A Systematic Approach towards Conformance Certification in Dynamic Systems

SBIR Topic N03-208[1]

13 December 2004

Douglas Wells, Principal Investigator
<d.wells@opengroup.org>

The Open Group, Suite 325
8 New England Executive Park
Burlington, MA 01803

+1 (781) 564 9206

1. Executive Summary

Abstract

This project has developed prototypes for a set of component certification procedures and contributed to the design of a test suite designed for use with adaptive, dynamic systems, such as TSCE. Based on the best practice processes embodied in The Open Group's existing testing and certification procedures, the new processes will work with the run-time resource management subsystem and will support an incremental and evolutionary process of certification of mission suitability. The new procedures will incorporate real-time latency and availability metrics within the acceptance criteria both for subsystems and for the overall system.

This portion of the project is sponsored under a Phase I SBIR contract. The purpose of this effort is to develop a set of component certification procedures and a test suite for use with adaptive, dynamic systems, such as TSCE. Based on the best practice processes embodied in The Open Group's existing testing and certification procedures, the new processes will work with the run-time resource management subsystem and will support an incremental and evolutionary process of certification of mission suitability. The new procedures incorporate real-time latency and availability metrics within the acceptance criteria both for subsystems and for the overall system. These procedures have extended existing commercial testing and certification practice when possible.

---

[1] This document is derived from the final report for U.S. Navy Contract N00024-04-C-4140, a Phase I SBIR project sponsored by the Naval Sea Systems Command (NAVSEA 05DP).

Several specific results are included in this report including:

- a layered certification and assurance structure for complex, composite systems
- illustrations of the overall certification process that would support this layered certification process
- a position statement that provides some insight into the use of dynamic resource management in high assurance systems.

This document also provides a conceptual plan for a Phase II effort based on the notional model included.

2. Introduction

Identification and Significance of Problem or Opportunity

Problem Description

Modern weapons systems are complex and distributed, involving multiple subsystems, each of which typically comprises multiple components. System components receive data from multiple sources of information, fuse that information, select a response, and then put that response into action. The sensors and actuators are often physically remote, connected via interruptible communication paths. And, of course, weapons systems are particularly subject to physical damage.

Combat is the principal capability of a warship; therefore the weapons systems must have characteristics which can be described as "mission critical." The overall system must react to environmental events in a timely, reliable, and verifiable manner. Defensive systems, in particular, must respond to unpredictable events, and therefore must provide highly reliable, highly available service even in severe conditions.

The primary goal in designing and building a weapons system is to ensure that that the users can defeat any enemy that they might encounter. Corresponding, the primary purpose of testing and certifying the computing infrastructure within that weapons system must be to assure that it fulfills its role in satisfying that goal. Historically, the development process for the computer hardware and software for military systems has consisted of several, largely sequential and independent steps[2]:

- The service (e.g., the Navy) creates a set of requirements for the system as a whole based on anticipated missions.
- The service partitions those requirements across the various subsystems, including the computing infrastructure.
- The service devises tests that verify that subsystem requirements are satisfied.
- The service advertises those requirements for bid and selects various contractors to develop the subsystems as well as an integrator.

[2] This description is, of course, vastly oversimplified

- The contractors design and develop the subsystems to comply with the requirements.
- The government tests and certifies each subsystem and provides it to the integrator.
- The integrator assembles the components.
- The government executes sea trials resulting in overall certification of the entire system.

Traditionally, this process has spanned a decade or more and left a broad trail of change orders that were expensive in both time and resources. These changes have originated from several causes, including:

- Errors discovered in original requirements.
- Improvements discovered during development.
- Mission changes due to evolving political conditions.
- Inability to satisfy requirements due to changes in commercial markets.

The difficulties arising from introducing these changes have been multiplied due to other complications, such as lack of institutional memory, particularly due to military assignment rotations and contractor personnel changes arising from variations in contractual funding and business changes, such as mergers and consolidation.

The introduction of COTS components has enhanced the capabilities of weapons systems in several ways:

- Commercial hardware is more capable, e.g., both processors and networks are faster
- There is more function and a greater selection in software components
- The components are easier to use and have enhanced interoperability.
- Commercial components are updated on a regular basis, resulting in improved hardware performance and enhanced software capabilities.
- There is a broad support infrastructure in place with vast amounts of knowledge available via vendors' web-sites.
- There is a wider body of development tools available, and there are more developers experienced in the use of those tools.

This has led to a dramatic enhancement in computing infrastructure capabilities, particularly in the area of adaptability. The commercial components have been designed to satisfy the requirements of multiple, diverse applications and often incorporate automatic fault tolerance facilities.

There has been a general recognition in the DoD community that the enhanced capabilities that were enabled by the introduction of COTS components have surpassed the capacity of the military testing and certification processes to assure mission success. SBIR topic N03-208 properly notes that the introduction of dynamic adaptation has led to a combinatoric explosion of potential configurations that would need to be tested in order to continue with the existing validation philosophy.

We note, however, that there are a number of other, less well-recognized problems that have also been introduced as part of the adoption of COTS components:

- Many commercial components optimize throughput, rather than the predictability and/or determinism required in real-time weapons systems.
- The computer hardware marketplace is oriented around Moore's Law.
- Hardware is replaced when it ceases to be competitive, and there is often no ability to acquire older components.
- The computer software marketplace is oriented around new features. Separate bug fixes are not available for earlier releases.
- Low error rate is not a significant criterion in the commercial marketplace, leading to a higher rate of bugs - particularly security-related bugs. Undiscovered bugs emerge when components are stressed in new ways, as would occur in defense applications.
- Software upgrades must be imported from the commercial marketplace - the Internet - in order to maintain hardware compatibility.
- If systems are to communicate with external systems, they may need to be connected to the Internet, leading to concerns about security vulnerabilities.
- The support tools were not designed for developing mission critical systems, and the developers are not experienced in the design of mission critical systems.

The result is that the deployed system is neither fish nor fowl. The use of adaptive, commercial components discredits the historic testing procedures, and the time-critical aspects of the mission critical applications invalidate assumptions in commercial validation procedures. In the end, we lack assurance that the system can perform its mission.

An Analysis of the Problem in the Context of Navy Applications

The Navy's Total Ship Computing Environment (TSCE) is a prime example of the problem set described above. TSCE is a complex, system-of-systems architecture based on tens or hundreds of separate computer nodes. In order to survive battle damage, these nodes are grouped into clusters and distributed across multiple watertight compartments. The clusters are interconnected via a relatively small number of high-speed network links. If a bulkhead compartment is compromised - or if a component fails for another reason, the system automatically recovers from the loss of equipment and is subsequently reconfigured to reestablish a fault-tolerant configuration. Exhaustive testing of this system is essentially impossible. While it might be feasible to restrict initial deployment configurations to a small (and therefore testable) subset, the reestablishment of fault tolerance depends upon a dynamic inventory of what is still working during an attack - an unpredictable situation[3].

