

# ISO/IEC JTC 1/SC 32 N 0335

Date: 1999-05-27

REPLACES: --

<p style="text-align: center;"><b>ISO/IEC JTC 1/SC 32</b></p> <p style="text-align: center;"><b>Data Management and Interchange</b></p> <p style="text-align: center;"><b>Secretariat: United States of America (ANSI)</b></p> <p style="text-align: center;"><b>Administered by Pacific Northwest National Laboratory on behalf of ANSI</b></p>
--

<b>DOCUMENT TYPE</b>	Proposed NP (proposal under review by WG or SC)
<b>TITLE</b>	SubprojectProposalforSQL/OLAPAmendment
<b>SOURCE</b>	ISO/IEC JTC 1/SC 32/WG 3
<b>PROJECT NUMBER</b>	32.03.04.00.00
<b>STATUS</b>	For approval by ISO/IEC JTC 1/SC 32
<b>REFERENCES</b>	
<b>ACTION ID.</b>	COM
<b>REQUESTED ACTION</b>	For approval by ISO/IEC JTC 1/SC 32
<b>DUE DATE</b>	
<b>Number of Pages</b>	6
<b>LANGUAGE USED</b>	English
<b>DISTRIBUTION</b>	P & L Members SC Chair WG Conveners and Secretaries

Douglas Mann, Secretariat, ISO/IEC JTC 1/SC 32

Pacific Northwest National Laboratory \*, 901 D Street, SW., Suite 900, Washington, DC, 20024-2115,  
United States of America

Telephone: +1 703 575 2114; Facsimile: +1 703 681 9180; E-mail: [MannD@battelle.org](mailto:MannD@battelle.org)

\*Pacific Northwest National Laboratory (PNL) administers the ISO/IEC JTC 1/SC 32 Secretariat on behalf of ANSI



**ISO**  
International Organization for Standardization



**ANSI**  
American National Standards Institute

**ANSI TC X3H2**  
**Database**  
**ISO/IEC JTC 1/SC 32**  
**Data Management and Interchange**  
**WG 3**  
**Database Languages**

**Project:** ANSI: 1234D — ISO: 1.32.3.4

**Title:** Subproject Proposal for SQL/OLAP Amendment

**Status:** Proposal for consideration by NCITS TC H2 and possible submission to ISO/IEC JTC1/SC32

**Author:** Jim Melton (USA), Fred Zemke (USA), Krishna Kulkarni (USA)

**Abstract:** ANSI and ISO work on “SQL3” recently culminated with the initiation of FDIS ballots on five parts of Database Language SQL (SQL/Framework, SQL/Foundation, SQL/CLI, SQL/PSM, and SQL/Bindings—collectively known as SQL:1999). Urgent market needs for extensions to SQL related to OLAP (On-Line Analytical Processing) would be poorly served by wholesale replacement of these parts of SQL:1999 with, say, SQL:2001. Those requirements may best be met by rapid progression of an amendment to the SQL standard that provides the language features in a form that can be readily merged into the next publication of appropriate parts of the standard.

**References:**

- 1) [FrameFDIS] ANSI X3H2-99-074 = WG3:YGJ-006, (*Final Committee Draft*) *Framework (SQL/Framework)*, March, 1999
- 2) [FoundFDIS] ANSI X3H2-99-075 = WG3:YGJ-007, (*Final Committee Draft*) *Foundation (SQL/Foundation)*, March, 1999
- 3) [CLI-FDIS] ANSI X3H2-99-\_\_\_ = WG3:YGJ-022, (*Final Committee Draft*) *Call-Level Interface (SQL/CLI)*, March, 1999

- 4) [PSM-FDIS] ANSI X3H2-99-076 = WG3:YGJ-008, (*Final Committee Draft*) *Persistent Stored Modules (SQL/PSM)*, March, 1999
- 5) [BindFDIS] ANSI X3H2-99-077 = WG3:YGJ-009, (*Final Committee Draft*) *Host Language Bindings (SQL/Bindings)*, March, 1999
- 6) [Rank] ANSI X3H2-99-009, *Rank Functions for OLAP*, Krishna Kulkarni, Hamid Pirahesh, Andy Witkowski, and Fred Zemke, 31 January, 1999
- 7) [Cube] ANSI X3H3-96-205R3 = DBL:MAD-032R1, *Super Sets (The Cube and Beyond)*, Frank Pellow, 11 October, 1996
- 8) [Plan] ISO/IEC JTC1/SC21 N8091, Elaboration of the Scope and Content of Project 1.21.3.4 (SQL3), dated 30 June, 1993, also distributed as document ISO/IEC JTC1/SC21/WG3 N1455, dated 20 November, 1992.
- 9) [OLAP-CWD] ANSI X3H2-99-146 = WG3:YGJ-061, (*Candidate Working Draft*) *Amendment for On-Line Analytical Processing Facilities (SQL/OLAP)*, March, 1999

## 1. Discussion

### 1.1. SQL Standard Status

Seven years of hard work recently culminated in the completion of Final Committee Draft (FCD) ballots and Editing Meetings for five parts of the SQL standard: Part 1 (SQL/Framework), Part 2 (SQL/Foundation), Part 3 (SQL/CLI), Part 4 (SQL/PSM), and Part 5 (SQL/Bindings). It is widely agreed that seven years was—and is—far too long for revision of the SQL standard, due in no small part to the rapidity of technological advances affecting products implementing that standard. These five parts recently were recommended for Final Draft International Standard (FDIS) ballot, which ballots are expected to succeed and result in publication of five parts of an International Standard in 1999.

