




XRootD




HTTP improvements and SciTags

Presented by Cedric Caffy on behalf of the EOS and XRootD team


EOS Workshop 2024
15/03/2024

HTTP protocol is more and more popular within WLCG

Amount of bytes read/written from/to EOS since the beginning of this year



WLCG Data Challenge 2024



WLCG Data Challenge 2024

The HTTP protocol in EOS

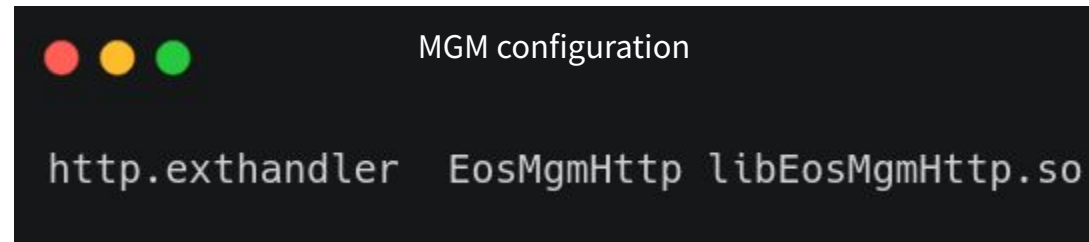
2 types of HTTP requests

- Webdav requests
 - Upload, download, deletion, checksum request, size request...
 - Verbs: PUT, GET, HEAD, DELETE...
- HTTP TPC (Third Party Copy) requests
 - An active server will PULL/PUSH data from/to another server
 - Verb: COPY

The HTTP protocol in EOS

Webdav requests handling

- EOS provides its own implementation of PUT, GET, HEAD, DELETE...
 - Via an XrootD HTTP external handler

A terminal window with a dark background and three colored window control buttons (red, yellow, green) in the top left corner. The title bar reads "MGM configuration". The terminal content shows the configuration line: `http.exthandler EosMgmHttp libEosMgmHttp.so`

```
MGM configuration
http.exthandler EosMgmHttp libEosMgmHttp.so
```

The HTTP protocol in EOS

HTTP TPC requests handling

- Uses the native XRootD HTTP TPC plugin

```
MGM configuration
http.exthandler EosMgmHttp libEosMgmHttp.so
http.exthandler xrdtpc libXrdHttpTPC.so
```

Some HTTP improvements

EOS

- FST - Chunk upload infinite loop in case no data is received from the client
 - Created a huge amount of logs and everlasting looping threads!

Some HTTP improvements

EOS

- FST - Chunk upload infinite loop in case no data is received from the client
 - Created a huge amount of logs and everlasting looping threads!

XRootD (Webdav)

- HTTP checksum handling
 - Support user-provided checksum implementation
 - Better error handling (checksum not supported, checksum on a non-existing file...)
 - Is compliant with the RFC 3230 - <https://www.ietf.org/rfc/rfc3230.txt>
 - Respects IANA registry for digest-algorithm names and value encoding (<https://www.iana.org/assignments/http-dig-alg/http-dig-alg.xhtml>)
 - Digest priority in the order they were received by the client

Some HTTP improvements

EOS

- FST - Chunk upload infinite loop in case no data is received from the client
 - Created a huge amount of logs and everlasting looping threads!

XRootD (Webdav)

- HTTP checksum handling
 - Support user-provided checksums
 - Better error handling (e.g. 404 for a non-existing file...)
 - Is compliant with [RFC 9530](https://www.rfcs.org/rfcs/9500/9530.txt)
 - Respects IANA registry for media types and value encoding (<https://www.iana.org/assignments/media-types/media-types.xhtml>)
 - Digest priority in the order they were received by the client

DEPRECATED!
(RFC9530)

Some HTTP improvements

XRootD (Webdav) - cont.

