NISTIR 4780

Guide for Specifying and Building CITIS with Data Management Standards

Leonard J. Gallagher
Joan M. Sullivan
Joseph C. Collica

U.S. DEPARTMENT OF COMMERCE
Technology Administration
National Institute of Standards
and Technology
Computer Systems Laboratory
Database and Graphics Group
Gaithersburg, MD 20899

U.S. DEPARTMENT OF COMMERCE
Barbara Hackman Franklin, Secretary

NATIONAL INSTITUTE OF STANDARDS
AND TECHNOLOGY
John W. Lyons, Director

NIST

NISTIR 4780

Guide for Specifying and Building CITIS with Data Management Standards

Leonard J. Gallagher
Joan M. Sullivan
Joseph C. Collica

U.S. DEPARTMENT OF COMMERCE
Technology Administration
National Institute of Standards
and Technology
Computer Systems Laboratory
Database and Graphics Group
Gaithersburg, MD 20899

February 1992

U.S. DEPARTMENT OF COMMERCE
Barbara Hackman Franklin, Secretary

NATIONAL INSTITUTE OF STANDARDS
AND TECHNOLOGY
John W. Lyons, Director

Guide for Specifying and Building CITIS with Data Management Standards

Leonard Gallagher
Joan Sullivan
Joseph Collica

National Institute of Standards and Technology
Information Systems Engineering Division
Gaithersburg, MD 20899, USA

- Abstract -

This paper complements the military specification for a Contractor Integrated Technical Information Service (CITIS) when that service is to be provided in an environment that supports other federal standards for data management systems. We present the current status of existing and emerging International Organization for Standardization (ISO), American National Standards Institute (ANSI), and Federal Information Processing Standard (FIPS) standards for database management and data dictionary systems, specifically Database Language SQL, Remote Database Access (RDA), and Information Resource Dictionary System (IRDS). We address the CITIS specification in terms of these data management standards and indicate how these standards may be specified or used to meet the requirements for various levels of service and functional requirements of CITIS. We conclude by identifying the benefits of data management standards in the CITIS architecture. In the Appendices, we describe the general content of each data management standard and discuss its applicability and availability. Where appropriate, we also address the availability of conformance test suites and future plans for enhancements and follow-on standardization efforts.

Keywords: (ANSI; CITIS; database; data dictionary; data model; FIPS; IRDS; ISO; MILSTD; network; RDA; relational; standard; SQL)

Note: This document was prepared as a report to the CALS Evaluation and Integration Office. Its publication as a NISTIR does not imply CALS endorsement of the conclusions or recommendations presented.

Table of Contents

1. Introduction

A Contractor Integrated Technical Information Service (CITIS) is a contractor-provided service that provides the government with authorized access to contractor databases and applications, including both business and technical information. CITIS encompasses all activities and functions that are necessary for the government to achieve practical use of digital data generated by the contractor.[1] Additionally, CITIS may include the provision of computer hardware and software by the contractor for government use.

The purpose of CITIS is to enable the government to acquire electronic access to data through a contractor-provided service rather than to take delivery of paper documentation. The CITIS specification defines the government's baseline functional requirements for integrated technical information services and should be tailored by the procuring agency for all acquisition programs that require extensive delivery of data. In order to ensure that the contractor provides the required services, the procuring agency may require the contractor to develop and deliver plans that detail the contractor's approach to providing such services.

One of the major objectives of this report is to identify appropriate data management standards that may be used in conjunction with CITIS to enhance the quality and usability of information provided. Proper use of data management standards should make it easier for the procuring agency to specify exactly the information required from the contractor and for the contractor to provide that data in a form that is more easily integrated into present and future government information processing systems. Section 2 describes data management fundamentals useful in viewing CITIS specifications. Section 3 provides an overview of the existing data management standards useful in a CITIS environment. These data management standards are Database Language SQL, Remote Database Access (RDA), and Information Resource Dictionary System (IRDS).

We present the current status of International Organization for Standardization (ISO), American National Standards Institute (ANSI), and National Institute of Standards and Technology (NIST) Federal Information Processing Standard (FIPS) adoption of the SQL, RDA, and IRDS standards as well as the status of emerging revisions and follow-on efforts. We address the CITIS specification in terms of these data management standards and indicate how these standards may be specified or used to meet the requirements for various levels of service and functional requirements of CITIS. We conclude by identifying the benefits of data management standards in the CITIS architecture and by noting the negatives if such standards are not specified.

[1] These introductory sentences are derived from the Scope section of a draft MIL-STD, MIL-C-CITIS, under development by the CALS-EDI Office, dated 31 March 1991.

Another major objective of this report is to provide guidance to contractors in building a CITIS with data management standards. Section 4 relates data management standards to various CITIS levels of service and provides examples of using specific data management standards at each level. Section 5 discusses how data management standards can support CITIS requirements.

A third major objective of this report is to provide an easy way for the procurement agency to specify their data management standards requirements for SQL, RDA, and IRDS along with their CITIS requirements. Section 3.3 gives an introduction to procurement wording and Appendix C provides suggested standard procurement wording and guidance on validation options.

In Appendix A we describe the general content of each data management standard (SQL, RDA, IRDS) and discuss its applicability, availability, and future plans for enhancements and follow-on standardization efforts. A similar analysis of other important standards for application portability and interoperability; e.g.,Government Open System Interconnection Profile (GOSIP), Portable Operating System Interface for Computer Environments (POSIX), Programmer's Hierarchical Interactive Graphics System (PHIGS), IRDS, Office Document Architecture (ODA), etc., is contained in a NIST publication, Application Portability Profile: The U.S. Government's Open System Environment Profile, OSE/1 Version 1.0, available from the U.S. Government Printing Office (GPO) as NIST Special Publication 500-187.

Appendix B presents the NIST program for validation and testing of language compilers and addresses the availability of conformance test suites for SQL, RDA, and IRDS. Appendix C provides candidate wording for database procurement, for software development, and for validation requirements and correction of deficiencies. Appendix D provides a list of standardization organizations and indicates where specifications can be obtained.

2. Data Management Fundamentals

Data management has traditionally employed a three-schema architecture to place itself in a data processing environment. A conceptual schema represents a high-level, enterprise-wide view of all data, data relationships (including rules restricting updates or cascading the effects of updates to related data), and the business processes that use and update the data. Generally an enterprise has only one conceptual schema. Changes in the details of computer implementation or the specific human users and application programs that access the data have no effect on the conceptual schema.

An external schema represents a logical view of the data as accessible to a set of human users and application programs; an enterprise may have many external schemas. An external schema is generally a small subset of the conceptual schema, and may have application-oriented synonyms for the data defined in the conceptual schema. Changes in computer implementation have no effect on the external schema; this facilitates migration of the data to other hardware and software environments. An external schema may change to accommodate changes in the use of the data.

An internal schema represents a physical view of the data as stored on persistent storage devices. An enterprise may have many internal schemas to provide efficiency on a variety of hardware and software environments. Conceptual and external schemas are independent of the structures and access methods of any underlying file system; in contrast, an internal schema may be heavily dependent on file structures and access methods.

Each schema is constructed according to the rules of a data model (e.g., the relational data model). The data model prescribes not only the rules for defining data structures, but also the rules for interpreting and manipulating the data structures.

The conceptual schema may consist of a very large collection of object types and their interrelationships; no single application program will require access to all the objects described by the conceptual schema. In contrast, an external schema may consist of a simple "record-oriented" view of a single object type; a third generation programming language can easily process data described by such a schema. The conceptual schema itself may be so complex that it must be maintained by specialized software such as an IRDS. The IRDS may also be required to manage the mappings between the conceptual schema and the different external schemas, and the relationships among data and processes.

In the Integrated Weapons System Database (IWSDB) described in previous Computer-aided Acquisition and Logistic Support (CALS) reports, we assume two or more local data processing environments (e.g., government and contractor sites) communicating with one another at the highest conceptual schema levels. This ensures that the communicating environments share a common understanding of business semantics of the data. This is a desirable goal that depends heavily on data management standards not yet fully developed. At the other extreme, in the absence of any data management standards, a government site and a heterogeneous contractor site can only communicate at the very lowest levels. Each site may be able to access files of data or "bulletin-board" views of data at other sites, or it may be able to pass parameters to application processes at remote sites, but it is not possible to access data described by a complex external schema without knowing details of the data model used to construct that external schema and a syntax or protocol for invoking operations on that data model. Unless there are standards for schemas, access to remote sites may be effectively limited to the use of very low-level, simple external schemas, limiting common understanding to very basic syntax and semantics. This low-level communication forces the application programs to perform many of the data structuring and re-structuring tasks that could be performed more effectively by the database management system. Also, data dictionaries must be maintained to ensure a common understanding of the data names and synonyms.

3. Applying Standards to CITIS Specifications

At the risk of over-simplification, we can say that, without data management standardization, government and contractor sites will be doomed to communicate only at the simple external schema level (e.g., file transfer), while future standards promise communication of complex objects. CITIS, with its four levels of service, allows gradual adoption of existing and emerging data management standards to move from simple external schema communication toward the desired goal of high-level conceptual object integration.

With existing or very near-term data management standards, primarily SQL with support from RDA and IRDS, it is possible for each local site to construct a "tabular" external view of its logical data structures that can then be accessed and manipulated by all other sites. With longer term emerging data management standards that support object-oriented and knowledge-based features, CITIS environments can evolve into the desired IWSDB with government and contractor sites communicating at the conceptual level with "seamless" integration of complex, structured data and supporting application services.

In the following subsections we address the CITIS baseline requirements and give an overview of existing data management standards and how these standards might be invoked in federal procurements. In Section 4 we describe the evolution from CITIS Level 1 to Level 4 services and in Section 5 we address CITIS functional requirements as stated in the CITIS specification.

3.1 CITIS baseline requirements

The CITIS specification defines the government's baseline functional requirements for integrated technical information services. These baseline requirements should be tailored by the procuring authority for each acquisition that requires extensive delivery of data. Along with each Request for Proposal (RFP), the government will provide a Government Concept of Operations (GCO) that details how it intends to meet the CITIS objectives to reduce overall technical information costs and to improve quality and timeliness.

The GCO will include a high level overview of the environment in which data is to be exchanged or shared among the government and its contractors, including requirements for contractor distribution of data to specified government users or for direct government access to contractor information databases.

