

# LFC LHCb



Barbara Martelli  
WLCG Workshop

CERN 22- 26 January 2007

# Outline

---

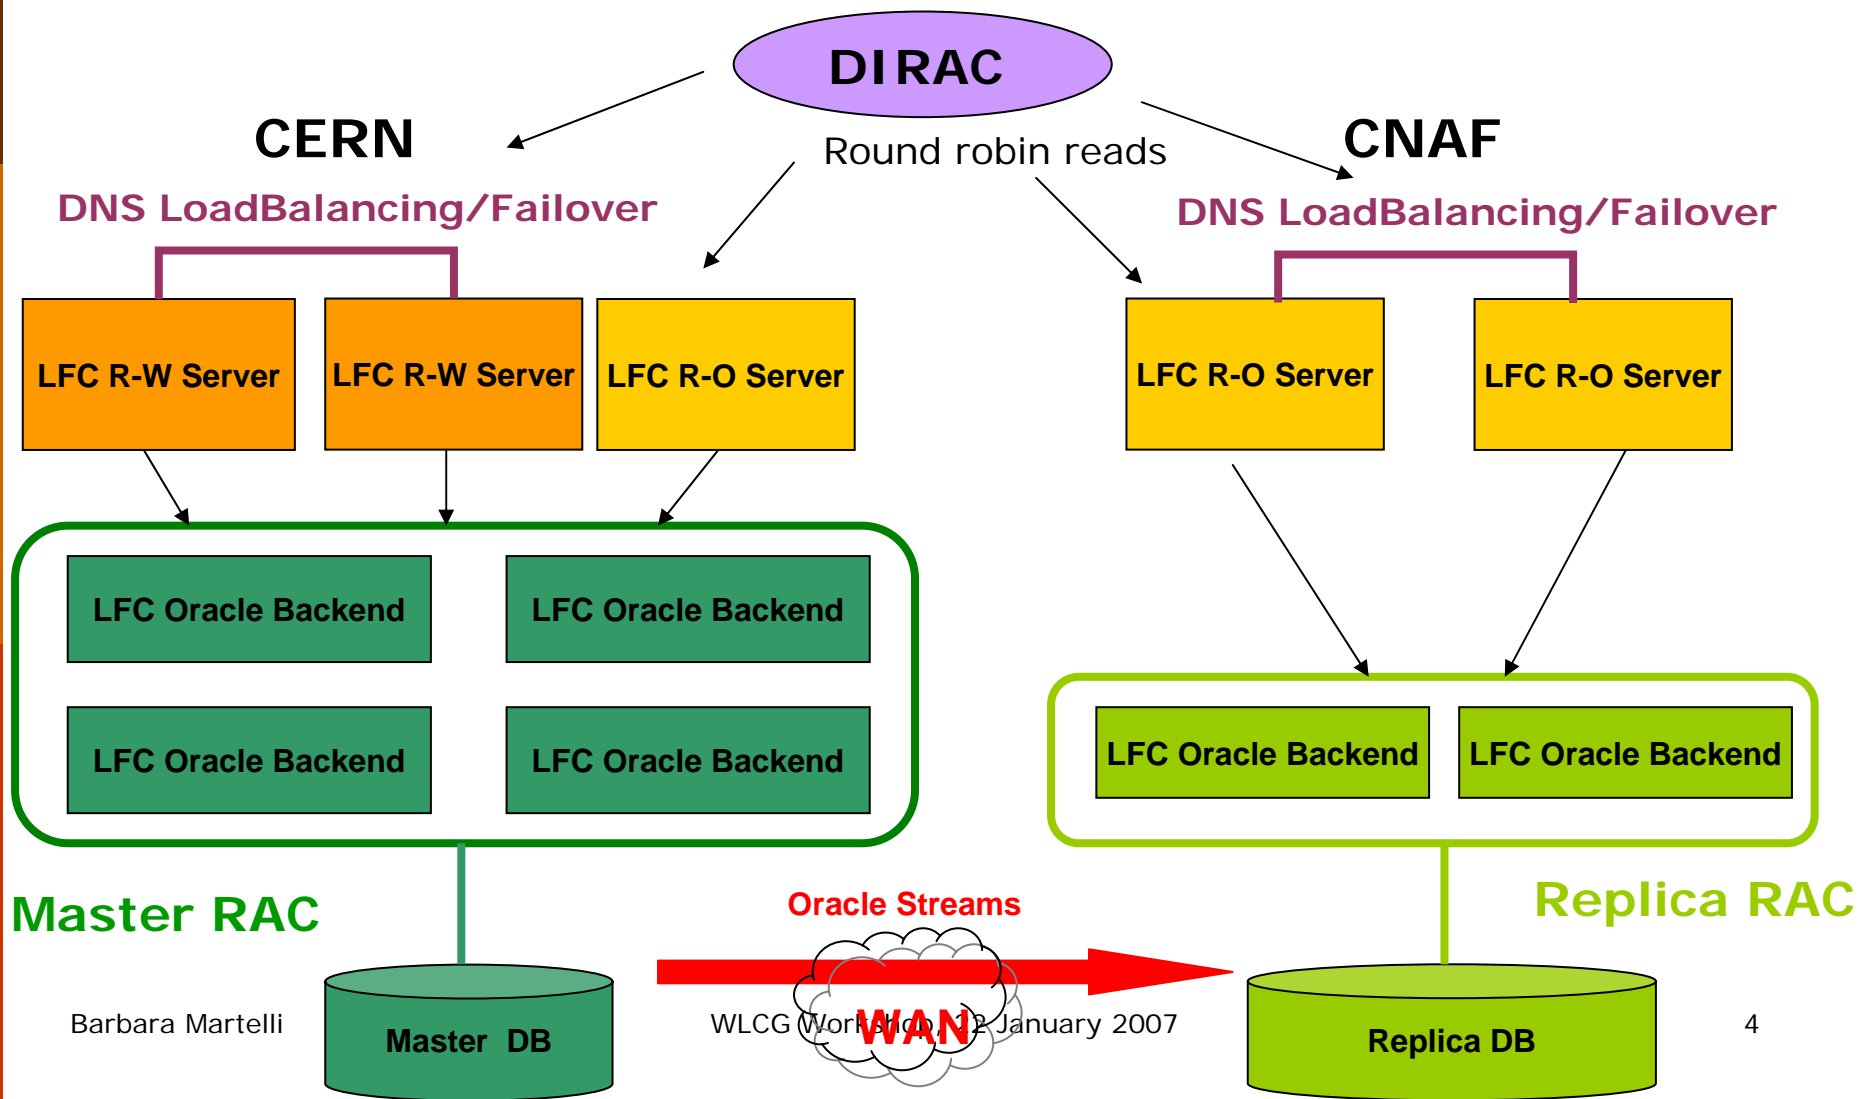
- LFC LHCb architecture and deployment
- Present usage
- Stress test on replication
- Conclusions

