

Storage Events: distributed users, federation and beyond

Paul Millar

(on behalf of the dCache team)

CHEP 2018 at Sophia, Bulgaria; 2018-07-09
<https://indico.cern.ch/event/587955/>



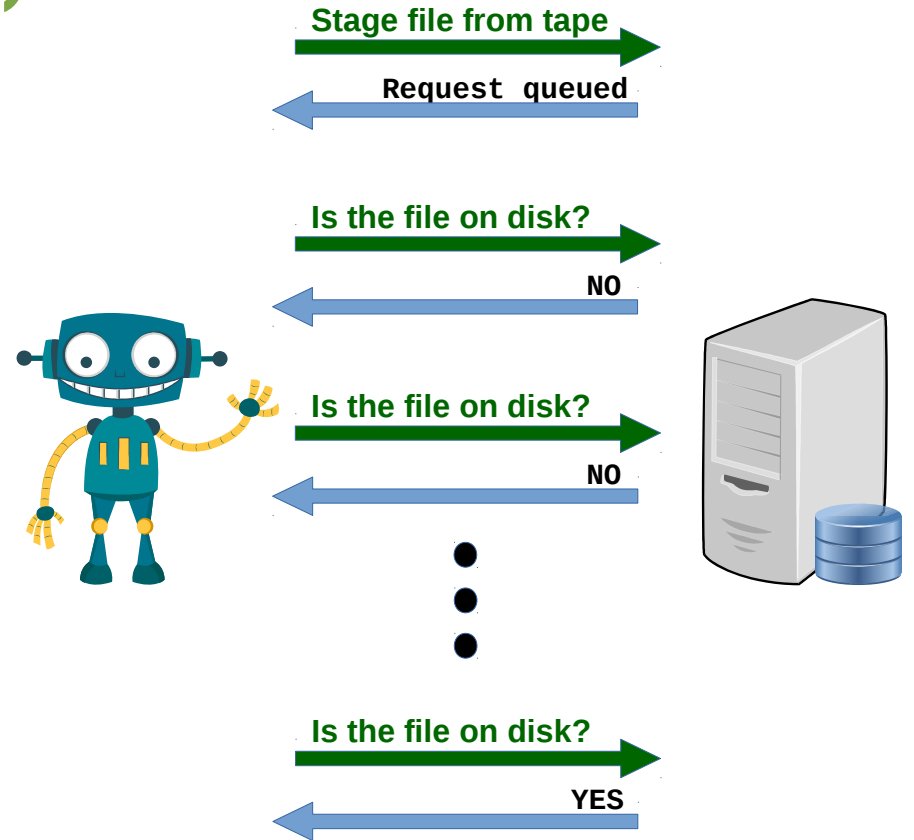
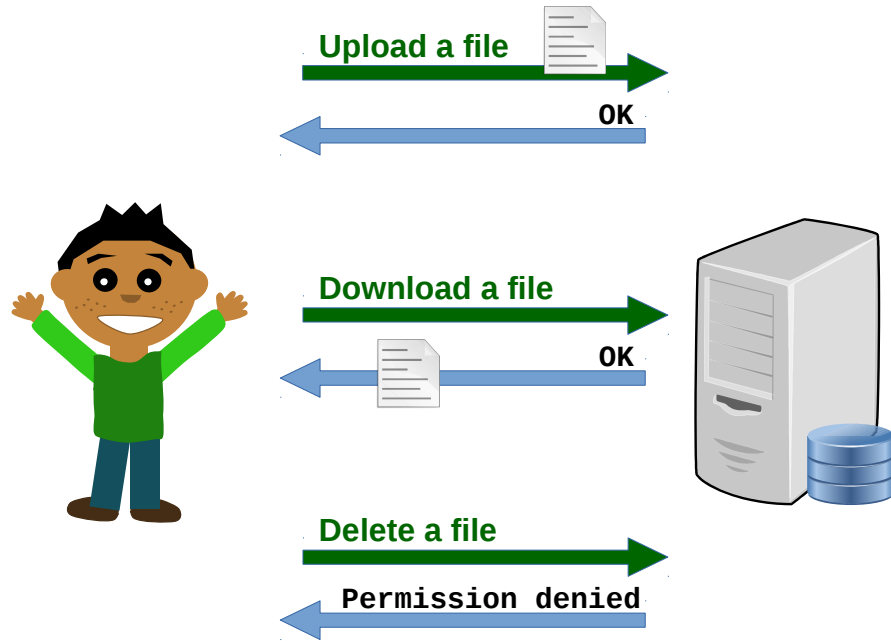
Nordic e-Infrastructure
Collaboration



