

# Arvados - Bug #11396

## Network saturation

03/30/2017 04:41 PM - Joshua Randall

<b>Status:</b>	New	<b>Start date:</b>	03/30/2017
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assigned To:</b>		<b>% Done:</b>	0%
<b>Category:</b>	Performance	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>			

### Description

Some of our extremely I/O intensive jobs are able to completely saturate the network links. When that happens, jobs fail. We have a particular workload which if we run it on our whole cluster, it fails every time (and fairly early on).

Some failure messages are from SLURM:

2017-03-16\_10:39:54 [z8ta6-8i9sb-2j5ypo3mug47vdw](#)

```
22464 154 stderr starting: ['srun', '--nodelist=humgen-05-08', '-n1', '-c1', '-N1', '-D', '/data/crunch-tmp', '--job-name=z8ta6-8i9sb-2j5ypo3mug47vdw.154.27774', 'bash', '-c', 'if [ -e \\/data\\/crunch\\/tmp\\/crunch\\/job\\/task\\/humgen\\/05\\/08\\.6 ]; then rm -rf \\/data\\/crunch\\/tmp\\/crunch\\/job\\/task\\/humgen\\/05\\/08\\.6; fi; mkdir -p \\/data\\/crunch\\/tmp\\/crunch\\/job\\/data\\/crunch\\/tmp\\/crunch\\/job\\/work \\/data\\/crunch\\/tmp\\/crunch\\/job\\/task\\/humgen\\/05\\/08\\.6 \\/data\\/crunch\\/tmp\\/crunch\\/job\\/task\\/humgen\\/05\\/08\\.6\\.keep && cd \\/data\\/crunch\\/tmp\\/crunch\\/job && MEM=$(awk \\'($1 == "MemTotal:") {print $2}\\' </proc/meminfo) && SWAP=$(awk \\'($1 == "SwapTotal:") {print $2}\\' </proc/meminfo) && MEMLIMIT=$(( ($MEM * 95) / (7 * 100) )) && let SWAPLIMIT=$MEMLIMIT+$SWAP && declare -a VOLUMES=() && if which crunchrunner >/dev/null ; then VOLUMES+=("--volume=$(which crunchrunner):/usr/local/bin/crunchrunner:ro") ; fi && if test -f /etc/ssl/certs/ca-certificates.crt ; then VOLUMES+=("--volume=/etc/ssl/certs/ca-certificates.crt:/etc/arvados/ca-certificates.crt:ro") ; elif test -f /etc/pki/tls/certs/ca-bundle.crt ; then VOLUMES+=("--volume=/etc/pki/tls/certs/ca-bundle.crt:/etc/arvados/ca-certificates.crt:ro") ; fi && exec arv-mount --read-write --mount-by-pdh=by_pdh --mount-tmp=tmp --crunchstat-interval=10 --allow-other --file-cache=5242880000 \\/data\\/crunch\\/tmp\\/crunch\\/job\\/task\\/humgen\\/05\\/08\\.6\\.keep --exec crunchstat -cgroup-root=\\/sys\\/fs\\/cgroup -cgroup-parent=docker -cgroup-cid=/data/crunch-tmp/crunch-job/z8ta6-ot0gb-k3bagiubyoaty7-0.cid -poll=10000 /usr/bin/docker run --add-host=api.arvados.sanger.ac.uk:172.17.180.10 --add-host=humgen-01-01:172.17.180.10 --add-host=humgen-01-01.internal.sanger.ac.uk:172.17.180.10 --add-host=humgen-01-02:172.17.180.11 --add-host=humgen-01-02.internal.sanger.ac.uk:172.17.180.11 --add-host=humgen-01-03:172.17.180.12 --add-host=humgen-01-03.internal.sanger.ac.uk:172.17.180.12 --add-host=humgen-02-01:172.17.180.13 --add-host=humgen-02-01.internal.sanger.ac.uk:172.17.180.13 --add-host=humgen-02-02:172.17.180.14 --add-host=humgen-02-02.internal.sanger.ac.uk:172.17.180.14 --add-host=humgen-02-03:172.17.180.15 --add-host=humgen-02-03.internal.sanger.ac.uk:172.17.180.15 --add-host=humgen-03-01:172.17.180.16 --add-host=humgen-03-01.internal.sanger.ac.uk:172.17.180.16 --add-host=humgen-03-02:172.17.180.17 --add-host=humgen-03-02.internal.sanger.ac.uk:172.17.180.17 --add-host=humgen-03-03:172.17.180.18 --add-host=humgen-03-03.internal.sanger.ac.uk:172.17.180.18 --add-host=humgen-04-01:172.17.180.19 --add-host=humgen-04-01.internal.sanger.ac.uk:172.17.180.19 --add-host=humgen-04-02:172.17.180.20 --add-host=humgen-04-02.internal.sanger.ac.uk:172.17.180.20 --add-host=humgen-04-03:172.17.180.21 --add-host=humgen-04-03.internal.sanger.ac.uk:172.17.180.21 --name=z8ta6-ot0gb-k3bagiubyoaty7-0 --attach=stdout --attach=stderr --attach=stdin -i --cidfile=/data/crunch-tmp/crunch-job/z8ta6-ot0gb-k3bagiubyoaty7-0.cid --sig-proxy --memory=${MEMLIMIT}k --memory-swap=${SWAPLIMIT}k --volume=\\/data\\/crunch\\/tmp\\/crunch\\/job\\/src\\:\\/data\\/crunch\\/tmp\\/crunch\\/job\\/src\\:ro --volume=\\/data\\/crunch\\/tmp\\/crunch\\/job\\/opt\\:\\/data\\/crunch\\/tmp\\/crunch\\/job\\/opt\\:ro --volume=\\/data\\/crunch\\/tmp\\/crunch\\/job\\/task\\/humgen\\/05\\/08\\.6\\.keep\\/by_pdh\\:\\/keep\\:ro --volume=\\/data\\/crunch\\/tmp\\/crunch\\/job\\/task\\/humgen\\/05\\/08\\.6\\.keep\\/tmp\\:\\/keep_tmp --volume=tmp "${VOLUMES[@]}" --env=CRUNCH_JOB_DOCKER_RUN_ARGS=\\ \\ \\ \\ --add-host=api\\.arvados\\.sanger\\.ac\\.uk\\:172\\.17\\.180\\.10\\ \\ \\ \\ --add-host=humgen\\/01\\/01\\:172\\.17\\.180\\.10\\ \\ \\ \\ --add-host=humgen\\/01\\/01\\.internal\\.sanger\\.ac\\.uk\\:172\\.17\\.180\\.10\\ \\ \\ \\ --add-host=humgen\\/01\\/02\\:172\\.17\\.180\\.11\\ \\ \\ \\ --add-host=humgen\\/01\\/02\\.internal\\.sanger\\.ac\\.uk\\:172\\.17\\.180\\.11\\ \\ \\ \\ --add-host=humgen\\/01\\/03\\:172\\.17\\.180\\.12\\ \\ \\ \\ --add-host=humgen\\/01\\/03\\.internal\\.sanger\\.ac\\.uk\\:172\\.17\\.180\\.12\\ \\ \\ \\ --add-host=humgen\\/02\\/01\\:172\\.17\\.180\\.13\\ \\ \\ \\ --add-host=humgen\\/02\\/01\\.int\\.ernal\\.sanger\\.ac\\.uk\\:172\\.17\\.180\\.13\\ \\ \\ \\ --add-host=humgen\\/02\\/02\\:172\\.17\\.180\\.14\\ \\ \\ \\ --add-host=humgen\\/02\\/02\\.internal\\.sanger\\.ac\\.uk\\:172\\.17\\.180\\.14
```

```
\\ \-\\-add\\-host\\=humgen\\-02\\-03\\:172\\:17\\:180\\:15\\ \-\\-add\\-host\\=humgen\\-02\\-03
\\.internal\\.sanger\\.ac\\.uk\\:172\\:17\\:180\\:15\\ \-\\-add\\-host\\=humgen\\-03\\-01\\:172\\
.17\\:180\\:16\\ \-\\-add\\-host\\=humgen\\-03\\-01\\:172\\:17\\:180\\:16\\ \-\\-add\\-host\\=humgen\\-0
0\\:16\\ \-\\-add\\-host\\=humgen\\-03\\-02\\:172\\:17\\:180\\:17\\ \-\\-add\\-host\\=humgen\\-0
3\\-02\\:172\\:17\\:180\\:17\\ \-\\-add\\-host\\=humgen\\-03\\-03\\:172\\:17\\:180\\:18\\ \-\\-add\\-host\\=humgen\\-03\\-03\\
:172\\:17\\:180\\:18\\ \-\\-add\\-host\\=humgen\\-04\\-01\\:172\\:17\\:180\\:19\\ \-\\-add\\-host\\=humg
