

# MonEx: An Integrated Experiment Monitoring Framework Standing on Off-The-Shelf Components

Abdulqawi Saif<sup>1,2</sup>

Alexandre Merlin<sup>1</sup>

Lucas Nussbaum<sup>1</sup>

Ye-Qiong Song<sup>1</sup>

<sup>1</sup>Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France

<sup>2</sup>Qwant Entreprise, Épinal, France

June 11, 2018

HPDC'2018 WORKSHOPS – *P-RECS'18*

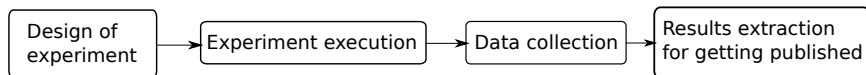
Tempe, AZ, USA





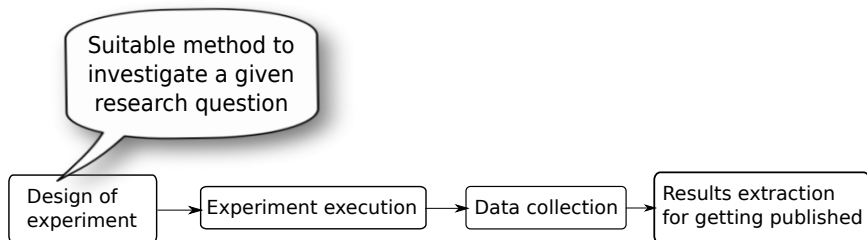
- Context
- EMF requirements
- EMF requirements Vs related work
- MonEx<sup>1</sup> EMF design
- Use case experiments of MonEx
- Conclusions

<sup>1</sup> <https://github.com/madynes/monex> (*tutorial provided*)

