

Structured logging for improved user experience with minimum additional infrastructure

Martin Cejp


CERN SY-EPC-CCS

(Electrical Power Converters group / Converter Controls Software)

CERN School of Computing 2022, Kraków, Poland

Background

- Electromagnets to steer/deflect beam → driven with DC or AC current → need for electrical power converters
 - Service for remote firmware upgrade
 - Non-interactive: controlled through SQL DB
 - Bulk deployments
 - 5000+ converters x ~5 programmable cards each x 1-2 campaigns per year
 - Limited time for upgrade campaigns
- Essential to have clear and precise feedback from updater service to power engineer




Firmware build

Configuration
database

Updater service

Modular power converter



Status feedback

Power engineer



Plaintext vs structured logging

- Text logs with fixed format (timestamp, severity, module, message)

```
[2021-10-06 13:05:55] [   INFO](area_worker): FgcWorker(Test_Validation4): job RPAA0.866.02.ETH8 removed from tasks
[2021-10-06 13:06:00] [  DEBUG](area_worker): RPAA0.866.02.ETH8 not in TODO job list
[2021-10-06 13:06:00] [   INFO](area_worker): (Test_Validation) job RPAA0.866.02.ETH8 added to queue
[2021-10-06 13:06:00] [  DEBUG](fgc_job_utils): RPAA0.866.02.ETH8: expected_data;
    slot: 2
    board: VS_STATE_CTRL
    devices:
        DB -> Device(Device='DB', Variant='DOWNLDBOOT_3', Var_Rev='303', API_Rev='303')
        MF -> Device(Device='MF', Variant='LPS4Q_130', Var_Rev='304', API_Rev='304')
    slot: 5
    board: VS_REG_DSP
    devices:
        DB -> Device(Device='DB', Variant='DOWNLDBOOT_3', Var_Rev='304', API_Rev='304')
        DEVICE_2 -> Device(Device='DEVICE_2', Variant='LPS4Q_130', Var_Rev='301', API_Rev='301')
        MF -> Device(Device='MF', Variant='LPS4Q_130', Var_Rev='301', API_Rev='301')
[2021-10-06 13:06:00] [  ERROR](area_worker): FgcWorker(Test_Validation4): failed to reprogram RPAA0.866.02.ETH8: err 39 not ready
[2021-10-06 13:06:00] [   INFO](area_worker): FgcWorker(Test_Validation4): job RPAA0.866.02.ETH8 removed from tasks
```

- Cheap to add – `logger.warning("Reply received %d seconds too late", delay)`
- User-unfriendly (remember: power engineer \neq controls expert)
- Logs usually stored somewhere on NFS or accessed via heavy-weight tool like Kibana
- **Becomes impossible to follow when multiple tasks are processed in parallel**

Plaintext vs structured logging

- Structured

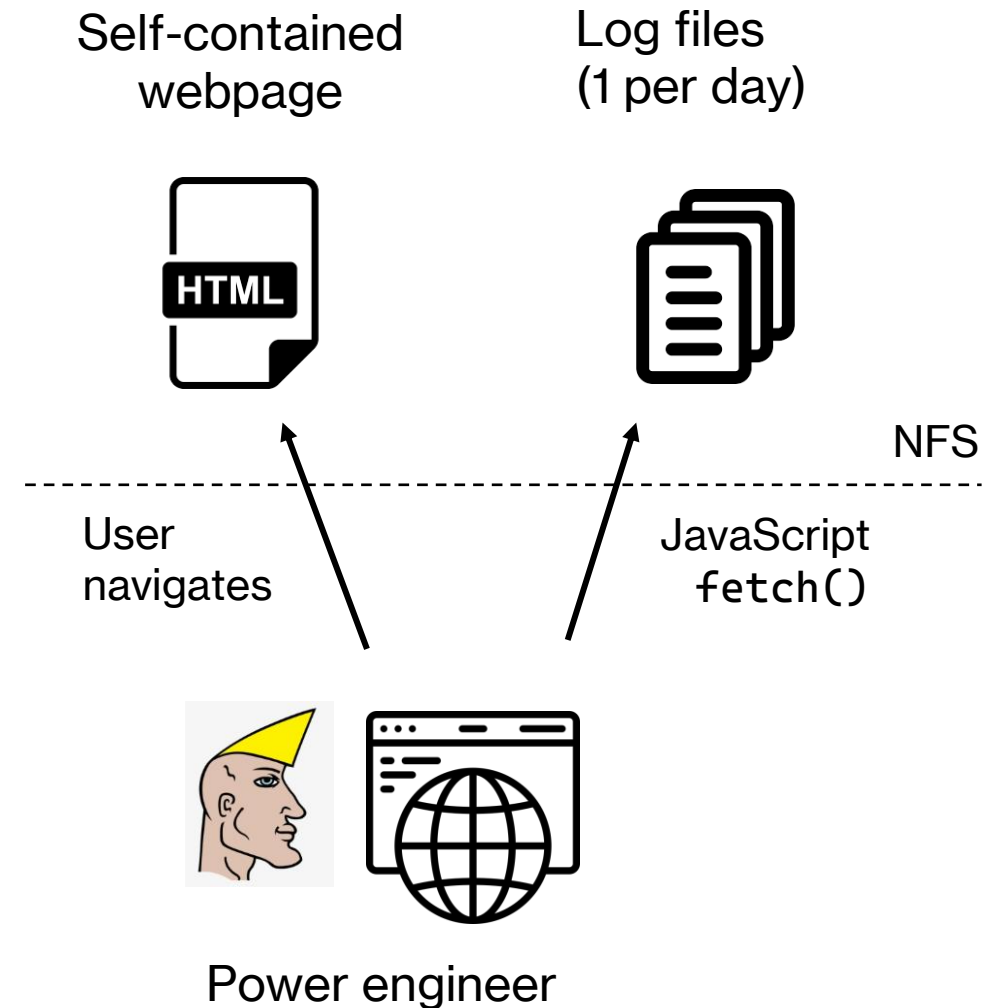
```
{“timestamp”: “2021-10-06T13:05:55”,  
  “severity”: “INFO”,  
  “message”: “job RPAA0.866.02.ETH8 removed from tasks”,  
  “type”: “job_removed”,  
  <message-specific attributes> } additional attributes depend on entry type  
}
```

