

ecFlow

Best Practices and Real-life Workflow examples

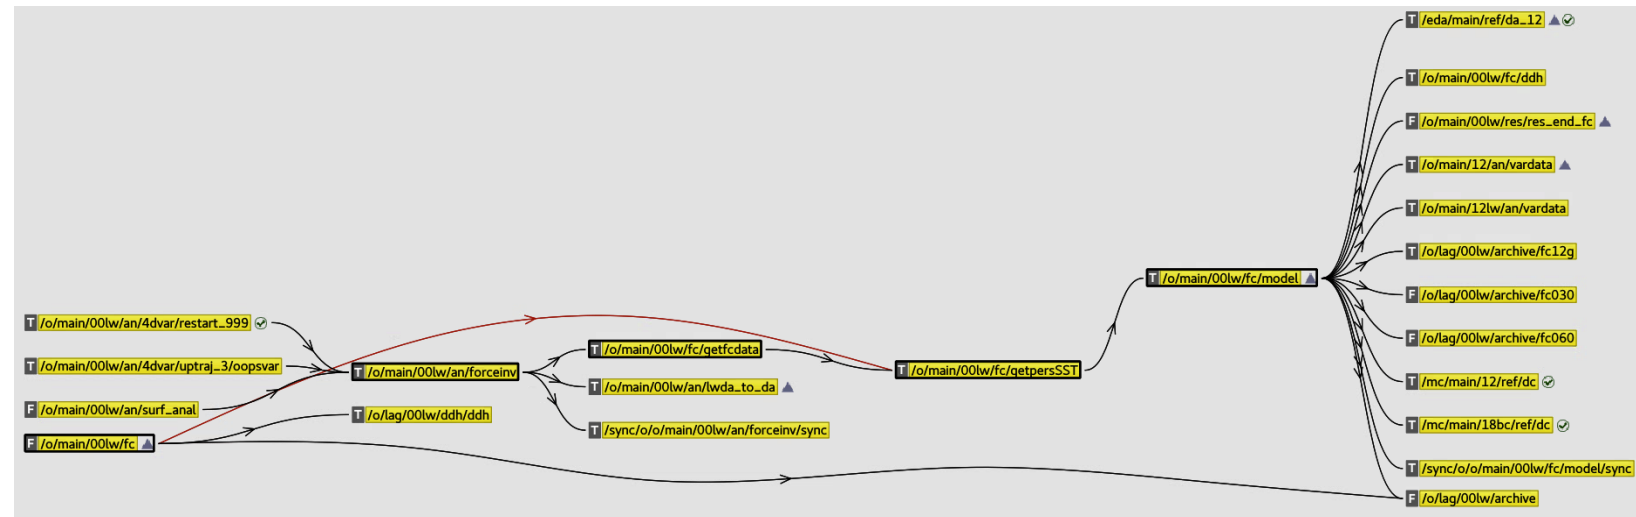
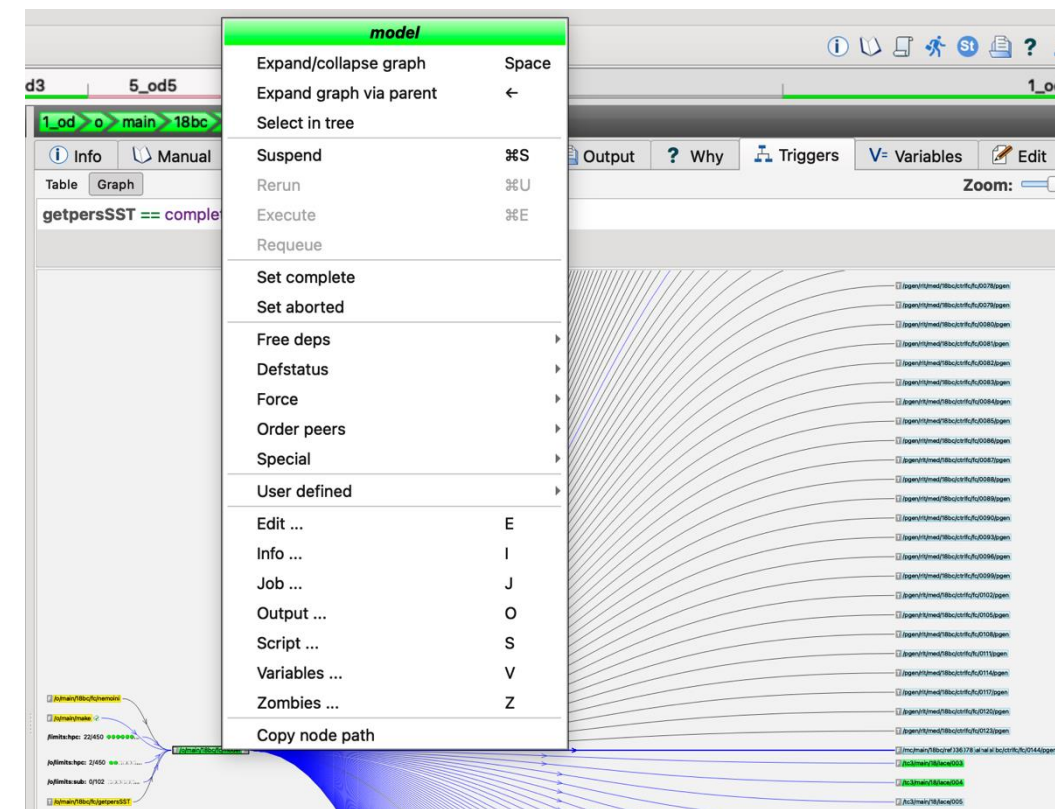

Axel Bonet, Iain Russell, Marcos Bento

ecFlow: Best Practices?

- Aims
 - How ecFlow can be used
 - Show few ecFlow use cases
 - Advertise [suites coding standards](#) manual
 - Hint toward [Pyflow](#) ecosystem
 - [A glossary](#)

- Overview

- Architecture
- Suites design
- Tasks wrappers design
- Monitoring & Observability
- Error Handling
- Security & Access
- Operational Tips




ecFlow best practices – servers' configuration

- From introduction, we know:

- Multiple servers per user is possible on one CPU
- Multiple users can share same CPU for their servers
- Multiple users sharing the same server is possible
- Setup a **whitelist** file to restrict access
 - use **TLS (SSL)** mode where needed,
 - use **password** for best security

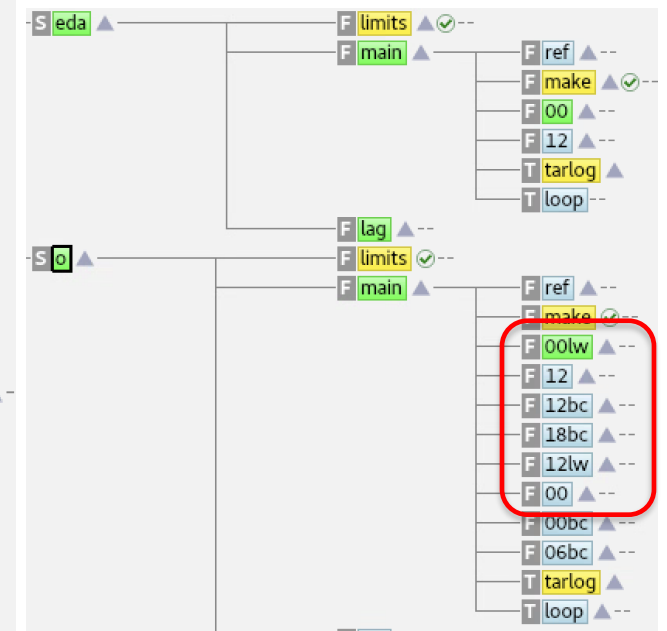
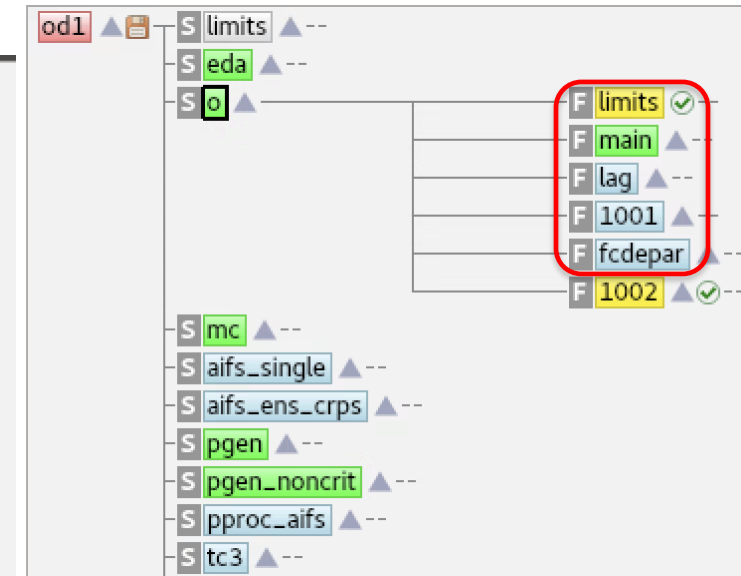
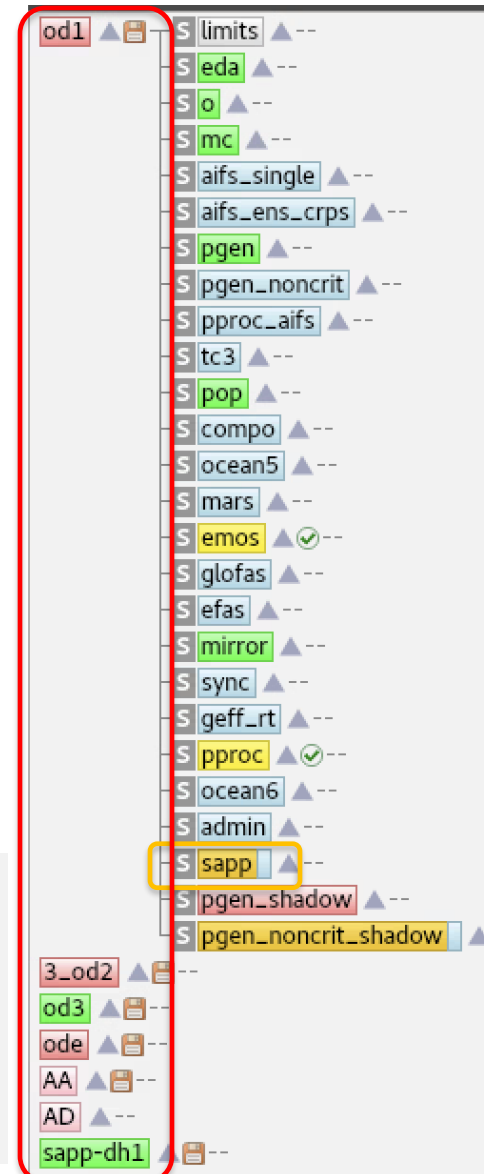
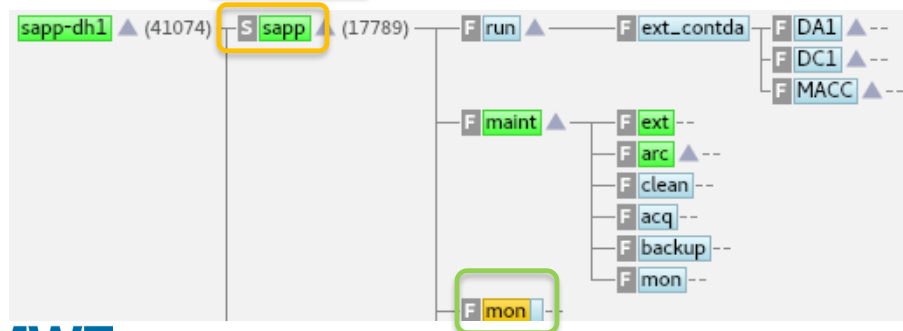
- Commissioning **one VM per server**

- **Jobs files** are local, transients (copied to HPC near job output by troika)
- **Tasks wrappers** on a mounted persistent file system
- **Checkpoint file** and **server log** are local (RD), shared (FSD), or under \$HOME/ecflow_server (users)
- VM and server health under the watch of **observability tools** ([Splunk](#), [Opsview](#), [Etcd](#), [Grafana](#), ...)
- Server **automatic restart** (RD, users) vs server restart in **halted mode** (production)
- A **logserver** is part of ecFlow sources, to run on the target host, to make output visible to the GUI
- UDP server on the same VM, REST/API on the same VM with memory requirements



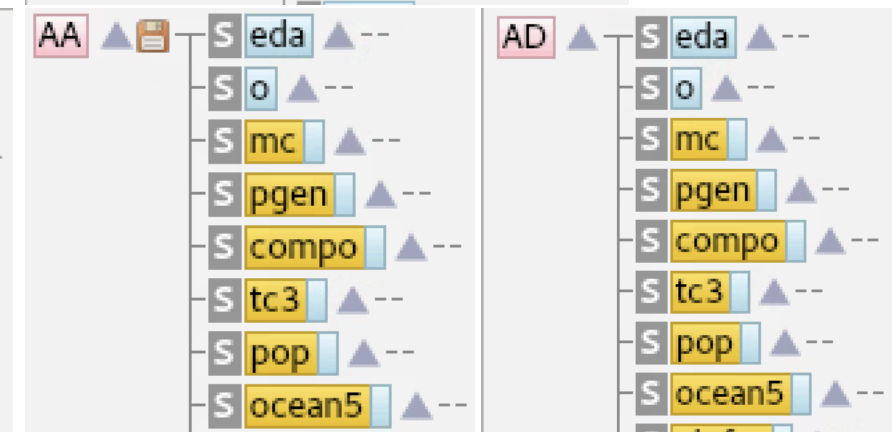
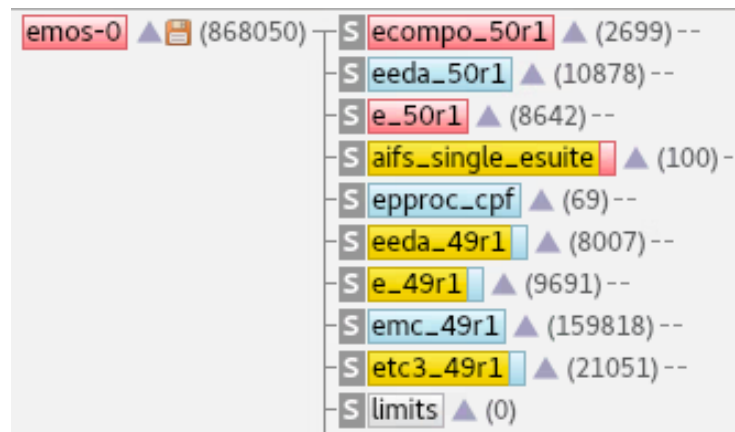

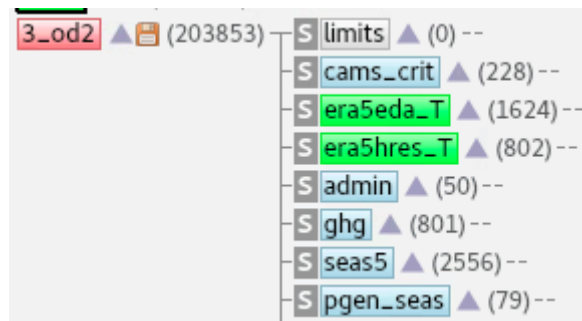
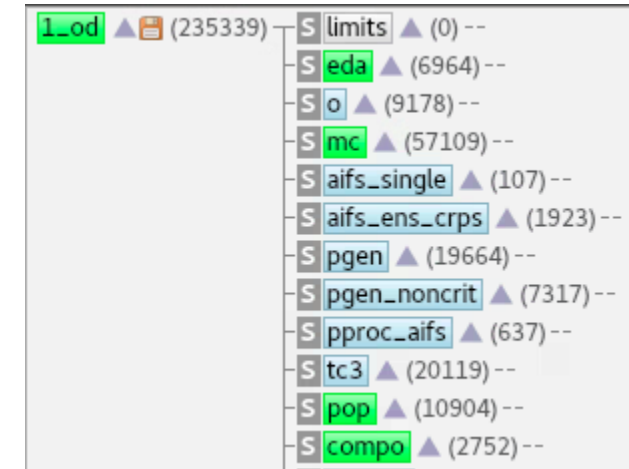
ecFlow in Operation: EMOS operational servers

- EMOS servers:
 - Display criticality: 1_od, 3_od2, 5_od3, 9_ode
- Suite structure reflects on-call criticality:
 - Main: critical path
 - Lag: archive, slow postprocessing
 - Other postprocessing families
 - Inner vs outer watchdog
- Suspended suites may be lively entities:
 - ECF_PASS: FREE, monitoring mode
 - Mirror suite: reflecting a suite on another server



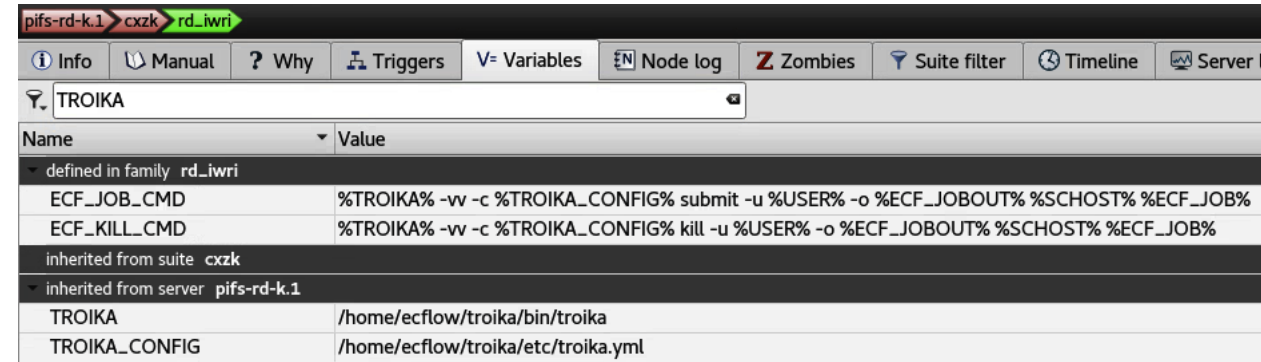
ecFlow best practices – servers' configuration - management

- A federation of servers working together:
 - visual **hierarchy**: top has priority
 - **load balance suites**:
 - od1 (daily run) od2 (monthly) od3 (special projects)
 - Ode (aka emos-0) for e-suites
 - Main and **backup servers** od1, AA, AD

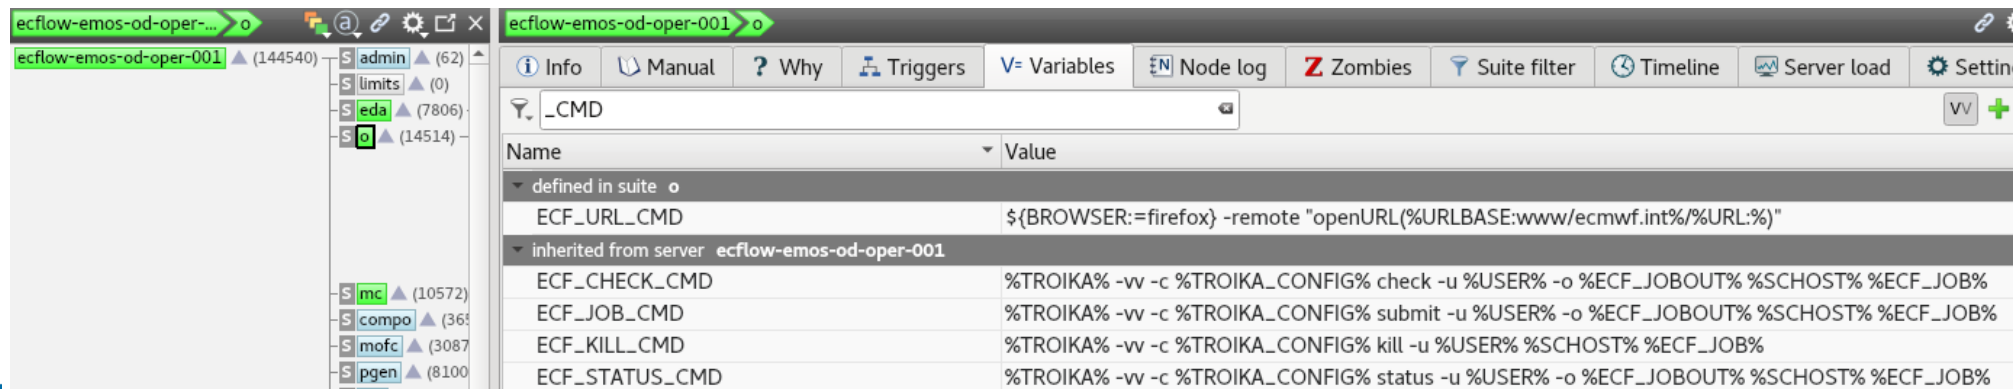



ecFlow: best practices - troika, a fine-tuning jobs submitter

- Troika is open-source, developed at ECMWF
- A system description with a Yaml file
- To Interact with remote queueing system
- Extra jobs tuning (MEM, THREADS, NPES)
- Run hooks (pre / post action)
- Allow deterministic + load balancing submit
- Troika is used in FD/RD/CD/MS workflows
- Extensible: connections (ssh, local), queuing system (Slurm, PBS, ...), hooks
- <https://github.com/ecmwf/troika>



