FED4FIRE
Federation for FIRE Plus

LiSCIo

Lightweight Self-adaptive Cloud-IoT Monitoring across Fed4FIRE+ Testbeds

ANTONIO BROGI

Department of Computer Science, University of Pisa, Italy

FEC9

Online, May 26th, 2021

LiSCIo

Experiment Description

Cloud-IoT Infrastructure Monitoring

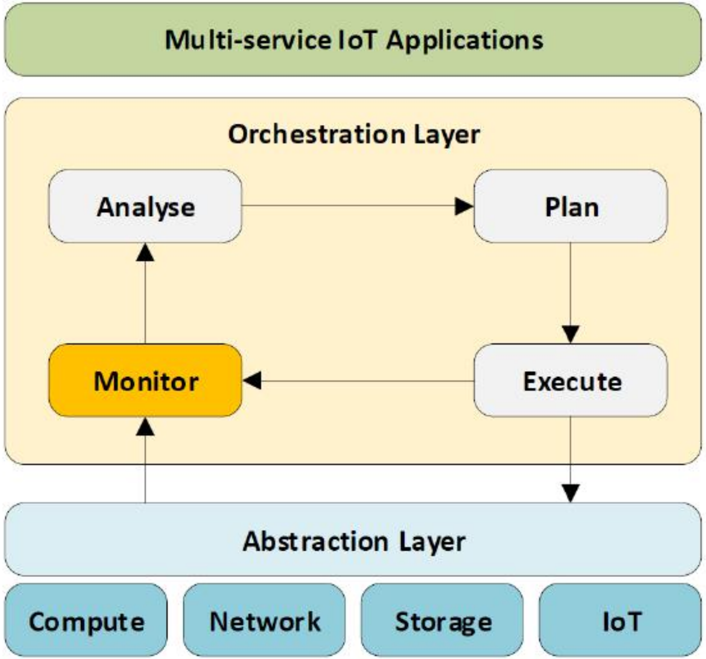

- **Cloud-IoT orchestration**
  - much work on **Analyse**
  - some work on **Plan** & **Execute**
  - less work on **Monitor**

- **Monitor** pivotal to decide
  1. where to deploy app services at first
  2. when/where to migrate app services

FogMon

- Open source lightweight fault-resilient monitoring service for Cloud-IoT infrastructures
- The service monitors:
  - hardware resource availability
  - end-to-end network QoS
- Hybrid overlay network, with $\sqrt{N}$ leaders
- Latency and bandwidth
  - intra-group measurements
  - inter-group estimates

    e.g. Lat(x,y) = Lat(x,L(x)) + Lat(L(x),L(y)) + Lat(L(y),y)
- Assessed in lab environment (13 nodes)

https://github.com/di-unipi-socc/FogMon

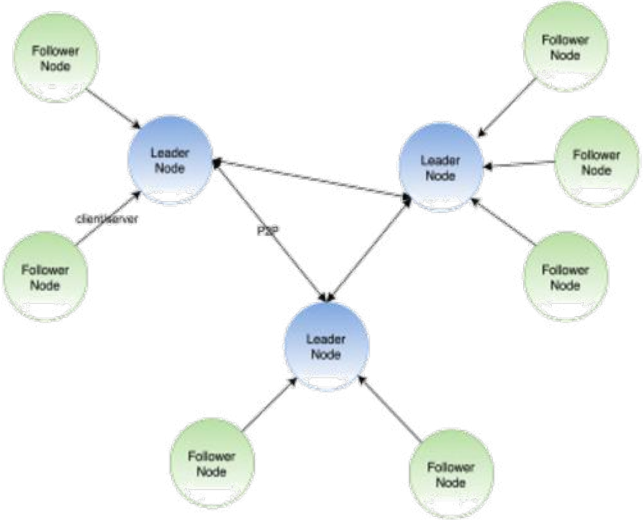

A. Brogi, S. Forti, M. Gaglianese. *Measuring the Fog, Gently*. ICSOC 2019.
S. Forti, M. Gaglianese, A. Brogi. *Lightweight self-organising distributed monitoring of Fog infrastructures*. Future Generation Computer Systems. 2021.

