Ptool: A Scalable Persistent Object Manager*

R. L. Grossman and X. Qin
Laboratory for Advanced Computing
University of Illinois at Chicago

1 Goals

Ptool is a scalable persistent object manager developed by the Database Computing Project at the University of Illinois at Chicago. It is designed to provide low overhead, high performance access to large numbers of objects distributed over a hierarchical storage system. It is also designed to operate in loosely and tightly coupled clusters of workstations, including wide area clusters.

2 Design and Architecture

Ptool is designed to provide low overhead, high performance management of persistent objects. It provides scalable access to large numbers of objects by supporting multi-level caching and migration within a multi-level storage hierarchy.

The physical design of ptool is based upon three concepts: segments, folios and stores. A *segment* is a continuous range of virtual memory that is managed by ptool. A segment may contain one or more objects or a piece of a large object. A *folio* is a physical collection of segments. A *store* is a physical collection of folios. Folios are maintained by ptool as files.

Ptool contains three different managers. The Persistent Object Manager uses the virtual memory system and mapping to manage persistent objects. The Segment Cache Manager uses the virtual memory system to map segments from a local disk to virtual memory. The Folio Cache Manager manages segments and folios from network storage and hierarchical storage to local disk. The Folio Cache Manager interfaces to any hierarchical storage systems compliant with the IEEE Storage System Reference Model and uses pre-emptive prefetching to improve performance.

The caching, migration and replication is designed so that not only can segments be moved between nodes in local area clusters, but also segments and folios can be moved to nodes in wide area clusters. In other words, in addition to moving queries close to the data, the data can also be moved closer to the queries when appropriate.

3 Applications

Ptool has been used as the persistent object manager for a number of applications:

— for scientific stores, ranging from a real time path planning algorithm in aeronautics to a visualization tool for vortex rings in computational fluid dynamics;

— for stores containing experimental data describing high energy physics particle collisions, as part of the PASS Project;

— for multimedia data, ranging from a store of medical images to a store of video sequences;

— for stores of financial data, ranging from time series data to the Budget of the U.S. Government.

4 Status

The ptool distribution is available on the internet and consists of several related tools: ptool32 is a simple, easy to use object manager which provides 32 bits of persistent addressable object space. Ptool64 is a companion tool for applications which require a 64 bit address space. Ptool was first released in June, 1993. The current version 0.6 has been ported to Sun Sparcstations running SUNOS 4.1.3, IBM RS/6000's running AIX 3.2, Intel 486 platforms running Linux 0.99, and the IBM SP-1.

*This research was supported in part by NASA grant NAG2-513, DOE grant DE-FG02-92ER25133, and NSF grants IRI 9224605 and CDA-9303433. For more information, contact Robert Grossman, Laboratory for Advanced Computing, m/c 249, University of Illinois at Chicago, 851 South Morgan Street, Chicago, IL 60607, grossman@eecs.uic.edu.