

# SAP R/3 on OS/390

## The **Red Hot** Application Meets **Cool Blue** Iron

BY TIM BOHLEN

SAP R/3's ability to support a number of different databases and operating systems and to allow a customer to be flexible in their choice of platform make it a viable choice for the OS/390 environment.



**FOR** a little over a year now, SAP R/3 (or simply R/3), now one of the most popular application suites and part of the new ERP (Enterprise Resource Planning) class of software, has been able to run using DB2/390 as the database. Strictly speaking, because of the SAP R/3 implementation of a “three-tiered” architecture (see Figure 1), R/3 from a functional standpoint “runs” on either UNIX or Windows NT servers. In fact, a DB2/390-based SAP R/3 system could almost be described as client/server/server, rather than the conventional client/server system. To fully describe SAP R/3 would take an article in its own right, but to summarize, it consists of comprehensive business functionality written in the proprietary language ABAP/4, a development environment and a compiler for this language. A comprehensive development environment, including code versioning and change management, is part of the system. A number of system management tools are incorporated, including database monitoring, job scheduling, output management, a graphical alert monitor, comprehensive logging, dump handling, buffer management, a comprehensive user security system, and built-in upgrade capability.

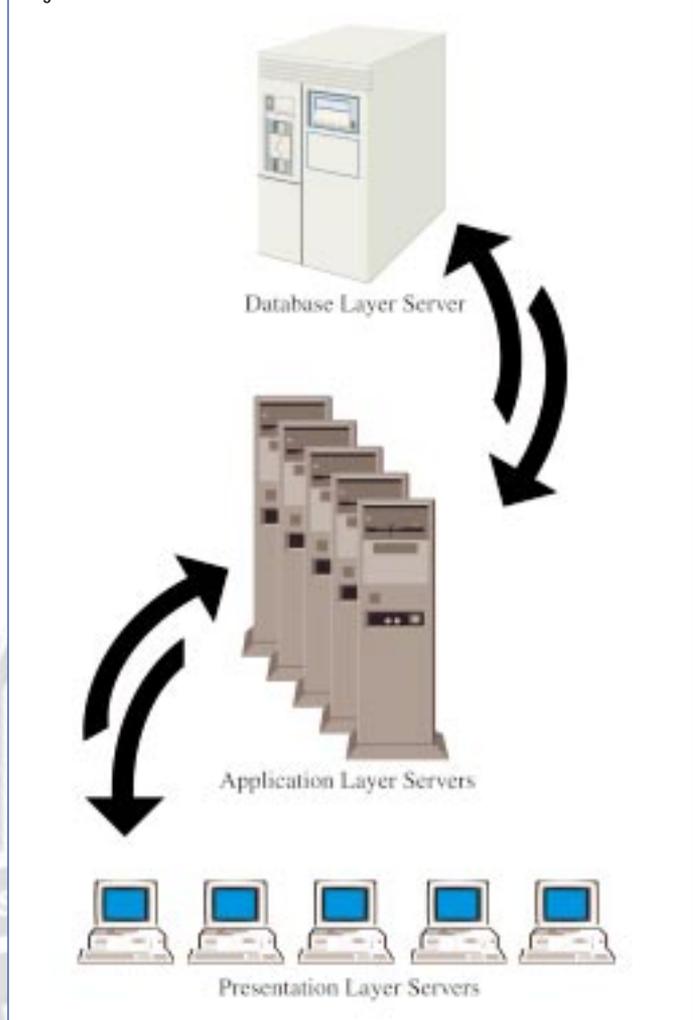
R/3 is essentially an IP- (Internet Protocol) based application, using a combination of TCP and UDP, along with some CPI-C over TCP/IP, as its primary protocols. To anyone with knowledge of the history of TCP/IP for MVS, this poses the question: How do you achieve the sort of throughput needed with TCP/IP for MVS to run an application platform that supports thousands of concurrent users? The Answer: Rewrite TCPIP/MVS. Well sort of. As we shall see, IBM has written an IP-enabling layer that achieves the throughput needed to support the volume of database traffic required for a large R/3 system. The solution also has a degree of scalability to overcome bandwidth limitations of a single ESCON or FDDI path.

