Department of Defense

High Performance Computing Modernization Program Information Environment Proposal

In Response to HPCMP Solicitation
4TS-TT-01-0001

Technical/Management Content
Volume I

High Performance Technologies, Inc.
4121 Wilson Boulevard, Suite 200
Arlington, VA 22203
DUNS: 784366544
March 7, 2001

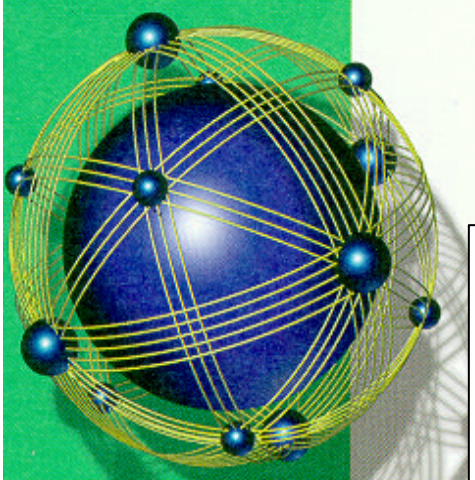

HPCMP
High Performance Computing Modernization Program

March 7, 2001

Aeronautical Systems Center (ASC)
Major Shared Resource Center (MSRC)
Attention: Charlotte Coleman, ASC/HPTI
2435 Fifth Street
WPAFB, OH 45433-7802

Dear Ms. Coleman,

High Performance Technologies, Incorporated (HPTi) teamed with the Northwest Alliance for Computational Science and Engineering (NACSE) is pleased to submit the attached proposal in response to DoD HPCMP Contract 4TS-TT-01-0001 Information Environment (IE). HPTi brings a solid history of support to the DoD HPCMP and looks forward to continuing our excellent support through this contract.

As a team, HPTi and NACSE are excited about the opportunity to further serve the DoD HPCMP mission. If you need any additional information concerning this proposal, please contact Mr. Tim Keenan at (703) 682-5322 regarding technical issues and Mr. Bill Edmunds at (703) 682-5329 for matters of a contractual nature.

Sincerely,

Mr. Tim Keenan
Partner

Table of Contents

High Performance Technologies, Inc. (HPTi), and the Northwest Alliance for Computational Science and Engineering (NACSE), have joined together to form a synergistic and dynamic team that utilizes the strengths of both. This team capitalizes on NACSE's expertise in web-based, distributed software development and usability engineering, and HPTi's expertise in program execution, integration and solution deployment.

In the proposal that follows, we provide the details of our approach for the development and deployment of an Information Environment that will satisfy the needs of the High Performance Computing Modernization Program. Three of the tools (Allocation Reporting, Queue & Process Status, and Allocation Management) have already been developed for NAVO by NACSE. Additional pieces have been partially developed by PET at ARL (under HPTi management), including user fill-in, queue process and status tools for GRD, and Kerberos authentication through a web browser. This proposal intends to leverage this existing groundwork by enhancing and expanding these development efforts and providing a uniform system and architecture framework in order to extend the tools to the entire HPCMP community.

The HPTi and NACSE team recognize the challenges facing the Major Shared Resource Centers (MSRC). The proposed Information Environment will turn those challenges into solutions through the development and deployment of a system that will provide increased efficiency in the allocation and utilization of computational resources, reduce the time spent on managerial tasks, increase the efficiency of managerial processes, and satisfy an increasingly diverse customer base.

1.0 Technical Approach

HPTi and NACSE propose to develop a HPCMP Information Environment (HPC-IE) that provides users with a secure web-based portal to allocation and utilization data, queue and process status data, and account management data, all through a consistent and intuitive user interface that meet the users' needs.

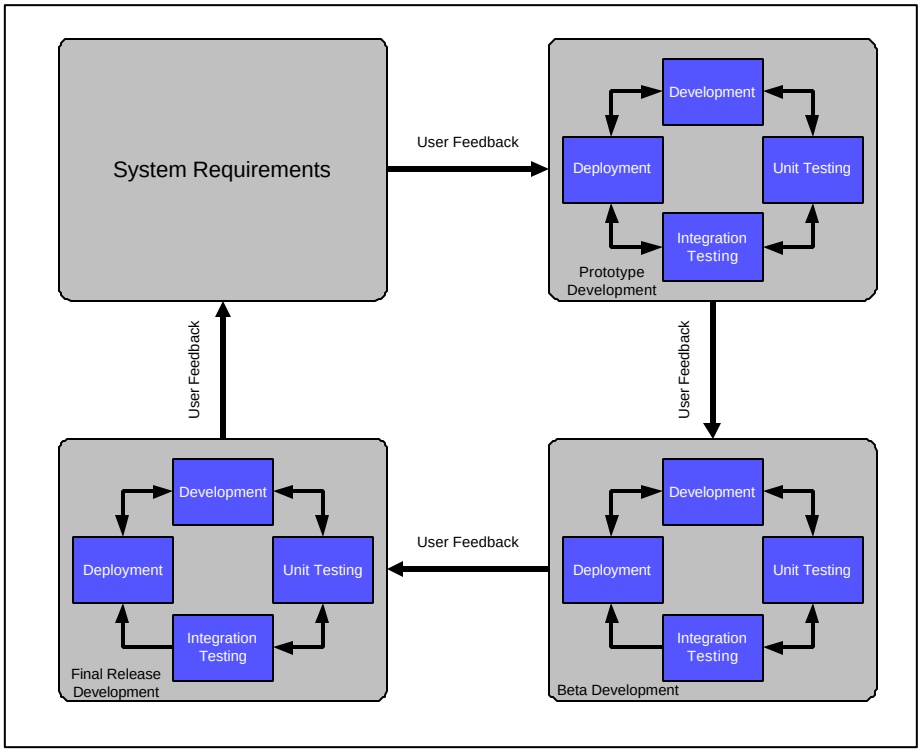

Figure 1: Development Approach

We will use an incremental development approach. Through a series of major milestone releases we will provide a prototype of the system, then continually build upon the system through incremental releases, improving the functionality and performance of previous releases based on feedback from system users. After a final release of the software, the cycle can begin again as the systems evolve to meet changing enterprise environments, expanding user needs, and technological advances. Our architecture is designed to support this continuing expansion and enhancement of the system without starting a new development cycle.

1.1 Overview of System Architecture

The proposed system separates the components of the HPCMP Information Environment into three functional areas: a Control Architecture, a Data Architecture, and a User Application Layer (Figure 2). The HPC-IE's Control Architecture manages the overall functioning of the system and the flow of data. The Data Architecture provides a scalable, flexible mechanism for dealing with the distributed and heterogeneous nature of HPCMP site data. The User Application Layer supports human access to the HPC-IE's functionality. This human access is accomplished via a series of Web-based user interfaces personalized to different user groups and an application architecture that can scale easily to additional types of functionality.

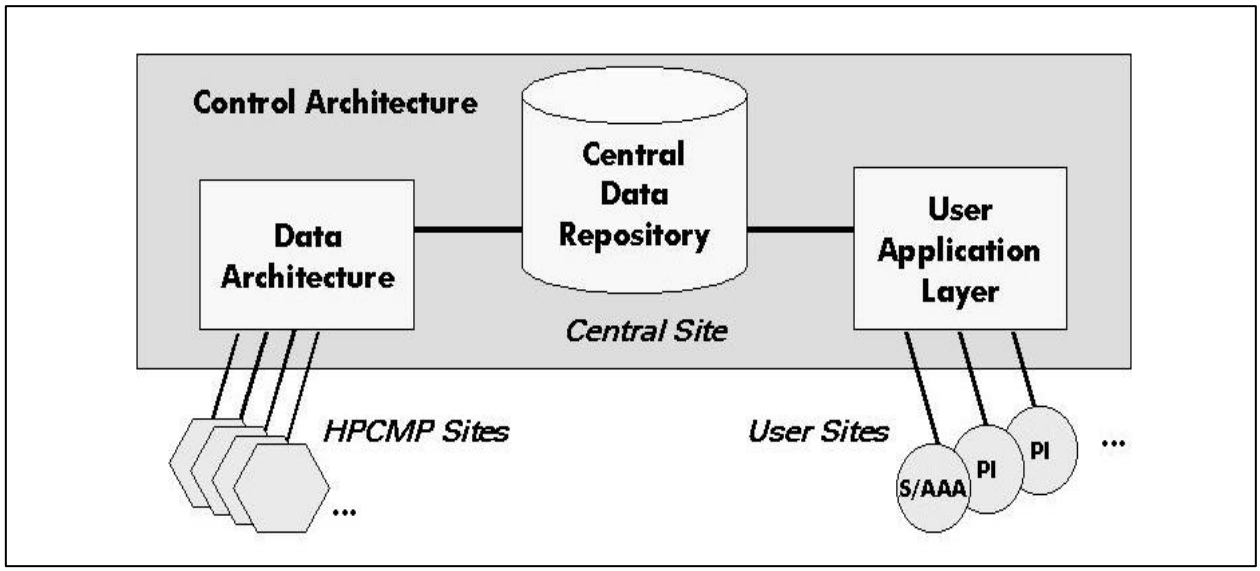

Figure 2: HPC-IE's System Architecture

There will be two separate, co-located IE servers, one for the data architecture and central database repository and another for the user application layer. This will allow for the easy addition of servers in the future as the demands on IE system increase. These servers will communicate securely over properly authenticated SSL-encrypted tunnels.

This architecture was chosen to maximize the flexibility and maintainability of the HPC-IE. Within each functional area, the architecture has been designed to support extensibility and

scalability, to support new types of information, to support new needs for user access, and to support new classes of users. By leveraging available off-the-shelf components, well-known networking protocols, and proven security mechanisms, the components will be robust and rapidly deployable.

1.2 HPC-IE Control Architecture

The Control Architecture (Figure 3) provides a functional substrate for the HPC-IE system. At the heart of the system is the Oracle 8i database management system (DBMS) providing a centralized repository for the operational data. The major components – central data repository (the IEDB), supervisor daemon, data collection tools, web server, and application server – are described individually below. They incorporate a number of features ensuring that the system is scalable to multiple HPCMP sites, as well as enhancing its overall maintainability. Security is maintained by Kerberos and SSL software, as discussed in a later section.

