Bursts of a Feather

David Paul, Andrey Kudryavtsev, John Bent

to share or not to share

a comparison of burst buffer architectures

John Bent, Seagate Government Solutions

Joint work with Bradley Settlemyer and Lei Cao, LANL

three places to add burst buffers

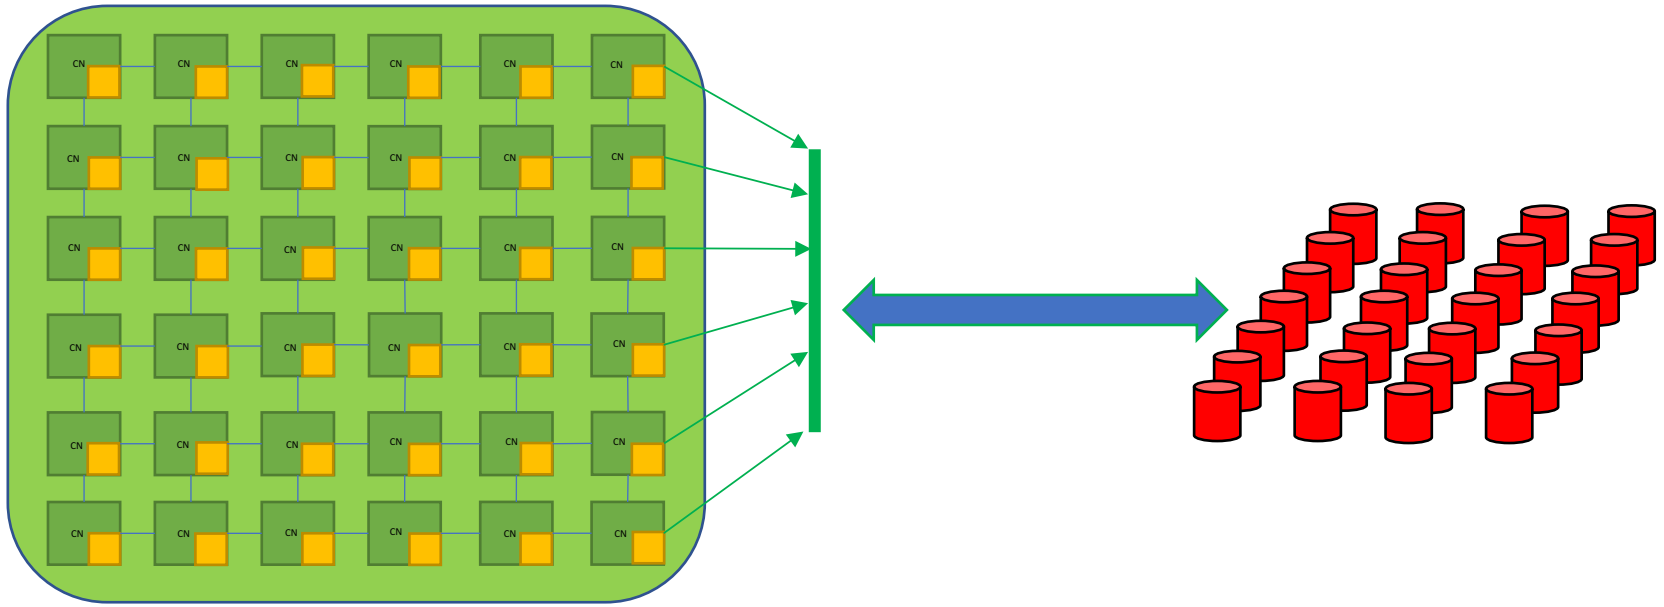

private, e.g. Cray/Intel Aurora @ Argonne