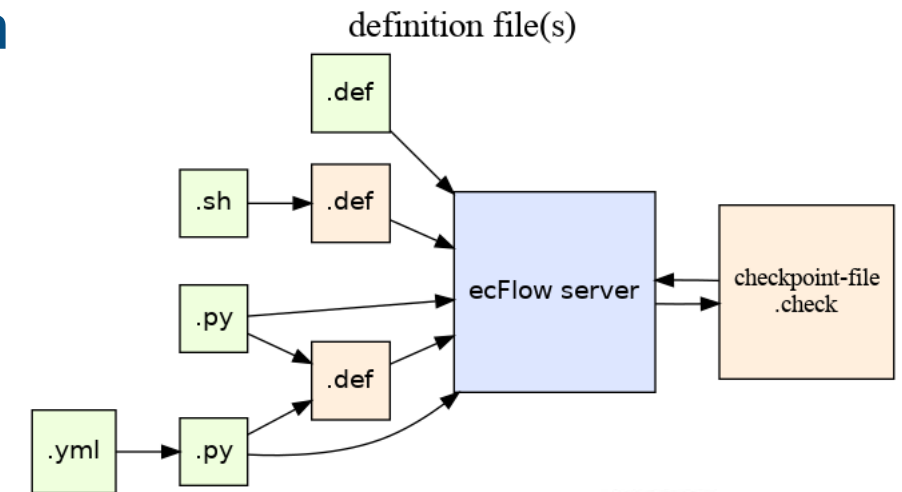
Name	Value
defined in family rd_iwri	
ECF_JOB_CMD	%TROIKA% -vv -c %TROIKA_CONFIG% submit -u %USER% -o %ECF_JOBOUT% %SCHOST% %ECF_JOB%
ECF_KILL_CMD	%TROIKA% -vv -c %TROIKA_CONFIG% kill -u %USER% -o %ECF_JOBOUT% %SCHOST% %ECF_JOB%
inherited from suite cxzk	
inherited from server pifs-rd-k.1	
TROIKA	/home/ecflow/troika/bin/troika
TROIKA_CONFIG	/home/ecflow/troika/etc/troika.yml



Name	Value
defined in suite o	
ECF_URL_CMD	\${BROWSER:=firefox} -remote "openURL(%URLBASE:www.ecmwf.int%/%URL:%)"
inherited from server ecflow-emos-od-oper-001	
ECF_CHECK_CMD	%TROIKA% -vv -c %TROIKA_CONFIG% check -u %USER% -o %ECF_JOBOUT% %SCHOST% %ECF_JOB%
ECF_JOB_CMD	%TROIKA% -vv -c %TROIKA_CONFIG% submit -u %USER% -o %ECF_JOBOUT% %SCHOST% %ECF_JOB%
ECF_KILL_CMD	%TROIKA% -vv -c %TROIKA_CONFIG% kill -u %USER% %SCHOST% %ECF_JOB%
ECF_STATUS_CMD	%TROIKA% -vv -c %TROIKA_CONFIG% status -u %USER% -o %ECF_JOBOUT% %SCHOST% %ECF_JOB%

ecFlow: Best Practices – Suites' design

- KISS: Keep it simple
- Pure text definition-file may be all you need
 - Download an existing suite: `ecflow_client -get`
- Shell suite definition, aka 'stream like definition' may be enough
- Python API: when definition file is no longer needed, yet...
 - Functional programming: no temporary objects, list comprehension
 - We can load/replace a node directly
- Python API:
 - Native `ecflow` API, `ecf` API
 - [pyflow](#): **Config as code design** with YAML file turned into a suite
 - [Wellies](#): DRY-design suites from YAML config
 - [PySuite](#) (OOD)
 - Check jobs creations from client side
 - Simulate: validate there is no deadlock from suite definition



PYFLOW

```
NAME = os.getenv("SUITE", "elearning")
DEFS = Defs()
SUITE = create(NAME)
DEFS.add_suite(SUITE)
DEFS.generate_scripts()
DEFS.simulate()
CLIENT = Client(os.getenv("ECF_HOST", "localhost"),
                 os.getenv("ECF_PORT", 3141))
CLIENT.replace(f"/{NAME}", DEFS)
```

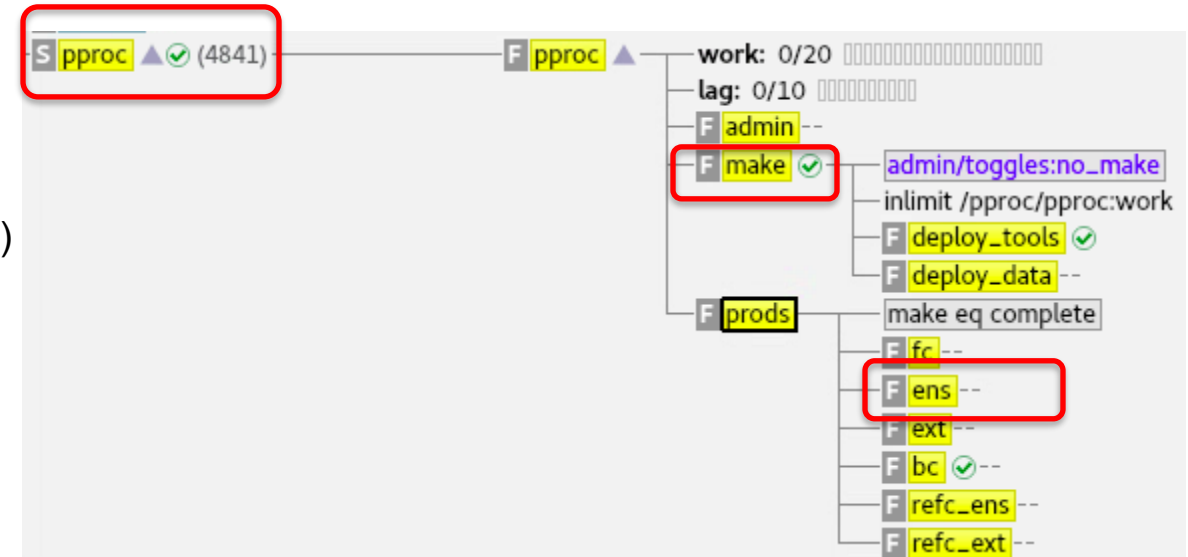
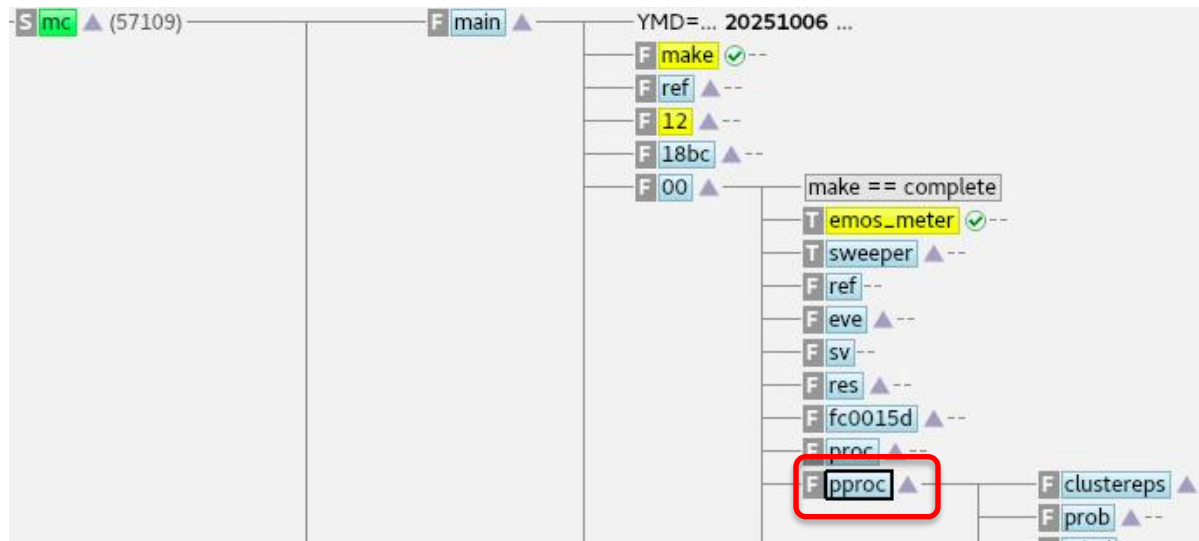


WELLIES

ecFlow: Best Practices – Suites' design - modularity

- **Modular design:**

- Design as **suite**, as **standalone family** (repeat inside), as **embedded family** (inherited repeat, variables): ex. pproc node
- Make/init/setup family v using binaries from a module
- We can move suite from server to server once loaded
 - ecflow_client –plug (provided it is ok with ECF_FILES/INCLUDE)
 - In fact, any node can be moved (! No active task within)



ecFlow: Best Practices – Suites' design – start processing

- **Waterfall** design vs **query** design

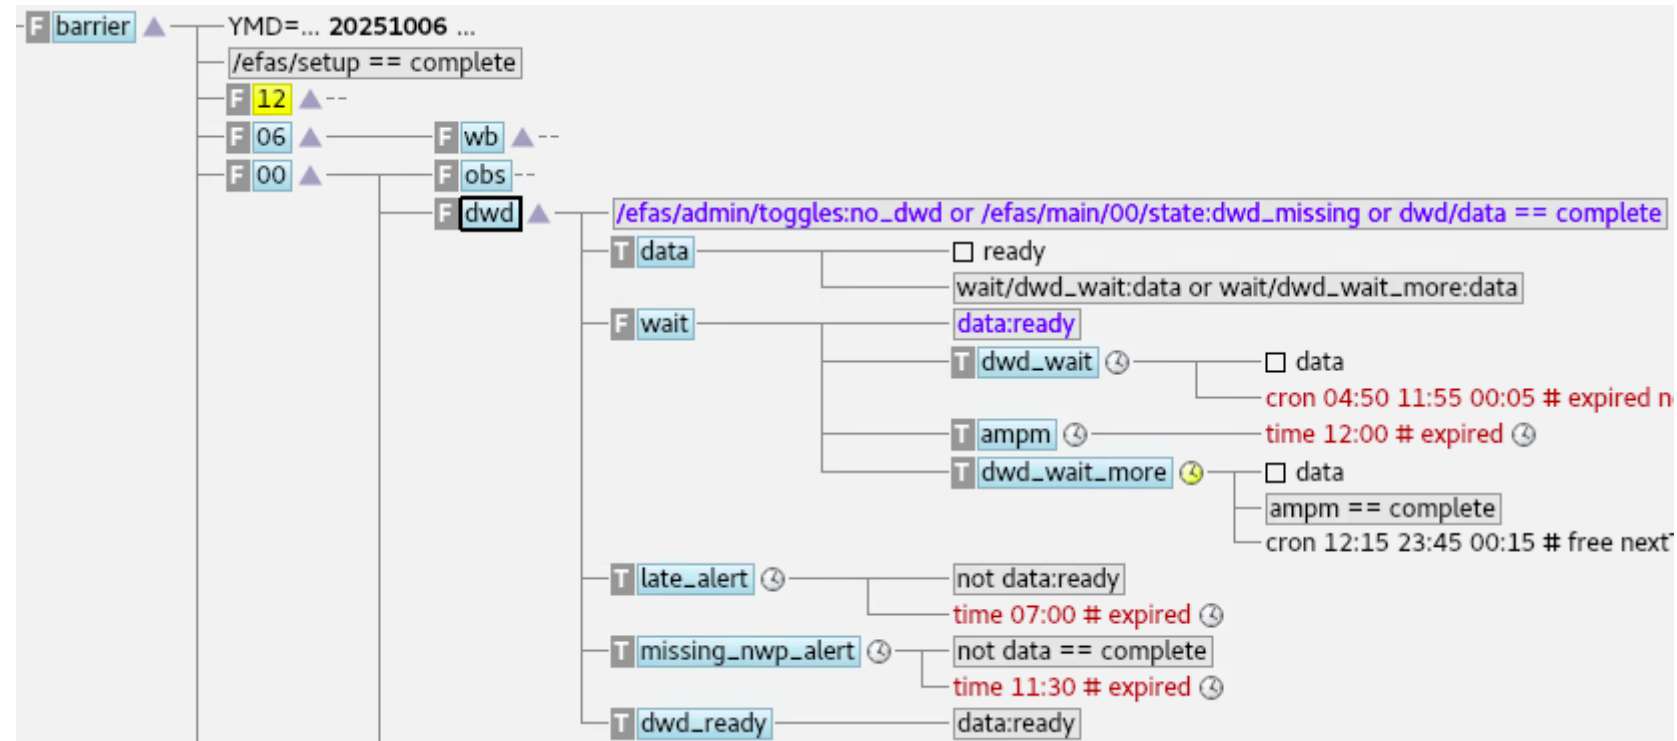
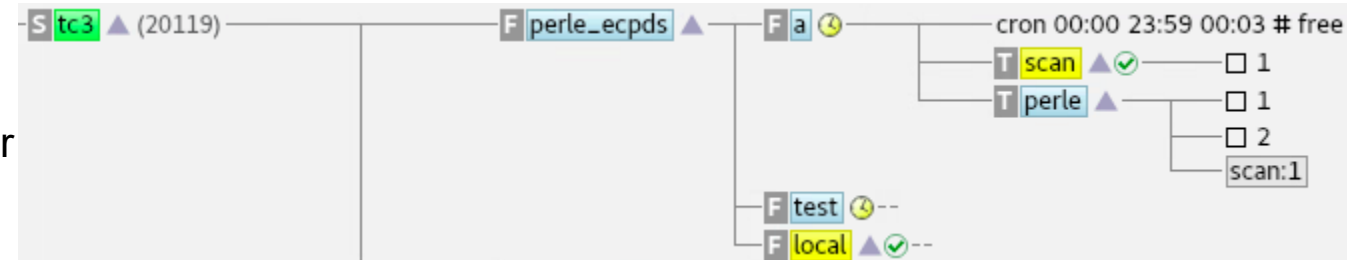
- Event sent by third party to start processing + trigger
 - Ecpds can be configured to do that on file reception

- **Barrier** family to query data availability in a time window

- A combination of
 - Event, Trigger, Complete, Cron
- It may raise multi-levels alarms

- ⚠ a time attribute?

- ✅ an [aviso attribute](#)



ecFlow: Best Practices – Suites' design – preliminary

- Once a preliminary suite is created:
 - Check jobs creations from client side
 - Detect issue with variables used in wrappers, yet may be undefined in the suite (or without default value)
 - Check all include files are present
 - Include loops can be avoided using **%includeonce**
- Simulate the suite from the client side
 - validate there is no deadlock from suite definition
- Check runtimes for the jobs
 - Find the right balance, number of parallel jobs vs user visible runtime
 - Use [bin packing](#) and/or [gnu-parallel](#)
- Add **limits** to prevent flooding with too many jobs
 - On top suite node: add **defstatus suspended**

ecFlow: Best Practices – Suites' design –a visual language

- Ecflow suite as a visual language (text definition example)

- "can you understand easily" ?

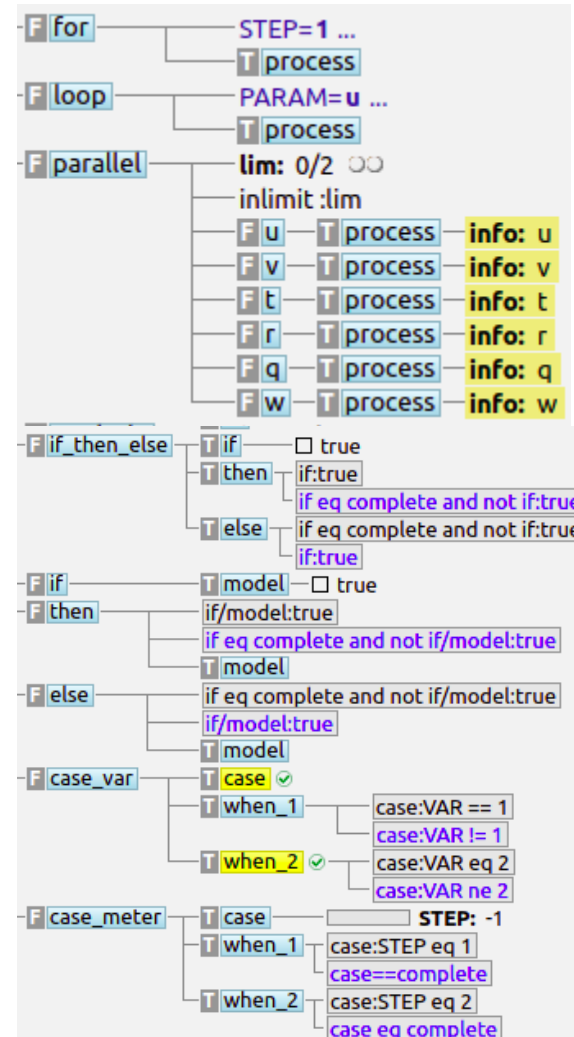
- Suite functionality
- System solutions for suite handling / robustness
- Vs a functionality hidden deep into a task
- One wrapper called multiple time (if model)
- Vs multiple wrappers

- Case block with a
 - Rigid selector (defstatus)
 - vs dynamic branch selection

- For block

- Inner loop
- Repeat outer loop
- 'exploded'

```
family case_var
  task case
    defstatus complete
    edit VAR '1'
  task when_1
    complete case:VAR != 1
    trigger case:VAR == 1
  task when_2
    complete case:VAR ne 2
    trigger case:VAR eq 2
endfamily
family case_meter
  task case
    meter STEP -1 48 48
  task when_1
    complete case==complete
    trigger case:STEP eq 1
  task when_2
    complete case eq complete
    trigger case:STEP eq 2
endfamily
```

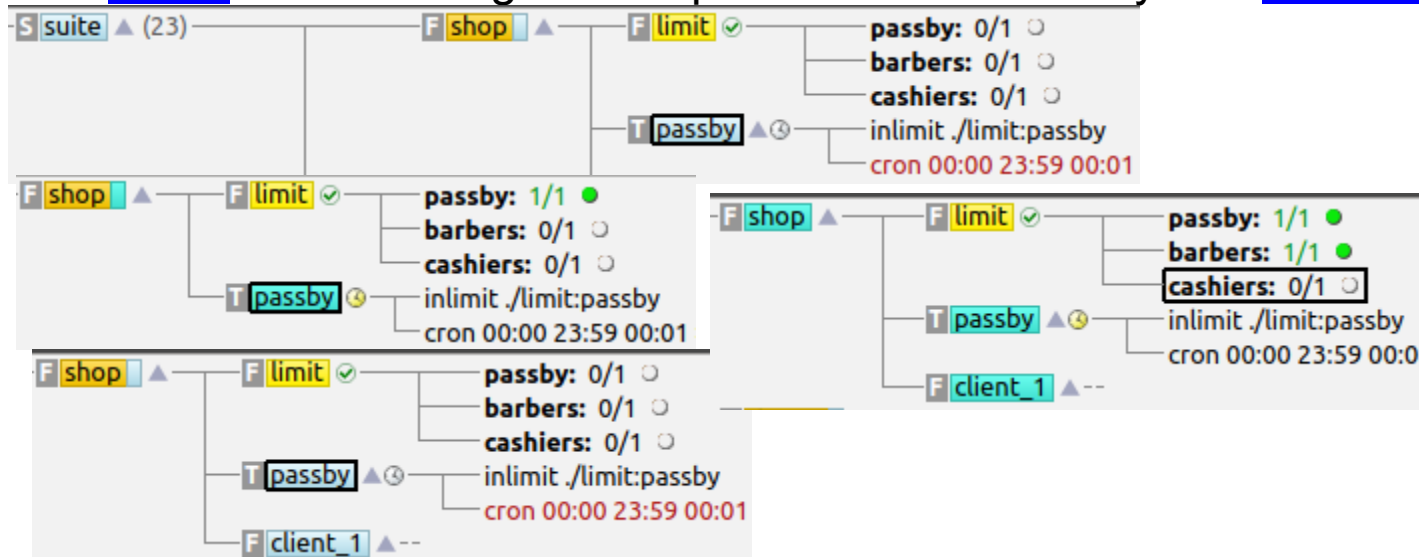


```
family if_then_else
  task if
    event 0 true
  task then
    complete if eq complete and not if:true
    trigger if:true
  task else
    complete if:true
    trigger if eq complete and not if:true
endfamily
family if
  task model
    event 0 true
endfamily
family then
  complete if eq complete and not if/model:true
  trigger if/model:true
  task model
endfamily
family else
  complete if/model:true
  trigger if eq complete and not if/model:true
  task model
endfamily
family for
  repeat integer STEP 1 240 3
  task process
endfamily
family loop
  repeat string PARAM "u" "v" "t" "r" "q" "w"
  task process
endfamily
```