1.2.1 Central Data Repository (The

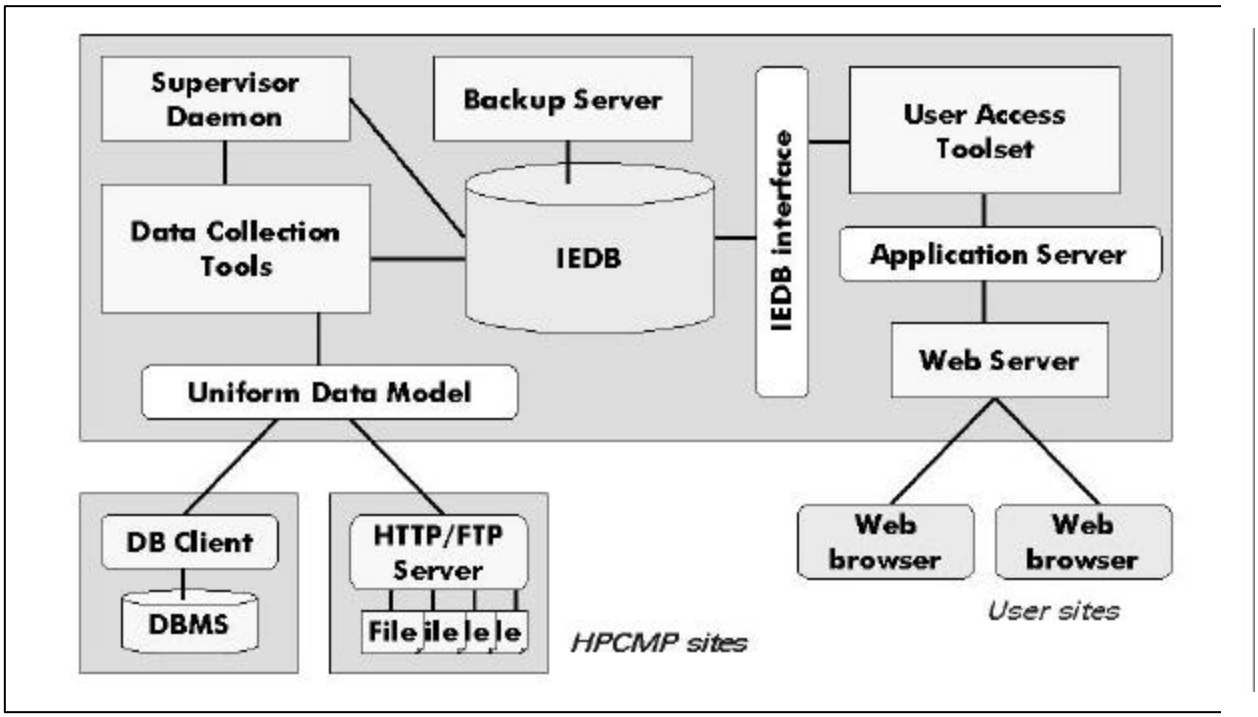

Figure 3: Control Architecture for the HPC-IE System

IEDB)

Central to the HPC-IE is a repository of HPCMP operational data that includes user and project data, resource allocation data, historical resource usage data, and metadata describing how the system interfaces with data sources at the individual HPCMP sites. Data will be maintained in a relational database, the IEDB (IE Database), which will be implemented using Oracle 8i.

Oracle was chosen as the IEDB DBMS for reasons of interoperability, security, and reliability. Oracle 8i is currently available to the HPCMO and to the MSRC's; it is therefore cost effective and it can be expected that MSRC staff is already familiar with the administration of Oracle databases. Uniform implementation of Oracle across the MSRC's provides better interoperability between the IE database and the databases at each individual site; even if this uniformity should exist only with the MSRC's, it will benefit the majority of the data transfer operations. Oracle has capabilities that will provide for many of the system requirements including support for large databases, Kerberos, PKI, and standby databases for backup purposes. Oracle, as a commercially supported database package, with a reputation for robustness and reliability in critical applications, is also well suited to 24x7 operations. A dedicated standby database will maximize uptime and protect against data loss. Database agents will automatically perform routine full and nightly incremental backups of the IEDB.

As the data are both sensitive and critical to operation of the entire HPC-IE, both data security and access security are important. Specific access security mechanisms are discussed in a separate section. However, a key element of our architecture is that neither users nor data providers will access the IEDB directly. Access to the IEDB – whether to insert, to update, or to

retrieve data – will be accomplished through the HPC-IE user toolset and the HPC-IE data collection tools.

1.2.2 Supervisor Daemon

The data management functions of the HPC-IE are overseen by a "supervisor" daemon. This operates in the background to perform basic support tasks such as:

- Initiation of periodic actions (e.g., data transfers, consistency checks)
- Self-checks to ensure the system is operating correctly (e.g., that necessary daemons are running)
- Self-correction where possible (e.g., restarting daemons)
- Resource detection, status, and availability reporting (e.g., monitoring site servers)
- Notification to HPC-IE administrators (via e-mail, with a pager gateway) of any problems requiring intervention

The supervisor daemon is designed to be configurable through metadata describing interactions with HPCMP sites and other information pertaining to HPC-IE operation. This structure provides flexibility in maintaining, upgrading, and expanding the HPC-IE system.

1.2.3 Data Collection Tools

Data from HPCMP sites are moved into the IEDB by a set of data collection tools. These processes are initiated at intervals specified in the site metadata. The use of metadata to control timing means that the periodicity of data harvesting may vary from site to site. The specific mechanisms used to harvest data are described in the "Data Architecture" section.

1.2.4 Web Server and Application Server

The Web server provides a single access point through which users gain access to HPC-IE functionality and data. It is responsible for initiating the authentication of HPC-IE users. Those

mechanisms are discussed under the "Security Section." Web services include a collection of user interfaces to individual user access tools; the interfaces share a common look-and-feel and provide consistent transitioning from one type of HPC-IE functionality to another. Each tool has multiple interfaces, which serve to tailor a varying level of tool functionality for different user groups (e.g., providing different user interfaces and access paths for S/AAAs than for PIs).

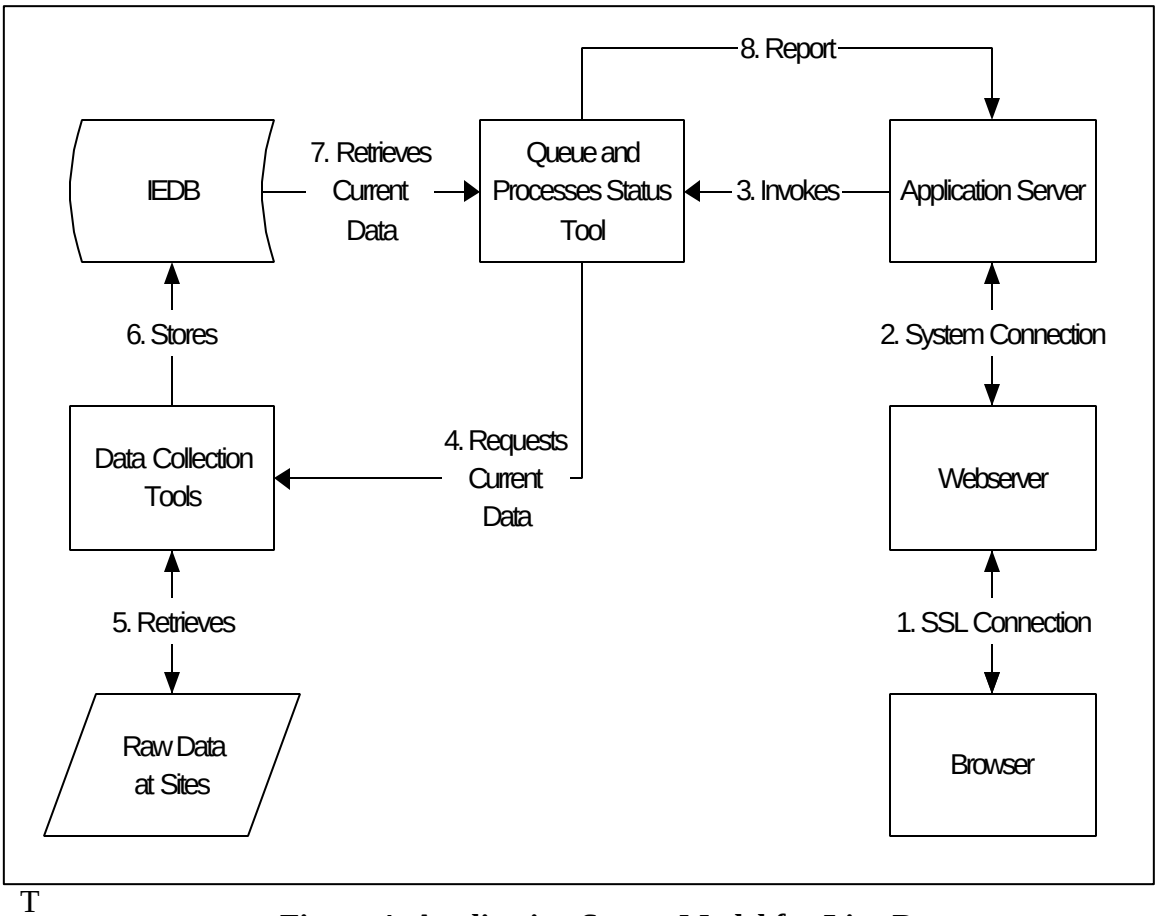

T

Figure 4: Application Server Model for Live Data

he

Web Server will be a robust, well-supported, secure product such as Apache with the mod_ssl extension. The Application Server will be developed by a collection of widely available, commercial languages and tools such as Practical Extraction and Report Language (Perl) and Database Interface (DBI). By using a widely available and supported commercial product, the

HPCMP staff will be able to easily support, customize, and if required later, to extend this system after deployment. These tools are open source and open standard tools and therefore are widely available and support numerous platforms. This is an ideal combination for a distributed, heterogeneous environment and will make the IE system portable to any architecture and database.

The primary component of the Web Server will be developed using Perl with the DBI module that will be used to access to the IEDB. The Perl DBI interface has built in driver support for numerous database systems including many commonly known databases, such as Oracle, Access and SQL Server.

The Application Server controls the invocation of the tools themselves. To launch a particular user access tool, the Web Server communicates the request with the Application Server, and the Application Server invokes the tool. Figure 4 shows the details of one example the Web Server and Application Server interaction. While such layering is not strictly necessary, we chose to separate the web and toolset functionality in order to provide flexibility and extensibility (e.g., adding new tools in response to future user requirements).