en\\-04\\-01\\:172\\:17\\:180\\:19\\ \-\\-add\\-host\\=humgen\\-04\\-02\\:172\\:17\\:180\\:20\\ \-\\-add\\-host\\=humgen\\-04\\-02\\:172\\:17\\:180\\:20\\ \-\\-add\\-host\\=humgen\\-04\\-03\\:172\\:17\\:180\\:21\\ \-\\-add\\-host\\
=humgen\\-04\\-03\\:172\\:17\\:180\\:21 --env=TASK_SEQUENCE\\=1 --e
nv=TASK_KEEPMOUNT\\=\\keep --env=JOB_PARAMETER_INPUTS_COLLECTION\\=dcc280be70e5d40106c22dcb5e600a
4d\\+7406818 --env=CRUNCH_SRC_COMMIT\\=794d64bc0ceb8bd4112397fd63bb3f97ab67e2b4 --env=TASK_QSEQUEN
CE\\=154 --env=CRUNCH_INSTALL\\=\\data\\crunch\\-tmp\\crunch\\-job\\opt --env=CRUNCH_GIT_ARCHI
VE_HASH\\=4b8cf8e1274f8d418cb494d3ea978152 --env=CRUNCH_REFRESH_TRIGGER\\=\\tmp\\crunch_refresh
trigger --env=ARVADOS_API_TOKEN\\=[...] --env=JOB_PARAMETER_INTERVAL_LISTS_COLLECTION\\=88bedd9be1
345b26e144067d6b985e27\\+10562 --env=CRUNCH_WORK\\=\\data\\crunch\\-tmp\\crunch\\-job\\work --
env=CRUNCH_TMP\\=\\data\\crunch\\-tmp\\crunch\\-job --env=TASK_TMPDIR\\=\\tmp\\crunch\\-job\\
-task\\-work\\humgen\\-05\\-08\\.6 --env=JOB_UUID\\=z8ta6\\-8i9sb\\-2j5ypo3mug47vdw --env=CRUNCH_
JOB_UUID\\=z8ta6\\-8i9sb\\-2j5ypo3mug47vdw --env=TASK_SLOT_NUMBER\\=6 --env=CRUNCH_SRC_URL\\=\\va
r\\lib\\arvados\\internal\\.git --env=TASK_SLOT_NODE\\=humgen\\-05\\-08 --env=JOB_SCRIPT\\=gat
k\\-genotypegvcfs\\.py --env=CRUNCH_NODE_SLOTS\\=7 --env=JOB_PARAMETER_REFERENCE_COLLECTION\\=a83bd
4e5a26a64612322f21515d93bab\\+6190 --env=CRUNCH_JOB_DOCKER_BIN\\=\\usr\\bin\\docker --env=TASK_
WORK\\=\\tmp\\crunch\\-job\\-task\\-work\\humgen\\-05\\-08\\.6 --env=TASK_KEEPMOUNT_TMP\\=\\ke
ep_tmp --env=ARVADOS_API_HOST\\=api\\.arvados\\.sanger\\.ac\\.uk --env=JOB_WORK\\=\\tmp\\crunch\\
-job\\-work --env=TASK_UUID\\=z8ta6\\-ot0gb\\-k3bagiubyoeaty7 --env=CRUNCH_SRC\\=\\data\\crunch
\\-tmp\\crunch\\-job\\src --env=HOME\\=\\tmp\\crunch\\-job\\-task\\-work\\humgen\\-05\\-08\\
.6 sha256\\:a4fa354645c849421c8bfc8da71c5b8adeldf1fe25792d196c59f88c11f5ceac /bin/sh -c 'python -c
"from pkg_resources import get_distribution as get; print "\\Using Arvados SDK version\\", get(\\
"arvados-python-client\\").version">&2 2>/dev/null; mkdir -p "/tmp/crunch-job-work" "/tmp/crunch-j
ob-task-work/humgen-05-08.6" && if which stdbuf >/dev/null; then exec stdbuf --output=0 --erro
r=0 \\data\\crunch\\-tmp\\crunch\\-job\\src\\crunch_scripts\\gat\\-genotypegvcfs\\.py ; el
se exec \\data\\crunch\\-tmp\\crunch\\-job\\src\\crunch_scripts\\gat\\-genotypegvcfs\\.py
; fi''']
2017-03-16_10:39:54 z8ta6-8i9sb-2j5ypo3mug47vdw
22464 154 stderr srun: error: Task launch for 17752.160 failed on node humgen-05-08: Communicatio
n connection failure
2017-03-16_10:39:54 z8ta6-8i9sb-2j5ypo3mug47vdw
22464 backing off node humgen-05-08 for 60 seconds
2017-03-16_10:39:54 z8ta6-8i9sb-2j5ypo3mug47vdw
22464 154 stderr srun: error: Application launch failed: Communication connection failure
2017-03-16_10:39:54 z8ta6-8i9sb-2j5ypo3mug47vdw
22464 backing off node humgen-05-08 for 60 seconds