# LHCb LFC read – only replica

---

- ❑ Why: LFC LHCb needs more scalability and fault tolerance.
- ❑ How: adding a read only LFC replica at CNAF. The CNAF slave is maintained **consistent** with the CERN master **at the database level** through the *Oracle Streams* replication.
- ❑ A *Streams* sustained rate of 170 KB/s is shown by preliminary throughput tests.
- ❑ LFC preliminary tests showed a sustained rate of 40 entries per second (entries inserted in the master and replicated at CNAF).

# LFC LHCb Deployment



# LFC Read Only Configuration

---

At LFC level, very simple configuration, in:

```
/etc/sysconfig/lfcdaemon
```

```
set
```

```
RUN_READONLY="yes"
```

Then the LFC will prevent write operations:

```
$ lfc-mkdir /grid/dteam/hello cannot create  
/grid/dteam/hello: Read-only file system
```

This is supported from **LCG-2\_7\_0**

At DB level we use an user with read only grants on LFC tables.

# LFC Present Usage


---

- ❑ Jan 8 to Jan 16 (9 days) LFC logs analyzed.
- ❑ A rough estimate of only 10 LFC API calls per second (in total over the 5 LFCs).
- ❑ Write operations are the 0.5% of the total traffic.
- ❑ Write operations are well distributed among the 2 LFC R/W servers
- ❑ Read operations seem to privilege CERN LFCs.


# Present Usage – Read Operations

Most used operations are *statg*, *getreplica*, *access*



Total of 5M read operations in 9 days  
(about 6-7 per second)



# Read Operations Share on LFCs




**getreplica**



**statg**

**access**



**CNAF's LFC less used**

Barbara Martelli


WLCG Workshop



# Present Usage - Write Operations

Most used operations are *addreplica*, *setfsizeg*, *creat*

Total of 121K write operations in 9 days  
(about 9 per *minute*)



# Stress Test Setup I

---

- ❑ Python script which adds files in */grid/lhcb/lfc\_test/cnaf*
- ❑ For each file 5 replicas are added.
- ❑ For each file the related replicas are added all in the same session (to avoid authentication overload in LFC frontends).
- ❑ Test duration: 25K seconds (~7 hours), 800K files, 2M replicas.

