
FCAST: scalable object delivery for the ALC and NORM protocols

draft-roca-rmt-newfcast-02

IETF 72 – Dublin, July 2008

Vincent Roca (INRIA)

Brian Adamson (NRL)



Outline

- situation/status
- design principles
- summary of the technical choices made
- next steps

Status


- what has been done so far?
 - initial I-D in March 2007
 - a call for volunteers
 - a few discussions on the mailing list till IETF'71
 - technical work began
- current version -02 lays the new foundations
 - as promised during IETF'71
 - still inspired from <draft-ietf-rmt-pi-alc-00.txt> (March 2000)...

Design principles

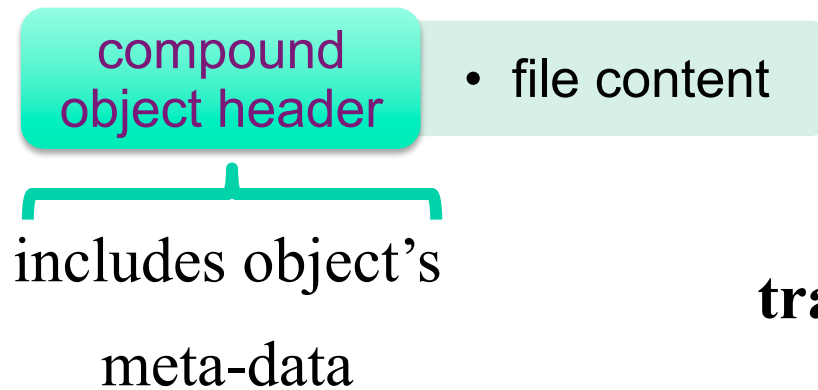
- “Keep It Simple and Stupid”
 - FCAST should remain easy to implement and lightweight to process
- don't try to address all possible use-cases
 - FCAST is perhaps more specialized than FLUTE
- design a solution for both ALC/LCT and NORM
 - FCAST considers the specificities of each protocol
 - FCAST/ALC and FCAST/NORM are almost the same...

Technical choices made

file + meta-data to send



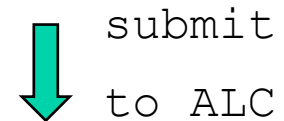
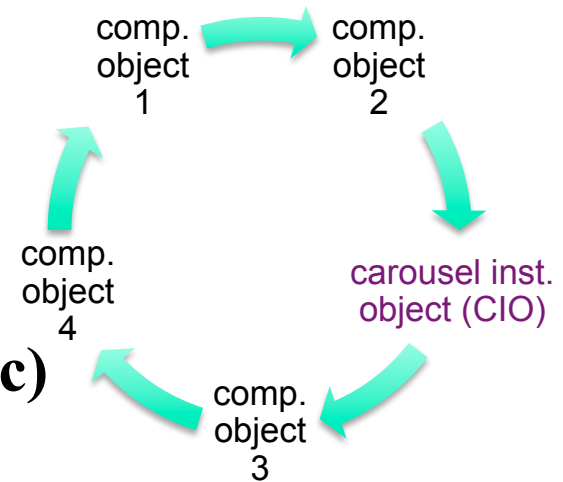
compound object



insert in the carousel



transmission carousel
(dynamic or static)



- all the files are sent as compound objects (no exception)!

Technical choices made... (cont')

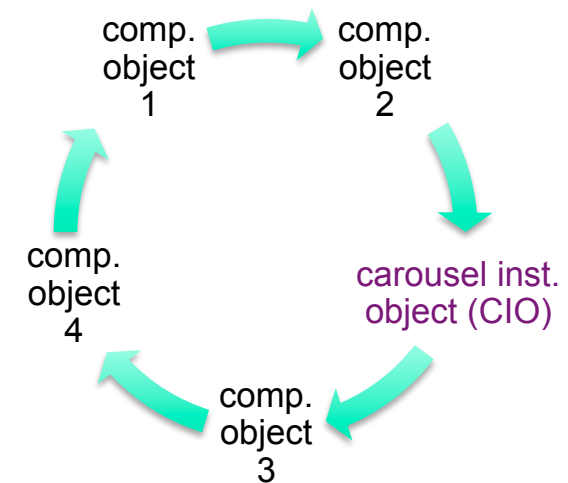
- carousel instance (CI)

- this is a fixed set of compound objects sent during a certain number of transmission cycles

- the user says:

*-here are the objects I want FCAST to send
-send them during three cycles*

Default: send each object once per cycle (optionally, send some of them more often for improved robustness).