Figure 4 illustrates the Application Server model for the Queue and Process Status Tool (Tool 2). This tool is unique in the system since it is the single component that will require near-real-time data. The other tools in the IE system will require periodic data retrieval. These tools are simpler since steps 4 through 6 will be accomplished using a separate process to retrieve the data and store it in the database. The remaining tools will directly access the database for their relevant data.

1.3 HPC-IE Data Architecture

The Data Architecture manages interactions between remote HPCMP sites and the various system components that must communicate with them. These include the tools that harvest data for the IEDB or operate on the remote data in order to update information. This section also describes the Uniform Data Model that ensures compatibility throughout the HPC-IE, as well as the measures taken to enforce data consistency in cases where data will be added manually.

1.3.1 Uniform Data Model

The design of a comprehensive information environment operating across distributed, heterogeneous data repositories requires some method of providing consistent access to those repositories. The participating HPCMP sites may maintain operational data using any number of various types of DBMS; they have also been given the option of maintaining the information in a hierarchy of simple, “flat” files.

Central to the Data Architecture there will be a staging area, implemented as a separate table in the IE database, where the sites may push appropriate data. This is the preferred model, as it allows for the most up-to-date information to be readily accessible to user tools and provides several other benefits. First, it reduces the burden for DC's who do not wish to maintain data on their own servers. Second, by providing one central location rather than having data collected from a site server to the IEDB, the chance for synchronization differences to occur will be greatly reduced. Lastly, the security of the individual sites is maximized. Since the data is being pushed out of the sites, rather than pulled from them, the critical point of security is on the IE system. This consolidates the security points from the participating sites to just the IE servers.

In order to achieve flexibility for the autonomous sites and scalability to large numbers of sites, we will employ a Uniform Data Model to provide a consistent, clearly documented interface to the target data repositories. This interface will be implemented as a library and used

by those HPC-IE components that require access to site data (i.e., the data collection tools and the allocation management tool).

A database of *site metadata* will define how the Uniform Data Model interacts with each site, and also define any transforms that must be performed on data from particular sites. The use of metadata, accessed dynamically at the time of data collection, to identify these idiosyncrasies will further enhance the scalability and maintainability of the Uniform Data Model.

1.3.2 Data Collection Mechanisms

In order for the Uniform Data Model to provide access to operational data from HPCMP sites with heterogeneous data repositories, it must implement means of collecting data from those sites. In doing so, it must strike a balance between the needs of individual sites and the need for a scalable, easily maintainable coordinating system. It is clearly desirable to allow participating sites as much flexibility as possible to maintain their own data in the schemes of their choice. A common denominator is required, however, in terms of the interfaces for ingesting site data. The HPC-IE will support two distinct mechanisms for harvesting data from participating HPCMP sites and ingesting the data into the IEDB:

1. Access to data stored at a given HPCMP site using commercial off-the-shelf DBMS software (i.e., the HPC-IE acts as a database client).
2. Access to XML or CSV flat files via HTTP or FTP (i.e., the HPC-IE acts as a transfer agent), two well-known protocols with robust implementations for both servers and clients.

For the site's internal use, the first mechanism is clearly superior, as it allows sites to develop a variety of data views that reflect their particular needs. Therefore, we will encourage sites to develop DBMS solutions and will provide guidance to assist them.

For both types of access, the preferred paradigm is that the HPC-IE "pull" data from the remote site at periodic intervals (specified in the site metadata). This requires that each site

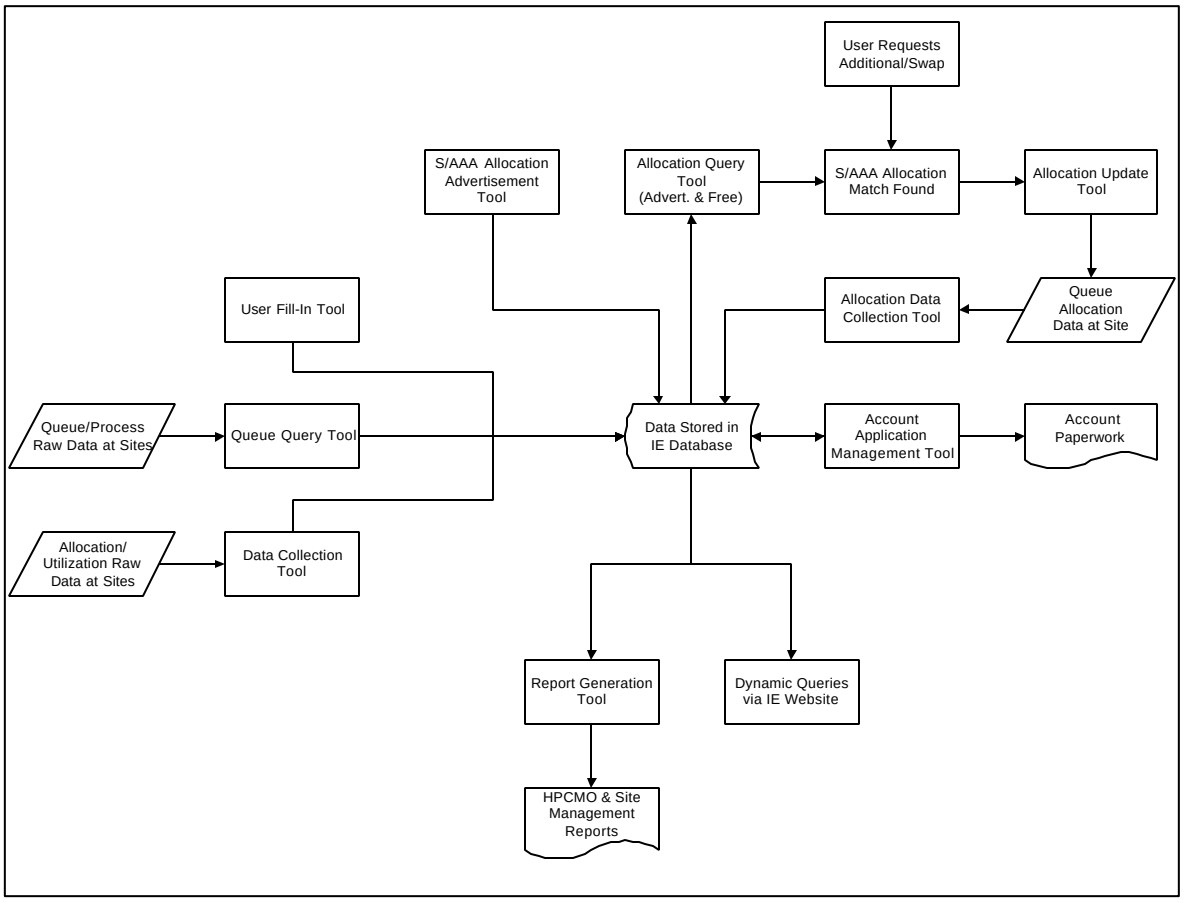

Figure 5: Data Flow Diagram

maintain servers with the capability of providing the necessary data on demand. The security implications of this approach are discussed below.

In the event sites are unable to, or do not wish to, maintain servers that can respond to pull mechanisms, we will provide an interface for allowing the sites to "push" data into the system. Direct database access cannot be allowed because of the security concerns in permitting remote access to the IEDB. Instead, remote sites will have to upload their XML or CSV flat files to an HTTP or FTP server at the central site, from which HPC-IE would harvest the data. Another

method is to use a database-to-database connection to the IE staging area to transfer the data for later importing into the IEDB.

Data collection will support the various formats at the sites via a set of conversion filters. XML will be the core non-database format utilized by IE. To support the various other formats, libraries will be used that will provide for conversion to and from XML.

The Uniform Data Model enhances scalability by providing a single layer at which data access mechanisms may be added. Since the mechanism(s) used by each participating HPCMP site are recorded in the site metadata, changes in access pattern are accommodated automatically.

Tools will be developed to import legacy data from participating sites into the IE system via asynchronous use of the Uniform Database Model. These data will be used as sample data during the design and testing of the system.

1.3.3 Special Handling of Queue and Process Status Data

Since the reporting of queue and process status will require access to near-real-time information from the HPCMP sites, these data will be retrieved on-demand as the queue/process status tool is activated (see Figure 4). In order to support this requirement, participating HPC sites must host servers that can provide the most recent data available. Services will be developed for the architecture, as the mechanisms for accessing queue and process data vary from machine to machine. The security considerations of these servers are discussed in a later section.

These services will routinely generate queue and process status data that will be written to the HPC site host server. This data will be pulled from the site servers and stored in IE as needed, namely when queue and/or process status is requested by the IE user. The reason for storing the data in IE is for scalability purposes and security. By storing the data in the IEDB,

subsequent requests that occur within a given time period will be able to access this data rather than repeatedly pulling the information from the sites. This will reduce the burden on the individual sites, as well as reducing the number of security mechanisms by utilizing the server pull method rather than creating an additional scenario. In addition, the services will process the data and create a less detailed version for the historical archive, which will then be pushed or pulled into IE system on a regular basis.

1.3.4 Special Mechanisms for Updating Site Allocations Data

In addition to retrieving data acquired from participating HPCMP sites, the allocation management tool must be capable of initiating the update of site allocation data. (This is necessary to maintain consistency between local site data and the central database as resource allocations are moved or modified.) The allocation management tool will initiate the updates via one of the two access mechanisms described earlier. In either case, the data transfer will be pushed from the HPC-IE, since sites will have no way of knowing when data should be pulled. Because they will communicate with remote sites via the Uniform Data Model, the selection of the mechanism will be made dynamically according to the information in the site metadata database, making it possible for sites to evolve their structures from flat files to databases.

Due to the geographical dispersion of the centers and the wide variance of time zones, allocation data is required to be in reserve by the S/AAA prior to the advertisement. When a sufficient match has been determined and agreed upon by the interested S/AAA's, all parties will be notified of the change. Where possible, the allocations will be updated by the IE system. Otherwise, notification will be given to the S/AAA's to perform the adjustments. In order to accommodate the differences in time zones, exchanges will be removed from reserve and be made available for use at a consistent local time for each center, as directed by the HPCMO.

Thus, some portions of the exchanges may be available before others. This will provide for synchronization of data between the sites and the IE system at any given time.

1.3.5 Maintaining Data Consistency

Any type of data repository is prone to errors. This is particularly an issue when data are entered manually. The incorporation of consistency constraints into a database schema provides data managers with an important tool for finding and correcting those errors. In developing the IEDB, we will identify and implement a collection of appropriate constraints and crosschecks. Consistency-checking tools launched by the Supervisor Daemon will report any discrepancies, both in data collected from participating HPCMP sites and data originating in the HPC-IE system.