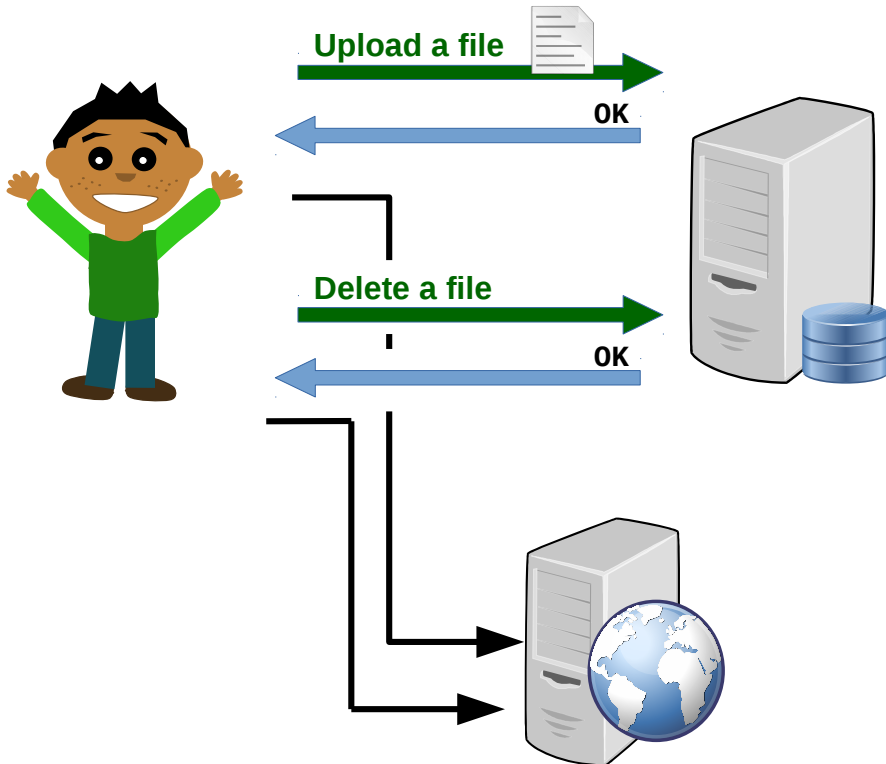
eXtreme DataCloud



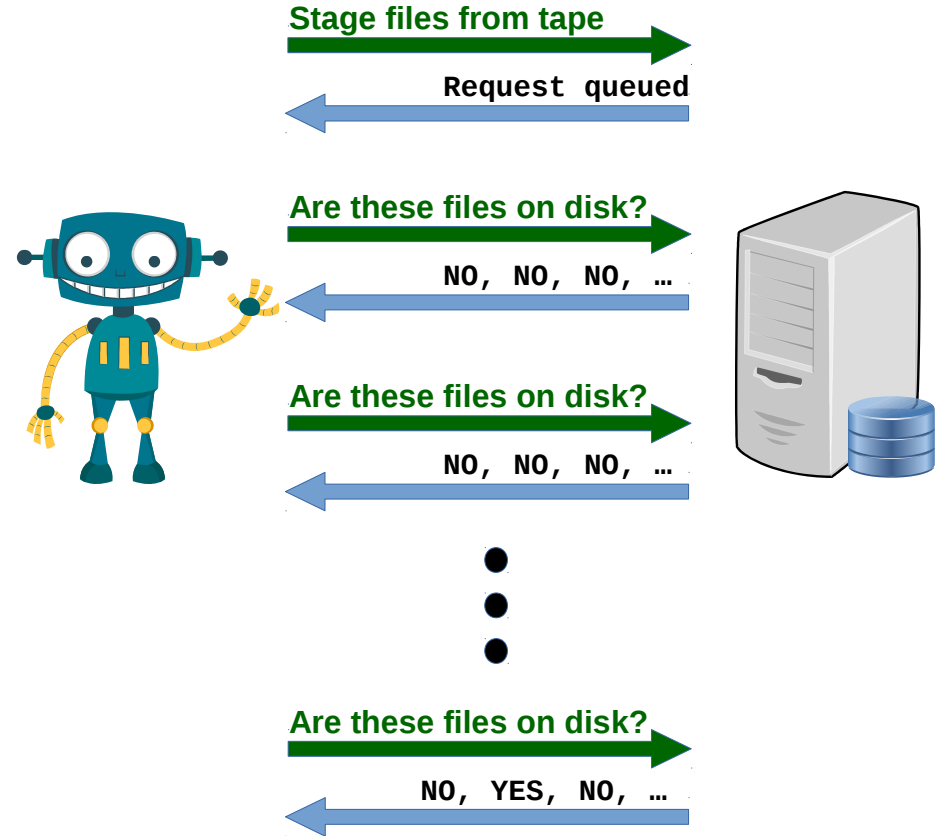
How storage is used (currently)



