



Michal Simon

# xrdcp primer



# Metalinks & ZIP archives

- **Metalinks are first class citizens**
  - They are treated as **virtual redirectors** (not just a list of alternative locations)
  - This means they can be used to **recover from errors on request-by-request basis**
  - Allow to **omit *kXR\_wait***
  - Can contain file checksum
- **ZIP archives are also first class citizens**
  - Currently we support **listing, append and extraction** from ZIP archives
    - Including decompression (only the default deflate algorithm)
  - In most cases it **feels like a directory**
  - Always contain file checksum

# Metalinks & ZIP archives

- **--zip**: extract file from ZIP archive
- **--zip-append**: append file to ZIP archive
- **--zip-mtln-cksum**: prefer Metalink checksum over ZIP checksum
  - It is useful if one of the replicas listed in a Metalink is a file packed in a ZIP archive

# Copy multiple files

- **--infile**s: the input file contains a list of files to be copied

```
xrdcp --infile ./list_of_files root://server.cern.ch//data
```

- **--recursive**: recursively copy a directory

```
xrdcp -r ./directory root://server.cern.ch//data
```

- In both cases multiple files are copied, each having a corresponding **copy job**

# parallel vs streams

- **--parallel:** defines the **number of threads** that are used to carry out copy jobs

```
xrdcp --parallel 3 ./dir3files root://server.cern.ch//data
```

- Consider an example where one is copying a directory containing 3 files
  - **By default xrdcp uses single thread to carry out copy jobs**, meaning the files will be copied sequentially one after another
  - Setting parallel to 3 will make all the 3 copy jobs run in parallel

# parallel vs streams

- **--streams**: defines the number of physical TCP connections used per channel

```
xrdcp -S 15 ./file.dat root://server.cern.ch//data
```

- by default there is **just one physical connection** (no matter how many copy jobs are there!) per data server
- **multiple streams are meant for data transfers over WAN** (mitigate performance hit due to package loss, extend TPC buffer size)
  - or to exceed the throughput limitation imposed by paged reads
- we only support reads over multiple streams (pulling)
  - we could support also writes but we don't have a use case

# parallel vs streams

- So what is the **XRD\_CPPARALLELCHUNKS**?
  - It is the **number of chunks ‘in-the-fly’**
  - What does it mean?
    - *xrdcp* IS NOT sending a write request and waiting for server response
    - Instead, *xrdcp* **sends asynchronously a number of chunks** (by default 4) and whenever a response arrives it pushes an additional chunk (maintaining this way constant number of chunks ‘in-the-fly’)



# sources

- **--sources**: defines the number of sources to be used for a single copy job (extreme copy)

```
xrdcp -y 4 root://redirector.cern.ch//data/file.dat ./dir
```

- By default *xrdcp* opens and copies data from a single source file
- Extreme copy allows to copy the file from multiple replicas at once
  - The list of replicas is obtained either using the *locate* request from a redirector or from a *Metalink*

# Transfer rate

- **--xrate**: limit transfer rate
- **--xrate-threshold**: if the transfer rate drops below given threshold force the client to use different source or if no more sources are available fail the transfer

# continue & retry

- **--continue**: allows to continue copying the file from the point where the previous copy was interrupted (i.e. due to timeout)
- **--retry**: defines how many times a failed copy job should be retried
- **--retry-policy <force | continue>**: defines the policy for retrying copy jobs
  - **force**: start from scratch
  - **continue**: continue from the point where the previous copy was interrupted

# Miscellaneous

- **--xattr**: preserve extended attributes
- **--rm-bad-cksum**: remove the target if checksum check fails
- **--allow-http**: the name gives it away (requires HTTP plug-in)

# Questions?

