THE POSSIBILITIES ARE INFINITE FUJITSU

WHITE PAPER

BS2000/OSD
UDS/SQL
Universal Database System

Issue June 2009

Pages 8

Over the last decades, database systems have evolved into an indispensable tool within operative organization. Initially, the crucial factor determining their use was the support of isolated work and business processes, while in recent years the integration of different information systems has become ever more important.

With this shift of emphasis, the flow of information within enterprises has become a competitive feature of the highest order. Operative work processes must be supported by high levels of performance, availability and reliability in data management. The success of a company is heavily dependent on how fast its changing data can be made available and on the extent to which it can be converted into high-quality information.

The data is held on a variety of platforms (PC server, departmental computer, central high-performance server), depending on the main purposes to which it is to be put. However, for integrated processing it must be possible to transfer it for productive use to other software and hardware environments.

When integrating applications of different systems and running on different platforms, it is essential to secure returns on the high investment costs in terms of the development of the individual systems, the know-how of the employees running and working with these systems, and the accompanying process organization.

As far as the individual systems are concerned, continuity is therefore at a premium.

As a means of handling business processes, internet technology has set the pace both across company borders and within single enterprises (intranet) as network technology. The utilization of data via client/server architectures is implemented on the basis of this network technology. Access to the server data through the client uses standard tools independent of location, hardware, software and procedures.

In this environment UDS/SQL on BS2000/OSD takes its place as a database server suitable for very large sets of data with extremely high demands in terms of performance and availability.

Contents

Data strategy with database philosophy

The way ahead in the IT world is clearly marked out: traditional systems for holding and managing data are giving way to databases integrated in client/server architectures.

This has undoubted benefits for the user:

- Applications of all kinds can be integrated at the user's workplace.
- All users access the same uniform data pool, which is therefore always up to date, for their own specific problems.
- The powerful desktop systems at the workplace guarantee a high level of user convenience through windows technology, mouse and graphics.
- Two or more users can make changes simultaneously and in parallel.
- Online access to corporate applications and databases can be handled conveniently with state-of- the-art technology for both local (LAN) and wide area networks (WAN).
- It is possible to administer a large number of users with different access profiles.
- Transaction processing offers effective protection against the loss of data, a problem with which individual users no longer have to concern themselves.
- Users have recourse to different access paths to the data, each tailored to the particular problem in hand.

Databases provide the foundation for secure, consistent and integrated information management within an enterprise.
The reasons are plain:

- No more data redundancy
  Integrated data management with database systems avoids data redundancy and the uncontrolled splitting of data sets.
- Parallel access
  Parallel access through a large number of users is made possible.
- Effective access control
  The data can be made available to different users, depending on particular requirements. Unauthorized access can be ruled out.
- Integration of stored data
  A high level of data integration is possible as compared with conventional storage methods.
- Consistency of stored data
  Database software guarantees the consistency of data even in the event of a crash, system errors, destruction of a volume or transaction abortion.

All this makes database software the cornerstone of modern commercial information systems, opening the way to the economic utilization and most effective combination of all the resources involved. The development of applications can be accompanied by high productivity, from data modeling to the technical specifications, right up to coding and testing.

The use of databases requires powerful and cost-effective information management:

- High performance level
  The performance level of the database system is a crucial factor, with regard both to the support of operative work processes and to the provision of results for dynamic queries from the enterprise-wide information system.
- High availability
  The enterprise-wide information system functions only if the database server can guarantee a high level of availability and failsafe operation.
- Autonomic
  The reliable automation of database operation is a key factor for cost-effective data management.
- Openness
  In networked, heterogeneous information systems, the data should be processed by the instance that is functionally best equipped to do so, regardless of how it is managed. Different individual procedures should be integrated at the workplace in a user-friendly manner.

Within this environment, UDS/SQL is a sophisticated, universally deployable high-performance database system for the cost-effective implementation of a wide range of requirements, especially in conjunction with large volumes of data where continuity in operative processes is of paramount importance.

Coexistence of network and relational data structures

The databases available on the world markets can be classified according to two basic criteria:

- Hierarchic and network data structures
  These data structures are particularly suited for the efficient processing of complex data relations. By providing the structure information for user data, it is possible to implement highly effective access in precisely defined ways, especially for network database systems.
- Relational data structures
  Relational data structures are more appropriate for dealing with spontaneous, flexible database queries. Data relationships are established dynamically during runtime. Users can access information on a "set-oriented" basis.