# Stress Test Setup II

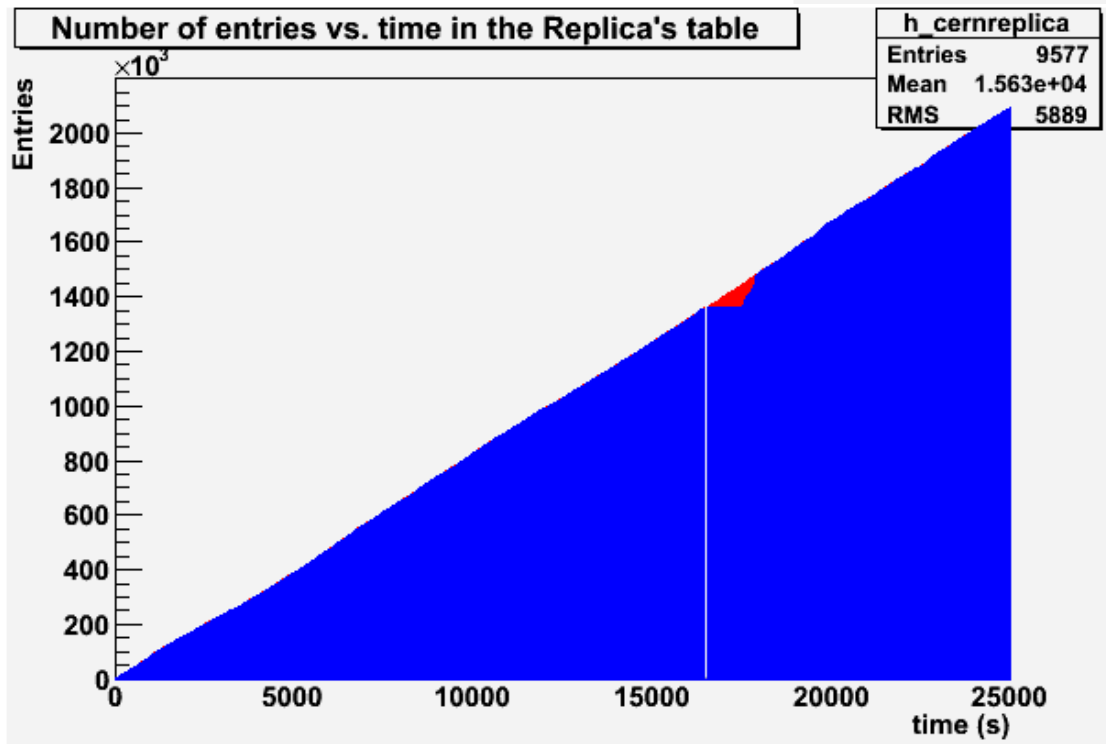
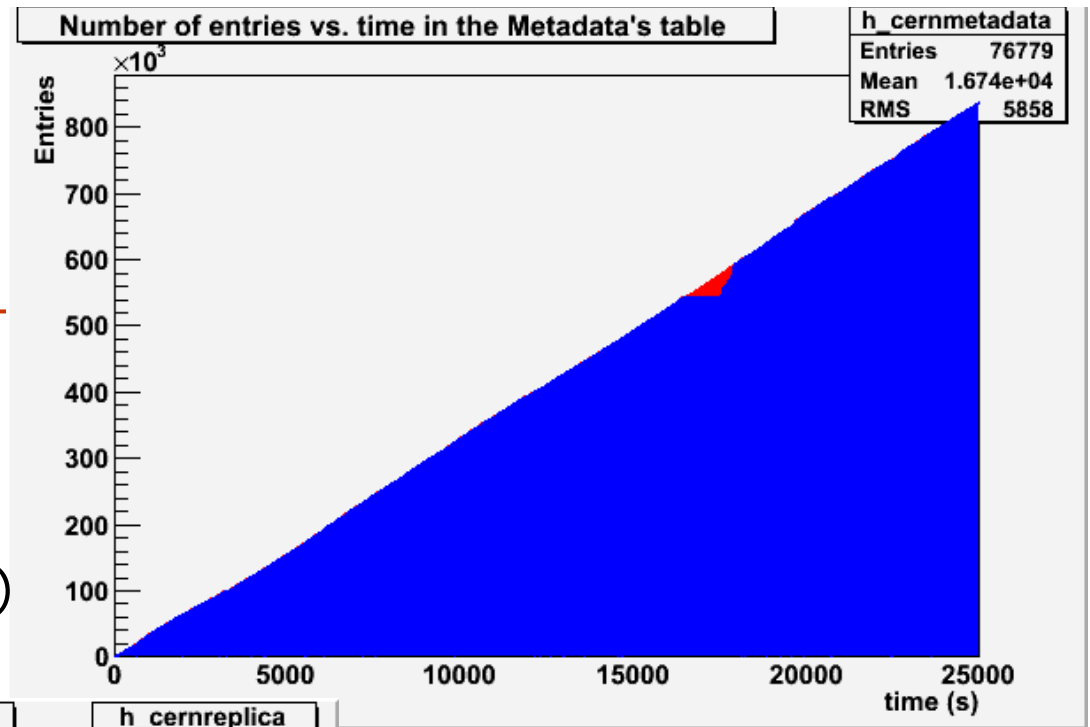
---

- ❑ At database level for each new file a row is inserted in the CNS\_FILE\_METADATA table. For each replica a row is added into CNS\_FILE\_REPLICA table.
- ❑ The script monitors the status of the master and slave database backends once a second. It counts the number of files and replicas present in each of them.
- ❑ We use the *strmmon* monitor to measure the replication latency at DB level.