The real-time deadlines associated with TSCE applications also introduce uncertainties into assuring correct system operation. Consider the situation where a ship is supporting a land-attack mission while simultaneously conducting AAW operations. Also consider that a

[3] This problem is discussed in more detail in a paper by the PI and engineers from Lockheed Martin and Software/Technology Development Corporation. See the DSN02 paper.

network packet might get dropped within the land-attack application, engendering a timeout and retransmission. How do we know that this change in application timing won't also trigger a failure in the AAW application that is sharing the same inter-bulkhead network link? An initial response might be to always reserve extra capacity in the critical components. Unfortunately, as is also noted in the DSN02 paper, sometimes inserting a faster component (or even reducing delays) can introduce increased latency and induce secondary failures in a real-time, fault-tolerant system. So, even if we were able to test all of the hardware configurations described above, we can't even enumerate the set of possibilities introduced by the deadlines and timeouts in real-time applications.

Finally, consider the situation where a ship has been at sea on a routine assignment and is returning to its home port. Due to a world crisis, however, the ship is diverted. Because normal hardware attrition has occurred the stock of spares is low, and it is replenished with freshly manufactured components from a visiting supply ship. Unfortunately, the COTS computer manufacture has chosen to "delight its customer" by "upgrading" the machines with "faster" processors. These processors may be faster in the sense that they have higher SPEC benchmarks but they might actually be slower when executing the AAW application. We now have a situation that even a complete recertification cycle would not detect - when using a traditional testing and certification paradigm.

Fortunately, our experience with adaptive systems as part of AFRL-sponsored research [WPDRTS99] [ICCRT98] and in DARPA's Quorum program [SPAWAR02a] [SPAWAR02b] [TAO01] indicates that most of these disastrous scenarios will not occur most of the time. Unfortunately, it also indicates that they will occur some of the time, and we don't know how to derive a number for that probability. Thus, we don't currently know how to incorporate the possibility into our design processes, and we therefore can't supply a metric of the level of assurance provided by our mission critical system. There appears to be no means for using traditional exhaustive testing techniques to develop such a metric.

3. Phase I Investigation

Overview

During the Phase I portion of this project, we investigated the use of existing commercial processes for testing and certification within the dynamic environment of the Navy's Total Ship Computing Environment. We examined several related issues:

- We examined the existing best practice procedures developed by The Open Group and widely adopted within industry with the intent of enhancing them for use in mission critical applications with real-time deadlines and availability requirements.
- We examined the operation of a hypothetical TSCE application as a use case in order to investigate how to increase the testability of TCSE applications for use in high assurance, mission critical systems.
- We investigated whether the information used by the resource management subsystem in dynamically managing a TSCE system could be used within the testing infrastructure in order to dynamically evaluate the system's level of assurance.

Originally, there were two goals for phase I. First, we were to identify procedures based on The Open Group's best practice COTS certification methods and procedures. These enhanced procedures would then be folded back resulting in a COTS certification system that would be applicable for systems with both components intended for use in the commercial marketplace as well as with components intended for mission critical applications in Navy ship-board systems. We also expected that components certified using these procedures would be acceptable for use in multiple systems certified by the Navy as mission capable.

The second result would be identification of a set of testing tools for use in testing mission critical systems that utilize these best practice certification procedures. We expected that these tools would be based on extensions to The Open Group's existing set of testing tools, such that the enhanced tools would be usable for components intended for the commercial marketplace as well as for use in mission critical applications in Navy ship-board systems that include both COTS components as well as purpose-built components.

Phase I Technical Objectives

As specified in the original proposal, the technical objectives of the Phase I effort were to:

1) Identify a set of performance criteria (e.g., execution time) that would be testable by straight-forward extensions to existing test suites and would also be applicable to an existing application that is relevant in a TSCE systems.

2) Develop an application model that predicts the performance of (a simplified version of) the application based on the performance criteria identified in objective 1.

3) Evaluate the difficulty and utility of adding performance criteria acceptance tests to existing test procedures, such as The Open Group's test suite for POSIX Realtime[4].

4) Evaluate the feasibility of modifying Navy system certification procedures to incorporate and build upon commercial component certification practices, such as those developed by The Open Group.

5) Identify a set of enhancements to existing certification procedures that might make them applicable for use in Navy certification procedures.

6) Investigate the effect of "composable" standards in mission critical systems - of incorporating one standard in another.

7) Create a development plan for a Phase II SBIR for developing testing and certification procedures for use in mission critical, high assurance systems.

Phase I Work Plan

The Work Plan for Phase I included the following work activities to investigate feasibility of developing enhanced testing tools and certification procedures for use in high assurance, mission critical systems.

4http://www.opengroup.org/testing/testsuites/TestSuiteIndex.htm#POSIX

1) Evaluate existing testing procedures to identify candidate evaluation criteria for use in real-time, mission critical systems. We would start with The Open Group's test suite for POSIX Realtime and identify relevant execution time criteria for representative APIs.

2) Develop prototype extensions to the POSIX Realtime test suite to characterize a performance-relevant metric and to incorporate that metric within the acceptance criteria for the test suite.

3) Identify a suitable application and/or implementation of the POSIX Realtime specification in order to exercise the extensions developed in work item 2.

4) Extend the results of work item 3 into an execution model for the application that correlates application performance and the metrics created by work item 2.

5) Evaluate the results of the work item 4 for the purpose of determining the suitability of using performance criteria in a test suite for high assurance, mission critical applications.

6) Identify and document the certification criteria used by the Navy in one or more hardware and or software components or subsystems.

7) Cast the results of work item 6 in the context of the industry certification procedures developed by and used by The Open Group.

8) Develop a hypothetical set of certification procedures that would achieve the Navy's goals for system certification, but are based on The Open Group's commercial certification procedures.

9) Evaluate the effects of having multiple layers of standards, particularly with respect to the issue of having varying compliance levels.

10) Evaluate the results of early work and develop a Phase II work plan and proposal.

Option:

11) If the option is approved, extend the investigation to explore the use of the testing procedures in the context of the TETware[5] testing harness. This would improve the integration with the ongoing Phase II SBIR work on the Open Tool Kit for Mission Critical Systems and would simply the integration into the existing test suite products from The Open Group.

Investigation

The Open Group has expertise in many areas. For this project, we created a team with project members from two critical areas. James Andrews and Deborah Schoonover work with external organizations to develop techniques and procedures for establishing testing and certification programs that meet the specialized requirements of those groups. They also use this knowledge and experience to adapt and extend our standard certification procedures. Originally, The Open Group (then operating as X/Open) only certified the UNIX® operating system. Today, it also provides certification services in areas such as Web-enabled

5http://tetworks.opengroup.org/Wpapers/TETwareRTWhitePaper.htm#Introduction

cellphones, software for K-12 schools, state lottery systems, and computer architecture frameworks. The result is that The Open Group has a basic template for certification services that can be straight-forwardly applied in many domains.