1.4 User Application Layer

The User Application Layer provides the user with a variety of tools supporting the access functionality of HPC-IE. They include the five tools outlined in the RFP: allocation reporting tool, queue and process status tool, allocation management tool, user fill-in tool, and account application management tool. The tools will be invoked via web interfaces accessed from the HPC-IE home page. Users will be authenticated by a common mechanism prior to the invocation of any tool; the tool will then serve as a proxy for the user, interacting with the IEDB through a common interface. This layering was chosen to isolate any idiosyncrasies associated with the IEDB implementation and to facilitate future modifications or extensions, in particular changes to the security architecture. The user interfaces will share a common look-and-feel to facilitate usability and to improve efficiency in use.

Various user groups will be defined by access privileges to the system, and can be easily extended to accommodate the needs of the program as they arise. Users with valid HPCMP

Kerberos principals will hold the most basic privileges to view information such as queue status, allocation reports, and also have access to the user fill-in tool in order to begin the process to obtain/renew accounts, as well as other user requests. Other groups will include those persons specifically listed in the IE access lists. They will be listed since they require additional privileges, and this listing will specify to what extent their privilege extends. These groups include, but are not limited to, S/AAAs, Site and Program Management, and Principle Investigators. The distinctions between groups will contribute to the user interface in order to accommodate the varying functionality required based upon the groups differing tasks.

1.4.1 Allocation Reporting Tool ("Tool 1")

This will provide different functionality to support the needs of different users (S/AAAs, PIs, system administrators, and HPCMO). As part of the user authentication process, the Application Server will identify the class of user based on his/her account information stored in the IEDB. Only interfaces appropriate to the user's data access privileges and informational needs will be presented. The user will be able to generate personalized reports based on any combination of selection parameters (user, project, machine, site, etc.), for any given time period (day, week, month, year, to date, etc.), and by any given time increment (hour, day, week, etc). This will be done using the allocation reporting tool portion of the user access toolkit to query the IEDB. The resulting data will be display in various formats including tabular, summary, and several types of charts. In addition to viewing and printing these reports through the web browser, the user will have the option to download the displayed data in a format suitable for importing into common spreadsheet or database software. This will be accomplished by utilizing the previously described data conversation filters and selecting the format the user would like to receive the data in (XML, CSV, etc.). Mechanisms will also be created by which the IE administrator can

configure regular data reports to be automatically downloaded to defined site servers and/or email addresses, without requiring any user intervention. The selection criteria for this will be controlled by metadata and is customizable to each site or user. Data for this tool will be loaded into the IEDB by the data collection tool on a regular basis defined by the site metadata, the frequency of which will be determined by usage requirements, but will occur at a minimum of once a day.

1.4.2 Queue and Process Status Tool ("Tool 2")

This tool will provide graphical and textual summaries indicating current activity at the participating HPCMP sites. It will also be possible for the user to define personalized views showing just those resources of interest to him or her. In addition, the information can be viewed across the sites to provide a more comprehensive picture of HPCMP activity. Limited data for historical purpose will be stored in the database; otherwise the data will be frequently accessed from the site servers. Special mechanisms to accomplish this have been previously described.

1.4.3 Allocation Management Tool ("Tool 3")

Different users will be presented with different interface functionality, based on the user's access privileges and needs. Because significant portions of the functionality overlap with the reporting tool, many interface components will be shared across the two tools; this will improve both tool maintainability and the consistency of the user interface. This system will allow S/AAAs to advertise any allocations freely available or for exchanging. Automated IE agents will attempt to make matches to existing advertisements when new ones are posted, and then notify the appropriate S/AAAs if it finds any that may be suitable. Once advertised these allocations can be searched based on hours, site, platform, and classification (challenge, unclassified, classified) to assist in the matchmaking process. Mechanisms for updating the

allocations when a match has been successfully completed have been previously described. Exchanges will be logged via the database by updating special tables for this purpose. These tables will be viewable through the user interfaces for those that have the necessary privileges, and will also be to be queried for varying levels of detail.

1.4.4 User Fill-in Tool ("Tool 4")

This tool will be the primary source of user-initiated input to the HPC-IE. As such, most of its structure will be devoted to enforcing the consistency and correctness constraints necessary to minimize the introduction of inadvertent errors. This tool and the account application management tool will share the same underlying data. After this information has initially been submitted, the user will be able to update it to make corrections as needed. Changes to the user information will be logged via the database by updating special tables for this purpose. These tables will be viewable through the user interfaces for those that have the necessary privileges. Upon submission of initial information and changes, the Principal Investigators, S/AAAs, and relevant center staff will be notified.

1.4.5 Account Application Management Tool ("Tool 5")

This will allow S/AAAs (and perhaps system administrators) to manage user accounts, but no interface will be available to other user classes. For areas of overlapping functionality, it will share interface components with the user fill-in tool. The interface components for this tool will depend greatly upon the user group that the IE user is in. S/AAAs, Principal Investigators, and center staff will have different interfaces to support the different usage requirements. This tool will operate on the data that has been input from the user fill-in tool. Additional data such as the "S/AAA only" information will be added via these tools to supplement the data acquired from the user fill-in tool for any account management purposes. Available resources will also be able

to be assigned to users and projects through this interface. Automated agents will also be developed that will notify appropriate individuals upon certain conditions. This will include tasks such as notifying the user and S/AAA when it is time to renew accounts, reminders when the user account will expire, and when all allocations have been used.

1.5 Usability Engineering

Consistent, well-engineered user interfaces will be provided for the IE tools. Established principles of usability engineering will be applied to maximize four aspects of IE system usage: ease-of-learning, ease-of-use, efficiency, and productivity. NACSE has significant experience in applying these principles to Web-based tool interfaces for the HPCMP PET program, including the queue status browser and resource allocations reporting and exchange tools that gave rise to the present RFP. NACSE has a proven track record in working with representative users in order to elicit what they want in tool interfaces. NACSE's development of the queue browser tool, and of the resource accounting database and exchange tools, began with extensive sessions of user meetings, interviews and surveys. Listening to what users want, rather than implementing tools that computer scientists want to build, ensures tools that meet user needs. We will focus on fast, reliable interfaces without unnecessary features. Since no special features are required this will provide the benefit of being viewable through SSL capable browsers and on all platforms.

1.6 Security

In a distributed information system, security is critical for protecting data from improper access, corruption, or loss. Since the HPC-IE is designed to be Web-accessible by a broad community of users, mechanisms such as firewalls or virtual private networks are not optimal (they are likely to have serious impact on system availability). As an alternative, we have chosen to protect the system at its points of access. There are three points of concern:

1. The Web Server: the web server will utilize plug-in Kerberos and SSL modules for authentication, authorization, and encryption. Users accessing the web server will be required to authenticate with Kerberos (using SecurID pre-authentication). Unfortunately, no commercial web browser supports Kerberos authentication directly, and implementing this capability is beyond the scope of the IE proposal, and is not cost effective. In IE system users will be required to enter their Kerberos and SecurID passwords in an SSL-encrypted browser session. Traffic between the browser and

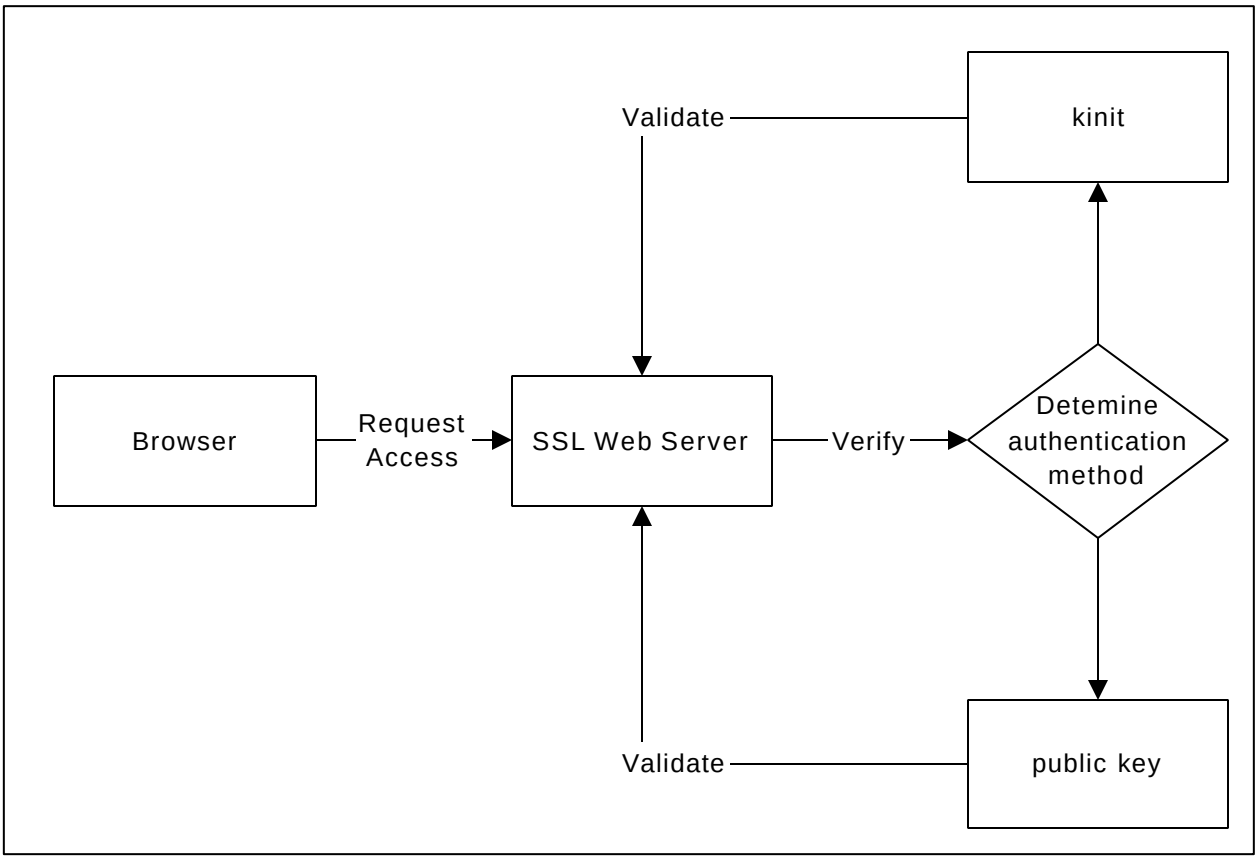

Figure 6: Authentication Mechanism

webserver will be encrypted for the entire session, and will be directly tied to the particular Kerberos principal.

