Operations	18
Chapter 4	Architecture Overview	21
4.1	Physical Model of Operations	21
4.2	JIFI Gateway Requirements	25
4.3	JMCIS LAN Extensions	28
Chapter 5	Gateway Assurance Strategy	30
5.1	Assuring Confidentiality.....	30
5.2	Protecting Against Tampering.....	33
5.3	Assuring Reliability.....	35
5.3.1	Reliability of the Pump.....	35
5.3.2	Reliability of Sybase Replication Technology.....	37
5.4	Assuring Good Performance.....	40
Chapter 6	Residual Risk.....	42
Appendix A.	Notation	44
Bibliography	46

Executive Summary

This document describes the strategy for developing the evidence that an extension of the Joint Maritime Command Information System (JMCIS) [24] satisfies its critical security requirements. The motivation for this work is the Multi-Level Security (MLS) Processing for Copernicus 6.3A Core Technology Program, the primary objective of which is to facilitate the development and implementation of the Copernicus architecture [37] by considering information security as a core element of the overall Copernicus design. The Copernicus concept can overcome many of the shortfalls of existing C⁴I support by eliminating wholesale message broadcast in favor of a smart push and warrior pull approach that permits maintaining a common tactical picture across a theater of operations. JMCIS provides a common operating environment for Naval tactical decision aids that supports this approach by replacing existing stovepipe systems with client-server architectures. Although JMCIS operation is currently limited to a local area, such as a facility on shore or aboard ship, it is an important step in providing an integrated, tailorable and flexible C⁴I support system for tactical mission planners.

JMCIS currently operates two distinct system high enclaves, one at SECRET/GENSER and one at TOP SECRET/SCI. Current security requirements severely restricts information flow between the enclaves, forcing the separate (redundant) processing of GENSER Naval messages by both the GENSER and SCI systems. Furthermore, relevant GENSER information created by GENSER analysts must be manually carried to the SCI enclave or be recreated by SCI analysts. This process is costly, error-prone and slow. If critical GENSER information is unavailable to the SCI analyst or is inconsistent with the SCI tactical picture, erroneous or contradictory observations may be made, possibly leading to failure of a mission. NRL Code 5540 is developing an extension of JMCIS, called JMCIS Information Flow Improvement (JIFI), to improve the timeliness and accuracy of GENSER information available to SCI JMCIS analysts while maintaining the security posture of the system. Since compartmented information is at risk, the Defense Intelligence Agency (DIA) is the accrediting authority, the Office of Naval Intelligence (ONI) is the certifying authority, and the DoD Intelligence Information Systems (DoDIIS) documents [14,15,16] provide guidelines for developing and maintaining a certifiable and creditable system.

Our approach to the development of JIFI is based on the need to present a clear and convincing argument, called the *assurance argument*, to persuade the accreditor that the risk of compromise is small enough to justify operating the system. The strategy for developing this argument is called the *assurance*

strategy and is the focus of this document. Our approach integrates security and system engineering to permit the explicit tradeoff of security requirements with other critical system requirements. We use the languages of Statemate [20,21,22,23], based on the formal theory of statecharts, as a rigorous foundation for illustrating the operational requirements and design of JIFI graphically. We use a variant of the Goal Structured Notation (GSN) [33,40,46] and the Assumptions/Assertions Framework [39] to state requirements in terms of Statemate primitives and refine them according to the Statemate decomposition. This permits tracing the security risk through the Statemate specification, thus strengthening the correspondence between the functional description and the security analysis. Finally, we use the Literate Assurance Approach [38] to help present the assurance strategy and argument in a manner convincing to certifiers and to ensure that the documentation is consistent with the actual specification and implementation.

This document focuses on a relatively small, but important, part of the larger Copernicus problem: improving information flow between the GENSER and SCI JMCIS enclaves aboard ship, i.e., in the aircraft carrier's Tactical Command Center (TCC). Each enclave contains a communication server, for processing and correlating external communication, a central data server (CDBS), for storing tactical data for access by JMCIS clients, and a set of operational facilities or workspaces. These facilities - including the Aircraft Carrier Intelligence Center (CVIC), the Combat Information Center (CIC), and the Ship's Signal Exploitation Space (SSES) - contain workstations that behave as JMCIS clients. Extending JMCIS to make GENSER information more readily available to the SCI enclave while maintaining a common tactical picture requires ensuring the security and consistency of the information and the reliability, recoverability and good performance of the underlying implementation.

The strategy that we adopt for extending JMCIS in a way that ensures these requirements are met is based on the SINTRA (Secure Information Through Replicated Architecture) paradigm [18]. This paradigm views databases in more classified enclaves as potential replica sites for data from less classified enclaves. Replicated data flows from lower enclaves to higher ones via simple one-way connections, yielding a high assurance MLS distributed system. The system high enclaves ensure discretionary security, i.e., the protection of information based on the identity and need-to-know of the user. The one-way connections are the only trusted component with respect to mandatory security, i.e., the protection of data based on the classification of the data and the clearance of the user. Applied to JIFI, this paradigm permits the use of the existing physical distribution of the GENSER and SCI enclaves and the development of a gateway between enclaves as the primary means for providing the enhanced (JIFI) function while ensuring information security. The gateway incorporates a one-way communications device, called the Pump, and an existing COTS database replication product, called the Sybase Replication Server, developed by Sybase. Goal structured graphs, documented using GSN, capture the strategy for achieving high security assurance balanced with requirements for reliability, recoverability, affordability and performance. Our approach is

consistent with DoDIIS recommendations to “focus on system-high client server operations with trusted interfaces to environments operating at different security levels.” [17]

The strategy that we have adopted for developing JIFI complements and exploits modern system design methods, which separate data management from data processing, and enables effective low-cost MLS operation within that paradigm.

THE JMCIS INFORMATION FLOW IMPROVEMENT (JIFI) ASSURANCE STRATEGY

Chapter 1 Introduction

1.1 Purpose

This document describes the strategy for developing the evidence that an extension of the Joint Maritime Command Information System (JMCIS) [24] satisfies its critical security requirements. JMCIS provides the common operating environment for Naval tactical decision aids operating either on shore or aboard ship. NRL Code 5540 is extending JMCIS to improve the information flow between users operating at different security levels while maintaining its security posture. We call the extended JMCIS system the JMCIS Information Flow Improvement (JIFI). JIFI is being developed as part of the Multilevel Security (MLS) Processing for Copernicus 6.3A Core Technology Program. An objective of this program is to demonstrate a capability for Navy mission planners operating at different security levels to access the data that they need in a timely and accurate manner and with high assurance that classified information is not compromised.

Mission planning takes place in a distributed environment where individual component commands in a theater of operations refine and execute mission objectives. Success of the mission planning process depends upon the ability of a component command to access quickly data that originates from diverse sources, such as other component commands, while ensuring that sensitive information is not compromised. Maintaining a common (i.e., consistent) and accurate picture of the tactical environment among these commands is paramount. While JMCIS, itself, is not a mission planning system, it is a platform for managing the tactical data needed by mission planning applications such as the Tactical Aircraft Mission Planning System (TAMPS) Version 6 [2]. JMCIS operation is limited to a local area such as a facility on shore or aboard ship, but may include users and operations at both the SECRET/GENSER and TOP SECRET/SCI levels. The goal of JIFI is to provide TOP SECRET/SCI mission planners with the SECRET/GENSER tactical data they need to do their job without compromising the confidentiality of that data.

1.2 Certification and Accreditation

Any modification or extension to JMCIS requires accreditation. The accreditation authority for JIFI is the director of the Defense Intelligence Agency (DIA) since compartmented information is at risk; the certifying authority is the Office of Naval Intelligence. DIA takes a site-based approach to security

certification and accreditation as outlined in the DoD Intelligence Information Systems (DoDIIS) guidelines [14,15,16]. Rather than evaluating information systems on a system-by-system basis, DIA's site-based approach certifies and accredits systems within a defined area, called a site, the security of which is managed by the Information System Security Officer (ISSO). Before a site's accredited baseline may be modified or extended, the ISSO in coordination with the certifying authority must conduct a security evaluation of the changes and provide accreditation recommendations to DIA.

DIA requires developers of systems that process intelligence information to develop a set of security documents that, once approved, forms an integral security certification record. These documents provide necessary information to define the accredited baseline for the site in which the system is to be integrated. The documents are listed below in the order in which DIA suggests that they be developed:

- System Security Concept of Operations --- describes the operation and architecture of the planned system identifying all of the intended users, their clearance levels, access approvals, and need-to-know authorizations.
- System Security Analysis --- identifies the risks associated with the operation of the system in its defined environment and the safeguards (countermeasures) used to counteract vulnerabilities
- System Security Requirements --- describes the security requirements mandated by the level of trust targeted for the system and the relevant standards and directives [7,12,36]
- System Test Plan --- presents a set of steps to prove satisfaction of each security requirement
- System Test Procedures --- presents a set of operational instructions to execute the steps identified in the test plan
- System Test Report --- presents the results of the execution of the test procedures and, if warranted, the certifying authorities approval to operate

This document provides information relevant to the first four DoDIIS documents. The refinement of this assurance strategy into an assurance argument will detail the System Test Plan and document the Test Procedures and Test Report.

1.3 Document Structure

Assurance that a system counters the threats of interest depends on the effectiveness of the security mechanism as well as the correctness of the system's design and implementation. A system implementation may correctly satisfy a set of security requirements, but may be easily subverted, for example, by crashing the system. Likewise, a system may be based on a very effective security mechanism, such as a non-bypassable Reference Monitor, but a programming error could allow low-level users to access high-level information.

This document defines a strategy for producing an effective and correct implementation of JIFI. Chapter 2 provides relevant background information concerning the problem that we are addressing and our approach to solving the problem. Chapter 3 defines an architecture-independent concept of operations for JIFI that is used to construct a model of operations. Chapter 4 identifies the architecture for JIFI based on the DoDIIS security mode of operation. Chapter 5 defines a strategy for gaining assurance that a major component of this architecture, the JIFI Gateway, is secure, reliable and a good performer. Finally, Chapter

6 analyzes the strength of the JIFI architecture, specified as the (residual) risk that remains after the components are implemented according to the derived requirements. The appendix to this paper reviews notation used in this document. A comprehensive graphical overview of the assurance strategy is given in a fold-out included at the end of the document.

Our approach to arguing the effectiveness of the JIFI architecture promotes a slightly different structure for the certification documentation than that advocated by DoDIIS. Our approach integrates security and system engineering to permit the explicit tradeoff security requirements with other critical system requirements. This reduces the redundancy, incompatibility and documentation maintenance problems that accompany separate security and development documents. We believe that this approach results in certification documentation that is easier to assess and change.

1.4 Prerequisites

Much of this document assumes only a basic understanding of information security and information system architectures. A high level understanding of the conceptual and physical models presented requires some understanding of graphical (CASE or CAD) design languages; a more detailed understanding requires familiarity with the languages of Statemate [20,21,22,23]. Readers will also find useful a basic understanding of the Goal Structured Notation [33,40] and the Assumptions/Assertions Framework [39], which form the basis for our method of requirements specification and risk analysis.

Chapter 2 Background

The end of the cold war has shifted the national security strategy from large-scale, global combat and containment to small-scale, regional conflict resolution. This new operational environment and reductions in military budgets and personnel have emphasized the need for joint service cooperation. The C⁴I for the Warrior program was developed to address the global C⁴I requirements for all services in joint operations. C⁴I includes electronic technology, warfare doctrine, personnel, procedures and facilities that support tactical command and control of warfighting units. The Copernicus Program [37] defines the Navy's role in meeting the objectives of C⁴I for the Warrior in today's dynamic tactical environment.

The overall objective of the MLS Processing for Copernicus 6.3A Core Technology Program is to facilitate the development and implementation of the Copernicus architecture by considering information security as a core element of the overall Copernicus design. This chapter briefly describes the Copernicus architecture and the role JMCIS plays in it. We describe the problem that JIFI addresses and our approach to developing JIFI's implementation and assurance evidence.

2.1 Copernicus

The Copernicus Program recognizes eight shortfalls of existing C⁴I support [25,37].

- **Command and Control Inflexibility:** the lack of support for defining flexible threat-based command and control doctrine;
- **Inefficient Communications Management:** the overloading of communication circuits and systems by using the same (scarce) communication bandwidth for both high priority operational traffic and low priority administrative traffic;
- **Dependence Upon Message Format:** the requirement that sites understand and parse human-readable narrative messages to extract relevant information and correlate with the current tactical picture;
- **Push/Broadcast of Information:** the wholesale broadcast of information to afloat units whether or not those units have a need for the information;
- **Dated Communications Technology:** the inadequacy of existing physical transmission systems for allowing tactical commanders to establish virtual circuits and to better manage available communications bandwidth;
- **Ambiguous Reporting:** the ambiguities that result due to the independent reporting of contacts by multiple sensors;
- **Limited Intelligence Infrastructure:** the inability to use existing intelligence networks to contact colleagues in State, CIA, DIA and industry who are working on the same problem but from a different perspective; and
- **Inefficient Intelligence Dissemination:** the slow dissemination of intelligence data in inefficient formats resulting in receipt of outdated information.

The Copernicus architecture describes the structure for a C⁴I support system that addresses the shortfalls of existing systems. It is divided into four distinct “pillars” as shown Figure 1: the Global Information Exchange Systems (GLOBIXS), the Tactical Information Exchange Systems (TADIXS), the Commander-in-Chief (CINC) Command Complex (CCC), and the Tactical Command Center (TCC). The CCC is located ashore and the TCC is located afloat. GLOBIXS is the communication system that the CCC uses to communicate with the outside world. TADIXS is the communication system that the TCC uses to communicate to the CCCs and other TCCs.




