Vampir

Technical Details

Scalable Trace Analysis

Performance data originating from large-scale parallel applications can easily reach volumes in the order of hundreds of gigabytes. Interactive visualization and analysis of such data volumes require a parallel and highly scalable analysis engine. The Vampir tool chain addresses this issue with a parallel analysis component called VampirServer. Developers can interact with VampirServer from their desktop using the Vampir GUI as a performance data browser.

Concept

VampirServer implements a client-server approach that separates trace visualization from trace analysis. The visualization of performance data is done by the Vampir GUI that connects to the server component, which performs the complex processing of performance data. This allows the bulk of performance data to be kept close to its origin, i.e., where the data was collected. The server is capable of analyzing concurrently different sets of performance data. The corresponding visualization can be carried out on any laptop or desktop PC connected to the Internet. The tool can efficiently handle performance data volumes in the order of hundreds of gigabytes representing tens of thousands of concurrent application tasks. An overview of the architecture is depicted in the figure to the right.

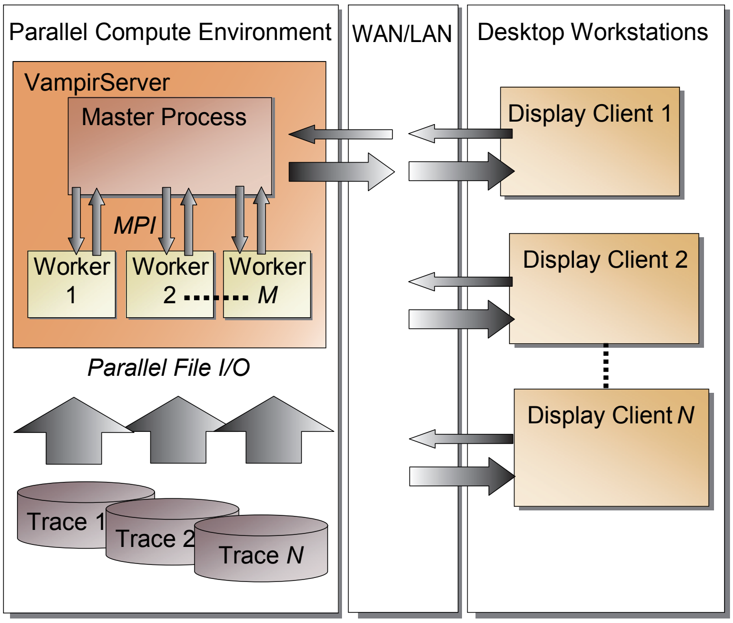

The distributed analysis architecture of VampirServer

Performance

VampirServer provides parallel event analysis algorithms that enable fast and interactive rendering of complex performance behavior. Very large data volumes can be processed due to the close integration with existing parallel production environments. The tool chain efficiently exploits available parallel compute power, high performance file systems, and large amounts of memory. The support for standard parallelization libraries such as MPI and Pthreads enables the analysis engine to operate on both distributed and shared memory systems.

Key Features

- Detailed visualization of large parallel applications‘ performance
- Scalable, parallel processing of performance data (event traces)
- Close integration with parallel production systems
- Distributed client-server approach
- No time consuming transport of large amounts of performance data to desktop needed
- Support of multiple performance data file formats, including the parallel Open Trace Format (OTF) and EPILOG

EU CONTACT
GWT-TUD GmbH
E-mail: sales@vampir.eu
Web: www.vampir.eu

U.S. CONTACT
ParaTools, Inc.
E-mail: info@paratools.com
Web: www.paratools.com

event.vampir.eu