- 1 entry per line -> “JSONL” format
 - Retains some advantages of plain text: new entries can be simply appended at end (contrast with classic JSON)
 - Rotate files same way as text logs (daily/hourly)
- Admittedly even worse to read than plaintext, *but*
- Easy to parse + transform → we can provide a domain-specific, intuitive view to users

How to display it?

- Custom web app – Flask/Node/Go/Java/...?
- ElasticSearch–Logstash–Kibana?
- Store & retrieve from a SQL database?
- All too complex! Our goal is zero infrastructure
- Take advantage of existing NFS*→HTTP bridge
 - Use to host both webpage + logs
 - Self-contained HTML page, JSONL data loaded by client-side script

* NFS = Network File System. Imagine a more primitive, 1990s version of AFS/EOS/Ceph, used heavily in the accelerator sector



What it looks like

Mon Feb 14 2022

Most recent first, entries grouped by system

2022-02-14 11:21:00
(44.6 seconds)

RPAEH.A7.1_010.2.TEST1

Status	Slot	Board	Device	Firmware in DB		Firmware in device	
✓ Up-to-date	2	VS_STATE_CTRL	DB	DOWNLDBOOT_3	rev 303	DOWNLDBOOT_3	rev 303
✓ Up-to-date	2	VS_STATE_CTRL	MF	LPS4Q_130	rev 306	LPS4Q_130	rev 306
✓ Up-to-date	5	VS_REG_DSP	DB	DOWNLDBOOT_3	rev 304	DOWNLDBOOT_3	rev 304
⚙ Programmed	5	VS_REG_DSP	DEVICE_2	LPS4Q_130	rev 311	LPS4Q_130	rev 307→311
⚙ Programmed	5	VS_REG_DSP	MF	LPS4Q_130	rev 310	LPS4Q_130	rev 306→310

2022-02-14 11:16:43
(3.4 seconds)

RPAEH.A7.1_010.2.TEST1

Status	Slot	Board	Device	Firmware in DB		Firmware in device		Messages
✓ Up-to-date	2	VS_STATE_CTRL	DB	DOWNLDBOOT_3	rev 303	DOWNLDBOOT_3	rev 303	
✓ Up-to-date	2	VS_STATE_CTRL	MF	LPS4Q_130	rev 306	LPS4Q_130	rev 306	
⚠ Warning	2	VS_STATE_CTRL	—					DID_NOT_SWITCH_TO_PROI
⚠ Not detected	5	VS_REG_DSP	DB	DOWNLDBOOT_3	rev 304	—		Error codes not explicitly handled by log viewer show up with degraded UX but no information is lost
⚠ Not detected	5	VS_REG_DSP	DEVICE_2	LPS4Q_130	rev 311	—		
⚠ Not detected	5	VS_REG_DSP	MF	LPS4Q_130	rev 310	—		
⚠ Warning	5	VS_REG_DSP	—					EXPECTED_BOARD_NOT_FOL

2022-02-14 11:05:06
(3.3 seconds)

RPAEH.A7.1_010.2.TEST1

Each row corresponds to structured log entry (10+ lines in old plaintext logs)


*Guiding principle:
Highlight changes carried out & potential problems; de-emphasize “background” state*

Data updated every ~15 seconds

Summary

- Treat your users to a nice UX
- Reusing existing infrastructure can reduce maintenance burden, enabling functionality that would otherwise be deemed not worth the complexity
 - This project was a bottom-up initiative motivated by seeing user suffering (+ JIRA tickets) first-hand
 - GitLab CI is another example of an extremely versatile, hosted tool
- Question complexity that we're told to take for granted
 - JS/CSS frameworks
 - Web application platforms
 - Elaborate build tools, microservice architectures, heavyweight databases ...

Power engineers
are empowered



Controls experts can tend
to other business