The other project members provided expertise in interoperability, particularly in real-time and distributed systems. David Lounsbury, John Spaulding, and Douglas have worked on numerous projects with industry and military organizations, including working with NSWC Dahlgren on technology for HiPer-D, which blazed the trail for many of the concepts that are being used in the Total Ship Computing Environment infrastructure. They have also developed and worked with techniques for dynamic resource management in these systems.

After initial briefings on the latest developments in each of the areas, the project team applied the corporate standard certification product template to the problem area of dynamic systems, such as are expected to be used in DD(X) ships. We first examined the TSCE concept and discussed some of the problem areas that had been previously uncovered relative to the difficulty of developing metrics for dynamic systems. We then reviewed the certification model and developed an initial set of similarities and analogous constituent components.

The primary results of this discussion were:

- a common understanding of the problem area, and
- an enumeration of existing certification programs where the underlying model has been extended in a way similar to what might be required in a Navy certification program.

Suitability for Purpose

The most striking result of this investigation was that most of the Navy's accreditation process is much more similar to The Open Group's best practice certification activities than it is to the product certification activities. In exploring this difference, we eventually came to describe the Navy's accreditation process as certifying "suitability for purpose," a concept that aided us in better understanding The Open Group's own certification programs and the relationship among them.

The Open Group has two basic structures for certification: products and best practice:

- Products are certified as conforming to a particular specification. The certification process usually includes successful execution of a test suite. In the case of anomalous test results or a user-initiated challenge, the resolution process might update the test process or "clarify" the specification.
- Individuals and organizations are certified as being knowledgeable and capable in a field. The certification process usually involves some knowledge testing and an attestation that the practitioner understands and follows the "best practices" known for the particular field.

The fundamental difference between these two types of programs is where the domain knowledge must exist, and thus who is responsible for proper system design. Consider the case of a system that monitors the temperature of a chemical process using a user written application operating on a UNIX system that communicates with a particular temperature sensor using an RS232-based serial communication protocol. If the application encounters a

problem with the UNIX *open* system call, the vendor of the UNIX system may be at fault. If the application can be shown to correctly use the *open* system call, the vendor must fix the system or risk losing its UNIX certification. If, on the other hand, the application fails to respond in time to a temperature anomaly, the vendor's UNIX certification is not at risk — because the UNIX specification does not include timing requirements. A system designer who was a TOGAF[6] certified practitioner[7], however, would be expected to have anticipated this sort of problem by understanding the problem domain and by working with other experts to anticipate and avoid such problems.

The issue that arises here is one of marketplace practicality: the size of the potential vendor community varies inversely with the degree to which that vendor has the required domain knowledge. At one extreme, we have the traditional military procurement model where the prime contractor is entirely responsible for the delivery of unique products; at the other extreme, we have the COTS commodity marketplace where the vendor knows nothing about the specific domain in which its product is used. Thus, in order to be able to leverage the COTS marketplace, we need to eliminate or reduce the uniqueness of the product.

System Composition and Assurance

In fact, this reduction of uniqueness is a normal part of the engineering process. A system engineer does not take the requirement of building a land-attack ship and immediately determine that it will require 2 million rivets, 50 computers, and some miscellaneous other parts. Instead, the original problem must be decomposed into multiple smaller, but simpler, problems, and then those problems are decomposed into even more, even simpler problems. The DD(X) system engineer may determine that each ship will require a gun system, a radar system, and various other components. Then requirements for each of those components must be determined. For example, the gun system might be required to fire 10 rounds a minute to a distance of 50 miles. The system engineer might also impose other requirements in order to promote overall system operation. For example, some smaller guns might be required to use 5 inch shells because that is what other ships use, and use of common supplies simplifies the global military logistics problem.

In many cases, the decomposition results in sub-problems that require different domain expertise. The ship system engineer might not understand how to cast gun barrels that can withstand the force required to throw a shell 50 miles, but it can specify requirements that the gun barrel must be able to do so. The vendor of the gun system, however, can be expected to know exactly how to cast such a gun barrel. It might be the case, however, that the gun vendor does not know how to turn iron ore into the type of steel that is necessary to create such a robust gun barrel.

The result is that the overall system is constructed as a layered structure of acquired components. At each such interface, the user and provider of the component have agreed as to the requirements that the component must satisfy. The provider of the component has had to determine that it can, in fact, construct a component that fulfills those requirements. The

---

[6] The Open Group Architecture Framework, http://www.opengroup.org/architecture/tpgaf8/

[7] http://www.opengroup.org/togaf/cert/cert_prodlist.tpl

acquirer of the component has also had to decide that it can fulfill the requirements for the component that it is building if each of its suppliers meets their requirements. At the top, the ship system engineer believes that it can build a properly functioning ship from the tree of components.

This engineering process is expensive, however, and introducing perturbations makes it even more expensive. If one supplier should fail to deliver its component, or if that component does not meet its specification, the user of that component can not meet its obligations. In a cost competitive environment, something has to give: the supplier will have to figure out a different way of creating the component, the acquirer will have to figure out how to use a different component. Someone will have to incur additional engineering costs. Often a replacement component will be more expensive, or it might otherwise fulfill the requirements less well — otherwise, why wouldn't the acquirer have chosen the replacement component originally. Usually, the schedule will slip for the component, and that slip might propagate all to way to the delivery date of the ship.

Each provider in this process has carefully prepared an estimate of its development costs and has usually pared its development effort as mach as possible in order to be selected as the "low cost" bidder. As a result, the participants attempt to protect themselves against these contingencies. Often the acquisition of components occurs between different companies or different divisions within a single company, in which case there is almost always a formal contract between the two parties — often with penalty clauses if either party fails to meet its obligations. Even if the development is spread across multiple divisions of a company, there may be formal obligations, with "career-limiting" penalties to those who fail to perform properly.

The problem with imposing such penalties, however, is that they are *ex post facto*. Money may transfer from one party to another, but the identification of a problem still doesn't occur until some deadline has been missed, or some component does not work properly with some other component — but the project is still late. In order to identify problems early, most development plans include processes and procedures for qualifying components prior to system completion. Typical examples include unit tests, qualified vendor lists, and audit teams. The goal is to identify — and correct — problems early, when it is assumed that the ramifications are minimized. A beneficial side effect is that the resolution to many problems need only involve two component developers.

Component Assurance

The practical result is that most large systems are complex compositions, with no one person or organization having detailed knowledge of every component. System assurance is based on a complex of web of trust among the various providers. Most of the individual assurances bonds are based primarily on one of two fundamental indicators of competence that a supplier can satisfy requirements:

> The simpler indicator is compliance to specification: Does the component adhere to an explicitly defined set of requirements? Will the component pass a set of procedures that test adherence to the specification.

