Poster Abstract: Acxiom's Capacity On Demand Framework

For the past five years Acxiom Corporation has be transitioning its IT infrastructure from a large SMP server model to a highly parallel, distributed Grid model. Because Acxiom's internal processing needs are somewhat different from traditional Grid applications custom monitoring and control software, called the Apiary, was written in house. Initially a number of CORBA based services were moved to arrays of nodes, numbering from 20 to 200 members each. These arrays, called Hives, are statically allocated from a Grid consisting of +6000 nodes. A joint research project involving the University of Arkansas and Acxiom has been investigating strategies for automatically allocating nodes to Hives based on processing load in an attempt to create a more dynamic Apiary environment. To help test these strategies Acxiom has added a new feature to its control software named the Capacity On Demand framework (COD).

COD works by detecting highly loaded nodes within the Hives it monitors and creating duplicates of those nodes, a process we call cloning. A pool of idle nodes, available for use as clones, is kept in special Resource Hives that are shared among a number of Service Hives. To clone a node an idle node is allocated from a Resource Hive, the application software and data that are present on the overloaded node are placed on the idle node, the service is configured and it is brought on line. Once a clone is online workload is balanced across all of the nodes in the affected Hive. This process is repeated until the load on the Hive is brought within acceptable margins. When the load on a Hive drops below a programmed lower threshold the cloned nodes are released back to the Resource Hive.

We present a demonstration of dynamic node allocation using real nodes in the Acxiom Grid. This was accomplished using decision model code from UA and the COD software, driven by a suite of test applications, also written by Acxiom developers. The test suite consists of a client Hive that can impose arbitrary network and CPU loads on a server Hive. The client Hive allows a workload profile to be accurately reproduced allowing different allocation strategies to be tested under consistent conditions.

The test workload used is taken from monitoring data for a real service running real processing jobs. The workload profile, consisting of several periods of peak demand lasting approximately 10 minutes each, runs for ~2 hours. Graphs (*) of the CPU load on the server (target) Hive show an entire 2 hour cycle without COD, with COD using initial settings, and after tuning the COD settings to prevent rapid release events. COD rapidly detects heavy load conditions and adds resources to bring the CPU load back within acceptable margins. When the load dissipates the added nodes are released back into the Grid.

By incorporating COD into its Grid environment Acxiom hopes to improve the utilization of its hardware resources and reduce overall operating costs. Funding for University participation in this project was provided through the Acxiom Laboratory for Applied Research.

* Graphs are present on the actual poster.

Bio info:

Brandon Willis is a developer in the Grid Architecture Technical Unit at Acxiom Corporation and is primary author of the COD framework software. B.S. in Computer Science, University of Central Arkansas.

Jesse Perkins is a developer in the Grid Architecture Technical Unit at Acxiom Corporation and is primary author of the COD test suite software. B.S. in Computer Science, University of Central Arkansas.

Alan Rainey is a Technical Advisor at Acxiom Corporation specializing in distributed software architecture, grid computing, and grid security. B.S. in Computer Science, University of Central Arkansas.

Doug L. Hoffman, P.I., is a developer at Acxiom Corporation and architect for Acxiom's Grid infrastructure software. B.S. Applied Mathematics, Florida Institute of Technology, M.S. and PhD in Computer Science, University of North Carolina at Chapel Hill.