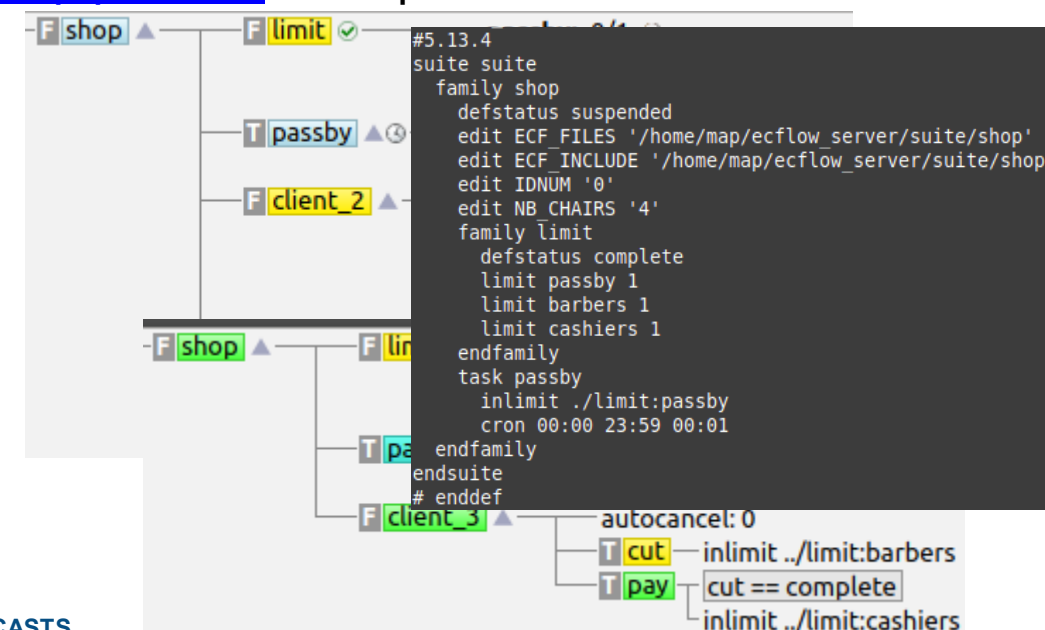
ecFlow: Best Practices – Suites' design – dynamic suites

- **Iterative design:** a suite as a lively entity where we add new families and task, and prune old branches
 - Only consistency is required in the definition file
 - A suite can be defined from multiple definition files
 - Ecflow_client –replace <node> <server>
 - Ecflow_client –delete <node> <server>
 - Provided no active / submitted tasks lie below to prevent zombies

Tasks can be designed to update their own family: the [barber shop problem](#) example



```
#!/%SHELL:/bin/bash%
%manual
%end
%include <%QSUB_H:pure%>
%include <%HEAD_H:trap.h%>
cd ${TMPDIR:=/tmp}
IDNUM=%IDNUM:0%
ECF_INCLUDE=%ECF_INCLUDE%
export PYTHONPATH=${ECF_INCLUDE}:${PYTHONPATH:=}
python3 $ECF_INCLUDE/shop.py -n $((IDNUM+1))
%include <%TAIL_H:endt.h%>
```

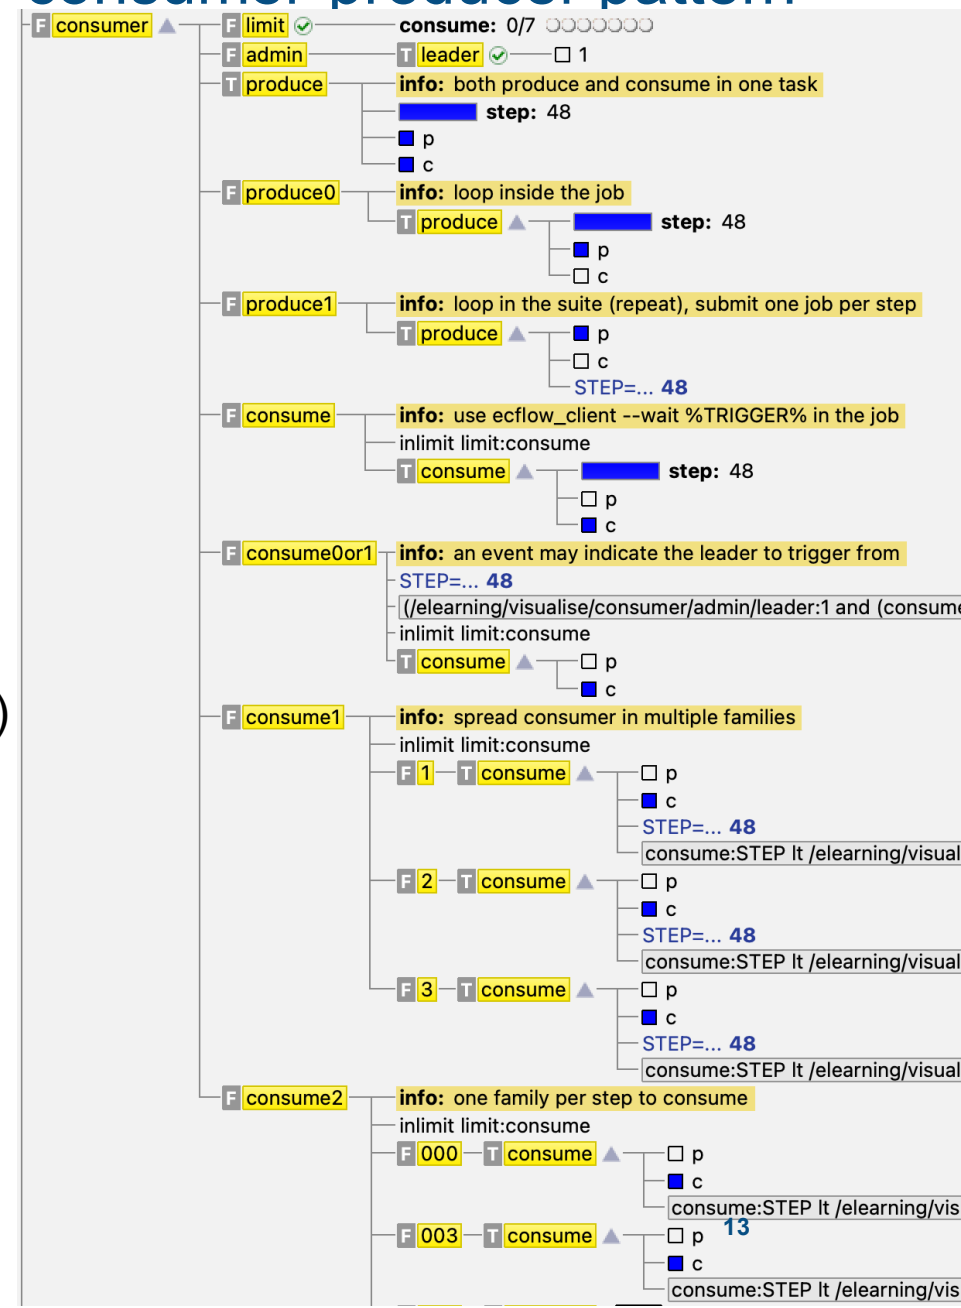


ecFlow: Best Practices – Suites' design –consumer-producer pattern

This example is part of the ecFlow eLearning [repo](#)

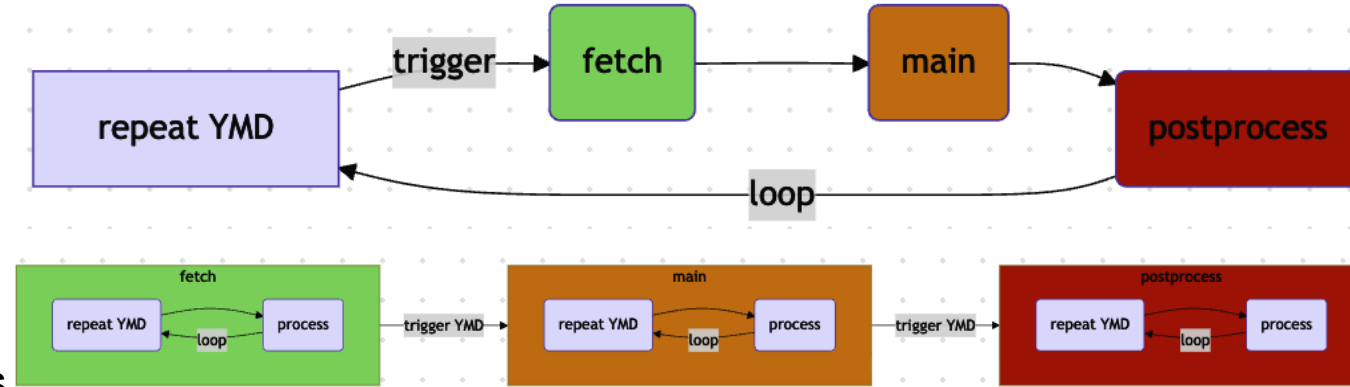
Among the multiple solutions we can find in ecFlow suites:

- a single task, do it all
 - Inner loop
 - Inner checkpointing to restart from previous state
- Separation of concern: producer vs consumer
 - Inner loop (produce0) vs outer loop (using repeat: produce1)
- Different approaches for a consumer:
 - We prefer to expand the loop with one family per step (consume2)
 - A limit to control the load
 - A failure won't affect another STEP
- It depends the project, the criticality, the available resources and support

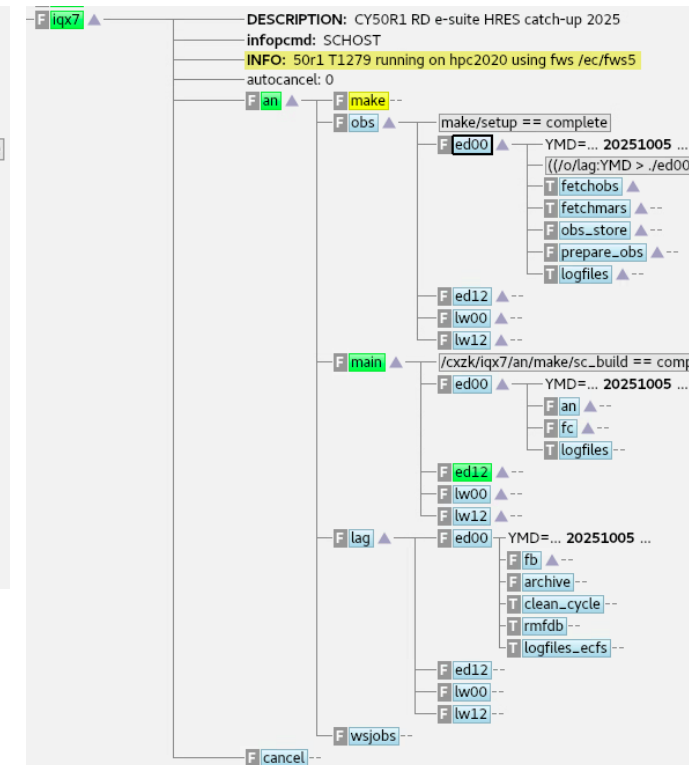



ecFlow: Best Practices – Suites' design - repeat

- **Cycling**: shared vs local repeat attribute
 - Possible in operation
 - Thanks to operators, to analysts to treat concerns
- **Loosely coupling** leads to more complex triggers.
 - Used in research, and few operational suites
 - fetch observations, process, archive
- **Repeat**
 - date / integer / enumerated / string / datelist
 - Extra variables for date (Julian, Dow, Doy)

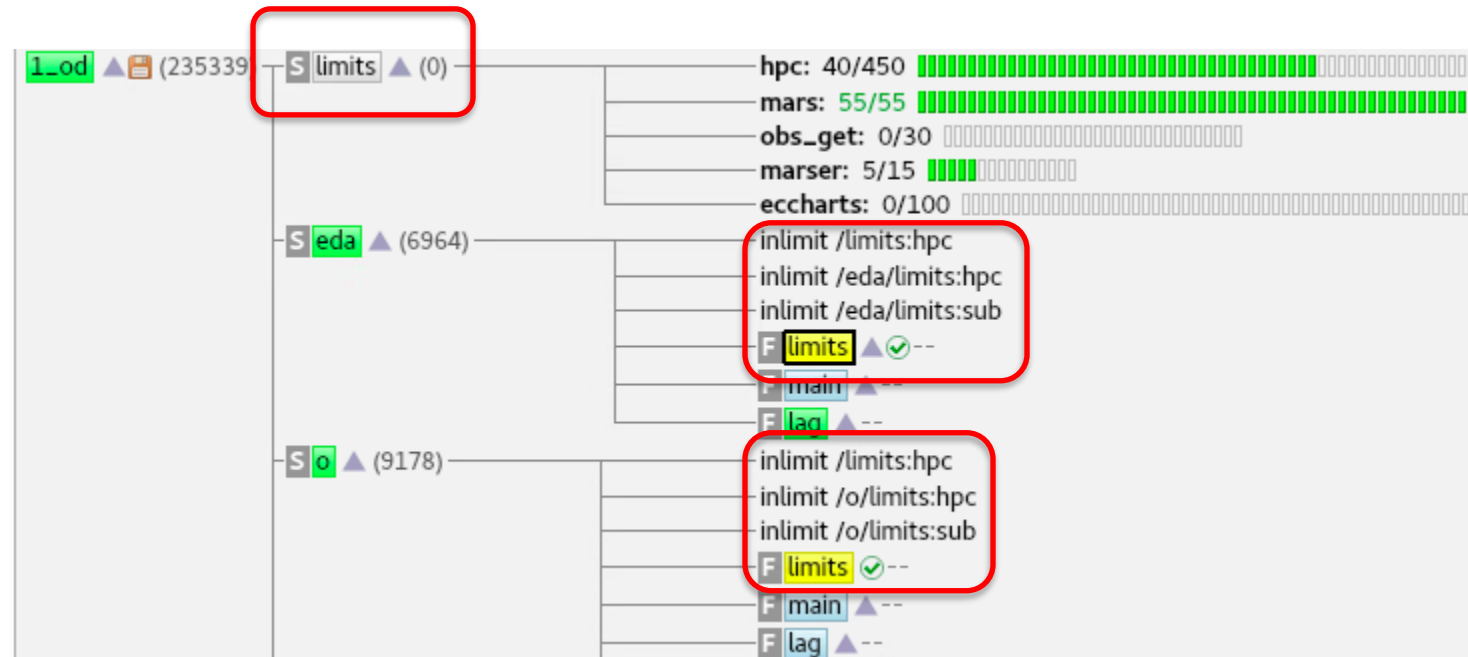


F main	YMD=... 20251007 ...
repeat date YMD 20241023 20321212 1	
Ⓞ FAMILY	main
Ⓞ FAMILY1	main
Ⓞ YMD	20251007
Ⓞ YMD_DD	7
Ⓞ YMD_DOW	2
Ⓞ YMD_JULIAN	2460956
Ⓞ YMD_MM	10
Ⓞ YMD_YYYY	2025



ecFlow: Best Practices – Suites' design - limits

- Resources: limit / inlimit / inlimit –s
 - Inlimit –s (tasks in submit state only) can refrain submission bursts
 - List all active/submit tasks under the limit
 - Set limit 0 from the GUI as a 'distributed one click suspend'



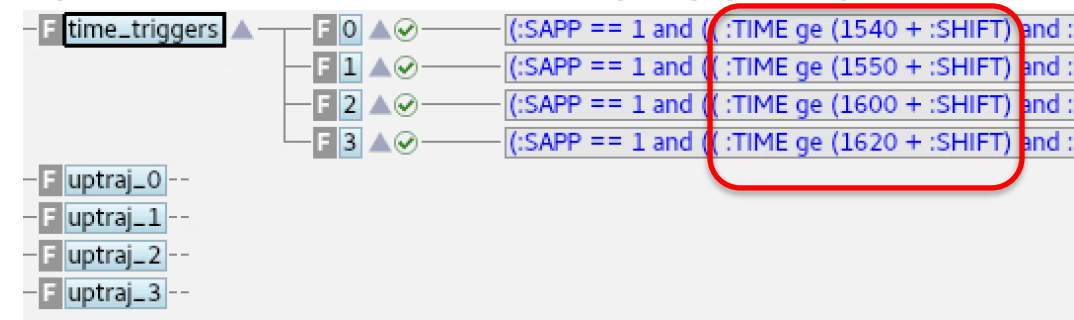
ecFlow: Best Practices – Suites' design – dependencies (triggers)

Group time dependencies in dedicated families + triggers

Easy replacement when schedule changes

Defstatus complete in catchup-mode

Dependencies: simple and explicit triggers (producer-consumer)



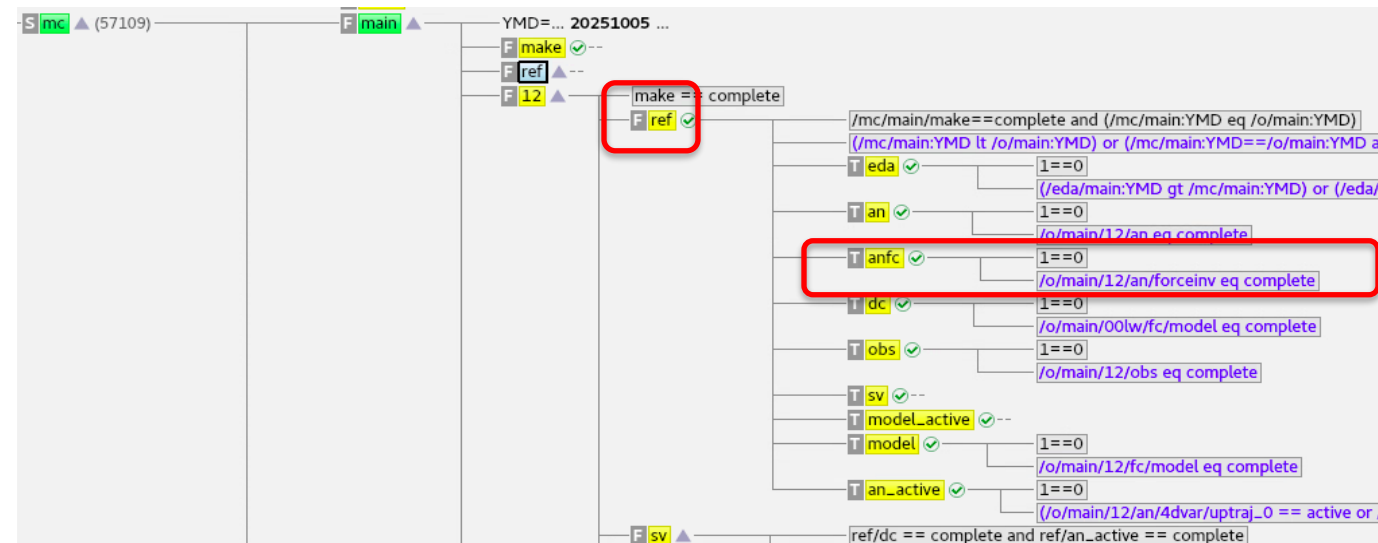
- Triggers can use node states, variables, limits, events, meters, late
- Expressions can use function cal::date_to_julian(/path/to:variable)

Group external triggers in dedicated families

- Dummy tasks easily replaced
- Set defstatus complete in standalone mode

'fake task' to collect complex triggers:

- **Trigger impossible** + complete attribute
- Can be replaced / updated
- Can be set 'defstatus complete'
- 'trigger by transitivity'