The other indicator is competence of the provider: Do the developers know what they are doing? Can they be expected to provide a component that works properly. Will they anticipate, identify and alleviate problems during the development and maintenance cycles?

Checking conformance to a well-defined specification is more straight forward in several ways. It requires less knowledge about the component on the part of the procurer, it is easier to include in contracts, and third parties can better component it, a concept that can become important if disputes should enter into the legal system where peer juries may need to attribute blame. There are several issues that may prevent the use of this method.

First, there must be a well-defined specification. Specifications are expensive to develop and require a significant investment of expert knowledge. Often these specifications are developed by industry groups, where the costs are shared by many participants. The Open Group supports many such efforts, including, for example, the UNIX and Application Response Measurement (ARM) standardization processes.

Second, the specification must adequately match the requirements for the component. The specification must define the proper requirements: if a bolt must mate with a metric nut, it does no good for the bolt to comply with English standards. The specification must also define sufficient capabilities: if a bolt must withstand loads of 500 pounds, it is not useful to acquire bolts that adhere to a standard that requires loads of only 100 pounds.

Third, complete conformance tests are expensive to develop. As a result, most compliance test suites actually only test a subset of requirements. Identifying that proper subset normally requires expert knowledge about the component and the development processes for the component, the avoidance of which was, of course, the goals of using such a test. Thus, most test suites are defined and developed by component developers with input from potential users.

Fourth, sometimes it is just not possible to develop a well-defined specification. Sometimes the requirements can not be clearly and specifically identified. For example, one can not write a requirement for an apple that “tastes good,” or one that works well in “baked goods.” Instead, one usually provides substitute requirements, such as wanting Golden Delicious or Cortland apples, which again transfers much of the required expertise back to the purchaser. Consider a requirement that a computer system be “scalable,” or adaptable to future industry trends. These are not precise, specific requirements.

Finally, sometimes it is not possible to test compliance to a specification in a reasonable way. For example, it might be a requirement that ship hull rivets should last 25 years, but how does one test this directly. If you place rivets in a seawater tank and wait 25 years, your results are not timely and almost certainly irrelevant. Again, a more practical alternative would be to specify an alternate set of requirements, for example, one might require that the rivets be galvanized, and conform to some relevant industry standard.

Provider Assurance

In selecting those alternative requirements, it would be nice if one could depend upon the expert knowledge of the candidate component suppliers, but checking the competence of a provider is difficult. If one is lucky, the choice is between selecting an expert provider and a run-of-the-mill provider. But again, the reason for going to an external supplier is often a local lack of expert domain knowledge. So, how does one select the best provider based on competence? Almost always, the selection process turns out to be reputation, largely based on past performance, and investigation of operational procedures.

Reputation can be determined based on contacts within the field and reference accounts provided by the supplier, but detailed analysis of internal development procedures would require detailed knowledge of technology used within the component. So, the review is usually based on good general engineering practice: Are the developers knowledgeable about their field? Do they have appropriate education, training, and certification? Are they active within the relevant professional groups, such as IEEE, ACM, ION, ASCE, etc? Do they have internal design reviews and quality assurance procedures?

Recently, customers have increased requirements for quality, and many technology and engineering companies have strived to satisfy this need. Most organizations have improved internal processes, and many have attained ISO 9000 certification, which requires explicit, documented processes. In doing so, a common concept of Best Practices has emerged. Best Practices are development and operational procedures that have been adopted by many practitioners in the field and are recognized as being suitable for adoption across the entire field. Usually, documentation of these Best Practices are publicly available. In many cases, such as in The Open Group's Architecture Forum, these Best Practices are the basis of certification of individuals and organizations. Often, such certification can provide substantial evidence of provider competence.

Assurance in Composite Systems

As noted earlier, the development of complex system is usually decomposed into subproblems, and then further into sub-subproblems, eventually resulting in a composite web of components, with no one person or organization having full knowledge of every aspect of every component. The provision of assurance must utilize the existence of this composite web, using the availability of common knowledge at each of the acquisition boundaries. Thus, system assurance is created by requiring each component provider to assure the quality of its sub-components. In fact, a common Best Practice is to utilize high quality suppliers.

Although the two types of assurance validation occur throughout the decomposition hierarchy, we note that higher levels tend to depend more on provider assurance, and lower levels tend to depend more on component assurance. For example, there is basically no way to verify that an airplane is safe by straightforward testing, so the FAA requires that all aircraft be certified under DO-178B, which is essentially a Best Practices certification. On the other hand, many vendors sell rivets, which can be analyzed in a Qualify Assurance Department incoming inspection procedure.

4. Phase I Results

The results of this phase of the project fall into two categories: those that could be applied to future work in assisting the Navy with certification of components in ships with dynamic, adaptive computing infrastructures, and those that we are already applying to existing projects within our company.

In spite of several requests within the Navy, we were unable to interest any Navy testing and certification groups in working closely with us to develop a plan for a Phase II project. In some cases, the groups had not been funded. In others, they did not have clear direction of how they should test and accredit future ships. It appears to be the case that the Navy's Open Architecture concept is subject to intense discussion within the Navy. We therefore developed a notional concept of ship accreditation based on our past experience in working with the HiPer-D group at NSWC Dahlgren. The supporting certification and assurance structure in include in this report, and an example component certification process to illustrate the concept is also included below. We have purposely omitted details as they would be incorrect. Our intent is to provide an indication that it would be possible to develop a structured, multi-layer component certification structure based on existing and future industry standards, such as POISIX, UNIX, CORBA, TCP/IP, etc. We also investigated the possibility of certification of components that would include dynamic, run-time verification of proper operation, with fallback to simpler modes of operation when necessary. A brief discussion on this is included in a position paper submitted to a DARPA Java Workshop, which is included below.

We have also applied the knowledge developed during this project to existing projects. Our analysis of application of real-time requirement within certified components has been applied to a companion Phase I SIBR project, Practical Test and Verification of Component Behavioral Characteristics (N03-088), which has been recommended for a Phase II effort. Our analysis of the relationship between component testing and use of Best Practices has resulted in a more encompassing certification diagram that covers both areas (and which is also included in this report). Finally, we have incorporated this knowledge into several ongoing standards and certification process development activities within the company, include projects for Safety Critical Java and IPv6 Internet Standards.

5. Opportunity

In spire of our inability to acquire a Navy partner in the current phase of the project, we still believe that there is a definite opportunity: The DD(X) ship is being developed; it is using a significant number of COTS components; and those COTS components will need to be updated and replaced on a time scale much closer to the several years commonly used in industry rather than the several decades traditionally used in the Navy.

The economic situation has not changed relative to the need to incorporate COTS components into military systems: There is a correlation among the three factors in adaptive systems: testing capability, dynamic adaptability, and operational assurance. Even if we should fail to make significant progress in specifying and testing adaptive systems, we can achieve higher operational assurance by limiting the dynamic, run-time adaptability (essentially trading off capability of adaptation against assurance). Correspondingly, for a particular level of operational assurance, the higher our capability for testing, the greater is

the run-time dynamic adaptability that we can allow. Thus, as our specification and testing procedures improve over time, the capability of the system will also improve - while maintaining the same level of confidence in system operation.