#### WHY CHOOSE R/3 ON OS/390?

The reasons for implementing SAP R/3 on OS/390 are numerous, but mainly revolve around a few key areas. The purpose of this article is to detail the technical aspects of the architecture, not to promote this platform over other choices. R/3 has been successfully implemented by SAP customers on all of the available operating system and DBMS combinations supported. Some aspects of this particular implementation that are considerations in making a platform choice are as follows:

**Scalability:** Because of the architecture of R/3, the size of a database subsystem, and more particularly the ability of the OS and DBMS to support this database, can be a limiting factor in the size of a single R/3 “instance.” An instance generally will function as a Development, Testing (Quality Assurance) or Production system. Most commonly this limitation concerns the Production instance. The “size” of a database has two distinct aspects: OLTP/Batch load and the physical size of the database. The first aspect concerns the CPU, memory and I/O requirements to support the R/3 application load running against the database. With the use of DB2 Data Sharing, using Sysplex features of S/390, up to 32 systems can run together to logically support one database. This can also be achieved with Oracle Parallel Server, but the comparison between this and DB2 Data Sharing is obviously way beyond the scope of this article. A physically large database also has special considerations for providing I/O bandwidth from the application to

Figure 1: SAP R/3 Architecture



numerous physical disk devices and for keeping the data organized for optimal performance.

**Availability:** Two features of this platform are specifically exploited by R/3: sysplex failover and DB2 online reorganizations. As well as providing the ability to grow a database beyond the footprint of one database server through DB2 data sharing, support for sysplex failover also allows R/3 application servers to tolerate a failure of one member of a data sharing group by supporting timely failover to another member in the group. Again, Oracle Parallel Server also supports this functionality, and again, I will not tread onto the dangerous ground of a comparison. The second feature of DB2 that is of interest to customers with large databases and/or high uptime requirements is online reorganizations. This feature allows tablespaces to be reorganized with minimal (though non-zero) “downtime,” as perceived by the application and hence, user.

#### R/3 ARCHITECTURE

To support a number of different databases and operating systems, and to allow a customer to be flexible in their choice of platform, R/3 has been implemented with a three-tier architecture (see Figure 1). The support of a variety of front ends, in SAP parlance the *presentation tier*, is testament to the distinct breakup of the application and presentation tiers. The application layer is unaware of what platform a

particular user is using, and the data flow is identical for any front end. The data flow between the presentation server and application server is in a SAP-specific format and is transported over TCP/IP. The ability of SAP R/3 to be configured with a wide variety of DBMS systems, now including DB2/390, is a product of an equally distinct breakup of the application tier from the database tier. This is achieved by using what SAP calls "Open SQL" throughout the supplied application code (written in ABAP/4). The Open SQL, in combination with a database interface layer (DBIF) used by UNIX or NT SAP processes to communicate to the database, allow the same ABAP/4 application code to run on R/3 systems with any supported DBMS.

This means that the Open SQL cannot exploit any database-dependent features, instead relying on the database-dependent

**R/3 is essentially an IP- (Internet Protocol) based application, using a combination of TCP and UDP, along with some CPI-C over TCP/IP, as its primary protocols.**

DBIF layer to make use of any features of the underlying DBMS. Within ABAP/4 code, it is possible to code "native" SQL calls, which are passed to the DBMS without the conversion from Open SQL. No SAP supplied code performs such native SQL calls (apart from additions to provide technical support for a DBMS, which will be automatically changed upon a platform movement).

Customers are discouraged from coding such calls except in exceptional circumstances, as these calls may result in unpredictable results if a platform/ DBMS change is undertaken in the future.

