Written by Michael Thompson, March 2003 TA000184DBS

Technology Infrastructure

Butler Group Subscription Services

Database Management Systems
TECHNOLOGY AUDIT

InterSystems
Caché v5

Abstract *Caché is a post-relational database, a multi-dimensional relational/object database complete with a full featured ADE. Utilising a Unified Data Architecture (UDA), Caché takes the best of both the relational and object world and incorporates them into a development solution that answers the data overloads of modern organisations. Caché applications are highly scalable and can be deployed utilising less hardware resource than many current database systems. Any organisation that is striving to cut development costs, without sacrificing reliability will find a use for Caché, and the ease with which applications can be developed will come as a pleasant surprise to many. Caché has to overcome one negative aspect, and that is the commitment to the relational model that is prevalent across all organisations. Caché provides access to relational systems, but the conceptual obstacle will be harder to overcome. Free evaluation of the product is available, with flexible licensing options for commercial application usage.*

STRENGTHS
- Incorporates the best of relational and the object model.
- Speedy development, with prototyping taking only days.
- Deploys with smaller hardware requirements than many other solutions.
- Proven scalability.

WEAKNESSES
- Caché has to overcome the obstacle of being other than relational.

FUTURE POTENTIAL
The growth of XML and the need to store persisted XML will start to expose the weakness of the relational model. Caché is a solution that will take away the pain that is coming in this area.

Butler Group

FUNCTIONALITY

Caché is marketed as a post-relational database; a terminology that does not fully describe the product or its functionality. At one level, Caché is a multi-dimensional database (but not in the OLAP data-cube sense) which utilises InterSystems' Unified Data Architecture (UDA) to present the Caché database as both a SQL92-compliant relational database (RDBMS) and also as an Object-Oriented (O-O) Object Database Management System (ODMBS).

The UDA allows developers to take advantage of both the relational and object models; developing applications that can leverage the specific strengths of each individual model and effectively ignoring the limitations imposed by reliance on a single model.

The debate on the relative merits/demerits of the RDBMS as opposed to the ODBMS is one of long standing, and has supporters on both sides, although the installed base for RDBMS is so large as to make many of the conceptual arguments in support of ODBMS irrelevant. The introduction of Caché into the equation makes many implementation arguments redundant.

From a development standpoint, Caché is amongst a growing number of solutions that create a database by way of application development. As opposed to creating logical and physical models of the database, and implementing applications on top of them, although this can be done with Caché if required.

What helps set Caché apart from others of this type is the logical and physical robustness of the database engine. Many of the solutions of this type have a real weakness when they are considered at this level, and also when consideration is given to database administration. Caché is a complete DBMS with all the functionality in terms of performance, scalability, security, and management that one would expect from a 'mainstream' DBMS, but requiring less database administration.

Although Caché is put forward as a DBMS it is worth looking beyond this to the development environment itself, with its IDE known as Caché Studio. This is a highly intuitive GUI-based IDE that makes application development both quick and simple. Even developers with little knowledge of object methodologies will be able to create true object classes; led through the process by a series of Wizards.

In this instance, simplicity does not go hand-in-hand with any form of reduction in functionality. Data objects created in Caché support encapsulation, multiple inheritance, polymorphism, embedded objects, references, collections, relationships, and BLOBs; all of which leads to intuitive modelling of complex data.

The combination of the multi-dimensional data engine and the UDA leads to many advantages as detailed below:

- Application performance is enhanced by the removal of the overhead caused by table joins in a relational model.
- There is full support for relational access to the Caché database, using ODBC or JDBC, with the dual benefit of allowing connectivity to SQL-based reporting tools and improving performance of legacy systems. These legacy systems can be accessed through the Caché Gateway.
- There is no impedance mismatching between objects and relational tables.

- Caché utilises its own bit-map indexes. The compressions used for the indexes allow them to be updated quickly enough to be used within a transactional environment, which meets the changing needs of database querying.
- Caché classes can be exposed as Java, COM, or C++, and there is also full support for exposure as EJBs. Caché has an implementation of persistence management for EJBs that removes the need for any manual mapping between Java classes and relational tables.
- Caché comes complete with two interoperable scripting languages (Caché ObjectScript and Caché Basic) for method coding and business logic scripting.