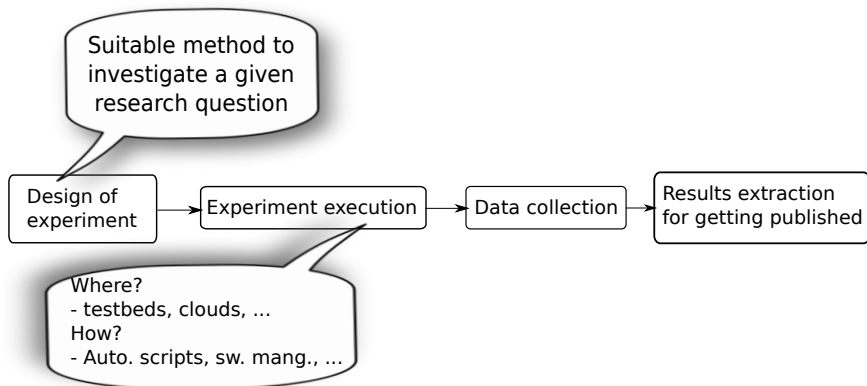


Abstracted life-cycle of a computer system experiment

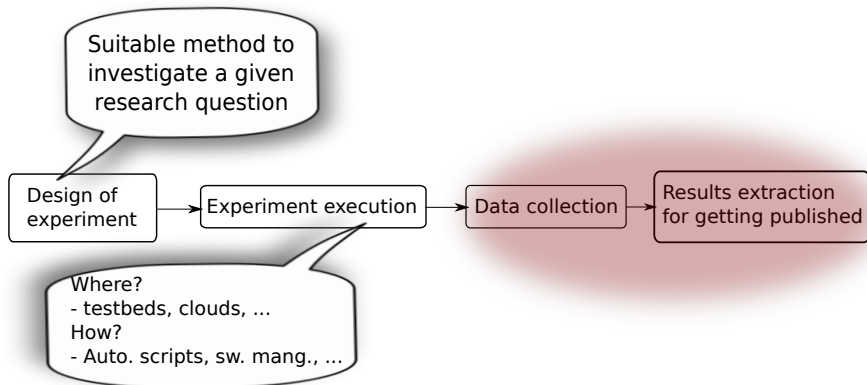
# Context



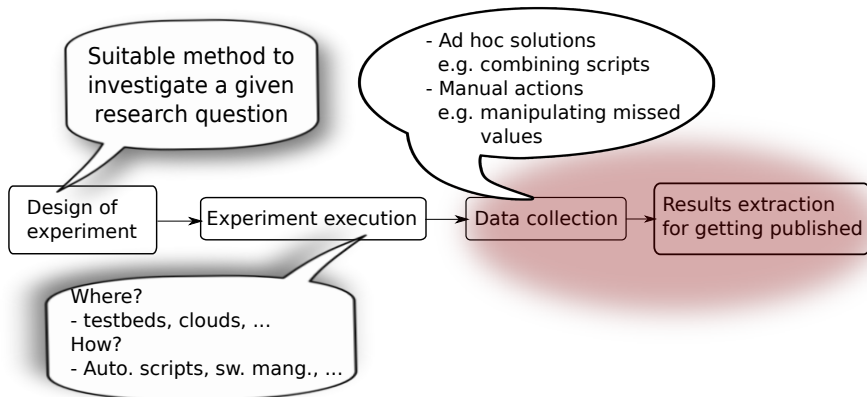
# Context



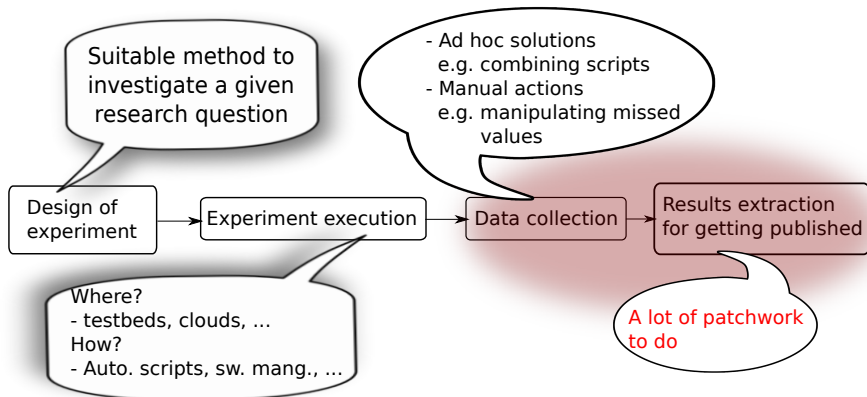
# Context

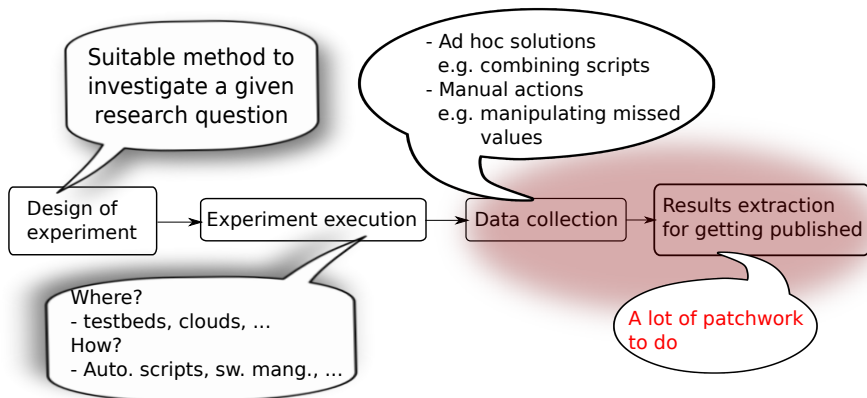


# Context



# Context





- **Objective:** enhancing data collection phase for better experiment portability and reproducibility

## Infrastructure monitoring services?

- Work from point of view of **infrastructure operators**
  - ◇ Ideal for **permanent detection of issues**
- Do not scale for **fine-grained measurements**
  - ◇ Typical for taking measurements every 5 minutes
- Have not a **definition of experiments**  
e.g. Experiments' time boundaries are not specified

## Experiment Monitoring Frameworks (EMFs)

- Monitoring **from experimenters perspective**
- Easy collection of experimental data
- Allow to **repeat experiment analysis** and comparisons
- Reducing experimenters' **effort** during experimentation  
*e.g. facilitating analysis by drawing figures*

## Infrastructure monitoring services?

- Work from point of view of **infrastructure operators**
  - ◇ Ideal for **permanent detection of issues**
- Do not scale for **fine-grained measurements**
  - ◇ Typical for taking measurements every 5 minutes
- Have not a **definition of experiments**  
e.g. Experiments' time boundaries are not specified

