Leibniz Computing Center Munich (LRZ) and Dresden University of Technology (TUD/ZHR)

Performance and Resource Profiling of User Jobs on the Hitachi SR8000-F1

Matthias Brehm, Holger Brunst, Wolfgang E. Nagel, Richard Patra, and Stephan Seidl
{ brehm | patra }@lrz.de { brunst | nagel | seidl }@zhr.tu-dresden.de
www.lrz.de www.tu-dresden.de

TeraFLOPS System with 112 Nodes and 1008 CPUs

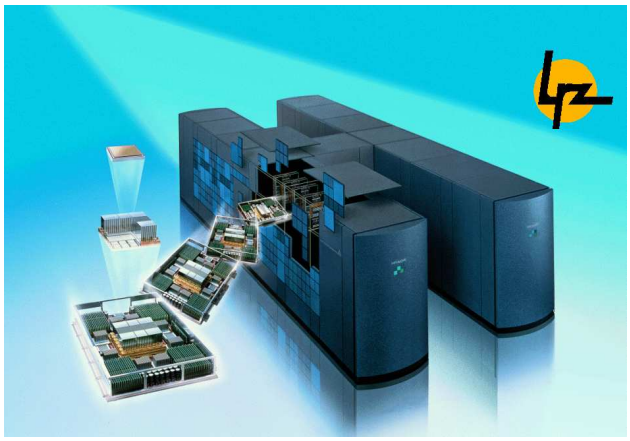

Motivation

Traditionally, VAMPIR is used as a tool for analyzing the performance of single programs.
Newer versions of VAMPIR come with features as grouping, filtering etc. to manage trace files from parallel applications with thousands of tasks, hardware performance counter information included. Consequently, VAMPIR offers itself to be a universal tool for monitoring and surveillance of very large systems, too.
To vizualize state and utilization of the HITACHI SR8000-F1, two data streams have to be translated into the VAMPIR file format. Firstly, this is high-level information from the batch queuing system, the partition manager, and the job accounting system. The source of the second stream is SAR data, containing accumulated system and processor information, i.e. resource consumption and hardware counters.

Resource Information Available on SR8000-F1

Number of nodes in use
User time
System time
Number of instructions per second
Number of FP instructions per second
Amount of free memory
Amount of wired memory
Number of load/stores per second
Numer of cache misses per second
Number of DTLB accesses per second
Number of system calls per second

System Time of the largest Job

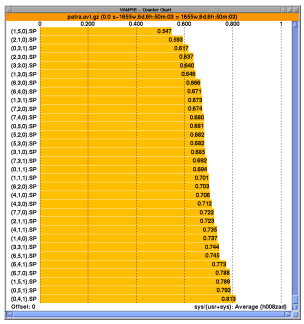

Performance of the largest Job

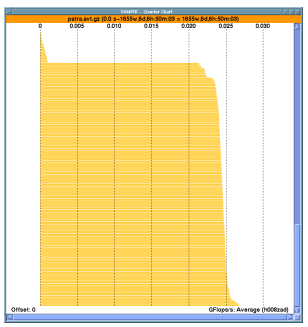

Solution to monitor a TeraFLOPS System

Map data onto VAMPIR Counters

NQS Database,
Partition Manager etc.
Information about Users, Jobs and Processes

Map data onto VAMPIR States

VAMPIR Trace File

Vizualize Data
Create Statistics
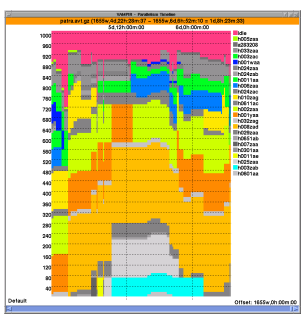
Parallellism on 1008 CPUs