When one considers Web application development specifically, Caché scores highly in terms of functionality and ease of use. Again, some of the specific functionality is detailed below:

- XML is fast becoming the *de facto* standard for application-to-application data sharing. One of the problems that developers face is the transformation aspects in terms of XML documents and traditional data stores. Caché provides automatic generation of XML documents and their DTDs or schemas from Caché classes as well as the reverse generation of classes from XML schemas.
- The world of Web services is also catered for by Caché, with the ability to publish any method as a Web service, including the generation of the WSDL.
- Classes may be exposed through Web Servers using Caché Server Pages (CSP) that reside and execute on the Caché data server. This provides performance advantages, and the centralisation of business logic with all the attendant benefits of speed and manageability of application change. CSPs can be created independently or directly from the most popular Web site development tool, Macromedia Dreamweaver, using an extension provided by InterSystems.
- CSPs can contain Caché Application Tags that execute functions either on the Caché data server or the Web browser. This helps split development between Web page design and Web page functionality, with the latter being enabled by coding the tags in the CSPs. Created tags can be customised for any specific Web-based application, but are also completely re-usable.
- Developers may script browser-based events to trigger procedures to be executed on the client or server. For example, when a field is validated they can avoid the need to submit the whole Web page, or to re-paint the page on the browser. This ensures reduced bandwidth and faster response times.

The combination of the UDA and multidimensional data engine provides a highly scalable architecture with the amount of hardware required significantly reduced when compared with other DBMSs. Installed implementations include distributed sites with tens of thousands of users. An example of a live implementation is with Partners Healthcare in the US which has implemented a series of more than 100 Caché applications and one terabyte production database servicing 33,000 users, although this does not imply any restriction to this number, simply being a real-world case.

The architecture is designed with scalability as a prime requirement, and it implements an Enterprise Cache Protocol (ECP), which allows an application to be spanned across a distributed server environment. Apart from scalability, this has an added advantage of reducing the level of administration.

DEPLOYMENT

Although Caché is easy to use, to fully exploit its features for enterprise-class application development, there is a need for skills in either Caché Object Script, or Caché Basic for server-side scripting. In real-world situations, there will also be a requirement for Java or .NET skills within the development organisation to integrate the Caché solution with existing application infrastructures.

Again, it is necessary to highlight the depth of functionality that is provided by Caché, and part of this functionality is to remove many of the development burdens associated with other ADEs. As with any solution of this type, the greater the initial skills of the developer, the quicker the result will be, but when one compares Caché with other ADEs, it becomes plainly obvious, within a very short space of time, that this is truly an intuitive environment. Whilst it would be plainly ridiculous to propound the theory that enterprise-class applications could be built by a complete tyro, the skills requirement is certainly reduced.

As an extension to this, it is also worthy of note that Caché presents an opportunity for developers familiar with one or other of relational and/or object models to extend that knowledge into the 'other' area with little additional training.

Detailing the length of time to develop an application is, naturally, dependent upon the nature and size of the application. However, the development time for a 'medium-size' system would be in the region of two to three months. The real positive aspect comes with the application prototyping, which would typically take only a few days. This provides a quick and simple route to managing larger deployments and reduces the costs associated with first-pass development.

InterSystems provide a comprehensive training program, with training for both developers and administrators. This training can be provided in a variety of formats, including classroom style, on-site standard, customised, and Web-based. The product itself comes with an interactive tutorial with examples for self-learning. InterSystems also run a certification program.

Support for the product is provided 24X7 on a global basis. Working closely with clients and customers is an important part of InterSystems strategy and ethos, as will be discussed in more depth in the following section.

The following large range of platforms are supported:

- Microsoft Windows 95, 98, NT, 2000, ME, and XP.
- OpenVMS.
- Tru64.
- HP-UX.
- AIX.
- RedHat Linux.
- SuSE Linux.
- Sun Solaris.

Interfacing with other database systems is by way of the provided ODBC and JDBC drivers; this gives Caché capability to external database sources.

PRODUCT STRATEGY

Caché as a database is primarily found embedded into applications, and as such the main route to market is through application partners. However, InterSystems does provide direct support for customers who choose to develop their application in-house.

With the majority of business carried out through application partners, InterSystems approach to this market is both flexible and indicative of the nature of the organisation itself.

Caché is available 'free' either on CD or by download from the InterSystems' Web site. This is a fully functional single-user version of the product with no time limitations. This gives any developer the opportunity to gauge its worth. Any applications built with the product require a licence key from InterSystems for deployment. The actual pricing for this can be based on number of users, CPU usage, or on a Web model. The pricing is flexible and forms an agreement between the developer and InterSystems.

Naturally, for larger implementations, this pricing arrangement would be worked out before full implementation started, but the business model does allow for smaller developments to take place with a valid pricing model available. This is a reflection of InterSystems working closely with all of its partners. The company sees that this is an effective route to market and recognises the benefits in terms of exposure that it gains.

When one considers other pricing models in the field of application development, where a large amount of money has to be paid up-front with no guarantee of success, InterSystems model has much to recommend it. Likewise, when development communities are considered, such as Delphi and Microsoft, there are strong market opportunities available.

New product developments may move InterSystems further upmarket, but as it is currently focused on increasing market penetration of Caché, it places no fixed minimum thresholds to its Partner agreements. The support and maintenance cost associated with the product, is calculated at 22% of licence cost. For this, clients receive 24X7 support and all software upgrades.

