



LFN Developer & Testing Forum

ONAP Stability/resiliency/Stress tests

What are we speaking about?

Natacha Chéreau, Morgan Richomme

Reminders

ONAP stability/resiliency/stress tests: what are we talking about?

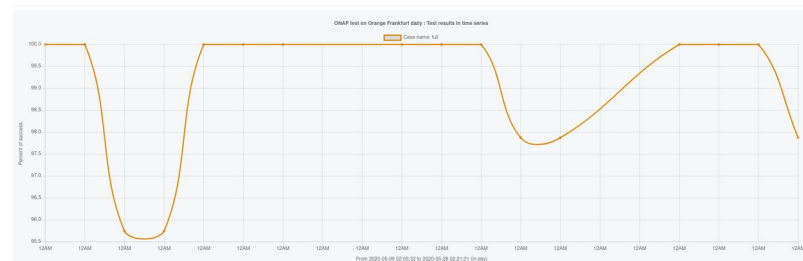
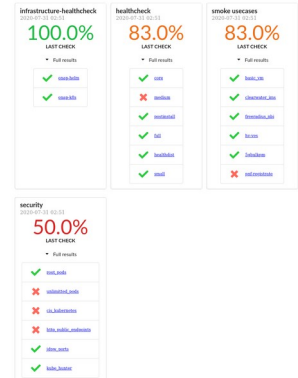
- ONAP Honolulu (8th version) in progress
- Stability/resiliency tests introduced in ONAP by Integration team since Beijing
- No real stress tests on ONAP from an end to end perspective

Stability tests (Beijing - Frankfurt)

- vFWCL use case run continuously during 72h
- Manual check of resources (kubectl top pod /nodes)
- ONAP Maturity Testing Notes created for the release Note
- More details on associated wiki page
- Beijing: <https://wiki.onap.org/pages/viewpage.action?pageId=33064037>
- Casablanca: <https://wiki.onap.org/display/DW/Casablanca+Stability+Testing+Instructions>
- Dublin: <https://wiki.onap.org/display/DW/Dublin+Release+Stability+Testing+Status>
- El Alto: <https://wiki.onap.org/display/DW/El+Alto+Stability+Run+Notes>
- Frankfurt: <https://wiki.onap.org/display/DW/Frankfurt+Stability+Run+Notes>

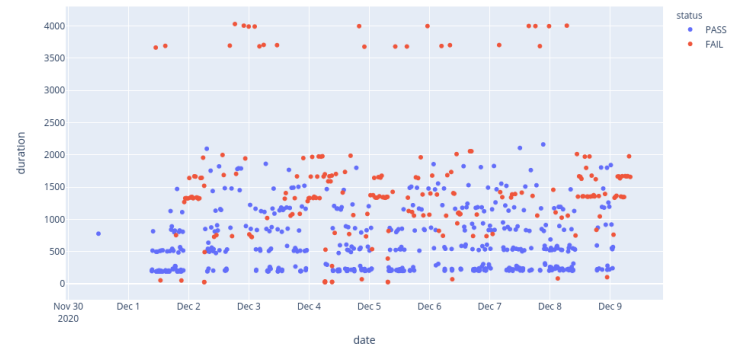

Stability tests: Daily CI/CD

- Introduced in Windriver lab
 - static config of dedicated Integration lab / accessible only through Windriver VPN
 - Monolithic jenkins scripts (installation and tests mixed)
 - Error rate increased over all the releases (some test suites were not maintained)
- Move to external public CI/CD in Frankfurt (gitlab-ci on gitlab.com)
 - Separation CI/CD and tests
 - Dockerization of the tests
 - Public reporting



Stability tests (Guilin)

- 1 day healthcheck control
- CI/CD follow up (with more ONAP dedicated tests – not only use cases)
- 1 week test running continuously basic_vm test (including prometheus/grafana monitoring)



Resiliency tests (Beijing – El Alto)


- Following procedure was performed
 - 1: Run vFWCL
 - 2: Kill a pod
 - 3: Re-Run vFWCL
 - 4: GOTO 1
- Casablanca resiliency testing:
<https://wiki.onap.org/display/DW/Casablanca+Maintenance+Release+Resiliency+Testing+Status>
- test run only once few days before the release
- Several “transient lab infra issues” ??

What can we really say about stability/resiliency for ONAP solution?

Critical questions

- Can ONAP survive the restart of a POD ? **most of the time yes, it depends of the pod..**
- Can ONAP survive a helm chart redeployment ? **most of the time yes but not 100% guarantee**
- Can ONAP survive the restart of a k8s controller ? **Not sure, as ONAP is not cloud native yet, the full promise of k8s is not met yet.**
- Can ONAP survive the restart of a k8s worker ? **not sure especially if some DBs are on this host. One more time the promise of k8s is not met.**
- Do we have figures on an ONAP solution for a given lab: number of simultaenous onboarding/instantiation/loop ? **not really**
- Can we perform a Backup&Restore (disaster scenario) ? **maybe with pain**
- Can we perform a smooth upgrade ? **clearly no**
- Can the community version be used in production ? **depend how brave you are but reasonably not yet**

