



CASTOR 2 Operations Infrastructure & Integration

Miguel Coelho dos Santos
CERN / IT



CASTOR Operations



❖ What are 'CASTOR operations'?

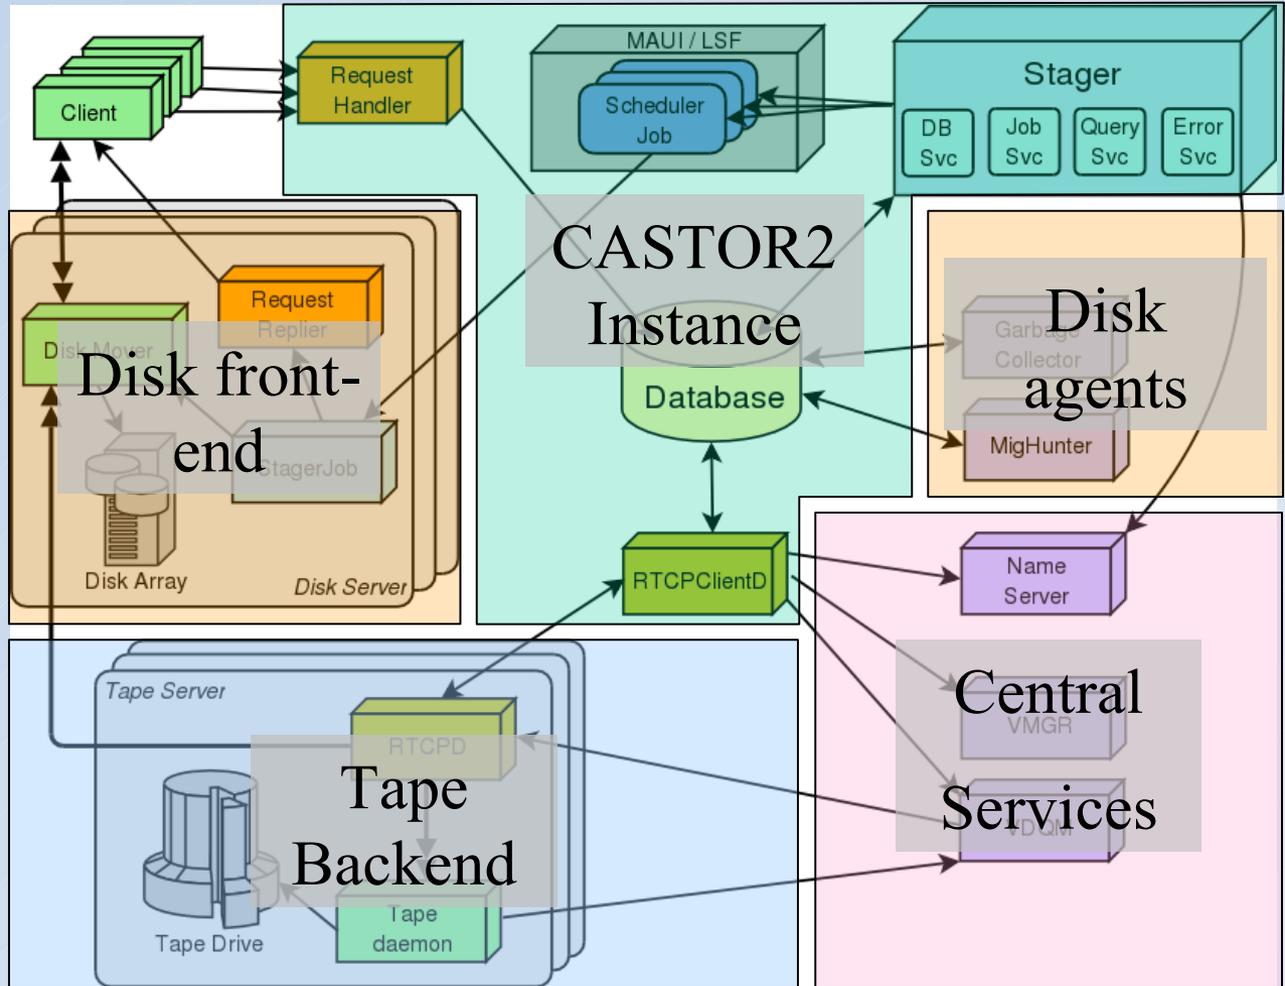
- Install, configure, manage different instances.

what is a CASTOR instance?