Objectives & Methodology

Testing, assessing and tuning the FogMon service
in medium- to large-scale settings
over heterogeneous Cloud and Edge resources
across two testbeds within the Fed4Fire+ federation

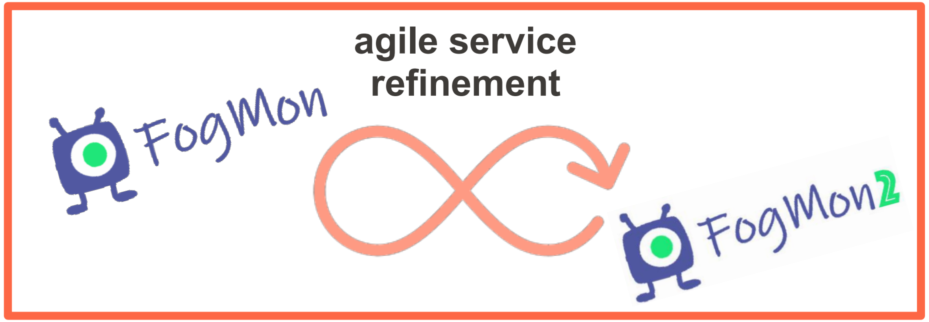

Experiment Setup and Plan

Measuring

- **footprint** on hardware and bandwidth
- **relative error** on measurements & estimates **against setup ground-truth** (configured via GRE-tunnels and tc)
- **time to reach stability**

on

- 3 types of **Follower Failures** (FF)
- 3 types of **Leader & Follower Failures** (LFF)
- 2 types of **Link Failures** (LF)
- **different numbers of Leaders** (LNT)
- **20 (S), 30 (M) and 40 (L) nodes across VIRTUALWALL and CityLab**
- **default** vs **reactive configurations**

	VIRTUALWALL (physical nodes)	CityLab (wireless nodes)
S	10	10
M	15	15
L	30	10

100+ experiments!

	Follower Failures			**Leader-Follower Failures**			**Link Failures**		**Tuning Number of Leaders**		
	FF 1	*FF 2*	*FF 3*	*LFF 1*	*LFF 2*	*LFF 3*	*LF 1*	*LF 2*	*LNT 1*	*LNT 2*	*LNT 3*
FogMon 1.0	S, M, L	S, M, L	S, M, L	S, M, L	S, M, L	S, M, L					
FogMon 1.1							S	S			
FogMon 2.0	S, M, L	S, M, L	S, M, L	S, M, L	S, M, L	S, M, L	S, M, L	S, M, L	S, M, L	S, M, L	S, M, L

LiSCIo

Project Results

FogMon tuning and refinement

FogMon 1.0

- added DNS compatibility
- fixed memory leaks
- fixed differential updates
- improved DB access
- improved dockerisation

FogMon 1.1

- fixed BW measurements
- 3x reduction of BW footprint

FogMon2

- parallelisation of latency & BW tests
- «bad news go fast» → faster reaction to link degradations and failures
- improved passive BW measurements (tuning of Assolo)
- improved accuracy of inter-group BW estimates

$$BW(x,y) = min(max\ BW(x,n(x)),\ max\ BW(y,n(y)),\ BW(L(x),L(y)))$$

Lessons Learnt on FogMon 2.0

- Average relative error on BW estimates from $\simeq$1000% to $\simeq$ 12%
- Average relative error (both on measurements and estimates): $\simeq$10%
- Reduced BW footprint by 70% (default) and 40% (reactive)
- Slightly increased HW footprint (RAM < 30MB, CPU < 4%)
- Reactive faster in identifying changes but more resource intensive
- Shorter time to stability in default configuration
- $\sqrt{N}$ Leaders good compromise in all configurations
- From TRL 4 (lab) to TRL 5 (relevant environment, 40 nodes)

- All experiments repeated on FogMon 2.0
- Further experiments to tune number of Leaders among $\frac{1}{2}\sqrt{N}, \sqrt{N}, 2\sqrt{N}$.