three places to add burst buffers

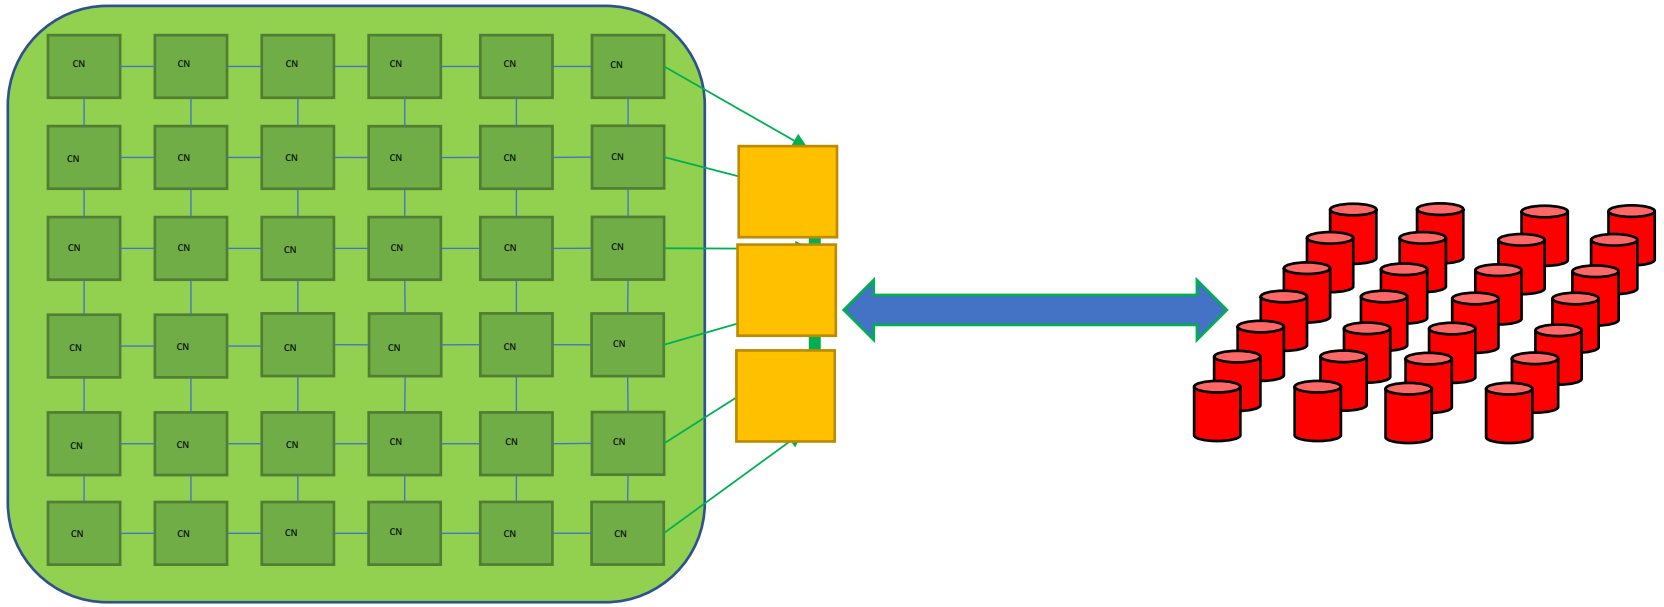

shared, e.g. Cray Trinity @ LANL

three places to add burst buffers

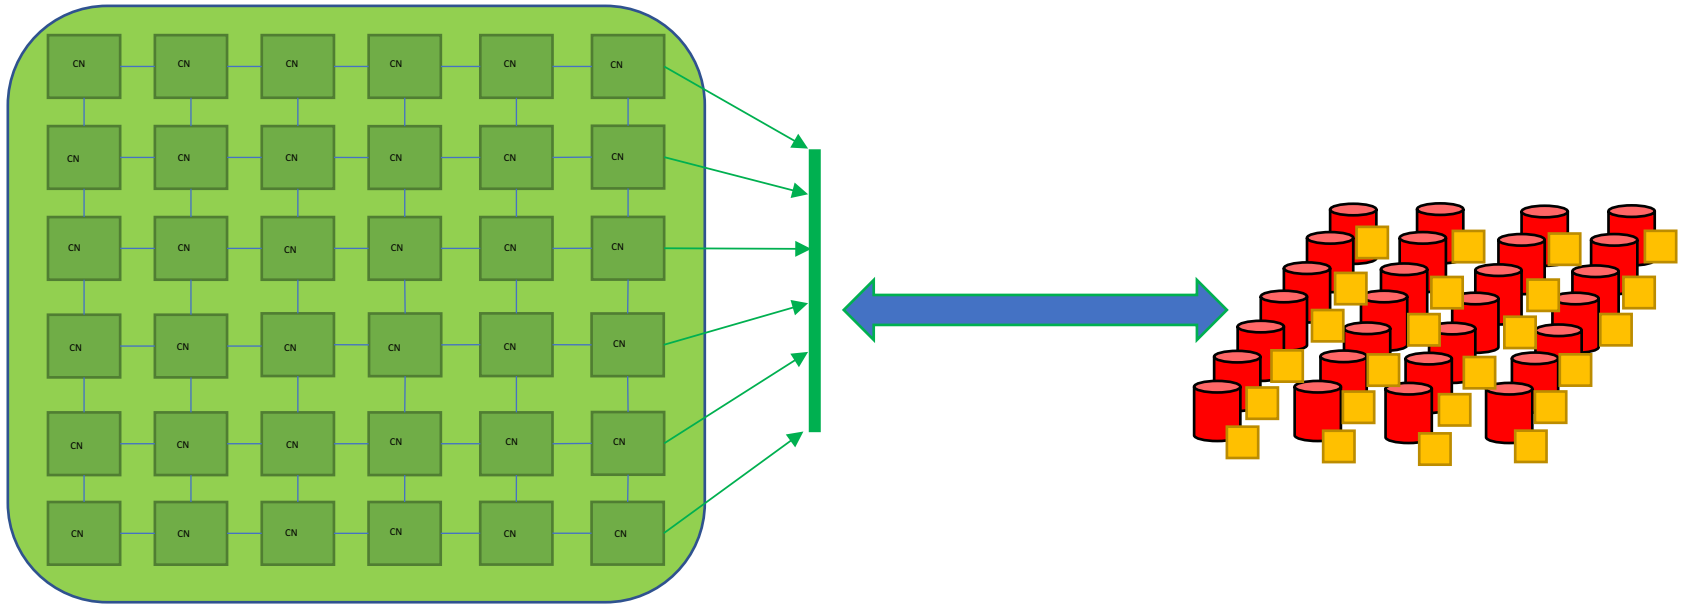

embedded, e.g. Seagate Nytro NXD

private

no contention
linear scaling
low cost
no network bandwidth

coupled failure domain
single shared file is difficult
small jobs cannot use them all

