Repack for CASTOR

Daniele Francesco Kruse

IT-DSS/TAB

Agenda

1. What & Why we repack
2. When do we need the job done
3. How will we do it
4. Features of the new Repack (2.1.13)
5. Problems encountered and solved
6. The repack command
7. Current results
8. Conclusions

What & Why

To ***repack*** a *tape* is to move data from that tape to one or more *other tapes*.

Extremely Important! Reasons:

1. Defragment
2. Move data away from faulty tapes
3. Move data to higher capacity tape media
4. Reuse media (TSM)

When

More than 100 PB to repack between 2013 and 2015! 3 GB/s needed on average!

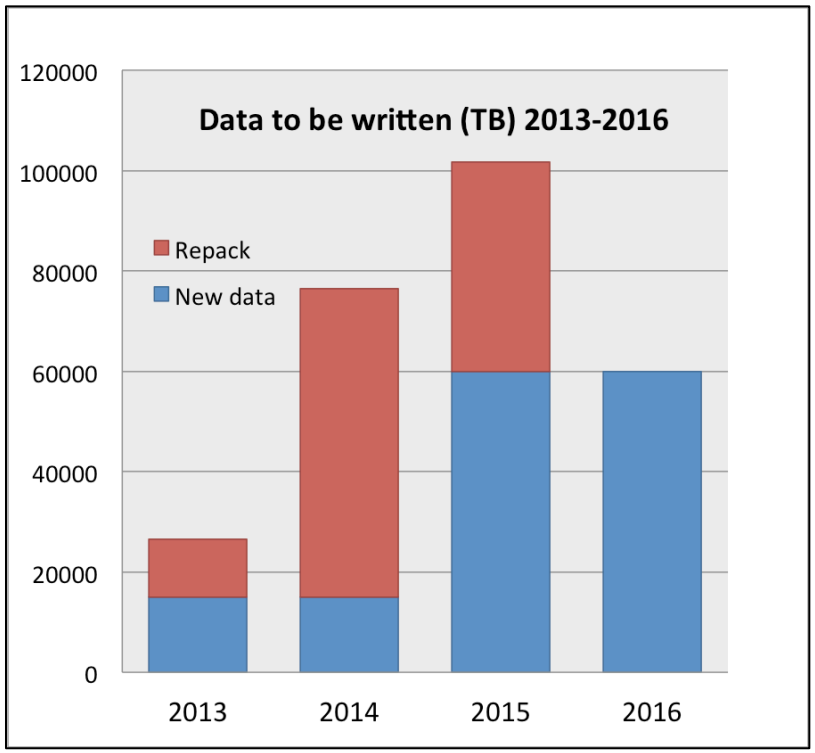

How

New ***repack system*** introduced in CASTOR version 2.1.12 and perfected in 2.1.13

- Faster startup
- More efficient usage of tapes
- Tape queuing (feeder no longer needed or needed only for very high level control)
- 2.1.13 is the current recommended version for *production* bulk repacks

2.1.13 Repack Features (1/2)

- Refactoring of the tape schema for recall

  (See *Sebastien*'s talk this afternoon)
- *Auto correction* of inconsistencies e.g.:
  - Dual copy file with only one segment
  - Files with wrong copy number
- Correction handling of real time

  modification and *deletion*
- 100% PL/SQL based

2.1.13 Repack Features (2/2)

- The DB job (no repack daemon!) logs its activity to `repackd.log`
- 1000s of tapes can be submitted
- Only 1 tape is pre-processed at a time (1M files/hour): `nslisttape` & stager initialization
- DB entities are created for all needed recalls and migrations
- Significant usage of DB link
- Bulk interfaces (less load on DB)

Problems Encountered & Solved

Issues encountered & solved during test runs:

1. Too many tape copies
2. Wrong copy numbers
3. File size and/or checksum mismatch
4. Logically deleted files

Refinement of the repack command: bug fixes, enhancements, interface improvements

The Repack Command

- Repack a list of tapes:

```
repack --bulkvolumeid tape_list.txt -o default
```

- Check repack status:

```
repack -s
```

- Check repack status of a single tape:

```
repack -S vid
```

- Check repack errors on a tape:

```
repack -e vid
```

- Abort a repack:

```
repack -R vid
```

Repacked data until now

These year's tests with the new repack:

- Around **3000 1TB** tapes (**2.5** PB of data)
- **1.1** PB in a single run (**1600** tapes)
- **1** PB/month (**10** PB/month needed)

SubmitTime	RepackTime	User	Machine	Vid	Total	Size	toRecall	toMigr	Failed	Migrated	Compl%	Status
25-Sep-12 17:17	7d 22:51	root	c2cernt3srv301.cern.ch	I17715	11066	810.76GiB	11064	0	0	2	0%	ONGOING
25-Sep-12 17:17	7d 22:51	root	c2cernt3srv301.cern.ch	I17716	9362	789.42GiB	9361	0	0	1	0%	ONGOING
25-Sep-12 17:17	7d 22:51	root	c2cernt3srv301.cern.ch	I17717	9615	684.10GiB	9613	0	0	2	0%	ONGOING
25-Sep-12 17:17	7d 22:51	root	c2cernt3srv301.cern.ch	I17718	10598	843.49GiB	10598	0	0	0	0%	ONGOING
25-Sep-12 17:17	7d 22:51	root	c2cernt3srv301.cern.ch	I17719	15512	844.94GiB	15512	0	0	0	0%	ONGOING
25-Sep-12 17:17	7d 22:51	root	c2cernt3srv301.cern.ch	I17720	12450	828.88GiB	12450	0	0	0	0%	ONGOING
25-Sep-12 17:17	7d 22:51	root	c2cernt3srv301.cern.ch	I17721	8858	715.80GiB	8857	0	0	1	0%	ONGOING
25-Sep-12 17:17	7d 22:51	root	c2cernt3srv301.cern.ch	I17728	23532	0.98TiB	23532	0	0	0	0%	ONGOING
25-Sep-12 17:17	7d 22:51	root	c2cernt3srv301.cern.ch	I17730	21268	569.74GiB	21268	0	0	0	0%	ONGOING
25-Sep-12 17:17	7d 22:51	root	c2cernt3srv301.cern.ch	I17731	17028	983.08GiB	17028	0	0	0	0%	ONGOING
...												
03-Oct-12 16:08	-	-	-	TOTAL	8657847	1.10PiB	6477990	3117	47	2176693	25%	ONGOING

Functionality is very satisfactory, need to improve *performance* of the underlying HW.

Conclusions

- **Tape** is NOT *dead*! (...as long as **Repack** is *alive*)
- The new ***repack system*** is out (2.1.13)
- Tested and used in *production* on a large amount of tapes: *functionality* is very good
- Software is stable and a good work horse
- *Performance*: need to cope with over 100 PB of data to be repacked over a period of 2 years