WebFOCUS Hyperstage

Bigger Data, Faster Queries, and Less Effort

WebFOCUS Hyperstage is an embedded data store that allows customers to dramatically improve the performance of their BI environment. Combining a column-oriented data store and Knowledge Grid architecture with Information Builders' WebFOCUS BI platform, Hyperstage delivers a self-managing environment that is optimized for reporting and analytics.

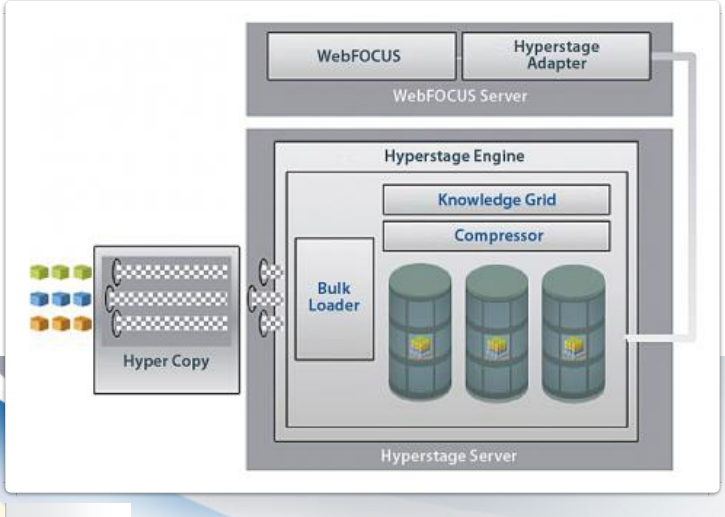

WebFOCUS Hyperstage

WebFOCUS Hyperstage is an OLAP database replacement. Rather than create a cube in a physical OLAP database, it aggregates the answer without the time-consuming process of loading the cube. Hyperstage combines the I/O advantage of in-memory analytics with an intelligent architecture that allows data to be stored on disk without sacrificing performance. It eliminates the need to create, build, and maintain indexes; partition data; or perform manual tuning. Instead, the built-in Knowledge Grid creates and maintains information about the data as it is being loaded and as queries are run.

Hyperstage at a Glance

- 90 percent less administrative effort than comparable solutions
- Less than half the cost of alternative solutions
- Supports up to 50 TB using a single server
- Ad hoc query response times are just as fast as anticipated queries
- Data compression rates of 10:1 to 40:1
- Modern column-oriented database organization

Efficient Data Management

As data is loaded into the Hyperstage environment, each field is immediately separated column by column, so that each item can be treated as a distinct table. The data within each column is accumulated into data packs of 65,536 values. Each data pack is compressed on an average ratio of 10:1, so that 1 TB of raw data can be stored on 100 GB of disk space.

Intelligent Query Processing

As a query comes through the Hyperstage Adapter, the Knowledge Grid determines which data packs are relevant to the query before decompressing any data. In some cases, the summary information already contained in the Knowledge Grid is sufficient to resolve the query. When this happens, results are dramatically accelerated because there is no need to access the detailed data store, and no need to decompress data packs.