shared

n-1 easy
data can outlive job
temporary storage if pfs offline
small jobs can use it all
decoupled failure domain
most flexible ratio btwn compute, burst, pfs

most expensive
interference possible

embedded

n-1 easy
data outlives job
small jobs can use it all
decoupled failure domain from app
low cost
most transparent

SAN must be provisioned for burst
interference possible
most transparent

the value of decoupled failure domains

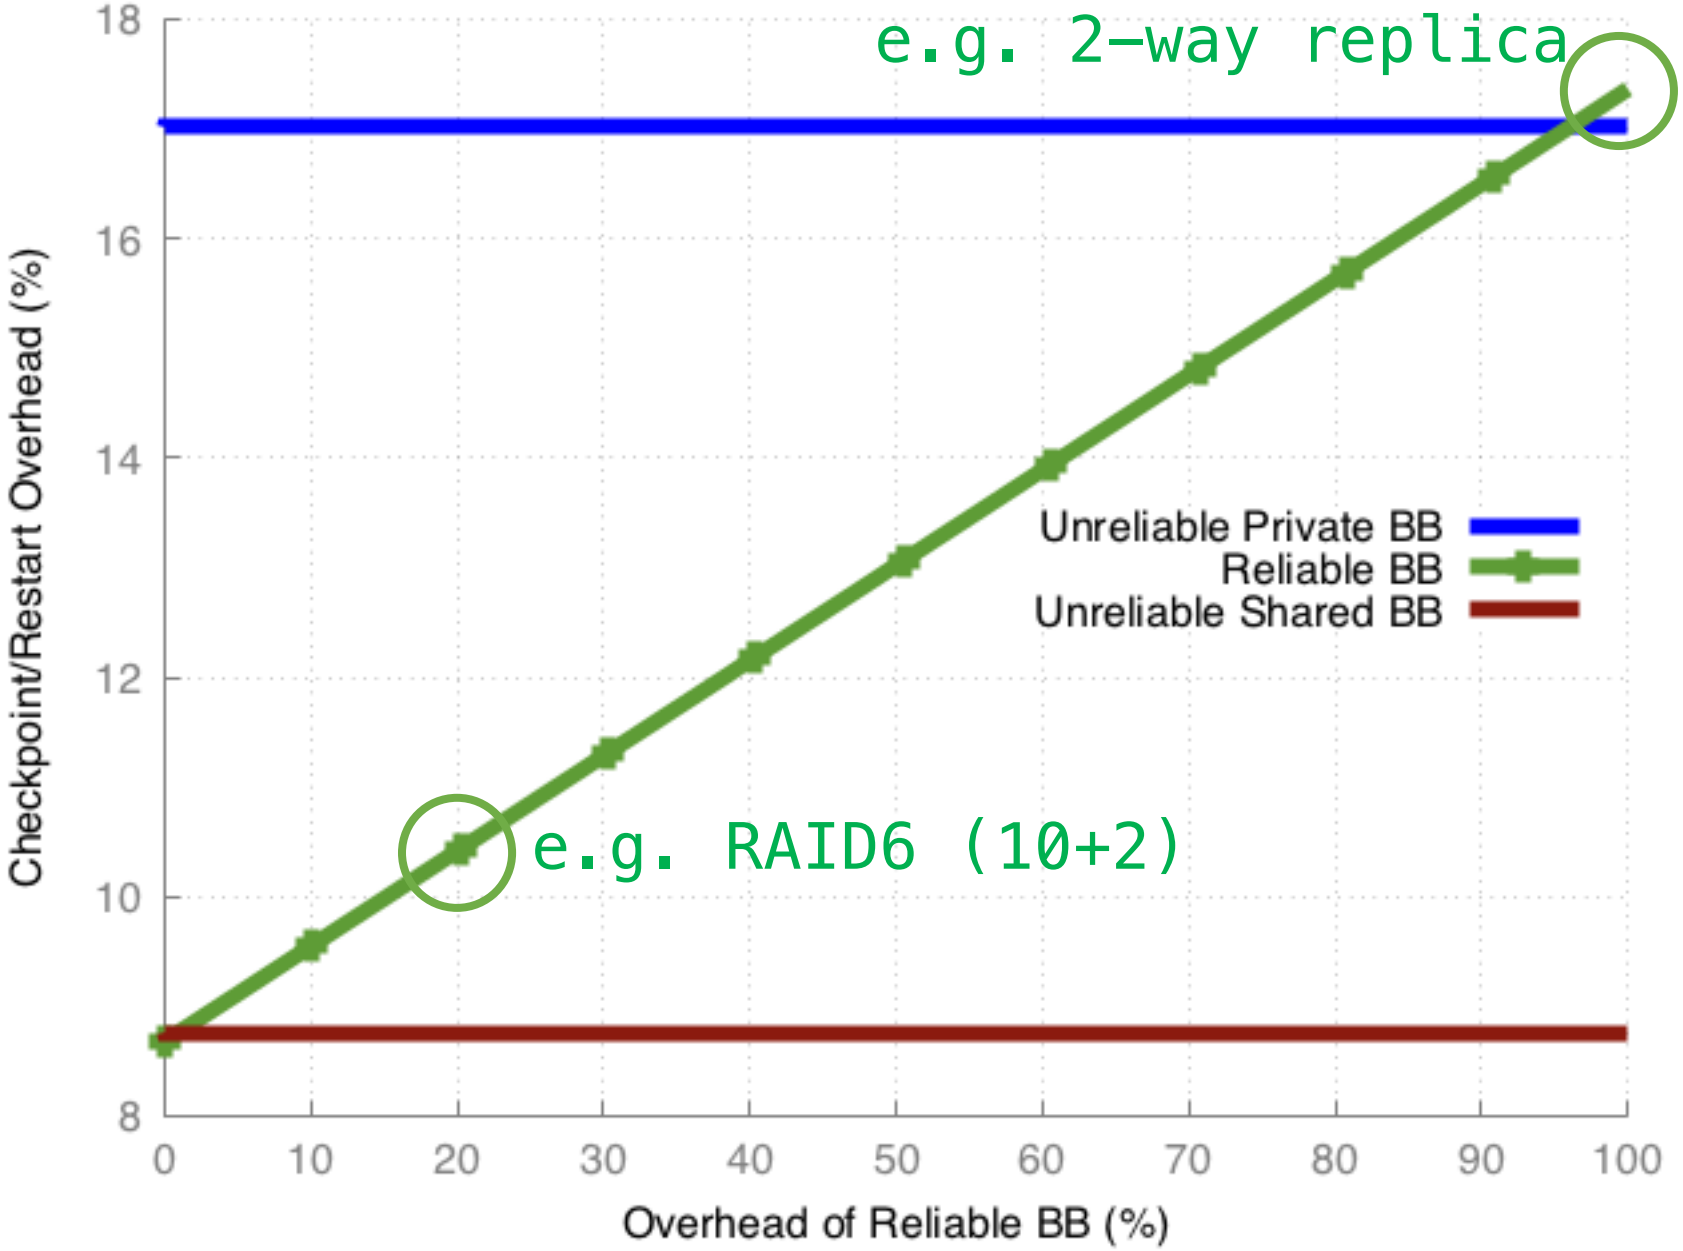

observations

shared doesn't need parity
private does
...
but then the
high private perf is lost

Bent, Settlemyer, et al.
On the non-suitability of
non-volatility.
HotStorage '15.

the value of shared for bandwidth

	Local Unreliable	Local 20% Parity	Shared Unreliable
Mean Ckpt Bw	206.8 GB/s		

simulation of APEX workflows running on Trinity

Lei Cao, Bradley Settlemyer, and John Bent. To share or not to share: Comparing burst buffer architectures. SpringSim 2017.

the value of shared for bandwidth

	Local Unreliable	Local 20% Parity	Shared Unreliable
Mean Ckpt Bw	206.8 GB/s	165.6 GB/s	

simulation of APEX workflows running on Trinity

Lei Cao, Bradley Settlemyer, and John Bent. To share or not to share: Comparing burst buffer architectures. SpringSim 2017.

the value of shared for bandwidth

	Local Unreliable	Local 20% Parity	Shared Unreliable
Mean Ckpt Bw	206.8 GB/s	165.6 GB/s	614.54 GB/s

simulation of APEX workflows running on Trinity

observation: capacity machines need shared burst buffers

Lei Cao, Bradley Settlemyer, and John Bent. To share or not to share: Comparing burst buffer architectures. SpringSim 2017.