# Population Rate

- CERN
- CNAF

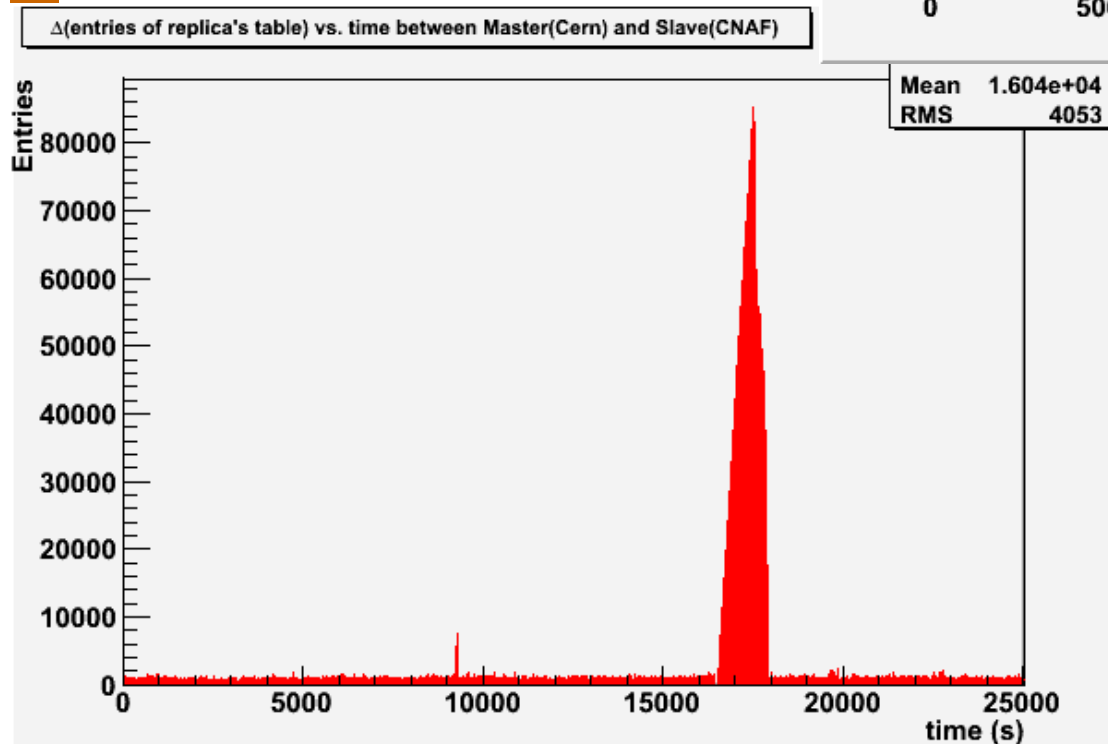
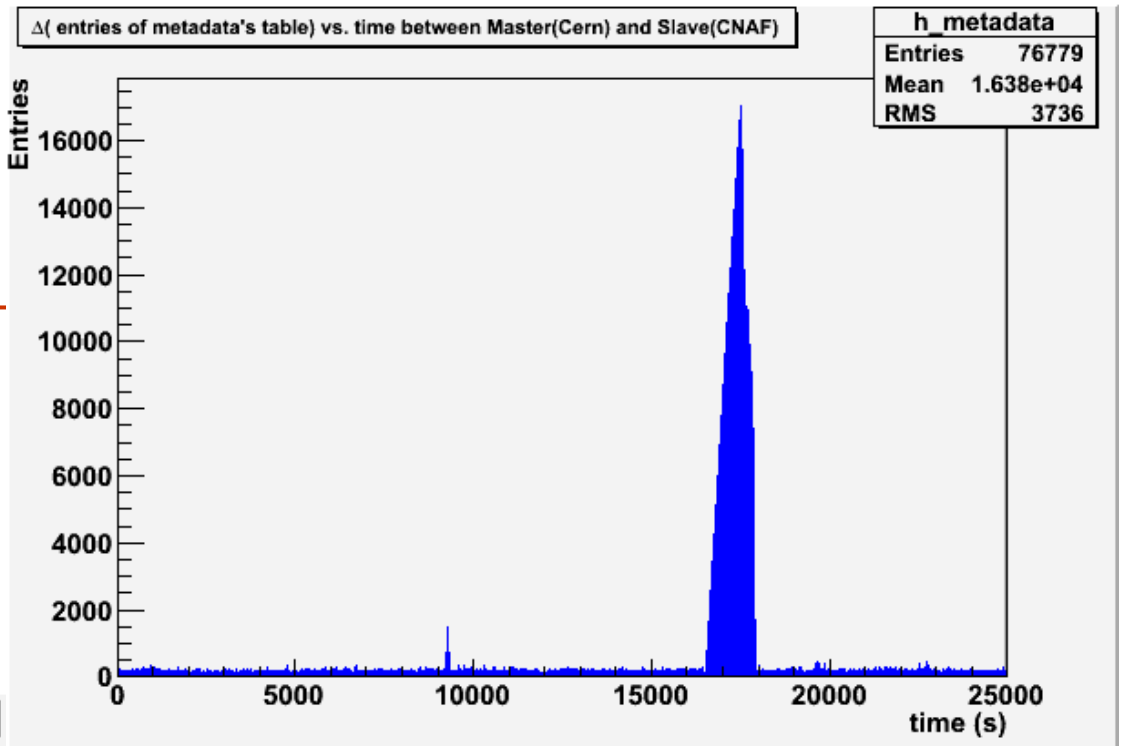
25000 seconds (7 hours test)



