SQL Best Practices

Join (SQL)

functions in a SQL statement". Inside Oracle APEX. Archived from the original on 2018-12-27. Larsen, Gregory A. (2009-10-29). "T-SQL Best Practices

Don't Use - A join clause in the Structured Query Language (SQL) combines columns from one or more tables into a new table. The operation corresponds to a join operation in relational algebra. Informally, a join stitches two tables and puts on the same row records with matching fields. There are several variants of JOIN: INNER, LEFT OUTER, RIGHT OUTER, FULL OUTER, CROSS, and others.

PostgreSQL

database management system (RDBMS) emphasizing extensibility and SQL compliance. PostgreSQL features transactions with atomicity, consistency, isolation,

PostgreSQL (POHST-gres-kew-EL) also known as Postgres, is a free and open-source relational database management system (RDBMS) emphasizing extensibility and SQL compliance. PostgreSQL features transactions with atomicity, consistency, isolation, durability (ACID) properties, automatically updatable views, materialized views, triggers, foreign keys, and stored procedures.

It is supported on all major operating systems, including Windows, Linux, macOS, FreeBSD, and OpenBSD, and handles a range of workloads from single machines to data warehouses, data lakes, or web services with many concurrent users.

The PostgreSQL Global Development Group focuses only on developing a database engine and closely related components.

This core is, technically, what comprises PostgreSQL itself, but there is an extensive developer community and ecosystem that provides other important feature sets that might, traditionally, be provided by a proprietary software vendor. These include special-purpose database engine features, like those needed to support a geospatial or temporal database or features which emulate other database products.

Also available from third parties are a wide variety of user and machine interface features, such as graphical user interfaces or load balancing and high availability toolsets.

The large third-party PostgreSQL support network of people, companies, products, and projects, even though not part of The PostgreSQL Development Group, are essential to the PostgreSQL database engine's adoption and use and make up the PostgreSQL ecosystem writ large.

PostgreSQL was originally named POSTGRES, referring to its origins as a successor to the Ingres database developed at the University of California, Berkeley. In 1996, the project was renamed PostgreSQL to reflect its support for SQL. After a review in 2007, the development team decided to keep the name PostgreSQL and the alias Postgres.

Microsoft SQL Server

Microsoft SQL Server is a proprietary relational database management system developed by Microsoft using Structured Query Language (SQL, often pronounced

Microsoft SQL Server is a proprietary relational database management system developed by Microsoft using Structured Query Language (SQL, often pronounced "sequel"). As a database server, it is a software product with the primary function of storing and retrieving data as requested by other software applications—which may run either on the same computer or on another computer across a network (including the Internet). Microsoft markets at least a dozen different editions of Microsoft SQL Server, aimed at different audiences and for workloads ranging from small single-machine applications to large Internet-facing applications with many concurrent users.

List of tools for static code analysis

coding practices in Perl. Padre – An IDE for Perl that also provides static code analysis to check for common beginner errors. TOAD – A PL/SQL development

This is a list of notable tools for static program analysis (program analysis is a synonym for code analysis).

MySQL Cluster

MySQL Cluster , also known as MySQL Ndb Cluster is a technology providing shared-nothing clustering and auto-sharding for the MySQL database management

MySQL Cluster , also known as MySQL Ndb Cluster is a technology providing shared-nothing clustering and auto-sharding for the MySQL database management system. It is designed to provide high availability and high throughput with low latency, while allowing for near linear scalability. MySQL Cluster is implemented through the NDB or NDBCLUSTER storage engine for MySQL ("NDB" stands for Network Database).

Transparent data encryption

Datacenter editions of Microsoft SQL Server, until it was also made available in the Standard edition for 2019. SQL TDE is supported by hardware security

Transparent data encryption (often abbreviated to TDE) is a technology employed by Microsoft, IBM and Oracle to encrypt database files. TDE offers encryption at file level. TDE enables the encryption of data at rest, encrypting databases both on the hard drive and consequently on backup media. It does not protect data in transit nor data in use. Enterprises typically employ TDE to solve compliance issues such as PCI DSS which require the protection of data at rest.

Microsoft offers TDE as part of its Microsoft SQL Server 2008, 2008 R2, 2012, 2014, 2016, 2017 and 2019. TDE was only supported on the Evaluation, Developer, Enterprise and Datacenter editions of Microsoft SQL Server, until it was also made available in the Standard edition for 2019. SQL TDE is supported by hardware security modules from Thales e-Security, Townsend Security and SafeNet, Inc.

IBM offers TDE as part of Db2 as of version 10.5 fixpack 5. It is also supported in cloud versions of the product by default, Db2 on Cloud and Db2 Warehouse on Cloud.

Oracle requires the Oracle Advanced Security option for Oracle 10g and 11g to enable TDE. Oracle TDE addresses encryption requirements associated with public and private privacy and security mandates such as PCI and California SB 1386. Oracle Advanced Security TDE column encryption was introduced in Oracle Database 10g Release 2. Oracle Advanced Security TDE tablespace encryption and support for hardware security modules (HSMs) were introduced with Oracle Database 11gR1. Keys for TDE can be stored in an HSM to manage keys across servers, protect keys with hardware, and introduce a separation of duties.