This means that we have an extant opportunity to leverage an improvement in testing and certification - and in the level of system assurance - based on the disruptive nature of the ongoing conversion to TSCE. Failure to exploit this leverage in testing and certification will result in a reduced operational capability for TSCE.

Benefits of Approach

It is clear that use of adaptive systems in mission critical systems is at a very early stage. Much research is ongoing, with significant support from both DARPA and ONR. Nonetheless, it is being applied to real applications in real systems - now. By leveraging this experience base, we can improve both the testing infrastructure and the capability of the mission critical applications.

In addition, our approach is incremental. As lessons are learned in developing adaptive applications and in using those applications in resource managed environments, such as TSCE, we can apply those lessons to the testing infrastructure. Similarly, as lessons are learned in testing mission critical systems, they can be leveraged into both application development and resource management. This spiral approach provides the following advantages:

- Improved potential for use of COTS real-time and fault tolerant components
- Reduced costs for testing due to more effective unit testing and less interaction between defects in real-time and fault tolerant components and overall system operation
- Less dependence on design decisions made many years before deployment. A system engineer can alter parameters for timeliness and failure recovery strategies after system integration — when their interactions can be experimentally measured and verified.
- Reduction in development costs and time due to increased reuse of components
- Improved mission effectiveness due to better tested, more dependable components
- More effective field upgrades. Modifications to adapt weapons systems to new battlefield situations can be performed with significantly less development cost — and particularly less time — because many of the modifications can be performed by adjusting component parameters. Even those that require software development will be simplified because of the separation of components.
- Reduced upgrade costs due to less dependency on the particular character of particular hardware and/or software components

6. Commercialization Strategy

The enhanced testing and certification facilities proposed here are a natural extension of The Open Group's existing capabilities, which have already been adopted by many potential suppliers of components into TSCE systems (and therefore into the fleet). Thus, we expect that extensions of our in-place sales, marketing, and operational structures can be used in

support of these enhanced testing and certification capabilities. In addition, the logical users of these products are already members of The Open Group and particularly the Real-Time and Embedded Systems Forum. Thus, we have established communication paths with them whereby we can work cooperatively to adapt these new products to their evolving needs.

What we must work on is establishing relationships and mutual understanding with the people and organizations who would specify, mandate, and manage the certification of military systems. Traditionally, this process has been operated by the individual military branches and often by operational divisions within those branches. There is, however, a growing pressure within DoD to establish common infrastructures and to foster interoperability within a "network-centric" operational environment. This movement is creating turmoil within the military relative to establishing standards and mandating development processes.

We are working to establish communications with these organizations via our membership base. For example, Lt. Col. Glen Logan, the co-chair of the Real-Time and Embedded Systems Forum is Deputy Director of the DoD Open Systems Joint Task Force and is assisting us in working within the DoD mission critical open systems community. Also, Elaine Babcock, a member of The Open Group's board of directors, is Chief of the Standards Integration Division, Defense Information Systems Agency (DISA) and is assisting us in working with the more conventional IT aspects of the DoD environment. We will continue working on establishing and fostering these relationships as part of this effort.

In addition, we believe that other industries with real-time, mission critical applications are natural candidates for these enhanced testing and certification capabilities. We will reach out to communities, such as process control and factory automation via our real-time product vendor members as well as by collaborations with our sister industry consortium organizations, such as the Object Management Group (OMG).

7. Phase II Development

We are not including a detailed Phase II Development plan due to the inability to identify a Navy accreditation group that could provide the necessary detail. We do believe that the notional layered certification model and illustrative certification process could be extended with little difficulty.

8. Notional Layered Certification Model

The team developed the following layered model of a notional certification program for a dynamic system, such as the computing infrastructure on DD(X). This model is only a straw man. The selection and identification of seven layers is incidental and the number is likely to change in future revisions.

This structure has been engineered to have two particularly interesting attributes:

- the top layer matches the Navy's existing ship accreditation program (at least to the level that we understand it), and
- the bottom layer matches The Open Group's existing, deployed, commercially proven certification program.

In between, there are several layers, each of which has relevant attributes:

- the layer specifies entities that are useful sub-components to higher levels, and
- the entities in a layer often incorporate sub-components from lower levels.

The introduction of layering provides several benefits:

- The lower the layer, the more likely it is that the products certified at that level can be used in other systems, both military and commercial.
- The existence of multiple layers increases the likelihood that the certification program for one layer can reuse the results of lower levels by simply requiring a component certified at the lower level.
- The lower down the model is, the more directly The Open Group's existing testing and certification infrastructure can be utilized to address the requirements at that layer.

In this particular layering model, there is direct synergy even as high as level 4. Above that, the team has an incomplete understanding of the methods and processes that are utilized for certification/accreditation. Our expectation is that the higher levels are more concerned with certification against best practice and acceptance testing based on broad requirements and usability constraints rather than simple testing against a detailed functional specification.

The team therefore decided to leverage its existing knowledge base to concentrate in developing ideas that can be applied in the lowered four layers initially. In parallel the team will endeavor to make contact with Navy and Navy contractor sources to better understand current practice for certification on the higher layers. The results of this research will enable the team to identify the specific added value that The Open Group can offer in this space in Phase II and Phase III SBIR projects.

The team identified a layered model of certification as follows. Under each layer we provide examples of certification criteria which are for illustrative purposes only. Real life certification criteria will almost certainly differ from these (and some are most likely classified!)`

- Layer 1: Ship Accreditation, e.g. DD(X)
  - ship floats (!)
  - able to defend itself against two concurrently attacking cruise missiles
  - able to monitor and guide 8 simultaneous missile launch
  - able to support 4 simultaneous land attack missions
  - other classified requirements
- Layer 2: Application Certification, e.g., Auto Special Defense
  - can detect attacking cruise missile at [classified] distance
  - can detect and launch weapons against two concurrently
  - attacking cruise missile in less than [classified] seconds
- Layer 3: Component Certification, e.g., Radar System
  - can detect a F16 at [classified] distance

    - can monitor [classified] concurrent tracks
    - can re-target beam in less than 500 milliseconds
- Layer 4: Commercial Component Environment Certification: e.g., Real-Time CORBA
    - underlying cluster has commercial component composite certification
    - end-to-end latency of less than 500 usecs using CORBA benchmark
    - availability >= 6 nines
- Layer 5: Specific Components have Enhanced Commercial Component Certification
    - underlying cluster has commercial component composite certification
    - node-to-node latency of less than 300 usecs
    - availability >= 7 nines
- Layer 6: Enhanced Commercial Component Certification. e.g., RT, reliable UNIX platform
    - commercial component certification
    - SPECmark >= TBS
    - user-level interrupt latency < 50 usec
    - mutexs use priority inheritance
    - availability >= 5 nines
- Layer 7: Commercial Component Certification, e.g., UNIX Platform
    - UNIX conformant at Profile 54

9. Certification Process Procedures

One result of this project was a reexamination of The Open Group's certification programs themselves. Previously each program had been the result of numerous tweaks over time as various operational problems had arisen and been addressed. This rapid expansion of existing and potential certification programs caused our certification team to reevaluate the business area and to examine it in the context of regularizing the various programs into a common, coherent process. The primary result has been a better internal understanding of why the programs have been successful at a fundamental level and how to apply the newly developed common process in new domains. It should also be noted that this reexamination did not produce any significant changes in operational characteristics of the existing programs, indicating that there had been a common, even if unidentified, process.