Hierarchic systems and network database systems are the main carriers of very large OLTP applications. The CODASYL data model was standardized around 1970 for network database systems.

UDS/SQL was developed on the basis of the CODASYL data model. Right from the start, a set-oriented access procedure designed to enable flexible queries on stored data was supported in addition to the standardized CODASYL interface.

The relational model, standardized through ISO and ANSI, has greatly increased in importance and has also been implemented in UDS/SQL. This makes UDS/SQL the first database system to incorporate both data models at once.

UDS/SQL uses a data description to automatically define the internal physical storage. This basic definition can be optimized with the aid of the Storage Structure Language (SSL), e.g. by splitting up the data areas according to

- access frequency or
- how it is grouped (cluster formation).

This optimizes performance and increases throughput, especially for extremely time-critical dialog applications.

Changing the physical storage structure has no effect on application programs.

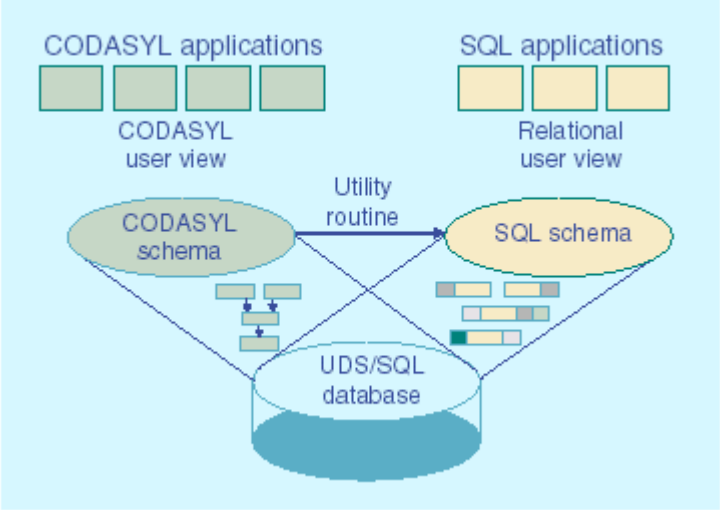

The use of COBOL-DML (Data Manipulation Language) enables highly effective data access, as it allows the best access path for the problem in hand to be defined as early as the program compilation stage. At runtime, efficient processing is supported by the structure information stored in the database.

The applications can also be programmed in languages other than COBOL, since a CALL-DML oriented to COBOL-DML enables data to be processed independent of the programming language.

With the aid of the program interface SQL, an ISO based database language used throughout the world, UDS/SQL allows users to program their applications on a purely relational basis. In this case, the CODASYL data is processed or queried by a "relational view" using SQL statements. The possibility of creating networked and relational database structures enables UDS/SQL to cover a wide variety of possible applications.

Optimal integration in BS2000/OSD

BS2000/OSD

UDS/SQL was developed with a particular eye to the resources of BS2000/OSD systems and constitutes the high-end OLTP system in the BS2000/OSD product range. Used in conjunction with other products, such as the openUTM transaction monitor, and with distributed databases, UDS/SQL combines convenience and speed with cross-platform transaction management.

UDS/SQL uses the standard BS2000/OSD procedures for installation, operation and long-term data backup.

UDS/SQL users therefore have access to the latest development status of all functions of the corresponding products, which in turn makes it possible to organize UDS/SQL on a cost-effective basis as it allows various tasks to be handled in procedures not tied to specific products.

Hardware

UDS/SQL is available for all BS2000/OSD business servers.

State-of-the-art disk peripherals can be used for UDS/ SQL applications. An extremely fast, non-volatile expanded storage (global storage) can be used for the backup file (RLOG file), which is especially time-critical.

UDS/SQL files can be maintained in mirrored form, which, especially in the case of remote mirroring, guarantees maximum compliance with the prevailing security and availability requirements, right up to catastrophe level. Mirroring can be based on hardware functions or on DRV (Dual Recording by Volume). If the data on mirrored disks has to be equalized when working with DRV, this can be done very quickly with the aid of UDS. However, with UDS/SQL it is also a simple matter to keep duplicates of logging files for fail proof operation, which again significantly contributes to boosting security.