- **umbrella triggers**: factorise trigger on parent node when possible
- Absolute vs relative path? Rigid vs fit for 'move family'

ecFlow: Best Practices – Suites' design – Cron

- We refrain from Cron in operation

- "a Cron task never ends"
- Not compatible with upper repeat attribute, unless ... it is managed like with the Barrier pattern
- A meter, event is reset on the requeue at completion, beware



- Yet sometimes a Cron is just enough:

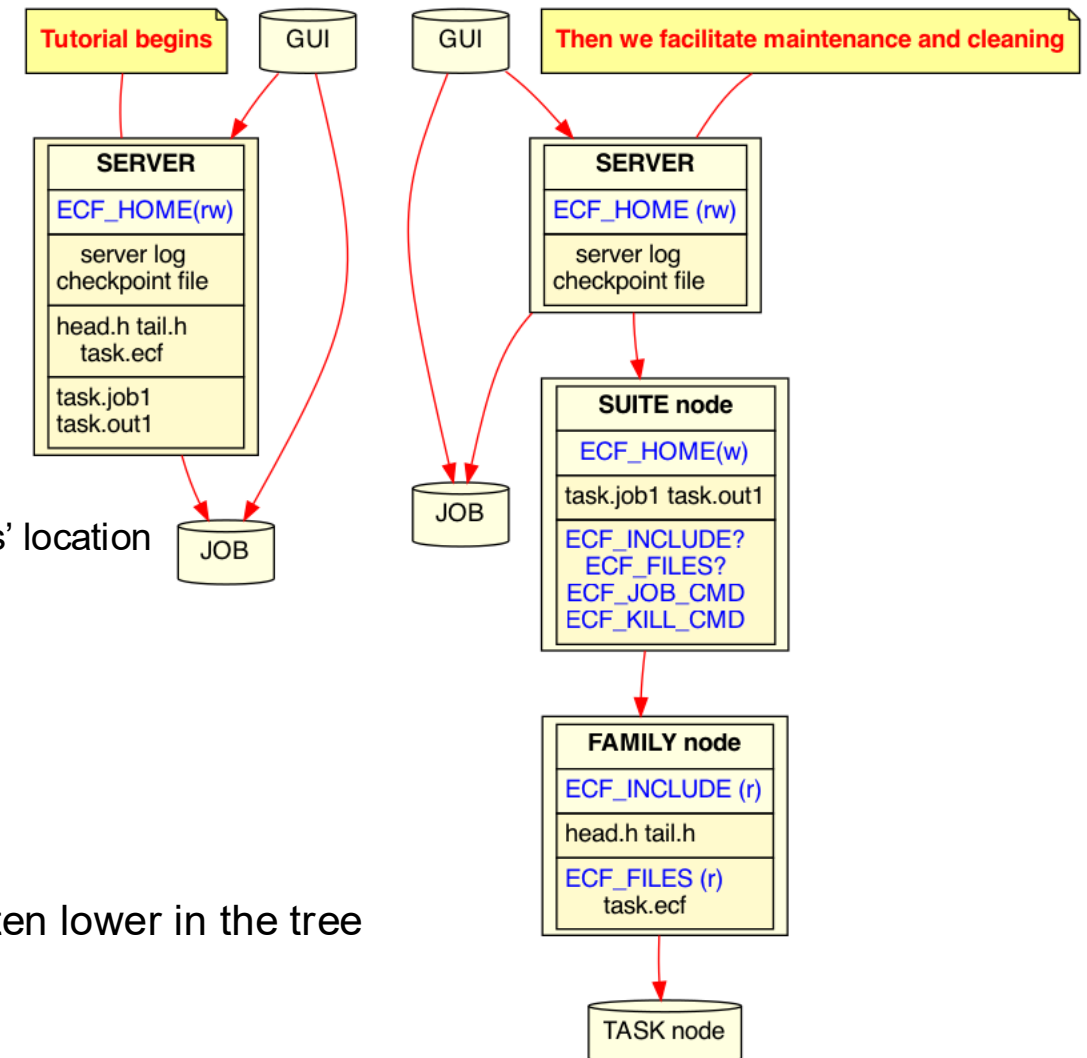
- Error handling: It has to be a very robust task. It would block on abort. designed 'not to fail' (send mail, external QC)
- Process accumulated input in the interval
- Minimum/NO operators action requested
- Can behave on ECF_TRYNO (job occurrence number) value
- Drawbacks of the Cron:
 - Output would be overwritten on next occurrence: an example where it is worth overwriting ecflow default variable

```
edit ECF_JOBOUT '%ECF_OUT%/%ECF_NAME%.%ECF_DATE%-%TIME%-%ECF_TRYNO%'
```

- Use date/time + repeat attribute or :TIME variable in a trigger expression as alternatives
- Cron is fit for simple administrative tasks

ecFlow: Best Practices, files locations, ECF_FILES, ECF_INCLUDE, ECF_HOME

- In the tutorial, ECF_HOME is used for everything:
 - Server logs, checkpoint
 - Tasks wrappers and headers
 - Job and output files
- Later:
 - Use ECF_FILES, ECF_INCLUDE variables (-r):
 - Can be refined lower in the tree for tasks templates and headers' location
 - Store wrappers/headers in read only access
 - ECF_HOME: where job files are (-w)
 - Where local outputs are by default (when ECF_OUT is absent)
 - ECF_OUT: preliminary path for remote jobs outputs
 - ECF_*_CMD variables: can be defined on top, and overwritten lower in the tree



ecFlow: Best Practices – Suites' design – ECF_MICRO

- default micro is %
- it can be changed **globally** from the definition file:

edit ECF_MICRO \$ # change to dollar / .def

helpful when using multiple languages in a suite

helpful when using script from different teams with different convention

- ECF_MICRO can be changed **locally** (and reverted)

in the wrapper/headers using

%ecfmicro ~

~ecfmicro %

- Change it in specific deliberate, documented cases
- Or **stick to simplicity using %**

```
#!/%SHELL:/bin/bash%
%manual
...
%end
%include <%QSUB_H:pure%>
%include <%HEAD_H:trap.h%>
%include <%BODY_H:pure%>
%ecfmicro ^
# block
^ecfmicro %
%include <%TAIL_H:endt.h%>
```

ecFlow: Best Practices – Suites' design – ECF_EXTN

A variable ECF_EXTN: the task wrapper extension

default is .ecf

it can be changed in the definition file

ex: edit ECF_EXTN '.epy'

useful when a suite uses multiple languages

even for shell wrappers it may facilitate a transition (ksh to bash)

we can design a pure shell/python script to be a valid task wrapper: edit ECF_EXTN .sh

(when there is no need for %include)

we can distinguish source-controlled wrapper (.ecf) vs generated wrappers (.ecg)

we can extend this practice

to distinguish template headers (.h) vs pure shell include (.sh)

where no preprocessing is expected (%includenopp <setup.sh>)

```
#!/usr/bin/env $SHELL:python3$  
$include <$QSUB_H:pure$>  
$include <$HEAD_H:head.epy$>  
$includenopp <$SCRIPT:pure$>  
$include <$TAIL_H:tail.py$>
```


ecFlow: Best Practices – Wrappers design trapping

- **Robust** jobs: designer duty is to preserve **trapping** (handle early exit and/or external signals reception)
 - Beware shell functions definitions with ksh:
 - `x() { cmd; }`
 - `function x {cmd; }`
 - in ksh, this will need to reload trapping
 - `%include <trap.h> ... trap 0`
 - Or use bash
- `set -eux -o pipefail`
 - Fail on error immediately (report abort)
 - Early exit must be trapped (`trap ERROR 0`)
 - Fail on undefined variables
 - Verbose mode logging each command
 - Set `-o pipefail` does NOT export in subshell
 - Refrain from using ssh inside, dedicate a task submitted on the other host (or use `-o ExitOnForwardFailure=yes`)
- **Consistency** and **Maintainability**: same `head.h` among suite tasks

```
#!/bin/bash
# head.h
set -e # stop the shell on first error
set -u # fail when using an undefined variable
set -x # echo script lines as they are executed
set -o pipefail # fail if last(rightmost) command exits with a non-zero status

ERROR() {
    # Clear -e flag, so we don't fail
    set +e
    wait
    ecflow_client --abort=trap --host=%ECF_HOST% --port=%ECF_PORT%
    # cleaning ?
    echo "environment was:"; printenv | sort
    trap 0 # Remove the trap
    exit 0 # End the script
}


# Trap any calls to exit and errors caught by the -e flag
trap ERROR 0

# Trap any signal that may cause the script to fail
trap '{ echo "Killed by a signal"; ERROR ; }' 1 2 3 4 5 6 7 8 10 12 13 15

export PID_RID=$$
ecflow_client --init=$PID_RID
```

ecFlow: Best Practices – Wrappers design - Manual

- Task documentation: block `%manual ... %end`
 - In the wrapper
 - In any included (pre-processed) script
 - We can delegate a suite variable for the filename for the manual
- A file `<node>.man` can provide a manual page for a suite or a family node
 - This can be a static or dynamic file
- ECF_URL_CMD is 'interpreted CMD variable' to open an URL from GUI
- What, why, When (delay), Where (information), Who
 - Title
 - Description
 - Input/output/variables
 - Procedures
 - Contact points



```
#!/usr/bin/env %SHELL:bash%
%manual
%include <%MANUAL_H:pure%>
%end
%include <%QSUB_H:pure%>
%include <%HEAD_H:head.h%>
%include <%SCRIPT:pure%>
%include <%TAIL_H:tail.h%>
```

```
tree . | grep -E '.*.man$'
.
├── compo.man
├── eda.man
├── fsobs.man
├── main.man
├── mc.man
├── mon.man
├── o.man
├── refc.man
├── thu.man
├── web.man
├── 4dvar.man
├── Aeolus_AMD_ifstraj.man
├── ifsmin.man
├── ifstraj_999.man
├── ifstrajbg.man
├── ifstraj.man
├── monitoring.man
├── oopsvar.man
└── pop.man
```

Name	Value
MANUAL_H	doit.man

Info Manual

a man page for doit task

ecFlow: Best Practices – Wrappers design - Rerun ability

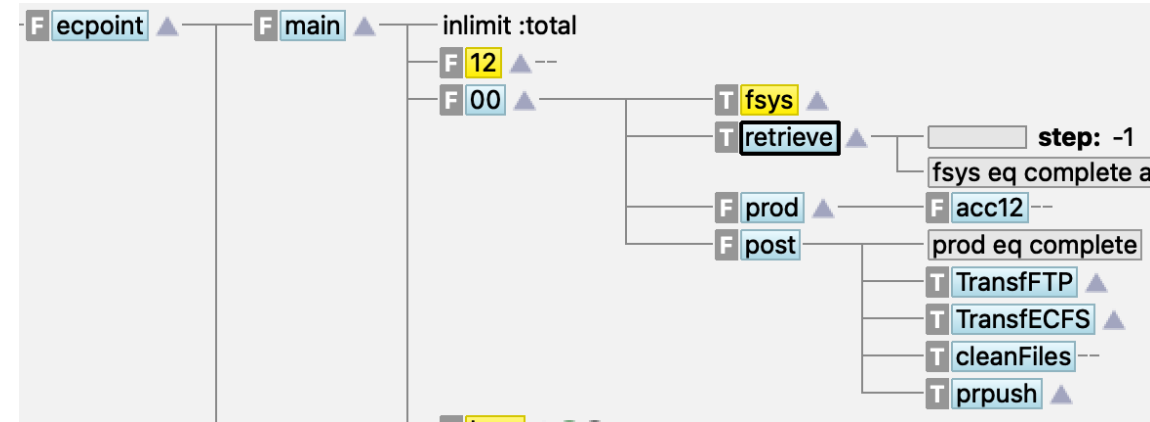
- **Idempotent:** "a task shall be rerunnable" as part of a robust design
 - Clean ground rerun mode for most tasks
 - Validated in pre-operational mode
- Multiple views: Retrieve – compute – postprocess - push pattern
 - May be one task in development
 - One task per function in operation
- **Checkpointing** for tasks?
 - Minimise rerun time (at the price of changing ETA)
 - Clean ground rerun mode by default vs 'thanks to a variable'
- Rerun mode: suites variables ECF_TRIES (total attempts) ECF_TRYNO (current occurrence)
 - Set ECF_TRIES to 1 for immediate abort, or increase it to allow automatic re-submit after abort
 - ECF_TRYNO is current try-number: jobs can be 'self-aware'
 - Jobs can change behaviour: verbosity, debug mode, silent mode, mail, contact alert system

ecFlow: Best Practices – Wrappers design – FAMILY1 TASK

- Anti-pattern ❌
- Refrain from using test in FAMILY, FAMILY1, TASK
 - A dedicated suite variable is better on the long term
- Separation of concern:
 - Running a task vs configuring a task
- Such script would be a generic script,
 - asking for links to provide the different tasks wrappers
 - and more maintenance (rather than focussing on the suite definition)

ecFlow: Best Practices – ecflow_client –wait <expression>

- A child command to query a trigger expression from the server to continue
- Aka 'a trigger in a job'
- Blocking: it may cause **live lock**
- **non blocking alternative:** use query
 - not a child command
 - status=\$(ecflow_client –query <node>)




Name	Value
defined in task retrieve	
TRIGGER	/mc/main:YMD gt ../../mc/main:YMD or /mc/main/12/sweeper:rmin gt 12

```
if [[ "%TRIGGER:1==1%" != "1==1" ]] ; then
    for Step in ${seq_retrieve}; do
        #PBS -l EC_memory_per_task=45000mb

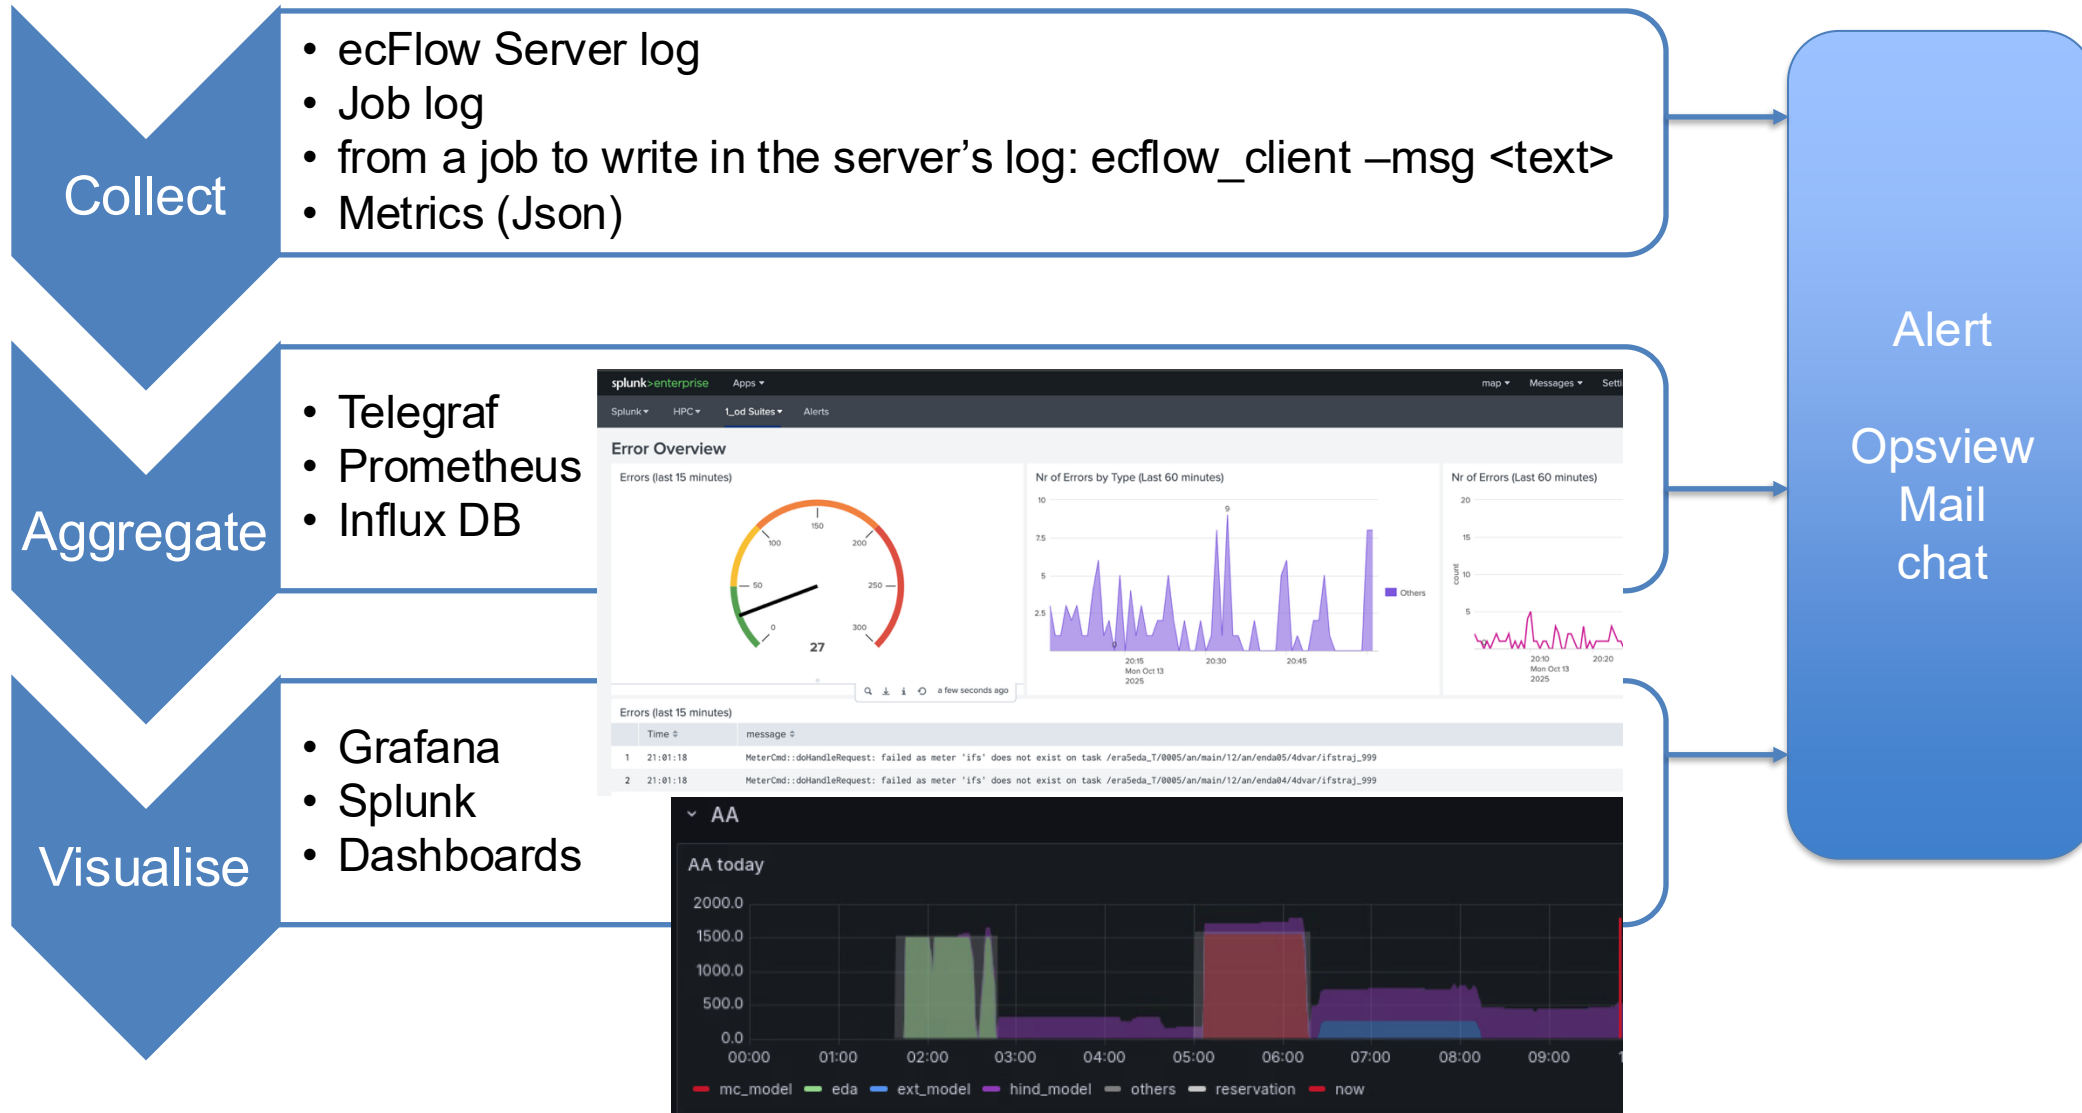
        %includenopp <cal_rainfall_ecmwf_ens_12h/MarsRetrieval.ksh>
        xmeter step $Step
        $ECFLOW_CLIENT --wait %TRIGGER:1==1%
    done
```

ecFlow: Best Practices – Wrappers design - Cleaning

- Work in WDIR / TMPDIR
 - Own job cleaning vs admin cleaning task vs system cleaning
 - ecflow_client –complete is not the end
 - Complete for ecflow
 - Cleaning may occur after complete was sent
- In conjunction with **monitoring** tools
 - detect trends, or infringed thresholds
- **Retention** policy:
 - Cleaning can use DELTA_DAY, to keep few days online
 - **Check** validity of archived data before cleaning
 - Move away data from **visibility** before cleaning
 - Roll back capability?
 - documentation
- Test cleaning in **esuite mode**