The two diagrams represent the view of the vendor preparing for submission and then during the actual submission process.

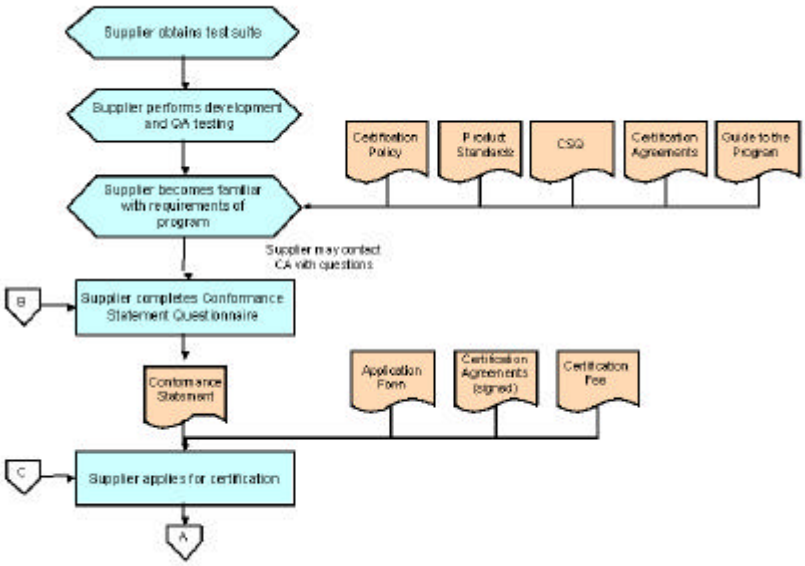

Figure 1: Vendor Preparation Process

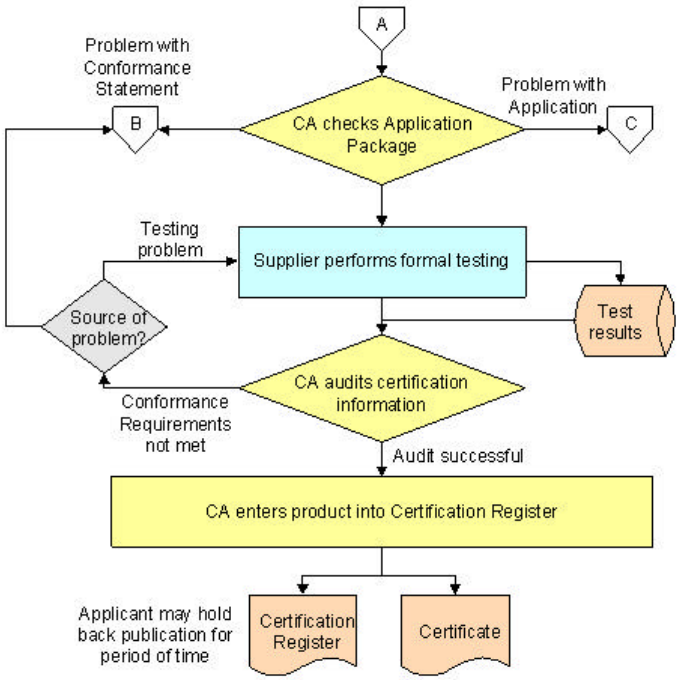

Figure 2: Submission Process

10. Best Practice Certification

The team analyzed both The Open Group's best practice certification model and our product certification model. The components included in the two models are similar. The only real difference between the two models is that the indicator for conformance for product certification is typically a standardized test suite and for best practice certification it is an assessment process and assessment report.

Below is a review of the best practice certification model (See Certification Process Procedures above). We have described each of the components of the model, and where applicable, indicated whether the component also applies to product certification. We have also provided information on whether the component might be applicable to the proposed Navy certification model, either as is or with modification.

Best Practice

This is the Best Practice standard against which an organization is to be certified. This model would apply equally well if this document is some other form of defined requirements. For product certification, this would likely be a technical standard document.

COTS certification is usually concerned with conformity to a specification rather than the more abstract concept of fitness for purpose. As the upper three layers of our notional layered certification model for the Navy do not represent COTS components, they may not lend themselves to a precise definition of technical requirements, We understand that at those levels best practices are used to qualify the sub-systems and systems rather than technical evaluations. For example, simulation of the components in an operational environment at mission time may be used to determine performance criteria, which in turn can be used to form certification criteria for procured sub-systems and systems as a whole. This is a best practice approach to certification rather a COTS approach. The rigorous processes that The Open Group has developed may be of assistance to the Navy in ensuring the repeatability, reproducibility, and objectivity of the application of the best practice.

At the lower four levels, traditional COTS style product certification could be directly applied to further the Navy Procurements against technical standards.

Organization implements best practice

For best practice certification, the integrator must design, develop, and implement the systems or subsystems against best practices by establishing the necessary processes and procedures. For product certification, the organization must design and implement the technical requirements, usually contained in a technical standard. The Open Group Quality assurance Best practice certification program contains procedures that may provide an outsourced resource to address the Navy's needs to police adherence to best practice requirements in design, development and acceptance.

Organization becomes familiar with program requirements

The organization not only has to understand and implement the best practice or technical requirements but also understand the requirements of the certification program itself. The Open Group has developed a set of documents that can aid the organization. The use of

customized versions of these standardized document templates could bring corresponding benefits to the Navy for procurement of components of the Warship.

Certification Policy

This document defines the policies, processes, and procedures for certification. It essentially defines the requirements for the certification program, and is applicable to both best practice and product certification. With greater understanding of the Navy requirements for the non-COTS procurement, a customized version of this document could form a standardized means for the Navy to communicate its design development, implementation, and acceptance testing requirements for the higher three certification layers. A customized version of The Open Group's Certification Policy document could also aid the Navy procurements of high integrity COTS components at the lower four layers. In particular, we believe that The Open Group's experiences in including interoperability criteria as part of its certification policy for some programs would meet many of the Navy needs with minimal adaptation.