Version & Config	Relative Error				Footprint				Time to stability
	Lat (intra)	*BW (intra)*	*Lat (inter)*	*BW (inter)*	*CPU %*	*RAM*	*BW (Tx)*	*BW (Rx)*	
FogMon 1.x (default)	6.2%	6.23%	10.23%	766.54%	1.4%	12.45 MB	213 KBps	218 KBps	692 s (11 m 32 s)
FogMon 1.x (reactive)	6.55%	5.72%	15.20%	1461.32%	2.25%	13.86 MB	298 KBps	300 KBps	349 s (5 m 49 s)
FogMon 2.0 (default)	2.31%	5.92%	11%	11.33%	1.41%	20.84 MB	65 KBps	66 KBps	572 s (9 m 32 s)
FogMon 2.0 (reactive)	6.44%	6.47%	11.78%	12.28%	3.56%	24.55 MB	167 KBps	173 KBps	396 s (6 m 36 s)

Experiment tooling

Topology Builder

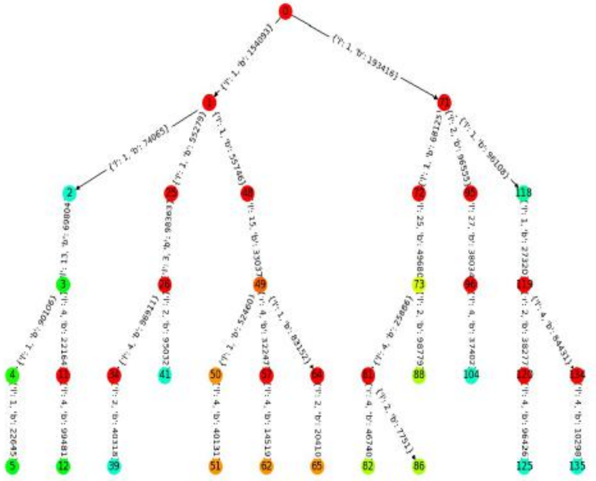

FogMonEye

WebGUI for monitoring LiSCIo experiments

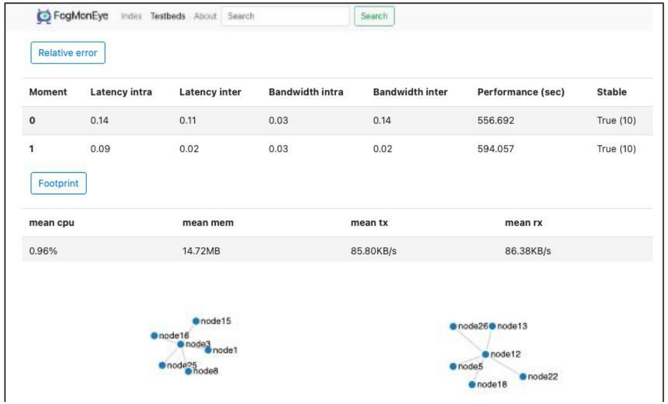

Moment	Latency intra	Latency inter	Bandwidth intra	Bandwidth inter	Performance (sec)	Stable
0	0.14	0.11	0.03	0.14	556.692	True (10)
1	0.09	0.02	0.03	0.02	594.057	True (10)

mean cpu	mean mem	mean tx	mean rx
0.96%	14.72MB	85.80KB/s	86.38KB/s

Business Impact

Impact on research

Leveraging Fed4Fire+ facilities, we have shown that FogMon:

- can be deployed across network boundaries and heterogenous computing capabilities
- detects and adapts to link failures
- exhibits low acceptable footprint on nodes and links at increasing infrastructure sizes (from 20 to 40 nodes)

\+ Improved, assessed and validated FogMon service

Not possible without large-scale heterogeneous federated testbeds!