- Making sure the service runs (Meets/satisfies client's needs)
- Provide user support

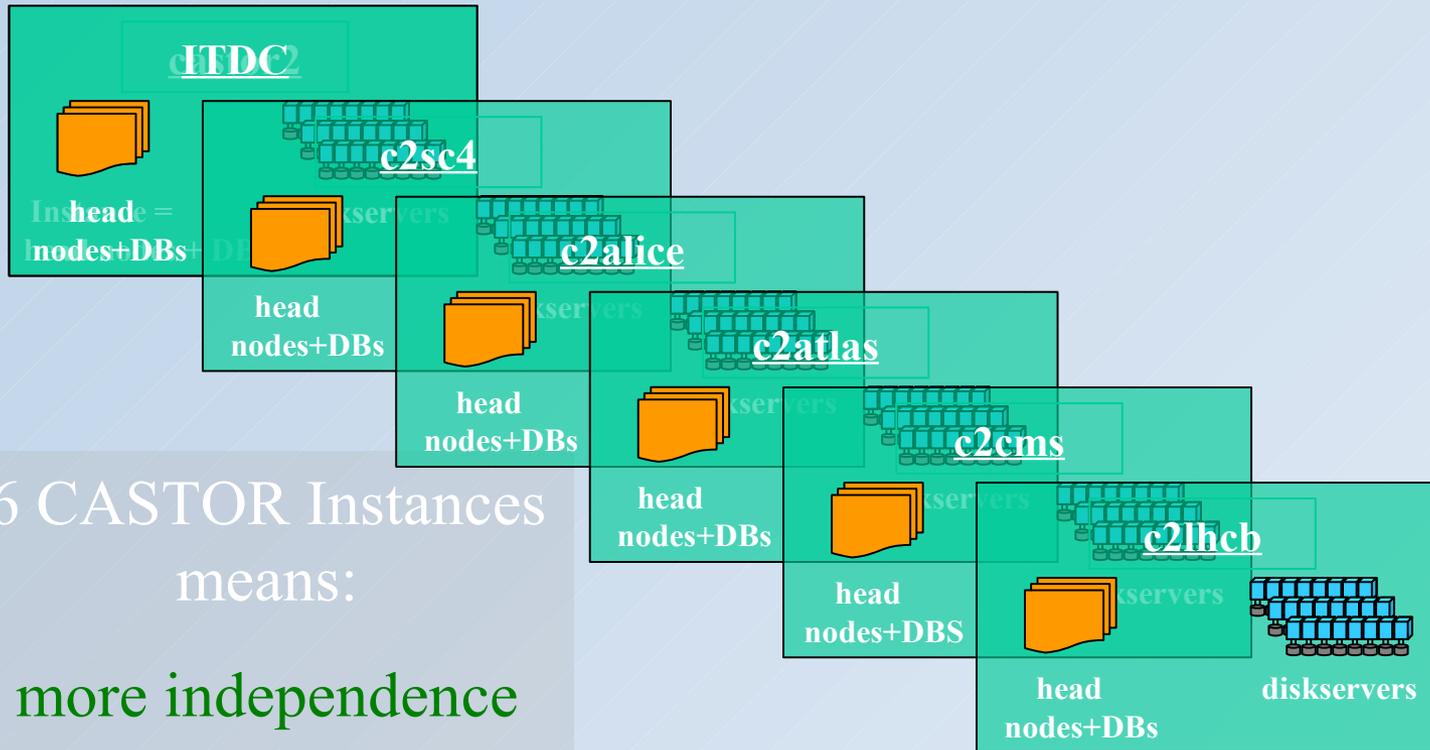


CASTOR Architecture





CASTOR Instances Infrastructure



6 CASTOR Instances
means:

more independence

less resource sharing

more capacity

more complexity



CASTOR Procedures



- ❖ Procedures are documented on the web (twiki)
 - Head node installation, configuration
 - Installing, putting a diskserver in and out of production
 - Svcclass, diskpool, tapepool reconfiguring
 - Drain and remove diskserver from an instance
 - Handle hardware failures
 - Power cut recovery
 - etc



CASTOR Procedures



- ❖ Regularly executed procedures have been automated
 - Restart stuck tape recalls
 - Clean tape copies of deleted disk copies (fixed in recent releases)
 - Generate LSF configuration files (keep cdb, rmmaster (stager), and LSF consistent).
 - Clean dead connections (if a batch job is killed the client is gone but connections stay open on CASTOR for a considerable amount of time using LSF slots)

- ❖ Work is ongoing:
 - Restart stuck migrations (almost ready)
 - Next big step is to fully automate diskserver installation