2. HPC-IE servers to which HPCMP sites push data: these servers will be secured by requiring HPCMP site servers to use Kerberos authentication – not because clients can access sensitive data, but because the legitimacy of incoming data must be verifiable. Database servers will limit direct access for this type of transfer, as previously described by the staging area. This minimizes the potential for data corruption by even well meaning remote clients. Transmissions between HPCMP site servers and the IE system will be encrypted.
3. The HPCMP site servers from which the HPC-IE pulls data: Kerberos authentication will be used to access the database, ftp, and http servers. The HPC-IE tools accessing these servers will use Kerberos credentials. The queue and process status tools must use this method since the push model cannot adequately serve the data, for the reasons outlined previously. The transmissions between HPCMP site servers and the IE system will be encrypted.

Other points of potential vulnerability are the IEDB database server itself, and the HPC-IE's Application Server. However, these are accessed only by other components of the HPC-IE system, and not by remote clients. Either Kerberos or PKI authentication over SSL-encrypted tunnels will be used between the various IE servers.

The IE system will be developed using the General Security Services API (GSSAPI) and other open standard API's where applicable to facilitate interoperability with various platforms and authentication systems. As a result the user and system authentication can easily accommodate a PKI capability.

User authentication will occur through a secure web interface. This interface will authenticate the user by passing required information to Kerberos for verification. Once the user

has been approved access, the Kerberos credentials will be stored and utilized for continuous transparent verification for the duration of the session or until expiration of the credentials. This method for secure web access has already been approved by HPCMO security teams and implemented at ARL. PKI authentication will be supported by automatic detection and verification of a users certificate and then bypassing this interface.

1.7 Post Deployment Technical Support

There are two categories of problems that may arise requiring technical support, software errors and miscellaneous errors. Technical support will be provided for the former by establishing a single point of contact. This individual will be available at any time for a period of six months after deployment. Documented procedures will be established for various problems and issues that the HPCMP IE administrator should follow before contacting this individual. These procedures will provide for contingency plans for hardware failure, security issues, etc. Following these procedures will allow for rapid resolution by providing the necessary information in order to diagnose and resolve any problems. Some of these procedures may also reduce the severity of the problem, such as activating the backup servers. Problems will be reported to this person through the HPCMP IE administrator(s). This individual will then verify the issue is a software problem, make the necessary arrangements required to resolve this problem within a suitable timeframe based upon the severity of the problem. This timeframe will be within 24 hours for critical problems and within 72 hours for non-critical problems.

The miscellaneous errors include such problems as software upgrades (Oracle, operating system, etc), hardware failures, unexpected site changes to data file formats and locations, and site servers and network downtime. These problems result in system downtime but are not a result of software development. The HPCMP IE Administrator will be responsible for resolving

this class of problems, as they are largely out of our control. However, the contractor will provide assistance with these errors when appropriate.

1.8 On-line Help

The system will provide on-line help accessible through the web browser, and will be available to the user at any screen throughout the IE interface. This on-line help will include thorough documentation for the user, IE administrator, and site data administrators.

2.0 Management Plan

Recognizing the inherent strengths of the teammates, the project is organized into two teams: the Development Team, led by NACSE personnel with their expertise in web-based, distributed software development and usability engineering, and the Integration and Deployment Team, led by HPTi personnel with their expertise in program execution, integration and solution deployment.

The Software Development team is responsible for the development of the system architecture, development of the component applications and their interfaces, and the unit level testing.

The Integration and Deployment team is responsible for the overall system integration, deployment, system level testing, and installation of the component applications, as well as the installation of interfaces between the software components and the data. Additionally, the Integration and Deployment team is responsible for post deployment support.

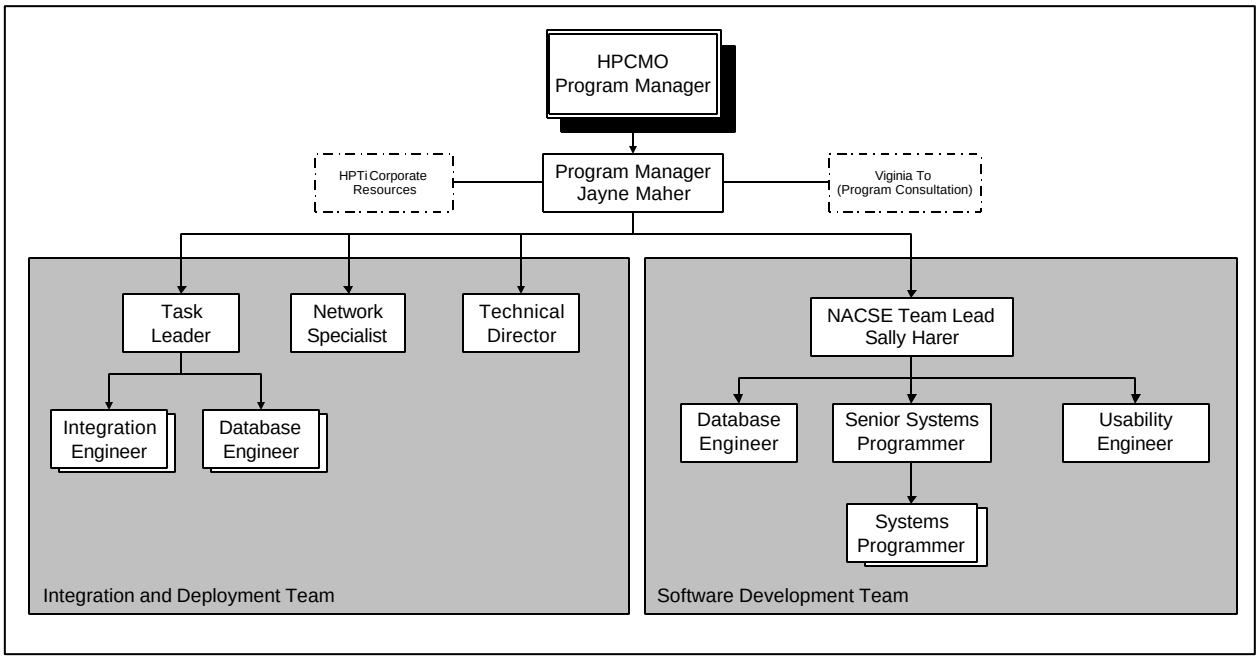

Figure 7: Project Organization

2.1 Implementation Plan

The Task Plan Schematic (Figure 8) shows the details of the development, testing, implementation, and deployment of the system, as well as where risk analysis tasks will occur in the process.

2.1.1 Testing Approach

As software is developed by the Software Development team, it is tested and integrated in a "Build, Test, Integrate" approach. This is reflected in the length of the testing and integration tasks for each release of the software. After unit testing is performed, the software is released to the Integration and Deployment team. The Integration and Deployment team will conduct independent integration testing at the Integration Testing Environment. This testing environment will be configured to replicate the target deployment environment. At the end of the project the Integration Testing Environment, less the licensed Oracle 8i DBMS, will be offered to the government to continue support of the system after final deployment.

When the system components are ready for a release, (prototype, beta, final), the Integration and Deployment teams will travel to identified locations and install the necessary system components for that location. Once the components are installed the team will conduct site acceptance testing and ensure that the installed components meet the acceptance criteria. The following paragraphs cover the major tasks of the implementation plan is detail.

2.1.2 Develop Detailed Project Plans – Task Plan Schematic 1.0

In this initial step of the project, detailed plans for project execution are developed in conjunction with the customer. These plans cover the detailed tasks that will be performed as part of the project, and how project controls will be established. Project controls are developed as appropriate for the scope and complexity of the project. However, as a minimum software

and documentation configuration, and cost controls are established immediately at the start of the project.

During this phase of implementation, schedule and other project risks are identified and mitigation plans are developed to offset those risks as required.

2.1.3 Design System Architecture – Task Plan Schematic 2.0

The second step of the project is to develop detailed system architecture. The system architecture ensures that the system components maintain a consistent interface with the user and that the system components meet user system performance expectations.

At this time system and deployment site infrastructure requirements are identified, as well as required hardware and commercial software components that will be provided by the government to support the system architecture.

2.1.4 Develop Prototype Applications – Task Plan Schematic 3.0

Concurrent with the development of the system architecture, the prototype component applications are developed. Component performance, functionality, and interface requirements are identified and incorporated into the system design. These prototype applications will then be developed, tested, and integrated by the Software Development team and handed off to the Integration and Deployment team. The Integration and Deployment team will perform system level testing on the release then deploy the software to the sites identified by the government.

2.1.5 Develop Beta Applications – Task Plan Schematic 4.0

From the prototype applications, the beta applications will be developed. These beta applications will enhance the prototype applications, and the enhancements to the applications will be based on the feedback received from the user communities. Once the beta applications

have been developed and undergone integration testing and system level testing, they will be deployed as upgrades to the existing system locations.

2.1.6 Develop Final Release Applications – Task Plan Schematic 5.0

The development of the component applications continues with the development of the final release applications. These applications are continued improvements of the initial prototype and beta release versions of the user applications. Based on continuing user feedback, improvements to the user interfaces application usability will be made. The final release applications will be undergo integration and system level testing prior to the execution of the final step in the development process.

2.1.7 Install and Integrate Full Release System – Task Plan Schematic 6.0

The final release applications are culmination of the development activities. The final release applications will undergo complete component and system testing and integration prior to final deployment. After the applications have completed the testing cycle they will be deployed to the system installation sites as final enhancements to the prototype and beta releases.