The GCO will contain a Contract Data Requirements List (CDRL) that expands and/or tailors each of the functional requirements given in the CITIS specification. The procuring authority will evaluate all functions requiring CITIS data and will specify appropriate data items and usage requirements for input to the CDRL. Each data item that is to be exchanged or shared among the government and a contractor will be the subject of a Data Item Description (DID) that defines its user output mode, distribution, format for delivery, and classification and sensitivity.

In response to an RFP a potential contractor will provide a plan for meeting the requirements identified in the Contract Data Requirements List (CDRL) and the Data Item Descriptions (DIDs). The plan will include details for acquiring, implementing, and maintaining the integration of functional processes and for documenting, implementing, and maintaining a Contractor Integrated Technical Information Service (CITIS) that gives government users access to both contractor generated and government furnished data. The contractor plan will specify the automated tools proposed to support each requirement and how the contractor will provide government access to the contractor maintained integrated database. In addition, the contractor plan will include a description of the agreement the contractor has with associate contractors, subcontractors, suppliers, and vendors for inclusion in the CITIS.

3.2 Overview of existing data management standards

Data management standards have been in existence since 1986 with parallel publication of FIPS for the Database Language SQL standard for relational database applications and the Database Language NDL standard for network model database applications. A data dictionary/directory standard followed in 1988 with publication of the Information Resource Dictionary System as a FIPS. A Remote Database Access specification is currently at Draft International Standard status and should be formally accepted in 1992 with conforming implementations to follow shortly thereafter.

The SQL standard has been the most successful of these standards and should be the primary focus of further development within CITIS. SQL is recommended for CITIS applications. The NDL standard has been much less successful and is no longer recommended as a viable alternative for new application development. NDL is not recommended for CITIS applications. The IRDS standard complements the SQL standard by providing services and utilities for management of metadata, that is, information about the data. The IRDS is recommended for CITIS applications. The RDA standard provides services and protocols for database interconnectivity between different implementations of database management systems and is especially effective in providing interoperability among conforming SQL implementations. RDA is recommended for CITIS applications.

The original SQL standard provides a data definition language (DDL), a data manipulation language (DML), and other associated facilities of the relational data model for creating and accessing tabular data structures in a single data processing environment. Later SQL revisions specify schema manipulation, integrity checking, exception handling, catalog tables, and enhanced data types and operations, as well as interfaces to multiple programming languages, interactive users, and remote sites. The emerging Remote Database Access standard specifies standard protocols for remote connections between a database "client" and a database "server" at separate nodes in a communications network to provide database interoperability in an open systems environment. The IRDS standard specifies services to catalog, document, manage, and use metadata (i.e.,information about data) and specifies computer software facilities to record, store, and process descriptions of an organization's significant data and data processing resources. See Appendix A, Data Management Standards, for references and more detailed descriptions about SQL, RDA, and IRDS.

3.3 Specifying SQL, RDA, and IRDS standards with CITIS

Standard procurement language for SQL is available in the General Services Administration publication Federal ADP & Telecommunications Standards Index. Federal agencies are encouraged to use this language, modified to suit their needs. Federal agencies should also read the applicable FIPS to select among the procurement options available.

The standard procurement language for SQL and IRDS is included in Appendix C. This wording has been modified to provide more explicit guidance for SQL and IRDS procurements. Since RDA is not yet a FIPS, there is no standard wording offered for RDA. Instead, Appendix C includes suggested wording to specify RDA using the emerging draft international standard.

Although a federal agency procuring a relational database or a data dictionary system must require conformance to FIPS 127-1 or FIPS 156, respectively, it is the option of the agency to specify how conformance will be evaluated.

> The option recommended by NIST's Computer Systems Laboratory (CSL) for SQL is formal testing by NIST or a NIST-accredited laboratory. At the present time formal testing by NIST is the only available option. However, at some future time, additional NIST-accredited testing laboratories may also perform formal testing.

The National Laboratory Accreditation Program (NVLAP) at NIST accredits testing laboratories.

The option recommended by CSL for IRDS depends on the timeframes of the procurement. A test suite for the IRDS is under development at NIST in 1991, with anticipated release in early 1992. For a period of time, IRDS vendors will be able to evaluate their own products in-house using the NIST IRDS test suite; however, there will be no testing service. Eventually, CSL will recommend formal testing for IRDS by NIST or an accredited laboratory.

The testing options for RDA are under consideration. It is expected that NIST will adopt RDA as a FIPS when it is formally approved by ANSI and ISO. Testing options should be available soon thereafter.

There are obvious benefits and economies of scale in centralized testing and reporting. Also, it is simpler for an agency to call for CSL testing in an RFP than it is to determine what other evaluation criteria will be used and to specify these criteria in the RFP.

Requiring CSL testing, even for SQL, is not always feasible. In general, it is not realistic for a small procurement to expect a vendor to undergo validation in order to sell a single license of the software. In addition, timing may be a problem for a vendor expecting to validate later than the procurement timeframes. Then there are always issues surrounding the applicability of test results from one hardware/software environment to other environments. An SQL implementation ported from a validated environment to the platform offered in a procurement should be retested. But formal validation may not always be required. A vendor's in-house test results may be sufficient.

Sometimes products are not fully conforming. Requiring that an SQL or IRDS implementation conform fully at the time of delivery may result in a noncompetitive procurement or may exclude low-bid products. The lack of available conforming tested products should not deter federal agencies from specifying FIPS 127-1 or FIPS 156 conformance. However, it does mean that a strategy must be developed to acquire products which will conform soon and which currently have an acceptable level of conformance.

Appendix C provides procurement wording for three validation options: "Delayed Validation," "Prior Validation Testing," and "Prior Validation." An agency may choose one of these options or may specify some other method of determining whether software conforms to the applicable FIPS.

Appendix B provides an overview of test suites and test services available for SQL, IRDS, and RDA. The availability of conforming products and the maturity of testing programs are two critical factors in deciding what procurement language should be used to acquire

FIPS-conforming products, at acceptable prices, and with timely delivery. Whereas availability of test suites, testing services, and conforming products may be a factor in selecting the validation option, the most critical issue is being able to determine if the product will meet performance and capacity requirements once it is in conformance with the standard. Benchmarking results may change drastically when nonconformities are corrected. The "prior validation testing" may provide some indication of risk, but only "prior validation" can give the best assurance that the agencies' system requirements can be met. The agencies should be advised to assess the risk involved when choosing other than "prior validation."

4. CITIS Levels of Service

The CITIS specification defines various levels of service to be applied to each of its requirements.

Level 1	--	Automated accession list
Level 2	--	Predefined query
Level 3	--	Ad hoc query
Level 4	--	Access to contractor applications

We discuss each of these levels in turn and indicate how each of them might be implemented with the use of appropriate data management standards.

Level 1 -- Automated accession list

The government and the contractor will agree on an automated accession list with a facility for locating data relevant to a given subject. At a minimum, the facility will include key words, identification numbers, locations, and short descriptions of each data item of interest and automated access to a description of the computerized services that are available to authorized users. As a preferred option, the service may provide a fully integrated, logical referencing system that assists the user in locating needed and related data sets. The automated accession list provides the underlying index, access, and delivery structure for all CITIS data and as such should be required for all Data Item Descriptions (DIDs) identified by the Contract Data Requirements List (CDRL).

Requirements at this level can be minimally satisfied with a contractor-maintained electronic bulletin board that captures the above information and instructs the user as to how to order or otherwise access desired items. It is possible that the electronic bulletin board can be improved to a point-and-click graphical user interface (GUI) over the contractor's underlying database if the government and the contractor can agree on the appropriate GUI standards (see NIST SP 500-187, Application Portability Profile, for the status of these emerging GUI standards) or if the contractor provides GUI equipment and software compatible with the contractor's system.

Each document, engineering drawing, contract status report, or other agreed data item specified in a DID will be electronically available and can be accessed according to directions given via the bulletin board. Authorized government users can then use a contractor-provided facility to electronically order data and indicate the preferred delivery method (e.g., hardcopy, magnetic or optical media, electronic transfer) as specified in the contract. If electronic transfer is specified, then the user can use file transfer protocols as specified in the Telecommunications sections of CITIS (e.g., GOSIP/FTAM) to transfer the desired data item from the contractor's file system to the user's local file system. Alternatively, the DID may require that the contractor automatically distribute specified items in the preferred delivery format to agreed distribution lists at specific times.

At this level of service there is very little required use of data management standards. Unless specifically defined by an output specification and required by contract, output shall be in contractor format. In most cases the government user will update his own contract tracking database (likely an SQL database or an IRDS directory) by hand after reading the appropriate reports or receiving the requested manuals or drawings. However, if some of the items defined by DIDs are tables of values (e.g., parts-list, price-list, directory-of-drawings, etc.), then the procuring authority may use the standard SQL table definition language to describe the data types and constraints of each column, constraints among columns, and referential constraints among tables. This is the approach taken for LSAR, MIL-STD-13882B. In such cases the DID may also specify, if appropriate, a table interchange format using ISO 8211 (DDF), ISO 8824 (ASN.1), SGML, or just plain ASCII text. The ASN.1 standard is preferred because it is the basis of definition for SQL data interchange in the emerging RDA standard (see Appendix A) for remote database access, but ASCII text interchange of numbers and strings with an agreed fixed-length format, or with agreed separators, is workable in many situations.

Level 2 -- Predefined query

In addition to all of the above facilities, this level of service provides the capability for authorized government users to perform predefined queries and extractions over contractor maintained data. A predefined query means that selection criteria and output format have been agreed contractually in advance via a DID, but the user is able to supply variable selection parameters. For example, a predefined query may produce an output graphic depicting progress on a specific project over a specified time interval; only the project-id and the start-date and end-date are entered by the user. This level of service supports fixed format output (e.g., hardcopy report, predefined screens, or digital file in accordance with MIL-STD-1840) but does not guarantee that extracted data files will be directly processable by existing government systems, unless output formats have been agreed in a specific DID.

Requirements at this level can be minimally satisfied with a contractor-maintained database that is accessible by authorized government users. In addition to being able to order documents and other contractual items as in Level 1 services, the service will provide immediate extraction and presentation of information identified by the predefined queries. It is likely that this level of service will require a point-and-click graphical user interface (GUI) over the contractor's underlying database. Thus the government and the contractor will agree on the appropriate GUI standards (see above) to enable government-owned workstations to take advantage of contractor-owned data presentation facilities, or the contractor will provide GUI equipment and software to the government user that is able to interface with the contractor's information processing system.