2017-03-16_10:39:54 z8ta6-8i9sb-2j5ypo3mug47vdw
22464 154 stderr srun: Job step aborted: Waiting up to 32 seconds for job step to finish.
2017-03-16_10:39:54 z8ta6-8i9sb-2j5ypo3mug47vdw
22464 154 stderr srun: error: Timed out waiting for job step to complete
2017-03-16_10:39:54 z8ta6-8i9sb-2j5ypo3mug47vdw
22464 155 stderr starting: ['srun', '--nodelist=humgen-05-09', '-n1', '-c1', '-N1', '-D', '/data/crunch
-tmp', '--job-name=z8ta6-8i9sb-2j5ypo3mug47vdw.155.27784', 'bash', '-c', 'if [ -e \\data\\crunch\\-t
mp\\crunch\\-job\\task\\humgen\\-05\\-09\\.6 ]; then rm -rf \\data\\crunch\\-tmp\\crunch\\-j
ob\\task\\humgen\\-05\\-09\\.6; fi; mkdir -p \\data\\crunch\\-tmp\\crunch\\-job \\data\\cru
nch\\-tmp\\crunch\\-job\\work \\data\\crunch\\-tmp\\crunch\\-job\\task\\humgen\\-05\\-09\\
.6 \\data\\crunch\\-tmp\\crunch\\-job\\task\\humgen\\-05\\-09\\.6\\.keep && cd \\data\\crunc
h\\-tmp\\crunch\\-job && MEM=${awk '\\($1 == "MemTotal:") {print $2}\\' </proc/meminfo} && SWAP=${aw
k '\\($1 == "SwapTotal:") {print $2}\\' </proc/meminfo} && MEMLIMIT=$(( ($MEM * 95) / (7 * 100) )) &&
let SWAPLIMIT=$MEMLIMIT+$SWAP && declare -a VOLUMES=() && if which crunchrunner >/dev/null ; then
VOLUMES+=("--volume=$(which crunchrunner):/usr/local/bin/crunchrunner:ro") ; fi && if test -f /et
c/ssl/certs/ca-certificates.crt ; then VOLUMES+=("--volume=/etc/ssl/certs/ca-certificates.crt:/etc
/arvados/ca-certificates.crt:ro") ; elif test -f /etc/pki/tls/certs/ca-bundle.crt ; then VOLUMES+=
("--volume=/etc/pki/tls/certs/ca-bundle.crt:/etc/arvados/ca-certificates.crt:ro") ; fi && exec arv
-mount --read-write --mount-by-pdh=by_pdh --mount-tmp=tmp --crunchstat-interval=10 --allow-other -
-file-cache=5242880000 \\data\\crunch\\-tmp\\crunch\\-job\\task\\humgen\\-05\\-09\\.6\\.keep
--exec crunchstat -cgroup-root=\\sys\\fs\\cgroup -cgroup-parent=docker -cgroup-cid=/data/crunch
-tmp/crunch-job/z8ta6-ot0gb-82cmb7pgdpktrd5-0.cid -poll=10000 /usr/bin/docker run --add-host=api.
arvados.sanger.ac.uk:172.17.180.10 --add-host=humgen-01-01:172.17.180.10 --add-host=humgen-01-01.i
nternal.sanger.ac.uk:172.17.180.10 --add-host=humgen-01-02:172.17.180.11 --add-host=humgen-01-02.i
```

```
nternal.sanger.ac.uk:172.17.180.11 --add-host=humgen-01-03:172.17.180.12 --add-host=humgen-01-03.i
nternal.sanger.ac.uk:172.17.180.12 --add-host=humgen-02-01:172.17.180.13 --add-host=humgen-02-01.i
nternal.sanger.ac.uk:172.17.180.13 --add-host=humgen-02-02:172.17.180.14 --add-host=humgen-02-02.i
nternal.sanger.ac.uk:172.17.180.14 --add-host=humgen-02-03:172.17.180.15 --add-host=humgen-02-03.i
nternal.sanger.ac.uk:172.17.180.15 --add-host=humgen-03-01:172.17.180.16 --add-host=humgen-03-01.i
nternal.sanger.ac.uk:172.17.180.16 --add-host=humgen-03-02:172.17.180.17 --add-host=humgen-03-02.i
nternal.sanger.ac.uk:172.17.180.17 --add-host=humgen-03-03:172.17.180.18 --add-host=humgen-03-03.i
nternal.sanger.ac.uk:172.17.180.18 --add-host=humgen-04-01:172.17.180.19 --add-host=humgen-04-01.i
