

RESOURCES  
SOFTWARE AND TECHNOLOGIES  
STATUS OF DATA IN DB  
ACTIONS AND RECOMMENDATIONS

# OPERA DATABASE EXPORT AND PRESERVATION

# RESOURCES

Physical machine received at CCIN2P3 with 2 TB disk

- Dedicated to OPERA for 2016/2017 for Oracle DB dump
- Not part of Linux batch clusters

Account activated at CERN for data preservation

- Access to EOS (disk) and CASTOR (tape) granted
- Some discussion ongoing for best data transmission technology
  - `scp`, `rsync` easy to use, require some scripting to assess transfer is OK
  - `xrdcp` – not sure if it can be used on the Lyon machine
  - WLCG (GRID-based) – would need EGI authentication

# SOFTWARE AND TECHNOLOGIES

OracleDumpManager (.NET/Mono exe)

Set of bash scripts

Full chain activated with a single command or event list

Input format:

| EVENT       | BRICK   | EVTTYPE | PRIORITY | ID_FEEDBACK      | PRONGS |
|-------------|---------|---------|----------|------------------|--------|
| 10240014359 | 1123062 | NC      | 1        | 1000010011679071 | 5      |
| 10310004085 | 1050094 | CC      | 1        | 1000010013001023 | 1      |
| 10209046135 | 1134018 | NC      | 1        | 1000010012731215 | 1      |
| 10270007687 | 1046052 | CC      | 1        | 1000010013000435 | 2      |
| 10233027637 | 1138277 | CC      | 1        | 1000010013003253 | 1      |
| 10317043625 | 1120844 | CC      | 1        | 1000010012735188 | 2      |
| 10232004243 | 1054025 | CC      | 1        | 1000010013001144 | 6      |
| 10316032251 | 1152030 | CC      | 1        | 1000010013022657 | 5      |
| 10284013149 | 1028243 | CC      | 1        | 1000010015819976 | 2      |
| 10180044583 | 1063291 | CC      | 1        | 1000010012110182 | 3      |
| 10144050777 | 1107570 | CC      | 1        | 1000010004782325 | 8      |
| 10120018681 | 1009831 | CC      | 1        | 1000010006911467 | 3      |
| 10133066579 | 1138613 | NC      | 1        | 1000010012731276 | 7      |

# SOFTWARE AND TECHNOLOGIES

Each event is completely contained in a single directory

Create dump directory exp\_evXXXXXXXXXX\_bkYYYYYYYYYYY

Extract event-related electronic data

Extract brick-related data (ALL – will cause duplication for multi-events)

Extract feedback views for LAST feedback for that event

Convert to ASCII (in parallel with extraction of next event – generate directory with name ascii\_exp\_evXXXXXXXXXXXX\_bkYYYYYYYYYYY)

Copy over Internet (after ASCII, in parallel with extraction of next event)

All actions are logged and extraction log is saved  
Single corrupt rows or fields are documented and skipped

# SOFTWARE AND TECHNOLOGIES

## Extraction test facts & figures

Binary file size: 69 GB

ASCII file size: 96 GB

Extraction time: 700 min

ASCII Conversion time: 1000 min

Network transfer time: computed from transfer speed – 2283s at 0.56 Gbps  
(measured)

Approximately **553** days to copy the whole DB

Parallel extraction will be needed

# STATUS OF DATA IN DB

The following slides contain queries that have been generated on purpose to run on the **opera** account of the Central DB

Local Database administrators are **encouraged** to use these queries to recover details on a brick-by-brick and event-by-event basis

The materialized views named

DBEX\_.....

are recomputed **daily** and available for you to check the status of your laboratory, accessible from your local **operapub** account, e.g.:

```
select * from opera.dbex_invalid_datasets@opfra where laboratory = 'SALERNO'
```