At this level of service there is only limited required use of data management standards. SQL and IRDS may be used at both government and contractor sites, and the government may specify the semantics of predefined queries using SQL syntax, but there is no requirement for direct processing of such queries. Even though there is an implicit assumption that the contractor can provide relational or other data model views of certain contract-specific data items, the only requirement is to provide data management capabilities that are functionally equivalent to those specified in the predefined query.

As an example of predefined queries, consider the government requirement to monitor contract deliverables in "cost plus" types of contracts. The contracting officer knows that a given contract consists of a number of projects, each with a project budget and a timeframe for completion. A project consists of a number of tasks, each with an estimated budget. The contractor assigns employees to each task and is required to state the number of hours each month that an employee works on each task. The contractor must maintain this information and file a report to the government contracting officer each month. The government contracting officer may assume the logical existence of the following relational tables as views over the contractor's proprietary information:

Project (Proj_Id, Proj_Title, Begin_Date, Due_Date, Proj_Budget)

Task (Task_Id, Task_Title, Proj_Id, Task_Budget)

Employee (Emp_Id, Emp_Name, Salary)

Timesheet (Emp_Id, Task_Id, Month, Hours)

These tables could represent an audit checking requirement over any CALS project, such as EDMICS or DSREDS/EDCARS. There may be a large number of predefined queries that the government wishes to specify over this base information. With the assumption of relational tables, some predefined queries can be made specific by using Database Language SQL to specify the intended semantics. For example, a government user may desire periodic reports on tasks or projects that are over budget, or a list of all employees that have worked on a certain task at a certain period of time, or a list of project titles that include a task with a task_title that includes "MIL-STD" as a substring, etc. Each of these queries can be made specific by defining them in a DID as SQL queries over the base information tables, even if those base tables do not physically exist in the contractor's information system.

The contracting officer may wish to specify an SQL view over the base information tables to identify total monthly salary expenses for each task as a ratio of the task budget. The following "Salary_Expenses" view specifies a relational join over the four base tables in order to capture this information for each month:

```
CREATE VIEW Salary_Expenses
     (Proj_Id, Task_Id, Month, Salaries, Budget_Ratio)
AS SELECT  Project.Proj_id, Task.Task_Id, Timesheet.Month,
     SUM(Salary/2000*Hours), SUM(Salary/2000*Hours)/Task_Budget
     FROM  Project, Task, Employee, Timesheet
     WHERE  Project.Proj_Id  = Task.Proj_Id
             AND  Task.Task_Id  = Timesheet.Task_Id
             AND  Employee.Emp_Id  = Timesheet.Emp_Id
     GROUP BY  Project.Project_Id, Task.Task_Id, Month
```

A data item description (DID) on the contract data requirements list (CDRL) may require that the contractor provide a "query-by-example" interface to the Salary_Expenses view definition, as follows:

Salary_Expenses

Proj_Id	Task_Id	Month	Salaries	Ratio
:Proj	:Task or ALL	:Month or <:Month		
			Total_Sal	Tot_Ratio

This template allows a government user to specify any one or more of the three variables: :Proj for Project_Id, :Task for Task_Id, and :Month for Month. As a shorthand, the user may specify "ALL" for Task_Id to mean that all tasks within a project should be considered or may specify "<" as a prefix to the :Month variable to indicate that cumulative information up to and including that month is desired. The DID would specify the semantics for each input alternative. For example, if ":Proj" and "ALL" and "<:Month" are specified, then the following predefined queries would be executed and displayed in the above screen format:

```
SELECT  Proj_Id,Task_Id,:Month, SUM(Salaries), SUM(Budget_Ratio)
      FROM  Salary_Expenses
      WHERE  Proj_Id = :Proj
              AND Month  <=:Month
      GROUP  BY Proj_Id, Task_Id

Total_Sal  = SELECT  SUM(Salaries)  FROM  Salary_Expenses
             WHERE  Proj_Id = :Proj AND Month  <=:Month

Tot_Ratio  = Total_Sal / (SELECT  Project_Budget  FROM  Project
           WHERE  Proj_Id = :Proj)
```

There is no requirement that the contractor be able to process any of these SQL queries. Instead, the contractor would be required to provide a screen interface that accepts as input the various specified parameter combinations and produces as output all of the information specified by the queries.

Additional requirements in the DID requirements list may specify that the output of some of these queries be represented in diagrams or bar charts or some other graphical representation that can be copied into monthly report documents. In all cases the output will be in a contractor specified format unless otherwise agreed to as a specific DID output format requirement. Thus there is no guarantee that output formats will be directly processable by existing government systems unless data interchange standards have been agreed to in the DID. It may still be necessary, as it is in Level 1 environments, for the government user to copy desired pieces of information from the contractor provided outputs to the government's information systems.

Level 3 -- Ad hoc query

In addition to all of the above facilities, this level of service provides the capability for government users to process contractor generated or maintained data and to make ad hoc requests for information. This level of service assumes the capability to generate ad hoc combinations, sorts, and separations of the data. It also assumes that the contractor has integrated all of the data from its own sources and from various subcontractors into a logical representation of all data so that procedures to propagate changes and updates can be propagated as necessary across the whole database. In addition, the contractor shall provide access to the metadata through data dictionary and directory facilities so that the user is able to locate all data items and construct ad hoc queries. Level 3 access will typically require a sophisticated user and thus may not be appropriate for the functional user with only an occasional need for access.

This level of service allows the government to improve its procedures and respond to any new requirements without additional demands for data on the contractor. It provides the capability for on-line review, comment, acceptance, and approval of contractor-generated data, with such updates and comments included in the contractor database and made available to originators of the data and all other authorized users. The contractor-proposed plan to establish this level of service will also provide mechanisms to establish appropriate audit trails to identify the sources of any modifications and to maintain configuration control.

This level of service requires enforced implementation of data management standards because interoperability of government, contractor, and subcontractor data management facilities is assumed. Using the terminology of RDA-client and RDA-server defined above, government systems will most often play the role of an RDA-client accessing remote data at contractor-maintained RDA-server sites. In some cases, when the government makes government-generated data available to contractors, government systems will play the role of RDA-server making that data available to contractor RDA-client sites.

The remote database access (RDA) standard discussed in Appendix A.2 makes achievement of Level 3 requirements possible, at least for data that can be represented in relational tables. In most cases, relational systems will be the basis for data dictionary and directory facilities at contractor sites and the RDA/SQL Specialization will be the protocol for interchanging data manipulation requests and returned data. We expect that commercial implementations of the RDA/SQL Specialization will be in place about the same time that CITIS becomes part of government procurements, so it is reasonable for government RFP's to require SQL conforming systems and RDA/SQL Servers at each contractor site. The contract data requirements list (CDRL) and the individual data item definitions (DIDs) will specify requirements for RDA Server views. The contractor plan in response to the CDRL will indicate how RDA/SQL protocols will be used to provide the required integration among government, contractor, and subcontractor information processing facilities.

As an example, consider the same tables of information discussed in the Level 2 services above. For Level 3 services the government RFP may require that the four base tables (i.e., Project, Task, Employee, Timesheet) and one derived table (i.e., Salary_Expenses) be available as RDA Server views over appropriate contractor and/or subcontractor databases. If this is the case, the government can use its own IRDS to assist in the development of application programs to access the data and integrate the data with other government information systems. The government can use its own GUI presentation facilities to present the data to users. There is no requirement for the government to specify in advance what queries may be needed over the base information tables and there is no requirement for the government and the contractor to reach agreements on the GUI standards to be used or for the contractor to supply workstations and/or software to government users.

In other situations, the government may not require specific views of the data. Instead the government will include a Government Concept of Operations (GCO) as part of an RFP and the contractor will propose RDA Server views as part of its CITIS plan in response to that RFP. For example, a contractor may propose to represent all CITIS forms, documents, data item definitions (DIDs), and requirements in relational tables with an RDA/SQL Server interface. Authorized government users may browse through contractor-produced views of contract-specific data and use SQL queries to filter and rearrange the data to put it into the form most appropriate for further government access.

At the present time, the RDA standard specifies interchange protocols for transmitting rows of data from a server site to a client site, provided that the data items in the rows are either numbers or character strings. Near term RDA follow-on specifications will extend the data types handled to all of those specified in the forthcoming SQL2 specification; i.e., fixed and variable length character strings, fixed and variable length bit strings, fixed and floating point numerics, dates, times, timestamps, and intervals. Later RDA follow-on specifications will provide interchange mechanisms for the user defined abstract data types (ADTs) specified in the emerging SQL3 standard. RDA protocols do not by themselves provide interchange mechanisms for other CITIS data objects, so standards for IGES, STEP, CGM, SGML, EDI, ODA/ODIF, CCITT, etc. will remain critical for transmitting agreed object definitions among various government and contractor sites.

Level 4 -- Access to contractor applications

In addition to all of the above facilities, this level of service provides capabilities for government authorized users to use application software available on the contractor system for processing and analysis of data. For example, the contractor may have available powerful simulation, presentation, and analysis tools that generate and use data objects that are not easily transmitted between sites. The service shall enable users to specify the data to be acted on by the application (possibly by an ad hoc SQL query on the contractor's database) and to establish such input parameters, limits, and options as are appropriate to the application. In some cases the service will provide the capability for storing the results of such processing in the contractor's database or, if appropriate, for transmitting the results in an agreed output format to specific government sites.

From a data management perspective, this level of service can be viewed as an integration of Level 2 services for pre-defined access to contractor facilities and Level 3 services for ad hoc data management. Level 2 services could require that a specific data analysis tool be invoked over specific data with the result returned to the government user in an agreed format. Now combined with Level 3 services for remote interoperability, an authorized government user may specify via an SQL query what data is to be analyzed, direct that data to a chosen application tool, analyze the data through the eyes of that tool (e.g., sophisticated expert systems and design analysis tools), and specify where the result should be stored for further access by other tools. The interoperability capabilities provided by SQL and RDA allow unprecedented integration of contractor applications from various contractor and subcontractor sites.

This level of service will be able to take maximum advantage of SQL3 support for user-defined abstract data types (ADTs). The emerging SQL3 specification (see Appendix A.1) includes object-oriented capabilities over ADTs such as methods, object identifiers, subtypes, inheritance, and polymorphism, thereby making SQL3 a complete language for creating, managing, and querying persistent objects.

The existing IRDS provides the capability of describing contractor applications developed in third generation programming languages; the government can then develop third generation programs which may use SQL and RDA to run the contractor-supplied applications. In the future, the emerging IRDS2 specification is expected to be capable of defining objects constructed by means of object-oriented programming languages. With SQL3 to manage persistent objects, IRDS2 to describe object-oriented programs, and RDA to provide remote interoperability, the government user will be able to invoke contractor-supplied applications and to integrate them with government information systems in a nearly "seamless" manner.

