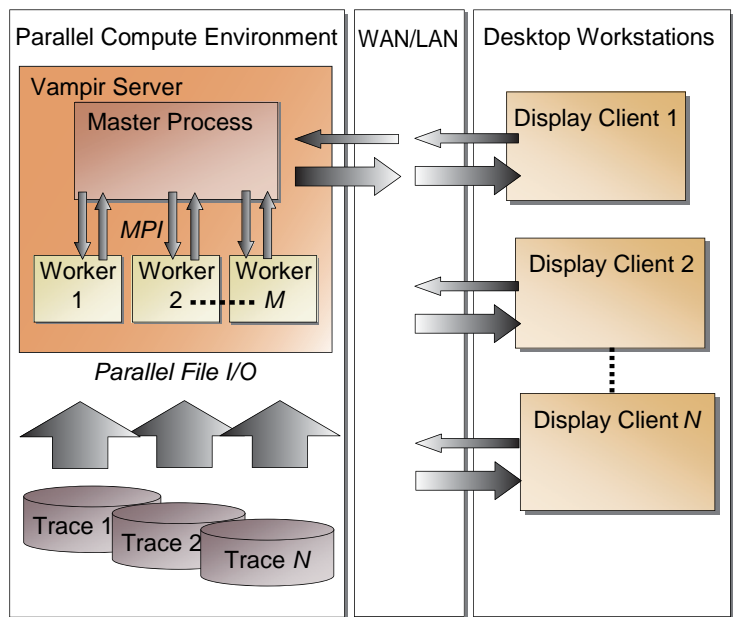


Performance data originating from large-scale parallel applications can easily reach volumes in the order of hundreds of gigabytes. Interactive visualization and analysis of such data volumes require a parallel and highly scalable analysis engine. The Vampir tool suite addresses this issue with a parallel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

Concept

VampirServer implements a client-server approach that separates trace visualization and analysis. The visualization of performance results is done by the Vampir GUI that connects to the server component, which performs the complex processing of performance data. This allows bulky performance data to be kept close to its origin, i.e. where the data was collected. The server is capable of analyzing different sets of performance data at the same time. 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 thousands of concurrent 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 which 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 suite 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.

Key Features

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



Contact

GWT-TUD GmbH
 Chemnitzer Str. 48b
 01187 Dresden, Germany
E-mail: sales@vampir.eu
Web: www.vampir.eu

U.S. Contact

ParaTools, Inc.

info@paratools.com
 www.paratools.com