







Going to Light Speed with DataWarp

An Administrators Perspective



Tina Declerck and Dave Paul
CUG 2016 – May 10, 2016

- **Hardware – 144 nodes**
 - **2 SSDs per node**
 - **4 devices - nvme**
 - **Intel P3608**
 - **Ability to increase endurance (DWPD)**
 - **Decreases available space**
 - **NERSC configured with 10 DWPD – default 3**

DWPD

DataWarp configuration



- **Uses DVS to project to compute nodes**
 - Each DW node is a DVS server
 - Limits access to GPFS in CLE 5.2
- **DW scheduler daemon**
 - Runs on sdb
- **ReSTful API – gunicorn**
 - Runs on mom/login node
 - Uses nginx as the http server

- **Assigned 2 ways**
 - Per job
 - Persistent
- **#DW directives in job scripts**
 - Private mode
 - Striped
 - Type: currently only scratch supported
 - How much space needed

- Pools define a set of DataWarp nodes with a specific configuration
- DataWarp supports multiple pools
 - Native SLURM does NOT
- Granularity is configured at the node and pool levels
 - Pool granularity defines the smallest unit that can be allocated per node

- **Session**
 - Equates to a job ID
- **Instance**
 - DataWarp space allocated to a job or persistent over many jobs
- **Fragment**
 - Portions of the instance on each node allocated to it

But wait, there's more...



- **Configuration**
 - Defines how a DW instance is used
- **Namespace**
 - A configuration can have 0 or more namespaces
 - Basically a directory or folder in a scratch configuration

We're not done yet...



- **Registration**
 - Binds a session with a configuration
 - Maintains information for stage-in/stage-out
- **Activation**
 - An available instance configuration on a set of nodes

Putting it all together

NERSC



96 node job
DW striped
Type=scratch

8 node job
Type=private

Type=scratch
persistent
striped

POOL

General problem solving - dwstat



- sess state token creator owner created expiration nodes
- 2520 CA--- myBBname CLI 33333 2016-02-19T13:45:33 never 0
- **3041** CA--- u1_bb1 CLI 11111 2016-03-02T15:01:01 never 0
- 6185 CA--- 2128492 SLURM 55555 2016-05-09T07:13:58 never 96

- inst state sess bytes nodes created expiration intact label public confs
- 2234 CA--- 2520 212.91GiB 1 2016-02-19T13:45:33 never true myBBname true 1
- **2550** CA--- **3041** 1.04TiB 5 2016-03-02T15:01:02 never true u1_bb1 true 1
- 5534 CA--- 6185 1.87TiB 9 2016-05-09T07:13:58 never true l6185-0 false 1

- conf state inst type access_type activs
- 2505 CA--- 2234 scratch stripe 0
- **2821** CA--- **2550** scratch stripe 0
- 5811 CA--- 5534 scratch stripe 1

- reg state sess conf wait
- 5877 CA--- 6114 5764 true
- 5890 CA--- 6131 5773 true
- 5943 CA--- 6185 5811 true

- activ state sess conf nodes mount
- 5732 CA--- 6185 5811 96 /var/opt/cray/dws/mounts/batch/2128492/ss

- frag state inst capacity gran node
- 61382 CA-- 2234 212.91GiB 4MiB nid00457
- 73696 CA-- 2550 212.91GiB 4MiB nid02249
- 73697 CA-- 2550 212.91GiB 4MiB nid00205
- 73698 CA-- 2550 212.91GiB 4MiB nid01801
- 73699 CA-- 2550 212.91GiB 4MiB nid00014
- 73700 CA-- 2550 212.91GiB 4MiB nid01169
- 165142 CA-- 5487 425.81GiB 4MiB nid01418

- ns state conf frag span
- 49200 CA-- 2505 61382 1
- 52607 CA-- **2821** 73696 5
- 59484 CA-- 5764 165142 129
- **States**
 - **Goal: C – create or D – destroy**
 - **Setup: A – actualized or – non-actualized**
 - **Condition: F – fuse blown or – fuse intact**
 - **Status: T – transitioning or – stable or blocked**
 - **Spectrum: M- mixed or – not delayed**

scontrol show burst



Name=cray DefaultPool=wlm_pool Granularity=218016M TotalSpace=872936064M UsedSpace=234803232M
StageInTimeout=86400 StageOutTimeout=86400 Flags=EnablePersistent,TeardownFailure
GetSysState=/opt/cray/dw_wlm/default/bin/dw_wlm_cli

Allocated Buffers:

Name=u1_bb1 CreateTime=2016-03-02T15:01:01 Size=1090080M State=allocated UserID=user1(11111)
Name=u2_space CreateTime=2016-05-09T11:00:43 Size=1090080M State=allocated UserID=user2(22222)
Name=myBBname CreateTime=2016-02-19T13:45:33 Size=218016M State=allocated UserID=user3(33333)
Name=u4_Test2 CreateTime=2016-05-05T18:31:36 Size=654048M State=allocated UserID=user4(44444)
Name=u4_Test CreateTime=2016-05-05T16:01:02 Size=654048M State=allocated UserID=user4(44444)
Name=u5_30TB CreateTime=2016-05-05T14:31:08 Size=31612320M State=allocated UserID=user5(55555)

Per User Buffer Use:

UserID=user1(11111) Used=1090080M
UserID=user2(22222) Used=1090080M
UserID=user3(33333) Used=218016M
UserID=user4(44444) Used=1962144M
UserID=user5(55555) Used=31612320M

Job hung with processes in 'D' state



- Node stuck completing (most likely admin down if using Alps)
 - With SLURM log into the node to see what the problem is
 - Process hung in 'D' state on a DW instance
 - Get the job information and look at:
 - 'dwstat sessions' to find the session id
 - 'dwstat instance' to find the instance id
 - 'dwstat fragments' | grep <instance id>
 - Find the MDS node
 - Drain the node and reboot to clear the issue

- **Dwstat shows a ‘D’estroy indicator that doesn’t clear**
- **“scontrol show burst” (SLURM) where “allocation” size=0 or state=teardown.**
- **Once the DW-server is rebooted most recovery issues are handled by the DWS software without need for further intervention.**

Problem w/ size=0



- **Silent problem**
- **Registration stuck in 'D' state and either T or M**
- **Dwcli rm activation - wait to see if that clears the issue**
- **Dwcli update registration - -id <num> - -no-wait**
 - **Can cause data loss if all data isn't staged out**

- **SMW – log per dw node + log for sdb**
 - `/var/opt/cray/log/p0-current/dws`
 - `/var/opt/cray/log/p0-current/console & message`
 - **Grep dw and xfs to see information**
- **On mom/login nodes**
 - `/var/log/nginx`

- **To create or destroy a persistent instance a compute node must be allocated**
- **Existing issues**
 - **Symbolic links don't work**
 - **If there is an empty directory in the stage-in directory the stage-in will fail**
- **If max writes/day is reached the node will be set to read-only (ro)**
- **Check status of an SSD with xtcheckssd**

This work was supported by the
Director, Office of Science, Office of
Advanced Scientific Computing
Research of the U.S. Department of
Energy under contract No.
DEAC02-05CH11231.



National Energy Research Scientific Computing Center