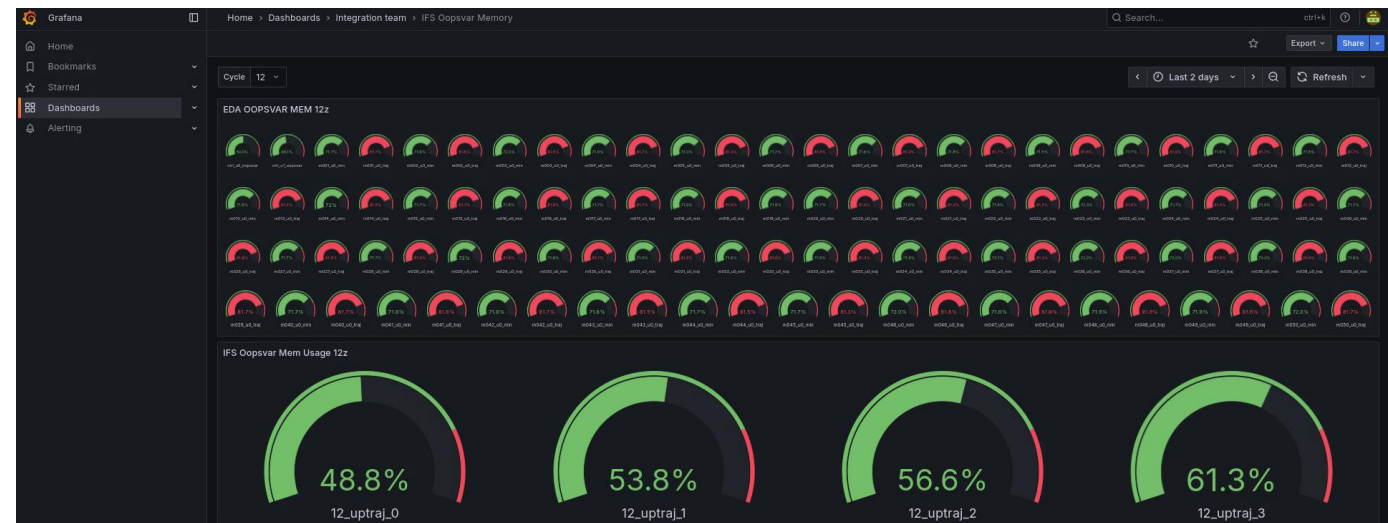



ecFlow: Best Practices – Monitoring & Observability



ecFlow: Best Practices – Monitoring & Observability

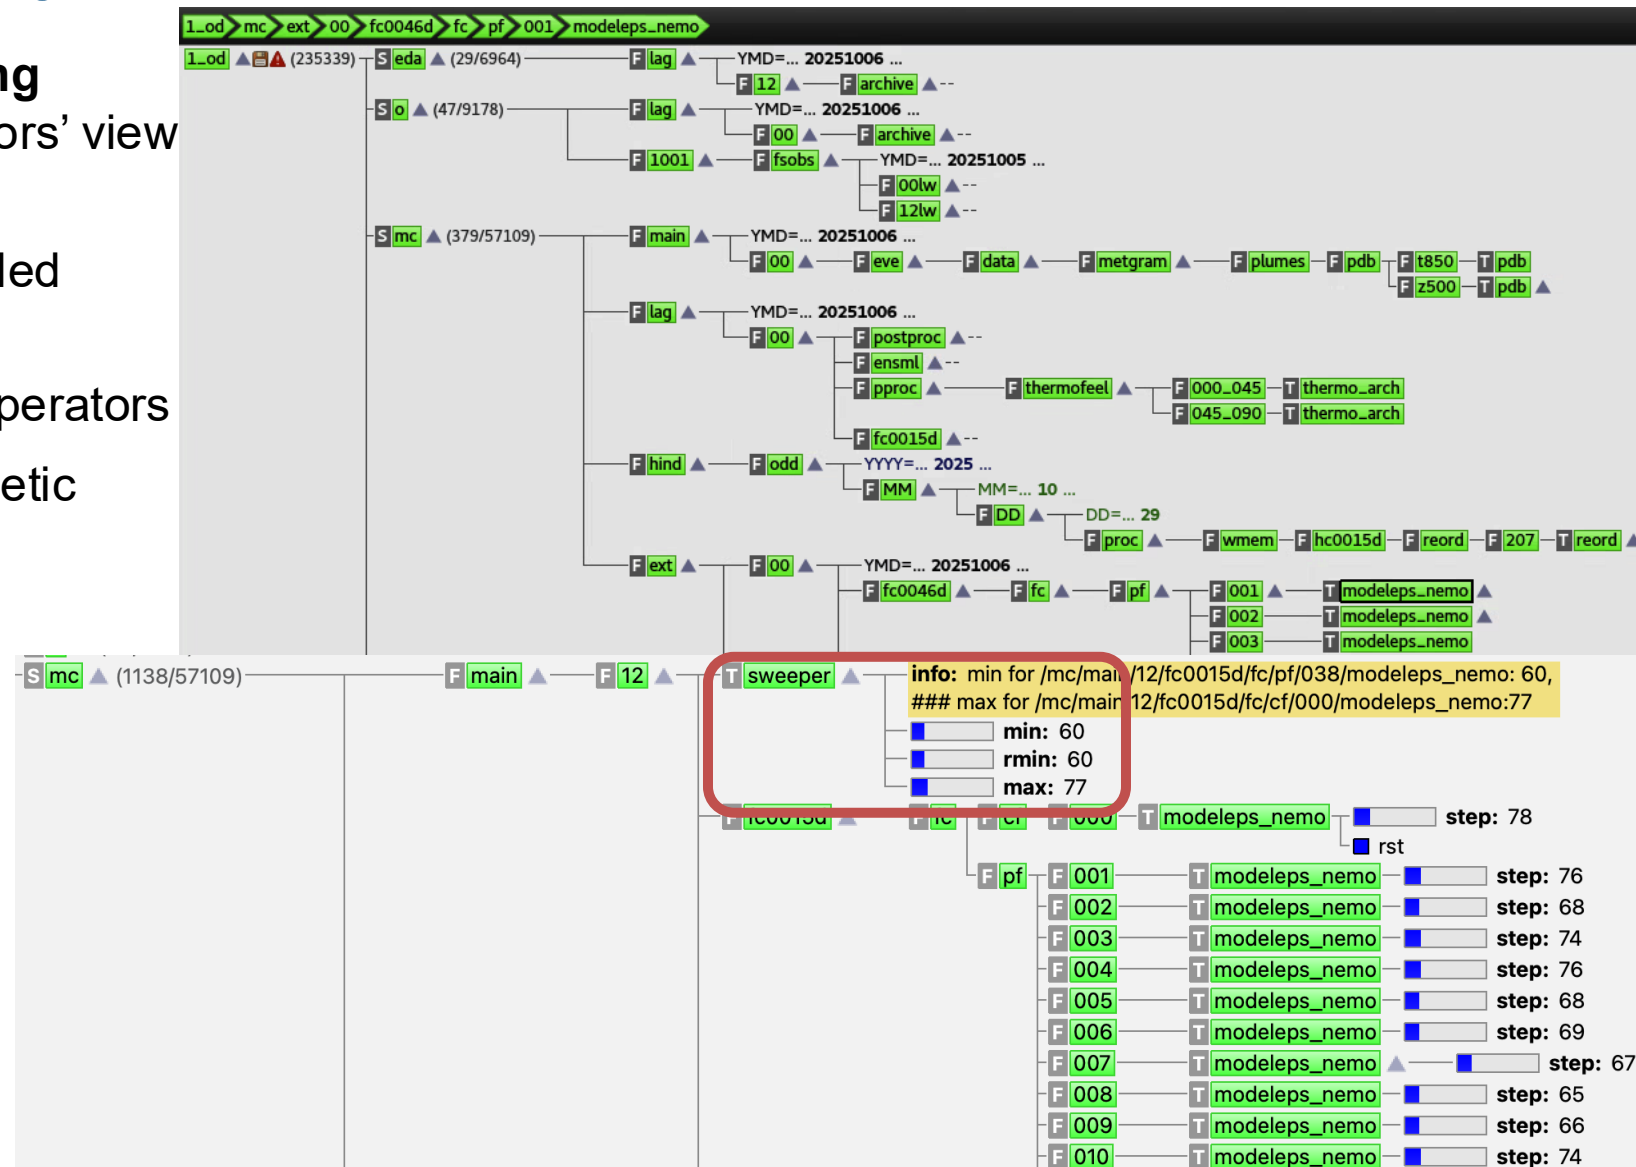
- Contact the server server:
 - `ecflow_client -ping`
 - `ecflow_client -info`
 - `ecflow_client -query <node>`
 - `ecflow_client -get`
 - to dump server content, may offer status when log is not accessible
- Or better use the **checkpoint** file than hammering a server with too frequent queries



Best practices: Monitoring - Operators view

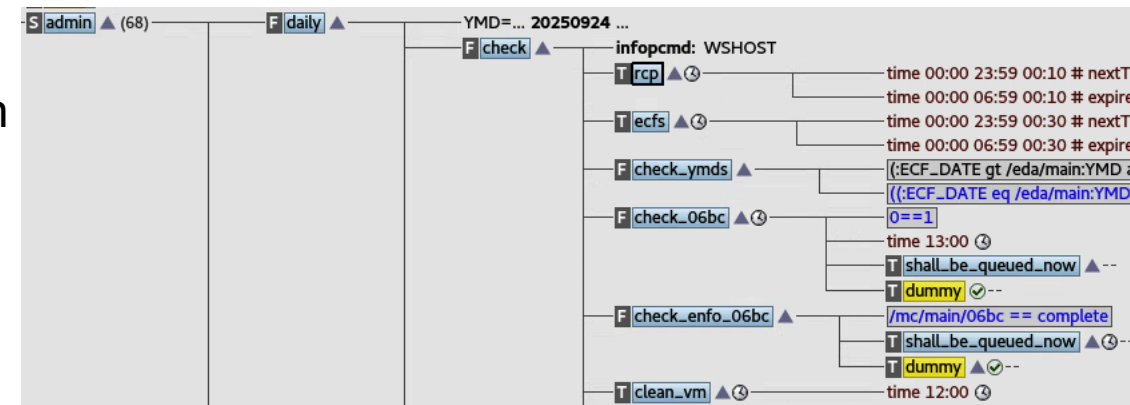
- Keep in mind **visibility** and **monitoring** purpose from the design phase: Operators' view

- Operators display only submitted/active/aborted/suspended nodes and few attributes
- Top is priority for ecflow and for operators
- Locate up 'summary tasks' (synthetic task): ex. sweeper



ecFlow: best practices – Monitoring - admin suite

- An admin suite per server
 - Snapshot checkpoint file 10min, 30min
 - Daily looping the server logs + archive
 - Warnings / cleaning for disk space
 - Host ‘outer watch dogs’ for suite



ecFlow: Best Practices – Errors handling

Detect

- Exit code (trap)
- Timeout: **late attribute**
- **Watch dogs**: inner, outer dedicated tasks to turn abort when a problem is detected
- Active watchdog (sweeper) vs passive watchdog (dummy task with trigger and complete attribute)
- Exception ?
- Monitoring the monitor

classify

- Light v critical:
 - **Transient**: managed setting ECF_TRIES > 1 to overcome glitches
 - **Resource**: adjust a limit, set limit 0 to 'suspend service'
 - **Critical**: disk switch, cluster switch, handover ecflow server
- Track error patterns

Take action

- Retry: **Automatic rerun** (ECF_TRIES)
- Automatic (task managed) **change of job behaviour** (according to its ECF_TRYNO)
- **Compensation**: family handover, server handover
- **Escalation**: Warn with mail, chat, Opsview
- Incident **response plan**: STHOST, SCHOST, FDB config, checkpoint recovery, server switch

ecFlow: Best Practices – Security & access

Authentication

- Certificates
- SSL/TLS communications ENABLE_SSL compilation option
- Password file

Authorisation

- RBAC: Whitelist for server read vs read-write access
- ACLs with ecFlow in preparation
- Hardening ecFlow server: VM with local logs, checkpoint, config scripts
- Isolation: docker container, network
- Server owner vs job owner
- Communication on a fixed port 3141

Audit Monitoring

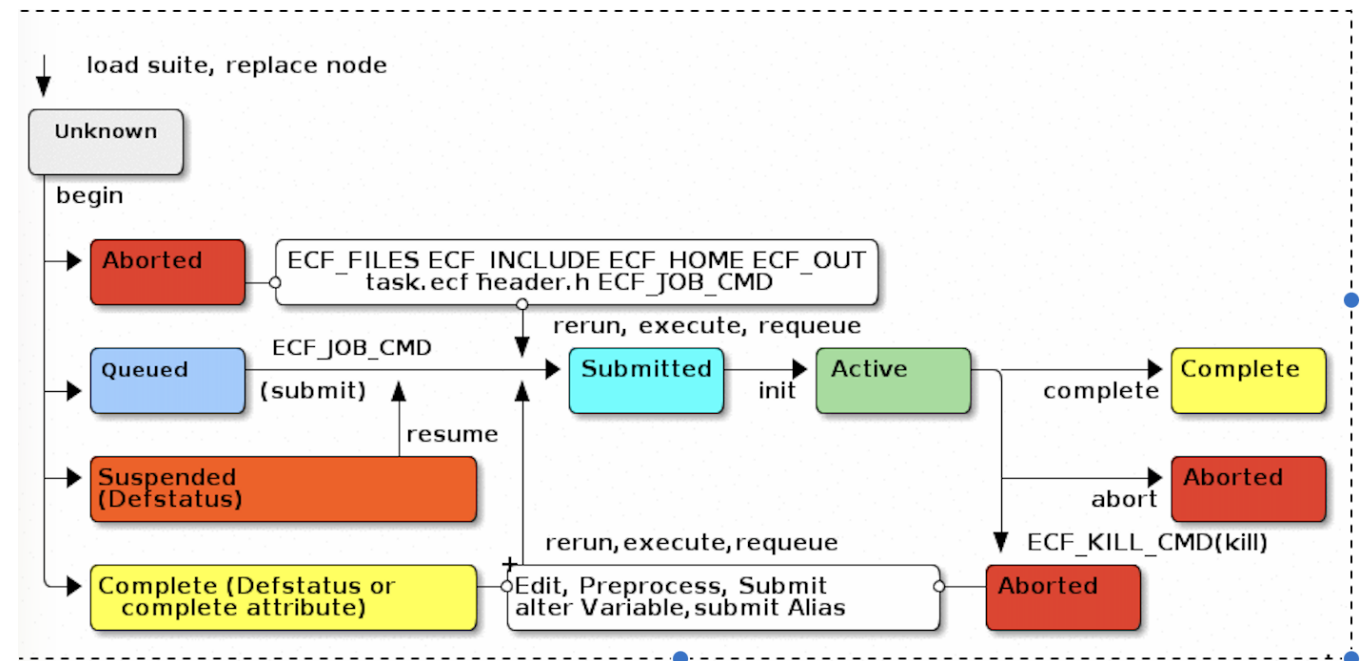
- Logging: Server logs, jobs logs, tar logs into ecfs
- tracing
- Alerts
- Check compliance by peer-review and/or third party
- NEVER RUN ecFlow server AS ROOT

ecFlow: Best Practices – Operational tips – requeue?

- We distinguish

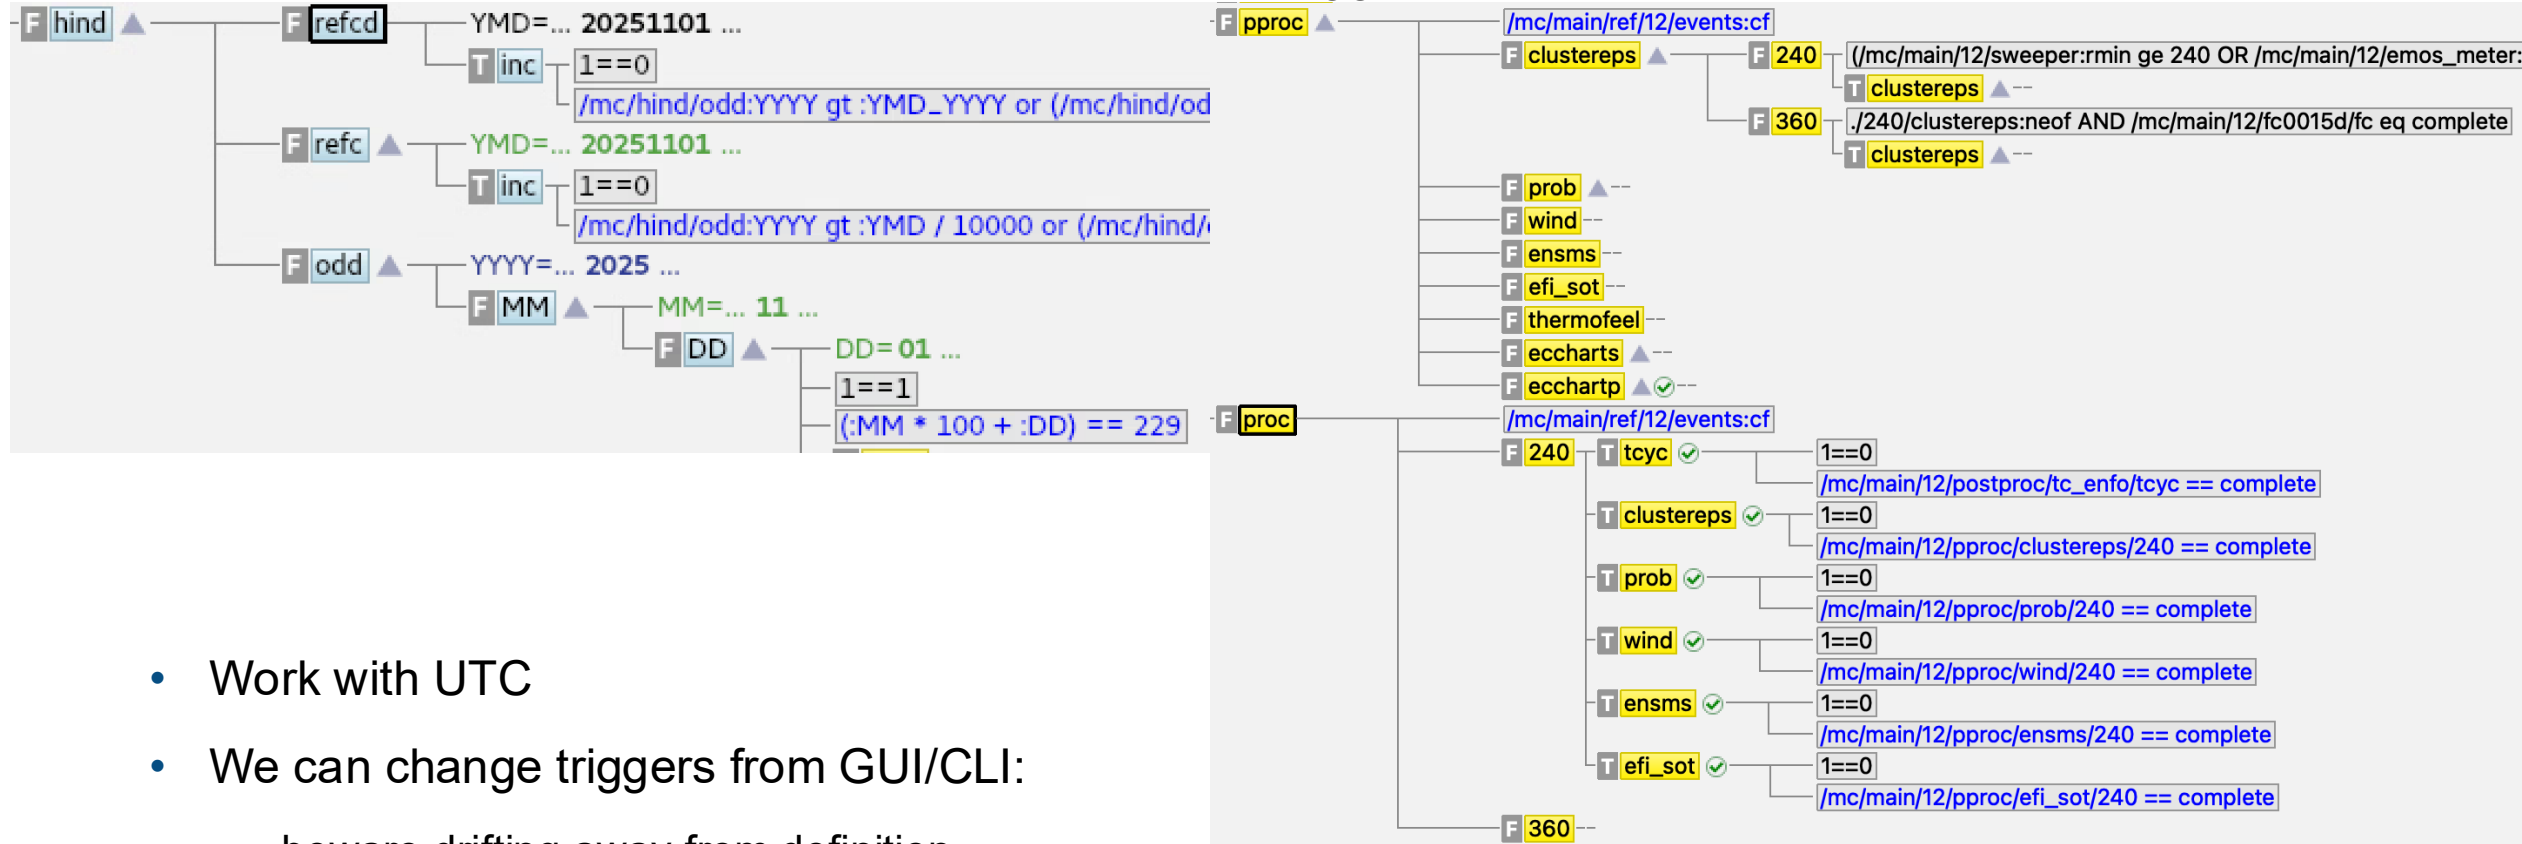
- Rerun: aka force queued (suite/family node),
 - ECF_TRYNO **increment** and honour limits, triggers, date/time dependencies
- Execute: run jobs without condition
 - ECF_TRYNO **increment**
- Requeue: resets
 - ECF_TRYNO (previous output **overwrite**),
 - repeat, date/time attributes
- Requeue aborted: menu on Family/Suite node
 - Requeue only aborted tasks below