nternal.sanger.ac.uk:172.17.180.19 --add-host=humgen-04-02:172.17.180.20 --add-host=humgen-04-02.i
nternal.sanger.ac.uk:172.17.180.20 --add-host=humgen-04-03:172.17.180.21 --add-host=humgen-04-03.i
nternal.sanger.ac.uk:172.17.180.21 --name=z8ta6-ot0gb-82cmb7pgdpktrd5-0 --attach=stdout --attach=s
tderr --attach=stdin -i --cidfile=/data/crunch-tmp/crunch-job/z8ta6-ot0gb-82cmb7pgdpktrd5-0.cid -
-sig-proxy --memory=${MEMLIMIT}k --memory-swap=${SWAPLIMIT}k --volume=\\data\\crunch\\-tmp\\cru
nch\\-job\\src\\:\\data\\crunch\\-tmp\\crunch\\-job\\src\\:ro --volume=\\data\\crunch\\-tmp
\\crunch\\-job\\opt\\:\\data\\crunch\\-tmp\\crunch\\-job\\opt\\:ro --volume=\\data\\crunch
\\-tmp\\crunch\\-job\\task\\humgen\\-05\\-09\\6\\.keep\\by_pdh\\:\\keep\\:ro --volume=\\dat
a\\crunch\\-tmp\\crunch\\-job\\task\\humgen\\-05\\-09\\6\\.keep\\tmp\\:\\keep_tmp --volume=
/tmp "${VOLUMES[@]}" --env=CRUNCH_JOB_DOCKER_RUN_ARGS\\=\\ \\-\\-add\\-host\\=api\\.arvados\\.sang
er\\.ac\\.uk\\:172\\.17\\.180\\.10\\ \\-\\-add\\-host\\=humgen\\-01\\-01\\:172\\.17\\.180\\.10\\ \\
\\-\\-add\\-host\\=humgen\\-01\\-01\\.internal\\.sanger\\.ac\\.uk\\:172\\.17\\.180\\.10\\ \\-\\-add
\\-host\\=humgen\\-01\\-02\\:172\\.17\\.180\\.11\\ \\-\\-add\\-host\\=humgen\\-01\\-02\\.internal
\\.sanger\\.ac\\.uk\\:172\\.17\\.180\\.11\\ \\-\\-add\\-host\\=humgen\\-01\\-03\\:172\\.17\\.180\\.
12\\ \\-\\-add\\-host\\=humgen\\-01\\-03\\.internal\\.sanger\\.ac\\.uk\\:172\\.17\\.180\\.12\\ \\-
\\-add\\-host\\=humgen\\-02\\-01\\:172\\.17\\.180\\.13\\ \\-\\-add\\-host\\=humgen\\-02\\-01\\.int
ernal\\.sanger\\.ac\\.uk\\:172\\.17\\.180\\.13\\ \\-\\-add\\-host\\=humgen\\-02\\-02\\:172\\.17\\.
180\\.14\\ \\-\\-add\\-host\\=humgen\\-02\\-02\\.internal\\.sanger\\.ac\\.uk\\:172\\.17\\.180\\.14
\\ \\-\\-add\\-host\\=humgen\\-02\\-03\\:172\\.17\\.180\\.15\\ \\-\\-add\\-host\\=humgen\\-02\\-03
\\.internal\\.sanger\\.ac\\.uk\\:172\\.17\\.180\\.15\\ \\-\\-add\\-host\\=humgen\\-03\\-01\\:172\\
.17\\.180\\.16\\ \\-\\-add\\-host\\=humgen\\-03\\-01\\.internal\\.sanger\\.ac\\.uk\\:172\\.17\\.18
0\\.16\\ \\-\\-add\\-host\\=humgen\\-03\\-02\\:172\\.17\\.180\\.17\\ \\-\\-add\\-host\\=humgen\\-0
3\\-02\\.internal\\.sanger\\.ac\\.uk\\:172\\.17\\.180\\.17\\ \\-\\-add\\-host\\=humgen\\-03\\-03\\
:172\\.17\\.180\\.18\\ \\-\\-add\\-host\\=humgen\\-03\\-03\\.internal\\.sanger\\.ac\\.uk\\:172\\.1
7\\.180\\.18\\ \\-\\-add\\-host\\=humgen\\-04\\-01\\:172\\.17\\.180\\.19\\ \\-\\-add\\-host\\=humg
en\\-04\\-01\\.internal\\.sanger\\.ac\\.uk\\:172\\.17\\.180\\.19\\ \\-\\-add\\-host\\=humgen\\-04\\
-02\\:172\\.17\\.180\\.20\\ \\-\\-add\\-host\\=humgen\\-04\\-02\\.internal\\.sanger\\.ac\\.uk\\:1
72\\.17\\.180\\.20\\ \\-\\-add\\-host\\=humgen\\-04\\-03\\:172\\.17\\.180\\.21\\ \\-\\-add\\-host\\
=humgen\\-04\\-03\\.internal\\.sanger\\.ac\\.uk\\:172\\.17\\.180\\.21 --env=TASK_SEQUENCE\\=1 --e