## Experiment Monitoring Frameworks (EMFs)

- Monitoring **from experimenters perspective**
- Easy collection of experimental data
- Allow to **repeat experiment analysis** and comparisons
- Reducing experimenters' **effort** during experimentation  
*e.g. facilitating analysis by drawing figures*

# EMF requirements

- Experiment-focused
  - ◇ name, time boundaries, metrics
- Independent of experiments
- Independent of testbed services
  - ◇ experiments' portability
- Scalable
  - ◇ in terms of metrics, nodes, measurements' frequency
- Low impact
- Easy deployment
- Controllable by users
- Real-time monitoring
  - ◇ Early detection of issues
- Publishable figures
  - ◇ Aspects of published figures
- Archival of experiments data
  - ◇ Analysis after a while

# EMF requirements

- Experiment-focused
  - ◇ name, time boundaries, metrics
- Independent of experiments
- Independent of testbed services
  - ◇ experiments' portability
- Scalable
  - ◇ in terms of metrics, nodes, measurements' frequency
- Low impact
- Easy deployment
- Controllable by users
- Real-time monitoring
  - ◇ Early detection of issues
- Publishable figures
  - ◇ Aspects of published figures
- Archival of experiments data
  - ◇ Analysis after a while

# EMF requirements vs related works

Infrastructure monitoring tools e.g. Munin	Testbed-provided measurement services	Vendetta[1]	OpenStack's <i>Gnocchi</i>	OML[2]
---	--	-------------	-------------------------------	--------

[1] Rensfelt, O. *et al.* Vendetta-a tool for flexible monitoring and management of distributed testbeds. In TridentCom 2007.

[2] Singh, M. *et al.* ORBIT Measurements framework and library (OML) : motivations, implementation and features. In TridentCom 2005.

# EMF requirements vs related works

	Infrastructure monitoring tools e.g. Munin	Testbed-provided measurement services	Vendetta[1]	OpenStack's <i>Gnocchi</i>	OML[2]
Experiment-focused	-	-	-	-	-

[1] Rensfelt, O. *et al.* Vendetta-a tool for flexible monitoring and management of distributed testbeds. In TridentCom 2007.

[2] Singh, M. *et al.* ORBIT Measurements framework and library (OML) : motivations, implementation and features. In TridentCom 2005.

# EMF requirements vs related works

**Vendetta** → Code modification is necessary

**OML** → def. of measurement points. Compilation within OML code

	Infrastructure monitoring tools e.g. Munin	Testbed-provided measurement services	Vendetta[1]	OpenStack's <i>Gnocchi</i>	OML[2]
Experiment-focused	-	-	-	-	-
Independent of experiments	+	+	-	+	-

[1] Rensfelt, O. *et al.* Vendetta-a tool for flexible monitoring and management of distributed testbeds. In TridentCom 2007.

[2] Singh, M. *et al.* ORBIT Measurements framework and library (OML) : motivations, implementation and features. In TridentCom 2005.

# EMF requirements vs related works

Vendetta → It is a management tool too

	Infrastructure monitoring tools e.g. Munin	Testbed-provided measurement services	Vendetta[1]	OpenStack's <i>Gnocchi</i>	OML[2]
Experiment-focused	-	-	-	-	-
Independent of experiments	+	+	-	+	-
Independent of testbeds services	+	+	-	+	+

[1] Rensfelt, O. *et al.* Vendetta-a tool for flexible monitoring and management of distributed testbeds. In TridentCom 2007.

[2] Singh, M. *et al.* ORBIT Measurements framework and library (OML) : motivations, implementation and features. In TridentCom 2005.

# EMF requirements vs related works

**Infra. mon.** → Typical for 5-10 m measurements

**Vendetta** → No info. about frequency of measurements

	Infrastructure monitoring tools e.g. Munin	Testbed-provided measurement services	Vendetta[1]	OpenStack's <i>Gnocchi</i>	OML[2]
Experiment-focused	-	-	-	-	-
Independent of experiments	+	+	-	+	-
Independent of testbeds services	+	+	-	+	+
Scalability	+ -	+	+ -	+	+

[1] Rensfelt, O. *et al.* Vendetta-a tool for flexible monitoring and management of distributed testbeds. In TridentCom 2007.