M. Gaglianese, S. Forti, F. Buti, F. Paganelli, A. Brogi. *Lightweight Self-adaptive Cloud-IoT Monitoring across Fed4FIRE+ Testbeds (LiSCIo) [Dataset]. Available on Zenodo:* http://doi.org/10.5281/zenodo.4682987 (2021)

Gained knowledge & competences

- Experience in using federated distributed testbeds
- New expertise gained in tools for testbed configuration and setup (JFed, Fabric, tc, GRE-tunnels, etc.) and their automation
- Stronger competences in the fields of next-gen Cloud-IoT infrastructures (from monitoring to service placement to SE and virtualisation techniques)

Other scenarios

- Incorporate FogMon in other solutions for the adaptive placement of next-gen applications over the Cloud-IoT continuum
- Collaborate with existing H2020 projects, e.g.

- New collaborations and partnerships in view of next EU project calls
- Publications in international conferences and journals

  M. Gaglianese, S. Forti, F. Buti, F. Paganelli, A. Brogi. *Lightweight Self-adaptive Cloud-IoT Monitoring across Fed4FIRE+ Testbeds.* (WIP.)

Follow-up activities

FED4FIRE

Follow-up Experiments

- Further engineer FogMon (hybrid Cloud-Edge architecture for storing data)
- Alternative methods for estimating bandwidth (e.g. matrix completion) or to identify failures (e.g. ML)
- Improved topology restructuring and Leader selection

New initiatives

- National/international projects on automating the whole lifecycle of applications in the Cloud-IoT continuum
- Start-up spin-off business on Cloud-IoT infrastructure monitoring?

Contamination Lab Pisa

AZIENDA SPINOFF DELL'UNIVERSITÀ DI PISA

LiSCIo

Feedback

Procedure & Administration

- Very adequate and sustainable level of work for admin / feedback / docs
- Lean documents to fill in from proposal to reporting
- Very useful conference calls and meetings

Setup of the experiment

- Only 8 weeks to setup and run the experiment for the first time (installing Docker, learning GRE-tunneling, tc and fabric)
- Another 2 weeks to tune experiment code and automatise whole setup and data collection process
- JFed greatly simplified the above: single point of contact & service

Setup of the experiment (cont.)