- Read range-request support (Thanks to David Smith !)
 - `curl -H "Range: bytes=0-50, 100-150" -X GET https://xrd-server.cern.ch//my/file.txt`
- SciTags support

Some HTTP improvements

XRootD (HTTP-TPC)

- Minor bug fixes
- SciTags support

SciTags

What is the use case?

- Understand **Who** and **How** the Research and Education (R&E) network is being used by the scientific community
 - Especially when critical links are overloaded, impacting workflows and data transfers

SciTags

What is a SciTag?

A 16 bits unsigned integer

$65 \leq \text{SciTag} \leq 65535$

SciTags

What is a SciTag?

- Provided by the user

SciTags

What is a SciTag?

- Provided by the user
- SciTag = $\text{experimentID} \ll 6 \mid \text{activityID}$
 - $\text{experimentID} = \text{SciTag} \gg 6$
 - $\text{activityID} = \text{SciTag} \& 0x3F$

SciTags

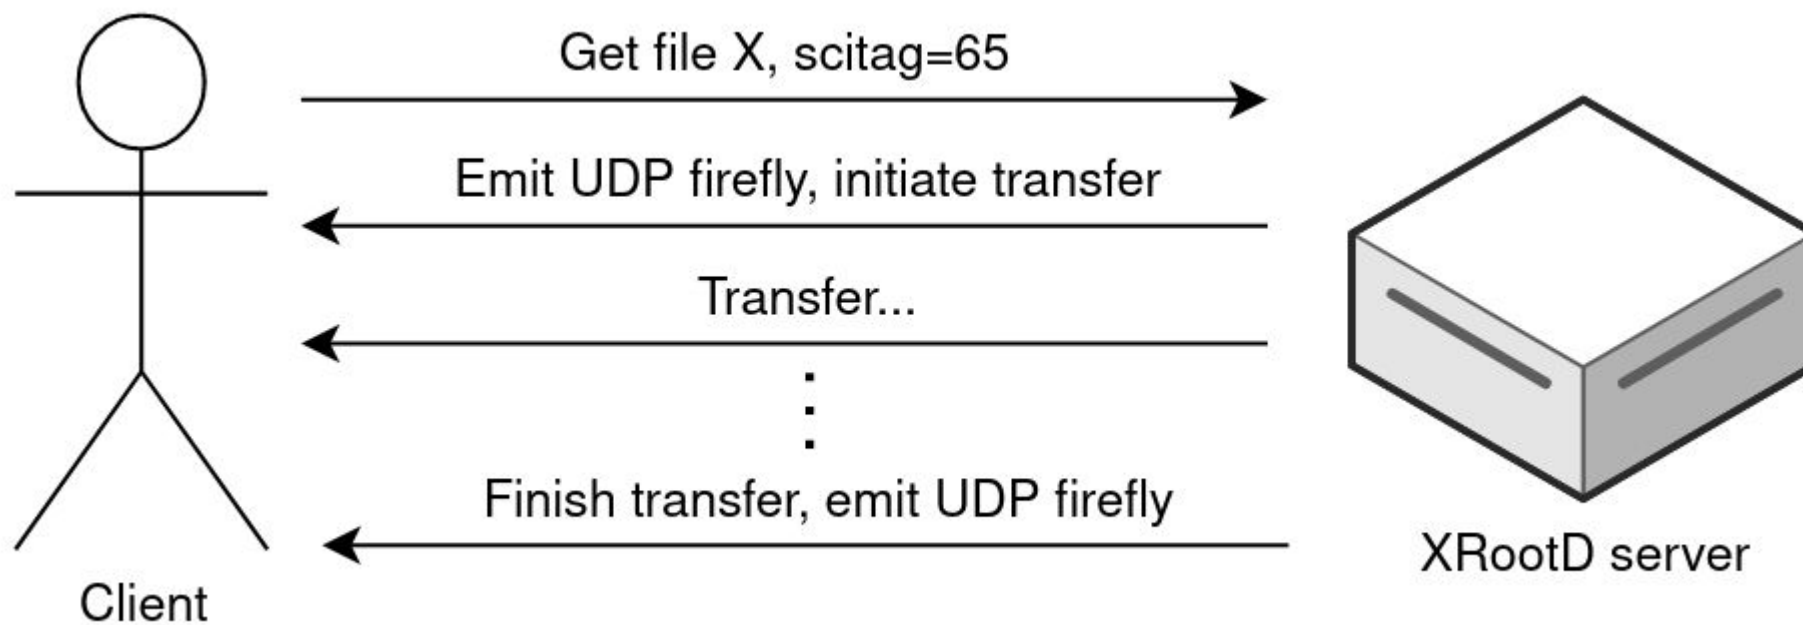
What is a SciTag?