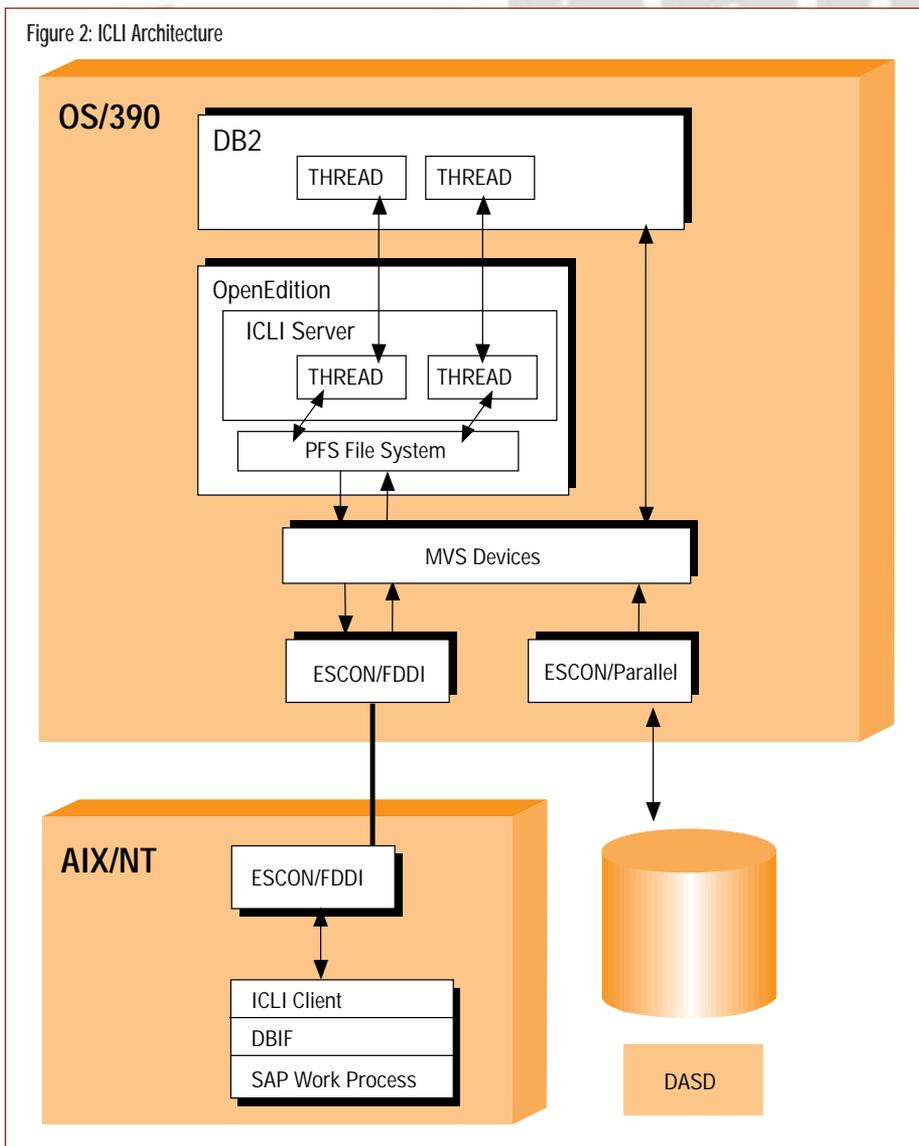
The application layer of R/3 is scalable, which means multiple servers can be implemented to support the application load. These servers can be either UNIX- or NT-based, with servers of both operating systems able to run concurrently. All servers must have IP connectivity to the single database. A unique enqueue (ENQ) process on one server, the Central Instance, performs coordination of the database activity among the servers. The ENQ process performs this activity by coordinating lock entities defined within the SAP application, in memory within the UNIX/NT process. As such, the ENQ process performs no database calls, and with all potential contention resolved by the application, allows DB2 to be configured with row level locking on all application tablespaces and lock escalation disabled. Processing is distributed among the multiple application servers to support the various work types, dialog (online), batch, printing and R/3's asynchronous update processing.

The contents of an SAP R/3 DB2 database are technically accessible using a variety of DB2-based tools and utilities. However, because SAP uses its own internal data dictionary to define its structure and views of data and the contents of the dictionary are potentially dynamic, direct data access should never be performed. Access to data must only be performed through the UNIX/NT application servers, which implement all of the rules of the data dictionary, SAP's ENQ or locking mechanism, and SAP's integrity rules. Some users may bemoan the loss of tools, but the risks and dynamic nature of the data structure make the enforcement of this rule necessary. Some read-only access is possible, but there are a lot of complications to overcome, depending on the particular data, such as the presence of multiple logical tables within single physical tables, called pool and cluster tables.