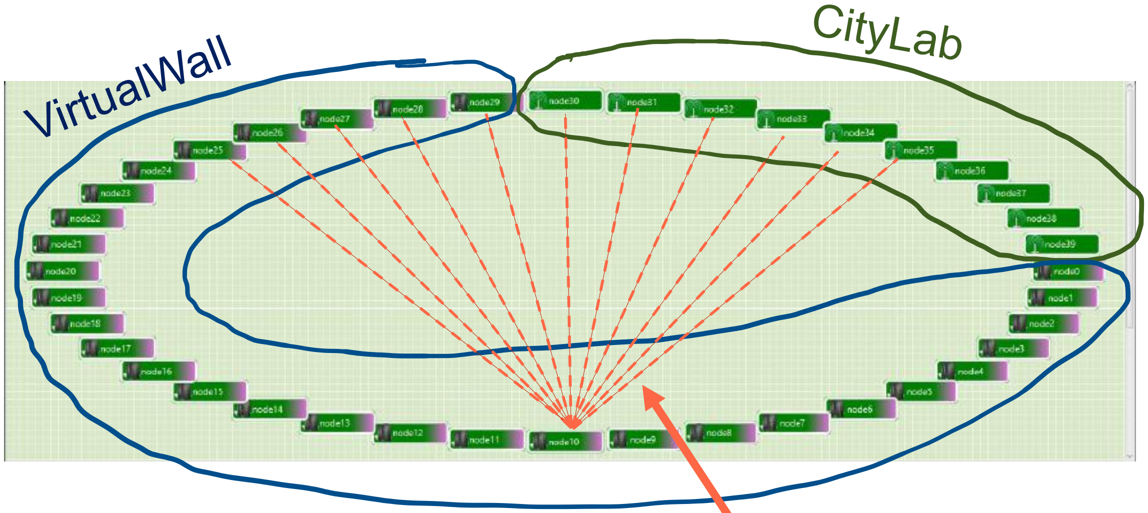

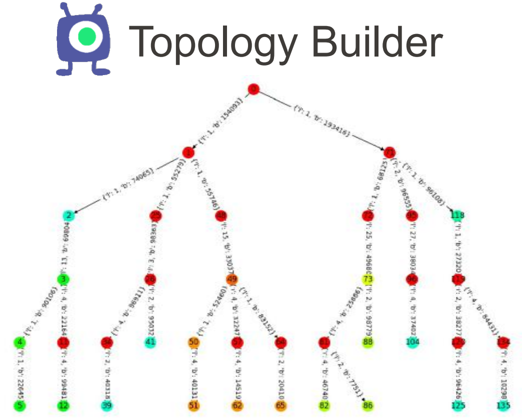

- Topology builder creates all N^2 **end-to-end connections** between all nodes across testbeds **via tc and GRE-tunnels**
- Mimicks a **hierarchical Cloud-IoT network** with Edge (at CityLab), transport (at both testbeds) and Cloud (at VIRTUALWALL) nodes, with lifelike latencies and bandwidths

Fed4Fire+ Portfolio

- All experiments requirements were fulfilled (even during an imec event, we could use other federated resources!)
- Fed4Fire+ could consider extending link impairment facilities of Virtualwall to other testbeds
- Fed4Fire+ could allow users to specify the characteristics of an arbitrary number of end-to-end links instead of groups of end-to-end links (via bridges) through JFed

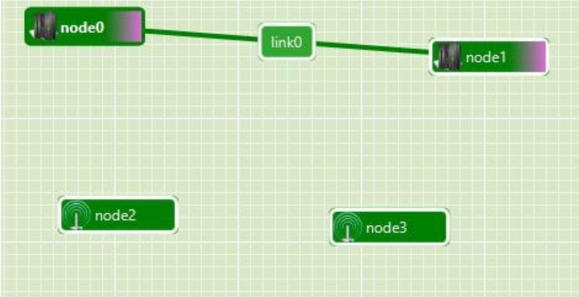

Could use our automated topology builder as a starting point!

Documentation & Support

- Documentation was very useful and complete (perhaps, some tutorials on GRE-tunneling and tc could be added)
- Very responsive, exhaustive and on-point online support from testbed owners/managers

Other feedback

FogMon 2.0 and all the opensource tooling of LiSCIo remain among **the set of services that can be used by future experimenters at Fed4Fire+**

Thanks are due to ...

Marco Gaglianese

Stefano Forti

Federica Paganelli

Francesco Buti

for all their work!

Daniel van den Akker
Johan Bergs
Bart Braem
Brecht Vermeulen

for all their support!

THANKS FOR YOUR ATTENTION

WWW.FED4FIRE.EU

Co-funded by the European Union

Co-funded by the Swiss Confederation

This project has received funding from the European Union's Horizon 2020 research and innovation programme, which is co-funded by the European Commission and the Swiss State Secretariat for Education, Research and Innovation, under grant agreement No 732638.

Parameters	Description
Report time (leader-to-follower communication period)	The time between two reports sent by a Follower to its Leader
Test time	The interval between two iterations of latency and bandwidth tests
Leader Check	The number of report time intervals before a Follower looks for a closer Leader
Latency time	The time between two latency tests on the same link
Bw Time	The time between two bandwidth tests on the same link
Silent period	The time before a non-responding Follower is considered dead by a Leader
Prop time (leader-to-leader communication period)	The number of seconds between gossiping rounds among Leaders
Sensitivity (of followers)	Threshold relative difference on average and variance for a measurement to be included in a report.
Measurement Window (size of data aggregation)	Number of measurements that are kept to compute averages.

Table 1 Configurable parameters in FogMon.

Parameters	Default	Reactive
Report time (leader-to-follower communication period)	30 s	15 s
Test time	30 s	15 s
Leader Check	8	4
Latency time	30 s	15 s
Bw Time	600 s	120 s
Silent Period	90 s	45 s
Prop time (Leader-Leader communication period)	20 s	10 s
Sensitivity (of followers)	15%	10%
Measurement Window (size of data aggregation)	20	10

Table 3 Experimental configurations of FogMon.