to share or not to share

a comparison of burst buffer architectures

CUG BoF
Cray Datawarp
System Perspective

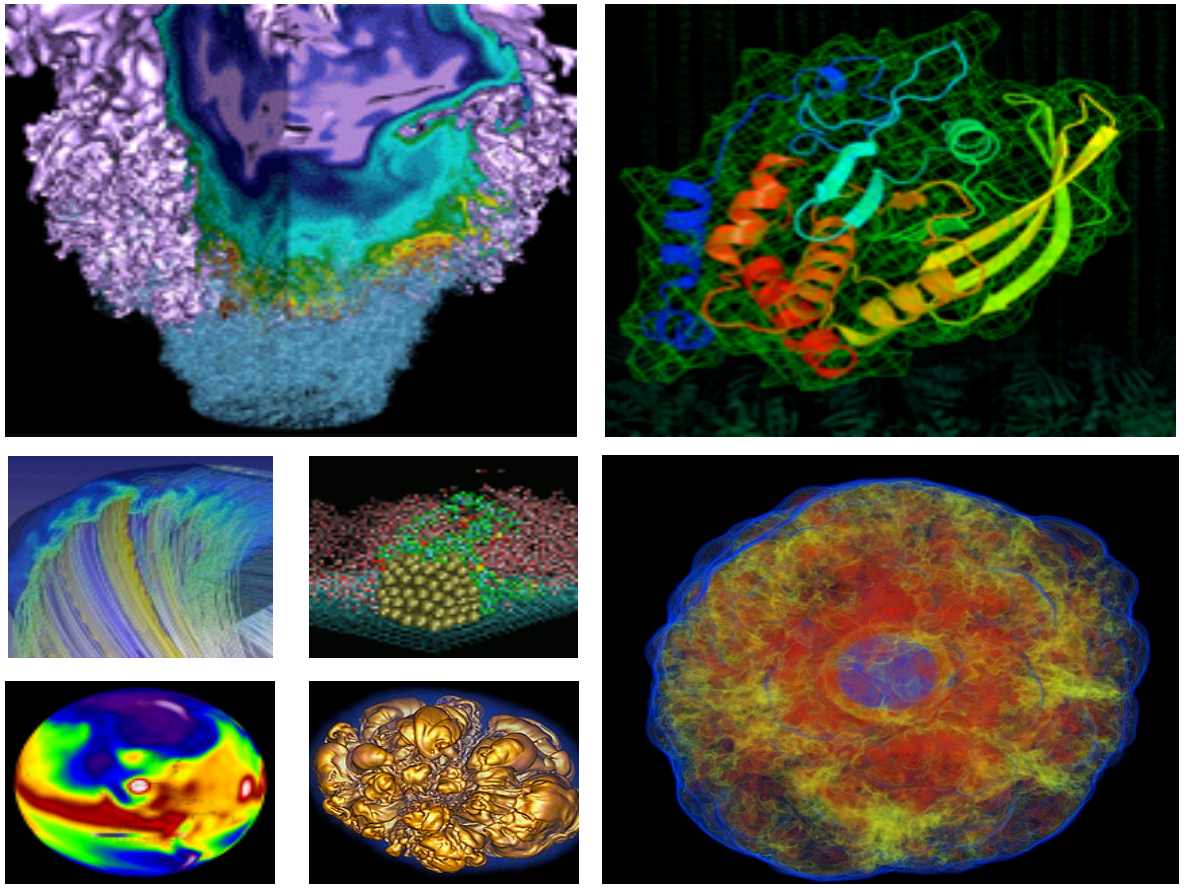

David Paul
Computational Systems Group
Lawrence Berkeley National Lab
DPAUL@LBL.GOV

May 11, 2017

Cori / NERSC-8 / XC40

- System specifics:
  - 9,688 KNL nodes
  - 2,004 Haswell nodes
  - 27PB Lustre Parallel Filesystem - $SCRATCH
  - Global GPFS - $HOMEs, $PROJECTs, S/W, Modules, etc.
    - Mounted on all NERSC systems
  - 288 Datawarp servers (576 Intel SSDs, two DW servers/blade)
  - Burst Buffer of 1.5PBs @ ~**1.6 TB/sec, 12.5M IOPS**