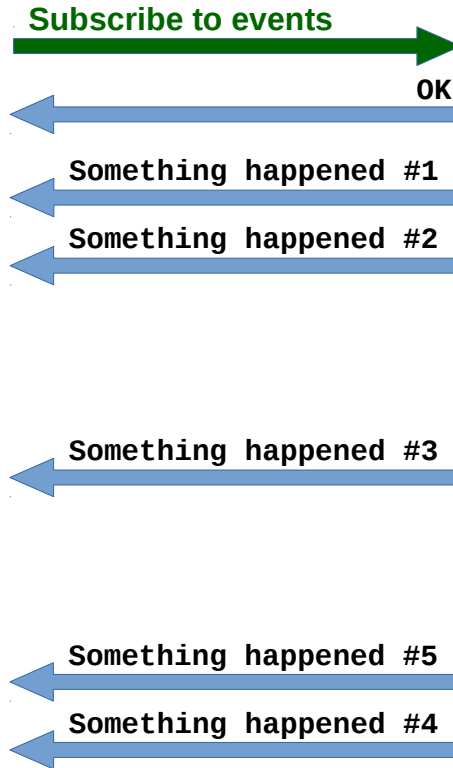

The problems...




Rucio/LFC/...





New way of interacting with storage: storage events



New solutions to old problems:

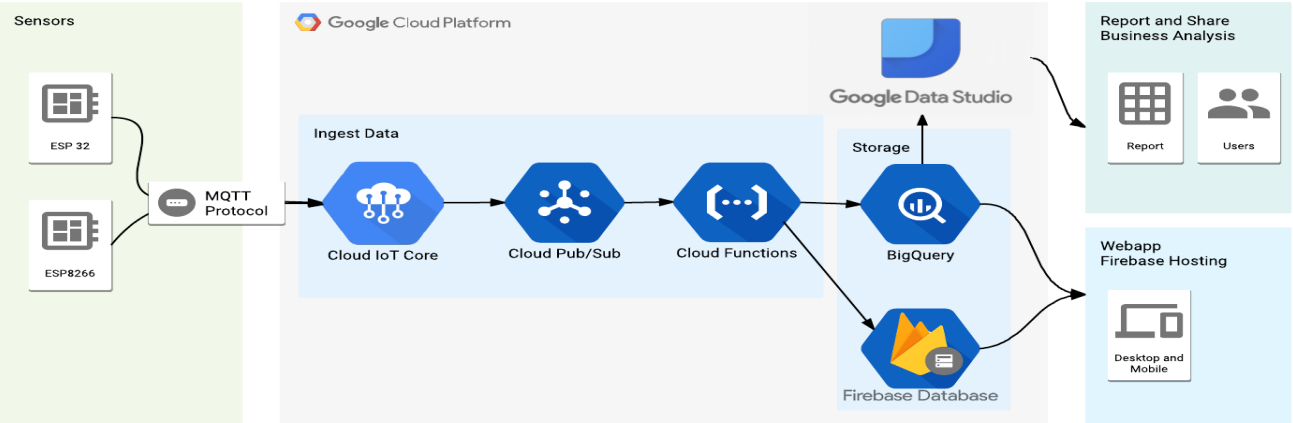



New solutions to old problems:



Comparison: it's what industry is doing...