Figure 1: The Copernicus Architecture

The CCCs use a Metropolitan Area Network (MAN) to group theater-level command centers on shore. The MAN interfaces with Local Area Networks (LANs) at each command center. The CCC continually upgrades its databases with strategic, tactical, logistical, administrative, and technical data that it obtains over GLOBIXS from shore sensor nodes, weather facilities, analytic nodes, higher-echelon authorities and other CCCs. GLOBIXS will permit establishing virtual networks that are customized by the CCC to respond to specific threats. The CCC makes its tactical databases available to the TCCs via TADIXS.

The TCC groups the tactical centers for a Battle Group together via a LAN (or possibly a number of LANs for different communities of interest). The TCC provides the Battle Group Commander, or in the case of joint operations the Multi-Force Commander, with access to tactical communications, tactical displays of track data, fused intelligence, fleet status and normal administrative information. The TCC can access information available at the CCC via virtual networks provided by TADIXS. Like GLOBIXS, TADIXS can be tailored to accommodate specific tactical situations faced by the commander, so that the information needed can be accessed quickly and efficiently.

The open systems client-server approach advocated by Copernicus will allow users to retrieve data from remote or local data sources as they need it. Instead of always broadcasting data to end users,

common data servers are updated and tactical commanders can access that data, or not, based upon the tactical environment and their own set of priorities. Data is correlated prior to its insertion in the common data server and is stored in a format to promote efficient processing. Combined with leading edge communication and networking technology, the Copernicus concept can overcome many of the shortfalls of existing C⁴I support system by eliminating wholesale message broadcast in favor of a smart push and warrior pull approach. This approach reduces duplicate reporting, reduces bandwidth requirements, improves command and control flexibility, and promotes a consistent tactical picture among command centers ashore and afloat.

The promise of Copernicus will only be realized if the supporting technology can be identified and, if necessary, refined. Many of the critical building blocks already exist, but significant hurdles have yet to be cleared. Primary among these hurdles is the move from an architecture where Naval messages are transmitted at the discretion of the sender to a client-server architecture where tactical commanders can request from common data servers the information that they need. This involves replacing the current technology used for communication between ship and shore, which is based upon Naval message broadcast and manual bulk update, with the TADIXS technology. Significant progress is being made on board ship using JMCIS to replace stovepipe architectures with client-server architectures. Although this does not solve the larger problem, it is a fundamental step in developing an integrated, tailorable and flexible C⁴I support system for tactical commanders.

2.2 JMCIS

JMCIS integrates Naval command and control applications to provide a common operating environment for tactical decision aids supporting track management, data correlation, communication and tactical display. The JMCIS single security level (system high) environment is depicted in Figure 2: JMCIS function is distributed across a LAN of workstations; JMCIS data is centralized into a single repository, called the Central Data Base Server (CDBS) [26,27]. Naval messages are received and processed by a Communication Server and the CDBS is updated appropriately. These system high environments, or *enclaves*, are usually either SECRET/GENSER or TOP SECRET/SCI. The rest of this document refers to SECRET/GENSER and TOP SECRET/SCI as GENSER and SCI, respectively, for simplicity.




Figure 2: The JMCIS Single Level (System High) Environment

JMCIS, as currently implemented, has not been approved (nor is appropriate) for multi-level operation. Nevertheless, accreditors have approved a limited form of communication between the GENSER and SCI enclaves. As illustrated in Figure 3, information can be transmitted in both directions between GENSER and SCI. In the SCI to GENSER direction, a trusted user operating the Esprit/Opus sanitizer determines what collateral information in the SCI enclave can be downgraded to the GENSER level. In the GENSER to SCI direction, the Communication Server operating in the GENSER enclave broadcasts GENSER message traffic from the GENSER communications network to the SCI Communication Server for correlation into the SCI CDBS. The SCI Communication Server provides no acknowledgement of the correct receipt of the GENSER data since this would complicate accreditation by providing a channel for leaking SCI data to the GENSER enclave.

2.3 The Problem

The primary problem in Navy (and DoD) tactical C⁴I systems is the maintenance of a common tactical picture where multiple distributed sources of information must be correlated and each source may have its own view of the tactical environment. Information classified at different levels and stored on system high systems contributes to the difficulty of maintaining a common tactical picture. The Copernicus and C⁴I for the Warrior programs are looking for innovative technologies to migrate DoD systems to cooperative, distributed multi-level secure computing. This requires advances in many technology areas, which are being investigated in Code 5540's 6.3A Program. To limit the scope of our effort, we decided to concentrate on a relatively small, but important, part of this larger problem: improving information flow in JMCIS shipboard operations, i.e., in the aircraft carrier TCC.

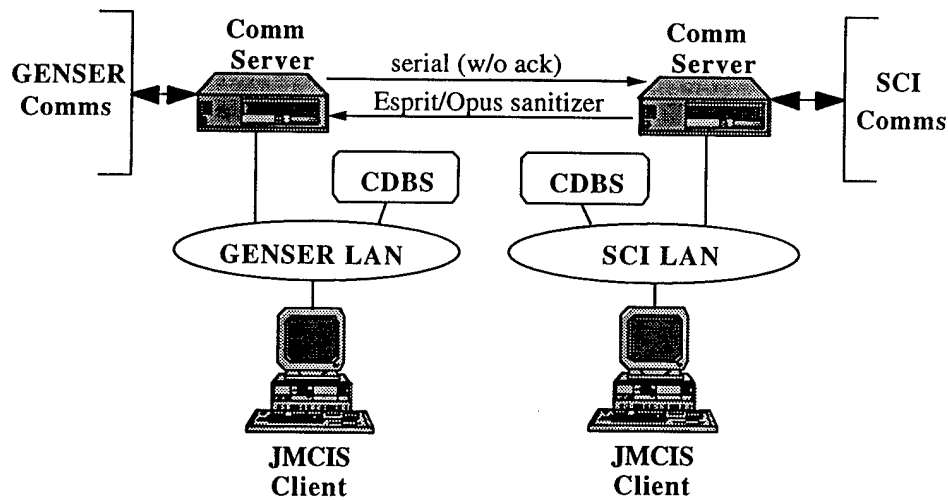


Figure 3: GENSER/SCI JMCIS Interconnection

Current operations in the TCC require that GENSER Naval messages be processed separately by both the GENSER and SCI systems. No acknowledgement is provided to the GENSER Communication Server that the SCI database correctly received the GENSER messages, due to the risk of leakage of SCI data to the GENSER enclave. Unfortunately, this may lead to loss of critical GENSER updates to the SCI CDBS. Furthermore, relevant GENSER information created by GENSER analysts must be manually carried to the SCI enclave (e.g., by disk or tape) and loaded into the SCI CDBS or be recreated by SCI analysts. This process is costly, error-prone and slow. If critical GENSER information is unavailable to the SCI analyst or is inconsistent with the SCI tactical picture, erroneous or contradictory observations may be made, possibly leading to the failure of a mission. The goal of JIFI is to improve the timeliness and accuracy of the tactical GENSER information available to SCI JMCIS analysts while maintaining the security posture of the system.

2.4 Approach

Our approach to the development of JIFI is based on the need to present a clear and convincing argument, called the *assurance argument*, to persuade the accreditor that the risk of compromise is small enough to justify operating the system. This document defines the strategy for developing the assurance argument for JIFI. This assurance strategy will be elaborated and refined throughout JIFI's development, yielding the assurance argument.

The integration of security engineering and system engineering is fundamental to our approach. This integration allows the explicit tradeoff among security and operational requirements and the development of an assurance argument that has a strong correspondence with the system implemented. We use the languages of Statemate [20] as a rigorous foundation for graphically illustrating the operational requirements and design of JIFI. The Statemate toolset, based on the formal theory of statecharts [19],

allows modeling system behavior and graphically executing this model to test its validity. The Statemate specification provides a formal structure for the assurance strategy and argument. Critical requirements are stated in terms of the Statemate primitives and refined according to the Statemate decomposition to strengthen their correspondence with system specifications. An overview of the notation used in Statemate charts is given in the appendix. The detailed definitions of elements used in the Statemate specification are provided in the JIFI Elements Dictionary [35].

We use a variant of the Goal Structured Notation (GSN) to represent the assurance strategy graphically. An overview of the modified GSN syntax [40] is given in the appendix. GSN was originally developed to represent overviews of safety arguments for safety-critical systems. We have extended GSN to improve readability for security assurance strategies and to support analysis using the Assumptions/Assertions Framework [39]. Within this framework, *assertions* are statements about the security that a particular INFOSEC discipline (computer security, communications security, administrative security, personnel security and physical security) is required to provide. *Assumptions* document requirements that one discipline places on another. For example, the computer security discipline may assume that its users are cleared for the most sensitive information that it processes; the personnel security administrator must ensure that procedures are performed for clearing users to that sensitivity level. Each assumption about some security discipline should match an assertion for another discipline; a gap in this mapping indicates a vulnerability.

Goals in the GSN syntax map to assertions in the Assumptions/Assertions Framework; assumptions map to assumptions in the Framework. To match assumptions with their validating assertions, we number goals and identify the numbers of the validating goals for each assumption after the statement of the assumption in the goal structured graph. If an assumption can not be validated, the letter V is used to indicate a vulnerability. Goals are numbered according to the decomposition of the goal structured graph as presented in this paper. That is, the *i*th goal structured graph presented in this paper starts with the goal *i.1* at the root of the graph. Subgoals of this root are numbered *i.2*, *i.3*, etc. going from left to right and top to bottom with one exception: goals that form the root of subsequent goal structured graphs are leaves of the *i*th goal structured graph and are numbered *i+1.1*, *i+2.1*, etc. In the case of a goal structured graph (say the *j*th) that has multiple root nodes, the root nodes are numbered *j.1*, *j.2*, etc. Thus, the *i*th goal structured graph presented in this paper always starts its numbering at the root with *i.1*. This numbering scheme minimizes the chance that a change to a particular goal structured graph will cause changes to the numbering of other goal structured graphs.

Our approach also uses the Literate Assurance Approach [38] to help present the assurance argument clearly to system certifiers. Literate programming is a methodology that supports the development and presentation of computer programs in a manner that promotes human, rather than computer, understanding [29]. Literate programming tools [41] take, as input, a literate program and generate both the formatted documentation of the program appropriate for a human and the list of instructions appropriate for

a computer. The Literate Assurance Approach extends literate programming techniques and tools beyond traditional programming to encompass the entire assurance argument. Literate programming and specification documentation tools can be used to present integrated formal and informal specifications and verifications in a coherent manner. Since all development products are generated from the same source this approach ensures that the documentation is consistent with the actual specification and implementation. From a certifier's perspective, this is valuable assurance evidence.

Chapter 3 Problem Definition

This chapter defines an architecture-independent concept of operations for JIFI that is used to construct a model of operations. This operational model forms the context for specifying JIFI's critical INFOSEC requirements. Section 3.1 describes the current operations of the tactical command with a focus on the Battle Group Tactical Center. Section 3.2 describes the primary problem with this operation and a framework for solving the problem. Section 3.3 describes an abstract model of JIFI operations that will form the basis for future refinement.

3.1 Current Operations

3.1.1 Tactical Command Organization

Tactical decisions are made at the highest echelon of command appropriate for the current tactical environment and level of conflict. Figure 4 provides an overview of the command relationships involved with making tactical decisions for joint force and Naval operations. During periods in which measured response is required, decisions may be made at a national command level and propagated to lower echelons. During intense wartime action or in unstable tactical environments, the authority to make decisions and develop mission plans may be delegated to lower level, joint force or theater commands [31].

Automated support for making tactical decisions is most valuable when those decisions have to be made very quickly, such as in an ongoing wartime campaign. The Unified Commander-in-Chief (USCINC) usually orders strikes against land targets and sets constraints on how the strikes are to be performed. The USCINC designates a Joint Force Commander responsible for planning and coordinating missions using the joint forces assigned to him in support of the USCINC's strike objectives. Facilities at the CINC Command Complex (CCC), including those of the Joint Task Force Center and the Joint Intelligence Center, support the decisions that have to be made by providing the USCINC and Joint Force Command with intelligence regarding enemy locations and capabilities, targeting information and imagery.

The Commander of the Naval component of the joint force (USCINCFLEET) insures the readiness of the fleet and deploys assets to designated areas. The Fleet Command Center serves as the center for gathering information concerning the composition and weapon loadout of individual Battle Groups. The USCINCFLEET works with the Joint Force Command to lay out the options for meeting strike objectives and

to make decisions on the best course of action. These decisions are organized as an Air Tasking Order (ATO) and sent to the Carrier Battle Force Commander.