Apart from training, InterSystems does not provide other chargeable services, such as pre-development work, or deployment strategies; all of which are handled by its development partners.

InterSystems have a number of partnerships that support the Caché product; chief amongst these are:

- EDS Healthcare.
- ISOFT.
- McKesson.
- Misys Asset Management.
- Torex.

InterSystems also have key technology partnerships with:

- HP.
- IBM.
- Sun Microsystems.
- Intel.

Competition in the market of ADE products is intense, but InterSystems crosses over both that and the 'traditional' market. Key competitors in both areas would be recognised as Oracle and Microsoft, and from a development viewpoint any number of top-line vendors would have to come under consideration.

However, when the full functionality of Caché is realised and understood, many of the pure-play ADEs are lacking when it comes to database implementation, with a large cost associated with further development in this area.

Both the UDA and the multi-dimensional database engine support a scenario of growing data volumes, higher throughputs, and limitations on infrastructure expenditure. A solution that answers all these issues has to be recognised as worthy of consideration in the market place. With ROI such a major concern, the ability to develop new Web-based data-centric applications and deploying these on existing infrastructures shows a positive opportunity for InterSystems, and one that is unlikely to change in the foreseeable future.

This market opportunity has to be tempered with one negative factor; as with the pure O-ODBMS before it, the introduction of any new database model is met with scepticism and a degree of resistance both from other vendors and also some 'experts'. In respect of Caché, Butler Group firmly believes that this scepticism should be put aside in favour of the overwhelming evidence of current implementations within both the healthcare and financial markets (both notoriously demanding of transaction intensive applications).

The other factor that should be considered is that Caché is not *per se* a 'new' model, rather it is a co-joining of two existing models, and it is this that forms the strength, as it utilises proven methods implemented within current DBMS solutions.

▶ COMPANY PROFILE

Founded in 1978, InterSystems is a privately held company, headquartered in Cambridge, MA. It has regional offices across the globe, including, but not limited to:

- Australia.
- Brazil.
- France.
- Germany.
- Israel.
- Japan.
- Russia.
- United Kingdom.

As can be seen from the limited list of offices shown above, InterSystems is a truly global organisation.

Although the company has been in existence for over 25 years, it entered the current market with the release of Caché in 1997. It now employs over 450 people, and in association with its licensed VARs, it serves over 4 million licensed users across 88 different countries.

Of its 450 employees, there is a split in the following areas:

- 15% in Research and Development (R&D).
- 38% in Sales and Marketing.
- 16% in Support and Services.
- 31% in Marketing, Product Management, Administration, etc.

The R&D effort is based solely around the product and does not touch on any specific application work, which is handled by the development partners and community.

As a privately held company, InterSystems does not release any financial information, nor does it release any figures as regards to geographical split of revenues. It has, however, indicated to Butler Group, that it has been consistently profitable since its foundation.

Key clients that use Caché are:

- Ameritrade Holding Corporation.
- Cogent.
- Integrated Software Solutions.
- Great Ormond Street Hospital.
- NHS Information Authority.
- Partners Healthcare.
- The Department of Justice in Berne, Switzerland.
- US Federal Government.
- Whitbread.

As an indication of scale, the Partners Healthcare in the US which has implemented a series of more than 100 Caché applications and one terabyte production database servicing 33,000 users served by six Caché data servers, and 11 application servers. Despite the size of the implementation, the entire Caché network is serviced by less than four full-time DBAs.

▶ SUMMARY

InterSystems' Caché is one of those rare technology solutions that, 20 years down the line, all and sundry will take for granted and the benefits of which will be obvious to everyone. However, nothing comes easily, and there is no quick fix to getting a non-relational database message out into the market. The first crack in the seemingly invincible armour of the 'relational is the answer to everything' brigade has to be proven technology. This Caché has in abundance.

Without a doubt, Caché represents a major shift in the way that applications will be developed and deployed in the future, and with the growing importance of XML, Caché will be well placed to take advantage of many new market opportunities.

▶ CONTACT DETAILS

InterSystems Corporation
One Memorial Drive
Cambridge, MA 02142
US

Tel: +1 (617) 6210600
Fax: +1 (617) 4941631

www.InterSystems.com

InterSystems
InterSystems House, Tangier Lane
Eton, Windsor, SL4 6BB
UK

Tel: +44 (0)1753 855450
Fax: +44 (0)1753 855290

www.InterSystems.co.uk

For more information on Butler Group's Subscription Services, contact:

Europa House, 184 Ferensway, Hull, East Yorkshire, HU1 3UT, UK
Tel: +44 (0)1482 586149 Fax: +44 (0)1482 323577 www.butlergroup.com