Rerun	Ctrl+U
Execute	Ctrl+E
Requeue	



ecFlow: Best Practices – Operational tips

- Fake families/tasks to offer shuffled view for simple triggers



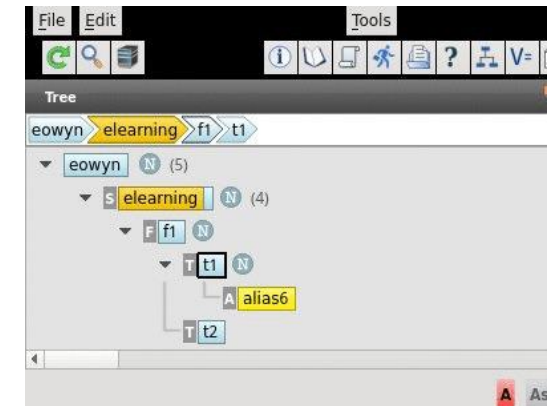
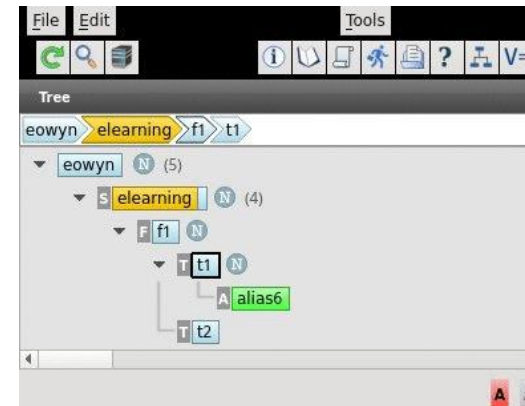
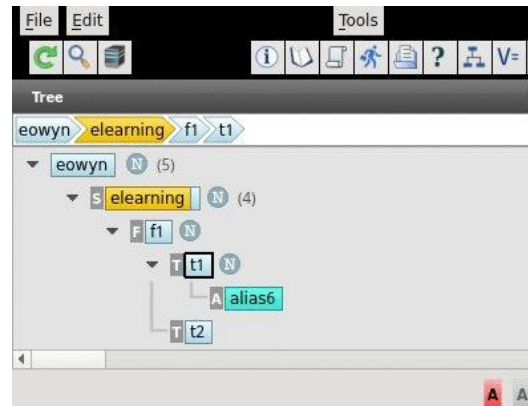
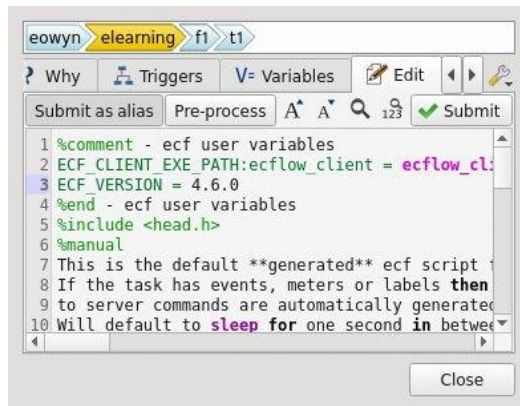
- Work with UTC
- We can change triggers from GUI/CLI:
 - beware drifting away from definition
 - better prefer node replace for valid, source-controlled definition
- Keep a maintenance time window: load balance suites over multiple servers

ecFlow: Best Practices – Operational tips

- Wrappers/headers can be updated with no need to replace the suite / node
 - Changing ECF_FILES, ECF_INCLUDE can point to another scripts' location
 - Atomic wrapper/headers update may occur when the task is queued, away from peak time
 - We deploy the script update from the SCM into ECF_FILES or ECF_INCLUDE directory, logging the change
 - We can run the task as an alias with DELTA_DAY -1

ecFlow: Operational tips - Alias

- An Alias can be created from GUI
 - Edit->Submit As Alias
 - Test a variation from the code without scripts alteration
 - Test a variation from variable(s)
 - A .usrN file is created and submitted
 - This file can be modified to submit the same alias again
 - Multiple aliases can be created for the same task



ecFlow: Operational Tips - Zombies

- Jobs are submitted with variable **ECF_PASS** set to pseudo-random value by ecflow server
- ⚠ **ATTENTION**
- Jobs are defined with unique identifiers **ECF_HOST-ECF_PORT-ECF_NAME-ECF_PASS**
 - ECF_PASS **mismatch** leads to a zombie
 - Simple (not best) action: set ECF_PASS FREE from ecflow_ui (Free password) to allow communication
 - Analyse logs, understand the cause (multiple init, child command after complete, system issue ...)
 - Clear Flag
 - Terminate / Kill / Fob off / Delete / Rescue Adopt a child command on a backup server (child will look for a list of foster parent thanks to ECF_HOSTFILE).

The screenshot displays the ecFlow web interface. At the top, a job flow diagram shows a 'map' node branching into 'F lorenz' and 'F otc_course'. 'F lorenz' further branches into 'infopcmd: SCHOST' and 'T compute'. The 'T compute' node is highlighted with a red 'Z' icon, indicating it is a zombie. Below the diagram, a navigation bar includes tabs for 'Info', 'Why', 'Variables', 'Server log', 'Node log', 'Zombies', 'Suite filter', 'Timeline', and 'Server load'. The 'Zombies' tab is active, showing a table of zombie jobs.

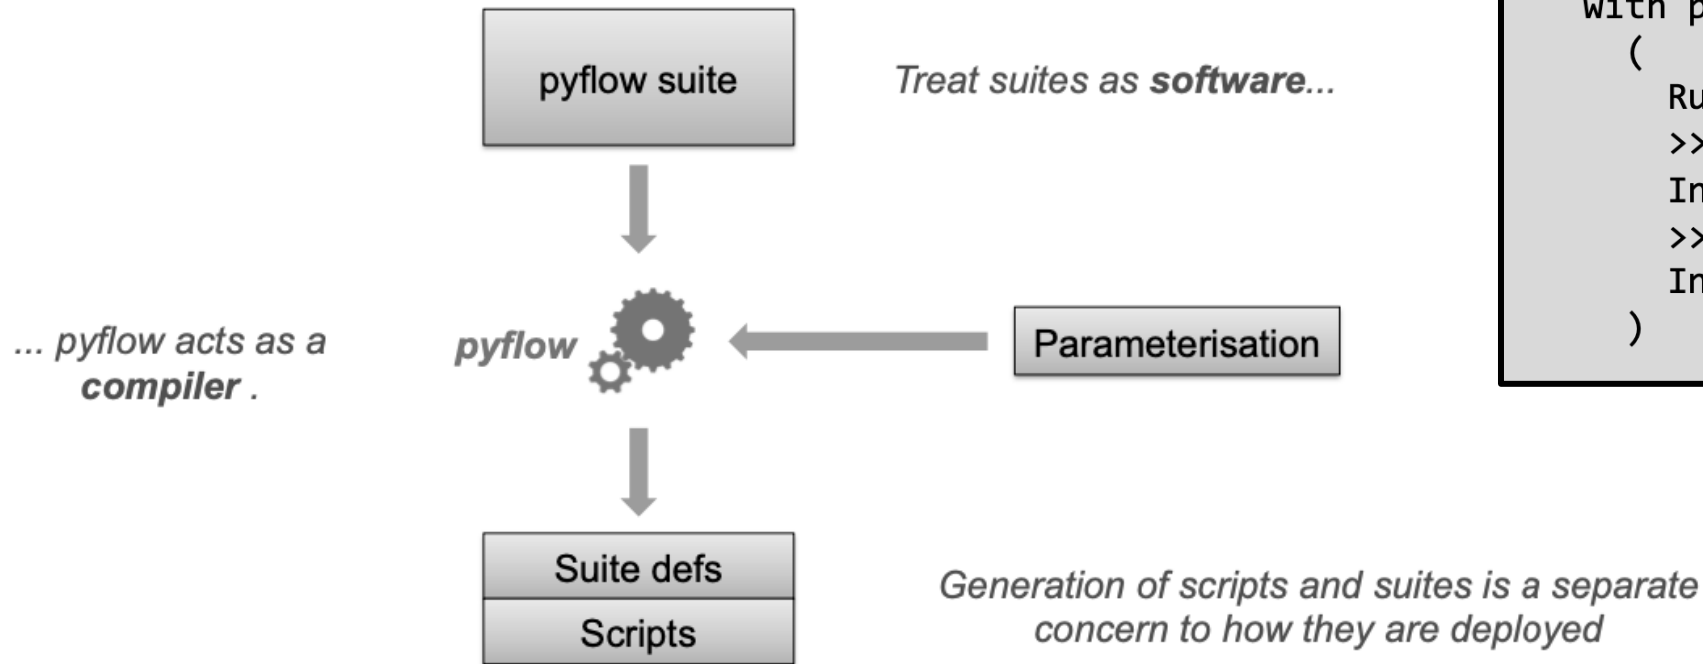
Path	Type	Duration	Allowed	Password	Pid	Host	Try no	Action	Child cmd	Ca
/map/lorenz/compute	user	65 s	300 s	p6Sodvo6	1867810	ac6-183.bullx	3	auto-block	init	5

ecFlow: Best Practices – Operational tips

- Multiple tabs with ecflow_ui vs multiple windows
 - Operators' view tab (only active/submitted tasks)
- Multiple nickname for same server:
 - Different settings: dev/test/prod

Best tools: Pyflow, a pythonic interface to ecFlow

The Tao of Pyflow



```
with pf.Suite('suite'):


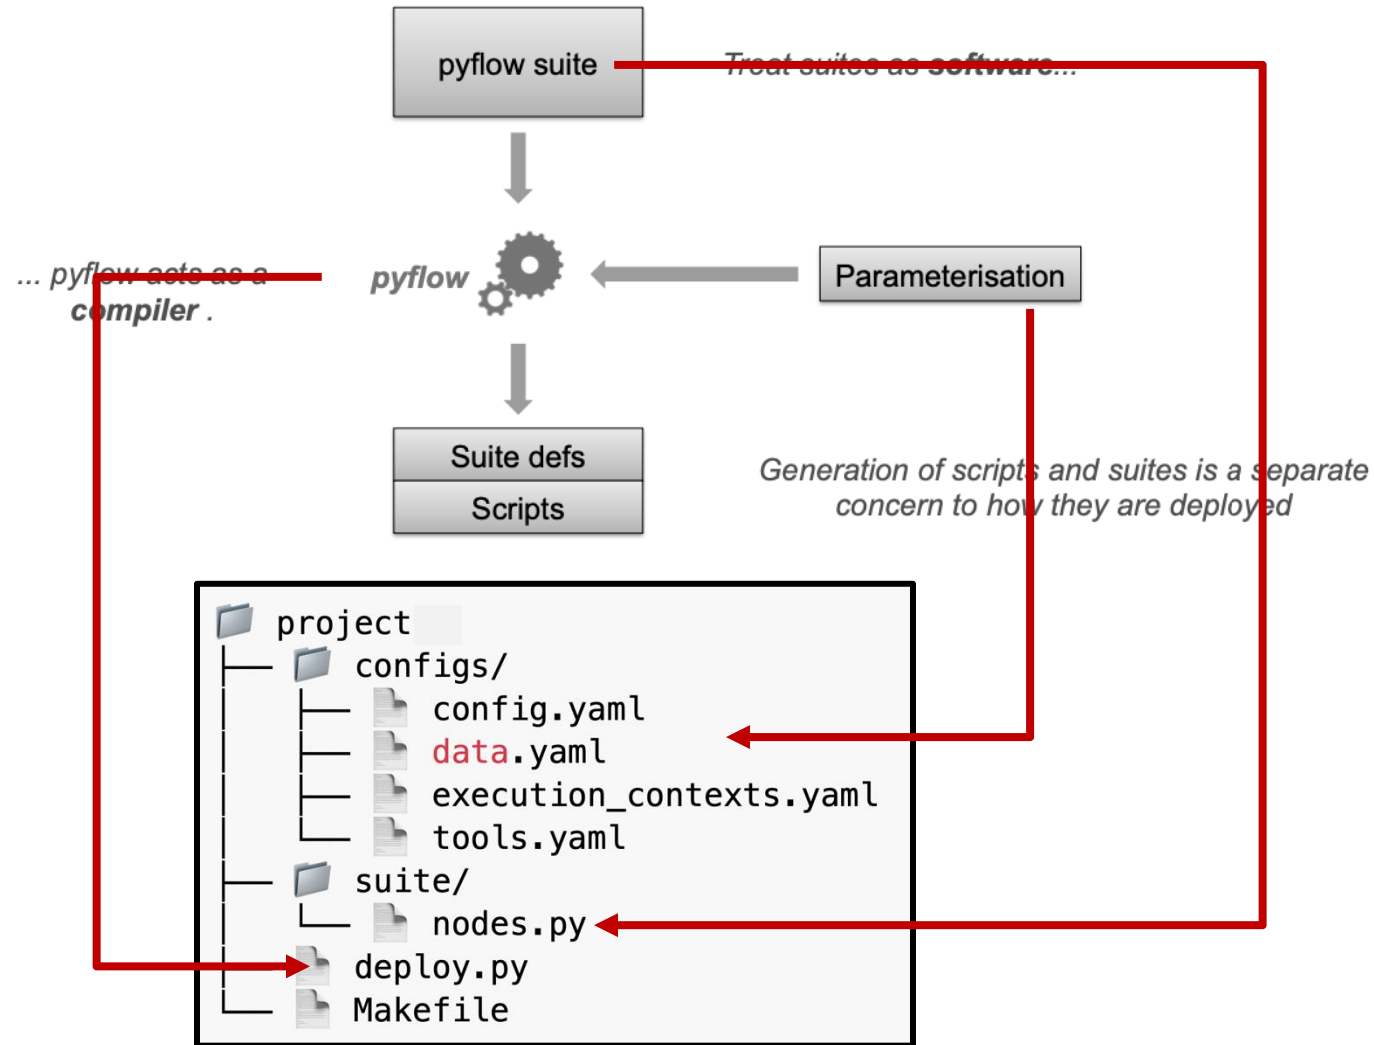
    DeploymentFamily(config)

    with pf.Family('tests', FLAG=123) as f:
        (
            RunUnitTests(config)
            >>
            IntegrationTest1(config, f.FLAG)
            >>
            IntegrationTest2(config, f.FLAG)
        )
```



PYFLOW

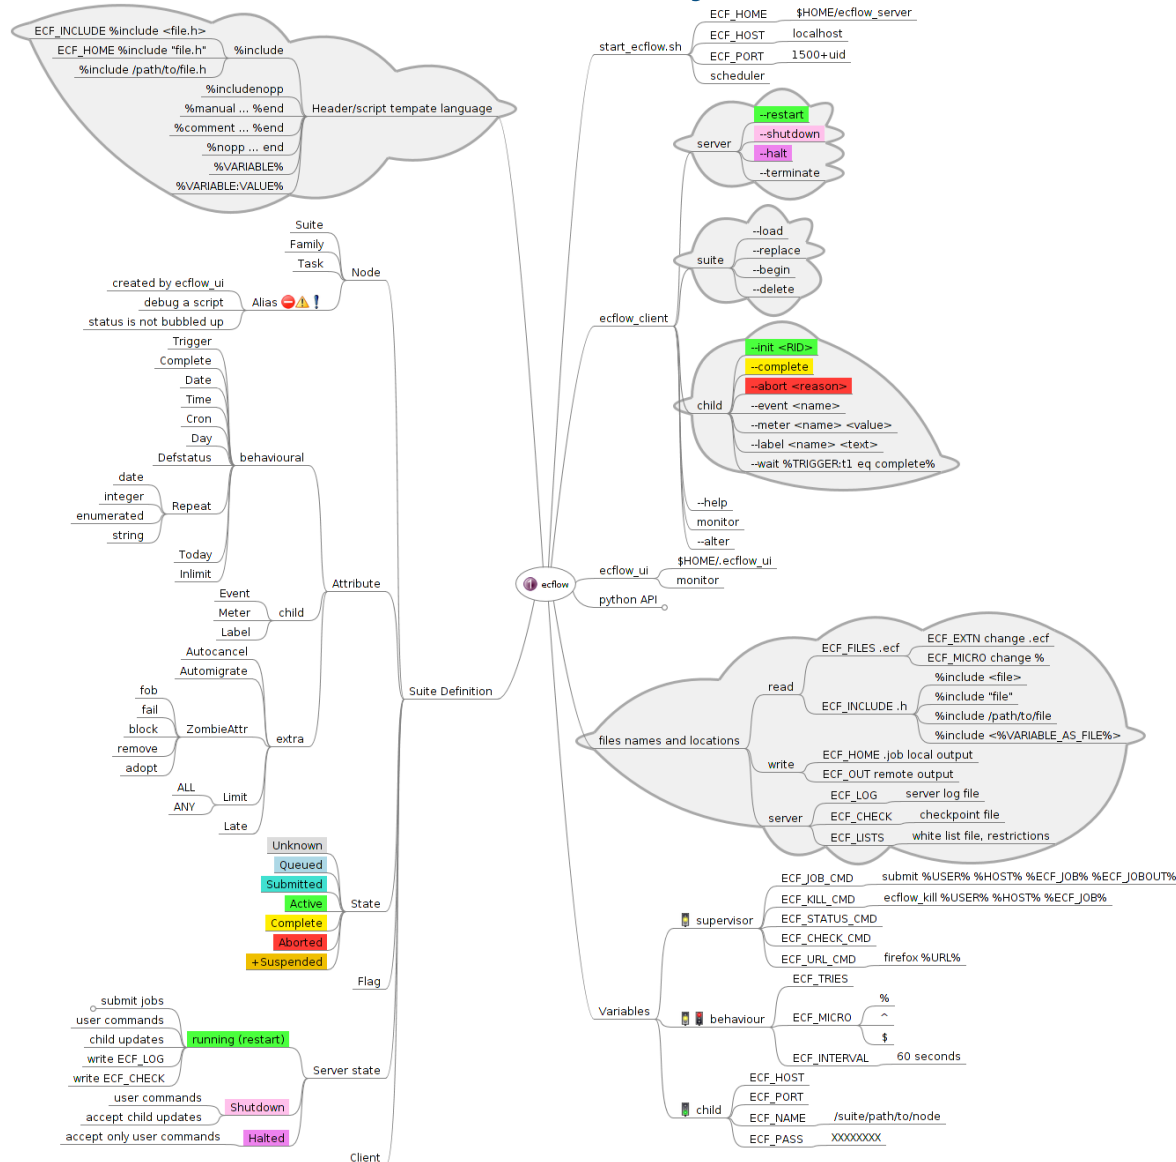
ecFlow: best tools, Wellies: build consistent suites through YAML configuration



ecFlow – best practices – check-list

- Re-runability
- Look after critical data - HA systems, backups
- Limit number of languages used
- Be careful with error trapping
- All variables should to be set (use default values %VARIABLE:default_value%)
- Use a generic user - identify operations
- Works on multiple systems
 - ECF_JOB_CMD
- Design based on constraints
 - Staff availability
- Avoid accessing off-line data in critical path
- Avoid NFS mounted files or unsafe file-systems (SCRATCH)
- Tasks can be serial or parallel
 - don't do serial things in parallel tasks
- Use generic directories to simplify cleaning and always clean up!
- Check task runtimes
- Keep output and job files
- Always use a SCM system and test
 - Test ecflow server/suites

Thanks for your attention!