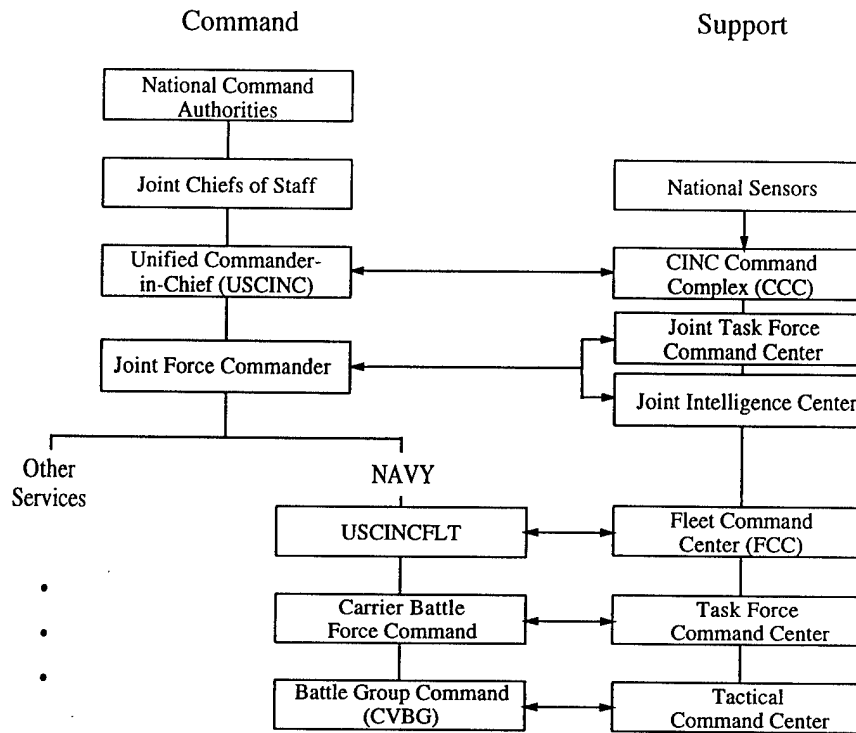


Figure 4: Tactical Command Relationships

The Task Force Command Center serves as the center for monitoring the execution of the ATO and status of the Battle Groups assigned. The Battle Group Commander performs the duties assigned to him in the Tactical Command Center (TCC). Decision support systems, such as JMCIS, underlie a Battle Group's capability to refine and carry out the Commander's orders.

3.1.2 Structure and Operation of the Battle Group

The Battle Group Commander commands the TCC under the direct authority of a Battle Force Commander or as the Naval Component Commander to a Joint Force Commander. Battle Group Commanders are generally responsible for

- exercising command and control of assigned operation forces;
- assessing and predicting tactical situations and readiness;




Figure 5 TCC Facilities' Structure and Operations

- engaging hostile forces as authorized;
- assessing hostile battle damage and redirecting assets as required and authorized; and
- providing humanitarian assistance and civilian relief as authorized and required.

A typical configuration of a Battle Group TCC facilities, which is depicted in Figure 5, includes

- portions of the Aircraft Carrier Intelligence Center (CVIC), which we call CVCIC-GENSER,
- the Combat Information Center (CIC), and
- a set of Ready Rooms ($RR_1...RR_n$)

in the GENSER enclave and

- portions of the CVIC, which we call CVCIC-SCI, and
- the Ship's Signal Exploitation Space (SSES)

in the SCI enclave. As shown, facilities in the GENSER enclave are networked together by a LAN; facilities in the SCI enclave are networked together by a LAN that is physically separate from the GENSER LAN. The only connections between the two enclaves is via serial lines interconnecting the GENSER and SCI Communication Servers, the function of which will be discussed further in the next section. The CVIC, including both CVCIC-GENSER and CVCIC-SCI, is typically protected to the SCI level, but it has workstations that operate at GENSER as well as SCI. Of course, only GENSER workstations are connected

to the GENSER LAN.¹ The CVIC and SSES facilities are bound by the requirements for Sensitive Compartmented Information Facilities (SCIF) [13].

3.1.2.1 Tactical Data Management

The GENSER and SCI JMCIS LANS each support a client-server environment that has its own Communication Server, for processing and correlating external communications to and from the CCC, and its own central data server (CDBS), for storing tactical data for access by JMCIS clients. Each facility contains a set of workstations that behave as JMCIS clients. JMCIS clients invoke application programs that track ships and ground forces in support of the development of both defensive and offensive plans, which may include air strikes, refinement of forces in consideration of the local environment, and actions to obtain more information. The Communication Server and client applications analyze military message traffic, satellite imagery, and data from other sources to develop a coherent picture of some part of the world. This picture is presented to the user as an annotated map. Users at workstations can click on map objects to get more information. Friendly platforms are colored *blue*, enemy *red*, and neutral *white*. Incoming information cannot be processed completely automatically, since data are often ambiguous: human intervention may be needed to recognize that two messages refer to the same ship by different names or to recognize a meaningful pattern in a combination of sensor reports. The resolved information and successfully parsed messages are stored in the CDBS.

Typically, before leaving port, an aircraft carrier's TCC is brought up-to-date with the current tactical picture by installing in each CDBS a base load of data (e.g., location of own/enemy forces, military asset capability/status, and other intelligence information). The SCI base load contains duplicates of all tables contained in the GENSER base load plus certain cryptologic tables that are more highly classified. Updates to the CDBS can come externally via Naval message broadcast from the CCC or locally from analysts working at a JMCIS client workstation in the TCC. The general philosophy behind CDBS updates is that existing information should not be deleted, but should be extended so that a historical record of information is maintained and available, e.g., to track the movement of forces over some period of time (see [44, page 14]). With this philosophy in mind, both external and local updates, generally, modify only Tactical Extension Tables, which are external to the base load. The only updates that modify the base load are external messages (or bulk updates) in Integrated Data Base Transaction Format (IDBTF); these updates are provided by DIA and are typically distributed to the intelligence facilities of the Fleet Command Center (either Atlantic Intelligence Center (AIC) or Joint Intelligence Center, PACific (JICPAC)).

External GENSER and SCI message traffic received from the CCC update their respective CDBSs similarly. Messages received by either enclave are parsed by that enclave's Communication Server and correlated with existing information to determine the necessary updates, if any. In addition to being

¹ A Supplemental Plot (Supplot) facility is often used to augment the SCI functions in the CVIC, but for simplicity, is not shown here.

processed by the GENSER enclave, GENSER messages received by the GENSER Communication Server are transmitted over an RS-232 serial line to the SCI Communication Server for processing and correlation by the SCI enclave. Additional SCI intelligence information used in this correlation may cause the updates to the SCI CDBS to differ from the updates to the GENSER CDBS. Transmissions over the serial line provide no acknowledgement of data received (i.e., blind write-up), resulting in the potential for lost data.

GENSER and SCI JMCIS workstations can communicate in a constrained manner over the connections between the GENSER and SCI message servers. Local updates to the GENSER CDBS by GENSER analysts do not automatically get sent to the SCI CDBS. The GENSER analyst can, however, format the update as a Naval message (e.g., using the Naval Intelligence Processing System (NIPS) [27]) and send it, via the serial line, to the SCI Communication Server for processing and correlation. Similarly, information in the SCI enclave may, on a case-by-case basis, be put in message format and sanitized for transmission to the GENSER enclave by the Esprit/Opus sanitizer, which must be controlled by a human guard. This sanitized information is automatically correlated into the GENSER CDBS if it is consistent with the GENSER tactical picture. Information that is inconsistent with the current picture is queued as an update recommendation to the GENSER operator, which he may accept or reject [5, section 1.5.3.1].

3.1.2.2 Tactical Decision Making

Tactical decision aids are most thoroughly exercised when our forces are actively engaged in wartime conditions. In times of peace, with little or no hostilities, decision aids primarily support monitoring activities. During times of heightened tensions, however, they support both monitoring and mission planning activities in the process of refining ATOs sent down by higher echelon commands. The ATO is received as a GENSER Naval message from the CCC. Refining the ATO involves many activities including target development, weaponing, asset management, plan development, plan evaluation, and report production. Most of the high level planning regarding allocation of forces and assets to meet strike objectives is performed in the CVIC-GENSER facility. Typically this planning involves coordinating with the Joint Force Command to resolve conflicts, request support and gain approval regarding forces assigned to targets. Once these assignments are made, mission planners need to determine the logistics of moving tactical assets along a route or within a theater of operations to conduct combat.

Planners and analysts in the CVIC-GENSER keep apprised of own force and enemy force movements through the GENSER communications network and, locally, through communication with the CIC. Analysts at workstations in the CIC evaluate, correlate, report upon and respond to data received by shipboard (GENSER) sensors and communications. The CIC acts as the center for tasking shipboard weapon systems and for communicating with other ships and aircraft locally. CIC analysts update the GENSER CDBS, as appropriate, making this information available to CVIC-GENSER analysts and mission planners.

Analysts working in the CVIC-SCI are responsible for merging the SCI and GENSER tactical pictures. SSES analysts consolidate SCI intelligence received locally from shipboard sensors with external sensor data received from the CCC. This intelligence, derived from signal and communication intelligence sources, includes value-added locational reports for enemy surface, air, subsurface and land platforms; and capabilities, intentions and status estimates for enemy platforms and troop movement. SSES analysts make this information available to CVIC-SCI workstations by entering it into SCI CDBS.

The GENSER tactical picture is viewed as the master because GENSER data is disseminated more broadly than SCI data [5 section 1.5.2]. Given the SCI information available to them, the CVIC-SCI analysts must validate and, when appropriate, augment the GENSER analyst's view of the tactical environment. To do this they need a complete view of the GENSER tactical picture. The GENSER and SCI CDBSs were installed with the same base load before leaving port. All subsequent GENSER updates from the CCC sent to the GENSER Communication Server were also sent to the SCI Communication Server via the serial line. Unfortunately, updates to the GENSER CDBS by GENSER analysts are not automatically sent to the SCI CDBS. These updates usually involve data concerning the movement of our own or enemy troop movements, called order of battle data, within the theater of operations.²

Order of battle data for air, missile, radar, anti-aircraft artillery and ground forces are contained in the Integrated Data Base (IDB) of CDBS [26]. Updates to the IDB by GENSER enclave analysts are stored in the Tactical Extension Tables of the GENSER CDBS. This information is currently added to the SCI CDBS by either

1. recording the information on a portable storage medium, such as magnetic tape, hand-carrying it to the SCI enclave and performing a bulk transfer;
2. manually re-entering the information at an SCI workstation; or
3. running a batch process, if the GENSER CDBS was updated automatically by that batch process.

The serial connection between the two LANs is not used for this transfer due to the potentially large volume of data involved and the slow speeds at which the (RS-232) serial connection operates. Also, use of the serial connection would require the data to be encoded in Naval message format, a cumbersome process that would only need to be undone to be entered into the SCI CDBS. The stringent time constraints common to strike planning requires near real-time access to data, which precludes the use of the existing serial line.

If SCI analysts have data that is inconsistent with the GENSER tactical picture, they may start a new SCI-only entry to the SCI CDBS or, if appropriate, they may inform the GENSER operator of the discrepancy. In the first case, key decision makers can move to the SCI enclave to analyze the data and make informed decisions. In the second case, the SCI data can be sanitized and forwarded to the GENSER LAN, or the SCI analyst can recommend a modification to GENSER data without actually passing the SCI data that suggested the need to make the change. GENSER analysts can then make more informed

decisions regarding the refinement of the ATO. Providing SCI analysts with a complete view of the GENSER tactical picture allows them to make informed decisions about the tradeoffs associated with the sanitization of the extended tactical picture for access by GENSER mission planners.

Decision makers refine an ATO into a coherent mission plan by providing guidance such as target identification and location, required levels of damage, desired routes, and critical timing constraints to experts that specialize in the use of particular weapons. These experts describe the details of how the weapons will be used to carry out strike objectives. Weapon loadout plans are developed to support scheduled missions. For example, experts on the use of fighter aircraft use TAMPS to define mission plans for strike and support aircraft. Weapon specialists, e.g., fighter pilots, become familiar with and fine-tune these plans in available Ready Rooms to prepare for execution of the plan. Unless authority for ordering strikes has been delegated to the Battle Group Command, finalizing Battle Group strike plans depends upon their review and approval by higher authorities [31, page 3-13]. Relevant portions of the plans can be sent to authorities at the CCC in the form of a Naval message. Once final approval is received, the plan may be executed as scheduled.

3.2 Required Additional Capability

Relying on personnel to manually transfer GENSER order of battle data from the GENSER CDBS to the SCI CDBS or to re-create it on the SCI CDBS is unacceptable, particularly in a real-time strike planning environment. In addition to the expense of using scarce human resources for this task, installing updates to the SCI CDBS in this way is slow and error-prone. Order of battle data describes a military organization's status or combat potential and, in times of intense action, can be extremely dynamic. Maintaining a credible, accurate and timely threat database throughout a common theater of operation is commonly recognized as the Achilles' heel of the strike planning process [1,32]. Dissimilar threat databases reduce the effectiveness of strike planning and place the success of the mission at risk.

Integrating data received by local sensors into TCC databases is vital to gaining an accurate and up-to-date representation of the battlefield, which is necessary for successful mission planning and execution. However, observations made using GENSER and SCI shipboard sensors complicates the maintenance of a common threat database across the theater of operations. Local observations, such as battle damage assessment or target relocation, must flow up to the CCC before all tactical units can be informed. Coordinating updates in this fashion is necessary to ensure proper integration of intelligence in a multi-source environment.³ Within the TCC, data received from shipboard sensors at the GENSER level

² Enemy troop movements (often called threat order of battle) are often more highly classified than own force movements due to the sensitivity of the sources of that information. This information may be classified at the SCI level and, therefore, already be available to SCI analysts.

³ Currently, a consistent tactical picture authority, e.g., Force Over-the-Horizon Track Coordinator (FOTC), is tasked with arbitrating inputs/inconsistencies and disseminating the tactical picture to other command facilities both vertically and laterally [5, section 1.5.3].

must be made available to SCI analysts in a timely manner. This problem, which is local to the TCC, is the one that this project addresses.

As shown in the goal structured graph in Figure 6, the overall objective of this project is to improve information availability in JMCIS (Goal 1.1). More specifically, within the context of the TCC shown in Figure 5, we must make updates to the GENSER CDBS available to SCI analysts (in the SCI CDBS) in a timely, secure and cost-effective manner (Goal 1.2). Since compartmented intelligence information is at risk, the Defense Intelligence Agency's (DIA) requirements (DoDIIS [14]) identify what assurance evidence is necessary. The evidence required partially depends upon the architecture chosen and so cannot at this stage of refinement be pinpointed.

