THE EMPRESS HETEROGENEOUS DISTRIBUTED DATABASE AS A VEHICLE FOR ADVANCED METEOROLOGICAL APPLICATIONS

Njai Wong
Empress Software Inc.
Ontario, Canada

ABSTRACT

The Empress heterogenous distributed database has a number of leading edge features that make it an appropriate vehicle for developing advanced meteorological applications.

Meteorological information is collected at thousands of sites around the world, in hundreds of countries, and even from space. The data collection is highly dispersed and relies on a multitude of different hardware and software environments to collect, store and process it. The data itself may be in character, numeric, binary or date format. Weather applications range from nowcasting which require huge arrays of supercomputers to medium range forecasting using mini-supercomputers to establish historical trends for television weather service. Empress, the world's most advanced database, has the functionality and capability to work in each of these different environments and link them together seamlessly using the Empress heterogeneous fully distributed database capability.

The paper will outline the functionality of Empress and map that into the functionality required to build advanced meteorological applications.

1. INTRODUCTION

In 1990 the meteorological research paper entitled, "Environmental Database For The Naval Environmental Operational Nowcasting System" stated that:

> "A database which handles a wide variety of meteorological and oceanographic data has been described. This system is "user-friendly" in terms of providing a simple and logical access to diverse data types from application programs as well as interactive browse. New environmental data are easy to add because of the use of generic storage types. Storage volume and I/O time is minimized by data packing. The in-house software effort is reduced and portability enhanced by the use of a mainstream commercial data-management system". (Jurkevics et. al., 1990)

The commercial database management successfully used in that meteorological research project was Empress, the Heterogeneous Distributed Relational Database Management System. Since that time, Empress has evolved to keep up with the sophisticated demands of meteorological applications.

This paper outlines the current functionality of Empress that makes it well suited for advanced meteorological applications under the following headings:

2. Corporate and Product Philosophy
3. Heterogeneous Fully-Distributed Databases
4. Binary, BLOB, Multimedia, Bulk Data Types
5. User-Defined Functions and Operators
6. Performance
7. Product Architecture
8. Future Developments
9. Conclusions

2. CORPORATE AND PRODUCT PHILOSOPHY

Empress Software Inc. was founded in 1979 by John Kornatowski and Ivor Ladd two researchers in relational database management systems at the University of Toronto with the mission statement:

"To develop, market and support the world's most advanced data management software"

The company's market philosophy has been in database development software for scientific and engineering applications. The Empress product has been sold world wide to discriminating and knowledgeable developers.

The Empress product runs on UNIX, VMS, DOS, Cray UNICOS and certain real-time operating systems. The Empress family consists of Relational Database Management Systems, RDBMS utilities, SQL, Interfaces to C, Fortran F77 and ADA, 4th Generation Language Application Generator and Report Writer.

3. HETEROGENEOUS FULLY-DISTRIBUTED DATABASES

When local weather is affected by what has happened remotely, local information is no longer sufficient for weather prediction. For example, the cold temperature in Britain today is caused by the abnormal low fronts in Siberia. Therefore, the ability to access information from different sites, cities or countries becomes very important. Not only are there many sources of data but there is a growing amount of it. Today, we are dealing with terabytes of data. Empress's ability to store and process

large volumes of distributed data becomes important. The Empress heterogeneous distributed database management system allows you to:

- Access and use data anywhere on the network
- Choose the best computer hardware for a specific site or location
- Work seamlessly with other computer hardware architectures within the same network
- Configure networks freely with optimum combinations of supercomputers, mainframes, mid-range computers and workstations
- Configure networks freely with optimum combinations of clients, servers, multi-client multi-server and fully distributed

4. BINARY, BLOB, MULTIMEDIA, BULK DATA TYPES

There is much meteorological data such as, satellite images, gridded data, and latitude/longitude/time data that do not fit into traditional data processing data types. When using other database management system, this type of data would be stored as operating system files with absolutely no data management control over data access and integrity. The Empress database management system allows the meteorological user to store anything the computer can store as a datatype within the Empress database management system without imposing size limitation on your data. With Empress, you have the power to:

- Store and manipulate variable-length binary and text data, including images, voice, and sound
- Operate on data in an object-oriented environment while preserving relational model performance
- Store binary satellite images and gridded data and retrieve and manage them consistently

5. USER-DEFINED FUNCTIONS AND OPERATORS

Functions and operators provided by traditional database management systems are: sum, average, count, maximum, minimum, greater than, less than, equal to, greater or equal to, less or equal to, and not equal. For meteorological applications other functions and operators are required, such as standard deviation, distance calculation based on latitude and longitude, or temperature colder than 10 °C. With Empress you can:

- Use the extended scientific and math functions within Empress
- Link your system math library into the Empress executables

- Add your own user written functions and operators
- Make your own functions and operators available to all the Empress modules such as SQL the C interface, 4GL, etc.

6. PERFORMANCE

Empress's superior performance with all types of data is appreciated by the meteorological community. This has been proven many existing meteorological applications including the ability to capture weather data in real-time.

7. PRODUCT ARCHITECTURE

From its inception, Empress was designed and built around a virtual machine architecture. All Empress functionality works through that virtual machine. This resulted in a single interface to all operating system calls making Empress highly portable among computer manufacturers different hardware and operating systems. In addition to this design, the Empress source code was constructed as highly layered and strongly structured modules. This strong structuring guaranteed that all related functionality at a particular layer would be clustered thus ensuring fully reusable code. This architecture allows Empress application to:

- Adopt new technologies quickly
- Extend new functionalities easily
- Be highly portable to 32 and 64 bit environments

8. FUTURE DEVELOPMENTS

Future development at Empress that are of particular interest to the meteorological community include.

- Security
- Additional Object-oriented extensions
- Application internationalization
- GUI 4GL

9. CONCLUSION

A well-engineer software product like Empress allow the meteorological users to:

- Use workstations supercomputers and mid-range machines to configure a true heterogeneous fully distributed relational database management system spanning 32 and 64 bit architecture

- Access data anywhere on the network as if the data is on the local system and still maintain the autonomy of each site and the integrity of all the data
- Store and process large volume of scientific data
- Add application specific functions and operators to manipulate these data
- Tune the system to obtain superior performance

Thus the Empress Database Technology works effectively today for existing meteorological applications and is engineered to evolve for meteorological applications of the future.

10. REFERENCES

Jurkevics, A., Titus, R and Clark, J., 1990: "Environmental Database For The Naval Environmental Operational Nowcasting System" Naval Oceanographic And Atmospheric Research Laboratory, Monterey, California.

Wong, N., Ladd, I., 1993: "Software Portability Issues as Illustrated by the Empress RDBMS and 4GL" Supercomputers at Work 1993 Fall Proceedings Kyoto, Japan.