Conformance Requirements

Each of The Open Group's certification programs has a document that precisely defines what the requirements for certification are. The requirements could be derived from a technical standard(s), a best practice(s), a combination of these, or any requirement howsoever documented. This applies to both best practice and product certification and we believe could benefit the Navy by providing a standard way for it to communicate its requirements unambiguously to its suppliers, for both COTS and bespoke systems.

Conformance Statement Questionnaire

Each of The Open Group's certification programs uses a questionnaire to enable the organization to define the precise way in which it has implemented the conformance requirements, and is applicable to both best practice and product certification. When completed, this document becomes the conformance statement. This document provides a succinct way for heterogeneous COTS components to be integrated with confidence and sets the criteria for interoperability trials. This would appear to be of great benefit to Navy particularly for multiple program COTS components.

Certification Program Guide

This is a user document that helps guide an organization through the certification process, explaining in detail what needs to be done and how to perform the various steps in the process, and is applicable to both best practice and product certification. A similar approach could bring benefits to the Navy as a consistent means to communicate its certification requirements to its supplier of either COTS or Bespoke components

Certification Agreement

This is the legal contract between the organization and the Certification Authority, governing the terms and conditions of the certification service provided by the Certification Authority as well as the obligations of the organization. It will also contain the warrantee of continued conformance. It also is applicable to both best practice and product certification The team believes that some form of warrantee of continued conformity would provide additional assurance for COTS components procured by the Navy. This could either be by referencing

preexisting COTS certification programs such as those operated by The Open Group or through The Open Group providing expertise to Naval procurement specific programs, for example at the upper three layers where bespoke systems are typically procured.

Organization registers for certification

At this point the organization completes a registration form, and signs the certification agreement and submits it together with any fees to the Certification Authority. The registration will be checked for completeness by the Certification Authority and the organization informed if there are problems. The registration for certification is now formal. It is presumed that this would happen at all levels of the proposed certification model, with the Navy, or its agent, fulfilling the role of the Certification Authority.

Fees

Commercial certification programs are, generally speaking, designed to be self-sustaining; that is, the consumer of the services pays for those services. Other models are possible, for example, the certification program may be sustained by a third-party. We suspect that even if there is not a direct financial cost for the supplier in certification and accreditation by the Navy, there are indirect costs in supporting acceptance trials. In the case of COTS components, where some of the costs of meeting the certification requirements can be spread over the entire market place for a COTS product (beyond just military use), both the supplier and the Navy will enjoy cost reductions related to demonstrating functional conformity and interoperability.

Organization completes Conformance Statement Questionnaire

Questions would include:

Administrative

A precise definition on what is to be certified, its scope, and environment. Note that in commercial certification this is commonly more than a specific instance of a product or practice, but commonly applies to a family of products or practices, which conform over time. The warrantee of conformance is used to legally enforce this.

Options

Where the best practice or technical standard has optional requirements, this is a statement as to whether the option is supported.

In cases where the organization has discretion on what is implemented, this captures how particular requirements have been implemented.

Other information

Typically information required for configuring the indicator of conformance, and/or interpreting its results

Conformance Statement

The completed conformance statement questionnaire is known as the Conformance Statement. In The Open Group's COTS certification, this is a public document, available to

all potential procurers of the certified product or practice. This could probably remain the case for the lower four layers of the proposed layered certification model, but it is probable that at the highest three layers, the Conformance Statement would only be available to the Navy as procurer, as some information contained within the Conformance Statement could be classified or restricted. The team envisages that the Navy may find merit in making use of customized versions of the templates in non-COTS procurement and can directly make use of The Open Group's public Conformance Statement database for COTS components, much as military procurers do in the COE POSIX-based Platform Certification program.

Organization gathers required documentation & completes checklist

In the certification of best practices, certain documents may be required for assessment purposes to ensure that best practice requirements have been met. While the Conformance Statement documents the optional requirements supported, it does not address those requirements that are mandatory. The checklist is where the organization asserts that all the requirements, whether option or mandatory, have been implemented and provides information on where evidence can be found in the documents provided to support this assertion. This component, specific to best practice certification, is only applicable to layers 1-3. The checklist is a confidential document in The Open Group's best practice certification programs. There is no analogy to this document for the COTS certification programs.

Assessor performs On-site Assessment

This applies only to best practice certification. Assessment consists of accredited experts interviewing key practitioners and examining records and documents in order to verify whether all requirements are met, as well as to identify any shortcomings with meeting the conformance requirements. Again there are security considerations that may apply to this approach when applied to the Navy's certification needs. However it is presumed that both suppliers and Navy personal have occasion to interact on the supplier's site. The extent to which the use of The Open Group's expertise and procedures would benefit the Navy needs further discussion.

Conformance Requirements met?

The assessor will make formal recorded observations during the on-site assessment. At the completion of the on-site assessment, the assessor will identify whether any of these observations represent non-compliance with the best practice and make a recommendation on the assessment outcome. The assessment report includes the recommendation together any identified non-compliance. The assessment report will be sent to the Certification Authority for review. The team suspects that an activity similar to this already takes place at levels 1-3, and is thus applicable. The extent to which the use of The Open Group's expertise and procedures would benefit the Navy needs further discussion.

At levels below this, where product certification may apply, this determination on whether the conformance requirements are met is done based on a combination of the test results, the conformance statement, and any supporting information such as agreed interpretations to the test results or the base specification. This activity takes place at the Certification Authority's site and the organization is informed of the outcome by email. Should the Navy choose to make use of preexisting COTS certification programs such as those operated by The Open

Group, the Navy will be freed of some of the assessment and acceptance costs of procuring COTS components.

Organization & Assessor agree action plan

If the conformance requirements have not been fully met, then corrective action will be necessary. In the case of best practice certification, the nature and timing of this corrective action will be agreed with the organization on-site and certification may or may not proceed pending the completion of the corrective action plan (dependent upon the nature and severity of the corrective actions). In the case of product certification, certification cannot case take place until the conformity issues are addressed. The team suspects that these concepts may be familiar to the Navy and its suppliers and be currently taking place in some form.

Organization implements corrective actions

Self-explanatory.

Assessor performs follow-up

In the case of best practice certification, the assessor will follow-up with the organization to determine if it has satisfactorily implemented the corrective actions. This follow-up may be done remotely via examination of resubmitted materials. In some cases, follow-up assessment will be required. In either case, the assessor will update the assessment report with the results of the follow-up.

In the case of product certification, this reassessment would take place at the Certification Authority's premises. Again the team suspects that these concepts will be familiar to the Navy and may currently be applied.

CA audits assessment report

This would applies only to best practice certification. In this final step, the Certification Authority reviews the assessment report to ensure that the correct assessment process has been followed. There is no equivalent for product certification.

CA works with Assessor and Organization to determine next steps

If the Certification Authority determines during it's review of the assessment report that the conformance requirements have not been met, or that there is some problem with the assessment process, the Certification Authority will work with all parties to resolve the issues so that certification may proceed. This applies only to best practice certification.