nv=TASK_KEEPMOUNT\\=\\keep --env=JOB_PARAMETER_INPUTS_COLLECTION\\=dcc280be70e5d40106c22dcb5e600a
4d\\+7406818 --env=CRUNCH_SRC_COMMIT\\=794d64bc0ceb8bd4112397fd63bb3f97ab67e2b4 --env=TASK_QSEQUEN
CE\\=155 --env=CRUNCH_INSTALL\\=\\data\\crunch\\-tmp\\crunch\\-job\\opt --env=CRUNCH_GIT_ARCHI
VE_HASH\\=4b8cf8e1274f8d418cb494d3ea978152 --env=CRUNCH_REFRESH_TRIGGER\\=\\tmp\\crunch_refresh_
trigger --env=ARVADOS_API_TOKEN\\=[...] --env=JOB_PARAMETER_INTERVAL_LISTS_COLLECTION\\=88bedd9be1
345b26e144067d6b985e27\\+10562 --env=CRUNCH_WORK\\=\\data\\crunch\\-tmp\\crunch\\-job\\work --
env=CRUNCH_TMP\\=\\data\\crunch\\-tmp\\crunch\\-job --env=TASK_TMPDIR\\=\\tmp\\crunch\\-job\\
-task\\work\\humgen\\-05\\-09\\6 --env=JOB_UUID\\=z8ta6\\-8i9sb\\-2j5ypo3mug47vdw --env=CRUNCH_
JOB_UUID\\=z8ta6\\-8i9sb\\-2j5ypo3mug47vdw --env=TASK_SLOT_NUMBER\\=6 --env=CRUNCH_SRC_URL\\=\\va
r\\lib\\arvados\\internal\\.git --env=TASK_SLOT_NODE\\=humgen\\-05\\-09 --env=JOB_SCRIPT\\=gat
k\\-genotypegvcfs\\.py --env=CRUNCH_NODE_SLOTS\\=7 --env=JOB_PARAMETER_REFERENCE_COLLECTION\\=a83bd
4e5a26a64612322f21515d93bab\\+6190 --env=CRUNCH_JOB_DOCKER_BIN\\=\\usr\\bin\\docker --env=TASK_
WORK\\=\\tmp\\crunch\\-job\\-task\\work\\humgen\\-05\\-09\\6 --env=TASK_KEEPMOUNT_TMP\\=\\ke
ep_tmp --env=ARVADOS_API_HOST\\=api\\.arvados\\.sanger\\.ac\\.uk --env=JOB_WORK\\=\\tmp\\crunch\\
-job\\-work --env=TASK_UUID\\=z8ta6\\-ot0gb\\-82cmb7pgdpktrd5 --env=CRUNCH_SRC\\=\\data\\crunch
\\-tmp\\crunch\\-job\\src --env=HOME\\=\\tmp\\crunch\\-job\\-task\\work\\humgen\\-05\\-09\\.
6 sha256\\:a4fa354645c849421c8bfc8da71c5b8adeldf1fe25792d196c59f88c11f5ceac /bin/sh -c 'python -c
"from pkg_resources import get_distribution as get; print \\\"Using Arvados SDK version\\\", get(\\
"arvados-python-client\\").version">&2 2>/dev/null; mkdir -p "/tmp/crunch-job-work" "/tmp/crunch-j
ob-task-work/humgen-05-09.6" && if which stdbuf >/dev/null ; then exec stdbuf --output=0 --erro
r=0 \\data\\crunch\\-tmp\\crunch\\-job\\src\\crunch_scripts\\gat\\-genotypegvcfs\\.py ; el
se exec \\data\\crunch\\-tmp\\crunch\\-job\\src\\crunch_scripts\\gat\\-genotypegvcfs\\.py
; fi'']
2017-03-16_10:39:54 z8ta6-8i9sb-2j5ypo3mug47vdw
22464 155 stderr srunk: error: Task launch for 17752.161 failed on node humgen-05-09: Communicatio
n connection failure
2017-03-16_10:39:54 z8ta6-8i9sb-2j5ypo3mug47vdw
22464 backing off node humgen-05-09 for 60 seconds
2017-03-16_10:39:54 z8ta6-8i9sb-2j5ypo3mug47vdw
22464 155 stderr srunk: error: Application launch failed: Communication connection failure
```

```
2017-03-16_10:39:54 z8ta6-8i9sb-2j5ypo3mug47vdw
22464 backing off node humgen-05-09 for 60 seconds
2017-03-16_10:39:54 z8ta6-8i9sb-2j5ypo3mug47vdw
22464 155 stderr srunit: Job step aborted: Waiting up to 32 seconds for job step to finish.
2017-03-16_10:39:54 z8ta6-8i9sb-2j5ypo3mug47vdw
22464 155 stderr srunit: error: Timed out waiting for job step to complete
```