Datawarp Architecture Concept

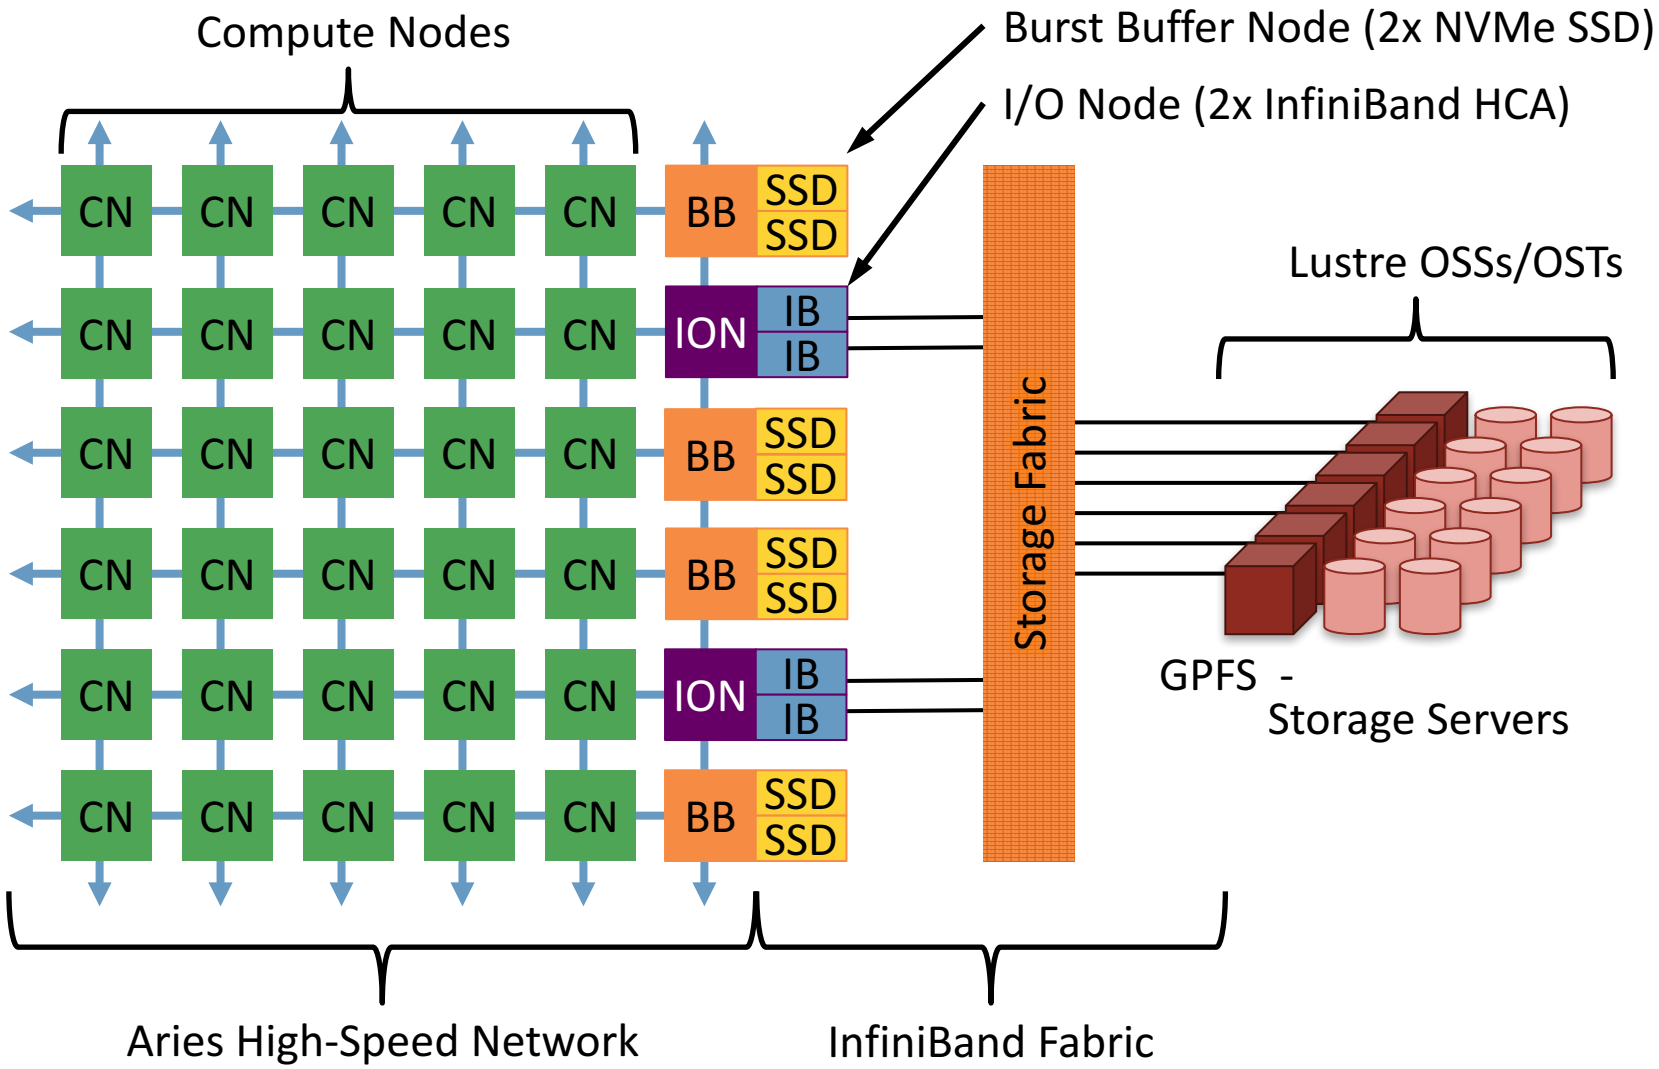

The 'moving pieces'…

WLM & DWS Interaction

- WLM (PBS, M/T, Slurm)
- Cray WLM Commands
- DataWarp Service
- DVS
- kdwfs
- XFS
- LVM
- xtnhd (scalable messaging)