The TimeFinder functionality in Symmetrix systems can be used, for example, to provide real data at great speed for tests to be carried out in parallel to productive operation.

openUTM

Nowadays, networks incorporating tens of thousands of users working at the same time are nothing unusual.

For these environments, too, the transaction monitor openUTM and the database system UDS/SQL offer perfect conditions for creating OLTP applications. The combination of coordinated and fine-tuned resource management, use of multiprocessing/ multithreading technologies, and distribution of the workload over several parallel processes guarantees very high throughput rates and short response times, as well as efficient processing and outstanding speed. This makes openUTM and UDS/SQL ideally suited for application scenarios that have proved beyond the capabilities of other OLTP systems.

Thanks to their coordinated transaction management, openUTM and UDS/SQL can guarantee the reliable processing of all business transactions, even if they involve complex processes with a number of dialog steps. openUTM handles all communication tasks with the users, while at the same time managing an unlimited number of applications. In open, heterogeneous IT infrastructures it is used as the basis for innovative client/server and web solutions. The tried and tested functions for mission-critical UDS/SQL applications are thus also available for the heterogeneous client/server world (BS2000/OSD, UNIX and Windows systems) as an integrated component of openSEAS.

Moreover, access to other database systems (e.g. on IBM computers in conjunction with the LU6.2 protocol) can of course also be implemented within a transaction.

The management of licences and access rights for UDS/SQL and openUTM is based on a common user model. This makes it possible to avoid security gaps in the organization of applications. Access authorization can be granted on a very specific level, being restricted first to selected users or clients and then either to a complete application, to particular sets of data, or even to individual operations within an application.

The transactions are completely synchronized, thus guaranteeing the permanent consistency of data and processing logic and taking due account of the client level. Any problems in the server application are dealt with transparently for the user and processing is restarted at precisely the spot in the application where it was interrupted and with an unchanged data status.

As a result, there is no need after a crash and subsequent restart to be plagued with the question: Was my application processed properly or has it been lost?

openNet Server

BS2000/OSD fits in perfectly to heterogeneous network environments. Via the openNet Server the data managed by UDS/SQL can be used in the open network with different procedures. The Communication Manager openNet Server enables communication via TCP, ISO and NEA transport protocols in LAN and WAN networks, and this in turn means that distributed operation via UDS-D, regardless of the concrete network technology involved, is also possible.

DRIVE/WINDOWS

The product DRIVE/WINDOWS was designed for the development of commercial OLTP (OnLine Transaction Processing) applications, which can also be distributed to different client/server architectures.

For the creation of such applications, DRIVE/ WINDOWS offers high-quality language resources and services which dramatically reduce the work required of the developer in the use of system-support interfaces to the openUTM transaction monitor, to the UDS/ SQL database system or to the user interface systems.

This leads to considerable time savings and productivity gains in the programming of UDS/SQL applications.

The language facilities supported by DRIVE/WINDOWS include the integration in the language of SQL statements for the definition of databases (DDL) and data access (DML), powerful statements for graphical screen formats (MS-Windows), comprehensive statements for generating reports and forms, resources and automated facilities for transaction control, and distribution options in client/server configurations.

NetExpress

NetExpress BS2000/Micro Focus can be used for the development and maintenance of UDS/SQL applications written in COBOL on Windows Systems.

Preliminary work, including testing for productive use on BS2000/OSD systems, is carried out on the Windows system.

Both COBOL-DML and CALL-DML calls are supported as interfaces for database access. For testing on the workstation, the database structures (schema and sub schema DDL) and test data can be imported from the BS2000/OSD system. Both batch and openUTM applications can be managed in this way, with openUTM and UDS/SQL transactions being coordinated in the same way as on BS2000/OSD. In addition to the testing tools generally available in NetExpress, such as the Animator, the graphical UDS-TRACE is available for monitoring each individual database access, the interactive modification of the parameters, and the querying of information from the currency table.

Added power through modern technology

The particular strength of UDS/SQL lies in is outstanding performance characteristics and versatile tuning options. In addition, its multitasking and multithreading architecture, in conjunction with its sophisticated cache technology, enable extremely short response times, even with unusually high throughput requirements.

Multithreading allows the server processes, where the actual database access takes place, to deal with incoming jobs efficiently and with low resource consumption. UDS/SQL internally bundles jobs so as to optimize processing, particularly in connection with heavy workloads. For individual users this means that jobs competing for processing resources do not have any noticeable effect, even with a high level of utilization.