Tricky because of rmmaster



CASTOR Operations

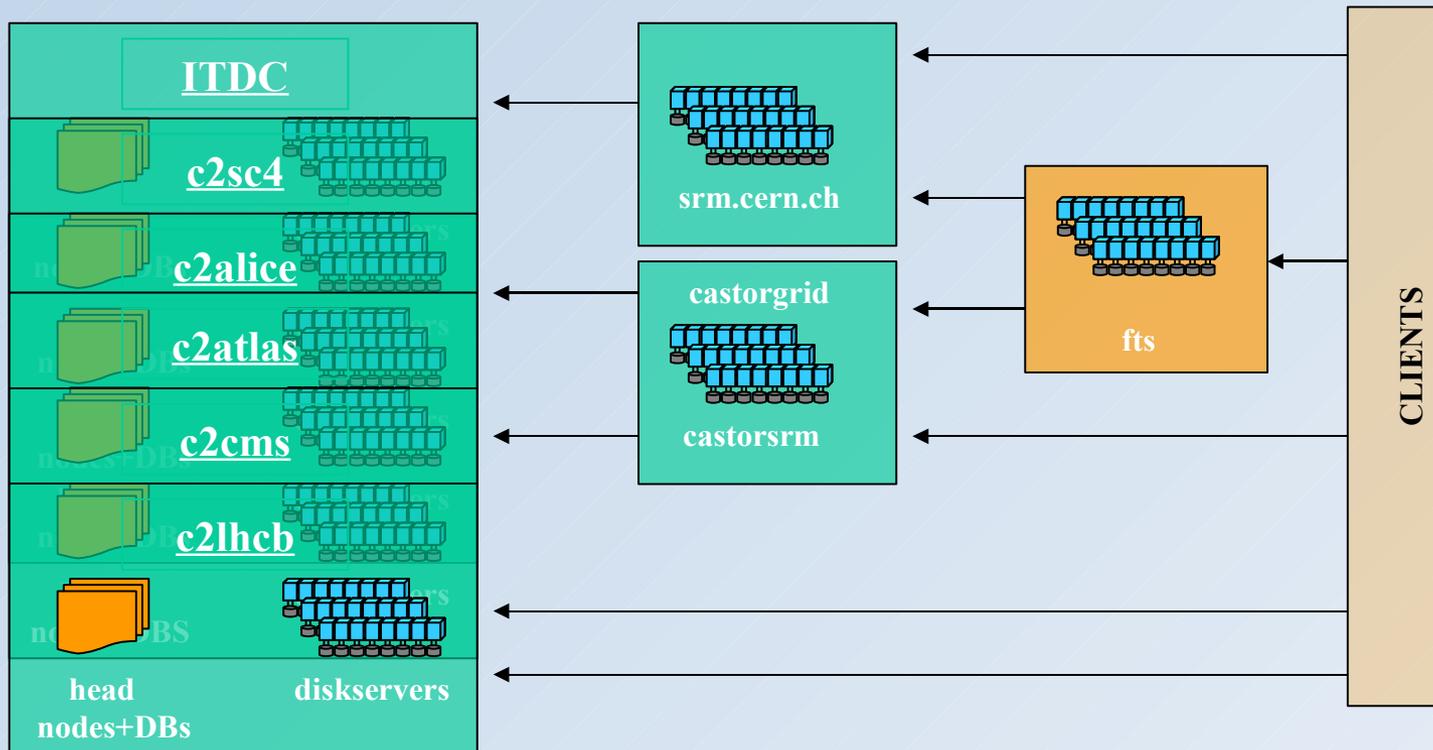


❖ What are 'CASTOR operations'?

- Install, configure, manage different instances
- Making sure the service runs (Meets/satisfies client's needs) **what is the CASTOR service?**
- Provide user support



Service Infrastructure





CASTOR service



Access protocols / tools:

rfio and root

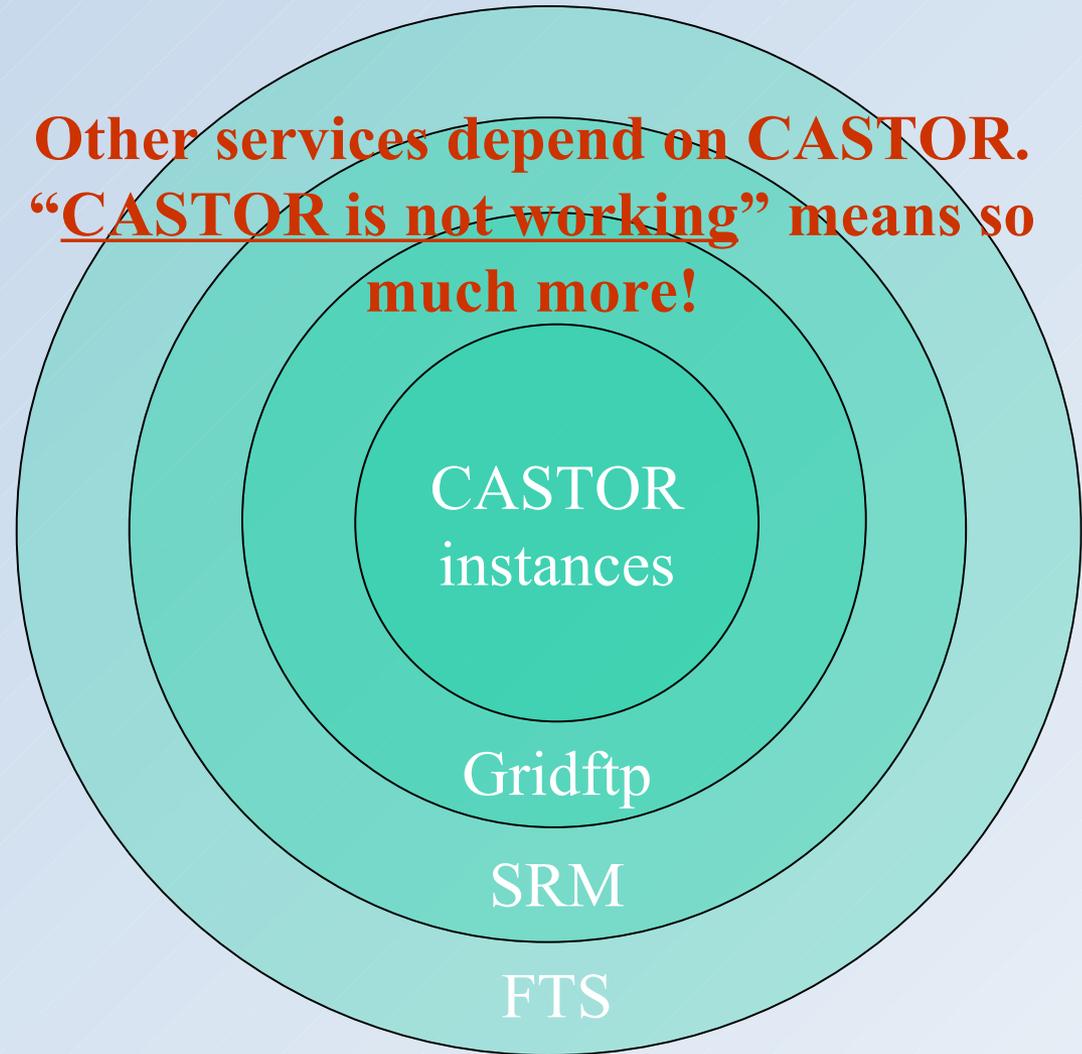
globus-url-copy

srmcp

LCG tools

**No 3rd level support
for SRM v1.1 ...**

**Other services depend on CASTOR.
“CASTOR is not working” means so
much more!**

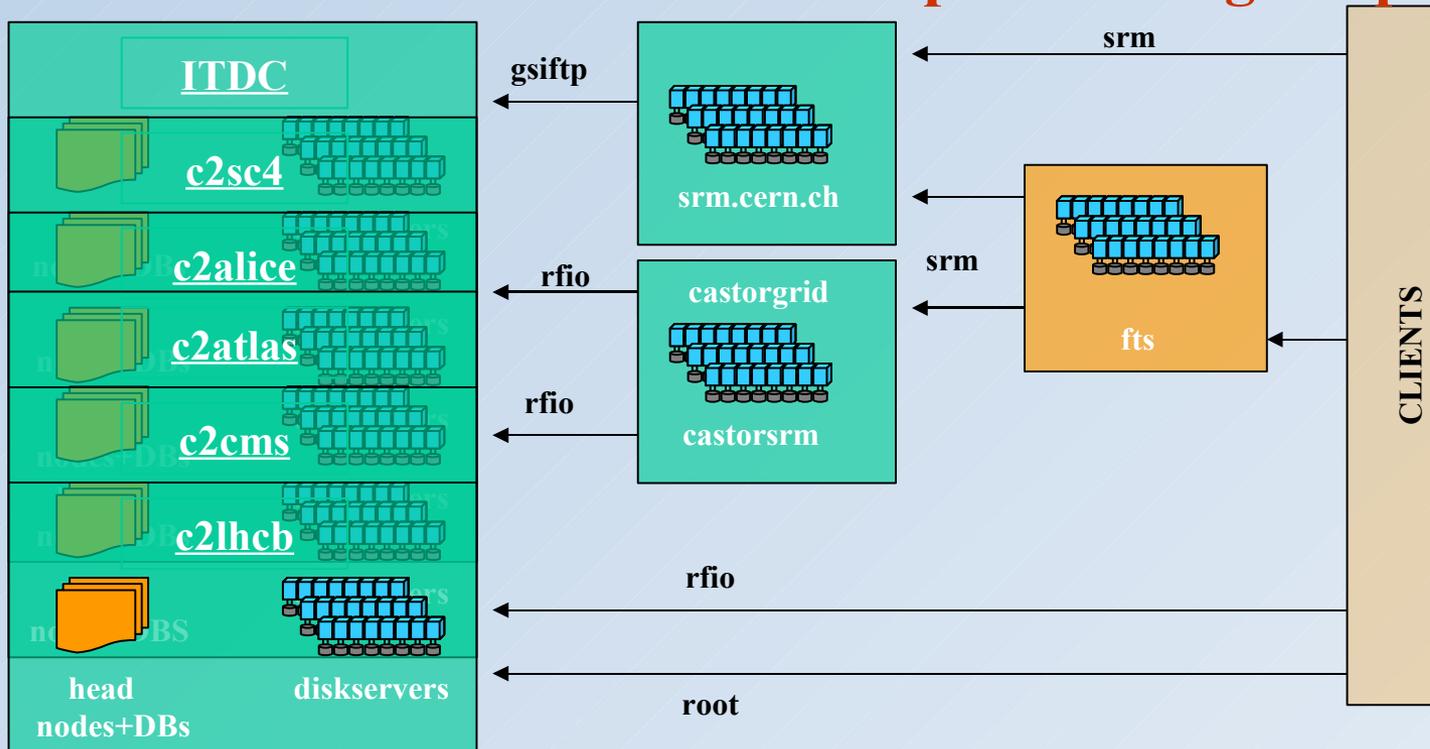




Service Infrastructure

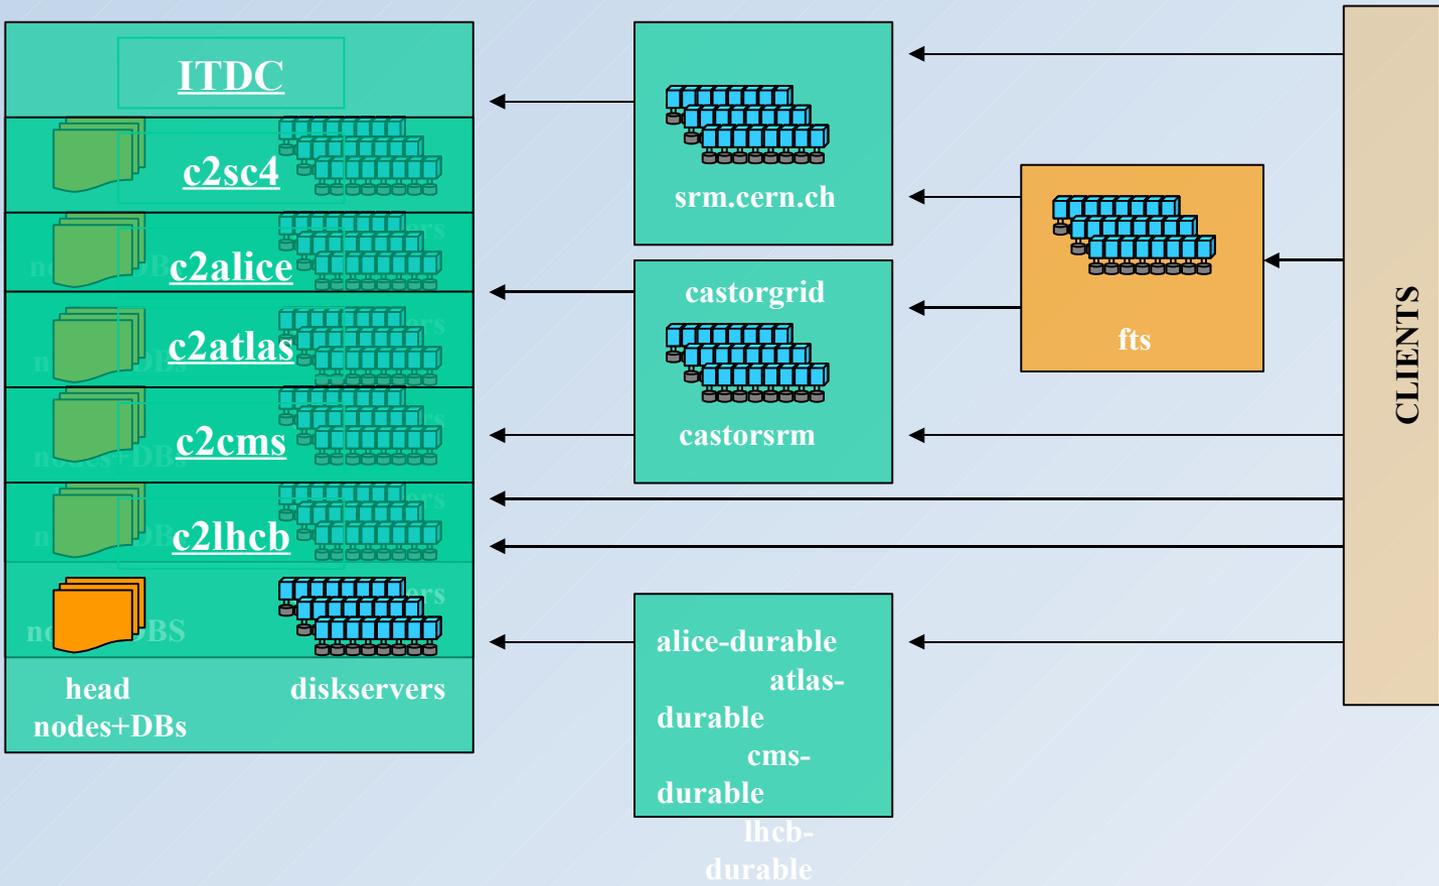


To make up for SRM v2 that we should be running, in the very near future, 4 srm v1.1 durable endpoints will go to production!





Service Infrastructure





CASTOR Monitoring



- ❖ Monitoring is a work in progress

- ❖ Two big reasons for monitoring:

- Problem detection / prevention
- Accounting

- ❖ Monitoring infrastructure:

Log parsing is minimal, except for SRM v1.1 and Gridftp

- Sensors collect metric data from stager and DLF DBs
- LEMON service is used to run the sensors, transport and store metric data