5. CITIS Functional Requirements

The CITIS specification contains a number of functional requirements for information management and planning. Included among these are configuration and access controls that ensure the accuracy and integrity of all information, maintains relationships among various versions of data, allows migration of data from one status to the next, correlates data objects to the appropriate application and support environment, and provides appropriate audit and accountability controls. In the following sections, we address how data management standards are applicable to these various requirements.

5.1 Data access control

CITIS will be characterized by remote access to large, shared, distributed data banks which contain a combination of privileged, proprietary, unclassified, sensitive, and classified information. CITIS shall protect information with data rights and distribution limitations from unauthorized access and distribution. The service shall provide controls to ensure that each user is able to access only those data elements and applications for which that user is authorized access. An access rule set shall be used to control access to data based on such parameters as contract requirements, CITIS level of service, information type, information acquisition strategy, status level, type of access, classification and proprietary rights limitations, distribution limitation, and requestor's access profile.

The existing SQL-1989 standard supports data access control requirements by providing a simple security model for declaring and managing database access control. The model consists of three main entities: Objects, Actions, and Users. Objects are things defined in a database schema definition, actions are operations on objects, and users are the entities that invoke an action on an object. A privilege is an authorization of an action on an object to a user. A privilege is represented by an ordered tuple:

(Grantor, Grantee, Object, Action, Grantable)

The grantor is a user who is already authorized to perform the action on the object and who is authorized to grant that authorization to other users. The grantee is the user who receives the authorization from the grantor. Grantable is a yes/no variable indicating whether or not the grant authorization is passed from the grantor to the grantee.

When an object is created, some user is designated as the owner of the object. The owner is authorized to perform all actions on the owned object and to grant this authorization to other users. No user other than the owner of an object is authorized to perform any action on that object unless an explicit sequence of privileges can be traced from the owner to that user. The grantor of a privilege is also authorized to revoke that privilege from the grantee at a later time, thereby cascading a revoke action through the maze of subordinate privileges.

The emerging SQL3 specification defines an enhanced facility for database security management that builds upon the existing Grant and Revoke mechanism. It extends the security model entities to include Roles in addition to Objects, Actions, and Users. A role is a named collection of authorized actions on objects. The existing Grant and Revoke mechanism for privileges is used to assign privileges to and from roles. A user identified as the role administrator is authorized to manage roles and to grant and revoke roles from individual users. In addition, roles may be granted to other roles so that role hierarchies, without cycles, may be established to facilitate groups of application privileges.

A role authorization is represented by an ordered tuple:

(Grantee, Role, Admin)

The grantee is the user, or role, receiving the role authorization from the role administrator. Role identifies the role by its role name, which is unique in the database environment. Admin is a yes/no variable indicating whether or not the role administration authorization is passed from the role administrator to the grantee. Role authorizations differ from privileges in that the grantor of the role authorization is not a persistent component of the privilege. Thus the role authorization remains intact even if the role administrator changes. A role authorization can only be revoked by a current administrator for that role. Roles are not revoked as the side-effect of some other action as privileges sometimes are.

Roles support the management of privileges needed to execute individual applications. Authorizing a user to run an application often involves many dozens of separate privilege authorizations. Instead of database administrators granting and revoking privileges on individual tables to individual users, the privileges for each application can be assigned to a role and the roles granted and revoked from individual users. In practice, the simplicity of role management should result in fewer human errors and thus more effective security management. Roles also avoid the problem of security breaches caused by one user receiving privileges to execute several different applications but having the cumulative effect of those privileges much greater than intended. With roles, only one role is enabled at any one time, and the definer of the role has complete control over the set of privileges available to the user at any one time. A user is allowed to change from one authorized role to another using a Set Role statement.

Roles also support database enforcement of military security levels such as "confidential", "secret", etc. Each functional role, such as Payroll Clerk or Personnel Clerk, may also be distinguished by a security classification so that a separate role exists for a Payroll_Clerk_Confidential or a Personnel_Clerk_Secret, etc. In many cases roles can be used to significantly reduce the need to assign a security level classification to each data item in the database. Instead, only complete objects such as documents, drawings, memos, or strategies will require a security classification, whereas the constituent parts may not -- since only appropriate roles will have access to the constituent parts.

The access control facilities in the existing SQL standard are more than sufficient to address CITIS access control requirements. The use of SQL for database definition and access will make it possible to specify and maintain an integrated collection of privileges for various users scattered throughout a communications network. The role definition capabilities in the emerging SQL3 specification will provide additional convenience and flexibility for access control.

5.2 Data integrity control

The contractor shall document, in the CITIS implementation planning, steps taken to ensure the integrity, accuracy, and consistency of data. Data should be generated once and made available for use across all operations without requiring redundant data occurrences. Access to obsolete versions of the data shall be prohibited.

The SQL-1989 standard offers an optional "integrity enhancement facility" for maintaining data accuracy and consistency. This facility consists of column constraints, row constraints, table constraints, and referential integrity constraints. All of these constraint options become required facilities in the forthcoming SQL2 standard.

Column constraints allow the user to specify restrictions on the data values that can be stored in a column of a table. These include the typical range constraints of many programming language data types as well as more complex Boolean expressions using the predicate calculus and arithmetic and string expressions defined in SQL. Row constraints allow the same type of predicates to be specified among column values for individual rows in a table. Table constraints allow subqueries to be used so that values in one row of a table can be related to values of a different row in the same table or to values from rows of a different table.

The referential integrity constraint makes it possible to declare a "foreign key" in one table that references the "primary key" of another, possibly the same, table. The implementation is required to ensure that no rows in the foreign key table have a foreign key value different from some primary key value in the primary key table. These facilities make it possible to maintain internal consistency in a large relational database by ensuring that modifications to primary key values are properly propagated to the remainder of the data base. Other "cascade" facilities in the forthcoming SQL2 standard ensure that delete and update operations in the primary key table are properly cascaded to foreign key values in the foreign key tables. These facilities make it possible to maintain internal consistency in a large relational database and to ensure that there are no dangling references to non-existent objects in the database.

The emerging SQL3 standard specifies enhanced capabilities for database integrity control via "assertions" and "triggers". An assertion is an integrity constraint that is triggered by a specific database action. Assertions are more powerful than the static integrity checks in SQL2 because they introduce the notion of before and after images of the database and because they can be activated at specific times. A trigger is a sequence of database actions that is initiated by a specific database action. In the current SQL3 specification, a trigger can be activated by an insert or delete statement on an underlying table, or by an update on specified columns. When a trigger is activated, a When clause determines if the list of actions should be executed. Assertions and triggers may be used for sophisticated access and integrity control over application objects.

The integrity control facilities in the existing SQL standard are more than sufficient to address CITIS integrity control requirements. The use of SQL for database definition and access will make it possible to specify and maintain an integrated collection of "static" integrity controls for data distributed across a communications network. The assertion and trigger capabilities specified in the emerging SQL3 standard will provide additional convenience and flexibility for "dynamic" integrity control as that need arises.

5.3 User training and orientation

The contractor shall provide training, orientation, and familiarization services for government users of CITIS and shall maintain and revise/update these services to keep current with system changes. For each CITIS level of service available to the user, training will provide users with knowledge of system capabilities, operating parameters, documentation, and hands-on experience in accessing, manipulating, and copying or otherwise transferring authorized CITIS data.

An obvious advantage of using data management standards is the significant reduction in training costs. A feature once learned will be usable across all data repositories maintained by all contractors and subcontractors. System changes will be minimized because future standards will build upon and always be upward compatible with existing standards. Even vendor supplied "extensions" (e.g., presentation and report writer tools) will derive from basic facilities in the underlying standards.

5.4 CITIS transition requirements

It is the intent of the government to ensure a viable option for the smooth transition of data and information from one fully operational CITIS operated by the contractor to an equally useful and responsive service operated by the government or its designated representative. Transition must emphasize continuity and consistency during transition and ensure responsiveness, affordability, and integration subsequent to transition. The contractor shall prepare a transition plan to ensure a smooth transition of CITIS databases and shall provide required demonstrations to validate system transition.

Smooth transition of databases from one contractor to another is virtually impossible without data management standardization. Using SQL and RDA the database schemas and/or data occurrences from one contractor's database can easily be transferred to another. If applications are developed in standard programming languages with standardized SQL calls to underlying databases, then complete applications will also be transportable from one contractor to another. If the applications are adequately documented in IRDS, then maintenance of them by the new contractor will be much simpler and more reliable. Even more valuable than database transition, may be the database interoperability achieved with use of data management standards. It may be unnecessary to actually move data and applications from one site to another if the data and applications are available to all participants in a contract effort. Such application sharing is a major goal of Level 4 CITIS.

Another requirement for CITIS transition is the delivery of a data dictionary. The contractor is required to assist the government in the development of appropriate data directory(ies) and data dictionary(ies) for management of data subsequent to transition. Use of the IRDS in the delivery of CITIS services would fulfill this requirement without additional effort or costs.

5.5 Distributed database requirements

Data management requirements for CITIS applications will likely exceed the capabilities of existing database management systems. Future CITIS applications will require a logically integrated database of diverse data (e.g. documents, graphics, alphanumeric records, complex objects, images, voice, video) stored in geographically separated data banks under the management and control of heterogeneous database management systems. An over-riding requirement is that these various data managers be able to communicate with each other and provide shared access to data and data operations and methods under appropriate security, integrity, and access control mechanisms.

Distributed database management, in its most demanding definition, implies totally integrated, distributed data, under the coordinated control of multiple heterogeneous database management systems. True distributed database is still a research consideration and, because it requires cooperating concurrency control managers, "standardized" distributed database management may be some time away.

In the absence of standards for "true" distributed database, the standard for Remote Database Access (RDA) can be used, along with other OSI standards, to establish a remote connection between a database client and a database server. The RDA communications protocols are defined in terms of OSI standards for Association Control (ACSE), Remote Operations (RO), Transaction Processing (TP), and Commitment, Concurrency and Recovery (CCR). With these standards, near term distributed processing and interoperability are possible. Distributed processing, a weaker definition than distributed database, implies client/server access to remote sites, with full capabilities for data definition and data manipulation, and for standardized transfer of data parameters and query responses across the communication line, but with management of multiple remote sites the responsibility of the client process. The TP standard (ISO/IEC 10026) allows both one-phase and two-phase commit protocols. This allows the client component of one database management system to access remote data from multiple remote sites, and then present a view of that data to a local user as if it were local data. In this way, standardized distributed processing can be used to simulate true distributed database.