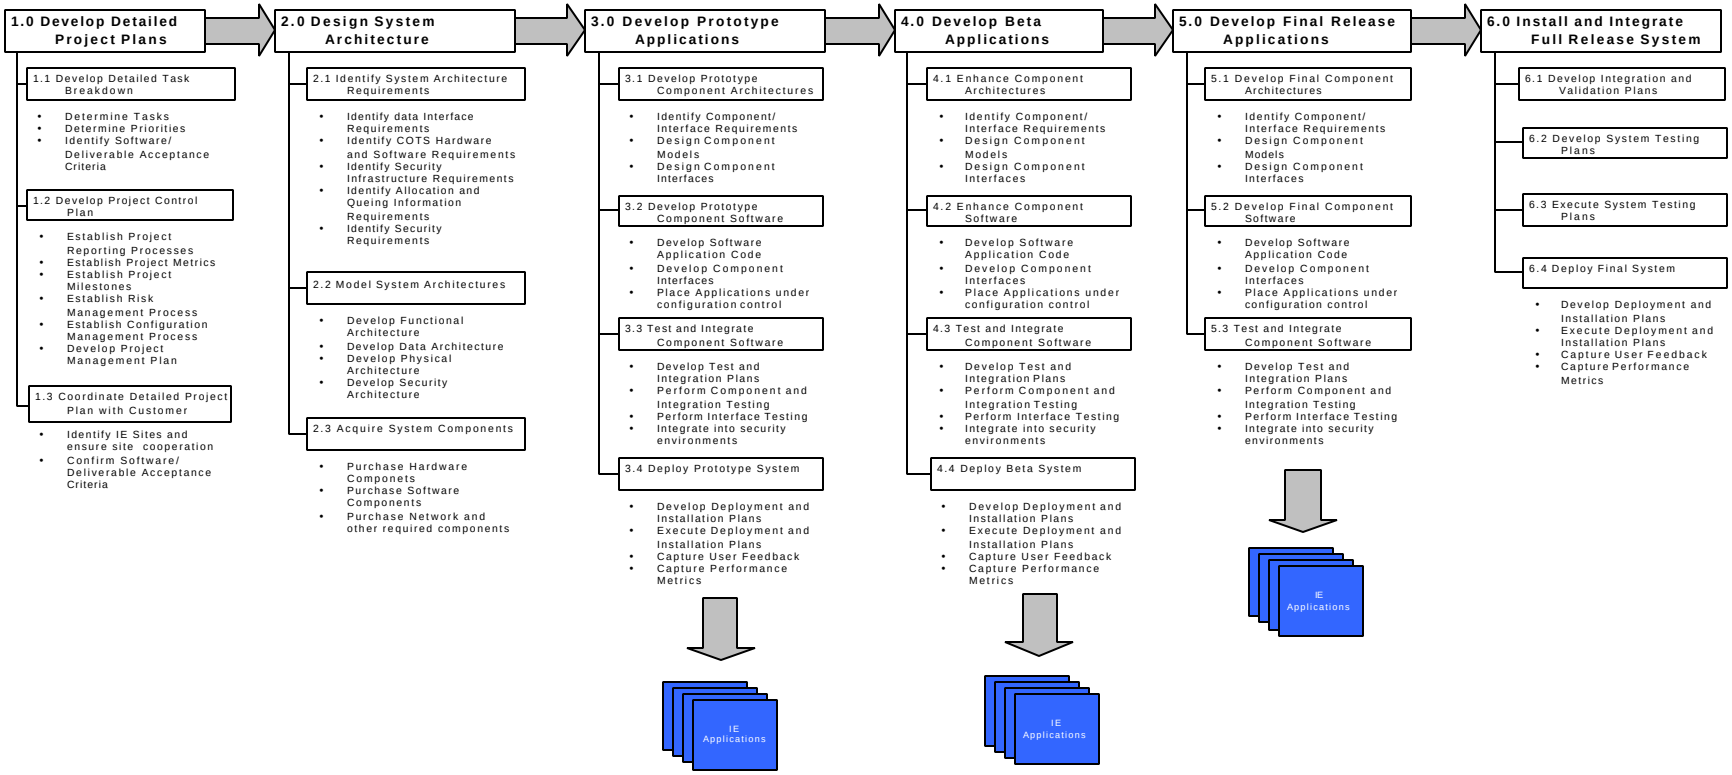

Figure 8 Task Plan Schematic

2.2 IE Development Schedule

Figure 9 shows the system development schedule. The schedule is a 365 day schedule with major milestones at 90, 180, and 300 days for the Prototype, Beta, and Final Releases of the user application tools respectively.

2.3 Deployment Plan

2.3.1 Site Survey and Plan Development

Our deployment practice is to first conduct a site survey based site-specific information provided by the government. We will interview and coordinate with the site government and

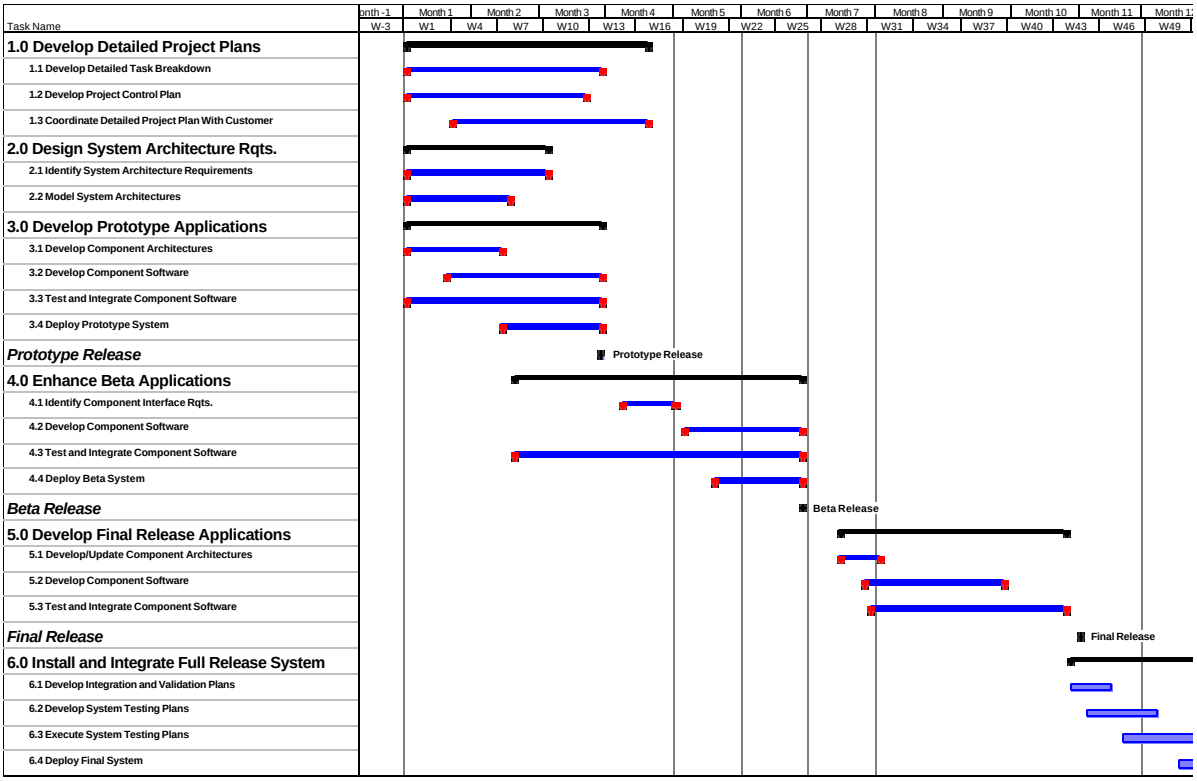

Figure 9: IE Development Schedule

contractor staff, and from this information develop installation and deployment plans unique to each location, and unique to the tools being deployed. Particular attention is paid to identification of schedule risks. As part of the Deployment and Installation plans, mitigation plans are developed to offset those risks.

Special attention is given to the existing site configuration and any external criteria are verified during the site survey effort. A facility baseline is developed that establishes the infrastructure necessary to support the preparation, installation, and integration of the target

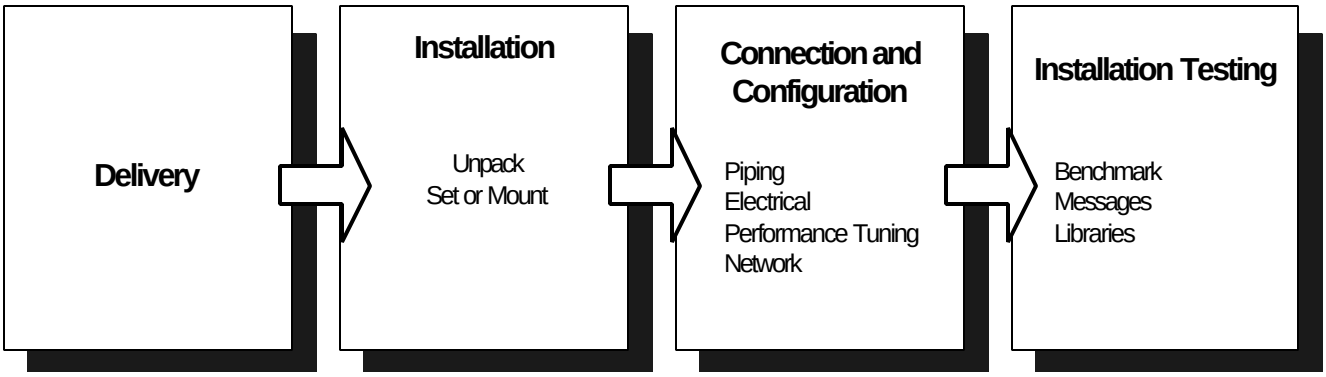

Figure 10: System Integration and Deployment

configuration. Our deployment schedule is predicated on the assumption that the deployment sites will possess the infrastructure required to support the release of software being deployed, and that if there are shortfalls in required infrastructure the government will procure and install the required supporting hardware or software. The schedule is also predicated on the assumption that the government will facilitate the mitigation of differences in local site policies such as accounts and security and deployment of non-COTS software.

2.3.2 Delivery

After the Deployment and Installation plans are developed, project personnel will work with the government and implementation site personnel to coordinate delivery of the government

furnished equipment required for that particular installation. It is assumed that it is the responsibility of the government to ensure required equipment is located at the installation site at the time of deployment. Additionally, the site staff will be available to assist as necessary with the installation of the system components.

2.3.3 Installation

The Integration and Deployment Team will travel to each site and install system components as specified in the Deployment and Installation Plan. In order to accomplish this task in a timely and efficient manner, it is assumed that the government will provide necessary access to each site where the system is deployed.

2.3.4 Connection and Configuration

As the system components are deployed to the specific sites it is the responsibility of the task leader to ensure the requirements specified in the Deployment and Installation plan are satisfied. Typical requirements may be establishing interfaces to other components and libraries. System components, especially network and communications interfaces, various security and access logging scripts, and other required configurations are verified from checklists in the Deployment and Installation plan. The Integration and Deployment team will place deployed components of the system under configuration control as they are deployed.

