



White Paper

20X Faster SQL. How Does Caché Do That?

Customers who switch to Caché from relational databases report that their average performance is up to 20 times faster, running on the same hardware, with no changes to the application. What is it about Caché that lets applications run so fast?

A) Multidimensional Data Engine

Caché stores data in the form of multidimensional arrays instead of in a series of flat tables. Data stored in a multidimensional structure is, by definition, already cross-referenced by as many parameters as the programmer chooses. It is possible to access or update data in a multidimensional database without performing the complicated and time-consuming joins required by relational database systems. Processing overhead is greatly reduced, so applications run much faster.

InterSystems makes the multidimensional data engine even more efficient by employing the concept of sparse arrays. Null cells within the Caché database do not take up any disk space. Less time is required to search for data or load it into local memory, because there are no “empty” data cells.

B) Transactional Bit-Map Indexing

In a bit-map index, a property of a class (column of a table) is described by an array of bits: for every possible value there is a string of bits representing each object in the class (row in the table). The bit is “1” if the object has that property value, “0” if it does not. The advantage of bit map indexes is that complex queries can be processed by performing Boolean operations (AND, OR) on the indexes – creating a temporary index that says exactly which instances (rows) fit the query conditions, without searching through the entire database.

Although bit-map indexes are known to significantly improve complex query response, they traditionally have been slow to build, and thus too cumbersome to use in anything other than data warehousing applications. However, Caché’s new bit-map indexing capability is nimble enough to use even with rapidly

changing transactional data.

C) Enterprise Cache Protocol

InterSystems' Enterprise Cache Protocol (ECP) can increase application performance in distributed systems. It is optimized for the thin-client architectures that are most commonly used today.

When information is requested across the network, the reply data package includes not only the desired piece of data, but also the entire database block where that data was stored. The natural data relationships inherent to objects and Caché's multidimensional data model make it likely that this "extra" data is related in some way to the original request.

The "extra" data is cached on the application server, where it is available to satisfy subsequent requests from the client, or in fact from any client connected to the application server. ECP automatically takes care of maintaining database locks, propagating changes back to the data server, etc.

The performance and scalability benefits of ECP are dramatic. Clients enjoy faster responses because they frequently use locally cached data. And network traffic is greatly reduced, so any given network can support many more application servers and clients.