Honolulu

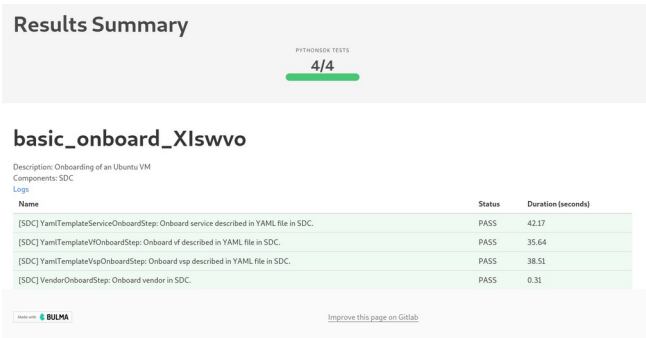


Stability/resiliency KPI definition

- Integration team started defining some KPIs
- Operational teams contacted to draw a realistic scenario
- Small developments to reuse existing use cases and build scenario to demonstrate the KPI
- Introduction of weekly CI/CD chains to perform the tests on weekly master (not only the 3 days before the release...one additional arguments for the projects to release more frequently...)
- Wait for additional Operational team feedback

Onboarding campaign 1/2

- Run simultaneous basic_onboard tests
- KPI ~ 10 simultaneous onboarding looks reasonable
- Performed manually beginning of January (10 series of 5 simultaneous onboarding and 5 series of 10 simultaneous onboarding)
- To be included as a weekly test in weekly CI/CD chains



5 parallel onboarding (10 series)

criteria \ Serie	1	2	3	4	5	6	7	8	9	10	Global
Success rate (%)	100	100	100	100	100	100	100	100	100	100	100
Min duration	27:39	10:24	10:15	10:26	11:18	07:42	07:54	08:05	08:35	08:20	07:42
Max duration	27:43	10:36	10:17	10:27	11:22	07:53	08:00	08:19	08:42	08:44	27:43
Average duration	27:41	27:40	10:16	10:27	11:20	07:48	07:58	08:12	08:39	08:42	11:09
Median duration	27:41	10:26	10:16	10:27	11:19	07:49	07:59	08:13	08:40	08:38	09:30
Comments/Errors	/	/	/	/	/	/	/	/	/	/	/

10 parallel onboarding (5 series)

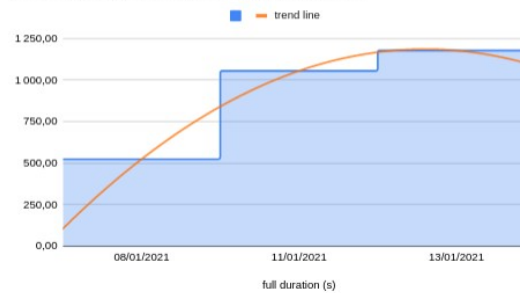
criteria \ Serie	1	2	3	4	5	Global
Success rate (%)	100	100	100	100	100	100
Min duration	16:04	15:24	16:32	19:40	19:07	15:24
Max duration	16:22	17:10	17:36	20:01	19:50	20:01
Average duration	16:15	16:51	17:23	19:52	19:46	18:00
Median duration	16:20	17:08	17:33	19:53	19:38	17:33
Comments/Errors	/	/	/	/	/	/

Onboarding campaign 2/2

Full duration mean over time (series of 5)



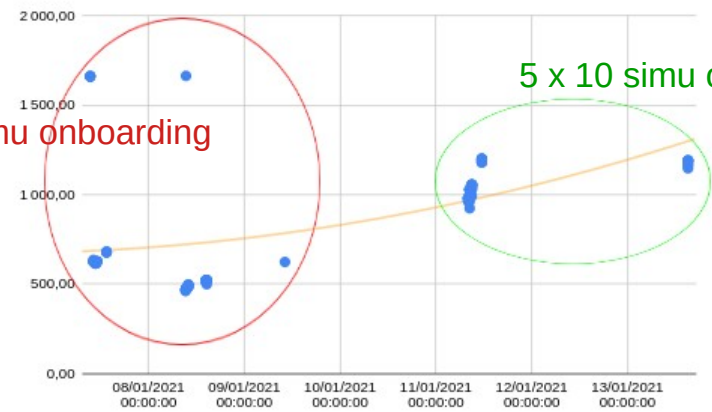


Full duration mean over time (series of 10)



10 x 5 simu onboarding


5 x 10 simu onboarding




- 10 simultaneous onboarding KPI OK
- The error seen in CI/CD not reproducible – no race conditions seen
- Number of existing models has no impact – duration depends on the state of the cluster
- Duration seems to increase linearly with the number of parallel instances
- Additional tests could be done to find a limit that is linked to Hardware/lab capabilities
- Campaign to be planned for basic_vm to stress the SO/SDNC

Resiliency

- Disaster scenario needed (reboot of a worker/controller)...
- Reused of k8s chaos framework to be tested in Honolulu – any help welcome
- Resiliency scenario could be run on master-weekly, master-guillin chains
- Disaster Backup&Restore will be also interesting to evaluate the gap between the k8s promise and the reality



kube-monkey



Chaos Mesh®



Conclusions

- Lots of work done since Amsterdam
- Still lots of work ahead to get more stability resiliency insights
- **Open Questions**
 - shall we work on that at the community level (more resources needed) or shall we consider third party partners to do that assuming that Service Providers will have different backup&restore policies/strategies leveraging existing procedures/tools
 - K8S/ONAP boundaries regarding stability/resiliency not crystal clear..cloud readiness is a must have to leverage k8s features...shall we perform such tests until cloud readiness level is acceptable ?



OLF NETWORKING

LFN Developer & Testing Forum