Future standards for true distributed database, and multi-vendor "cooperating" concurrency control, can be built on top of existing SQL and RDA standardization efforts. Therefore, the SQL and RDA efforts are vital building blocks for future distributed database management standards.

6. Conclusions

CITIS is a contractor-provided service that makes available to the government authorized access to contractor databases and applications, including both business and technical information. CITIS enables the government to acquire electronic access to these databases and applications. Data accessed through this service may include any combination of information required by the government from the system prime contractor(s), associate contractors, subcontractors, suppliers, and vendors.

When the government specifies CITIS in the requirements of a procurement, the government can also specify data management standards that are to be used in conjunction with CITIS. These data management standards enhance the quality and usability of the information provided, and aid transition to future CITIS environments. Data management standards provide facilities for defining, managing, and protecting data. The data management standards recommended for use with the CITIS specifications are Database Language SQL, Remote Database Access (RDA), and Information Resource Dictionary System (IRDS).

This report provides an overview of the capabilities and status of these data management standards, and analyzes their application to CITIS. The SQL standard in particular is leading the way toward unprecedented portability of database applications. The emerging RDA standard promises to complete the link among SQL products from different vendors, thereby leading to true open systems interconnection and interoperability. The existing IRDS standard provides utilities for managing metadata. Emerging specifications for future revisions of SQL, RDA, and IRDS promise enhanced facilities to support object-oriented and knowledge-based applications in a distributed processing environment.

This report analyzes how data management standards relate to various CITIS levels of service and functional requirements, and provides guidance to contractors in building a CITIS with data management standards. The report concludes that SQL and IRDS are appropriate and helpful at all levels of service, even if they are not always required to achieve minimum CITIS requirements. SQL is particularly appropriate to address the CITIS functional requirements for access and integrity control.

Beginning at CITIS Level 3 and continuing into Level 4, it is virtually impossible to achieve CITIS objectives without requiring SQL and RDA as the basis of communication and interoperability. With emerging standards for SQL3 and IRDS2, the "seamless" interoperability goals of CITIS Level 4 are well within reach.

Finally, this report provides guidance to federal agencies in specifying SQL, RDA, and IRDS data management standards requirements with their CITIS requirements. This report includes candidate wording for procurement, software development, validation requirements, and correction of deficiencies for SQL, RDA, and IRDS.

Appendix A -- Data Management Standards

A.1 Database Language SQL

Existing Specifications: ISO 9075:1989

ANSI X3.135-1989
ANSI X3.168-1989

FIPS 127-1 February 2, 1990

Emerging Specifications: ISO/IEC DIS 9075:199x, SQL Revision
ANSI dpANS X3.135-199x, SQL Revision
(Expected 1992)

Sponsoring Organizations: ISO/IEC JTC1, ANSI X3, and NIST

Description

Database Language SQL specifies data definition, data manipulation, integrity checking, and other associated facilities of the relational data model. In addition, SQL specifies components that support access control, programming language interface, and data administration. The basic structure of the relational model is a table, consisting of rows and columns. Data definition includes declaring the name of each table to be included in a database, the names and data types of all columns of each table, constraints on the values in and among columns, and the granting of table manipulation privileges to prospective users. Tables can be accessed by inserting new rows, deleting or updating existing rows, or selecting rows that satisfy a given search condition for output. Tables can be manipulated to produce new tables by cartesian products, unions, intersections, joins on matching columns, or projections on given columns.

SQL data manipulation operations may be invoked through a cursor or through a general query specification. The language includes all arithmetic operations, predicates for comparison and string matching, universal and existential quantifiers, summary operations for max/min or count/sum, and GROUP BY and HAVING clause to partition tables by groups. Transaction management is achieved through COMMIT and ROLLBACK statements.

The standard provides language facilities for defining application specific views of the data. Each view is the specification of database operations that would produce a desired table. The viewed table is then materialized at application execution time.

The SQL standard provides a Module Language for interface to other languages. Each SQL statement may be packaged as a procedure that can be called and have parameters passed to it from an external language. A cursor mechanism provides row-at-a-time access from third generation programming languages since they can only handle one row of a table at one time.

Access control is provided by GRANT and REVOKE statements. Each prospective user must be explicitly granted the privilege to access a specific table or view using a specific statement.

The optional SQL Integrity Enhancement facility offers additional tools for referential integrity, CHECK constraint clauses, and DEFAULT clauses. Referential integrity allows specification of primary and foreign keys with the requirement that no foreign key row may be inserted or updated unless a matching primary key row exists. Check clauses allow specification of inter-column constraints to be maintained by the database system. Default clauses provide optional default values for missing data.

An Embedded SQL specification (ANSI X3.168) provides the SQL interface to programming languages, specifically Ada, C, COBOL, FORTRAN, Pascal, and PL/I, by specifying the effect of syntactically embedding SQL statements into otherwise conforming application programs. Applications may thereby integrate program control structures with SQL data manipulation capabilities. The Embedded SQL syntax is just a shorthand for an explicit SQL Module accessed from a standard conforming programming language.

Applicability to CITIS

The purpose of a database language standard is to provide portability and interoperability of database definitions and database application programs among conforming implementations. Database Language SQL is appropriate for all database applications where data will be shared with other applications, where the life of the application is longer than the life of current equipment, or where the application is to be understood and maintained by programmers other than the original ones. The SQL standard is particularly appropriate for database applications requiring flexibility in the data structures and access paths of the database. It is desirable both for applications under production control and when there is a substantial need for ad hoc data manipulation.

Use of the SQL standard is appropriate in all cases where there is to be some interchange of database information between systems. The schema definition language may be used to interchange database definitions and application specific views. The data manipulation language provides the data operations that make it possible to interchange complete application programs. Used with the RDA remote database access standard or with a generic data interchange standard such as ASN.1, data occurrences may also be transferred in a standard manner.

Product availability

Numerous implementations of the existing SQL standard exist on all levels and brands of equipment. Vendors are vigorously implementing additional features from the proposed 1992 revision. An SQL validated processor list is available from NIST (see next paragraph).

Conformance testing

Version 2.0 of the NIST SQL Test Suite has been publicly available since December 1989, and a formal SQL test service was instituted by NIST in April 1990. Version 3.0 of the NIST SQL Test Suite was released in January 1992. The SQL test suite measures conformance to both required and optional features of FIPS 127-1. See Appendix B, NIST Validation Testing, for additional information about SQL testing and Appendix C for SQL procurement wording.

Future plans

An emerging SQL2 specification, with features for schema manipulation, dynamic SQL, exception handling, enhanced integrity constraints, transaction management and data administration is expected to be adopted by ANSI and ISO in 1992. Further enhancements, including tools for the support of object-oriented and knowledge based systems, are expected in the 1995 time frame. The following two paragraphs give a more detailed description of emerging SQL2 and SQL3 specifications.

Emerging SQL2 Specification

A substantial upward compatible enhancement of the existing SQL standard, often called SQL2, has already been specified by the ANSI and ISO SQL development committees. It will standardize a number of SQL features not included in the original specification because they were not commonly available in SQL products. The technical specification for SQL2 is quite stable; only a very few features are controversial and subject to modification. SQL2 is intended to be a superset of SQL-1989 that replaces the existing SQL standard. Review copies of SQL-199x are available after May 1991 from GLOBAL Engineering (reference dpANS X3.135-199x) or from OMNICOM (reference ISO/IEC DIS 9075:199x, SC21 N5739).

SQL2 significantly increases the size of the SQL language to include a schema manipulation language for modifying or altering schemas, schema information tables to make schema definitions accessible to users, new facilities for dynamic creation of SQL statements, and new data types and domains. Other new SQL2 features include outer join, cascade update and delete referential actions, set algebra on tables, transaction consistency levels, scrolled cursors, deferred constraint checking, and greatly expanded exception reporting. SQL2 also removes a number of restrictions in order to make the language more flexible and orthogonal.

SQL2 was registered as an ISO/IEC Draft International Standard (DIS) and as an ANSI dpANS in early 1991 and is currently undergoing an ISO/IEC national body ballot for final adoption as an international standard. Parallel processing to adopt the identical specification as an American National Standard is taking place in the United States. Adoption as a revised FIPS PUB 127 is expected in calendar year 1992.

Emerging SQL3 Specification

A second substantial SQL enhancement, often called SQL3, is under development by ANSI and ISO SQL specification committees, with publication expected in the 1995 time frame. SQL3 is a forward looking SQL enhancement that intends to provide additional facilities for managing object-oriented data and for forming the basis of "intelligent" database management systems. It includes generalization and specialization hierarchies, multiple inheritance, abstract data types and methods, object identifiers, polymorphism, triggers and assertions, support for knowledge based systems, recursive query expressions, additional data administration tools, standardized database export/import facilities, and progress toward distributed database management. Proposals are under consideration that would make SQL a computationally complete language with stored procedures, looping and branching control structures, and Ada-like exception handling. These features are preliminary and subject to significant modification and improvement before final adoption. An early draft of these specifications should be available from OMNICOM after February 1992 as a Working Draft SQL Revision.

A.2 Remote Database Access (RDA)

Existing Specifications:	None.
Emerging Specifications:	ISO/IEC CD 9579-Part1 Generic RDA (Expected late 1992)
	ISO/IEC CD 9579-Part2 SQL Specialization (Expected late 1992)
Sponsoring Organization:	ISO/IEC JTC1/SC21

Description

Remote Database Access (RDA) provides standard protocols for establishing a remote connection between a database client and a database server. An emerging RDA standard will standardize distributed processing in a "client/server" SQL environment; i.e., an environment with standard-conforming RDA/SQL "servers" at each remote node in a communications network. RDA specifies a two-way connection between a client and a server, as well as transfer syntax and semantics for SQL database operations. The client is acting on behalf of an application program or remote process, while the server is interfacing to a process that controls data transfers to and from a database. The communications protocols are defined in terms of OSI standards for Association Control (ACSE), Remote Operations (RO), Transaction Processing (TP), and Commitment, Concurrency and Recovery (CCR). The goal is to promote distributed processing by standardizing the interconnection among SQL database applications at non-homogeneous sites.

The RDA specification consist of two parts, a Generic RDA for arbitrary database connection and an SQL Specialization for connecting databases conforming to Database Language SQL. Both parts were registered as ISO/IEC Draft International Standard (DIS) in 1991 and will soon undergo ISO/IEC national body ballots for final adoption as international standards. Formal approval is expected in late 1992.

