



Thinking about an
annual release DRP
(like DR2) from a
bird's eye view.



U.S. DEPARTMENT OF
ENERGY

Bird's eye DRP run metric estimates:

Previous DRP campaign metrics with similar stack pipelines:

DRP	RTN	Year Mos	visits	det/visit	tracts	steps/stages	core-hours	storage	days
DC2		2 days	300	150	2	steps 1-7	5K	10TB	2d
RC2		3 days	423	100	3	steps 1-7	10K	25TB	3d
DP1		3 days	1800	9	22	steps 1-7	10K	25TB	3d
DP0.2	rtn-039	2022 1-6	20K	150	150	steps 1-7	2.5M	2.5PB	180d
PDR2	rtn-063	2023 5-9	14K	100	400	steps 1-7	1.5M	2.5PB	120d
(projection)									
DR 1		2026 1-9	240K	189	8000	stages 1-4	40M	50 PB	150d
DR 2		2027 1-9	500K	189	10000	stages 1-4	80M	100 PB	300d

Rough numbers in an early DRP (DR1-DR2)

800 exposures/night, 300 nights/year = **240,000 visits/year**. Each visit has ~**200** (189) detectors (**50M raws/year**). Sky is divided into **10K tracts** of 2.5 sq degrees/tract (25,000 sq degree LSST footprint) Early DR depth will scan the sky **N=10** times/year.

Stage 1 (ISR, PSF) runs on $200 \times 1M/4 = 50M$ quanta wallclock time is 2 min/quanta → 100M minutes = **2M core-hours**. If we have **3 DFs x5000** cores, this would take running full out at 100% efficiency: **~6 days**.

Stage 2 (photometric calibrate (fgcm), astrometric calibration (gbdes), finalCharacterization (time chewing step): Currently takes almost as half as long to run stage 2 as running stage 3 (may not be able to use all 15K cores, say 10K): **8M core-hours: ~35 days**.

Rough numbers in a full year DRP (early years)

800 exposures/night, 300 nights/year = **240,000 visits/year**. Each visit has ~**200** (189) detectors (**50M raws/year**). Sky is divided into **10K tracts** of 2.5 sq degrees/tract (25,000 sq degree LSST footprint) Early DR depth will scan the sky **N=10** times/year.

Stage 3 (Coadds) generated for 10K tracts of 100 patches each (**1M patches**) wallclock time is 8x that of step 1, or perhaps **16M core-hours** (with 10K cores): **~66 days** .

Stage 4 (DIA, light curves) runs on patches x visits (1M x 10 visits/patch = 10M quanta) Takes a similar amount of time to Stage 3 (currently **16 M core-hours**, all 15K cores): **45 days**.

Total 4 stages DR2 DRP: 6 + 35 + 66 + 45 = 152 days.

Requirement: Must run in 200 days. Risks: cores-filling, data movement, restarts.

Bird's eye DRP run metric estimates:

Previous DRP campaign metrics with similar stack pipelines:

DRP	RTN	Year Mos	visits	det/visit	tracts	steps/stages	core-hours	storage	days
DC2		2 days	300	150	2	steps 1-7	5K	10TB	2d
RC2		3 days	423	100	3	steps 1-7	10K	25TB	3d
DP1		3 days	1800	9	22	steps 1-7	10K	25TB	3d
DP0.2	rtn-039	2022 1-6	20K	150	150	steps 1-7	2.5M	2.5PB	180d
PDR2	rtn-063	2023 5-9	14K	100	400	steps 1-7	1.5M	2.5PB	120d
(projection)									
DR 1		2026 1-9	240K	189	8000	stages 1-4	40M	50 PB	150d
DR 2		2027 1-9	500K	189	10000	stages 1-4	80M	100 PB	300d

A DRP ‘grouping’ example

DRP	RTN	Year	Mos	visits	det/visit	tracts	steps/stages	core-hours	storage	days
DR 1		2026	1-9	240K	189	8000	stages 1-4	40M	50 PB	150d

Stage 1: 50M quanta 2 min/job: 8M x 6 (1,1,3,1) at (RAL, LANCS, FR, US) reasonable max
Group size might be 800K jobs/group → 10-30 groups for CM to manage/site. **6 days**

Stage 2: 1M quanta 8 hours/job: group size might be 500K jobs/group → 2-6 groups at each
DF: **35 days**

Stage 3: 1M quanta 20 hours/job: group size might be 20 tracts/group (2K quanta/group)
need 500 groups (100,100,300,100) at (RAL,LANCS,FR,US) managed: **66 days**

Stage 4: 10M quanta 1 hr/job, group size similar to step1 (800K jobs/group), 2-6
groups/DF: **45 days**

2x3 multisite step-by-step

Step1

Step2a

Step2b

fan-in-step2c

Step2c

fan-out-step2c

Step2d

Step2e

Step3

fan-in-step7

Step7

Step4

Step5

Step6

Make-Collections-of-everything

Generate-Resource-Usage

fan-in-plotsoutputs

Make-collection-of-everything-at-USDF

Load Plot Navigator

Move-ahead-of-step2c (can be called via finalJob at end of step2a)

```
move-from-ukdf-step2cin.sh
```

```
export CONFIG_FILE=rucio_register.cfg
```

```
export BUTLER_REPO=hsc_pdr2_multisite
```

Needs:

```
rucio-register \
```

```
-r \${BUTLER_REPO} \
```

```
-c HSC/runs/PDR2/VVDS/w_2025_04/DM-48570/step2a_output/ukdf \
```

```
-t visitSummary \
```

```
-d Dataset/HSC/runs/PDR2/VVDS/w_2025_04/DM-48570/step2a_output/ukdf/2cinputs \
```

```
-C \${CONFIG_FILE} \
```

```
-s 30
```


2x3 multisite step-by-step (Rucio move sample)

RUCIO add-rule to launch data copy:

```
rucio add-rule  
hsc_pdr2_multisite:Dataset/HSC/runs/PDR2/VVDS/w_2025_04/DM-48570/ukdf/2cinputs  
2 LANCS_BUTLER_DISK\|SLAC_BUTLER_DISK
```

Check if dataset is done moving (try this in a loop):

```
rucio list-dataset-replicas  
hsc_pdr2_multisite:Dataset/HSC/runs/PDR2/VVDS/w_2025_04/DM-  
48570/ukdf/2cinputs
```