- Castor Status Pages use altered Lemon Status Pages to download information and display it (but the same framework)

We monitor services and processes, for example: stagers, svcclasses (diskpools), tape recalls and migrations.



CASTOR Monitoring



❖ LEMON framework basics:

- An agent runs on a node collecting metric data and forwarding it to a DB
- The LEMON DB stores data from a series of metrics
- The Lemon Status Pages (LRF) are configured in order to group hosts in clusters/subclusters and aggregate some metrics (definable)

❖ CASTOR and LEMON

- Besides normal node information, service information is collected and forwarded to LEMON. Data is reported on behalf of a stager/svcclass/diskpool/tapepool by c2xxxxsrv05 (DLF head node)
- A version of LRF adapted for CASTOR downloads all metrics
- Instead of hosts->subclusters->clusters, aggregation of information is done from diskserver->svcclass(diskpool)-> stager with service information being added at the appropriate level
- Display pages are then adapted to show information differently



CASTOR Monitoring



❖ General page

Castor instance: c2atlas

Diskpool	Total Size (TB)	Occupancy (TB)	Usage (%)	fs count	hostcount	Recall Queue	Migration Queue	Staged Files
analysis	5.4	3.5	64.8	3	1	0	0	0
atldata	5.4	0	0	3	1	0	0	0
default	74	25.5	34.5	45	15	969	466	551472
t0merge	10.1	3.4	33.7	7	2	0	0	0
wan	28.7	7.4	25.8	18	6	0	2610	19644
Total	123.6	39.8	32.2	76	25	969	3076	571116

Castor instance: c2cms

Diskpool	Total Size (TB)	Occupancy (TB)	Usage (%)	fs count	hostcount	Recall Queue	Migration Queue	Staged Files
cmsprod	21.8	10.4	47.7	12	4	0	0	18922
default	65.2	33.3	51.1	43	13	70	882	339628
wan	32.6	23.8	73	18	6	1	575	46701
Total	119.6	67.5	56.4	73	23	71	1457	405251

CASTOR1 limit was about 200k.

During SC3, when all experiments shared an instance, it reached 1M.



