

VampirServer

The [Vampir suite](#) is a set of tools for profiling MPI applications. Vampir itself is a proprietary GUI-based trace viewer, capable of displaying and analyzing traces stored in [OTF \(Open Trace Format\)](#).

Any valid OTF trace can be viewed in Vampir; the best way to generate such a trace is using the VampirTrace utility. (You may also be able to generate an OTF trace using other profiling/tracing utilities, such as TAU.) This page describes the use of the Vampir utility. For details on using VampirTrace to generate traces, see the software map entry for that utility.

Note that although Vampir and Intel Trace Analyzer are essentially different forks of the same software, they are incompatible with respect to trace file formats. Specifically, ITA reads traces stored in STF (Structured Trace Format), a proprietary format owned by Intel, and cannot read OTF traces; similarly, Vampir can no longer read STF traces.

Usage

VampirServer is the distributed analysis part of the Vampir suite. As such, a familiarity with using Vampir in it's standard, serial mode is advised. For information on this, please see our software page for Vampir on rajin.

We recommend launching VampirServer in an interactive PBS batch job. To do this, you need to submit the job to the queue:

```
qsub -I -lwalltime=... -ln_cpus=... -lmem=...
```

where the walltime, number of CPUs, and memory have been appropriately selected (I recommend using a multiple of 16 for n_cpus and then twice that in GB for the memory request - for example, 32 CPUs and 64GB of memory). If you want the job to start quicker, considering using the express queue, although note that it consumes SUs at a significantly faster rate. Once the job has started, you need to load the vampir module:

```
module load vampir/8.3.0
```

and then you can launch the VampirServer processes using

```
vampirserver start
```

Once it has started the VampirServer processes, you will get output similar to this:

```
Launching VampirServer...  
/apps/Modules/modulefiles/(3):ERROR:105: Unable to locate a modulefile for 'openmpi//opt/OTF2-1.4-64'  
VampirServer 8.3.0 (r9073)  
Licensed to Australian National University  
Running 15 analysis processes... (abort with vampirserver stop 18103)  
VampirServer listens on: r643:30005
```

Take note of the last line, since this gives you the hostname (r643) and port (30005) that VampirServer is listening on. Now, you can open a new SSH session to rajin and launch Vampir as per the instructions on it's software page - however this time you would select "Remote File" and enter the hostname and port from above in the fields provided, select "socket" as the connection type, and then select "connect". Then it will ask you for the trace file as usual.

If you're using the alternative method of specifying the path to the trace file on the command line, you should enter it in the form host:port:path when launching Vampir, and use an absolute path. For example, for the VampirServer instance above, you would use

```
vampir r643:30005:/path/to/file.otf
```