#### ICLI: THE UNIX-OS/390 GLUE

ICLI, or the Integrated Call Level Interface, consists of two pieces, an ICLI server and an ICLI client (see Figure 2), and is delivered in OS/390 as part of the Application Enabling (AE) component. The ICLI server runs in the OpenEdition environment,

Figure 2: ICLI Architecture



requiring OpenEdition customization and setup to run. There are two possible transport mechanisms for ICLI, both of which use an OpenEdition PFS (Physical File System) to implement a UDP-based IP stack in OpenEdition. One of these is HPDT, or High Performance Data Transfer, which exploits either the OSA-2 card available in CMOS-based mainframes and FDDI links to a LAN or can be routed across ESCON subchannels. HPDT is required for Windows NT application servers and can optionally be used for AIX application servers. The second protocol is called Enhanced ESCON (EESCON), UDP+ or UDP for short. This supports communication over ESCON and involves installing a driver in AIX to provide emulation of a 3088 (CTC).

After the definition of the PFS appropriate to the chosen protocol and an IPL to activate the OpenEdition changes, one of two OpenEdition commands is issued to configure the UDP stack. For EESCON, *configstk* is run using a parameter file called *configstk.conf*, and for HPDT *oeifconfig* is run, with configuration details passed directly as parameters. The purpose of both of these commands is to configure and initialize a separate IP stack in the OpenEdition environment for use by the ICLI server. The parameters used define the relationship between the IP addresses assigned to the application servers for UDP (a different address to the existing TCP/IP address) and the MVS devices defined for either ESCON or OSA-2. Both EESCON and HPDT stacks can run simultaneously, but an ICLI server instance can only use one of the stacks at a time. Multiple ICLI servers can be run allowing a number of redundancy options to recover from failures.

On the AIX application server, SMIT, a menu-driven interface to AIX system commands, is used to define the EESCON and/or HPDT devices. Prior to the use of the EESCON, subchannel executable code for a Network Interface Driver (NID) that supports this protocol must be transferred from the OS/390 Application Enabling library. After configuration and a system restart, the UDP connection is verified by a ping. Following the establishment of communications, the ICLI server(s) can be started on OS/390 as an OpenEdition task, started as a JOB, a started task or started by an OpenEdition script. An ICLI server then "listens" on a specified port number for connection

requests from ICLI clients associated with processes that are started on AIX/NT.

### THE R/3 DATABASE

The installation of R/3 involves the creation of a DB2 subsystem with prescribed DSNZPARM settings, and the loading of this subsystem with data through an AIX (or NT) process, R3INST for R/3 Version 3.x, or R3SETUP for R/3 Version 4.x. The OS/390-specific install process downloads the ICLI client, prompts for details of the host connection (host name for UDP and the ICLI port number), creates a number of OS/390 jobs to set up DB2 plans and security, and uses ftp to submit these (using an authorized OS/390 signon). The SAP install process creates the DDL to create stogroups, databases and tablespaces and then creates and loads the tables and indexes. For an installation of R/3 version 3.1H on DB2/390, 346 tablespaces, more than 9,000 tables and more than 10,000 indexes are created.

---

**One customer that recently installed SAP R/3 will implement a production system with a database totaling around 4TB, while a number of other customers are implementing systems capable of supporting more than 5,000 users.**

---

In active use, the ICLI server internally creates a thread for each work process on AIX or NT and uses the Call Attach Facility (CAF) to connect to and create threads within the requested DB2 subsystem. All data in the database is stored in ASCII, but the DB2 optimizer expects EBCDIC SQL statements. The ICLI client converts all of an ASCII SQL statement (but not the data) to EBCDIC, and subsequently, the ICLI server converts any LITERALS within the SQL back to ACSII, to perform correct query comparisons. The bulk of the traffic between ICLI client and server is ASCII data, which is not translated. While somewhat complex, this process ensures that a minimal amount of conversion is performed overall.

### DB2 FEATURES EXPLOITED

Various features of DB2 are exploited by R/3. Some of these features are new in V5.1, while others have been available since DB2 V4.