CASTOR Monitoring



❖ Operations page

```
[ c2alice instance ][ consistency ]
```

Diskpool	Total Size (TB)	Occupancy (TB)	Usage (%)	fs count	hostcount	Recall	Queue	Migration Queue	Staged Files
alimdc	2.9	2.2	75.9	3	1	0	8	793	
default	7	4.5	64.3	14	5	16	202	275498	
spare	43.5	0	0	24	8	0	0	0	
wan	16.3	1.2	7.4	9	3	0	0	30302	
Total	69.7	7.9	11.3	50	17	16	210	306593	

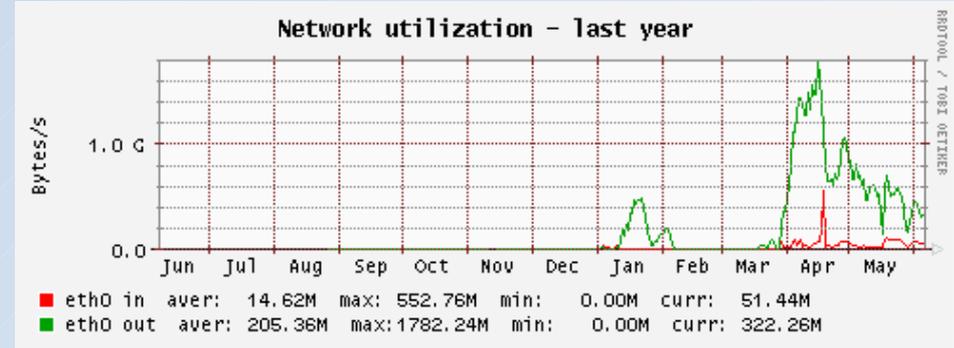
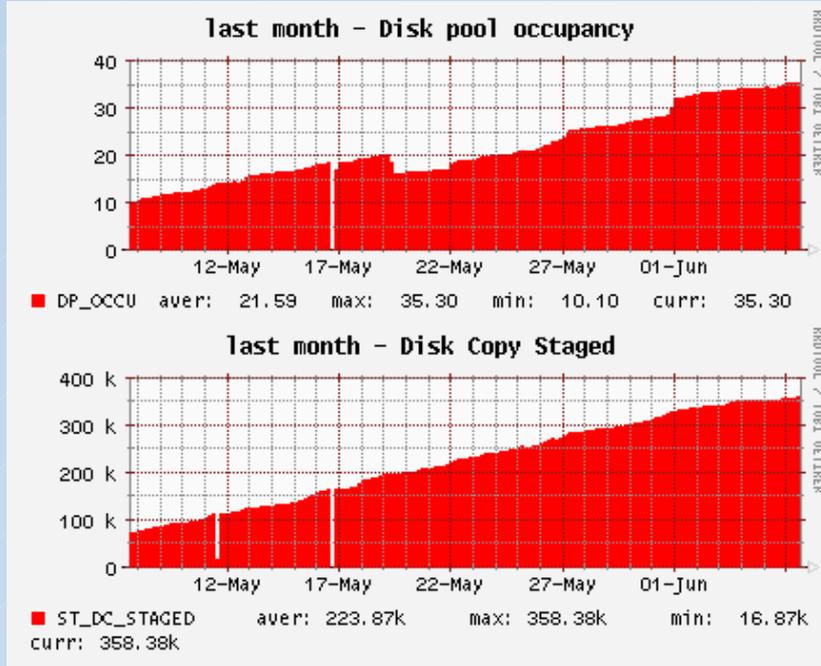
CDB, LSF, rmmaster consistency check

Instance page: stager data and Svcclass data aggregation

Svcclass page: disk copies, tape copies, streams, subrequests, diskserver data aggregation (network or disk usage for example)



❖ Some examples:



SC4 network activity

ATLAS wan service class: disk pool occupancy and count of staged files



CASTOR Operations



❖ What are 'CASTOR operations'?

