

CDNI Logging Extensions

Ben Rosenblum
CDNI WG, IETF 117
July 26th, 2023

Motivation

- RFC 7937 supports a limited set of log fields in a format which is not wholly supportive of the current needs of CDN operators and content providers
- In addition to new record field definitions, requirements exist for:
 - Common transport mechanisms, both push and pull, batch and streaming
 - Record formats and log file formats in JSON and CSV
 - Log file archives
 - More extensive log file metadata, configuration of log file naming
 - Field obfuscation, truncation, and other transformation

Log records

- Define a common set of fields with clear definitions that meet the current needs of participants
- Specify standard sets of fields and log formats (record types) that cover various needs: billing (minimal), operations, troubleshooting.
- Allow optional inclusion of arbitrary HTTP headers if permitted by dCDN capabilities
- Allow advertisement and configuration of expectations for log file publishing intervals, naming, other metadata (log source IDs, group IDs, etc..)

Containers

- Move beyond the ELF format defined in RFC 7937
- File formats (JSON, protobuf) that are in line with what participants are using in production today
- Log archives (tar) for transmission of multiple files
- Advertisement objects and configuration metadata allow control of the logging container (size, time range)

Transports

- Define objects to allow advertisement and configuration of standard file and streaming transports as an addition to the Atom-based index format of RFC 7937
 - S3
 - SFTP
 - Kafka
- Add endpoint discovery for RFC 7937 feeds
- Provide for both push (dCDN->uCDN) and fetch (uCDN->dCDN)

Transforms

- Permit control of transformations to accommodate privacy, security or other requirements for obfuscation of log data on a field-by-field basis
- Support both advertisement (inform uCDN of transformations to log data) and configuration metadata (control dCDN generation of log records)
- Allow combinations of transform configuration (e.g., truncate followed by encrypt)

Draft Status

- Working on converting our collaborative document to an IETF draft, didn't make it in time for IETF 117
- Might need to split this into multiple drafts (currently ~70 pages in Google Docs)
- Other aspects still need work will have to be covered by a future draft:
 - Filtering
 - Aggregation
 - Sampling
 - Custom fields, types, and format