- ▣ Sustained rate of 84 entries/s on CNS\_FILE\_REPLICA
- ▣ Sustained rate of 36 entries/s on CNS\_FILE\_METADATA
- ▣ Total: 120 entries/s
- ▣ Only once CNAF's slave become misalignmed for about 1 hour

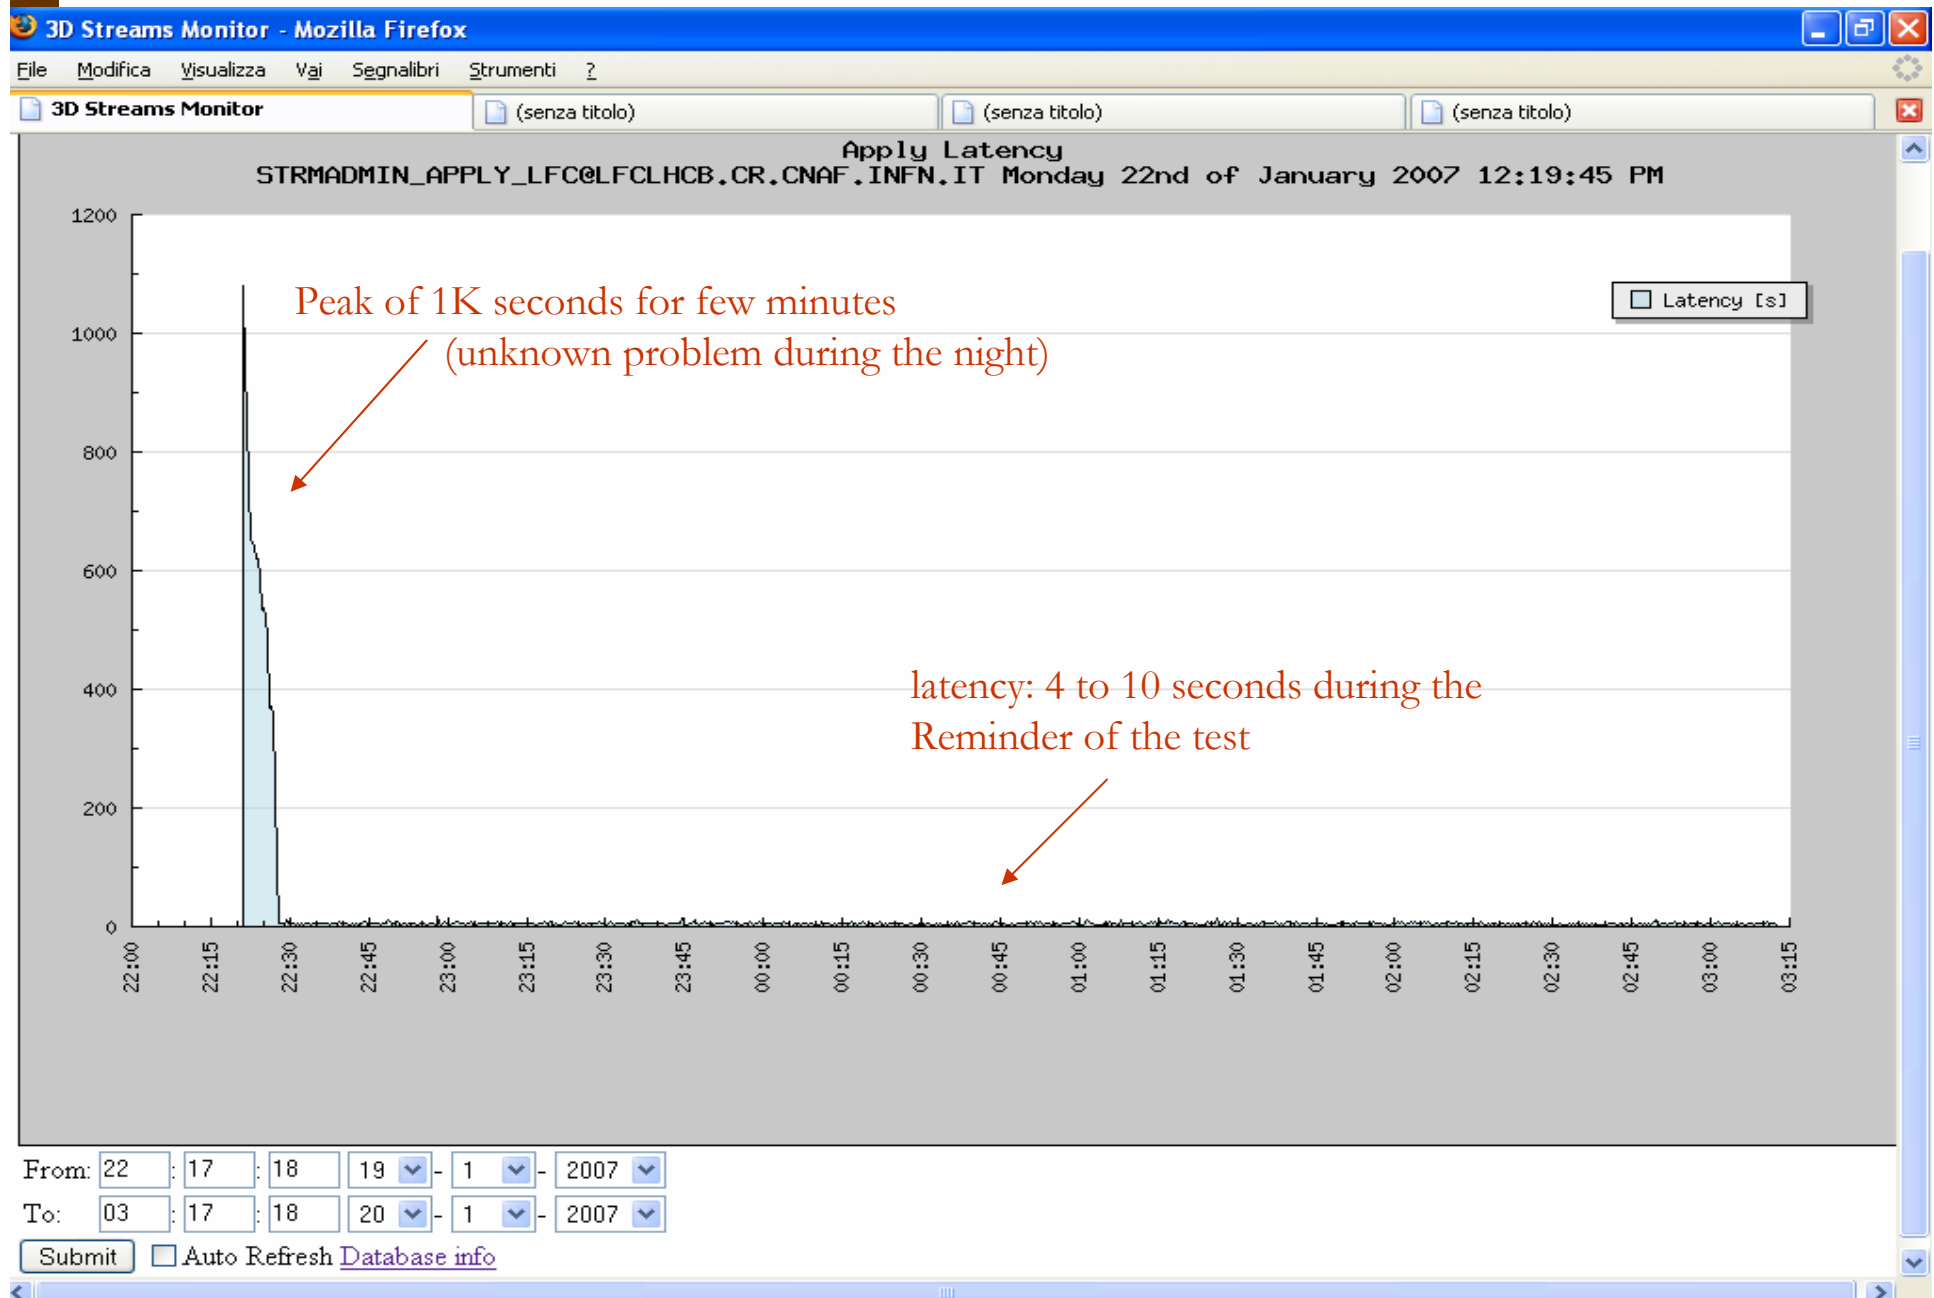
007