Generic RDA

The proposed Generic RDA standard provides an RDA Service Interface and an RDA Communications Element that exist at both the client and server sites. The Generic RDA Service Interface consists of service elements for association control, for transfer of database operations and parameters from client to server, for transfer of resulting data from server to client, and for transaction management. Association control includes establishing an association between the client and server remote sites and managing connections to specific databases at the server site. Transaction management includes capabilities for both one-phase and two-phase commit protocols. A common mistake is to assume that RDA is a distributed database specification. Although distributed extensions are under consideration, existing RDA is only a two-way service and protocol specification.

The RDA Communications Element at the client site converts RDA service requests into the appropriate underlying RO, ACSE, and TP protocols as part of an open systems interconnection. The RDA Communication element at the server site converts received protocols into requests to its underlying database management system.

The Generic RDA service does not specify the syntax or semantics of database operations sent from client to server. Instead, the standard assumes the existence of a language specialization (e.g., IRDS or SQL) that specifies the exact transfer syntax for standard operations.

SQL Specialization

The RDA/SQL Specialization complements the Generic RDA standard for use when a standard conforming SQL data manager is present at the server location. The client site may also have an SQL conforming data manager, but this is not required. The client processor transforms user requests into the appropriate standard protocols for transmission across the network to the server site. SQL data operations are sent as character strings conforming to the SQL language and are packaged in an RDA/ASN.1 envelope that allows for embedded parameter values to be send with the data operation.

The result of an SQL statement will contain status code and exception code parameters and may contain one or more rows of data in response to a query. The transfer syntax for all such data is specified in ASN.1 as part of the SQL Specialization standard.

Applicability to CITIS

RDA is used to establish a remote connection between a database client, acting on behalf of an application program, and a database server, interfacing to a process that controls data transfers to and from a database. The goal is to promote interconnection of database applications among heterogeneous environments, with emphasis on an SQL server interface. It is expected that the RDA/SQL Specialization will become the basis for all interconnection among SQL database management products from different vendors. Interconnection among database products from the same vendor will likely continue to use vendor specific communication and interchange forms.

Product availability

There are no known existing commercial RDA implementations, but many SQL vendors are planning to have conforming client and server products available before final adoption as a standard in the 1992 time frame. The SQL Access Group publicly demonstrated interoperability among multiple SQL implementations on July 16, 1991, in New York City.

Conformance testing

RDA will likely become a future optional part of conformance testing for GOSIP. At the present time, RDA can be tested indirectly using the NIST SQL Test Suite with application programs at the client site and data at the server site. See Appendix B, NIST validation testing, for additional information, and Appendix C for RDA procurement wording.

Future plans

Enhancement projects for distributed database and stored database procedures have already been proposed to ISO. Extensions to support new features in evolving SQL standards are under development.

A.3 Information Resource Dictionary System (IRDS)

Existing Specifications: ANSI X3.138-1988

FIPS 156 April 5, 1989

Emerging Specifications: ISO/IEC DIS xxxx:199x,IRDS Framework
ANSI dpANS X3.138-199x, IRDS Revision
(Expected 1995);
Export/Import File Format, ANSI dpANS X3.195-199x
(Expected late 1991 or 1992);
Services Interface, ANSI dpANS X3.185-199x
(Expected in 1992).

Sponsoring Organizations: ISO/IEC JTC1, ANSI X3, and NIST

Description

IRDS services consist of utilities and systems necessary to catalog, document, manage, and use metadata (information about data). IRDS specifies a computer software system that provides facilities for recording, storing, and processing descriptions of an organization's significant data and data processing resources. The IRDS includes the functions performed by data dictionary systems or information repositories. The standard specifies two user interfaces: the full syntax and semantics of a Command Language, and the semantics of a menu-driven Panel Interface.

Applicability to CITIS

An information resource dictionary system is intended to serve as the central repository of current information about an organization's data. Implementations of this standard are required in information resource management applications that are either developed or acquired for federal government use. These applications include development, modification, and maintenance of manual and automated information systems throughout the lifecycle; support to an agency defined data element standardization program; support to records, reports and form management; and management of schema and subschema definitions for database management systems.

The specifications of this standard are not applicable to those dictionary systems embedded within a database management systems. This is the case when the dictionary function is part of the data definition function of the DBMS, and the dictionary data is stored as part of the database for the database management system.

The next few years should see nominal changes in the current standard. Related standards efforts are specifying additional functionality. Two on-going efforts will make IRDS active (i.e., provide interfaces to other software). First, a draft proposed IRDS Export/Import file format (dpANS X3.195-199x) is expected to become an ANSI standard late in 1991 or in 1992 and a FIPS in 1992; this will be appropriate for schema and metadata interchange among IRDS compliant databases, among IRDS and CASE tools with repositories or dictionaries, and between IRDS and application programs. Second, a draft proposed IRDS Services Interface (dpANS X3.185-199x) is expected to become an ANSI standard in 1992; this will be appropriate for metadata interchange with a database management system, and between IRDS and application programs.

Product availability

Commercial implementations have been developed, but their quality has not yet been determined. A prototype implementation is available from CSL which contains a large subset of IRDS functionality.

Conformance testing

Conformance tests are currently under development for testing conformance to FIPS 156. A beta version, limited to the IRDS Command Language, is planned to be released for testing in the first half of calendar year 1992. The official release of the test system is planned to be released later in 1992. See Appendix C for IRDS procurement wording.

Future plans

Related standards work will provide a software interface, enhanced functionality and capability to manage object-oriented data structures, and provide for communication of information between applications and other data management tools. This new functionality will constitute a major revision to the standard; the revision is commonly called IRDS2 and is expected to be available about 1995.

Appendix B -- NIST validation testing

NIST's Computer Systems Laboratory assists federal agencies in procuring conforming products by providing various test suites and test services. Each FIPS has a different history in the development of the test suite(s), the availability of testing services, the policies for certification, the treatment of nonconformities, and the reporting of test results.

In the area of data management standards, the Software Standard Validation Group at CSL provides testing for Database Language SQL as well as the programming languages COBOL, FORTRAN, Ada, and Pascal. Test suites for programming language C and the IRDS are currently under development by CSL.

The testing programs for the programming languages COBOL, FORTRAN, Ada, and Pascal have been in place for some time. Test results are published in the quarterly Validated Products List. Certificates of validation are issued for tested products which meet the published criteria. For Ada, certificates are issued only for products without deficiencies. For the other programming languages, certificates are awarded even for products "with nonconformities" but are not reissued or extended unless the nonconformities are removed. Products with nonconformities are duly noted on the certificate.

The compiler testing program is very successful. Vendors routinely schedule their compilers for annual testing to maintain their certification. Federal agencies, as well as private industry and the international community, refer to the Validated Products List in making procurement decisions.

The testing program for SQL began in April 1990 and is in a "Trial Use Period." During this period, testing procedures and policies are under review and may change. Also, during this period, the list of tested products is growing. With each release of the Validated Products List, additional SQL implementations are shown conforming to the standard. This puts pressure on the remaining SQL implementations to be tested in order to compete for federal procurements.

As a testing laboratory, CSL will validate both production software and pre-release software. The pre-release software, which may be offered for procurements when it attains production status, is tested early so that SQL vendors can comply with RFP requirements for testing well in advance of delivery. It is also helpful for federal agencies planning an RFP to know about product availability in order to formulate a strategy for getting conforming SQL implementations as early as possible at off-the-shelf prices.

B.1 SQL Testing

A suite of automated validation tests for SQL implementations was developed in-house by CSL and is currently available from CSL. The original Version 1.1 of the NIST SQL Test Suite was offered in August 1988. Version 1.2 was released in May 1989, followed by Version 2.0 in December 1989. Version 3.0 was released in January 1992 and will be used in the testing service in 1992. Each release increases the test suite's coverage of the standard.

Version 2.0 contains eight programming language test suite types: embedded (preprocessor) COBOL, embedded FORTRAN, embedded C, embedded Pascal, module language COBOL, module language FORTRAN, module language C, and module language Pascal. An Interactive SQL test suite is also included. For each test suite type there is an optional testing module which evaluates conformance to the Integrity Enhancement Feature. Version 3.0 contains additional test suites for embedded Ada and module language Ada.

Over 50 organizations hold licenses for Version 2.0 of the test suite. SQL implementors purchase a license to use the NIST SQL Test Suite to assist in evaluating the conformance of their products to FIPS 127-1. After in-house testing, and hopefully after correcting deficiencies, SQL implementors may request CSL to perform a formal on-site validation.

Typically, an SQL implementation will be tested in all the programming language interfaces which are supported. If the Integrity Enhancement Feature is supported, the optional testing modules for that feature will be executed. The decision of which testing modules to execute is based on the vendor's claim of conformance. The vendor completes a "Supplier's Declaration of Conformance" to identify the SQL interfaces to be tested and the hardware/software environment in which the testing takes place. This information will be published in the Validated Products List.

The vendor may also complete a "Derived Validation" declaration to identify other hardware/software environments where the vendor certifies that identical test results will be obtained. This claim must be based on testing (actual execution of the test suite or other testing) and is restricted to processors with binary level compatibility. This claim, if accepted, will be published in the Validated Products List under the category of "other environments."

The decision to validate is often a marketing decision, although the technical work to conform may have begun years before. Since considerable resources may be required for a requestor to prepare for a formal validation, it is usually not performed until required by a Request for Proposal (RFP) or until the advertising benefits of a validation can be justified. Although it is possible for an SQL product to conform but not to have been formally tested, formal testing is the user's best assurance of a vendor commitment to achieve conformance to the required standards. Vendors are also required to fix any detected errors within one year after testing is completed.

Certificates for SQL conformance are not offered during the Trial Use Period. Instead, a Registered Validation Summary Report is produced to document the results of the testing. At the end of the trial use period, criteria for certification will be published and certificates will be offered. In the interim, federal agencies are advised (1) to require a Registered Validation Summary Report, (2) to review the reported deficiencies for impact on current programs, and (3) to specify timeframes for correction of the deficiencies.

Given that there is high interest in conformance among SQL vendors, most of the major SQL products have been evaluated in-house. It is reasonable to assume that vendors know how well their products perform on the NIST SQL Test Suite, even though the products have not been validated by CSL. In-house test results may be shown by vendors to customers requiring a demonstration of conformance to FIPS 127-1. These results are valid, providing that testing has been conducted according to approved procedures and that the reporting does not contain any misrepresentations.