**While others are from API clients:**

```
2017-03-16_10:41:59 z8ta6-8i9sb-2j5ypo3mug47vdw
22464 154 child 27774 on humgen-05-08.6 exit 233 success=
2017-03-16_10:41:59 z8ta6-8i9sb-2j5ypo3mug47vdw
22464 154 ERROR: Task process exited 233, but never updated its task record to indicate success and record its output.
2017-03-16_10:47:02 API call /job_tasks/z8ta6-ot0gb-k3bagiubyoeaty7 failed: 500 SSL read timeout:
2017-03-16_10:47:02 SSL read timeout: at /usr/share/perl5/Net/HTTP/Methods.pm line 256
2017-03-16_10:47:02 at /usr/local/lib/perl/5.14.2/Net/SSL.pm line 222
2017-03-16_10:47:02 Net::SSL::die_with_error('LWP::Protocol::https::Socket=GLOB(0x2deea18)', 'SSL read timeout') called at /usr/local/lib/perl/5.14.2/Net/SSL.pm line 230
2017-03-16_10:47:02 Net::SSL::__ANON__('ALRM') called at /usr/local/lib/perl/5.14.2/Net/SSL.pm line 233
2017-03-16_10:47:02 eval {...} called at /usr/local/lib/perl/5.14.2/Net/SSL.pm line 233
2017-03-16_10:47:02 Net::SSL::read('LWP::Protocol::https::Socket=GLOB(0x2deea18)', 'HTTP/1.1 504 Gateway Time-out\x{d}\x{a}Server: nginx/1.10.2\x{d}\x{a}Date: Th...', 1024, 0) called at /usr/share/perl5/Net/HTTP/Methods.pm line 256
2017-03-16_10:47:02 Net::HTTP::Methods::my_readline('LWP::Protocol::https::Socket=GLOB(0x2deea18)', 'Status') called at /usr/share/perl5/Net/HTTP/Methods.pm line 343
2017-03-16_10:47:02 Net::HTTP::Methods::read_response_headers('LWP::Protocol::https::Socket=GLOB(0x2deea18)', 'laxed', 1, 'junk_out', 'ARRAY(0x2909438)') called at /usr/share/perl5/LWP/Protocol/http.pm line 378
2017-03-16_10:47:02 LWP::Protocol::http::request('LWP::Protocol::https=HASH(0x2d9c228)', 'HTTP::Request=HASH(0x2dd44b0)', undef, undef, undef, 180) called at /usr/share/perl5/LWP/UserAgent.pm line 192
2017-03-16_10:47:02 eval {...} called at /usr/share/perl5/LWP/UserAgent.pm line 191
2017-03-16_10:47:02 LWP::UserAgent::send_request('LWP::UserAgent=HASH(0x2d4b818)', 'HTTP::Request=HASH(0x2dd44b0)', undef, undef) called at /usr/share/perl5/LWP/UserAgent.pm line 274
2017-03-16_10:47:02 LWP::UserAgent::simple_request('LWP::UserAgent=HASH(0x2d4b818)', 'HTTP::Request=HASH(0x2dd44b0)', undef, undef) called at /usr/share/perl5/LWP/UserAgent.pm line 282
2017-03-16_10:47:02 LWP::UserAgent::request('LWP::UserAgent=HASH(0x2d4b818)', 'HTTP::Request=HASH(0x2dd44b0)') called at /usr/share/perl5/Arvados/Request.pm line 58
2017-03-16_10:47:02 Arvados::Request::process_request('Arvados::Request=HASH(0x27f3218)') called at /usr/share/perl5/Arvados/ResourceMethod.pm line 104
2017-03-16_10:47:02 Arvados::ResourceMethod::execute('Arvados::ResourceMethod=HASH(0x2c9c5f8)', 'uuid', 'z8ta6-ot0gb-k3bagiubyoeaty7', 'job_task', 'Arvados::ResourceProxy=HASH(0x2d371f0)') called at /usr/share/perl5/Arvados/ResourceProxy.pm line 15
2017-03-16_10:47:02 Arvados::ResourceProxy::save('Arvados::ResourceProxy=HASH(0x2d371f0)') called at /var/www/arvados-api/shared/vendor_bundle/ruby/2.1.0/gems/arvados-cli-0.1.20170217221854/bin/crunch-job line 1192
2017-03-16_10:47:02 main::reapchildren() called at /var/www/arvados-api/shared/vendor_bundle/ruby/2.1.0/gems/arvados-cli-0.1.20170217221854/bin/crunch-job line 1029
2017-03-16_10:47:02 at /usr/local/lib/perl/5.14.2/Net/SSL.pm line 233
2017-03-16_10:47:02 Net::SSL::read('LWP::Protocol::https::Socket=GLOB(0x2deea18)', 'HTTP/1.1 504 Gateway Time-out\x{d}\x{a}Server: nginx/1.10.2\x{d}\x{a}Date: Th...', 1024, 0) called at /usr/share/perl5/Net/HTTP/Methods.pm line 256
2017-03-16_10:47:02 Net::HTTP::Methods::my_readline('LWP::Protocol::https::Socket=GLOB(0x2deea18)', 'Status') called at /usr/share/perl5/Net/HTTP/Methods.pm line 343
2017-03-16_10:47:02 Net::HTTP::Methods::read_response_headers('LWP::Protocol::https::Socket=GLOB(0x2deea18)', 'laxed', 1, 'junk_out', 'ARRAY(0x2909438)') called at /usr/share/perl5/LWP/Protocol/http.pm line 378
2017-03-16_10:47:02 LWP::Protocol::http::request('LWP::Protocol::https=HASH(0x2d9c228)', 'HTTP::Request=HASH(0x2dd44b0)', undef, undef, undef, 180) called at /usr/share/perl5/LWP/UserAgent.pm line 192
2017-03-16_10:47:02 eval {...} called at /usr/share/perl5/LWP/UserAgent.pm line 191
2017-03-16_10:47:02 LWP::UserAgent::send_request('LWP::UserAgent=HASH(0x2d4b818)', 'HTTP::Request=HASH(0x2dd44b0)', undef, undef) called at /usr/share/perl5/LWP/UserAgent.pm line 274
```

```
2017-03-16_10:47:02 LWP::UserAgent::simple_request('LWP::UserAgent=HASH(0x2d4b818)', 'HTTP::Request=HASH(0x2dd44b0)', undef, undef) called at /usr/share/perl5/LWP/UserAgent.pm line 282
2017-03-16_10:47:02 LWP::UserAgent::request('LWP::UserAgent=HASH(0x2d4b818)', 'HTTP::Request=HASH(0x2dd44b0)') called at /usr/share/perl5/Arvados/Request.pm line 58
2017-03-16_10:47:02 Arvados::Request::process_request('Arvados::Request=HASH(0x27f3218)') called at /usr/share/perl5/Arvados/ResourceMethod.pm line 104
2017-03-16_10:47:02 Arvados::ResourceMethod::execute('Arvados::ResourceMethod=HASH(0x2c9c5f8)', 'uid', 'z8ta6-ot0gb-k3bagiubyoeaty7', 'job_task', 'Arvados::ResourceProxy=HASH(0x2d371f0)') called at /usr/share/perl5/Arvados/ResourceProxy.pm line 15
2017-03-16_10:47:02 Arvados::ResourceProxy::save('Arvados::ResourceProxy=HASH(0x2d371f0)') called at /var/www/arvados-api/shared/vendor_bundle/ruby/2.1.0/gems/arvados-cli-0.1.20170217221854/bin/crunch-job line 1192
2017-03-16_10:47:02 main::reapchildren() called at /var/www/arvados-api/shared/vendor_bundle/ruby/2.1.0/gems/arvados-cli-0.1.20170217221854/bin/crunch-job line 1029
2017-03-16_10:47:02 at /usr/share/perl5/Arvados/ResourceProxy.pm line 15
```

