

知 R5500 G6搭配H20模组做SPEC CPU压测失败

CPU 驱动安装升级 张玉林 2024-11-07 发表

组网及说明

GCC版本8.5.0, 推荐11.3.0
OS和内核版本
[root@localhost ~]# cat /etc/redhat-release
CentOS Linux release 8.2.2004 (Core)
[root@localhost ~]# uname -r
4.18.0-193.6.3.el8_2.v1.6.x86_64
GPU驱动版本550.54.14

告警信息

问题描述

R5500 G6+H20模组+Intel 8468V, 做SPEC CPU压测时, 1 copy正常, 192 copies失败

```
### Binaries
Running 583.bwaves_r_train base ic19.lul.lin-core-avx512-rate-20280386_revA (1 copy) [2024-10-24 23:13:46]
Running 589.cactu64_r_train base ic19.lul.lin-core-avx512-rate-20280386_revA (1 copy) [2024-10-24 23:14:06]
Running 588.namd_r_train base ic19.lul.lin-core-avx512-rate-20280386_revA (1 copy) [2024-10-24 23:14:14]
Running 510.povray_r_train base ic19.lul.lin-core-avx512-rate-20280386_revA (2 copy) [2024-10-24 23:14:52]
Running 511.povray_r_train base ic19.lul.lin-core-avx512-rate-20280386_revA (1 copy) [2024-10-24 23:15:08]
Running 519.blw_r_train base ic19.lul.lin-core-avx512-rate-20280386_revA (1 copy) [2024-10-24 23:15:08]
Running 521.wrf_r_train base ic19.lul.lin-core-avx512-rate-20280386_revA (1 copy) [2024-10-24 23:15:11]
Running 526.blender_r_train base ic19.lul.lin-core-avx512-rate-20280386_revA (1 copy) [2024-10-24 23:15:20]
Running 527.cam4_r_train base ic19.lul.lin-core-avx512-rate-20280386_revA (1 copy) [2024-10-24 23:17:33]
Running 538.imagick_r_train base ic19.lul.lin-core-avx512-rate-20280386_revA (1 copy) [2024-10-24 23:17:45]
Running 544.nab_r_train base ic19.lul.lin-core-avx512-rate-20280386_revA (1 copy) [2024-10-24 23:18:01]
Running 549.fotonik3d_r_train base ic19.lul.lin-core-avx512-rate-20280386_revA (1 copy) [2024-10-24 23:18:51]
Running 554.roms_r_train base ic19.lul.lin-core-avx512-rate-20280386_revA (1 copy) [2024-10-24 23:19:15]
Running 597.specrand_r_train base ic19.lul.lin-core-avx512-rate-20280386_revA (1 copy) [2024-10-24 23:20:38]
Access: 2060.bwaves_r, 1007.cactu64_r, 1508.namd_r, 1x510.povray_r, 1x511.povray_r, 1x519.blw_r, 1x521.wrf_r, 1x526.blender_r, 1x527.cam4_r, 1x538.imagick_r, 1x544.nab_r, 1x549.fotonik3d_r, 1x554.roms_r, 1x597.specrand_r
binaries: selected: 583.bwaves_r, 589.cactu64_r, 588.namd_r, 510.povray_r, 519.blw_r, 521.wrf_r, 526.blender_r, 527.cam4_r, 538.imagick_r, 544.nab_r, 549.fotonik3d_r, 554.roms_r, 597.specrand_r
compiling binaries
Up to date 583.bwaves_r base ic19.lul.lin-core-avx512-rate-20280386_revA
Up to date 589.cactu64_r base ic19.lul.lin-core-avx512-rate-20280386_revA
Up to date 588.namd_r base ic19.lul.lin-core-avx512-rate-20280386_revA
Up to date 510.povray_r base ic19.lul.lin-core-avx512-rate-20280386_revA
Up to date 519.blw_r base ic19.lul.lin-core-avx512-rate-20280386_revA
Up to date 521.wrf_r base ic19.lul.lin-core-avx512-rate-20280386_revA
Up to date 526.blender_r base ic19.lul.lin-core-avx512-rate-20280386_revA
Up to date 527.cam4_r base ic19.lul.lin-core-avx512-rate-20280386_revA
Up to date 538.imagick_r base ic19.lul.lin-core-avx512-rate-20280386_revA
Up to date 544.nab_r base ic19.lul.lin-core-avx512-rate-20280386_revA
Up to date 549.fotonik3d_r base ic19.lul.lin-core-avx512-rate-20280386_revA
Up to date 554.roms_r base ic19.lul.lin-core-avx512-rate-20280386_revA
Up to date 597.specrand_r base ic19.lul.lin-core-avx512-rate-20280386_revA
Setting up Run Directories
