

Intermittent Job Failures Due to InfiniBand Issues

Category: Resolved

Problem: Since the March 21 dedicated time, Pleiades users have reported new MPI/InfiniBand errors.

Status: Resolved

Actions:

05.09.12 - The `subnet_timeout` parameter for `ib0` was reverted back to `18`. This has eliminated the `ib0: timing out` messages in the system log files and there have been no further reports of job failures with `inet_connect to host failed` messages.

If you are still experiencing this issue, contact the NAS Control Room: (800) 331-8737, (650) 604-4444, support@nas.nasa.gov.

Background:

During the March 21, 2012 Pleiades dedicated time, 1,728 Harpertown nodes were removed in preparation for new hardware. At the same time, the operating system was upgraded to use OpenFabrics Enterprise Distribution (OFED) 1.5.4.1. Since then, a few users have reported the following error messages, which were not seen before:

With the SGI MPT library:

```
MPI ERROR: inet_connect to host port ##### failed 110
MPI: daemon terminated: - job aborting
```

In addition, NAS system administrators have noticed the following new messages in the system log files:

```
r34i0n9 kernel: [25813.752772] ib0: timing out; 3 sends not completed
```

Such error messages have all been traced to intermittent IB timing-out issues. It is not clear whether this points to new IB problems or that the new OFED simply does a better job of

logging IB timing-out issues. The systems staff have noticed more IB issues lately and are working to resolve them.

Article ID: 308

Last updated: 18 May, 2012

Updates on Issues -> Resolved -> Intermittent Job Failures Due to InfiniBand Issues

<http://www.nas.nasa.gov/hecc/support/kb/entry/308/?ajax=1>