# Difference between entries in Master LFC and CNAF's Slave

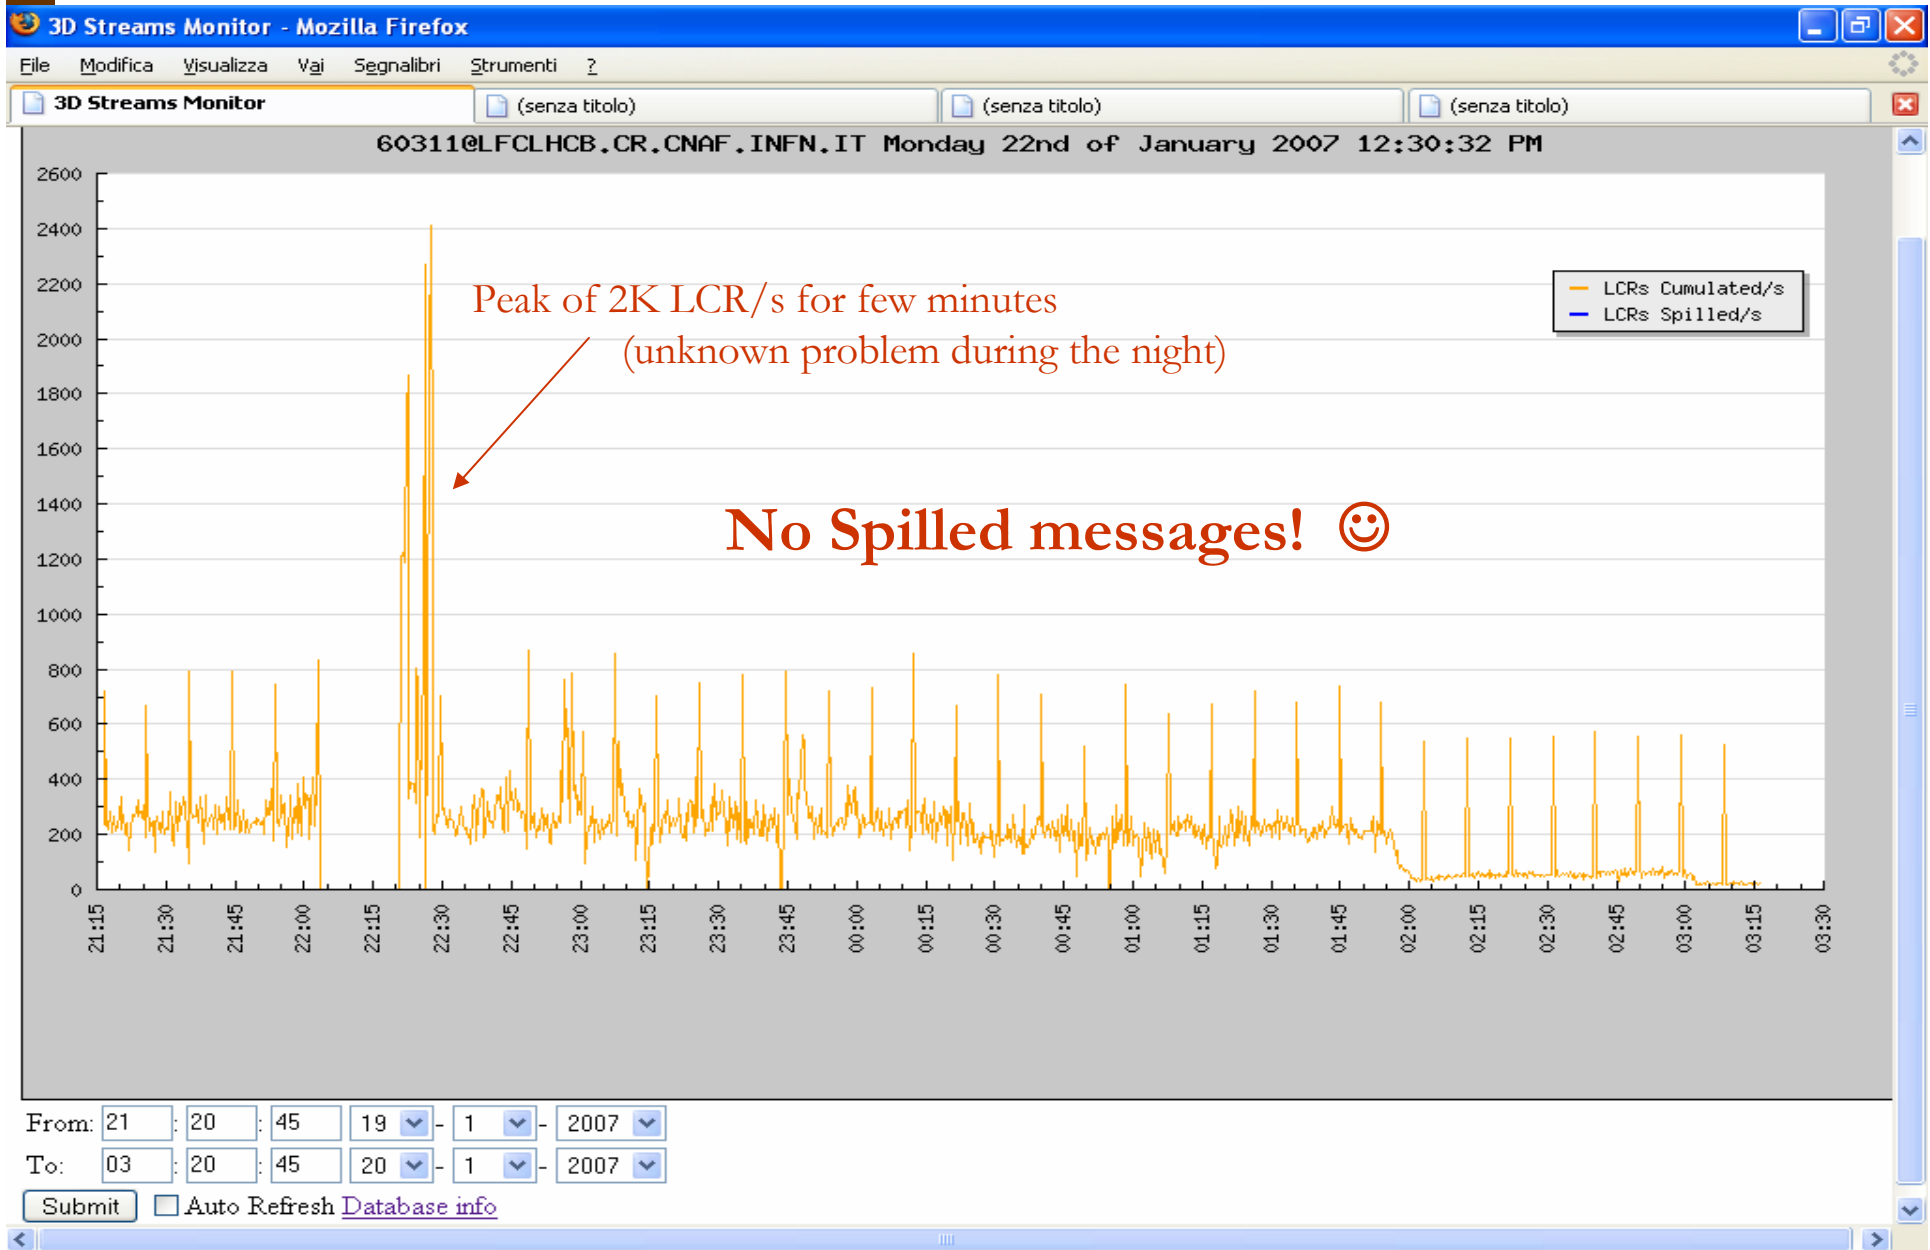


- Usually not higher than 200 entries in CNS\_FILE\_METADATA A and 1000 in CNS\_FILE\_REPLICA
- Peak of some K entries for about 1hr
- CNAF's DB succeeded to apply all missing entries.

# Strmmon View: Replication Latency



# Strmmon View: Spilled Messages



# Conclusions

---

- ❑ LFC LHCb deployment is stable.
- ❑ Sustained rate of 120 entries/s has been achieved
- ❑ NO SPILLING has been detected in case of database misalignment between master and slave. The messages are accumulated by the master only in the archived log.
- ❑ The Master-slave synchronization has been exercised.
- ❑ LFC at CNAF are less used for read operations than the CERN ones. We should analyze this anomalous behavior.
- ❑ Many thanks to Angelo Carbone and Gianluca Peco who made the stress test for LHCb.