To work around this problem, I limited the set of crunch nodes to only those that are also running keep stores and have what basically amounts to a fabric interconnect between them (because they are all connected to a single 10Gbps switch). In that configuration, the same job runs without issue.

Our configuration has:

- 12 nodes each with their own 10Gbps interface on the same 10Gbps switch, one master/keep node and 11 crunch/keep nodes.
- 16 crunch nodes on their own 10Gbps blade chassis switch with a single 10Gbps uplink to the other switch (this is the bottleneck)

I think it would be better if Arvados could cope with situations in which it is possible to saturate network links. Tom Clegg and I thought of several possibilities for how the system might be made to better handle these situations (not intended to be an exhaustive list):

- keepstore: detect the situation when links to certain keep clients appear to be saturated and back off the send rate (client GET)
- keep client: same as above but for PUTs
- run a keep proxy on each node and implement some bandwidth throttling rules in it rather than in the clients
- put nodes that are on the far side of a bottleneck behind a transparent squid proxy that throttles keep bandwidth
- use QoS/CoS rules in the network to deprioritize keep traffic (difficult given the network equipment in this particular situation)
- make API server/clients better able to handle slow / high latency links
- make SLURM better able to handle slow / high latency links
- keepstore configuration: globally limit max buffers so that the total across the cluster cannot saturate any links (I could do this now and it should work but at a substantial performance hit for nodes not behind the bottleneck)
- keepstore: allow setting buffer limits (or more directly set bandwidth throttling limits) between the keepstore and configurable sets of clients

#### Related issues:

Related to Arvados - Bug #11460: [SDK] avoid interfering with socket open/clo...

**In Progress**

## History

### #1 - 03/30/2017 04:42 PM - Joshua Randall

Might be useful in the context of running a squid proxy between keep clients and keepstores:

- [https://github.com/frispete/squid\\_dedup](https://github.com/frispete/squid_dedup)
- <http://wiki.squid-cache.org/Features/DelayPools>

### #2 - 03/30/2017 04:44 PM - Joshua Randall

Possibly not relevant, but the example job that failed under saturation conditions ended up failing with the following log output:

```
2017-03-16_10:47:25 salloc: Relinquishing job allocation 17752
2017-03-16_10:47:25 salloc: Job allocation 17752 has been revoked.
2017-03-16_10:47:25 close failed in file object destructor:
2017-03-16_10:47:25 sys.excepthook is missing
2017-03-16_10:47:25 lost sys.stderr
```

### #3 - 08/28/2017 09:46 PM - Tom Morris

- Target version set to Arvados Future Sprints

### #4 - 07/06/2021 09:24 PM - Peter Amstutz

- Target version deleted (Arvados Future Sprints)