# STATUS OF DATA IN DB

## CS publication statistics

### DBEX\_CS\_PUBLICATION\_STATUS

| RUNYEAR | LAB    | ASSIGNED | FLAGSRECEIVED | DATARECEIVED |
|---------|--------|----------|---------------|--------------|
| 2007    | LNGS   | 1        | 1             | 0            |
| 2007    | Nagoya | 4        | 3             | 0            |
| 2008    | LNGS   | 1406     | 1359          | 805          |
| 2008    | Nagoya | 1330     | 1301          | 654          |
| 2009    | LNGS   | 2794     | 2735          | 1563         |
| 2009    | Nagoya | 2962     | 2869          | 1436         |
| 2010    | LNGS   | 3069     | 2894          | 1553         |
| 2010    | Nagoya | 2987     | 2887          | 1544         |
| 2011    | LNGS   | 3561     | 3336          | 1297         |
| 2011    | Nagoya | 3303     | 3164          | 1603         |
| 2012    | LNGS   | 2513     | 2373          | 986          |
| 2012    | Nagoya | 2447     | 2345          | 954          |

# STATUS OF DATA IN DB

## CS results published, missing candidates

### DBEX\_CS\_RESULTS\_MISSCANDS

```
select id_cs_eventbrick, result_status, sum(nvl2(id, 1, 0)) as ncands,
tb_brick_assign.cs_assign from
(select id_cs_eventbrick, result_status, id from tb_cs_results
left join tb_cs_candidates on id_eventbrick = id_cs_eventbrick)
inner join tb_brick_assign on tb_brick_assign.id_eventbrick = id_cs_eventbrick
group by id_cs_eventbrick, result_status, cs_assign)
```

| RESULT_STATUS             | CS_ASSIGN | CSDOUBLETS | CSWITHCANDS |
|---------------------------|-----------|------------|-------------|
| B2B_FASTUNPACK            | JP        | 8          | 0           |
| B2B_FASTUNPACK            | UE        | 8          | 0           |
| B2B_NOSCAN                | UE        | 2          | 0           |
| BACK_TO_DETECTOR          | JP        | 2558       | 28          |
| BACK_TO_DETECTOR          | UE        | 3721       | 112         |
| BACK_TO_DETECTOR_NO_CS    | JP        | 2525       | 1           |
| BACK_TO_DETECTOR_NO_CS    | UE        | 2445       | 25          |
| BLACK_CS_DEVELOP          | JP        | 159        | 10          |
| BLACK_CS_DEVELOP          | UE        | 351        | 0           |
| CS_CAND_OK_DEVELOP        | JP        | 6632       | 5951        |
| CS_CAND_OK_DEVELOP        | UE        | 5745       | 5699        |
| CS_CAND_OK_FAST_UNPACK    | JP        | 315        | 67          |
| CS_CAND_OK_FAST_UNPACK    | UE        | 43         | 40          |
| NO_COSMIC_RAYS_DEVELOP    | JP        | 414        | 141         |
| NO_COSMIC_RAYS_DEVELOP    | UE        | 454        | 414         |
| WRONG_CS_HANDLING_DEVELOP | UE        | 33         | 11          |

JP: many CS without candidates

UE: in some cases another CS was flagged with results to trigger re-extraction

# STATUS OF DATA IN DB

## Brick data consistency checks

**Question: do we have scanning/feedback data for events that are at least connected?**

Main query

```
select laboratory, event, brick, result_status, idcs, located, deadmaterial, passing, ds_any_done,
decode(sum(nvl2(sb1.id, 1, 0)),0,0,1) as issb, decode(sum(nvl2(vo2.id, 1, 0)),0,0,1) as isvol,
decode(sum(nvl2(re3.id, 1, 0)),0,0,1) as isrec from
(select laboratory, event, brick, result_status, idcs, located, deadmaterial, passing, ds_any_done, id as idop from
(select laboratory, event, brick, result_status, idcs, located, deadmaterial, passing, ds_any_done from
xv_event_location_detail where started > 0 and connected > 0)
 left join tb_proc_operations on id_parent_operation is null and id_eventbrick = brick
)
left join tb_scanback_paths sb1 on sb1.id_processoperation = idop
left join tb_volumes vo2 on vo2.id_processoperation = idop
left join tb_reconstructions re3 on re3.id_processoperation = idop
group by laboratory, event, brick, result_status, idcs, located, deadmaterial, passing, ds_any_done
```