The same key is used to encrypt columns in a table, regardless of the number of columns to be encrypted. These encryption keys are encrypted using the database server master key and are stored in a dictionary table

in the database.

Snapshot isolation

systems, such as InterBase, Firebird, Oracle, MySQL, PostgreSQL, SQL Anywhere, MongoDB and Microsoft SQL Server (2005 and later). The main reason for its

In databases, and transaction processing (transaction management), snapshot isolation is a guarantee that all reads made in a transaction will see a consistent snapshot of the database (in practice it reads the last committed values that existed at the time it started), and the transaction itself will successfully commit only if no updates it has made conflict with any concurrent updates made since that snapshot.

Snapshot isolation has been adopted by several major database management systems, such as InterBase, Firebird, Oracle, MySQL, PostgreSQL, SQL Anywhere, MongoDB and Microsoft SQL Server (2005 and later). The main reason for its adoption is that it allows better performance than serializability, yet still avoids most of the concurrency anomalies that serializability avoids (but not all). In practice snapshot isolation is implemented within multiversion concurrency control (MVCC), where generational values of each data item (versions) are maintained: MVCC is a common way to increase concurrency and performance by generating a new version of a database object each time the object is written, and allowing transactions' read operations of several last relevant versions (of each object). Snapshot isolation has been used to criticize the ANSI SQL-92 standard's definition of isolation levels, as it exhibits none of the "anomalies" that the SQL standard prohibited, yet is not serializable (the anomaly-free isolation level defined by ANSI).

In spite of its distinction from serializability, snapshot isolation is sometimes referred to as serializable by Oracle.

Database

majority use SQL for writing and querying data. In the 2000s, non-relational databases became popular, collectively referred to as NoSQL, because they

In computing, a database is an organized collection of data or a type of data store based on the use of a database management system (DBMS), the software that interacts with end users, applications, and the database itself to capture and analyze the data. The DBMS additionally encompasses the core facilities provided to administer the database. The sum total of the database, the DBMS and the associated applications can be referred to as a database system. Often the term "database" is also used loosely to refer to any of the DBMS, the database system or an application associated with the database.

Before digital storage and retrieval of data have become widespread, index cards were used for data storage in a wide range of applications and environments: in the home to record and store recipes, shopping lists, contact information and other organizational data; in business to record presentation notes, project research and notes, and contact information; in schools as flash cards or other visual aids; and in academic research to hold data such as bibliographical citations or notes in a card file. Professional book indexers used index cards in the creation of book indexes until they were replaced by indexing software in the 1980s and 1990s.

Small databases can be stored on a file system, while large databases are hosted on computer clusters or cloud storage. The design of databases spans formal techniques and practical considerations, including data modeling, efficient data representation and storage, query languages, security and privacy of sensitive data, and distributed computing issues, including supporting concurrent access and fault tolerance.

Computer scientists may classify database management systems according to the database models that they support. Relational databases became dominant in the 1980s. These model data as rows and columns in a series of tables, and the vast majority use SQL for writing and querying data. In the 2000s, non-relational databases became popular, collectively referred to as NoSQL, because they use different query languages.

Active record pattern

name column has the value "gearbox". The SQL command used might be similar to the following, depending on the SQL implementation details of the database:

In software engineering, the active record pattern is an architectural pattern. It is found in software that stores in-memory object data in relational databases. It was named by Martin Fowler in his 2003 book Patterns of Enterprise Application Architecture. The interface of an object conforming to this pattern would include functions such as Insert, Update, and Delete, plus properties that correspond more or less directly to the columns in the underlying database table.

The active record pattern is an approach to accessing data in a database. A database table or view is wrapped into a class. Thus, an object instance is tied to a single row in the table. After creation of an object, a new row is added to the table upon save. Any object loaded gets its information from the database. When an object is updated, the corresponding row in the table is also updated. The wrapper class implements accessor methods or properties for each column in the table or view.

This pattern is commonly used by object persistence tools and in object–relational mapping (ORM). Typically, foreign key relationships will be exposed as an object instance of the appropriate type via a property.

Snake case

variables and subroutines Oracle SQL and PL/SQL, for all unquoted identifiers (tables, columns, indexes, constraints, PL/SQL variables, constants, procedures/functions

Snake case (sometimes stylized autologically as snake_case) is the naming convention in which each space is replaced with an underscore (_) character, and words are written in sentence case. It is a commonly used naming convention in computing, for example for variable and subroutine names, and for filenames. One study has found that readers can recognize snake case values more quickly than camel case. However, "subjects were trained mainly in the underscore style", so the possibility of bias cannot be eliminated.

A variation is screaming snake case, where words are written in all caps (stylized as SCREAMING_SNAKE_CASE). This convention is used for constants in programming languages like C/C++, Python, Java, PHP, as well as for environment variables.