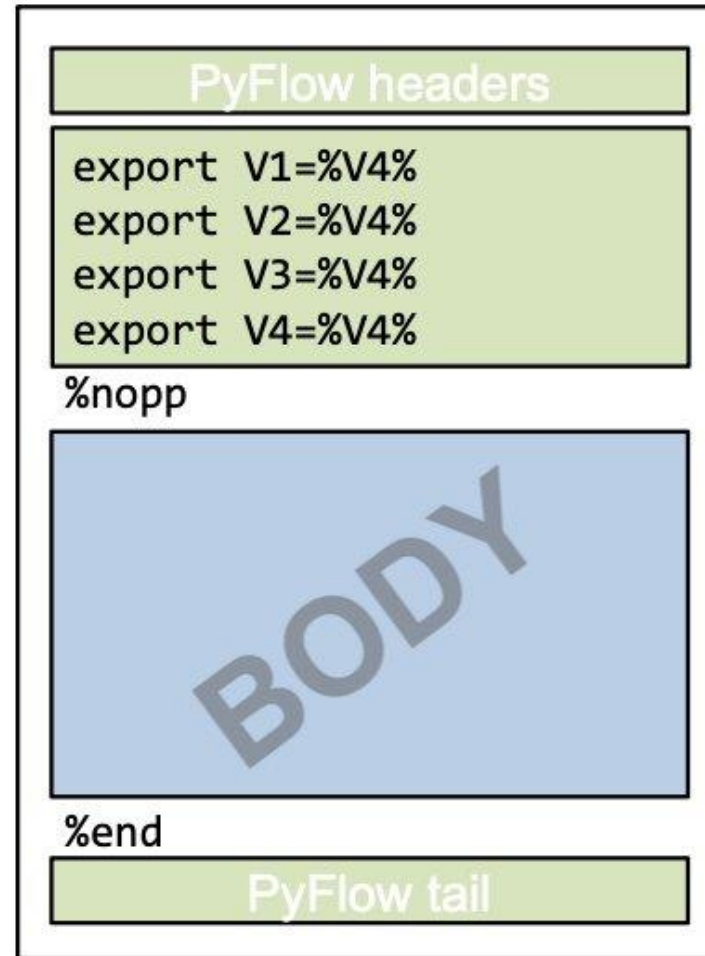


Practice with eLearning Jupiter lab [notebook](#)
ecFlow flashcard and [Quizz](#)?

Time for Questions?

Gratitude to Corentin Carton de Wiart for pyflow/Wellies aspects,
Gratitude to Christopher Barnard for the monitoring aspects
Image on slide 26 is thanks to Google Gemini

Pyflow task wrapper format



ecFlow: definition-file, checkpoint-file, nodes + attributes

Checkpoint-file written by ecflow_server

- a definition file
- defs, enddef, history additional keywords
- recent values for states and variables, next run time in comment

Nodes:

- suite, family, task
- (endsuite, endfamily, endtask)

Attributes can be classified in multiple ways:

- Active/passive (task requeued)
- Related with child command or not
- Behavioural: defstatus, complete

Looping

- repeat, Cron, time, today, date, day, defstatus, autocancel

Scheduling attributes

- trigger, complete, limit, inlimit

Informational attributes

- label, zombie

used in jobs

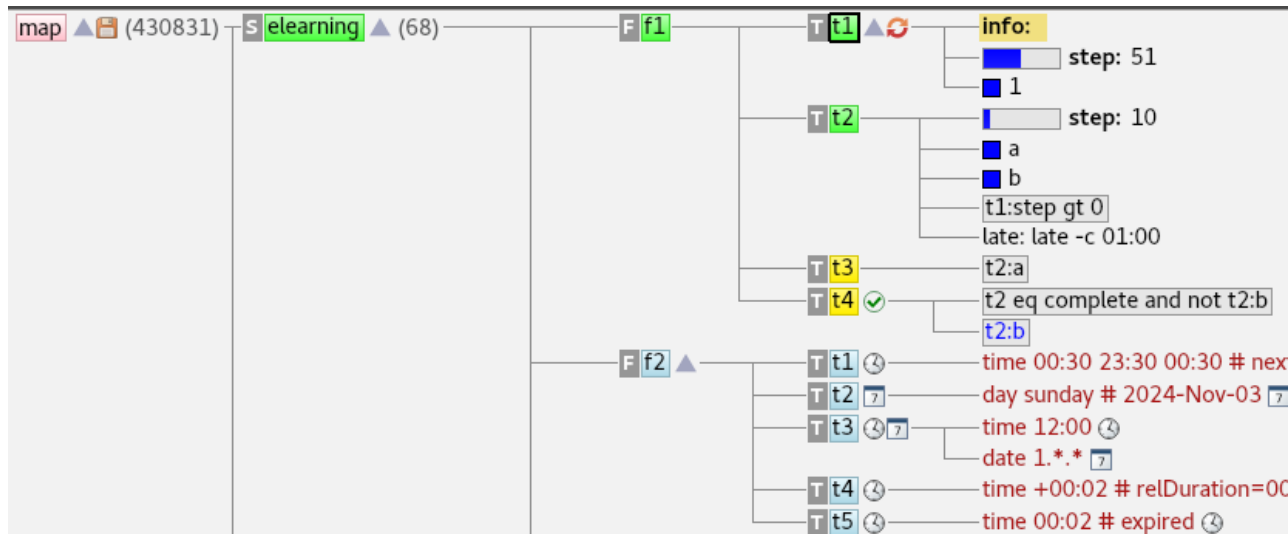
- edit (variable)

used in trigger

- Node status, variable, event, meter, limit, late

ecFlow: Inheritance status v variables v dependencies

- Status inheritance is bubbling up
 - A suite or family node reflects most important status
 - server node status can be
 - **Halted**: accept only user commands
 - **Shutdown**: accept user and child commands
 - **Running**: additionally, jobs can be submitted
- variables inheritance is top down
 - A Variable can be redefined lower in the tree
 - Lowest value prevail for jobs creation
- dependencies can be defined on any node
 - Trigger, complete, time, date, Cron attribute
 - All conditions must be true to create a job
 - High dependency will hide the lower
 - Trigger, complete attribute are instantaneous
 - Date, time, Cron attribute have memory



ecFlow: Tasks wrappers / Tasks headers

key variables

- ECF_EXTN: wrapper extension .ecf .sh .py
- ECF_FILES: wrappers location (r)
- ECF_INCLUDE : headers location (r)
- ECF_HOME: where .job are created (w)

Tasks wrappers

- a template script
- describe generic or specific work to do

Tasks headers

- head.h / qsub.h / tail.h
- %include <%QSUB_H:qsub.h%>

ECF_MICRO % character: variable/block/keyword

- %VARIABLE:default_value%
- manual, nopp, comment,
- include, includenopp
- global scale or locally in the template script: %ecf_micro \$

Tolerance for failures (hardware and software):

- ECF_TRIES: number of automatic rerun
- ECF_TRYNO: job instance number
- Watchdog task to handle known issues

```
1 #!/bin/bash
2 # a task wrapper file to be turned into a job by ecflow
3 # include file located in ECF_INCLUDE directory: qsub + trapping (abort) + in
4 %include <head.h>
5 %manual
6     manual section
7 %end
8 %comment
9     comment section
10 %end
11 # we may need to include a header file, WITHOUT preprocessing
12 %includenopp <compute.sh>
13 %nopp
14     # no preprocessing in this section
15 %end
16
17 echo a variable %STEP% with no default value shall be found in py-def
18 # edit STEP 120 # for example, expected in definition file
19
20 echo a variable %PARAM:Z% with a default value Z, can be omeclieftted in py-def
21 # call ecflow client --complete # cleanup:
22 %include <tail.h>
```

```
1 #!/bin/bash
2 %include <qsub.h>
3 set -e # stop the shell on first error
4 set -u # fail when using an undefined variable
5 set -x # echo script lines as they are executed
6 # Defines the variables that are needed for any communication with
7 export ECF_PORT=ECF_PORT # The server port number
8 export ECF_HOST=ECF_HOST # where the server is running
9 export ECF_NAME=ECF_NAME # The name of this current task
10 export ECF_PASS=ECF_PASS # A unique password
11 export ECF_TRYNO=ECF_TRYNO # Current try number of the task
12 export ECF_PID=$$ # record the process id. Also used for
13 # zombie detection
14 # Define the path where to find ecflow client
15 # make sure client and server use the "same" version.
16 # Important when there are multiple versions of ecflow
17 export PATH=/usr/local/apps/ecflow/ECF_VERSION/bin:$PATH # on HPC
18 # export PATH=$PATH:/usr/local/apps/ecflow/bin:/usr/local/bin
19 # Define a error handler
20 ERROR() {
21     set -e # Clear -e flag, so we don't fail
22     wait # wait for background process to stop
23     ecflow_client --abort-trap # Notify ecflow that something went
24     # wrong, using 'trap' as the reason
25     trap 0 # Remove the trap
26     exit 0 # End the script
27 }
28 trap ERROR 0
29 # Tell ecflow we have started
30 ecflow_client --init=$$ ; set -m
```

```
1 wait
2 ecflow_client --complete
3 trap 0
4 exit 0
```


ecFlow: child commands

- ecflow_client called from a job
- 4 variables:
 - ECF_NAME: path for the node in the definition tree
 - ECF_HOST,
 - ECF_PORT,
 - ECF_PASS:
 - unique pseudorandom key for current job.
 - Zombie flag is raised when incorrect.
 - set to FREE to rescue a child, or in monitoring mode

```
ecflow_client --help child
```



Update status:

- init <jid>
- complete
- abort <reason>

Update attribute:

- meter <name> <value>
- event <name>
- label <name> <msg>

Embedded trigger:

- wait <expression>

Write into server log:

- msg <text>

Get an item from a list: queue


- queue <name> <list> # def-file

ecflow_client: child

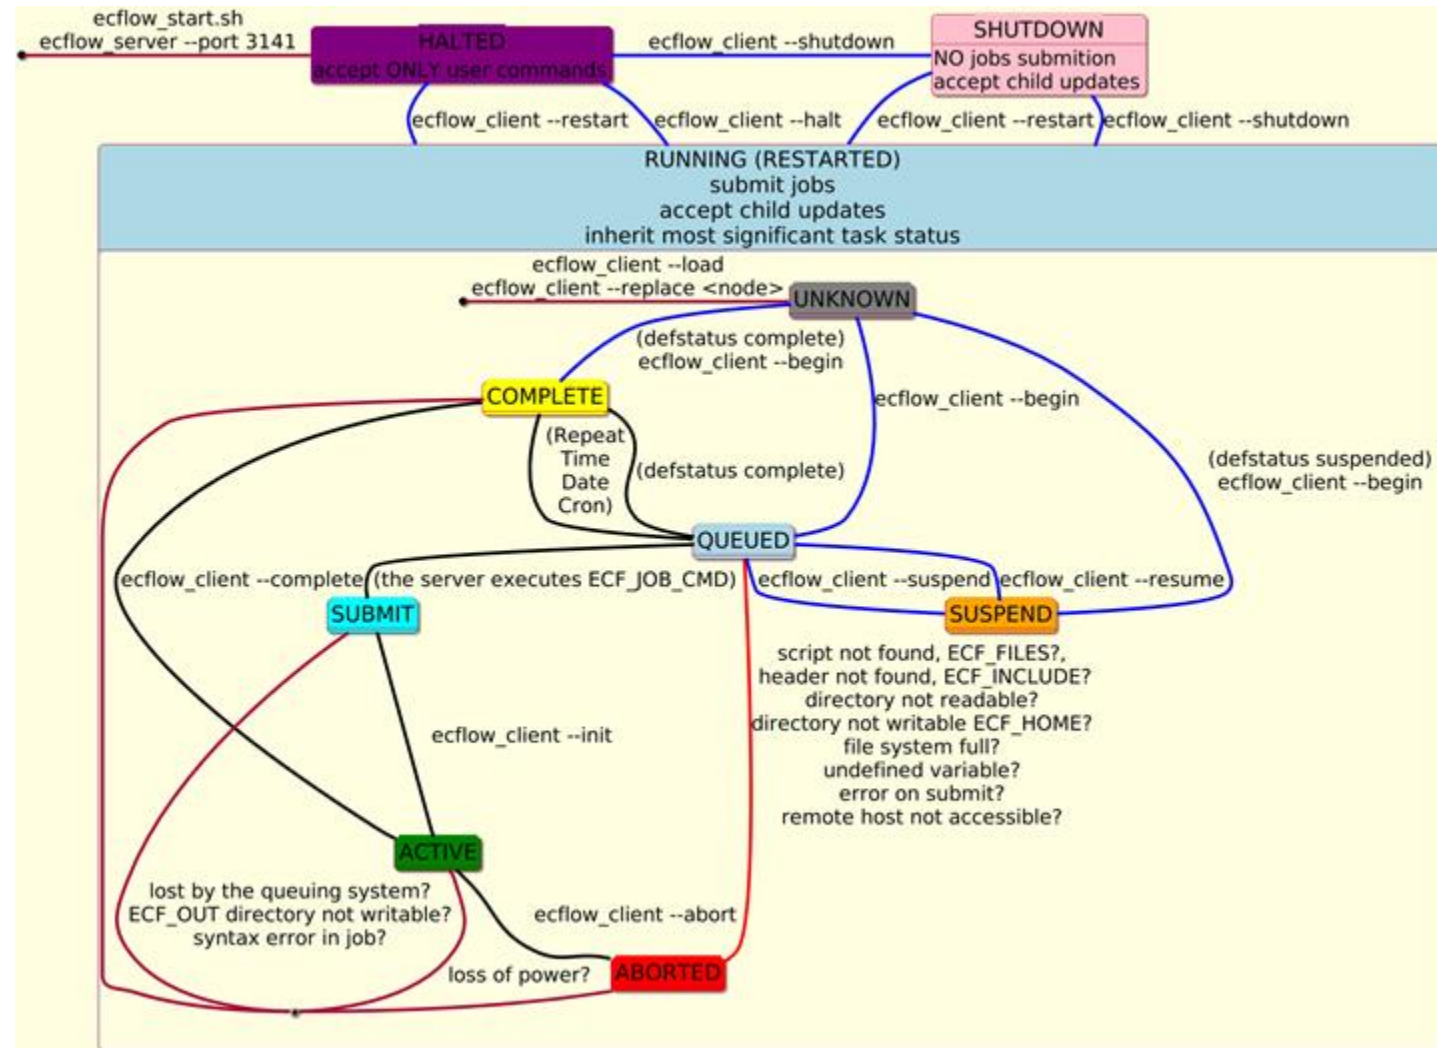
- Update status: **--init** \$JOBID **--abort** "\$REASON" **--complete**
- Attribute update: **--event** <name> **--meter** <name> <value> **--label** <name> "\$message"
- Child commands are 'privileged' i.e. blocking
 - Some may prefer UDP (for labels, meter, event when not critical)
 - **--wait** \$expression
 - Is blocking embedded trigger, that may cause a live lock
 - Some will prefer the non blocking **--query** \$node to check status or a variable
 - token=\$(ecflow_client **--queue** <name>)
 - get a token, from a queue attribute (definition)
- When parent server is not responding:
 - **ECF_HOSTFILE**=/path/to/file
 - The child will look if this variable is present, if the file is accessible, and will contact the listed alternative servers for adoption. Backup server shall already have a (suspended) suite matching for this job, zombie **icon** shall be displayed, and the task shall be listed in the **zombies' panel**

ecFlow: users use case

- ecFlow server is hosted in a dedicated VM
 - ping ecflow-gen- $\{\text{USER}\}$ -001
- ecflow_ui is run on VDI (or laptop, or HPC)
- Jobs are submitted on HPC
- \$HOME is common between VM, HPC, VDI
 - .check, .log under \$HOME/ecflow_servers
 - File ecf.lists to grant or refrain access (rw/r/none)



ecFlow: statuses diagram



Ecflow: python – Perl –ruby – bash – na(t)ive example

- A generic wrapper: body.ecf
- **Headers** included are defined **as variables** in the definition
- **ECF_MICRO** may change (% default, used for ksh/bash)
 - ^ for Perl
 - \$ for python
- Script extension can change using **ECF_EXTN**:
 - .ecf (default), .ecg (generated template), .ech (preprocess to shell), .erb (to ruby), epy, epl
 - .rb (pure ruby), .pl, .py, .sh
- **Payload** defined as SCRIPT variable
 - To be pre-processed with %include
 - Or not, using %includenopp
- An empty file (touch): **pure**
- QSUB_H empty for local submit, or can be defined as slurm.h, pbs.h, qsub.h provided as needed

```
1 #!/usr/bin/env %SHELL:bash%
2 # generic task template - wrapper
3 %include <%QSUB_H:pure%>
4 %include <%HEAD_H:head.h%>
5 %include <%SCRIPT:pure%>
6 %include <%TAIL_H:tail.h%>
```

Ecflow: python – na(t)ive example

- Python def / head / body / tail

```
1 #5.14.1
2 suite starter
3   defstatus suspended
4   family example
5     edit ECF_MICRO '%'
6     family language
7       edit ECF_FILES '/path/to/example/lang'
8       edit ECF_INCLUDE '/path/to/example/test/lang'
9       edit BBIN '/home/linuxbrew/.linuxbrew/bin/'
10      family python
11        edit SHELL 'python3'
12        edit ECF_MICRO '$'
13        edit ECF_EXTN '.epy'
14        edit HEAD_H 'head.epy'
15        edit TAIL_H 'tail.py'
16        edit ECF_JOB_CMD '$ECF_JOB$ 1> $ECF_JOBOUT$ 2>&1'
17        task doit
18          edit SCRIPT 'hello_world.py'
19      endfamily
20
```

```
1 #!/usr/bin/env $SHELL:python3$
2 $include <$QSUB_H:pure$>
3 $include <$HEAD_H:head.epy$>
4 $includenopp <$SCRIPT:pure$>
5 $include <$TAIL_H:tail.py$>
```

```
1 #!/usr/bin/env python3
2 print("Hello world");
```

```
1 # managed by atexit
```

```
1 child.report("complete")
```

```
1 class Child(object):
2     # ...
3     def signal_handler(self, signum, frame):
4         """ catch signal """
5         print("Aborting: Signal handler called with signal ", signum)
6         self.report("abort", "Signal handler called with signal " + str(signum))
7
8     def __exit__(self, exc_type, exc_value, traceback):
9         self.report("abort", "__exit__")
10
11     def report(self, msg, meter=None):
12         """ communicate with ecFlow server """
13         if msg in ("stop", "complete"):
14             self.client.child_complete()
15             self.client = None
16             sys.stdout.flush()
17             sys.stderr.flush()
18             print("#MSG: stop")
19             sys.exit(0)
20         elif msg in ("abort",):
21             self.client.child_abort()
22             self.client = None
23             raise Exception(msg)
24         elif meter:
25             self.client.child_meter(msg, meter)
26         else:
27             self.client.child_label("info", msg)
```

```
1 #!/usr/bin/env $SHELL:python3$
2 def excepthook(exctype, value, traceback):
3     if exctype == KeyboardInterrupt:
4         if child:
5             child.report("abort", "keyb")
6     else:
7         sys.__excepthook__(exctype, value, traceback)
8         if child:
9             child.report("abort", "gen")
10
11 class Child(object):
12     def __init__(self):
13         import signal
14         print("#MSG: kill: ssh %s kill -15 %d" % (os.uname()[1], os.getpid()))
15         for sig in (
16             signal.SIGINT,
17             signal.SIGHUP,
18             signal.SIGQUIT,
19             signal.SIGILL,
20             signal.SIGTRAP,
21             signal.SIGIOT,
22             signal.SIGBUS,
23             signal.SIGFPE,
24             signal.SIGUSR1,
25             signal.SIGUSR2,
26             signal.SIGPIPE,
27             signal.SIGTERM,
28             signal.SIGXCPU,
29             signal.SIGPWR,
30         ):
31             signal.signal(sig, self.signal_handler)
32         self.set_client()
33
34     def set_client(self):
35         self.client = ecflow.Client()
36         host = "$ECF_HOST:$"
37         self.client.set_host_port(host, int("$ECF_PORT:0$"))
38         self.client.set_child_pid(os.getpid())
39         self.client.set_child_path("$ECF_NAME$")
40         self.client.set_child_password("$ECF_PASS$")
41         self.client.set_child_try_no(int("$ECF_TRYNO$"))
42         self.client.child_init()
43         self.client.set_child_timeout(20)
44
```