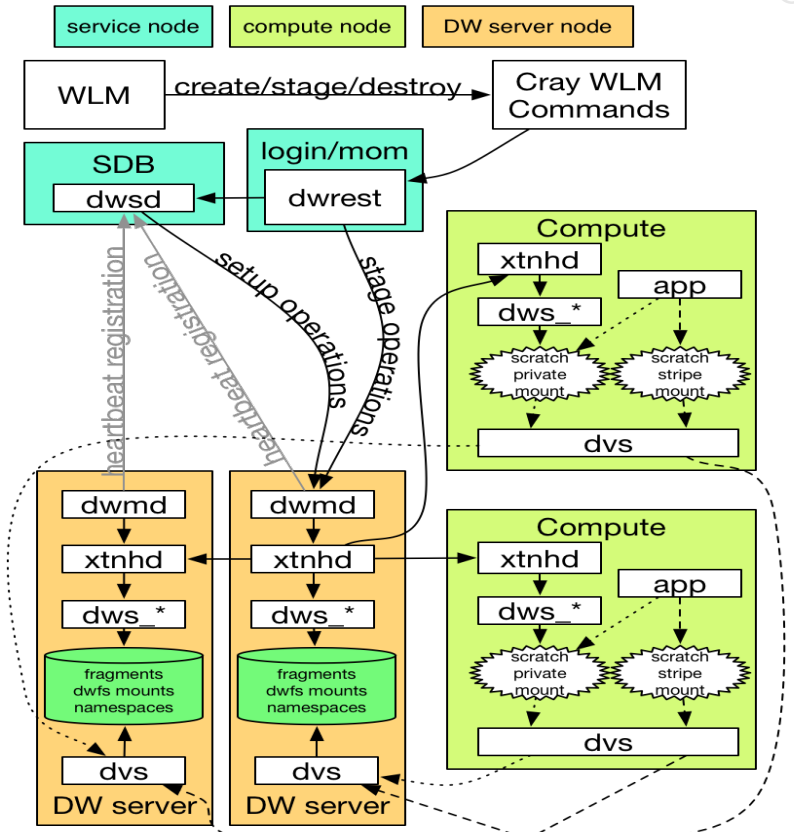

Datawarp Terms

- **DWS** DataWarp Service software for managing and configuring the SSD I/O installation
  - **Pool** subset of DW-servers with a common allocation granularity (ex. 200GB), ex. wlm_pool
  - **Session** typically created by a DWenabled WLM job (Token = JobID or Name)
  - **Instance** an object representing a user's request for disk space (ex. 600GB)
  - **Fragment** a piece of an instance, as it exists on a DW-server (ex. 3 @ 200GB)
  - **Configuration** an object representing how the space is to be used (scratch, striped, private, etc.)
  - **Namespace** represents the metadata (called tree) and data (called data), i.e. a Filesystem (DWFS)
  - **Registration** an object for linking together a session and a configuration
  - **Activation** an object representing where a configuration is to be used (i.e. mounted)
- **DWFS** DataWarp Filesystem a Cray Filesystem that supports staging of files and striping data across multiple DW-servers.
- **DVS** Data Virtualization Service Cray's I/O forwarding software (enables GPFS & DWFS I/O on HSN)
- **XFS** Filesystem used to persist data flowing through DWFS
- **DW-Server** – DVS server with SSDs, DWS, DWFS, XFS, LVM and access to a PFS

SLURM configuration for Datawarp (very simple)

- **slurm.conf : BurstBufferType=burst_buffer/cray**
- **burst_buffer.conf :**
  - **DefaultPool**: name of the pool used by default for resource allocations
    - **wlm_pool**
  - **AltPoolName**: allows for different storage configurations (ex. Granularity size)
  - **DenyUsers**: list of user names and/or IDs prevented from using burst buffers
  - **Flags EnablePersistent**: allows users to create/destroy persistent burst buffers
  - **Flags TeardownFailure**: remove DW allocation on job failure
- **QoS/TRES** – control user access, user quotas, usage and report them

SLURM configuration for Datawarp (very simple)

- **slurm.conf : BurstBufferType=burst_buffer/cray**
- **burst_buffer.conf :**
  - **DefaultPool**: name of the pool used by default for resource allocations
    - **wlm_pool**
  - **AltPoolName**: allows for different storage configurations (ex. Granularity size)
  - **DenyUsers**: list of user names and/or IDs prevented from using burst buffers
  - **Flags EnablePersistent**: allows users to create/destroy persistent burst buffers
  - **Flags TeardownFailure**: remove DW allocation on job failure
- **QoS/TRES** – control user access, user quotas, usage and report them