- Provided by the user
- SciTag = $\text{experimentID} \ll 6 \mid \text{activityID}$
 - $\text{experimentID} = \text{SciTag} \gg 6$
 - $\text{activityID} = \text{SciTag} \& 0x3F$
- Mapping
 - $\text{experimentID} \leftrightarrow \text{experiment name}$
 - $\text{activityID} \leftrightarrow \text{activity name}$
 - Can be found here: <https://www.scitags.org/api.json>

```
{
  "expName": "atlas",
  "expId": 2,
  "activities": [
    {
      "activityName": "default",
      "activityId": 1
    },
    {
      "activityName": "perfSONAR",
      "activityId": 2
    },
    {
      "activityName": "Data Brokering",
      "activityId": 3
    },
    {
      "activityName": "Data Consolidation",
      "activityId": 4
    }
  ]
}
```

SciTags

Workflow



SciTags

UDP firefly

- It's a UDP packet emitted at the beginning and at the end of a transfer by the XrootD server (per socket)
 - Received by a configured endpoint!

SciTags

UDP firefly

- It's a UDP packet emitted at the beginning and at the end of a transfer by the XrootD server
 - Received by a configured endpoint!

UDP firefly payload

- syslog facility header with severity: Informational: 6 and facility: Local0: 16
- JSON document that conforms to the JSON firefly schema described here:
<https://www.scitags.org/schemas/v1.0.0/firefly.schema.json>

SciTags

UDP firefly JSON payload

```
{
  "version": 1,
  "flow-lifecycle": {
    "state": "end",
    "current-time": "2024-02-20T16:13:56.626818+00:00",
    "start-time": "2024-02-20T16:13:55.507480+00:00",
    "end-time": "2024-02-20T16:13:56.626818+00:00"
  },
  "usage": {
    "received": 117441692,
    "sent": 1174
  },
  "netlink": {
    "rtt": 0.665
  },
  "context": {
    "experiment-id": 1,
    "activity-id": 1,
    "application": "http"
  },
  "flow-id": {
    "afi": "ipv6",
    "src-ip": "XXXX:XXXX:d00:16::18a",
    "dst-ip": "YYYY:YYYY:d00:16::18a",
    "protocol": "tcp",
    "src-port": 57284,
    "dst-port": 2001
  }
}
```

SciTags

Deployment

- Configuration
 - On the FST - the MGM is metadata only!

```
FST configuration
xrootd.pmark use firefly scitag
xrootd.pmark domain any
xrootd.pmark ffdest firefly-endpoint.cern.ch:10514
```

SciTags

Tests

- Deployed on the EOS CMS FSTs before the WLCG data challenge 2024
 - UDP fireflies emitted toward an endpoint in es.net

Total Flows per Exp/Act




Conclusion

Different HTTP improvements were done over the last year

- Chunk upload bug fix, checksum handling, range read requests handling

SciTags and UDP fireflies have been implemented in both EOS and XRootD

- A SciTag is unsigned 16 bits integer $65 \leq \text{SciTag} \leq 65535$ provided by the client
- A UDP firefly is emitted at the beginning and the end of each transfer
- Allow to see how the network is being used by the scientific community



home.cern