Setting up 583.bwaves_r refrate (ref) base ic19.lul.lin-core-avx512-rate-20280386_revA (192 copies): run_base.refrate_ic19.lul.lin-core-avx512-rate-20280386_revA.0800-0191
Setting up 589.cactu64_r refrate (ref) base ic19.lul.lin-core-avx512-rate-20280386_revA (192 copies): run_base.refrate_ic19.lul.lin-core-avx512-rate-20280386_revA.0800-0191
Setting up 588.namd_r refrate (ref) base ic19.lul.lin-core-avx512-rate-20280386_revA (192 copies): run_base.refrate_ic19.lul.lin-core-avx512-rate-20280386_revA.0800-0191
Setting up 510.povray_r refrate (ref) base ic19.lul.lin-core-avx512-rate-20280386_revA (192 copies): run_base.refrate_ic19.lul.lin-core-avx512-rate-20280386_revA.0800-0191
Setting up 511.povray_r refrate (ref) base ic19.lul.lin-core-avx512-rate-20280386_revA (192 copies): run_base.refrate_ic19.lul.lin-core-avx512-rate-20280386_revA.0800-0191
Setting up 519.blw_r refrate (ref) base ic19.lul.lin-core-avx512-rate-20280386_revA (192 copies): run_base.refrate_ic19.lul.lin-core-avx512-rate-20280386_revA.0800-0191
Setting up 521.wrf_r refrate (ref) base ic19.lul.lin-core-avx512-rate-20280386_revA (192 copies): run_base.refrate_ic19.lul.lin-core-avx512-rate-20280386_revA.0800-0191
Setting up 526.blender_r refrate (ref) base ic19.lul.lin-core-avx512-rate-20280386_revA (192 copies): run_base.refrate_ic19.lul.lin-core-avx512-rate-20280386_revA.0800-0191
Setting up 527.cam4_r refrate (ref) base ic19.lul.lin-core-avx512-rate-20280386_revA (192 copies): run_base.refrate_ic19.lul.lin-core-avx512-rate-20280386_revA.0800-0191
Setting up 538.imagick_r refrate (ref) base ic19.lul.lin-core-avx512-rate-20280386_revA (192 copies): run_base.refrate_ic19.lul.lin-core-avx512-rate-20280386_revA.0800-0191
Setting up 544.nab_r refrate (ref) base ic19.lul.lin-core-avx512-rate-20280386_revA (192 copies): run_base.refrate_ic19.lul.lin-core-avx512-rate-20280386_revA.0800-0191
Setting up 549.fotonik3d_r refrate (ref) base ic19.lul.lin-core-avx512-rate-20280386_revA (192 copies): run_base.refrate_ic19.lul.lin-core-avx512-rate-20280386_revA.0800-0191
Setting up 554.roms_r refrate (ref) base ic19.lul.lin-core-avx512-rate-20280386_revA (192 copies): run_base.refrate_ic19.lul.lin-core-avx512-rate-20280386_revA.0800-0191
Running Binaries
Running (43) 583.bwaves_r refrate (ref) base ic19.lul.lin-core-avx512-rate-20280386_revA (192 copies) [2024-10-24 23:20:48]
Running (43) 589.cactu64_r refrate (ref) base ic19.lul.lin-core-avx512-rate-20280386_revA (192 copies) [2024-10-24 23:25:20]
Running (43) 588.namd_r refrate (ref) base ic19.lul.lin-core-avx512-rate-20280386_revA (192 copies) [2024-10-25 00:00:54]
Running (43) 510.povray_r refrate (ref) base ic19.lul.lin-core-avx512-rate-20280386_revA (192 copies) [2024-10-25 00:05:50]
Running (43) 511.povray_r refrate (ref) base ic19.lul.lin-core-avx512-rate-20280386_revA (192 copies) [2024-10-25 00:31:40]
Running (43) 519.blw_r refrate (ref) base ic19.lul.lin-core-avx512-rate-20280386_revA (192 copies) [2024-10-25 00:40:36]
Running (43) 521.wrf_r refrate (ref) base ic19.lul.lin-core-avx512-rate-20280386_revA (192 copies) [2024-10-25 00:55:11]
Running (43) 526.blender_r refrate (ref) base ic19.lul.lin-core-avx512-rate-20280386_revA (192 copies) [2024-10-25 01:16:03]
root@localhost:~#
```