SLURM – (# scontrol show burst)

```
Name=cray DefaultPool=wlm_pool Granularity=82496M TotalSpace=1192325G FreeSpace=1078057728M UsedSpace=128858752M
  AltPoolName[0]=dev_pool Granularity=20624M TotalSpace=47693G FreeSpace=47693G UsedSpace=0
  AltPoolName[1]=sm_pool Granularity=20624M TotalSpace=476930G FreeSpace=476930G UsedSpace=0
  Flags=EnablePersistent,TeardownFailure
  StageInTimeout=86400 StageOutTimeout=86400 ValidateTimeout=5 OtherTimeout=300
  GetSysState=/opt/cray/dw_wlm/default/bin/dw_wlm_cli
  Allocated Buffers:
    JobID=4363770 CreateTime=2017-03-31T01:56:25 Pool=wlm_pool Size=1072448M State=staged-in UserID=userA(11002)
    Name=XT3 CreateTime=2017-02-27T11:57:37 Pool=(null) Size=329984M State=allocated UserID=userB(12023)
    JobID=4365735 CreateTime=2017-03-31T06:47:11 Pool=(null) Size=0 State=allocated UserID=userD(13734)
    Name=tiger CreateTime=2017-03-29T01:18:15 Pool=wlm_pool Size=10312G State=allocated UserID=userE(11642)
    Name=Mybb2 CreateTime=2017-03-24T20:14:11 Pool=(null) Size=114339456M State=allocated UserID=userC(19499)
  Per User Buffer Use:
    UserID=userA(11002) Used=3217344M
    UserID=userB(12023) Used=742464M
    UserID=userE(11642) Used=10312G
    UserID=userC(19499) Used=114339456M
```

DWS dwstat (administrator focused)

```
ctl1:~ # dwstat most
    pool units  quantity      free     gran
dev_pool bytes  46.58TiB  46.58TiB 20.14GiB
 sm_pool bytes 465.75TiB 465.75TiB 20.14GiB
wlm_pool bytes   1.14PiB      1PiB 80.56GiB

 sess state      token creator owner             created expiration nodes
12273 CA---     myBBT1     CLI 52023 2017-02-27T10:12:40      never     0
12290 CA---        XT2     CLI 52023 2017-02-27T11:57:37      never     0
13883 CA--- celestebb2     CLI 69499 2017-03-24T20:14:11      never     0

14123 CA---  octotiger     CLI 61642 2017-03-29T01:18:15      never     0

14259 CA---    4365615   SLURM 33734 2017-03-31T06:13:46      never  2048

inst state  sess     bytes nodes            created expiration intact      label  public confs
4239 CA--- 12273  80.56GiB     1 2017-02-27T10:12:40      never intact     myBBT1  public     1
4253 CA--- 12290 322.25GiB     4 2017-02-27T11:57:37      never intact        XT2  public     1
4627 CA--- 13883 109.04TiB   198 2017-03-24T20:14:12      never intact celestebb2  public     1

4709 CA--- 14123  10.07TiB   128 2017-03-29T01:18:15      never intact  octotiger  public     1

4744 CA--- 14253   1.02TiB    13 2017-03-31T01:56:22      never intact   I14253-0 private     1

conf state inst    type activs
4913 CA--- 4709 scratch      1

  reg state  sess conf wait
13968 CA--- 14259 4913 wait

activ state  sess conf nodes ccache                                                mount
13666 CA--- 14259 4913  2048     no /var/opt/cray/dws/mounts/batch/octotiger_striped_scratch
```

Problem Identification

- **Output from “squeue –l”**
- **SMW: console log**
- **SLURM log:**
  - slurmctld.log
- **Datawarp logs:**
  - Centrally to SMW with LLM consolidated by daemon name
    - /var/opt/cray/log/p#-<bootsession>/dws/
    - dwsd.yyyymmdd – scheduling daemon (typically on sdb node)
    - dwmd.yyyymmdd – DW-servers manager daemon
    - dwrest.yyyymmdd – dwgateway node(s)

Problem Identification

```
==> squeue -l | grep burst

  4705188     debug test-BB-   djbard  PENDING       0:00      1:00      1 (burst_buffer/cray: _create_persistent:
                                                                           Access mode private not supported for
                                                                           persistent instances)
```

```
corismw1:/var/opt/cray/log/p0-current/dws # grep kdwfs ../console-20170424

 c4-0c1s2n2 xtconsole 33429 p0-20170420t160534 [console@34] kdwfs: KDWFS protection limit(s) exceeded!
 c4-0c1s2n2 xtconsole 33429 p0-20170420t160534 [console@34] kdwfs: KDWFS protection limit(s) exceeded!
```

```
<slurm.log>

  _start_stage_in: dws_data_in for job 4650696 ran for usec=601646
  dw_wlm_cli --function data_in --token 4650696 --job /global/syscom/cori/var/cori-slurm-state/hash.6/job.4650696/script
  Error staging in session 14527 configuration 5059 path /global/cscratch1/sd/djbard/model0112.tar -> /: dwput error 405:
  METHOD NOT ALLOWED for url: https://c6-0c1s4n3:81/dw/v1/stages/14527/5059/files/ messages:
```

BB Areas of Interest - Feedback

- **Bare Metal Mode**
  - **“Give me #-of-BB-Servers and get out of my way...”**
- **Compute Application on BB-Server**
  - **“Give me #-of-BB-Servers and #-of-Compute/Threads”**
- **SWAP Mode**
  - DWS supports this now – use cases?
- **Other Use Cases:**
  - BoF feedback
    1. _
    2. _
    3. _
    4. _