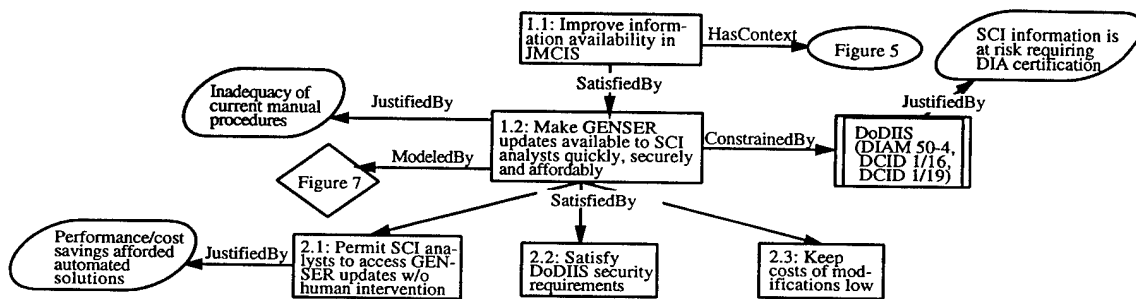


Figure 6 Top Level Requirements

In summary, an effective solution to the problem addressed will permit SCI analysts access to GENSER updates without the expensive and time-consuming human intervention currently required (Goal 2.1), will satisfy DoDIIS security requirements [7,10,12,14,17] (Goal 2.2), and will require relatively inexpensive modifications or extensions to JMCIS (Goal 2.3).

3.3 Abstract Model of Operations

This section presents an abstract model of JIFI operation as a focal point for future refinement. The Statemate specification shown in Figure 7 underlies the JIFI model. In the following textual overview of this model, names not introduced previously that appear in upper-case are primitives of Figure 7. Conceptually, JIFI includes all of the JMCIS function, including functions of both the clients and the servers. As the refinement of JIFI progresses, we increasingly focus on the data management (server) portions of JIFI since this is where JMCIS function is extended and refined. This and subsequent refinement is based on the Unified Build software architecture [43] which provides the core of JMCIS function.




Figure 7: Activity Chart Depicting JIFI Logical Architecture

Figure 7 reflects the operation and structure of JMCIS in the TCC. System inputs and outputs come from people (USERS) authorized to use JIFI and from the external communications network (COMMS_NET). People can gain access to JIFI only by logging in (SYS_CMDS). To log in, a person presents a user id and password and the system must authenticate the person as a valid user. Following successful authentication, the user can invoke operations (SYS_CMDS) to execute JIFI applications and receive results (SYS_RPT). Users monitor shipboard sensors and local communications and report findings appropriately. Messages received from the external communications network (MSG_IN) may also invoke JIFI operations based, in part, on the source of the message, which is identified in the message's header. System operations, invoked either by a received message or a valid user command, may cause messages to be transmitted (MSG_OUT) over the external communications network.

Internally, data processing (MNG_APP_CTL) and data management (MNG_DATA) functions are separated in JIFI just as they are in JMCIS. In fact, the internal activities of JIFI in Figure 7 can be related directly to the primary components of JMCIS (see Figure 2). : MNG_COMMS_CTL describes activities implemented in the JMCIS Communication Server, MNG_DB_CTL describes activities in the CDBS, and MNG_APP_CTL describes activities in the JMCIS client (e.g., JMCIS clients in Figure 5 embody the functions implemented on JMCIS workstations in CVIC-GENSER, CIC, CVIC-SCI, or SSES). A simplified view of this overall function is that users (USERS) invoke commands (SYS_CMDS) causing applications (MNG_APP_CTL) to be started. These applications may query (USR_QRY) the data manager (MNG_DB_CTL) resulting in an update of internal databases, a reply to the query (QRY_RPLY), and/or the generation of a message (BCST_MSG) to be broadcast (MNG_COMMS_CTL) over the

communications network (MSG_OUT). Results of this processing may be sent back to valid users (SYS_RPT). Messages received externally (MSG_IN) are processed by the Communication Server (MNG_COMMS_CTL) and, if appropriate, sent to the data server (MNG_DB_CTL) for correlation and update with existing data. This may result in data being sent back to a user/application or a message being broadcast over the communication network, as before.

Although Figure 7 can be viewed from a system high JMCIS perspective, it actually represents JMCIS function that has been extended to support users cleared to different levels, in our case, GENSER and SCI. This chart, therefore, represents an abstraction of a design that implements the required additional capability discussed in the last section. This is possible since it is a logical rather than a physical specification. In this multi-level view, the operations that a user may invoke, to view or modify objects, depend upon the user's clearance, the object's classification, and the roles for which the user is authorized. Notice that while Figure 7 suggests a client-server architecture, it permits many possible implementations within this scheme, e.g., use of a single trusted MLS database versus a physically distributed database in separate system high enclaves. Subsequent chapters describe the particular implementation for JIFI that we have adopted.

Chapter 4 Architecture Overview

This chapter identifies the architecture of the JIFI system based on the DoDIIS security mode of operation. Section 4.1 identifies the physical structure of the JIFI architecture and refines the abstract model of operations defined in the last chapter based on this structure. Section 4.2 outlines requirements for a device that serves as a trusted gateway between the GENSER and the SCI JMCIS LANS. Finally, Section 4.3 describes extensions to the JMCIS LAN implementations required to accommodate gateway processing within the constraints of the security requirements specified for JMCIS 2.1 [42].

4.1 Physical Model of Operations

The DoDIIS Developer's Guide [14] requires that systems processing SCI information be designed to operate in one of four modes: Dedicated, System High, Compartmented or Multi-Level. The security mode in which the system operates determines the security requirements that are mandated by DoDIIS for that system. Although the fact that we are dealing with two different security levels (not compartments) might suggest that the Multi-Level Mode is required, [17] permits a limited form of "multi-level" operation based on the interconnection of systems operating in System High Mode:

"In the near-term, the DoDIIS security architecture will focus on system high client-server operations with trusted interfaces to environments operating at different security levels. Trusted interfaces will be trusted host or workstation-based platforms (with network security enhancements) that provide electronic connections between the Top Secret SCI DOD Intelligence Community and systems operating in other security environments (e.g., collateral classified, law enforcement)."

Requirements for interfaces that span security boundaries are based on the particular application:

"The Accreditation Authority (along with cognizant Site ISSO) shall be responsible for approving the use of trusted interfaces to support mission and security needs."

The choice of security mode of operation is seen in the refinement of our goal structured graph in Figure 8. As shown by Goals 2.4 and 2.5, we choose to follow DoDIIS guidance extending JMCIS in a way that satisfies requirements for System High Mode operation and other application-specific requirements for a trusted interface. The trusted interface will be built as a gateway linking the JMCIS GENSER CDBS and SCI CDBS (Goal 3.1). Adopting this approach permits reusing the previous certification and accreditation of the JMCIS system high LANs, e.g., according to [3] and to avoid the cost of satisfying the more stringent

Multi-Level Mode requirements. There will, nevertheless, be some required modification of the data management portions of JMCIS that must be shown not to invalidate the previous certification (Goal 4.1).

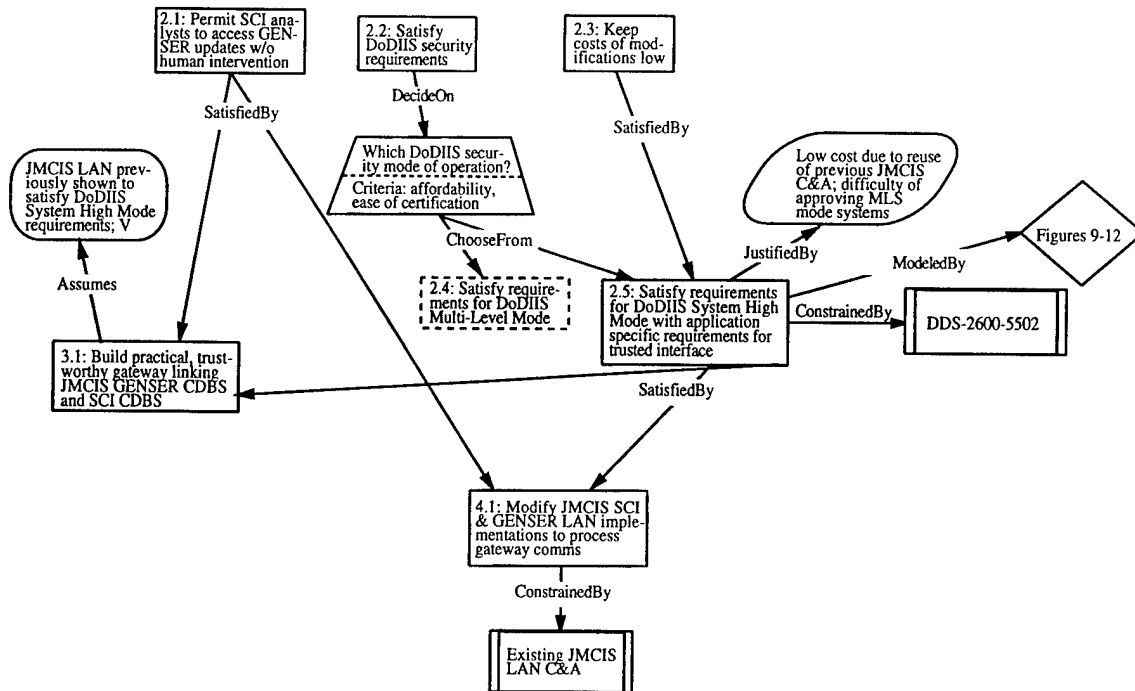


Figure 8: Identifying the DoDIIS Security Mode of Operation

Physical separation of the GENSER and SCI security domains is crucial to their operation in the System High Mode of operation. Figure 9 depicts the physical separation of JIFI function into two system high enclaves: L_ENCLAVE and H_ENCLAVE. Our intention is to preserve the existing distribution of GENSER/SCI data and processing in the JMCIS architecture as depicted in Figure 5. In the Statestate specification and henceforth in the textual descriptions, we refer to GENSER as Low and SCI as High, for generality. Thus, L_ENCLAVE will be the home of the GENSER JMCIS LAN and H_ENCLAVE will be the home of the SCI JMCIS LAN. The only link (information flow) between the two enclaves is the information flow permitted by the GATEWAY_MC module.⁴

⁴ The serial and sanitizer lines depicted in Figure 5 are not represented here. We expect that the serial line will eventually be phased out of the JMCIS configuration. The sanitizer line has no impact on our planned extensions.

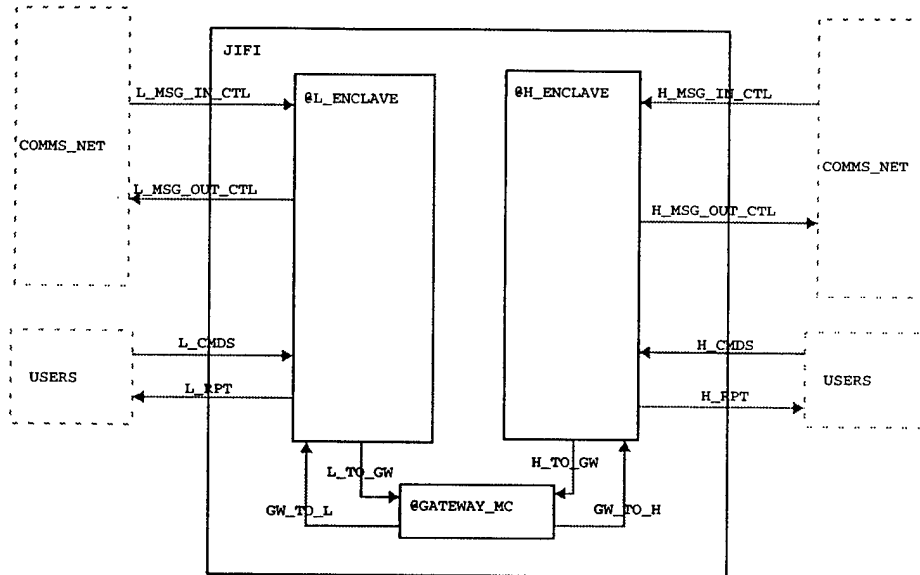


Figure 9: Module Chart Depicting JIFI Physical Architecture

The functional view of JIFI shown in Figure 7 must correspond to the physical view shown in Figure 9. At this high level, this correspondence is seen in the decomposition of the external information flows in Figure 7. Each information flow between JIFI and an external entity is split into two components, one representing the flow to/from L_ENCLAVE and one to/from H_ENCLAVE. For example, the information flow MSG_IN of Figure 7 consists of two components, L_MSG_IN_CTL and H_MSG_IN_CTL, in Figure 9. MSG_OUT, SYS_CMDS, and SYS_RPT are similarly decomposed. Intuitively, the USERS that can enter H_CMDS and receive H_RPT must be cleared for High, whereas USERS that can enter L_CMDS and receive L_RPT need only Low clearance. Likewise, the portion of COMMS_NET responsible for High traffic must be protected to High, whereas the portion responsible for Low traffic need only be protected to Low.

One level decomposition of the primitive activities in Figure 7 allows us to start mapping JIFI activities to JIFI modules. This mapping permits us to strengthen the correspondence between the functional and the physical views. Figure 10, Figure 11, and Figure 12 present the first level decomposition of the activities MNG_COMMS_CTL, MNG_DB_CTL and MNG_APP_CTL, respectively. The goal in this decomposition is to separate the JMCIS functions (communications management, database management, and application management) into their Low and High counterparts. For example, MNG_COMMS_CTL is partitioned into High communications processing, H<MNG_COMMS, and Low communications processing, L<MNG_COMMS. These activity charts decompose the commands sent and reports received by USERS into System Administrator (SA), Database Administrator (DB) and Operator (OP) classes. For example, H_CMDS is split up into H_SA_CMDS, H_DB_CMDS, and H_OP_CMDS.