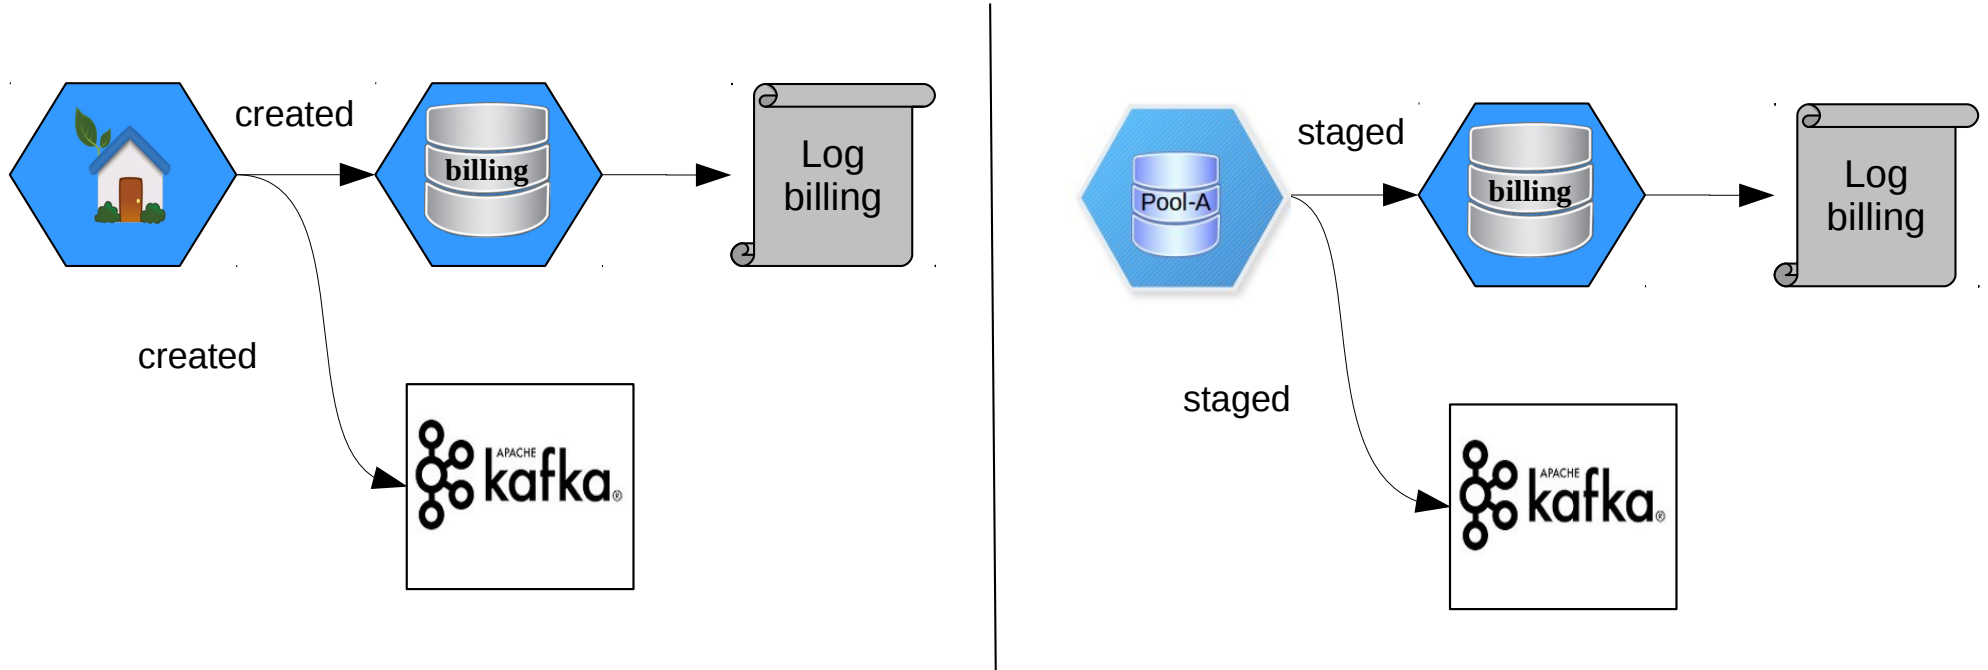
Example: Image Thumbnail Creation



Comparison: it's what Open-Source is doing...



dCache Storage Events: Kafka





dCache Server-Sent Events (SSE)