- encompasses transmissions in

- PUSH mode: **single** transmission cycle
- ON-DEMAND mode: **high number** of transmission cycles

Technical choices made... (cont')

- carousel instance **object** (CIO)

- this is a specific **compound** object type (l==1)

- goal is to **list** (≠ describe) the content of the CI:

carried
as CIO
content



- list of TOIs, e.g.: 1, 2, 3, 100-104, 200-203, 299

carried
as CIO
meta-data



- "Fcast_CIO_ID: value" identifies the instance
- "Content-Encoding: gzip" indicates the TOI list is gzip'ed
- "Fcast_CIO_complete: 1" indicates that no new objects will be sent (optional)

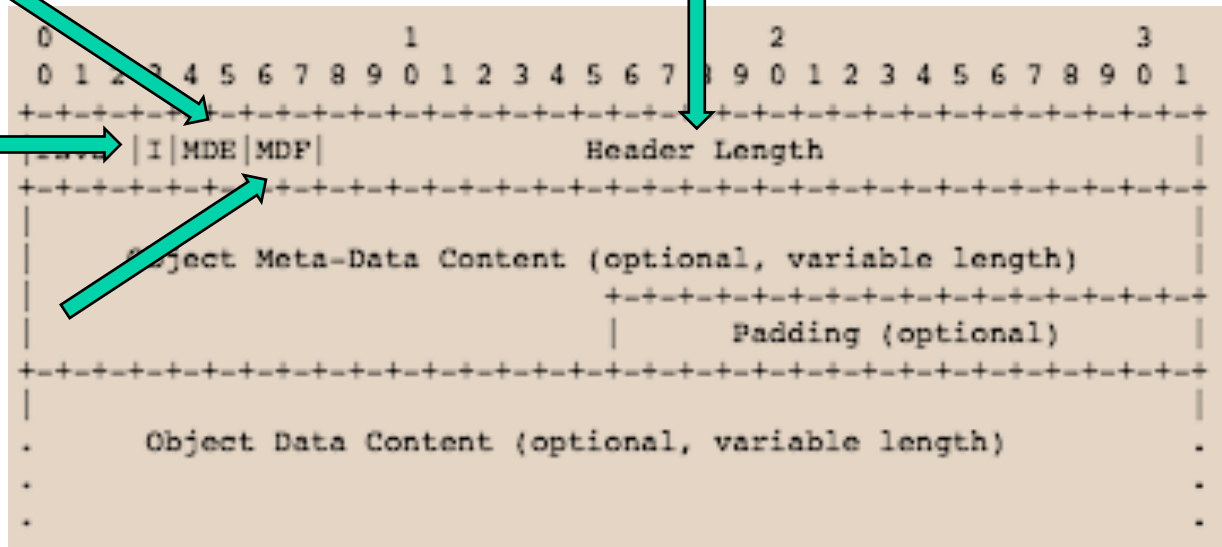
Technical choices made... (cont')

- compound object format

meta-data *encoding*
(e.g., gzip'ed)

up to $2^{24} - 4$ long meta-data field

I==1 with a CIO



meta-data *format*
(currently HTTP1.1)

Figure 1: Compound Object Header with Meta-Data

Technical choices made... (cont')

- HTTP1.1 meta-data format example:

```
Content-Location: example.txt <CR-LF>
```

- Other formats (e.g., XML) might be added in future
- header **prepended** rather than appended
 - with **NORM** or **ALC** in **PUSH** mode, enables a receiver to quickly get the header, process it, and decide to continue download or not
 - **useless with ALC in on-demand mode**
 - ... **at the price of a slightly more complex processing**
 - e.g., requires a scattering read method, like `recvmsg()` does

Technical choices made... (cont')

- several variants are possible W.R.T. meta-data:

1. send meta-data **out-of-band**, totally or partially

- compound object header is probably shorter (but ≥ 4 bytes)
- enables incoming packets filtering based on user's preferences

2. NORM: send meta-data in **NORM_INFO** messages, totally or partially

- idem...

{Not in -02... Will be added!}

Technical choices made... (cont')

- TOI considerations

- **ALC**: very **flexible** (several sizes, flexible management)
- **NORM**: more **directive** (16 bits, managed sequentially)
- **FCAST**: no reserved TOI for the CIO (\neq FLUTE)
 - **it's the receiver FCAST's role to identify it's a CIO, process it, and decide what to do...**
- **consequence:**

Since a receiver doesn't know in advance which TOI will be used for the following CIO, he **MUST NOT** filter out packets that are not in the CIO's TOI list

{Not in -02... Will be added!}

Technical choices made... (cont')

- sending a CIO with an empty TOI list
 - **useful to tell receivers that:**
 - **all the previously sent objects have been removed from the carousel**
 - comes in addition to LCT's B flag for improved robustness
 - **there is currently nothing new but the session is active**
 - provides a "heartbeat"
 - implicitly: new objects may be sent in the future, if the session is not "Complete"

{Not in -02... Will be added!}

What's next?

- finalize the technical content and proof-read it
 - **feasible for next IETF!**
- implement/test it
- and it will be done with FCAST

...it's perhaps a bit optimistic ;-)