Communications between MNG_COMMS_CTL and MNG_DB_CTL, and between MNG_DB_CTL and MNG_APP_CTL, are decomposed into their Low and High counterparts, as before.




Figure 10: Activity Chart Depicting Management of JIFI Communications

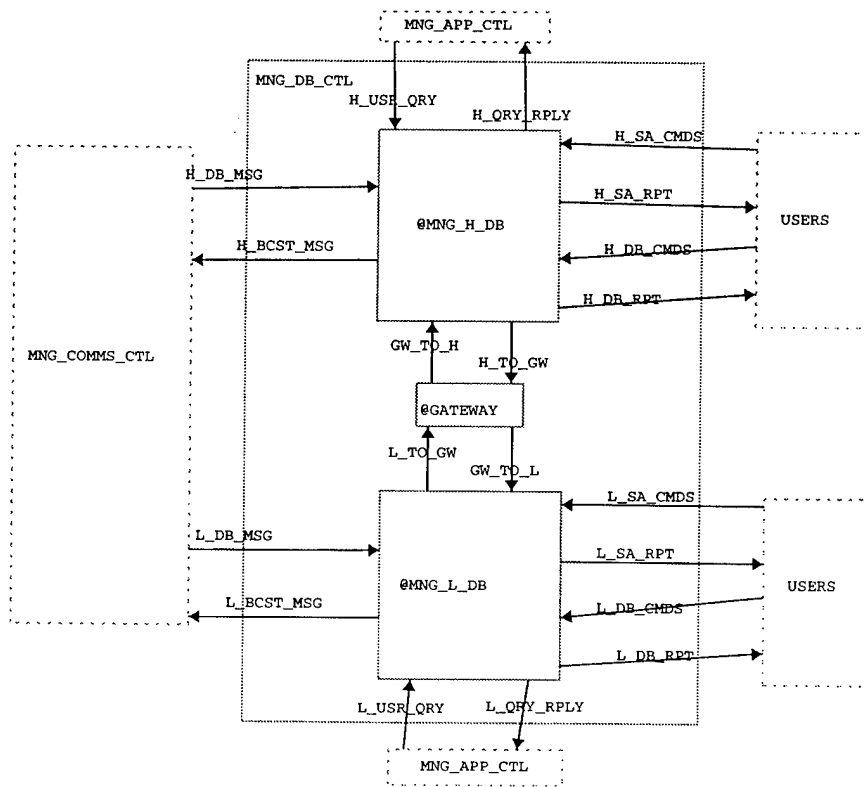


Figure 11: Activity Chart Depicting Management of JIFI Databases




Figure 12: Activity Chart Depicting Management of JIFI Applications

The mapping of activities to modules at the current level of decomposition is, for the most part, straightforward. Activities associated with the processing of Low data are mapped to L_ENCLAVE and those associated with High data are mapped to H_ENCLAVE. The GATEWAY activity of Figure 11 maps to the GATEWAY_MC module of Figure 9. GATEWAY_MC receives Low database transaction updates from MNG_L_DB via L_TO_GW and forwards these to the High database, MNG_H_DB, via GW_TO_H. The status of these updates in the High enclave is relayed to GATEWAY_MC via the H_TO_GW flow. In the context of Figure 5, GATEWAY_MC replicates update transactions of the GENSER CDBS to the SCI CDBS. At this point in the decomposition, the mapping of activities to the modules in which they are implemented is really only a grouping of functions into the two enclaves. As the details of the logical and physical architecture are refined in lower level specifications, this mapping will be made much more precise.

4.2 JIFI Gateway Requirements

As shown in the goal structured graph of Figure 13, the first step in deciding how to refine Goal 3.1 is to determine how the trusted gateway is going to be used. In general, we believe that information should be stored at its appropriate (lowest) security level while still providing near real-time and cost-effective access of Low information to High users. Therefore, we could either (Goal 3.2) securely replicate Low enclave SQL updates to the High CDBS or (Goal 3.3) allow High enclave users to query Low information securely, as needed. Since JMCIS High users, for the most part, know in advance what classes of information they need to perform their jobs, we choose the first option. This assumes that the Low side is the sole authority in the TCC for modifying Low information.




Figure 13: JIFI Gateway Requirements

Previous experience has shown that gaining high assurance of the security properties of a system in a cost-effective manner, as required for JIFI, requires isolating the security-critical function in small, reusable components (Goal 3.4). This function must protect High information from access by users only cleared to Low (Goal 5.1) while still providing reliable communication of update transactions from Low to High. This is possible by requiring that the security-critical component mediate all communications between Low and High. The security-critical function must be protected from unauthorized tampering, since this could violate the assumptions on which it is based. The component that we will use to isolate the security-critical function of JIFI is called the Pump, the design for which is specified in [28]. The next chapter describes the assurance requirements for the Pump based on the GATEWAY activity originally specified in Figure 11 and refined in Figure 14.

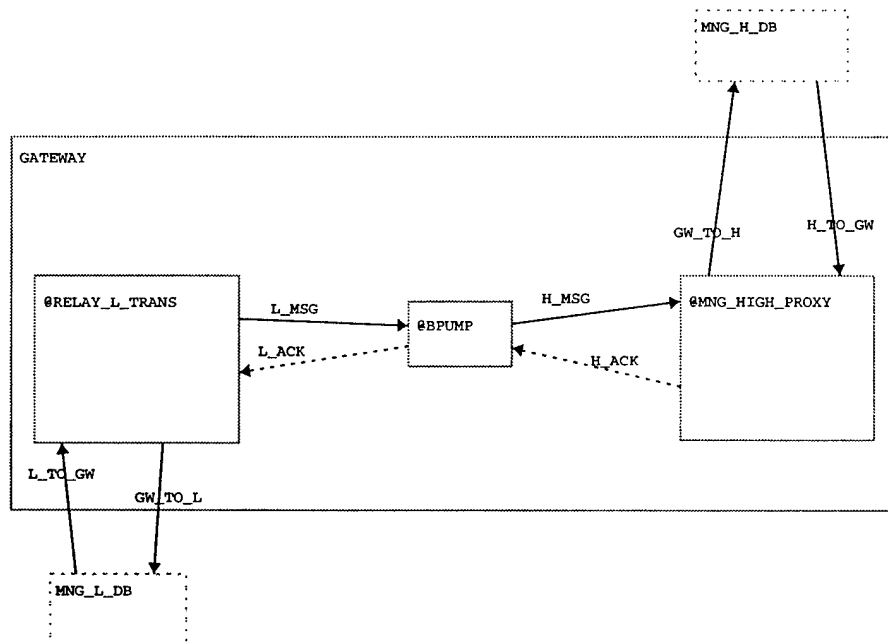


Figure 14: Activity Chart Depicting the JIFI Gateway

We use Sybase replication technology to reliably and quickly replicate Low CDBS update transactions to the High CDBS (Goal 3.5). As shown in Figure 15, replication is performed by two primary processes: the Log Transfer Manager (LTM) and the Replication Server (RS).⁵ The responsibilities of the LTM include reading the transactions from the transaction log of the primary database server (in our case, the Low CDBS) and sending them to RS using the Log Transfer Language (LTL). Subscriptions of tables that must be replicated are stored on the primary database server in the Replication Server System Database (RSSD). The responsibilities of RS include storing those transactions that update subscribed tables in stable storage and sending them to the replicate database server (in our case, High CDBS).

⁵ The special symbols representing the components of the Sybase replication technology in the figure were invented by Sybase and will not necessarily be familiar to readers not acquainted with their documentation.

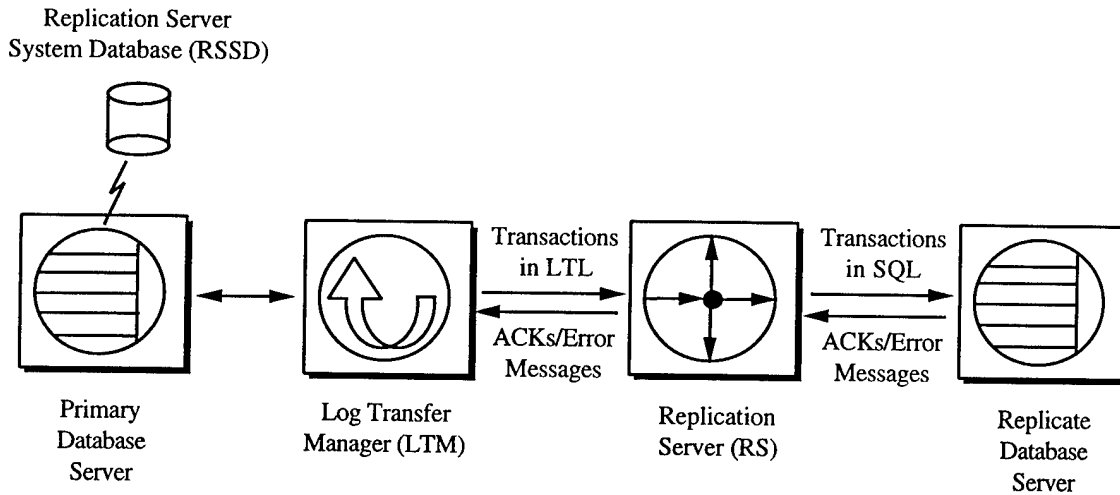


Figure 15: Sybase Replication Architecture Information Flow

Using Sybase replication technology to connect databases at different levels requires deciding where to insert the security critical function in the replication architecture. Figure 15 suggests two logical places: between LTM and RS (Goal 3.6) or between RS and High CDBS (Goal 3.7). The second of these options is chosen since the communication protocol between RS and High CDBS (which is based on SQL) is much better documented and less likely to change than the protocol between LTM and RS (which is based on LTL). The implementation for the resulting architecture must be reliable, in the sense that no update transactions being replicated may be lost or duplicated and no spurious transactions may be created (Goal 7.1). The implementation must also perform better than manual database update (Goal 8.1).

4.3 JMCIS LAN Extensions

As shown in Figure 10, the JIFI Gateway is connected to the Low database by the L_TO_GW and GW_TO_L channels. Similarly, the Gateway is connected to the High database by the H_TO_GW and GW_TO_H channels. The goal structured graph in Figure 16 shows that modifying the existing JMCIS LAN implementations (Goal 4.1) involves recording Low CDBS transactions to the transaction log (Goal 4.2), which flow over L_TO_GW, and incorporating into the High CDBS transactions received over the GW_TO_H channel (Goal 4.3).




Figure 16: JMCIS LAN Extensions

Satisfying Goal 4.2 should require little, if any modification to the Low enclave processing. Verification requires only inspection that updates to relevant tables are logged. Goal 4.3, on the other hand, requires creating separate tables within the High CDBS to store the Low replicate tables (Goal 4.4). These replicate tables must be read only by the High side to reduce the chance that update transactions for Low (sent through the gateway) cause errors. Such errors would require that replication to the High CDBS be discontinued until the problem is fixed manually, a time-consuming and expensive proposition at best. Recall, however, that initially the High base load contains duplicates of all tables contained in the Low base load. These duplicates are writable and thus diverge from the Low replicated tables as they are updated by High analysts. This imposes a requirement to mediate High CDBS queries to resolve differences between the two versions of the tables (Goal 4.5). Testing will be used to ensure the consistency of Low tables with the High replicates, as well as, the proper resolution of disparities between these tables in responding to queries on the High side.

Chapter 5 Gateway Assurance Strategy

This chapter presents an overview of the strategy for gaining assurance that the JIFI Gateway satisfies its security, reliability and performance requirements. In the last chapter, we described how the architecture isolates the security-critical function in a device called the Pump. Section 5.1 describes how confidentiality is achieved in the design and implementation of the Pump. Section 5.2 describes how the gateway architecture and system management policy protect against physical tampering with the Pump. Finally, Section 5.3 and Section 5.4 describe how we will gain confidence in the reliability and good performance of operation, respectively, of the Gateway implementation.

5.1 Assuring Confidentiality

The most difficult part of separating information of different security levels in JIFI is accomplished by preserving the separation inherent in the physical distribution of the Low and High enclaves. The addition of the Gateway connecting these enclaves introduces the potential for unsecure flows from High to Low. The capacity of such flows must be reduced to an acceptable degree while still ensuring the good reliability and performance needed to maintain a common tactical picture. This is the objective of developing a component, the Pump, that isolates the security-critical function to protect High information from access by Low users as shown in Goal 5.1 of the goal structured graph in Figure 17.