All SQL requests that SAP issues are either in the form of Dynamic SQL or an ASCII string of SQL syntax passed to the ICLI client. This means the plans bound for ICLI's CAF access are effectively empty but are named in such a way to allow a DB2 thread to be easily attributed to one of the many processes running on one of potentially many AIX/NT servers. The exclusive use of dynamic SQL ensures that the Dynamic SQL Caching (DSC) feature of DB2 V5 is an essential part of performance in a DB2/390 R/3 system. A recent maintenance level of DB2 V5 is required to ensure all known DSC problems are resolved. SAP maintains an accurate list of all known PTFs that have impacts on the R/3 environment and as such, are required in an SAP DB2 subsystem. A recent APAR (PQ07701) and associated PTF introduce a number of new IFCIDs to allow for analysis of DSC performance. The effect of retaining the prepared statements has an impact on the usage of the EDM pool of DB2. The presence of a potentially large number of long running threads in DB2 requires quite recent maintenance to ensure DB2 memory management does not cause the DBM address space to abend due to storage growth.

As mentioned previously, one feature of DB2 V5 that is of particular interest to R/3 implementations is online reorganization. An SAP database has very complex relationships of data within tables and this integrity is handled by the application logic. As a result, when one tablespace is stopped for reorg or recovery, the whole SAP system is effectively considered to be down. As a result, performing reorganizations to improve performance and/or to prevent unplanned outages due to tablespaces reaching extent limits can substantially affect the uptime of an R/3 system. This would particularly apply in the case of large tablespaces. The use of online reorgs during periods of low activity, which minimizes the length of time the tablespace is stopped, allows most routine downtime to be effectively avoided.

### OPERATIONAL CONSIDERATIONS

From an OS/390 perspective, R/3 utilizing DB2/390 as a database server is a complex

environment in that it requires the customization of the OS/390 OpenEdition environment (although this is mandatory for later releases of OS/390), the set up and verification of either HPDT or EESCON protocol for the use of ICLI, as well as exploitation of many features of DB2 V5. Potentially, for large and/or highly available systems, the complexity of data sharing and sysplex is also introduced.

However, in another sense, SAP simplifies a lot of tasks normally associated with an MVS environment. For example, SAP provides a security model within the application, negating the need to maintain security beyond what is necessary for systems staff and the tasks and processes involved. This in itself is quite simple and straightforward, although it will require some knowledge of the OpenEdition requirements of the security package used.

From the perspective of DB2, all the work performed by the SAP application appears to emanate from one user, SAPR3, who has open access to all of the SAP data. Workload management involves setting up an optimal environment for DB2 to carry out work, and for a productive environment, a separate LPAR for the SAP DB2 production workload is recommended. While this may not be practical for all customers, it allows predictable and accurate resource allocation and monitoring.

From a DB2 systems programmer's perspective, things, however, get somewhat

more complex, which is a classical understatement. As well as coping with monitoring new features like DSC, row-level locking and potentially DB2 data sharing, the organization and structure of the databases, tablespaces, tables and buffer pools is substantially more complex than most DB2 applications. As installed, only five buffer pools are defined for roughly 9,000 tables. While some immediate performance enhancing restructuring is recommended, most gains will come with tuning specific to the individual customer's workload and functional usage.

Tuning involves analyzing the most active tables and classifying them in terms of read activity, write activity and effectiveness of buffering. This allows particular buffer pools to be set up for each category of tables, and memory to be allocated in an optimal fashion amongst pools. While this is considered normal tuning, some unique features of SAP tables complicate the exercise, and the sheer number of tables adds yet more challenge. As the tables involved vary according to the particular combination of functions used by the customer, general guidelines for a tuning methodology are applicable, with a few rules of thumb existing for common tuning techniques.

#### **CONCLUSION**

In summary, a rapidly growing number of sites have chosen R/3 on DB2 for

OS/390 mainly because it allows them to build larger R/3 systems than ever before. One customer that recently installed SAP R/3 will implement a production system with a database totaling around 4TB, while a number of other customers are implementing systems capable of supporting more than 5,000 users. Using the features of DB2 outlined in this article, these sites can be confident that the R/3 system will be robust, and most importantly, capable of scaling upwards without striking any hard limitations until the systems reach a very large system size. 



**NaSPA member Tim Bohlsen works for SAP America and SCS America LLC as an SAP Basis consultant specializing in the architecture of large R/3 systems. He previously worked in Australia as an MVS systems programmer and SAP Basis consultant with BHP Information Technology. He can be contacted at [tim.bohlsen@sap-ag.de](mailto:tim.bohlsen@sap-ag.de).**

*©1998 Technical Enterprises, Inc. For reprints of this document contact [sales@naspa.net](mailto:sales@naspa.net).*