# STATUS OF DATA IN DB

## Located events, no volume data available

**DBEX\_LOCATED\_NOTSDATA** (located > 0 and deadmaterial = 0 and isvol = 0)

| LABORATORY     | BRICKS |
|----------------|--------|
| ANKARA         | 21     |
| BARI           | 59     |
| BERN           | 952    |
| BOLOGNA-PADOVA | 127    |
| DUBNA          | 32     |
| FRASCATI       | 21     |
| LEBEDEV        | 27     |
| NAGOYA         | 359    |
| NAPOLI         | 523    |
| ROMA           | 12     |
| SALERNO        | 67     |
| SINP-MSU       | 8      |

Some «located» events are not «decaysearched» hence not published yet, but that's not all

# STATUS OF DATA IN DB

## Decaysearched events, no volume data available

DBEX\_DECAYSEARCH\_NOTSDATA (located > 0 and ds\_any\_done > 0 and isvol = 0)

| LABORATORY     | BRICKS |
|----------------|--------|
| ANKARA         | 13     |
| BARI           | 28     |
| BERN           | 890    |
| BOLOGNA-PADOVA | 92     |
| DUBNA          | 28     |
| FRASCATI       | 20     |
| LEBEDEV        | 20     |
| NAGOYA         | 127    |
| NAPOLI         | 480    |
| ROMA           | 7      |
| SALERNO        | 21     |
| SINP-MSU       | 3      |

Publication not done or incomplete transfer

# STATUS OF DATA IN DB

## Decaysearched events, no feedback

DBEX\_DECAYSEARCH\_NOFEEDBACK (located > 0 and ds\_any\_done > 0 and isrec = 0)

| LABORATORY     | BRICKS |
|----------------|--------|
| ANKARA         | 5      |
| BARI           | 8      |
| BERN           | 428    |
| BOLOGNA-PADOVA | 30     |
| DUBNA          | 3      |
| FRASCATI       | 9      |
| LEBEDEV        | 4      |
| NAGOYA         | 121    |
| NAPOLI         | 24     |
| SALERNO        | 3      |

Publication not done or incomplete transfer

# STATUS OF DATA IN DB

## Bricks with no data at all

**DBEX\_BRICKS\_NODATA** (issb = 0 and isvol = 0 and isrec = 0)

| LABORATORY     | BRICKS |
|----------------|--------|
| ANKARA         | 14     |
| BARI           | 108    |
| BERN           | 703    |
| BOLOGNA-PADOVA | 148    |
| DUBNA          | 19     |
| FRASCATI       | 18     |
| LEBEDEV        | 16     |
| NAGOYA         | 1659   |
| NAPOLI         | 232    |
| ROMA           | 18     |
| SALERNO        | 147    |
| SINP-MSU       | 22     |

Data not strictly needed but we should have a policy defined

# STATUS OF DATA IN DB

## Decaysearched events, no data

**DBEX\_DECAYSEARCH\_NODATA** (`ds_any_done > 0 and issb = 0 and isvol = 0 and isrec = 0`)

| LABORATORY     | BRICKS |
|----------------|--------|
| ANKARA         | 5      |
| BARI           | 8      |
| BERN           | 361    |
| BOLOGNA-PADOVA | 8      |
| DUBNA          | 3      |
| FRASCATI       | 6      |
| LEBEDEV        | 4      |
| NAGOYA         | 112    |
| NAPOLI         | 23     |
| SALERNO        | 2      |

Publication not done or incomplete transfer

# STATUS OF DATA IN DB

## Passing-through/edgeout bricks, no scanning data

**DBEX\_PASSING\_NOSCANDATA** (passing > 0 and issb = 0 and isvol = 0)

| LABORATORY     | BRICKS |
|----------------|--------|
| ANKARA         | 11     |
| BARI           | 5      |
| BOLOGNA-PADOVA | 8      |
| DUBNA          | 6      |
| FRASCATI       | 3      |
| LEBEDEV        | 7      |
| NAGOYA         | 174    |
| NAPOLI         | 45     |
| ROMA           | 5      |
| SINP-MSU       | 6      |

Publication not done or incomplete transfer

# STATUS OF DATA IN DB

## Passing-through/edgeout bricks, no feedback data

DBEX\_PASSING\_NOFEEDBACK (passing > 0 and isrec = 0)

| LABORATORY     | BRICKS |
|----------------|--------|
| BOLOGNA-PADOVA | 5      |
| DUBNA          | 2      |
| LEBEDEV        | 2      |
| NAGOYA         | 174    |
| NAPOLI         | 14     |
| ROMA           | 4      |
| SINP-MSU       | 9      |