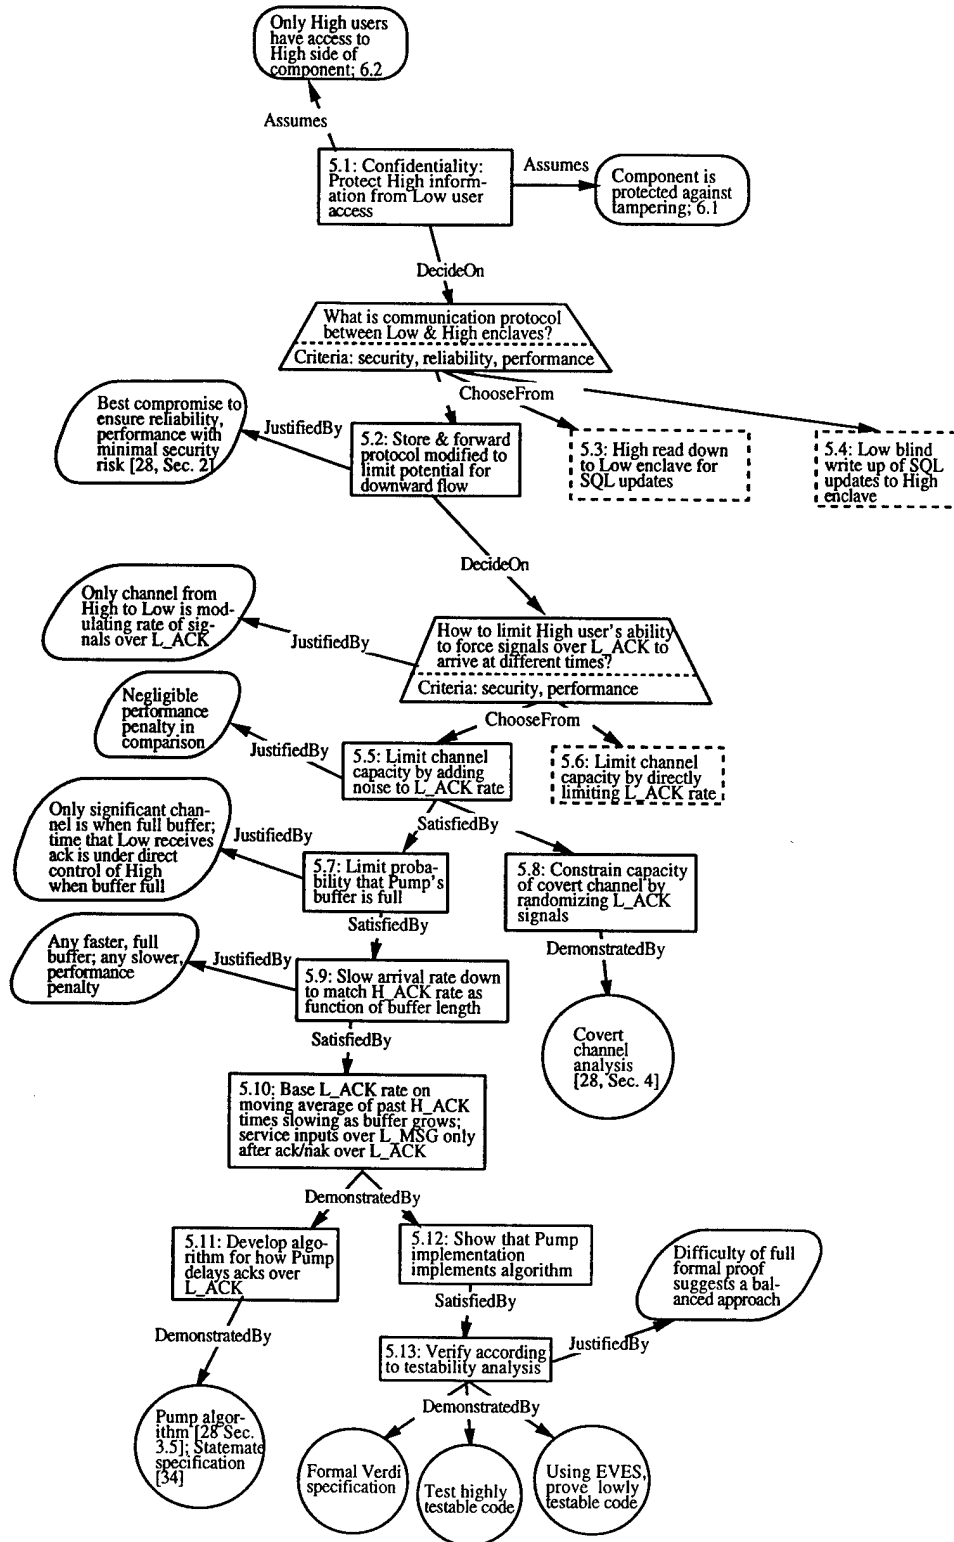


Figure 17: Assuring Confidentiality

The Pump's responsibility for ensuring confidentiality is limited to mandatory security. Kang and Moskowitz [28] analyze two inherently secure communication protocols between different security levels: a read down protocol (Goal 5.3) and a blind write up protocol (Goal 5.4). Their analysis shows that the read down protocol sacrifices performance for security and the blind write up protocol sacrifices reliability for security. They show that these sacrifices are unnecessary by describing the design for the Pump that is based on a store and forward protocol. The protocol is modified to limit the potential for downward information flow to an acceptable capacity, which is set on an application-specific basis when the device is configured (Goal 5.2). Although the Pump does not permit any direct communication from High to Low, information can flow using covert timing channels. Assurance that the Pump's design, set forth in [28], constrains the covert channels adequately without sacrificing reliability or performance is largely based on the argument made in that paper. This section outlines that argument and extends it to include a strategy for showing that the Pump implementation is a proper refinement of the Pump design. We use Figure 14 as the basis for discussion.

The timing channel through the Pump occurs when a High user forces signals over L_ACK to arrive to a Low user at different times. The modulation of L_ACK signals can be used to encode High information. Two options for limiting a High user's ability to exploit this channel are to add random noise to the L_ACK rate (Goal 5.5) and to directly limit the L_ACK rate (Goal 5.6). The second of these options could limit the covert channel to capacities acceptable, e.g., by the TCSEC [36], but can severely reduce system performance if communication is required and possible at faster rates. The first option limits the capacity of the channel with negligible performance impact (see Section 5.4 for requirements on performance analysis).

A conventional store and forward protocol that is used to allow a user to reliably transmit messages to a more highly cleared user suffers from a significant covert channel when the buffer is full. As described in [28], when the buffer is full, the time that a Low user receives a signal over L_ACK is under direct control of High. The Pump must therefore limit the probability that its buffer is full (Goal 5.7). The Pump should not completely prevent its buffer from ever becoming full, since this would allow the High user to signal information when the buffer is not full - a situation worse than the one with which we started. The strategy of avoiding a full buffer reduces the capacity of, but does not eliminate, the covert channel. The capacity of the resulting channel is constrained further by randomizing the rate of signals over L_ACK (Goal 5.8). The capacity of the resulting channel was analyzed in Section 4 of [28].

The approach to avoiding a full buffer is to slow the arrival of messages over L_MSG down to match the rate of acknowledgements over H_ACK as a function of the buffer length (Goal 5.9). In particular, the Pump buffers incoming transactions over L_MSG and bases the rate of acknowledgements over L_ACK on a moving average of past H_ACK times (Goal 5.10). Since the Pump services inputs over L_MSG only after acknowledgements over L_ACK, this slows down the arrival rate to match the H_ACK rate. Thus, as the High side slows down (or speeds up), so does the average rate of consumption of Low

transactions by the Pump. To further reduce the chance that the buffer is full, the rate of acknowledgements sent over L_ACK slows as the buffer gets larger.

Kang and Moskowitz [28] specify an algorithm that captures the constraints on the L_ACK rate as described above (Goal 5.11). The Pump design based on this algorithm has also been specified in Statemate [34]. The assurance strategy requires showing that the Pump implementation conforms to this algorithm (Goal 5.12). A technique known as testability analysis [45] will be used to determine the parts of the software implementation for which testing alone will provide effective verification and those parts that require additional analysis (Goal 5.13). We will analyze the lowly testable code using the EVES verification system and its associated Verdi specification language [4,30].

5.2 Protecting Against Tampering

The approach to protecting the security-critical function from tampering (Goal 6.1 of Figure 18) is to place the Pump in the SCI Facility (SCIF) that houses the JMCIS SCI LAN (Goal 6.2). This approach is effective because of the physical, administrative and personnel security requirements imposed on JIFI (Goals 6.3, 6.4, and 6.5). Some of these requirements, which are specified in this section, will already be enforced in the existing SCI JMCIS LAN; others will be new additions imposed due to the introduction of the Gateway. These requirements will extend the existing JMCIS Facilities Manual and thus, will not be refined further.

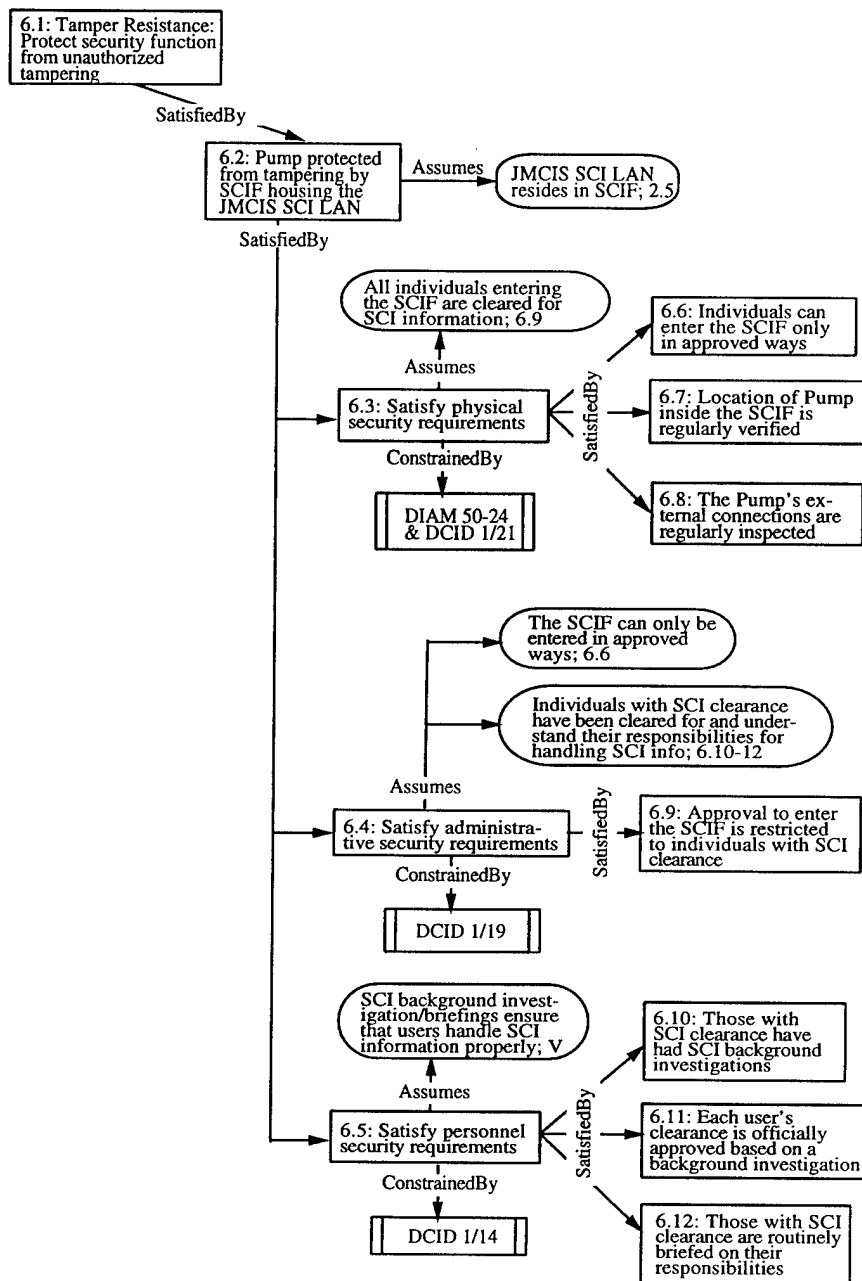


Figure 18: Protecting Against Tampering

Physical security, which for a SCIF is regulated by DIA Manual 50-24 [13] and DCID 1/21 [11] must ensure that individuals are able to enter the SCIF only in approved ways (Goal 6.6), e.g., through the vault doors. The location of the Pump inside the SCIF (Goal 6.7) and the Pump's proper connection to the JIFI Gateway (Goal 6.8) must be routinely inspected and verified. Physical security assumes that all individuals entering the SCIF are cleared for SCI information.

Administrative Security, which for SCI information is regulated by DCID 1/19 [10] requires assuring that which physical security assumes: approval to enter the SCIF is restricted to individuals with

SCI clearance. Administrative Security assumes that SCIFs can only be entered in approved ways and that individuals with SCI clearance have been cleared for and understand their responsibilities for handling SCI information.

The second assumption of Administrative Security is enforced through Personnel Security countermeasures. Personnel Security, which is regulated by DCID 1/14 [6], requires that those with SCI clearance have had SCI background investigations (Goal 6.10), that the SCI clearance of each JIFI user is officially approved based on that investigation (Goal 6.11) and that those with SCI clearance are routinely briefed on their responsibilities for protecting SCI information and the equipment that it processes (Goal 6.12). This is based on the fundamental assumption that an SCI background investigation and routine briefing on responsibilities ensure that users handle SCI information properly.

5.3 Assuring Reliability

We say that a system or component provides reliable communication if there is no loss, duplication nor spurious creation of messages that it is responsible for transmitting. As shown in the goal structured graph in Figure 19, if the Pump is reliable (Goal 9.1) and the Sybase replication technology is reliable (Goal 10.1) then the Gateway is reliable (Goal 7.1).