Use of cache technology means that superfluous I/O operations when accessing the databases can be avoided. In view of the impact it can have on system performance, the data required for failover and long-term archiving is managed particularly efficiently.

For the most part, the optimization is performed internally and automatically. The user only has to enter tuning parameters in cases where the application requirements are especially stringent. For example, it may be advisable to create separate buffers for particularly performance-critical databases in order to prevent the displacement of data by other databases working in parallel.

The multi-DB strategy allows an application program to access several databases at the same time without jeopardising transaction management. In this case, the data can be organized and used according to logical criteria, and it is possible to implement a transparent workload distribution for distributed processing.

Set-oriented queries are optimized on the basis of the current data volume, with particular emphasis on the optimization of the access paths contained in the concrete data structure. As early as at the data storage stage it is possible to check and thus guarantee the referential integrity within the data structures.

With over 32,000 record types per database and more than 2 billion records per type, there are virtually no limits to the size of the data sets that can be managed.

With the aid of a monitor, it is possible to keep track of the individual factors of particular importance for the performance of the database system, thus allowing potential bottlenecks to be detected in good time.

Setting up the database, data management and data backup are supported and simplified by powerful utility routines geared particularly to the processing of extremely large volumes of data.

UDS/SQL is therefore especially recommended

- in conjunction with extreme performance requirements
- when complex and extremely large volumes of data have to be managed
- when referential integrity in the data structures is essential.

Distribution of data and procedures

In order to distribute applications and data over the network, distributed systems require powerful but easy-to-use client/server system functions. The BS2000/OSD openNet server incorporates a product strategy designed to enable the cooperation of distributed applications and data in the network, as well as the optimization of the client/server solution according to a variety of criteria.

The distribution component UDS-D can be used with the openNet server to ensure the homogeneous distribution of the databases, which can be set up at various locations, e.g. at the place where they will be used most. The actual manner of their distribution is transparent for the application programs, while access to the relevant database is controlled by a dynamically updateable distribution table. From the user's point of view, however, the distributed setup appears as a system. Data consistency is guaranteed throughout the network by means of the "two-phase commit" protocol. With distributed data sets, complete local site autonomy is ensured at all times.

In the event of faults, every sub transaction can be used to actively restore the consistency of the global transaction.

UDS/SQL can also be used for the simple implementation of branch concepts. In this case, data is copied and made available for processing as shadow databases at different locations. The copies can be produced in the original database during operation and kept consistent with the aid of standard tools used for long-term backup.

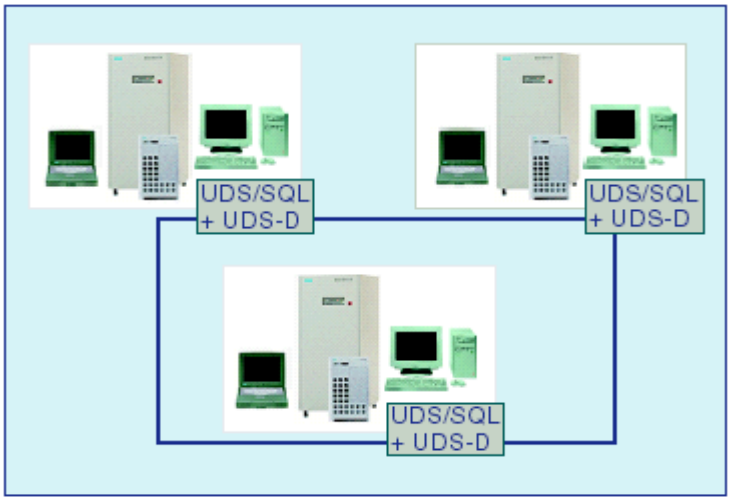

With the aid of openUTM it is possible to create distributed applications that can use different database systems in a heterogeneous computer landscape.

openUTM guarantees distribution with global transaction management with a high degree of reliability and speed. In an environment like this, UDS/SQL can demonstrate its outstanding prowess as a data server to the full.

High availability

The information stored in the databases is of extremely high value. As a result, equal if not greater importance is attached to the long-term backup of this information, which is frequently subject to changes and extensions.