Publication not done or incomplete transfer

# STATUS OF DATA IN DB

## Apparently export-ready events

**DBEX\_EXPORT\_READY\_CHECK1** (`ds_any_done > 0 and isvol > 0 and isrec > 0`)

| LABORATORY     | BRICKS |
|----------------|--------|
| ANKARA         | 22     |
| BARI           | 289    |
| BERN           | 311    |
| BOLOGNA-PADOVA | 317    |
| FRASCATI       | 42     |
| LEBEDEV        | 6      |
| NAGOYA         | 2944   |
| NAPOLI         | 196    |
| ROMA           | 62     |
| SALERNO        | 478    |
| SINP-MSU       | 15     |

This check alone is not sufficient, see more in the next slides

# STATUS OF DATA IN DB

## Probably export-ready events

### DBEX\_EXPORT\_READY\_CHECK2

(DS primary vertex location is contained in a TS volume at least 5×5 mm<sup>2</sup>)

| LABORATORY     | DECAYSEARCHED | TSAVAILABLE |
|----------------|---------------|-------------|
| ANKARA         | 35            | 15          |
| BARI           | 317           | 278         |
| BERN           | 1262          | 108         |
| BOLOGNA-PADOVA | 428           | 279         |
| DUBNA          | 28            | 0           |
| FRASCATI       | 64            | 16          |
| LEBEDEV        | 26            | 0           |
| NAGOYA         | 3080          | 2887        |
| NAPOLI         | 677           | 72          |
| ROMA           | 69            | 59          |
| SALERNO        | 500           | 463         |
| SINP-MSU       | 18            | 5           |

```
select laboratory, ev, brick, result_status, idcs, 1 as ds_any_done, decode(sum(case when maxz is null then 0 when maxz > posz and minz < posz then 1 else 0 end),0,0,1) as hasvol from
(select laboratory, ev, brick, result_status, idcs, lastidr, id_vertex, posx, posy, posz, id_volume, max(Z) as maxz, min(Z) as minz from
(select laboratory, ev, brick, result_status, idcs, idr as lastidr from
(select laboratory, event+0 ev, brick, result_status, idcs, located, deadmaterial, passing, ds_any_done from xv_event_location_detail where ds_any_done > 0)
left join (select id_eventbrick, event+0 as event, max(id_reconstruction) as idr from vw_feedback_reconstructions group by id_eventbrick, event) on id_eventbrick = ev)
left join vw_feedback_vertices vx on vx.id_eventbrick = brick and id_reconstruction = lastidr and isprimary = 'Y'
left join tb_volume_slices sl on sl.id_eventbrick = brick and minx < posx and maxx > posx and miny < posy and maxy > posy and (maxx - minx) * (maxy - miny) > 25e6
left join tb_plates pl on pl.id_eventbrick = brick and pl.id = id_plate
group by laboratory, ev, brick, result_status, idcs, lastidr, id_vertex, posx, posy, posz, id_volume
) group by laboratory, ev, brick, result_status, idcs
```

Also this check is not sufficient, see more in the next slides

# STATUS OF DATA IN DB

Special events (by DB tags)

2 τ, 24 charm

|         |
|---------|
| 1041378 |
| 1046908 |

|         |
|---------|
| 1014653 |
| 1020518 |
| 1027222 |
| 1034730 |
| 1035653 |
| 1037545 |
| 1044354 |
| 1048057 |
| 1059278 |
| 1065097 |
| 1066404 |
| 1073614 |
| 1077152 |
| 1078815 |
| 1079117 |
| 1082561 |
| 1085405 |
| 1107689 |
| 1110205 |
| 1118858 |
| 1127653 |
| 1140875 |
| 1142664 |
| 1150444 |

??