Clearly, there are benefits to customers in validation by a third party laboratory trained in the specific test method to be used. However, formal on-site validation can be expensive and time consuming. It is simply not feasible for CSL to validate every release of every DBMS product on every hardware/software platform offered for sale to the federal government. For SQL testing, CSL is investigating whether some combination of periodic CSL validation and interim vendor self testing might be the most effective way to assist federal agencies in buying conforming products.

B.2 RDA Testing

Testing for RDA is premature, since RDA is only a Draft International Standard. A FIPS PUB for RDA could possibly be finalized as early as August 1992, when ANSI and ISO standards are published.

The next version of GOSIP, GOSIP 3, may contain specifications for the "RDA Basic Application Context." "RDA Transaction Processing Application Context" will be deferred to GOSIP 4. The former consists of RDA and ACSE. (One-phase commit transaction management is provided by the RDA Service.) The latter consists of RDA, TP, ACSE, and optionally CCR. (When required, two-phase commit transaction management is provided by the TP service.)

There is considerable interest in RDA testing in the private sector. One vendor consortium, SQL Access Group, seeking to promote rapid standardization of RDA has already demonstrated limited interoperability of heterogeneous database management systems. This demonstration used prototype software based on the RDA specification. It is probable that this vendor consortium will pool resources to develop an RDA test suite.

RDA testing will be different from SQL testing. It may be desirable to have a laboratory with a "standard" server to assist in testing the RDA Client Interface of client software and a "standard" client to assist in testing the RDA Server Interface of server software. In order to test the SQL Specialization of RDA, another test module for SQL must be created to validate those features of the SQL2 standard which are required for RDA but are not included in the SQL-1989 standard and hence are not tested in the current NIST SQL Test Suite.

It is not yet clear whether RDA testing should fall under the scope of GOSIP testing and operate under NVLAP accreditation procedures.

B.3 IRDS Testing

A suite of automated validation tests is being finalized by CSL to test implementations of the IRDS Command Language. A similar system to test conformance to the IRDS Panel Interface is being designed. The CSL Command Language tests are planned to be released for beta testing in the first half of 1992. The first official release of the IRDS test suite is planned for the second half of 1992. Initially, vendors will use the IRDS test suite in-house to evaluate the conformance of their IRDS implementations. During this period, CSL will evaluate the availability of IRDS implementations and the feasibility of a test service.

Appendix C -- Procurement wording

Insofar as CITIS attempts to utilize vendors' existing databases and data dictionaries, it may not always be practical to require conformance to the FIPS. However, if additional general-purpose software is purchased with CITIS funds or if applications are developed for CITIS reporting requirements, then the appropriate FIPS should be specified.

If the federal agency requires remote database access from agency computers running (1) applications written in standard programming languages or (2) applications based on 4GL tools capable of remote database access, then the federal agency should specify an RDA/SQL interface to CITIS data. Additionally, if the federal agency requires an on-line ad hoc query capability for federal analysts trained in SQL or IRDS, then the federal agency should specify an Interactive SQL or an IRDS interface to CITIS data.

Other alternatives for SQL, such as pull-down menus and full-screen displays may be more desirable (although nonstandard), and may be acceptable if they do not create a training problem. Other alternatives for IRDS, such as an implementor-defined command language or a panel interface to nonstandard data dictionary functionality, may be acceptable as an interim solution. If the federal analyst needs to access several dissimilar CITIS databases, then this is a training problem, and the need for standards is more stringent. In general, it is to a vendor's advantage to use FIPS-conforming software whenever possible, for the same reasons the federal government uses standards: portability, interoperability, training, etc.

Remote database access of multimedia data objects (such as graphics, raster-scanned images, documents, video images, and voice files) raises compatibility issues. Many data objects can be satisfactorily exchanged through data interchange standards such as IGES (planned FIPS), STEP (Draft Proposed ISO 10303), CGM (FIPS 128), SGML (FIPS 152), and ODA/ODIF/ODL (ISO 8613:1989). Transmission of CITIS data objects using these data interchange standards is pointless, of course, unless the receiving federal computer has software to process (or store for later use) data objects in the applicable standard data interchange formats.

Today, transmission of multimedia data objects is typically accomplished through the use of homogeneous DBMS software or through special gateway software written for a specific pair of client and server systems. In the future, standards for remote database access and data interchange formats will enable broad remote database access across heterogeneous environments. However, it is likely that initial implementations will be limited in functionality and will lack performance optimization.

This section provides procurement language for SQL, IRDS, and RDA general-purpose software, as well as for applications based on these standards. The CITIS procurement should incorporate the following procurement wording, as appropriate.

C.1 FIPS 127-1 for Relational Database Procurement

If a relational database management system is to be procured, include the following wording in the RFP:

"Acquisition of Database Language SQL Processors"

"SQL language processors offered as a result of the requirements of which this is a part shall conform to the requirements in FIPS 127-1 Database Language SQL. These processors shall implement (1) all of the required language elements of FIPS 127-1 not previously covered by a waiver, (2) all of the FIPS 127-1 options specified elsewhere in this requirements document, as well as, (3) all default options required by Section 13 of FIPS 127-1, Special Procurement Considerations. These processors shall also implement any additional language elements specified elsewhere in this document [insert reference]. Claims of conformance to FIPS 127-1 shall be supported by [insert (choose one) 'Delayed Validation', 'Prior Validation Testing' or 'Prior Validation'] as specified elsewhere in this document [insert reference]."

Complete the following wording to select among various alternatives specified in Section 13 of FIPS 127-1.

"Special Procurement Considerations for FIPS 127-1"

		YES	NO
Integrity enhancement feature is required?		__	__
Programming language interfaces			
Embedded SQL interface to	Ada	__	__
Module Language SQL interface to	Ada	__	__
Embedded SQL interface to	C	__	__
Module Language SQL interface to	C	__	__
Embedded SQL interface to	COBOL	__	__
Module Language SQL interface to	COBOL	__	__
Embedded SQL interface to	FORTRAN	__	__
Module Language SQL interface to	FORTRAN	__	__
Embedded SQL interface to	Pascal	__	__
Module Language SQL interface to	Pascal	__	__
Interactive SQL required? (Indicate any alternative or additional requirements for direct invocation of SQL statements.)		__	__
Sizing for database constructs in 127-1 to be modified? (If YES, indicate alternate or additional minimum requirements.)		__	__
Character set and collating sequence required to be 95-character graphic subset of ASCII (FIPS 1-2)? (If NO, specify required character set and collating sequence or indicate whether implementor-specified character data values are acceptable.)		__	__

C.2 FIPS 127-1 for Development of Database Applications

If software is being developed or acquired for CITIS, include the following wording in the RFP:

"Development or Acquisition of Application Programs"

"When computer application programs are developed or acquired as a result of the requirements of which this is a part, and the FIPS 127-1 Database Language SQL is specified elsewhere in this requirements document [insert reference to FIPS SQL here], only the language elements of FIPS 127-1, as well as any additional language elements as specified elsewhere in this document [insert reference] shall be used. In these cases, processors used in developing such programs shall be validated."

C.3 FIPS 156 for Data Dictionary Procurement

"Acquisition of Information Resource Dictionary System (IRDS)"

"The data dictionary offered as a result of the requirements of which this is a part shall conform to the requirements set forth in FIPS 156 and shall implement all of the required modules not previously covered by a waiver, as well as any additional features as specified elsewhere in this document [insert reference here]. Claims of conformance to FIPS 156 shall be supported by [insert (choose one) 'Delayed Validation', 'Prior Validation Testing' or 'Prior Validation'] as specified elsewhere in this document [insert reference]."

C.4 FIPS 156 for Development of Data Dictionary Applications

If data dictionary applications are being developed or acquired for CITIS, include the following wording in the RFP:

"Development or Acquisition of Applications"

"When data dictionary applications are developed or acquired as a result of the requirements of which this is a part, and FIPS PUB 156 is specified elsewhere in this requirements document [insert reference to FIPS IRDS here], only the Command Language or Panel Interface of FIPS 156, as well as any additional features as specified elsewhere in this document [insert reference] shall be used. In these cases, processors used in developing such applications shall be validated."

C.5 DIS 9579 for Remote Database Access

"Acquisition of Remote Database Access Software"

"The database [insert FIPS SQL or other database server reference] {and/or database interface software [insert FIPS SQL or other database client/precompiler or user interface software]} offered as a result of the requirements of which this is a part shall have an RDA component which conforms to the RDA Server Interface {and/or RDA Client Interface} specification set forth in ISO/IEC DIS 9579-1 (Generic RDA) and 9575-2 (SQL Specialization) and shall implement all of the required language elements not previously covered by a waiver, as well as any additional language features as specified elsewhere in this document [insert reference here]. Claims of conformance to DIS RDA shall be supported by [insert acceptance criteria]."

C.6 DIS 9579 for Remote Database Access Application Development

If database client or user interface software is being developed or acquired for CITIS to access a remote database server which supports the RDA Server Interface, include the following wording in the RFP:

"Development or Acquisition of RDA Applications"

"When database client applications are developed or acquired as a result of the requirements of which this is a part, only the language elements which conform to the RDA Client Interface specification set forth in ISO/IEC DIS 9579-1 (Generic RDA) and 9575-2 (SQL Specialization), as well as any additional language elements as specified elsewhere in this document [insert references to RDA, SQL, programming languages, etc.] shall be used. In these cases, processors used in developing such applications shall be tested [insert acceptance criteria]."

If CITIS data is to be accessed remotely by applications which support the RDA Client Interface, include the following wording in the RFP:

"Development or Acquisition of RDA Servers"

"When database server applications are developed or acquired as a result of the requirements of which this is a part, the RDA component of the database server shall conform to the requirements of the RDA Server Interface specification set forth in ISO/IEC DIS 9579-1 (Generic RDA) and 9575-2 (SQL Specialization), as well as any additional language elements as specified elsewhere in this document [insert references to

RDA, SQL, etc.] shall be used. In these cases, processors used in developing such applications shall be tested [insert acceptance criteria]."

C.7 Choosing a Validation Option

A federal agency requiring conformance to FIPS 127-1 must specify (1) the acceptance criteria for conformance and (2) timeframes for delivery of conforming software. Three validation options are specified "Delayed Validation," "Prior Validation Testing," and "Prior Validation." However, federal agencies may choose to require other acceptance criteria. For example, an agency with a small procurement may decide that an in-house analysis of the product by the vendor, using the NIST SQL Test Suite, will be initially acceptable providing that the product is validated within 18 months and that all deficiencies are corrected at that time. The agency may then review the vendor-provided analyses to determine whether any existing deficiencies are unacceptable for interim use of the SQL software.