UDS/SQL incorporates a wide range of individual functions designed to meet these requirements:

- Data can be backed up using powerful standard BS2000/OSD tools.
- Data can be mirrored using hardware functions or DRV (Dual Recording by Volume). Should equalization be necessary in connection with DRV, this can be handled promptly by UDS/SQL.
- Using UDS/SQL, it is simple to keep duplicates of logging files for use in the event of a crash .
- Transaction management ensures that the data of all completed transactions is consistent.
- With distributed processing using UDS-D or openUTM, the "two-phase commit" protocol guarantees global data consistency.
- A restart that has been coordinated with openUTM guarantees the consistency of data and procedures, even when faults occur.
- The use of online copies and precautionary reading in of the data from long-term backups ensures that consistent data can be made available again very quickly, even when volumes have been destroyed.
- Powerful access control tools and a user model coordinated with openUTM can be used to effectively prevent the unauthorized modification of data.

Integration of procedures

The middleware suite openSEAS

With openSEAS and Oracle Application Server Fujitsu Technology Solutions offers a comprehensive product suite for application innovation: WebTransactions makes it possible to integrate UDS/SQL applications and data into the internet or an intranet/extranet without having to change the applications. The grid-enabled J2EE application server Oracle Application Server 10g offers a complete set of middleware services for the use and management of applications and web services, for the provision of personalized applications in portals and on portable devices as well as for controlling and automation of business processes (Enterprise Application Integration). The standard compliant JCA adapters of BeanConnect provide connectivity towards UDS/SQL applications based on the TP-Monitor openUTM.

WebTransactions

Nowadays, the integration of business sequences both within the enterprise and across company boundaries is viable only in conjunction with internet technology.

This is precisely what is offered by the openSEAS product WebTransactions, which starts with the existing system landscape and makes the user interfaces of the applications - which represent a considerable investment - available in the internet/intranet without having to change them.

This means that UDS/SQL applications can be added to the web with both openUTM and BS2000/OSD dialogs, and then serviced by a web browser. Without modifying the host application, e.g. Netscape Navigator or Microsoft Internet Explorer.

The first, extremely simple step can be a straightforward 1:1 conversion, in which the screen masks are integrated unchanged in the internet/intranet/extranet (see the following example with the original openUTM mask in an MT9750 emulation and the same contents presented via WebTransactions in Netscape Navigator).

However, it is still possible to further improve the screen masks and make them more attractive, or to modify the host application by combining or extending dialog steps to bring it more into line with actual requirements (dialog reengineering). It is even possible to combine several different applications under the same web interface and run them in parallel.

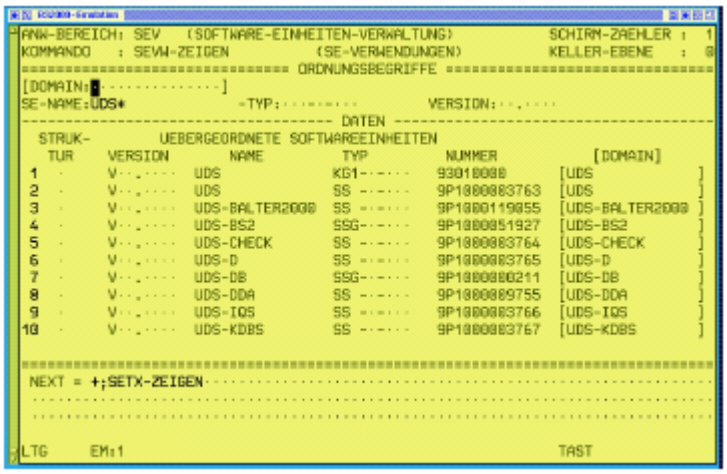

WebTransactions transforms the original, usually character-oriented interface of an application into a form that can be processed by a web browser. This conversion is based on so-called templates, which are generated automatically by WebTransactions for various mask systems. They serve as a starting point for global editing (e.g. setting the background color or inserting a company logo) and the enhancement of individual masks.

Software products such as web servers, firewall systems and encryption programs are available in connection with WebTransactions for accessing the internet or for setting up and running an intranet/extranet.