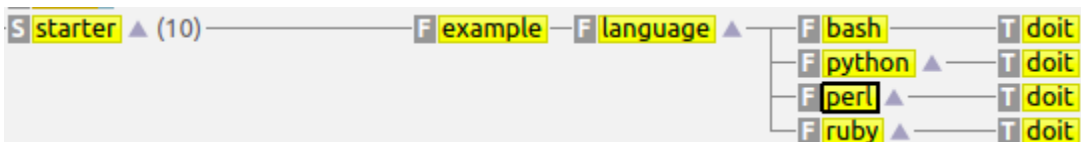
```
1 child = Child()
2 $includenopp <$SCRIPT:pure$>
3 child.report("complete")
```

Ecflow: perl - native example

- Perl

```
1 #5.14.1
2 suite starter
3   defstatus suspended
4   family example
5     edit ECF_MICRO '%'
6     family language
7       edit ECF_FILES '/path/to/example/lang'
8       edit ECF_INCLUDE '/path/to/example/test/lang'
9       edit BBIN '/home/linuxbrew/.linuxbrew/bin/'
10      family perl
11        edit SHELL 'perl'
12        edit ECF_MICRO '^'
13        edit ECF_EXTN '.epl'
14        edit HEAD_H 'head.epl'
15        edit TAIL_H 'tail.pl'
16        edit ECF_JOB_CMD '^ECF_JOB^ 1> ^ECF_JOBOUT^ 2>&1'
17        task doit
18          edit SCRIPT 'hello_world.pl'
19      endfamily
```

```
1 #!/usr/bin/env ^SHELL:perl^
2 ^include <^QSUB_H:pure^>
3 ^include <^HEAD_H:head.h^>
4 ^includenopp <^SCRIPT:pure^>
5 ^include <^TAIL_H:tail.h^>
```



```
1 #!/usr/bin/env perl
2 # head.pl
3 use strict;
4
5 # my $ECF_PORT=^ECF_PORT:0^;
6 $ENV{'ECF_PORT'} = "^ECF_PORT:0^"; # port
7 $ENV{'ECF_HOST'} = "^ECF_HOST:0^"; # host
8 $ENV{'ECF_NAME'} = "^ECF_NAME:0^"; # task path
9 $ENV{'ECF_PASS'} = "^ECF_PASS:0^"; # password
10 $ENV{'ECF_TRYNO'} = "^ECF_TRYNO:0^"; # job number
11 sub xinit() { system("^BBIN:^ecflow_client --init $$"); }
12 sub xabort() { system("^BBIN:^ecflow_client --abort $$"); }
13 sub xcomplete() { system("^BBIN:^ecflow_client --complete $$"); }
14 sub xmeter($$) { my $name=shift; my $value=shift;
15                 system("^BBIN:^ecflow_client --meter $name $value"); }
16 sub xevent($) { my $n=shift;
17                system("^BBIN:^ecflow_client --event $n"); }
18 sub xlabel($$) { my $name=shift; my $value=shift;
19                 system("^BBIN:^ecflow_client --label $name $value"); }
20 xinit();
21 eval ' ';
```

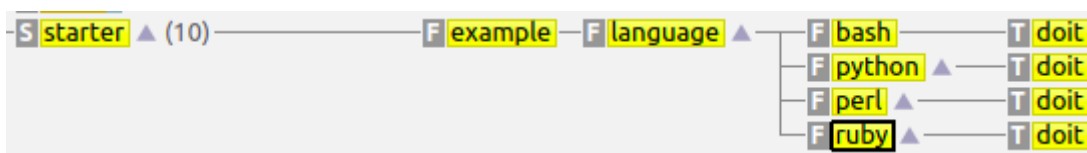
```
1 #!/usr/bin/env perl
2 print("Hello World\n");
```

```
1 # tail.pl
2 ;
3 if ($@){
4   print "# caught signal: $@\n";
5   xabort();
6   exit;
7 }
8 print "# the job is now complete\n";
9 xcomplete();
10 exit;
```


ecFlow: ruby – native example

- Ruby head / body / tail

```
1 #5.14.1
2 suite starter
3   defstatus suspended
4   family example
5     edit ECF_MICRO '%'
6     family language
7       edit ECF_FILES '/path/to/example/lang'
8       edit ECF_INCLUDE '/path/to/example/test/lang'
9       edit BBIN '/home/linuxbrew/.linuxbrew/bin/'
10      family ruby
11        edit SHELL 'ruby'
12        edit ECF_MICRO '^'
13        edit ECF_EXTN '.erb'
14        edit HEAD_H 'head.erb'
15        edit TAIL_H 'tail.rb'
16        edit ECF_JOB_CMD '^ECF_JOB^ 1> ^ECF_JOBOUT^ 2>&1'
17        task doit
18          edit SCRIPT 'hello_world.rb'
19        endfamily
20      endfamily
21    end
22  end
```



```
1 #!/usr/bin/env ^SHELL:ruby^
2 ^include <^QSUB_H:pure^>
3 ^include <^HEAD_H:head.h^>
4 ^includenopp <^SCRIPT:pure^>
5 ^include <^TAIL_H:tail.h^>
```

```
1 #!/usr/bin/env ruby
2 print("Hello World\n");
```

```
1 # tail.rb
2 # begin
3 rescue => e
4   puts "caught signal: #{e}"
5   xabort
6   exit
7 end
8
9 puts "the job is now complete"
10 xcomplete
11 exit
```

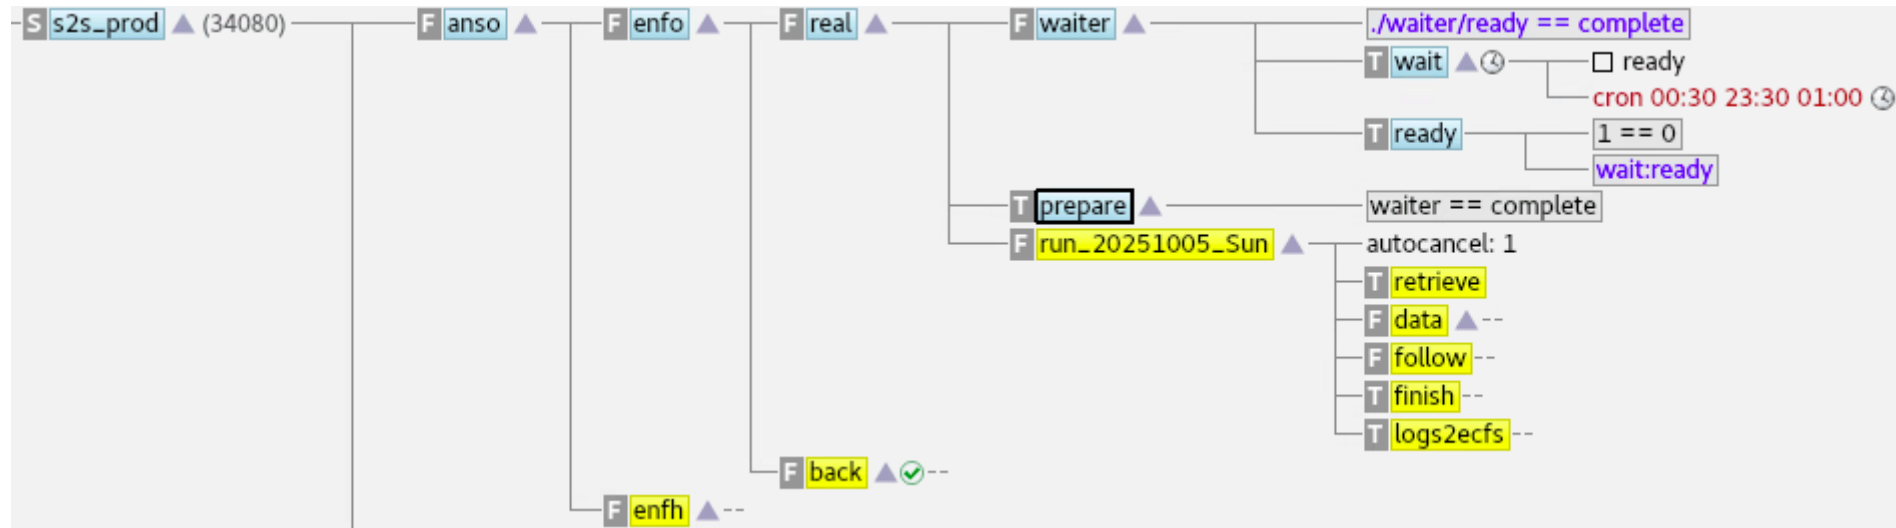
```
1 #!/usr/bin/env ruby
2 # coding: utf-8
3 # head.rb
4 # env variable for child commands
5 ENV['ECF_PORT'] = "^ECF_PORT:0^" # port
6 ENV['ECF_HOST'] = "^ECF_HOST:0^" # host
7 ENV['ECF_NAME'] = "^ECF_NAME:0^" # task path
8 ENV['ECF_PASS'] = "^ECF_PASS:0^" # password
9 ENV['ECF_TRYNO'] = "^ECF_TRYNO:0^" # job number
10
11 def xinit
12   system("^BBIN:^ecflow_client --init #{Process.pid}")
13 end
14
15 def xabort
16   system("^BBIN:^ecflow_client --abort #{Process.pid}")
17 end
18
19 def xcomplete
20   system("^BBIN:^ecflow_client --complete #{Process.pid}")
21 end
22
23 def xmeter(name, value)
24   system("^BBIN:^ecflow_client --meter #{name} #{value}")
25 end
26
27 def xevent(n)
28   system("^BBIN:^ecflow_client --event #{n}")
29 end
30
31 def xlabel(name, value)
32   system("^BBIN:^ecflow_client --label #{name} #{value}")
33 end
34
35 # init
36 xinit
37
38 begin
39
```

%include preprocessing directive

- `%include <file.h>`
 - include a file under ECF_INCLUDE directory
- `%include "file.h"`
 - Include a file below ECF_HOME directory
- `%include /path/to/file`
 - a hardcoded location
- `%include:` NOTE % MUST be first character of the line
 - Avoid complexity, it prevents echo "`%include <file>`"
 - Avoid ambiguity: `# %include <file>`
- `%include <%FILE_H:pure%>`
 - Filename can be provided by a suite variable, here FILE_H
 - `Edit({"FILE_H": "config.oper.h", })`
 - `Edit({"FILE_H": "config.test.h", })`
- `%includeonce <%FILE_H:pure%>`

ecFlow: Best Practices – Suites' design – dynamic families

- A task to decorate a suite with new families
 - Barber shop example
 - S2S: waiter/prepare + autocancel

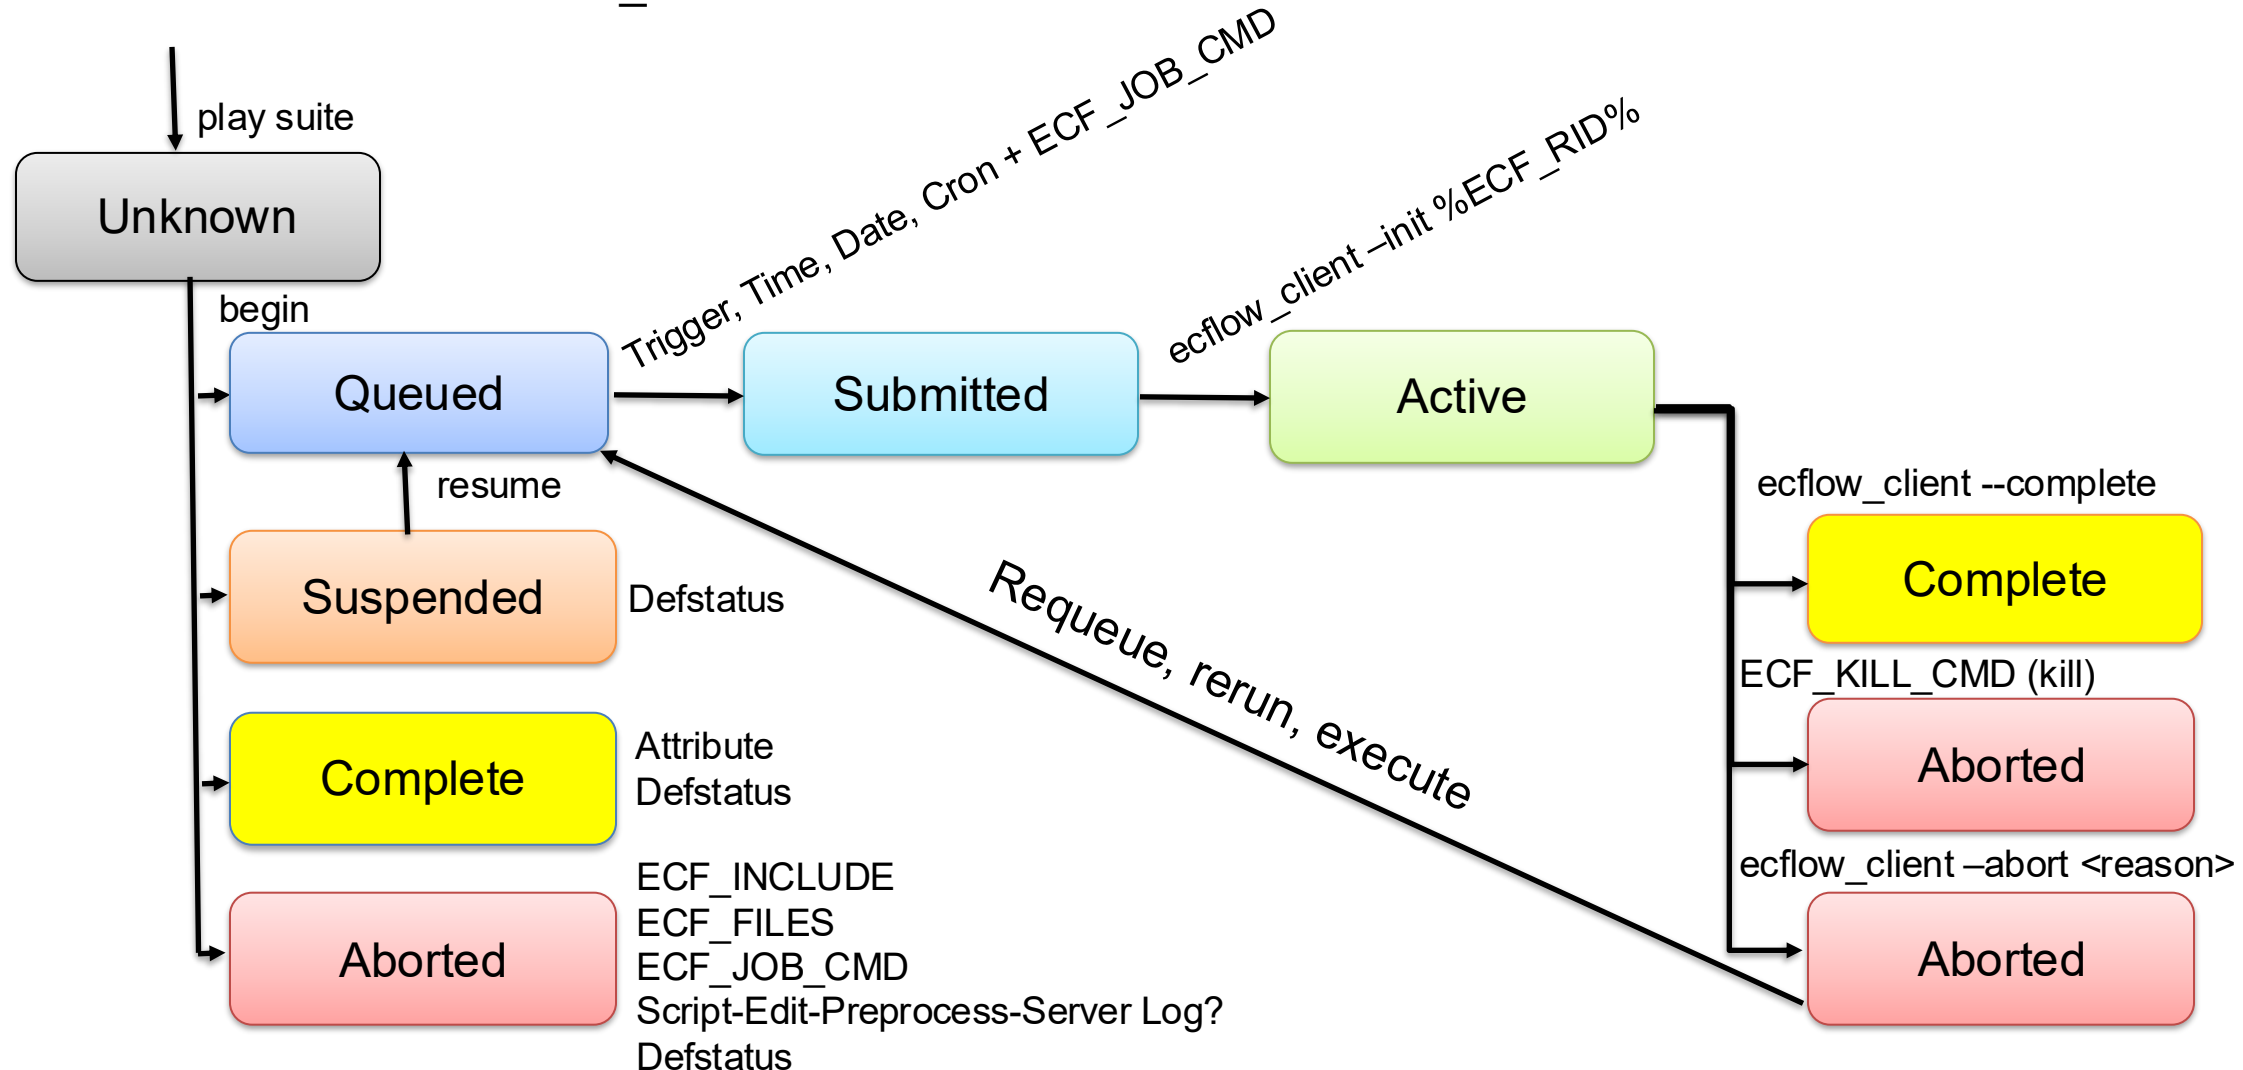


ecFlow: Hands-in

- Definition + scripts under SCM (git)
- ECF_FILES/ECF_INCLUDE may refer to the same directory
- One time vs cyclic script
 - Choose ignition strategy: time condition, event set by third party, a wait barrier
 - Choose looping strategy: asap, or delay looping (keep output visible + rerun capable)
- A simple wrapper: lorenz.ecf

ecFlow: Status

- Task v Node: ecflow_client v inheritance



Pyflow ecosystem links

<https://github.com/ecmwf/pyflow>

<https://github.com/ecmwf/pyflow-wellies>

<https://github.com/ecmwf/tracksuite>

<https://github.com/ecmwf/troika>

<https://pyflow-workflow-generator.readthedocs.io>

<https://pyflow-wellies.readthedocs.io>



WELLIES



PYFLOW

module load wellies/new