- Install, configure, manage different instances
- Making sure the service runs (Meets/satisfy client's needs)
- Provide user support



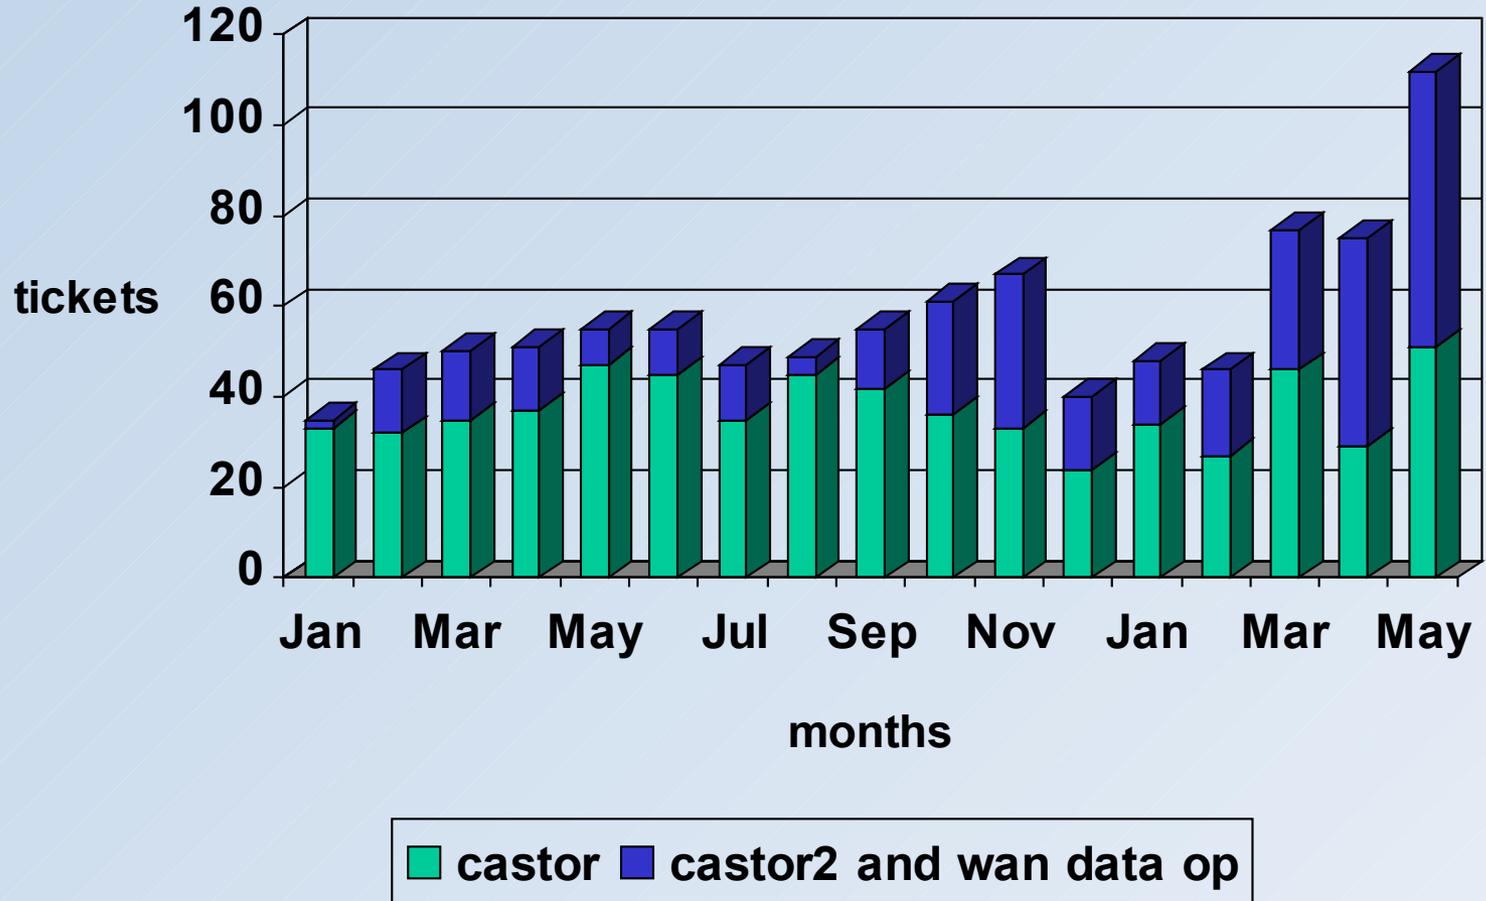
What is user support



- ❖ Normal questions on how to use CASTOR
- ❖ Expert question from the same power users
- ❖ Data challenges preparations: gather requirements and set in place the infrastructure
- ❖ Gridftp and SRM support:
 - Normal user questions
 - Gridmap-file user mappings
 - Problem detection



Ticket distribution per month





Problem handling



- ❖ Problem report flows are being addressed
 - Reorganization is ongoing
 - Separation of operations flows and development flows
- ❖ SMoD will start handling more and more tickets
 - Automation still needs improvements



Thank you



❖ Questions?