2.3.5 Installation Testing

The final step in the deployment process is installation testing. As components are deployed at the installation site they are placed under configuration control and tested to ensure functional, interface, and performance requirements are satisfied. If a component fails a test a discrepancy report is submitted to the development team and the problem is corrected. Once components have passed Installation Testing they are considered deployed to the site.

2.4 Personnel Plan

The HPTi and NACSE team is committed to ensuring that dedicated, trained and inspired personnel participate in the design, development, and deployment of the IE system. Each potential candidate for employment on the project will go through a rigorous screening process where that person's qualifications are evaluated to ensure they have the skills required to meet the technological challenges faced during the development and deployment of the system. Perhaps more importantly, the engineers, programmers, and other specialists who will work on the project will be motivated by those challenges and strive to provide the best system possible.

The resume of the Program Manager is provided below. Ms. Maher brings extensive managerial experience to project and has led large development efforts to successful conclusions in the past. She recently received the IRS CIO Award for Excellence for her current work on a nation-wide system deployment of large IRS Data Centers. Ms. Maher is especially adept at ensuring customers expectations and requirements are satisfied. She has extensive experience in ensuring the right personnel are working the right tasks on the project.

2.5 Key Personnel Resume

NAME: Jayne L. Maher

POSITION: Director

EDUCATION: Graduate Certificate, Information Systems, February 1990
American University, Washington D.C.
B.S. Business Management, June 1986
Virginia Polytechnic Institute, Blacksburg, VA.

GENERAL EXPERIENCE: Ms. Maher has thirteen (13) years experience in program management, software development, and systems integration of large complex systems.

- Program and Task management.
- Application development including: Visual Basic, Gupta SQLWindows, BASIC, COBOL and FORTRAN programming.
- Databases: Oracle, Informix, SQLBase, Access
- Systems analysis including systems life cycle development and requirements analysis.
- Technical documentation including complete systems life cycle documentation and proposal development.

EXPERIENCE HISTORY:

8/98 – Present: High Performance Technologies, Inc. (HPTi), Director

IRS Service Center Mainframe Consolidation (SCMC) PMO Support – Senior Technical Manager in charge of 20 program management specialists supporting the IRS PMO office. Functions performed as the HPTi Senior Technical Manager, include; task scheduling, budget reporting and tracking, contract negotiations, status reporting, and project staffing. The HPTi team is responsible for supporting the IRS in managing the consolidation of 10 service centers

into two computing centers. This management effort incorporates common management methodologies and tools for tracking program requirements, change requests tracking, issues and risks tracking, and status and executive reporting. The HPTi team is responsible for identifying, implementing and facilitating a number of the management committees and efforts used by the PMO to support this effort. In addition, HPTi has developed and is in the process of implementing a web-based program management tool that integrates program related data in to a centralized database that is used for executive reporting, status tracking, and executive decision making. This system integrates issues and risks, the program schedule (WBS), program requirements, and lessons learned into a centralized program management tool.

1/97 – 8/98: DynCorp I&ET, Project Manager of Software Development

Provided technical and managerial support in the development and implementation of the Navy BUPERS Defense Personnel Records Imaging System-Electronic Military Personnel Records System (DPRIS-EMPRS). The integrated solution supported over 9 Unix servers, 1500 client workstations with a Visual Basic client application, and a COTS package to provide document storage and access. Primary responsibilities included technical management for the software development team consisting of 7-8 analysts/programmers. Additional responsibilities included supervising and coordinating the integration of independent system modules developed by subcontractors.

10/94 – 1/97: DynCorp Corporate MIS, Project Manager/Software Engineer

Provided technical and managerial support to projects within DynCorp Corporate MIS. Primary responsibilities included: project management, application development, training and technical documentation. Secondary responsibilities included: providing technical support on proposals.

JVAP Proposal Development - Project Manager, responsible for providing technical support and direction in developing a web-based information system to be used as part of the program management tool for the Joint Vaccine Accreditation Program (JVAP) for the DOD-JPO.

Enterprize Imaging System - Project Manager, worked with DynCorp's Information and Engineering Technology (I&ET) group to use their Enterprize Imaging System internally on projects within DynCorp.

TESCO Proposal Development - Project Manager, developed the Technical Implementation Plan used as part of the TESCO proposal to provide an integrated solution for equipment management and scheduling.

Ft. Rucker Electronic Document Management System - Project Manager, was responsible for the design and integration of an imaging solution to provide document management of Personnel Records within DynCorp.

2/90 – 10/94: Universal Systems Incorporated (USI), Project Manager/Systems Engineer

Provided technical and managerial support to multiple projects within the Document Management Center. Primary responsibilities included: project and task management, application development, training and technical documentation. Secondary responsibilities included developing imaging system demonstrations and marketing imaging system capabilities to prospective clients.

SECURITY CLEARANCE: DOD Secret Clearance, Issued June 1999

3.0 Past Performance

ARL PET

Name of Contracting Organization	Contract Number
Army Research Laboratory	DAHC94-96-C-0010
Contracting Officer	**Technical Representative**
Tom Crimmins Bldg 328 Aberdeen Proving Ground, MD 21005 410-278-6267	Andrew Mark Bldg 394 Aberdeen Proving Ground, MD 21005 410-278-9761
Contract Type	**Total Contract Value**
CPFF	$4,486,577
Performance	

Description of Work: As subcontractor to Raytheon, HPTi provides the leadership and management of the Programming Environment and Training (PET) component of the Army Research Laboratory's (ARL) Major Shared Resource Center (MSRC). HPTi's overarching objective is to help ARL in ensuring the MSRC's HPC capabilities are utilized effectively by DoD researchers to support the warfighter's current and future mission needs. PET resources provide the computational science and programming environment foundation for DoD researchers, and provide an environment of support and outreach to ensure DoD researchers advantageously use the MSRC's capabilities in their programs. At ARL, HPTi's PET program accomplishes this by:

- Improving the ARL programming environment through directed investments in computational sciences applications software, visualization software, and software engineering tools. HPTi provides specialized support and development in the Distributed Interactive Computational Environment for Computational Fluid Dynamics research done at the MSRC, as well as other scientific visualization expertise.
- Providing leadership in the DoD High Performance Modernization's Program Initiative in Metacomputing. HPTi has led a DoD working group in defining the HPCMP's Metacomputing Program and has identified partners with NASA's Information Power Grid and the Globus and Legion Programs.
- Leveraging the capabilities of the best in industry and academia through active technology transfer of tools, algorithms, knowledge and expertise. Activities in this PET area involve the collection, review, and dissemination of the best programming tools available.
- Providing outstanding Graduate, Post-Graduate, and skills training support utilizing innovative delivery mechanisms such as training courses, seminars, workshops and international conferences.

- Leading efforts in Information/Communications (I/C), to address software technologies to support information dissemination and collaboration within the HPC research community. I/C works with the leads of the various Computational Technology Areas (CTA) to identify needs and provide expertise, tools, and training in the required software technology areas.

A few of the notable successes of the ARL PET program under HPTi's leadership are:

- *Engagement of PM SADARM.* A serial SADARM scene generation code was ported to a 32-procesor SGI O2K and parallelized. The run time decreased from nine hours to nine minutes. This success led PM SADARM to seek the ARL PET to help them to develop a virtual prototyping initiative for SADARM. This was pursued by HPTi under a CRADA via the Armaments RDEC.
- *Immersive visualization capability.* Testers can now validate the data gathered from field tests in ways they never could before. This capability was used successfully in a fluid tanker test.
- *TANGO.* The ARL PET program, in collaboration with Syracuse University, successfully developed and deployed a synchronous distance learning technology, TANGO. It has been used successfully in many training programs across multiple DoD sites.
- *Collaboration with ACCESS.* The ARL PET program has teamed with NSF Alliance programs to develop and sponsor a series of training programs at the Alliance ACCESS Center in Ballston, VA. This offers opportunities for DoD users to collaborate and interface with users involved in NSF programs.

NAVO PET

Name of Contracting Organization	Contract Number
Logicon	DAHC94-96-C-0008
Contracting Officer	**Contracting Officer Representative**
Brenda Spence, Administrative Contract Officer, Naval Oceanographic Office, Building 9134, Stennis Space Center, MS 39529 228-688-4640	Eleanor Schroeder, PET Director, Naval Oceanographic Office, Building 9134, Stennis Space Center, MS 39529 228-688-4669
Contract Type	**Total Contract Value**
CPFF	$2,845,725

Performance

Description of Work: As a subcontractor to Logicon, we performed user requirements analysis with respect to software tools, programming environments, integration of data from distributed sources, and remote access/navigation of data sources. For example, we identified the users' needs for three tools related to general HPCMP resources, and then developed those tools:

- A Queue Status Web Page (joint project with SDSC) enables researchers to maximize their effectiveness in using HPC resources by informing them of current availability and predicted timelags associated with job queues. The interface is "personalizable" so that each researcher can limit displays to the particular machines and queues in which he/she is interested. To ensure that system performance is not degraded even when many users are accessing the page, data ingestion is carried out at changeable intervals and cached for reference.
- The Resource Allocations Database is a relational (Oracle-based) database that maintains information on HPCMP users, projects, resource allocations, and usage. Multi-level, Web-based interfaces allow different classes of users (e.g., S/AAAs versus PIs) to see different views that reflect their concerns as well as their rights to access the data. Special features support report generation (corresponding to the formats identified as most useful by different classes of users) and single-keystroke capability to download the data into spreadsheets. Data is ingested automatically from raw machine accounting records on a daily basis.
- The Resource Allocations Exchange extends the concept of the resource database to allow HPCMP users and their allocation authorities to manage machine allocations more efficiently. Based on an in-depth analysis of user requirements and allocations processes, the tool allows users to post "offers" of machine allocations that they are willing to exchange for allocations on other machines, finds matching offers, then manages the authorization messaging necessary to enact the exchange. It also performs historical tracking to provide information on what "exchange rates" proved viable.

These efforts demonstrate experience with identifying HPCMP user requirements, conducting reviews with users and S/AAAs from MSRCs and DCs, and developing and carrying out usability testing regimes.

Usability engineering techniques were also applied to the development of standards for HPC debuggers; improving access to distributed archives of computational model results; and defining requirements for HPC software environments.