Trademark License Agreement signed?

In certification programs that use a trademark in association with certified products and practices, the Certification Authority will validate that a signed Trademark License Agreement is in place prior to awarding certification. A trademark is used in commercial certification of both products and practices to facilitate self-policing by the market. A certified product or practice can be associated with a trademark to provide ready and easy indication of its certification status. It is envisaged that at the lower four layers, the Navy may be able to make use of these preexisting COTS certification programs as an entry criteria for its certification requirements. In the case of certification at the top three layers of

the model, we suspect that the certification activity is bespoke and specific to warship certification needs. The team considers that use of a trademark will not be appropriate or valuable for these higher layers. However the legal requirements for continued conformity would be extremely valuable to the Navy and should be enshrined in the certification agreement at all layers of the model.

Organization signs Trademark License Agreement

As stated this is applicable to preexisting COTS certification programs, which may apply to the Navy's needs, but a Navy specific trademark does not seem appropriate.

CA certifies practice and enters it into Certification Register

As stated this is applicable to preexisting COTS certification programs, which may apply to the Navy's needs, but a Navy-specific Certification Register may not be appropriate.

Certification Register, Certificate, Logo

This is applicable to preexisting COTS certification programs, which may apply to the Navy's needs.

11. Position Paper on Resource Management in High Assurance Systems

Douglas Wells, the Principal Investigator on this project, was invited to submit an abstract for a DARPA-sponsored Workshop on Java in Real-Time and Embedded Defense Applications, which was held on July 13. He submitted the following position statement, which was defended at the workshop[8]:

Managing Resources in Multilevel Secure and Safety Critical Systems via Middleware Schedulers

This paper proposes the use of middleware schedulers in Java and CORBA environments to address the requirement in multilevel secure and safety critical systems for high degrees of predictability in order to reduce timing covert channels and high assurance. These schedulers are organized within an overall hierarchy that allows coordinated scheduling across an extensive distributed system.

There is an emerging trend for computing platform standards to incorporate an explicit tasking, often with an explicit scheduler component and interfaces to that scheduler. The existence of standardized scheduling interfaces in Real-Time Specification for Java and Dynamic Scheduling for CORBA, and the emerging application level scheduler for POSIX offers an opportunity to improve the effectiveness of systems requiring resource partitioning, such as multilevel secure (MLS) and safety critical (SC) systems.

It is generally recognized that uncoordinated scheduling of interacting resources leads to sharing conflicts and reduced performance. As more defense applications utilize distributed resources, it is increasingly necessary to utilize a platform-level or even

[8] The full presentation is available from the DARPA support contractor's web site, which is password protected. Please contract the Principal Investigator for information on accessing this site.

system-level scheduling model. Several of these standards also allow control over resources beyond the traditional CPU control. Thus, these schedulers can be used to apply system resources to the most important system applications (explicitly and in a coordinated manner according to the current battlefield environment.

Use of shared resources in MLS systems can lead to covert timing channels, which are usually alleviated via the use of a portioning kernel. We have previously shown that a dynamic scheduler can provide increased effectiveness while limiting the bandwidth of timing channels. Thus, we propose to extend this concept to an MLS-aware scheduler that utilizes middleware schedulers to adaptively apply more resources in support of the immediately most important applications. This scheduler would be hierarchical and only a simple basic resource scheduler must be included in the security kernel.

Similar issues occur in SC systems where current high assurance methodology requires strong scheduling determinism, which conflicts with the emerging trend for adaptive resources, such as power-aware commercial CPUs and distributed computing platforms. We propose that middleware schedulers can be developed to constrain the behavior of systems resources to behavior modes that can be analyzed for use in certified systems.

12. Submitter Bibliography

[AFRS95] J. Goldberg, L. Gong, I. Greenberg, R. Clark, E. D. Jensen, K. Kim, D. Wells, Adaptive Fault-Resistant Systems, Technical Report CSL-95-02, SRI International, Menlo Park, CA, January, 1995.

[CSAC93] R. K. Clark, D. M. Wells, E. D. Jensen, T. F. Lunt, P. G. Neumann, I. B. Greenberg, and P. K. Boucher, Effects of Multilevel Security on Real-Time Applications, Proceedings of the Ninth Annual Computer Security Applications Conference, Orlando, FL, December 6-10, 1993.

[Dual-Use94] D. Wells, A Trusted, Scalable, Real-Time Operating System Environment, 1994 Dual-Use Technologies and Applications Conference Proceedings, Utica, NY, May 23-26, 1994.

[ICCRTS98] T. Lawrence, P. Hurley, T. Wheeler, A, Kanevsky, J. Maurer, P. Wallace, D. Wells. R. Clark, Quality of Service for AWACS Tracking, Fourth International Command and Control Research and Technology Symposium, Stockholm, Sweden, September 14-16, 1998.

[ISORC02] A. Wellings, R. Clark, D. Jensen, D. Wells, A Framework for Integrating the Real-Time Specification for Java and Java's Remote Method Invocation, Fifth IEEE International Symposium on Object-oriented Real-time distributed Computing, Washington, D.C., April 29-May 1, 2002.

[SPAWAR02a] J. Drummond, D. Wells, M. Rahman, Detecting Failures within Distributed Environments, SPAWAR Technical Report 1884, U.S. Navy SPAWAR Systems Center, San Diego, CA, April, 2002. J. Drummond, L. Coker, J. Marcelino, D. Wells, J. Carroll,

[SPAWAR02b] J. Drummond, L. Coker, J. Marcelino, D. Wells, J. Carroll, M. Rahman, A. Vadlamudi, M. Srivastava, Managing Quality of Service within Distributed Environments,

SPAWAR Technical Report 1883, U.S. Navy SPAWAR Systems Center, San Diego, CA, April, 2002.

[TAO01] D. Wells, Fault Management Based on Quality of Service Criteria, First Workshop on The ACE ORB (TAO), Washington University, St. Louis, MO, August 5-6, 2001.

[WDMS02] D. M. Wells, R. E. Bernstein, A. Vadlamudi, Testability of Complex, Middleware-Based Systems, Supplemental Volume of the 2002 International Conference on Dependable Systems & Networks, Washington, DC, June 23-26, 2002.

[WORDS96] D. Wells, Using Object Frameworks to Enable Real-Time and Dependability in a Modular Operating System, Second Workshop on Object-Oriented Real-Time Dependable Systems, Laguna Beach, CA, February 1-2, 1996.

[WPDRTS99] R. Clark, E. D. Jensen, A. Kanevsky, J. Maurer, P. Wallace, T. Wheeler, Y. Zhang, D. Wells, T. Lawrence, P. Hurley, An Adaptive, Distributed Airborne Tracking System, Seventh International Workshop on Parallel and Distributed Real-Time Systems, San Juan, Puerto Rico, April 12-13, 1999.

— end —