The following guidance is provided to assist agencies choose among the three validation options:

DELAYED VALIDATION--When an agency determines that the nature of the requirement is such that implementations of a FIPS may be offered that have not yet been tested for conformance to that FIPS the 'Delayed Validation' solicitation wording option shall be included.

PRIOR VALIDATION TESTING--When an agency determines that it is essential for implementations of a FIPS to be previously tested for conformance to that FIPS before being offered, and the nature of the requirements is such that an implementation of a FIPS may be initially offered that has not yet been fully validated (i.e.,implementation has not fully demonstrated compliance to the FIPS), the "Prior Validation Testing" solicitation wording option shall be included.

PRIOR VALIDATION--When an agency determines that it is essential for implementations of a FIPS to be validated (i.e., implementation has been tested and has demonstrated compliance to the FIPS) for conformance to that FIPS prior to being offered the 'Prior Validation' solicitation wording option shall be included.

In most cases, the validation option selected depends upon the availability of a test suite and a test service, as well as the availability of conforming products. Before specifying "Prior Validation" or "Prior Validation Testing," consult the Validated Products List to determine the availability of tested products and the degree of conformance. Whereas availability of test suites, testing services, and conforming products may be a factor in selecting the validation option, the most critical issue is being able to determine if the product will meet performance and capacity requirements once it is in conformance with the standard. Benchmarking results may change drastically when nonconformities are corrected. The "prior validation testing" may provide some indication of risk, but only "prior validation" can give the best assurance that the agencies' system requirements can be met. The agencies should be advised to assess the risk involved when choosing other than "prior validation."

For SQL, "Delayed Validation" or "Prior Validation Testing" are reasonable choices for 1991-1992. By 1993 the "Prior Validation Testing" and "Prior Validation" options should be use more frequently. For the IRDS, "Delayed Validation" is the only reasonable choice for 1991-1992. It is too early to determine whether "Prior Validation" or "Prior Validation Testing" should be required by 1993.

For RDA, products may become available in 1992. Since there is no existing test suite or validation service, customer assurance of conformance will be based on other methods: vendor demonstrations of interoperability with heterogeneous RDA clients and servers, vendor's test suites or documentation, customer's tests or analyses, customer success using RDA products for prototypes, etc.

C.8 Specifying Validation and Correction of Deficiencies

Modify the FIPS 127-1 and FIPS 156 requirement above to specify one of the three validation options (or some other acceptance criteria for conformance to FIPS 127-1), and include the appropriate validation specification in the RFP:

"Validation of FIPS Implementations"

"In addition to the **FIPS** implementation requirements specified elsewhere in this requirements document, all implementations of FIPS that are brought into the federal inventory as a result of this document for which validation is specified, and those implementations used by vendors to develop programs or provide services shall be validated using the official Validation System specified by the Computer Systems Laboratory (CSL). Validation shall be in accordance with CSL validation procedures for that FIPS. The results of validation shall be used to confirm that the implementation meets the requirements of the applicable FIPS as specified in this document.

To be considered responsive the offeror shall:

(1) Provide validated FIPS implementations through [select one - 'Delayed Validation', 'Prior Validation Testing' or 'Prior Validation'], as specified for each FIPS that is specified elsewhere in this requirements document."

(Option 1) - "Delayed Validation"

"The offeror shall certify in the offer that all implementations of FIPS, including FIPS options, offered in response to this document will be submitted for validation upon contract award with a request to have the validation completed at the earliest possible date permitted by the CSL validation procedures, or have been previously tested or validated and included on the current list of validated products maintained by the Computer Systems Laboratory (CSL). (The CSL list is periodically published when sufficient changes warrant.) Unless specified elsewhere, proof of submission for validation shall be in the form of a letter from CSL scheduling the validation, and the subsequent delivery by the offeror of a CSL-registered report or CSL certificate immediately upon receipt thereof. Proof of testing shall be provided in the form of a CSL registered validation summary report. Proof of validation shall be in the form of a CSL Certificate of Validation."

(Option 2) - "Prior Validation Testing"

"The offeror shall certify in the offer that all implementations of FIPS, including FIPS options, offered in response to this document have been previously tested or validated and included on the current list of validated products maintained by the Computer Systems Laboratory (CSL). Unless specified elsewhere, proof of testing shall be provided in the form of a CSL registered validation summary report. Proof of validation shall be in the form of a CSL Certificate of Validation."

(Option 3) - "Prior Validation"

"The offeror shall certify in the offer that all implementations of FIPS, including FIPS options, offered in response to this document have been previously validated and included on the current list of validated products maintained by Computer Systems Laboratory (CSL). Unless specified elsewhere, proof of validation shall be in the form of a CSL Certificate of Validation."

(2) Agree to correct all implementation nonconformance from the applicable FIPS reflected in the validation summary report not previously covered by a waiver. All areas of nonconformance must be corrected within 12 months [or specify another acceptable time frame] from the date of contract award unless otherwise specified elsewhere in this document. If an interpretation of the FIPS is required that will invoke the procedures set forth in Interpretation Procedures for Federal Information Processing Standards for Software (FIPS PUB 29-2), such a request for interpretation shall be made within 30 calendar days after contract award. Any corrections that are required as a result of decisions made under the procedures of FIPS PUBS 29-2 shall be completed within 12 months [or specify another acceptable time frame] of the date of the formal notification to the contractor of

the approval of the interpretation. Proof of correction in either case shall be in the form of a CSL Certificate of Validation or registered validation summary report for the corrected implementation. Failure to make required corrections within the time limits set forth above shall be deemed a failure to deliver required software. The liquidated damages as specified for failure to deliver the operating system or other software shall apply."

Appendix D -- Organizations and Standards Groups Referenced

ISO — International Organization for Standardization, recently entered into cooperative agreement with IEC via Joint Technical Committee 1 (JTC1) to jointly publish information technology standards. All ISO publications are available through ANSI. See OMNICOM for availability of early drafts.

IEC — International Electrotechnical Commission, recently entered into cooperative agreement with ISO via Joint Technical Committee 1 (JTC1) to jointly publish information technology standards.

ANSI — American National Standards Institute, 1430 Broadway, New York, NY 10018, publishes all American National Standards, phone: 212-642-4900, International: 212-642-4986.

X3 — An ANSI accredited committee for the development of Information Processing Standards in the United States. X3 standards are published through ANSI. See GLOBAL for availability of early drafts.

NIST — National Institute of Standards and Technology, develops Federal Information Processing Standards and guidelines (FIPS) for the United States government for sale through NTIS.

NTIS — National Technical Information Service, Springfield, VA 22161, is the sales agent for all FIPS documents, Sales phone: 703-487-4650.

GPO — Government Printing Office, Washington, DC 20402, is the sales agent for NIST reports, such as the NIST Special Publication 500 and 800.

GLOBAL — Global Engineering, Inc. A commercial firm that has a special relationship with X3 to provide copies of draft proposed ANSI/X3 standards before formal publication, phone: 800-854-7179.

OMNICOM — OMNICOM, Inc. A commercial firm that has a special relationship with ISO/IEC JTC1 to provide copies of draft International Standards before formal publication, phone: 800-666-4266.

NIST-114A
(REV. 3-90)

U.S. DEPARTMENT OF COMMERCE
NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY

BIBLIOGRAPHIC DATA SHEET

1. PUBLICATION OR REPORT NUMBER
NISTIR 4780

2. PERFORMING ORGANIZATION REPORT NUMBER

3. PUBLICATION DATE
FEBRUARY 1992

4. TITLE AND SUBTITLE

Guide for Specifying and Building CITIS with Data Management Standards

5. AUTHOR(S)

Leonard Gallagher, Joan Sullivan, Joseph Collica

6. PERFORMING ORGANIZATION (IF JOINT OR OTHER THAN NIST, SEE INSTRUCTIONS)

U.S. DEPARTMENT OF COMMERCE
NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY
GAITHERSBURG, MD 20899

7. CONTRACT/GRANT NUMBER

8. TYPE OF REPORT AND PERIOD COVERED

9. SPONSORING ORGANIZATION NAME AND COMPLETE ADDRESS (STREET, CITY, STATE, ZIP)

10. SUPPLEMENTARY NOTES

11. ABSTRACT (A 200-WORD OR LESS FACTUAL SUMMARY OF MOST SIGNIFICANT INFORMATION. IF DOCUMENT INCLUDES A SIGNIFICANT BIBLIOGRAPHY OR LITERATURE SURVEY, MENTION IT HERE.)

This paper complements the military specification for a Contractor Integrated Technical Information Service (CITIS) when that service is to be provided in an environment that supports other federal standards for data management systems. We present the current status of existing and emerging International Organization for Standardization (ISO), American National Standards Institute (ANSI), and Federal Information Processing Standard (FIPS) standards for database management and data dictionary systems, specifically Database Language SQL, Remote Database Access (RDA), and Information Resource Dictionary System (IRDS). We address the CITIS specification in terms of these data management standards and indicate how these standards may be specified or used to meet the requirements for various levels of service and functional requirements of CITIS. We conclude by identifying the benefits of data management standards in the CITIS architecture. In the Appendices, we describe the general content of each data management standard and discuss its applicability and availability. Where appropriate, we also address the availability of conformance test suites and future plans for enhancements and follow-on standardization efforts.

12. KEY WORDS (6 TO 12 ENTRIES; ALPHABETICAL ORDER; CAPITALIZE ONLY PROPER NAMES; AND SEPARATE KEY WORDS BY SEMICOLONS)

ANSI; CITIS; database; data dictionary; data model; FIPS; IRDS; ISO; MILSTD; network; RDA; relational; standard; SQL

13. AVAILABILITY

☒ UNLIMITED

□ FOR OFFICIAL DISTRIBUTION. DO NOT RELEASE TO NATIONAL TECHNICAL INFORMATION SERVICE (NTIS).

□ ORDER FROM SUPERINTENDENT OF DOCUMENTS, U.S. GOVERNMENT PRINTING OFFICE, WASHINGTON, DC 20402.

☒ ORDER FROM NATIONAL TECHNICAL INFORMATION SERVICE (NTIS), SPRINGFIELD, VA 22161.

14. NUMBER OF PRINTED PAGES
49

15. PRICE
A03