IRS/SCMC EDMIS

Name of Contracting Organization	Contract Number
Internal Revenue Service	TIRNO-00-D-00009
Contracting Officer	**Contracting Officer Representative**
Ms. Debbie Smith Accenture LLP 11951 Freedom Drive Reston, VA 20190 (202) 622-1528	Mr. Jerry Yaros IRS-PMO-SCMC Offices Unisys 8008 Westpark Drive McLean, VA 22102 (703) 556-5271
Contract Type	**Total Contract Value**
CPFF	$5.9 million

Performance

Description of Work: The Internal Revenue Service (IRS) is about to complete what is probably the largest and most successful consolidation of dispersed mainframe service centers in the Department of the Treasury. The technology consulting team of High Performance Technologies, Inc. (HPTi) continues to play a pivotal role in working with the IRS' Service Center Mainframe Consolidation (SCMC) Program Management Office (PMO) to complete this $1.5 billion dollar effort. Originally, the IRS had ten mainframe service centers located throughout the Nation to handle all IRS taxpayer-processing requirements. In seeking to modernize the service and its technology, the IRS embarked on an ambitious program to consolidate its ten service centers into two state-of-the art computing centers with upgraded, Y2K-compliant technologies which would handle all IRS taxpayer processing requirements nationwide and provide the foundation for future modernization efforts.

The goal of the PMO, with the assistance of the HPTi team, has been to manage this key SCMC consolidation with other partners such as Unisys and IBM. The team is providing critical services in the support and management of major technology initiatives and platform upgrades essential to the IRS' capability to function within the post-consolidation technical environment. As the completion of the Program's third year approaches, the joint efforts of the IRS and the HPTi Team demonstrate substantial success in achieving the consolidation goals. To date, all ten of the IRS service centers have been folded into the new dual center concept without apparent difficulties in transition and delivery of services, as well as establishing a sound foundation for the ongoing technology modernization effort.

The HPTi Team works with the IRS PMO to identify and implement those Best Practices that apply to the Federal Government in areas including:

- ? Risk Management and Oversight;
- ? Executive Reporting;
- ? Program Audit Management;
- ? Technical Oversight and Coordination;

- ? Business Process Improvement;
- ? Strategic Planning;
- ? Transition Planning;
- ? Program Scheduling; and
- ? Change management and Control.

In creating this relationship, the PMO worked with its team members to create a practical strategy that incorporates the best of methodologies, applications, and Internet to achieve a platform for promoting electronic government and serving taxpayer needs in the 21st century. This aggressive tactic included incorporating current and future Internet and Intranet technologies to facilitate the flow of knowledge management in a decentralized SCMC PMO. The Enterprise Directorate Management Information System (EDMIS) is a first step towards reaching this goal and helping the IRS' PMO to continue as a recognized U.S. Government Center of Excellence.

HPTi designed, developed and implemented EDMIS, as a web-based program management tool, with a centralized database, to support activities such as issues and risk tracking, requirements tracking, change and configuration management and tracking, action item tracking, lessons learned, and schedule management. Each of these functions was designed to provide an integrated view into the status of the program allowing the project office to manage and control the consolidation effort. EDMIS provides a three-layer application architecture with business objects providing the layer between the database backend and the web front-end.

Based on the capabilities and success of EDMIS for the SCMC PMO, other organizations within the IRS have or are currently adopting it as the system of choice for supporting the Program Management functions within IRS Information Systems (IS). For example, as part of the Andersen team, HPTi is providing program and project management support to the System Development (SD) division within the IRS. As part of this effort, HPTi used the EDMIS system as a baseline to develop a Risk and Issues Management System (RIMS) for use by SD. Like EDMIS, RIMS provides an integrated management tool for use as part of the overall program/project management methodology. RIMS was developed using the same system architecture as the one used for EDMIS.

For both IRS systems, EDMIS and RIMS, HPTi was responsible for providing the requirements analysis, design, testing, implementation, user training, system documentation, and is currently performing maintenance on the production systems.

FBI/ITB

Name of Contracting Organization	Contract Number
Federal Bureau of Investigation	#J-FBI-99-037
Contracting Officer	**Contracting Officer Representative**
Chrisann Wirtz Federal Bureau of Investigation J Edgar Hoover FBI Building 935 Pennsylvania Ave, NW Washington, DC 20535 (202) 324-3000	Dr. Joseph Kielman Federal Bureau of Investigation J Edgar Hoover FBI Building 935 Pennsylvania Ave, NW Washington, DC 20535 (202) 324-1533
Contract Type	**Total Contract Value**
FFP	$8.8 Million

Performance

Description of Work: HPTi assists the FBI in its efforts to design and implement a secure information-sharing infrastructure. This support spans the definition of enterprise, information, and technical architectures to the prototyping, piloting, and implementation of components of the target architectures to support the ongoing investigative mission of the FBI. Our experience with development of the FBI's Technical Architecture Definition (TAD) demonstrates our ability to develop large, complex, technical architectures, and specifically demonstrates an ability to creatively apply advanced IT technologies to support requirements across a broad spectrum of areas, from WAN/LAN networking and communications, to multimedia information management, to scalable servers, to security. Equally important, our development of a standardized, enterprise-wide information architecture demonstrates our understanding of the need for the use of common data elements across applications to promote interoperability as the FBI moves forward to upgrade its infrastructure. Finally, the enterprise architecture initiative will provide a blueprint that connects and guides the architecture components of the FBI along with the overarching processes for implementing and updating their strategic plan.

As part of the technical architecture initiative tasking, HPTi reviewed the various technologies currently available in the areas of telecommunications, hardware, data management and software. Along with the current technologies available, future technology trends were identified and incorporated into the assessment. From this evaluation of current and future hardware and software, alternatives were analyzed, evaluated, and presented to the FBI. The resulting Technical Architecture (TA) incorporated a wide range of technologies including ATM, SMP parallel servers at distributed sites, Netware and NT for local networking, and data modeling/management suites. Through these technology investigations and ongoing research, we have a corporate knowledge of technology industry trends and can help guide key technology decisions.

A fundamental element to providing guidance to the FBI's decision makers is the ability to demonstrate the potential and operational readiness of key technologies and their ability to meet short-term requirements, and their relevance to long-term infrastructure strategies. As part of the Office of the Chief Scientist (OCS) Infrastructure Testbed (ITB) program, HPTi

manages the activities of the Infrastructure Testbed, including personnel management, resource management, contract management, and scheduling. Our success with the investigation and implementation of technology solutions is due in large part to our extensive knowledge in specific technical disciplines such as ATM and Gigabit Ethernet switching, Oracle database development, SMP enterprise systems, Java development, scalable document imaging solutions, multimedia storage, search, and retrieval, heterogeneous networking (Unix, NT, Novell), etc. One role of the ITB is the ongoing investigation of leading edge technologies, which keeps our technologists current across many disciplines.

These skills have been demonstrated through:

- ? Development and implementation of NT and UNIX, Oracle-based multimedia capture, management, and search prototypes and pilots.
- ? Advanced WAN switching, and implementation of a NetWare, NT, and UNIX heterogeneous networking solution. The ITB LAN supports NT, NetWare, and UNIX workstations and servers, all of which are interconnected to allow file sharing and network printing. In addition, in coordination with the Major Projects Section (MPS) and Operations and Maintenance Section (OMS), HPTi assisted in the planning of the FBI's migration from a private, token ring based LAN/WAN infrastructure to one which uses fast Ethernet, ATM WAN, and encryption components to interconnect with the Justice Consolidated Network (JCN). We evaluated the solutions of leading vendors' solutions for standards compliance and performance in the target architecture.
- ? Flexible, scalable document imaging solutions that have been deployed to Strategic Information Operations Center to assist the FBI in two high profile cases utilizing the resources of the ITB.
- ? Development and implementation of a formal performance modeling, simulation, and evaluation methodology that has been used successfully to guide LAN networking strategies, investments to improve large dataset visualization capabilities, secure WAN standards compliance and compatibility vendor choices, the cost/benefit of investment upgrades for a large data mining application, and the relative capabilities of UNIX and NT servers for web-based, multimedia data management systems.
- ? As an aid to technology exploration and evaluation, HPTi has ongoing relationships with key hardware and software vendors, allowing HPTi to continue to develop specialized knowledge regarding pertinent technologies. HPTi provides expertise in computer security related technologies, such as network firewalls, single sign-on, data encryption, intrusion detection systems, and NT, Unix, and Novell OS security issues. HPTi has supported development and implementation of a secure remote, dial-up capability, and has defined security procedures and policies for the ITB and validated hardware and software solutions against these policies.

Geographic Separation. HPTi supports the FBI's infrastructure testbed from two locations: FBI Headquarters, in Washington DC, and HPTi headquarters, in Arlington, VA. Prototypes and evaluations are developed and run in both locations.

Electronic Delivery. Technology evaluations, prototype designs, and other reports generated

as part of the Testbed Initiative are published to the OCS Website to make them accessible throughout the Office within the FBI.

Technical Quality and Accuracy. Performance metrics for the ITB NTET Program are defined by: SUM [(milestone met) X (scheduled date delivered) X (capabilities or performance)]

To date, all milestones have been met and deliverables accepted. In FY00 alone, this includes over 110 separate deliverables.

National Science Foundation

Name of Contracting Organization	Contract Number
National Science Foundation	ASC-95-23629
Contracting Officer	**Contracting Officer Representative**
Irene Lombardi, 4201 Wilson Blvd. Suite 1122, Arlington, VA 22230 703-292-8970	Dr. Richard Hirsh, 4201 Wilson Blvd. Suite 1122, Arlington, VA 22230 703-292-8970
Contract Type	**Total Contract Value**
Federal funds - restricted	$1,518,567

Performance

Description of Work: This project developed a number of web-based middleware and user-interface products. One user-level tool, Webterm, was adopted by Hewlett-Packard as part of its HP/UX operating system distribution.

Several products were associated with layered middleware for generating web-to-database interfaces that isolate platform-dependent aspects to maximize maintainability. Activities included:

- Design and implementation of robust middleware supporting a uniform data interface to hide the idiosyncrasies of common relational DBMS software packages
- Experience developing and deploying portable system-level middleware
- Packaging and distribution of software that can be successfully installed and used by untrained HPC users
- Experience applying usability engineering methodologies to the design of user-level middleware

As a result, a highly successful collaboration of scientists, engineers, and computer scientists has been established and has generated several software products. Many provide a middleware infrastructure upon which other software products can be built. A variety of user-level tools have also been developed to make scientific data accessible to users without their having to be aware of whether data are located locally or across the country.

Through the efforts of this initial project, the NACSE team has established a proven track record in (a) developing interfaces that scientists will actually use, (b) producing software that scientists can install and apply without special training or assistance, and (c) devising ways of minimizing user errors by trapping them before they take effect.