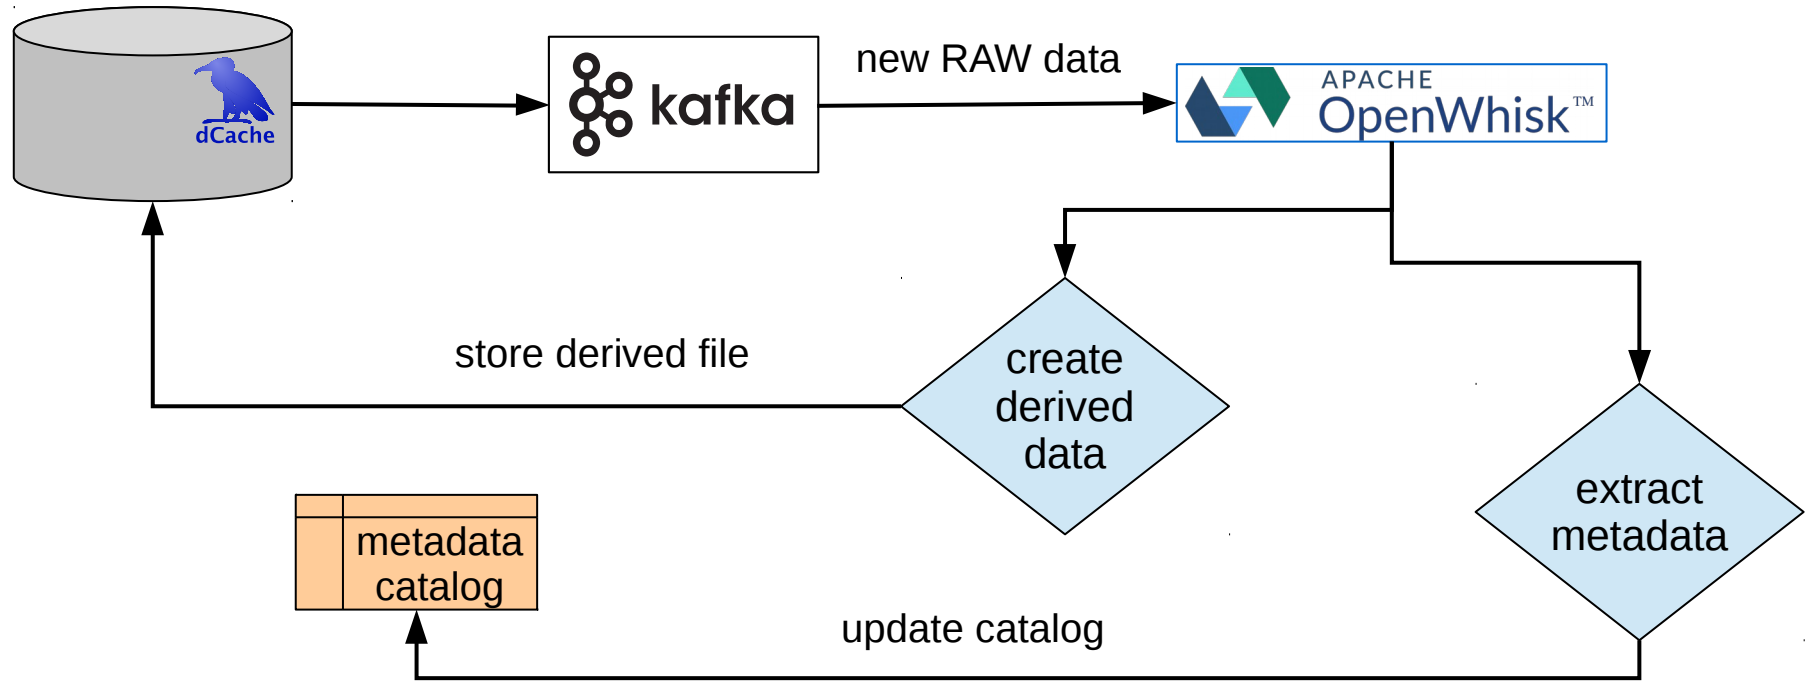
- Based on HTTP v1.1
 - HTML 5 **standard**
 - Support for many languages and web-browsers
 - Initially adding support for **inotify** events
 - (it's how Linux does namespace notification)
 - Plan to add:
 - Locality change notification: flush, stage, ...
 - Transfer-related events
 - QoS changes
-