It is also widely, if not universally, agreed that replacement of these parts of the standard in *too short* a period would do disservice to both product vendors and the marketplace (in part because of the confusion that results whenever a revision of a standard is published, and in part because there will have been insufficient time for the standard to take hold before it is replaced).

### 1.2. OLAP and SQL

[Cube] introduced the notion of specialized SQL language to support On-Line Analytical Processing (OLAP) capabilities. That paper proposed the specification changes necessary to provide two specialized grouping operations (CUBE and ROLLUP) that are particularly difficult for applications to accomplish without database engine support.

However, it is obvious to practitioners of OLAP that CUBE and ROLLUP are merely two of a number of language enhancements that are required to effectively support OLAP capabilities. There are a host of “statistical” operations that are required and not all of these can be effectively handled by application code without SQL engine support.

A recent paper, [Rank], in the United States outlined the changes required to begin supporting additional operations related to OLAP requirements. Although a motion to adopt the paper failed, the discussion indicated strong interest in understanding the requirements better and in pursuing the topic further—as long as the success of the 1999 revision of the standard was not impeded by developing the additional technology. Discussion also indicated significant interest in exploring various technical approaches to satisfying those requirements.

### 1.3. Breaking Standardization Ground

In order to permit enhancements to the SQL standard without undertaking the undesirable goal of revising the parts of SQL just now being processed in 1999 or further delaying progression of those parts, some creative thought was required.

A solution that presented itself almost immediately is development of an *amendment* to various parts of SQL:1999. An amendment has a format that is substantially similar to that of an incremental part as defined in the SQL:1999 Framework. That is, an amendment specifies changes to particular sections of an existing standard, including new or revised Concepts, new or revised Formats, and new or revised Syntax, Access, General, and even Conformance Rules. The principal difference between an amendment and an incremental part is that an amendment is developed with the explicit intent of merging it with the document or documents it amends before those documents are revised and republished, while an incremental part may never be so merged.

## 2. Needs

Database Language SQL (as specified in ISO/IEC FDIS 9075-1:1999, FDIS 9075-2:1999, FDIS 9075-3:1999, FDIS 9075-4:1999, and FDIS 9075-5:1999) has only limited support for On-Line Analytical Processing facilities (specifically, the CUBE and ROLLUP features defined in [Cube]). Considerable additional support is required by the marketplace. This is apparent from the significant market presence of companies specializing in providing data warehouse and other OLAP capabilities, either as extensions to existing SQL products, as layers building above existing SQL products, or as complete replacements for existing SQL products for use in specialized applications. The need for these facilities is increasingly urgent, although it has arisen recently enough that it was not feasible to specify them as part of SQL:1999.

The SQL Project, ISO project 1.32.3.4, is documented in [Plan]; although the project plan has not been revised in some time, it remains valid and active. The project plan already authorizes further development of Database Language SQL, with additional data types and associated operators clearly included amongst expected enhancements.

## 3. Existing Practice

In early 1999, OLAP capabilities are available in several forms. Various SQL database vendors have implemented a variety of OLAP-oriented features into their products; without a standard to guide them, the vendors' choices of features are incompatibly different between products with the predictable result of confused and frustrated customers who are increasingly "locked into" specific products. A number of after-market tools and layered products have been marketed in an attempt to provide more common OLAP capabilities without the restrictions imposed by the underlying SQL database systems; however, these tools and products have also suffered from the lack of a guiding standard and are at least as varied as the SQL database products' efforts.

In the vacuum left by lack of a standard response to the OLAP market needs, no vendor's solutions have emerged as a *de facto* standard, although there remains the possibility that some dominant database vendor will stumble into an opportunity to promulgate such a standard without the participation of other vendors or the customer base.

## 4. Expected Stability

In the last couple of years, the market requirements for OLAP capabilities and the technical solutions responding to those requirements have finally gelled. The requirements are much better understood today than they were only five years ago, and the database implementation facilities have improved enough to make responses to the requirements entirely feasible.

Therefore, it is expected that an amendment published in the very early years of the next decade will be widely implemented by SQL product vendors and that those implementations will continue for at least another decade (this prediction being based on the longevity of other SQL features specified and implemented over a decade ago). While it should be expected that additional OLAP-related facilities will be discovered and specified during the coming decade, they will be compatible with the facilities specified in the proposed amendment.

## 5. Program of Work

Specify an Amendment to 9075-1:1999, 9075-2:1999, and possibly 9075-3:1999, 9075-4:1999, and 9075-5:1999 to define additional operations—including concepts, syntax, and semantics—that respond to market requirements for SQL database support for On-Line Analytical Processing (OLAP) capabilities, with the intent of publishing this amendment in late 2000 or early 2001.

Such capabilities may include, but are not limited to, features such as:

- cumulative sums and averages
- ranking functions
- statistical functions, such as standard deviation and linear regression

## **6. Justification for Subproject Request**

Various criteria have been specified for the approval of program extensions (that is, subdivisions and minor enhancements) of existing projects. Those criteria relative to the subdivision of the existing SQL project [Plan] are satisfied as followed:

1. The rationale for the SQL/OLAP extension of Database Language SQL is given in the "Needs" and "Existing Practice" paragraphs above. This proposed work is within the scope of already-authorized work in the SQL project description for "additional...associated operators".
2. Consensus on the need to do this work is evident from the response to [Rank].
3. An initial candidate base document, [OLAP-CWD], has been produced and made available in preparation for the SC32 meetings in Matsue, Japan, in May, 1999. NCITS TC H2 and ISO/IEC SC32/WG3 are uniquely qualified to address and resolve issues specific to SQL implementations.
4. The SQL editor, Mr. Jim Melton, is willing to be the subproject editor for Database Language SQL — Amendment for On-Line Analytical Processing Facilities (SQL/OLAP).

**End of paper**