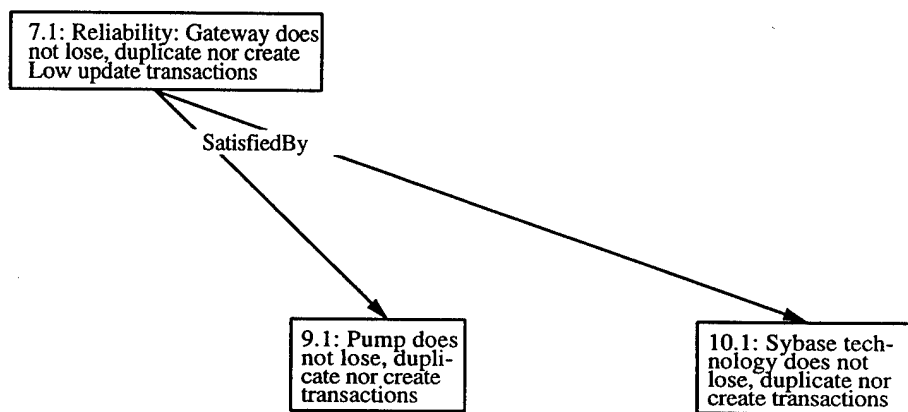


Figure 19: Assuring Reliability

5.3.1 Reliability of the Pump

Goal 9.1 of the goal structured graph in Figure 20 requires that the Pump not lose nor duplicate transactions sent to it by Low, nor create any spurious transactions of its own. Of course, there is a possibility that transactions received from Low are not acknowledged by the Pump, for example due to a full buffer, or that transactions sent by the Pump are not acknowledged by High, for example due to some failure of High. The transactions must be retransmitted by Low, in the first case, and by the Pump, in the second case. In the following decomposition of this goal, we assume that transactions sent to the Pump are uniquely identifiable, e.g., by assigning sequence numbers.

In the context of Figure 14, Goal 9.1 is satisfied if those transactions received over L_MSG are eventually sent over H_MSG in the same order received disregarding the duplicates retransmitted due to lack of acknowledgement (Goal 9.2). Only those transactions received may be transmitted; no spurious transactions may be introduced. Furthermore, all transactions over L_MSG must be acknowledged exactly once over L_ACK within τ_L time units of its last transmission; likewise, all transactions over H_MSG must be acknowledged exactly once over H_ACK within τ_H time units of its last transmission (Goal 9.3). The constants τ_L and τ_H are parameters set at configuration time. Finally, Pump operation must be automatically recoverable from system failure (Goal 9.4); media failure is beyond the scope of our analysis. The recoverability of the Pump design was analyzed in Section 3.2 of [28].

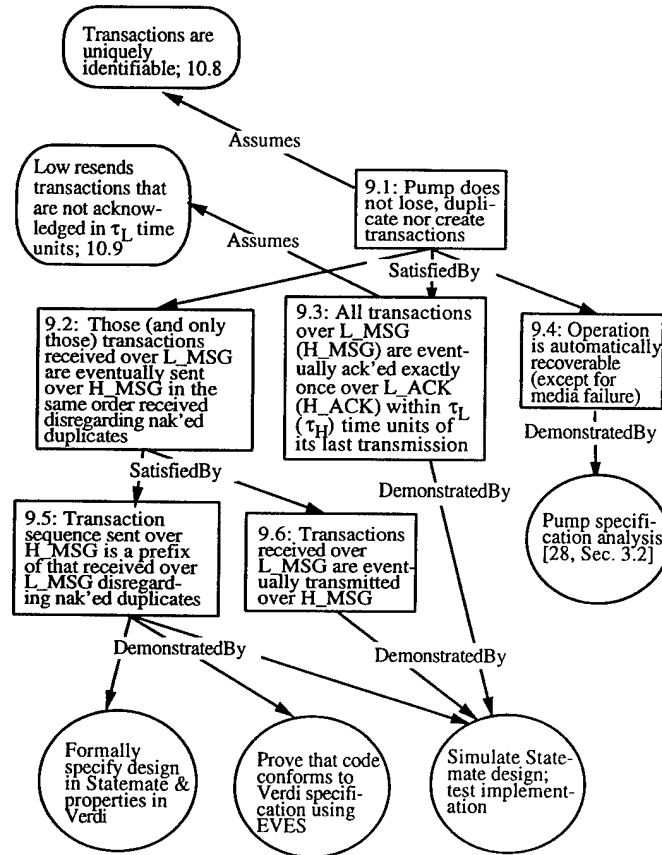


Figure 20: Reliability of the Pump

The decomposition of Goal 9.2 requires viewing the transactions that flow over L_MSG and H_MSG as sequences. Goal 9.5 requires that the sequence of transactions sent over H_MSG be a prefix of that received over L_MSG, when you ignore the duplication of messages that occurs due to lack of acknowledgment. Goal 9.6 requires that progress be made in the Pump's transmission of messages received over L_MSG. Goal 9.5 will be demonstrated by proving that the code conforms to its Verdi specification

using EVES, by simulating the Statemate design [34], and by testing the implementation. Since formal proof of liveness properties is much more difficult and expensive than for safety properties, Goals 9.3 and 9.6 will be demonstrated only through simulation and testing.

5.3.2 Reliability of Sybase Replication Technology

The reliability of Sybase replication technology requires the reliability of the individual components of the replication architecture as previously described in Figure 15. This is the objective of Goal 10.1 in the goal structured graph in Figure 21. The components of the Sybase replication architecture have uniform interfaces defined in the Sybase Open Server and Open Client protocols. When RS sends a transaction to a replicate database, it is sending information as an Open Client and expects the proper data exchange from an Open Server, the replicate data server. This exchange signals either a successful or unsuccessful completion of the request from the server. Likewise, when the replicate data server receives a transaction, it assumes the transaction is from an Open Client and sends its response in a format appropriate to a client.

Since we do not want to modify the Sybase software, interrupting the communication protocol between RS and High CDBS in the standard configuration requires introducing a wrapper to the Pump. This wrapper provides the responses to inputs that would be required if RS and High CDBS were directly connected. In particular, the Low part of the wrapper serves as a proxy for the High CDBS to RS (Goal 10.2) and, thus, must return responses to RS in Open Server data exchange format (Goal 10.7). In addition, the Low wrapper is responsible for assigning sequence numbers to individual messages sent to the Pump (Goal 10.8) and for resending messages that are not acknowledged in τ_L time units (Goal 10.9). The Low proxy is represented by the activity MNG_LOW_PROXY in the refinement of RELAY_L_TRANS shown in Figure 22.




Figure 22: Activity Chart Depicting Low Update Transactions to the Pump




Figure 23: Activity Chart Depicting Replication of a Database

Similarly, the High part of the wrapper serves as a proxy for the RS to High CDBS (Goal 10.3) and must relay transactions received from the Pump and respond to High CDBS in Open Client data exchange format (Goal 10.10). Transactions received from the Pump are stored in stable storage, when room is available and receipt acknowledged (Goal 10.11). The High proxy is represented by the activity MNG_HIGH_PROXY in Figure 14. Proper operation of these proxies will be verified through testing to ensure proper responses are supplied. This, of course, assumes that they are properly connected to the Pump.

The rest of the Gateway consists of the components of the conventional Sybase replication architecture, as shown in Figure 24. The RSSD is the database in the Low CDBS that is used to store the subscriptions to the tables being replicated (Goal 10.4). The LTM scans the transaction logs of Low CDBS and sends updates of tables marked for replication to RS (Goal 10.5). RS stores transactions for subscribed tables in stable storage and forwards them to Low Proxy when ready for input (Goal 10.6). The XFER_LOG and MNG_LTM activities of Figure 23 represent the LTM; the MNG_DB_RS activity represents the RS. Logically, there is one of these Activity Charts, i.e., MNG_DB_REP, for each database being replicated. Figure 22 shows the replication of two databases, DB1 and DB2, which might represent databases containing order of battle data such as IDB. Testing will demonstrate that these components reliably transfer to High CDBS the appropriate transactions applied to the Low CDBS.




Figure 24: Secure Replication

5.4 Assuring Good Performance

A high-level breakdown of the requirement that the Gateway improves performance over manual database update (Goal 8.1) is shown in the goal structured graph of Figure 25. The sufficient, but not

necessary, condition to achieve this are to ensure that the slow down of communication traffic due to the Pump as compared to a store and forward buffer is negligible (Goal 8.2) and that the Gateway Low update transactions available to High users in near real-time (Goal 8.3). The first of these goals was demonstrated by a performance simulation of the Pump design in Section 5 of [28]. Both goals will be demonstrated through analysis of the final implementation.




Figure 25: Assuring Good Performance

Chapter 6 Residual Risk

An overview of the complete goal structured graph representing the assurance strategy is given in the fold-out at the end of this document. Our assessment of residual risk assumes that the goals of this graph are satisfied and, where appropriate, demonstrated as required. The residual risk, therefore, arises from assumptions that are not validated by a goal, or set of goals, elsewhere in the graph. Assumptions that are so validated are signified by the list of validating goal numbers specified at the end of the assumption description. Goals that are not validated are vulnerabilities and are signified by the letter "V" at the end of the assumption description.

There are two assumptions that are not validated:

- JMCIS LAN previously shown to satisfy DoDIIS System High Mode requirements; and
- SCI background investigation ensures that users handle SCI information properly.







These assumptions represent the risk that remains after the derived requirements are satisfied.

Acknowledgments




This work benefited greatly from the contributions of Eather Chapman, Judy Froscher, David Goldschlag, Richard Hale, Myong Kang, Carl Landwehr, John McDermott and Ira Moskowitz from NRL; Mary Rock and Rodney Peyton from Kaman Sciences Corporation; and George Thomas from Sierra Cybernetics.

Appendix A. Notation




Activity Charts - the functional view

-  - an activity named X
-  - an activity X that is refined in a lower level activity chart also named X
-  - an instance Y of a generic (parameterized) activity X
-  - an activity Y with controlling state chart X.
-  - an activity X that is external to the chart being elaborated
-  - an activity X with a lower level description (mini-spec) of its behavior


Module Charts - The Physical View

-  - a module named X
-  - a module X that is refined in a lower level module chart also named X
-  - a module X that is external to the chart being elaborated

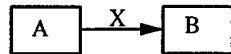
State Charts - The Behavioral View

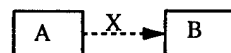
-  - a state named X
-  - a state X that is refined in a lower level state chart also named X
-  - a state X with a lower level description (static reactions) of its behavior

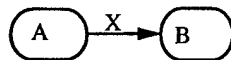
Data Stores










 - a place to store data item or control element X

Flows and Transitions

 - a flow of data item X from activity (module, or data store) A to activity (module, or data store) B

 - a flow of control element X signaled from activity (module, or data store) A to activity (module, or data store) B

 - a trigger X that causes transition from state A to state B

		SOURCE								
		Goal	Assumption	Justification	Choice	Goal Option Not Chosen	Solution	Model	Context	Constraint
										
D E S T I N A T I O N	Goal	SatisfiedBy Requires DemonstratedBy			ChooseFrom	SatisfiedBy Requires DemonstratedBy				
	Assumption	Assumes		Assumes	Assumes	Assumes	Assumes	Assumes	Assumes	Assumes
	Justification	JustifiedBy			JustifiedBy	JustifiedBy	JustifiedBy	JustifiedBy	JustifiedBy	JustifiedBy
	Choice	DecideUpon			ChooseFrom	DecideUpon				
	Solution	SolvedBy DemonstratedBy			ChooseFrom	SolvedBy DemonstratedBy				
	Model	ModeledBy			ChooseFrom	ModeledBy	ModeledBy		ModeledBy	ModeledBy
	Context	HasContext	HasContext	HasContext	HasContext	HasContext	HasContext	HasContext		HasContext
	Constraint	ConstrainedBy	ConstrainedBy	ConstrainedBy	ConstrainedBy	ConstrainedBy	ConstrainedBy	ConstrainedBy	ConstrainedBy	

Elements of the Assurance Strategy Notation and their Relationships

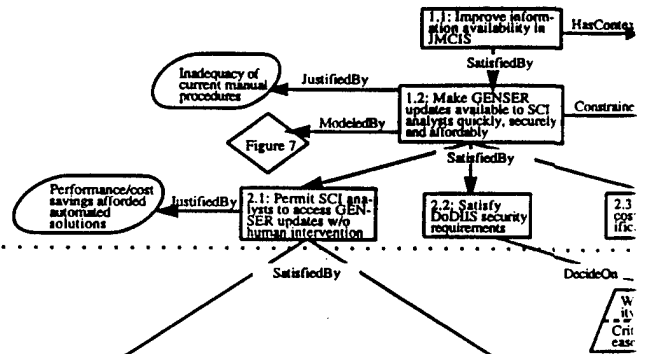
Bibliography

1. David Bulford. *TAMPS Data Base Study*. Space and Naval Warfare Systems Command, August 1993.
2. Naval Air Warfare Center. Tactical Aircraft Mission Planning System (TAMPS) Version 6.0: Concept of Operations. Technical report, Naval Air Warfare Center, July 1995.
3. The Mitre Corporation. *DDS-2600-5502-87, Security Requirements for System High and Compartmented Mode Workstations*, November 1987.
4. Dan Craigen, Sentot Kromodimoeljo, Irwin Meisels, Bill Pase, and Mark Saaltink. Reference Manual for the Language Verdi. Technical Report TR-91-5429-09a, ORA Canada, Ottawa, Ontario, September 1991.
5. Stanley Davis. JMCIS Database Concept of Operations. Technical report, Delphin Systems, 1995.
6. Defense Intelligence Agency (DIA). *DIA Manual 50-3, Physical Security Standards for Construction of Sensitive Compartmented Information Facilities..*
7. Defense Intelligence Agency (DIA). *DIA Manual 50-4, Security of Compartmented Computer Operations (C)*, June 1980.
8. Defense Intelligence Agency (DIA). *DIA Manual 50-5, Sensitive Compartmented Information (SCI) Contractor Administrative Security*, June 1980.
9. Director, Central Intelligence. *DCID-1/14 Minimum Personnel Security Standards and Procedures Governing Eligibility for Access to Sensitive Compartmented Information*, February 1987.
10. Director, Central Intelligence. *DCID-1/19 Security Policy for Sensitive Compartmented Information*, February 1987.
11. Director, Central Intelligence. *DCID-1/21 U.S. Intelligence Community Physical Standards for Sensitive Compartmented Information*, September 1987.
12. Director, Central Intelligence. *DCID-1/16 Security Policy for Uniform Protection of Intelligence Processed in Automated Information Systems and Networks (S)*, July 1988.
13. Defense Intelligence Agency (DIA). *DIA Manual 50-24, Security Policy for Using Communications Equipment in a {Sensitive Compartmented Information Facility (SCIF)*, August 1990.
14. Defense Intelligence Agency (DIA). *DODIIS Developer's Guide for Automated Information Systems Security in DOD Intelligence Information Systems*, November 1993.
15. Defense Intelligence Agency (DIA). *DODIIS Site Certifier's Guide*, November 1993.