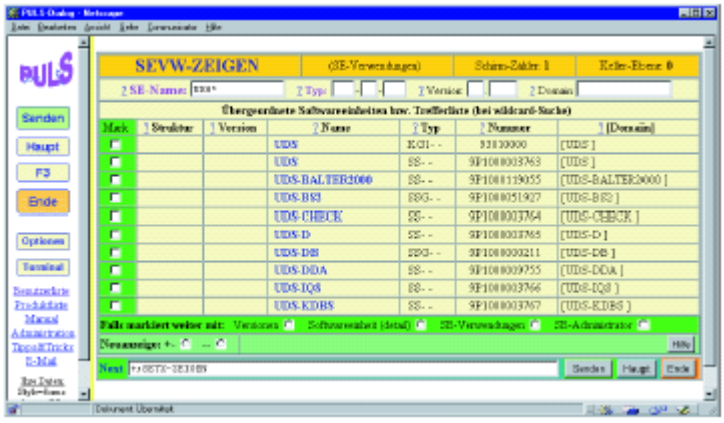

A user-friendly script language is available for designing the web interface. It comprises all the language resources of HTML as well as additional statements and control structures in Javascript syntax which are dynamically evaluated by WebTransactions during runtime. The resources of the template language make it possible to modify the appearance of a host application for presentation in the web.

The WebTransactions template facilities also enable the operation of one and the same application with several web interfaces with different styles. This option is used typically in applications offered in various languages for multinational companies or applications which the user wishes to access in varying degrees of detail. In addition, the WebTransactions template language can be used to combine or split up different dialog steps and service these steps from the templates.

Data selection and extraction

ODBC access to UDS/SQL databases

With the aid of ODBC tools, such as ODBC Rocket / gfs (Gesellschaft für Informationssysteme mbH), applications running under MS-Windows and UNIX based systems can also access data from UDS/SQL databases. For example, data from UDS/SQL can also be processed in MS-Excel spreadsheets under MS-Windows.

Multiple connections to different database systems can be set up at the same time, enabling the user to access the relevant database by means of the graphical user interface of the application program and using uniform function calls.

Data warehouses

Data from operative systems can be used in data warehouses without impairing its integrity or availability.

The data is selected, linked up, compressed, edited and organized such that it can be evaluated, correlated and represented. Data in the operative systems represents snapshots of the business status.

In other words, it reflects a given status but does not say anything about updates or trends in development.

Nevertheless, decisions are taken on the basis of this information. Procedures in the data warehouse ensure, for example, that the data is collected, stored and constantly updated on the same time axis, so that it forms a sound basis for such decisions.

Extraction of the operative data and the setting up of the actual database for the data warehouse are generally tool-supported. These ETL (Extraction, Transformation, Loading) tools access the operative data in UDS/SQL databases directly, e.g. via an ODBC connection (ODBC-Rocket) and make it possible to compress and combine this data and data from other data management systems. The tools also generate informative reports recording what they have done. In addition to the standard reports that dominated up to now, custom reports and user-friendly ad hoc reports are now also possible. However, since this generally involves the processing of enormous volumes of data, it places a considerable burden on the data management systems of the operative data. Thanks to its outstanding performance characteristics, UDS/SQL is more than up to the challenge.

UDS/SQL: a sound investment for the future

Database applications in central data centers have to service a large number of users and support high transaction rates. For these reasons the database systems used rely on optimal system integration to make use of the properties of the operating system and the software environment.

BS2000/OSD is a sound, secure and economical basis for applications of this sort. Its 370/390 architecture, multiprocessor capability, storage hierarchy management and security mirror disks equip it perfectly for the toughest challenges OLTP applications have to offer.

UDS/SQL is optimized in terms of throughput and top-of-the-class performance and works closely together with other BS2000/OSD components to get the most out of both sides.

If they are to meet the requirements of users aiming at long-term investments, database systems have to have some good growth arguments.

UDS/SQL has got them all:

- Long-term interface guarantee
- High software quality plus comprehensive service covering consulting, training, diagnostics and maintenance
- Processing of large volumes of data
- Excellent performance capabilities:
  high throughput for a large number of users, short response times, use of state-of-the-art hardware architectures
- Data security (even when hardware and software components fail) and network-wide "logical data consistency"
- Round-the-clock operation seven days a week: with fast restart capability following faults and the ability to switch operation to standby components
- Versatile facilities for protecting data against unauthorized access
- Database distribution over different systems with transparent access to local and remote data sets
- Incorporation of PCs with access to data on enterprise databases
- Use of database data through separate and autonomous procedures – for queries, analysis, further processing.

Published by department:

Frank Schützhold
Phone: ++49 89 3222 2627
frank.schuetzhold@ts.fujitsu.com
ts.fujitsu.com

Partner login
partners.ts.fujitsu.com