过程分析

跑speed和rate_int可以通过, 跑rate_fp失败, 调整散热模式无效, 卸载GPU驱动后未复现问题, 更换550.54.15版本的驱动后压测通过

解决方法

550.54.14版本的驱动在CPU压满的情况下, GPU驱动初始化不成功, 会一直有nv_open_q进程占用CPU资源, 且550.54.14版本的驱动H3C没有引入, 不确定还会有哪些问题, 550.54.15有引入, 建议使用该版本

测试步骤参考:

1. 安装 GCC11.3.0

```
1. wget http://ftp.gnu.org/gnu/gcc/gcc-11.3.0/gcc-11.3.0.tar.gz
2. tar xvf gcc-11.3.0.tar.gz
3. cd gcc-11.3.0
4. ./contrib/downloadprerequisites
5. mkdir build
6. cd build
7. ./configure --enable-checking=release --enable-languages=c,c++,fortran --disable-multilib
8. make -j 32
9. make install
```

2. 创建运行软链接

```
1. # 删除原有运行依赖，并指向/usr/local/bin/gcc，每次关闭 shell 后都需要重新运行
2. yum -y remove gcc g++ gcc-gfortran
3. rm -rf /usr/bin/gcc;rm -rf /usr/bin/g++;rm -rf /usr/bin/gfortran
4. ln -s /usr/local/bin/gcc /usr/bin/gcc
5. ln -s /usr/local/bin/g++ /usr/bin/g++
6. ln -s /usr/local/bin/gfortran /usr/bin/gfortran
7.
8. strings /usr/lib64/libstdc++.so.6 | grep CXXABI
9. rm -f /usr/lib64/libstdc++.so.6
10. ln -s /usr/local/lib64/libstdc++.so.6.0.29 /usr/lib64/libstdc++.so.6
11. strings /usr/lib64/libstdc++.so.6 | grep CXXABI
12. Export LD_PRELOAD=$LD_PRELOAD:/home/lh/gcc-11.3.0/x8664-pc-linux-gnu/libgfortran-5.3.0/libgfortran.so.5:/home/lh/jemalloc-5.3.0/lib/libjemalloc.so.2:/home/lh/gcc-11.3.0/x8664-pc-linux-gnu/libquadmath-5.3.0/libquadmath.so.0
```

3. 安装依赖库

```
1. yum install libnsl numactl
```

4. 安装 jemalloc5.3.0

```
1. wget https://github.com/jemalloc/jemalloc/archive/refs/tags/5.3.0.tar.gz
2. cp 5.3.0.tar.gz /home/spec
3. tar -zxvf 5.3.0.tar.gz
4. cd jemalloc-5.3.0/
5. ./autogen.sh
6. make && make install
```

5. 安装 spce cpu2017

```
1. mount cpu2017-1.1.8.iso /mnt
2. cd /mnt
3. mkdir /home/spec
4. ./install.sh -d /home/spec
```

运行跑分

1. 配置文件

配置 rate-hygon.cfg 文件

1. 将 Baseline Tuning Flags 中 COPTIMIZE、CXXOPTIMIZE、FOPTIMIZE 三行配置中 jemalloc 的位置
2. 改为实际安装 jemalloc5.3.0 的位置，如 -L/home/spec/jemalloc-5.3.0/lib/

配置文件信息可见 <https://www.spec.org/cpu2017/Docs/config.html#specmake>

2. 配置运行脚本

配置 gcc-rate-cpu2017.sh

1. 如果需要单独运行某个跑分程序（比如 500）：
2. `runcpu -c rate-hygon.cfg 500 -l -n 1`
3. 调试常用指令
4. # 只编译不运行，方便调试编译环境
5. `action runsetup`
6. # 将存在的编译文件删除
7. `action realclean`
8. # 重新编译
9. `rebuild`

runcpu 命令可见 <https://www.spec.org/cpu2017/Docs/runcpu.html#iter>

3. 运行

配置好 rate-hygon.cfg 文件和运行脚本后，将 cfg 文件放入~/spec/config 目录，将 gcc-rate-cpu2017.sh 放入~/spec 目录，运行即可。