# STATUS OF DATA IN DB

$\nu_e$  statistics

| LABORATORY     | RUNYEAR | TRIGGERED | STARTED | SCANNED | ANALYSED | DATASETS |
|----------------|---------|-----------|---------|---------|----------|----------|
| BERN           | 2008    | 1         | 1       | 1       | 1        | 0        |
| BOLOGNA-PADOVA | 2008    | 3         | 3       | 3       | 3        | 2        |
| NAGOYA         | 2008    | 6         | 0       | 0       | 0        | 0        |
| NAPOLI         | 2008    | 3         | 3       | 3       | 3        | 0        |
| BARI           | 2009    | 3         | 2       | 2       | 2        | 0        |
| BERN           | 2009    | 5         | 5       | 5       | 5        | 0        |
| BOLOGNA-PADOVA | 2009    | 2         | 2       | 2       | 2        | 2        |
| NAGOYA         | 2009    | 25        | 0       | 0       | 0        | 0        |
| NAPOLI         | 2009    | 6         | 5       | 5       | 5        | 0        |
| BARI           | 2010    | 2         | 2       | 2       | 2        | 0        |
| BERN           | 2010    | 6         | 6       | 6       | 6        | 0        |
| BOLOGNA-PADOVA | 2010    | 3         | 3       | 3       | 3        | 3        |
| DUBNA          | 2010    | 1         | 1       | 1       | 1        | 0        |
| FRASCATI       | 2010    | 1         | 1       | 1       | 1        | 0        |
| NAGOYA         | 2010    | 23        | 18      | 18      | 18       | 0        |
| NAPOLI         | 2010    | 4         | 4       | 4       | 4        | 0        |
| ROMA           | 2010    | 1         | 0       | 0       | 0        | 0        |
| SALERNO        | 2010    | 4         | 4       | 4       | 4        | 1        |
| BARI           | 2011    | 3         | 3       | 3       | 3        | 0        |
| BERN           | 2011    | 9         | 9       | 9       | 9        | 0        |
| BOLOGNA-PADOVA | 2011    | 2         | 2       | 2       | 2        | 2        |
| FRASCATI       | 2011    | 6         | 6       | 6       | 6        | 0        |
| NAGOYA         | 2011    | 19        | 15      | 15      | 15       | 0        |
| NAPOLI         | 2011    | 6         | 6       | 6       | 6        | 0        |
| SALERNO        | 2011    | 1         | 1       | 1       | 1        | 0        |
| BARI           | 2012    | 1         | 1       | 1       | 1        | 0        |
| BERN           | 2012    | 11        | 11      | 11      | 11       | 0        |
| BOLOGNA-PADOVA | 2012    | 4         | 4       | 4       | 4        | 4        |
| DUBNA          | 2012    | 1         | 1       | 1       | 1        | 0        |
| NAGOYA         | 2012    | 11        | 5       | 5       | 5        | 0        |
| NAPOLI         | 2012    | 9         | 9       | 9       | 9        | 0        |

133 completed events, not all  $\nu_e$ 's – missing datasets

# STATUS OF DATA IN DB

## Incomplete data transfers

### DBEX\_INVALID\_DATASETS

| LABORATORY     | DATASETS |
|----------------|----------|
| ANKARA         | 27       |
| BARI           | 8        |
| BERN           | 11       |
| BOLOGNA-PADOVA | 37       |
| DUBNA          | 6        |
| LYON           | 2        |
| LNGS           | 13       |
| ROMA           | 18       |
| SALERNO        | 22       |
| SINP-MSU       | 13       |

Interrupted transfer - data need to be deleted and published again

# ACTIONS AND RECOMMENDATIONS

Out of 6486 events declared “decaysearched”, we have no more than 4177 (64%) in the DB

If we manage to have all events, DB size will scale up to 86 TB, and export time will increase

It is later than we expected!

Indeed these estimates are optimistic because I cannot check correctness of data (and we know that at least flags are missing), only presence of “reasonable” datasets

I cannot check each event one-by-one, local responsables **can**

This remains the responsibility of each scanning laboratory

# ACTIONS AND RECOMMENDATIONS

Local scanning responsables should check data – DB views are there to help you!

Flag events

Passing CHECK1 AND CHECK2 means that at least an event “looks correct” – but check physics flags (tau, charm, nue) in feedback

I cannot know how you found each event – local scanning responsables have to decide whether to provide Scanback/TrackFollow/ScanForth in **addition** to **TotalScan** and **Feedback** (which are the minimum requirements for publication)

Complete publication

Retransmit interrupted data sets