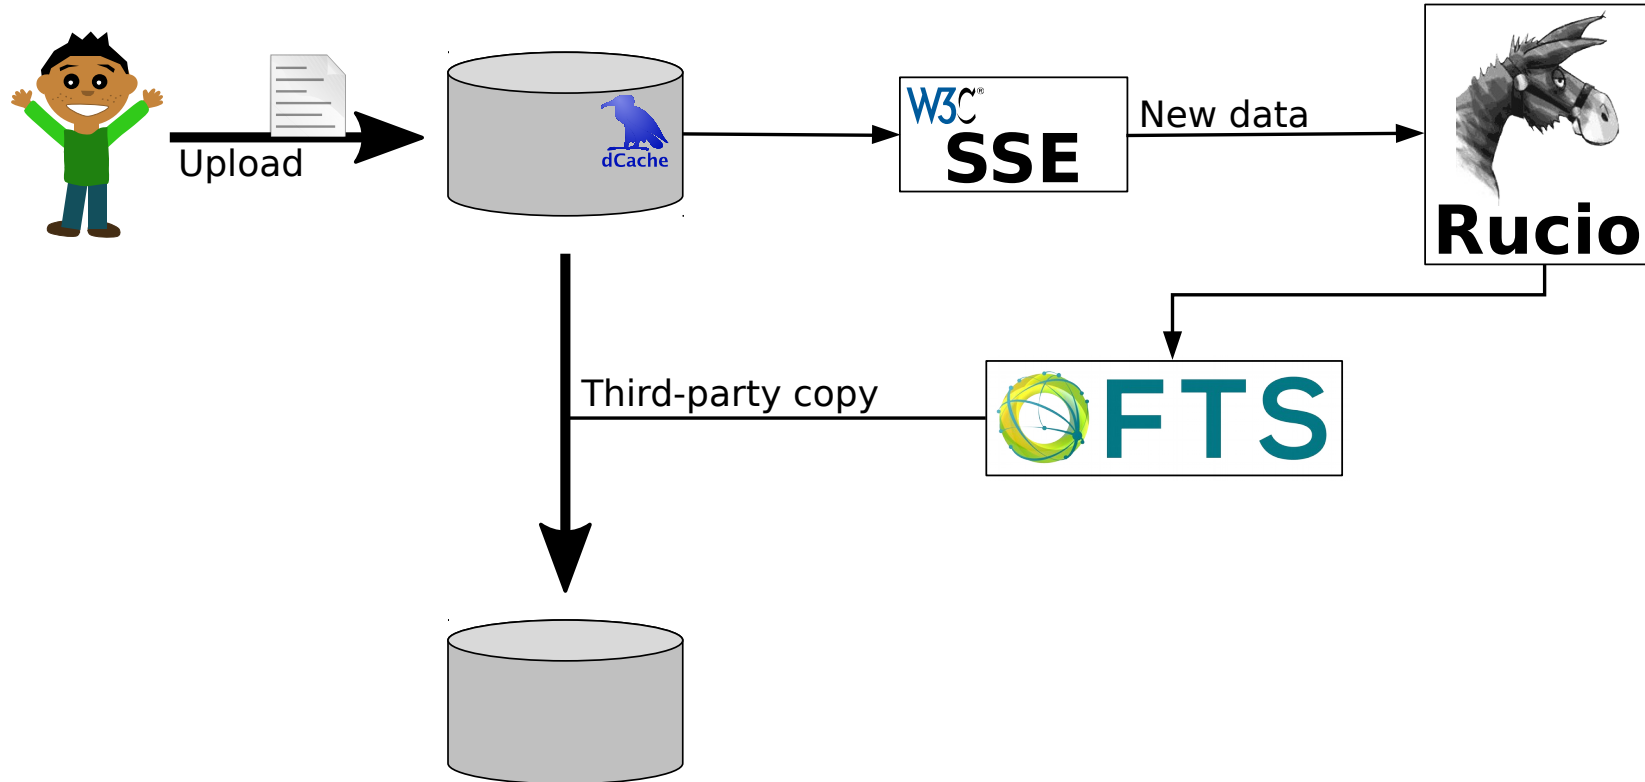
Cheat sheet: Kafka vs SSE

	 kafka	 SSE
Standard ...	Component	Protocol
What events does it see?	dCache internal events	Controlled
Main benefit	Easy integration	Built-in security
“Catch-up” storage	Memory & disk	Memory-only (currently)
Target audience	Site-level integration	Events for users

EOSC-Pilot demonstrator: EU-XFEL data ingest



Rucio demonstrator: automated replication with SSE



Future directions

- Complete SSE **inotify** support in dCache.
 - Add **additional events**, based on initial feedback.
 - Further explore **automated data workflow** (EU-XFEL usecase).
 - Work with **Rucio team** to explore SSE integration.
 - Work with dCache sites to deploy store events in **production**.
-

Conclusions

- We have presented **storage events**:
 - Notification of changes within the storage service
 - Several use-cases have been shown that illustrate potential **benefits to HEP**.
 - dCache provides two complimentary event delivery mechanisms: **Kafka** and **SSE**.
 - Work is on-going:
 - New event types to be added
 - Work with sites to bring this into production.
-

Thanks for listening!