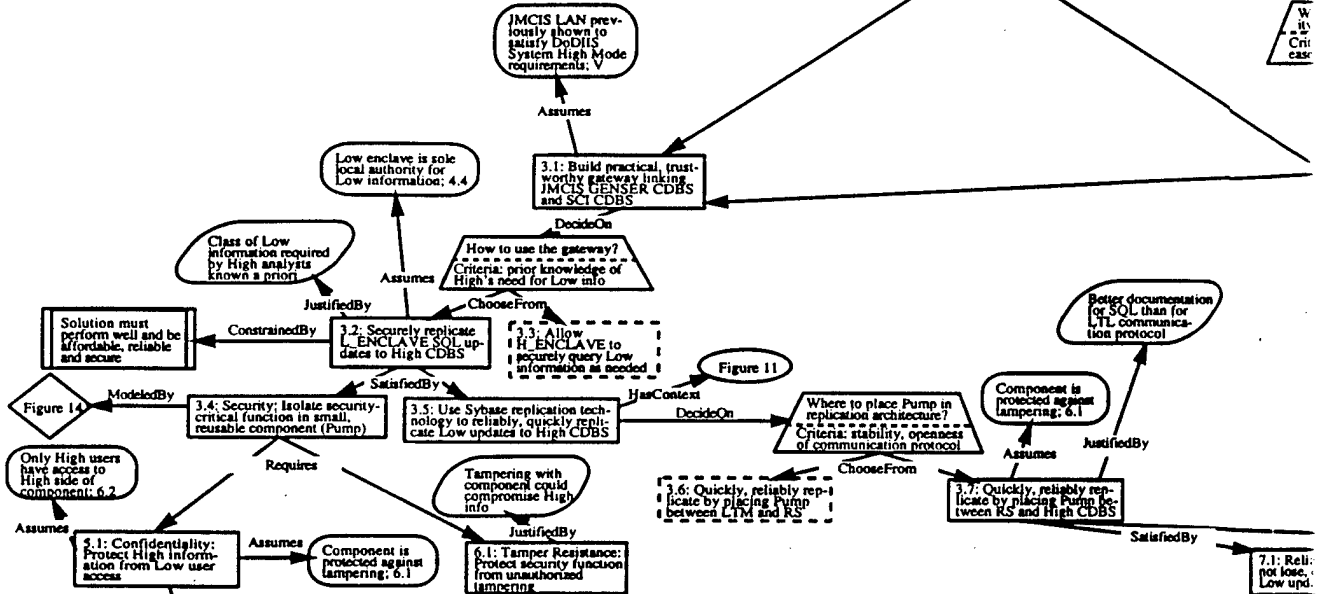
16. Defense Intelligence Agency (DIA). *DODIIS Site Information Systems Security Officer's Handbook*, November 1993.
17. Defense Intelligence Agency (DIA). *DODIIS Security Architecture Guidance and Directions*, September 1994.
18. Judith N. Froscher, Myong Kang, John McDermott, Oliver Costich, and Carl Landwehr. A Practical Approach to High Assurance Multilevel Secure Computing Service. In *Proc. 10th Annual Computer Security Applications Conference*, December 1994.
19. David Harel, Amir Pnueli, J. Schmidt, and Rivi Sherman. The Formal Semantics of Statecharts (Extended Abstract). In *Proc. 2nd IEEE Symposium on Logic in Computer Science*, pages 54-64, Ithaca, NY, 1987.
20. David Harel, Hagi Lachover, Amnon Naamad, Amir Pnueli, Michal Politi, Rivi Sherman, Aharon {Shtull-Trauring}, and Mark Trakhtenbrot. StateMate: A Working Environment for the Development of Complex Reactive Systems. *IEEE Transactions on Software Engineering*, 16(4):403--414, April 1990.
21. i-Logix Inc. *The Languages of StateMate*, January 1991.
22. i-Logix Inc. *StateMate 5.0 Features*, June 1993.
23. i-Logix Inc. *StateMate 6.0 Features*, May 1995.
24. Inter-National Research Institute Inc. Joint Maritime Command Information System (JMCIS) Common Operating Environment (COE). Technical Report SPAWARSYSCOM PD 60E, Space and Naval Warfare Systems Command, February 1994.
25. PRC Inc. Copernicus Architecture Functional Overview. Technical report, Naval Research Laboratory, April 1994.
26. PRC Inc. Navy Tactical Command Systems - Afloat: Client Database Document (CBD)) for Version 2.1. Technical Report PRC R-5305, Naval Command, Control, and Ocean Surveillance Center: NRaD Detachment, October 1994.
27. PRC Inc. NIPS 2.1 Applications. Technical Report PRC R-5368, Naval Command, Control, and Ocean Surveillance Center: NRaD Detachment, October 1994.
28. Myong H. Kang and Ira S. Moskowitz. A Data Pump for Communication. Submitted for publication; also available as NRL Memorandum Report 5540-95-7771, Naval Research Laboratory, Washington, D.C., December 1994.
29. Donald E. Knuth. Literate Programming. *The Computer Journal*, 27(2):97-111, May 1984.
30. Sentot Kromodimoeljo, Bill Pase, Mark Saaltink, Dan Craigen, and Irwin Meisels. EVES: An Overview. Technical report, ORA Canada, Ottawa, Ontario, February 1993.
31. Applied Physics Laboratory. Concept of operations for the Tomahawk Strike Coordination Module (TSCM). Johns Hopkins University Technical Report, February 1994.

32. Naval Research Laboratory. Trusted Tactical Aircraft Mission Planning Systems Guard Concept of Operations and Requirements Definition. Technical Report, Naval Research Laboratory, March 1995.
33. J.A. McDermid, S.P. Wilson, and P. Fenelon. ASAM-II: Concepts and Process. Technical Report ASAM-II/REQ/95.3 Issue 2.0, University of York, Department of Computer Science, University of York, Helsington, York YO1 5DD, UK, 12 March 1996.
34. Andrew P. Moore. A Behavioral Requirements Model for the NRL Pump. NRL Technical Memorandum 5540-021a:apm, Naval Research Laboratory, Washington, D.C., November 1995.
35. Andrew P. Moore and Mary F. Rock. JMCIS Information Flow Improvement (JIFI) Statemate Elements Dictionary. NRL Technical Memorandum 5540--004a, Naval Research Laboratory, October 1995.
36. National Computer Security Center, Ft. Meade, MD. *DoD 5200.28-STD, Trusted Computer System Evaluation Criteria*, December 1985.
37. Office of the Chief of Naval Operations. The Copernicus Architecture, Phase 1: Requirements Definition. Technical report, Space and Electronic Warfare, August 1991.
38. Charles N. Payne, Jr., and Andrew P. Moore. Increasing Assurance with Literate Programming Techniques. NRL Technical Memorandum 5540--276A:apm, Naval Research Laboratory, Washington, D.C., August 1995.
39. Charles N. Payne, Judith N. Froscher, and Carl E. Landwehr. Toward a Comprehensive INFOSEC Certification Methodology. In *Proc. of the 16th National Computer Security Conference*, pages 165--172, Baltimore, MD, September 1993.
40. Clive Pygott and Stephen Wilson. Uses of an Argument Framework; Its Data Model and Generic Features. Technical Report DRA/CIS3/PROJ/A53XL/95011/1.0, Defence Research Agency, Farnborough, Hampshire GU14 6TD, UK, 27 March 1996.
41. E. Wayne Sewell. *Weaving a Program: Literate Programming in WEB*. Van Nostrand Reinhold, New York, NY USA, 1989. ISBN 0-442-31946-0.
42. Space and Naval Warfare Systems Command (PMW 171). Joint Maritime Command Information System Version 2.1 (Sensitive Compartmented Information) Security Requirements Document. Technical report, Space and Naval Warfare Systems Command (PMW 171), August 1994.
43. Space and Naval Warfare Systems Command. Software Architecture Guide for the Unified Build (UB) Software Development Environment (SDE). Technical Report SPAWARSYSCOM SDE-SAG-2.0, Space and Naval Warfare Systems Command, December 1993.
44. Space and Naval Warfare Systems Command. *Navy Tactical Command System - Afloat (NTCS-A) Database Concept of Operations*, March 1992.
45. Jeffrey M. Voas and Keith W. Miller. Software Testability: The New Verification. *IEEE Software*, pages 17--28, May 1995.
46. S. Wilson, J.A. McDermid, P.M. Kirkham, and P. Fenelon. The Safety Argument Manager: An Integrated Approach to the Engineering and Safety Assessment of Computer Based Systems. In *Symposium and Workshop on Engineering of Computer-Based Systems*, pages 198--205. IEEE Computer Society Press, March 1996.

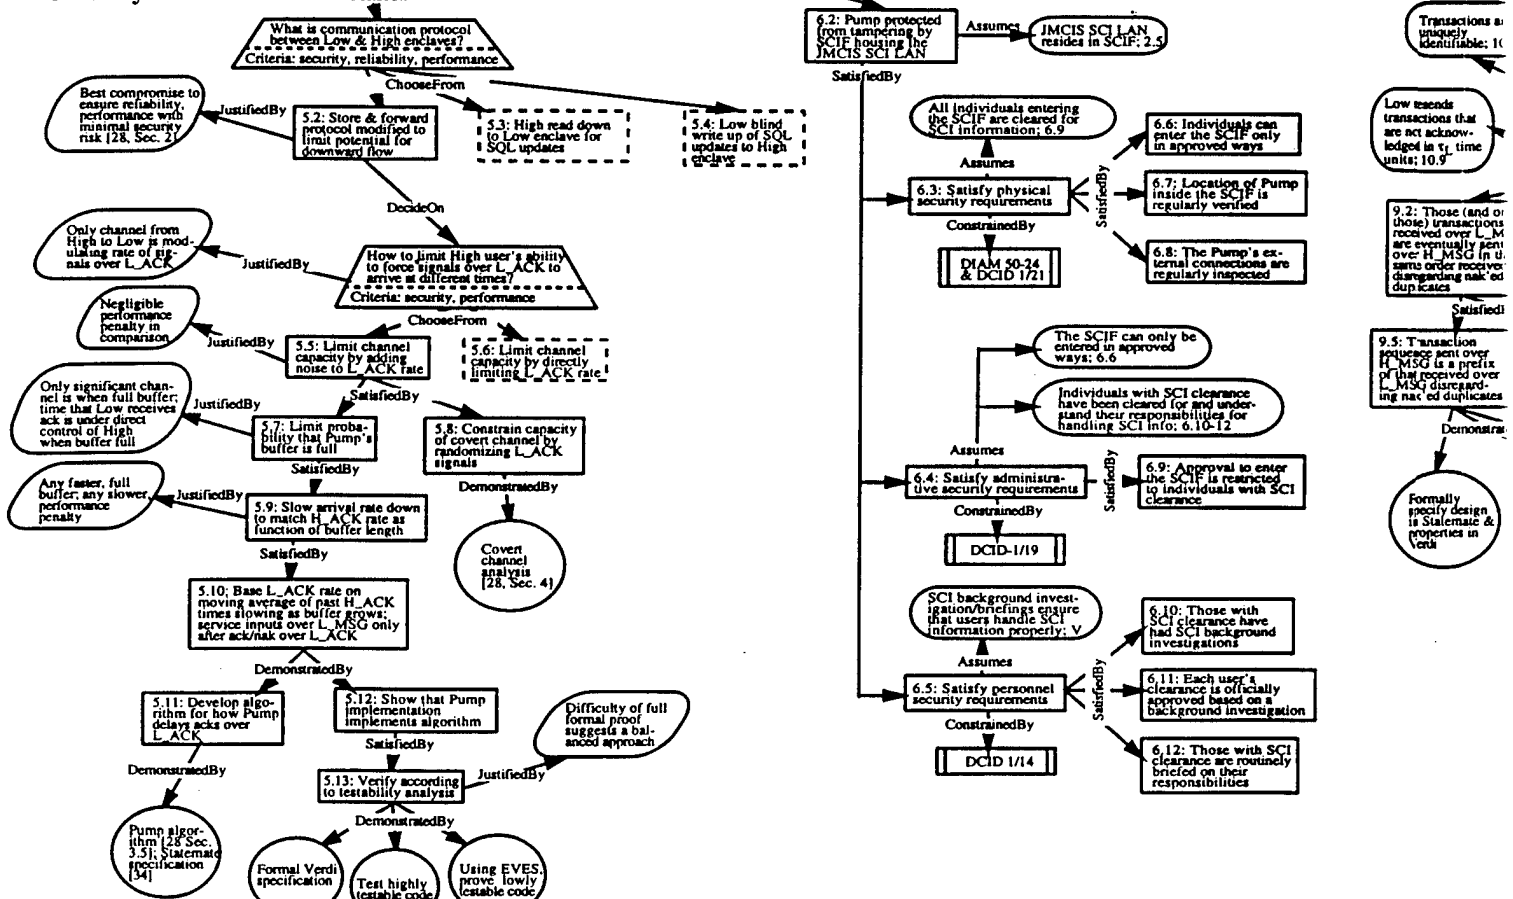
The Problem



The Architecture



Gateway Assurance



Assurance Strategy