[2] Singh, M. *et al.* ORBIT Measurements framework and library (OML) : motivations, implementation and features. In TridentCom 2005.

# EMF requirements vs related works

**Infra. mon.** → Some with heavy nature (e.g. threads)

**Vendetta** → Parsing events in-place

	Infrastructure monitoring tools e.g. Munin	Testbed-provided measurement services	Vendetta[1]	OpenStack's <i>Gnocchi</i>	OML[2]
Experiment-focused	-	-	-	-	-
Independent of experiments	+	+	-	+	-
Independent of testbeds services	+	+	-	+	+
Scalability	+ -	+	+ -	+	+
Low impact	-	+	-	+	+

[1] Rensfelt, O. *et al.* Vendetta-a tool for flexible monitoring and management of distributed testbeds. In TridentCom 2007.

[2] Singh, M. *et al.* ORBIT Measurements framework and library (OML) : motivations, implementation and features. In TridentCom 2005.

# EMF requirements vs related works

Testbed-provided services → Designed for specific infrastructure

	Infrastructure monitoring tools e.g. Munin	Testbed-provided measurement services	Vendetta[1]	OpenStack's <i>Gnocchi</i>	OML[2]
Experiment-focused	-	-	-	-	-
Independent of experiments	+	+	-	+	-
Independent of testbeds services	+	+	-	+	+
Scalability	+ -	+	+ -	+	+
Low impact	-	+	-	+	+
Easy deployment	+	-	+	+	+

[1] Rensfelt, O. *et al.* Vendetta-a tool for flexible monitoring and management of distributed testbeds. In TridentCom 2007.

[2] Singh, M. *et al.* ORBIT Measurements framework and library (OML) : motivations, implementation and features. In TridentCom 2005.

# EMF requirements vs related works

Infra. mon. & Testbed-provided services → Only by infrastructure operators

	Infrastructure monitoring tools e.g. Munin	Testbed-provided measurement services	Vendetta[1]	OpenStack's <i>Gnocchi</i>	OML[2]
Experiment-focused	-	-	-	-	-
Independent of experiments	+	+	-	+	-
Independent of testbeds services	+	+	-	+	+
Scalability	+ -	+	+ -	+	+
Low impact	-	+	-	+	+
Easy deployment	+	-	+	+	+
Controllable	-	-	+	+	+

[1] Rensfelt, O. et al. Vendetta-a tool for flexible monitoring and management of distributed testbeds. In TridentCom 2007.

[2] Singh, M. et al. ORBIT Measurements framework and library (OML) : motivations, implementation and features. In TridentCom 2005.

# EMF requirements vs related works

	Infrastructure monitoring tools e.g. Munin	Testbed-provided measurement services	Vendetta[1]	OpenStack's <i>Gnocchi</i>	OML[2]
Experiment-focused	-	-	-	-	-
Independent of experiments	+	+	-	+	-
Independent of testbeds services	+	+	-	+	+
Scalability	+ -	+	+ -	+	+
Low impact	-	+	-	+	+
Easy deployment	+	-	+	+	+
Controllable	-	-	+	+	+
Real-time monitoring	+	+	-	-	-

[1] Rensfelt, O. *et al.* Vendetta-a tool for flexible monitoring and management of distributed testbeds. In TridentCom 2007.

[2] Singh, M. *et al.* ORBIT Measurements framework and library (OML) : motivations, implementation and features. In TridentCom 2005.

# EMF requirements vs related works

Infra. mon. & Testbed-provided services → RRDtool

	Infrastructure monitoring tools e.g. Munin	Testbed-provided measurement services	Vendetta[1]	OpenStack's <i>Gnocchi</i>	OML[2]
Experiment-focused	-	-	-	-	-
Independent of experiments	+	+	-	+	-
Independent of testbeds services	+	+	-	+	+
Scalability	+ -	+	+ -	+	+
Low impact	-	+	-	+	+
Easy deployment	+	-	+	+	+
Controllable	-	-	+	+	+
Real-time monitoring	+	+	-	-	-
Producing publication-quality figures	-	-	-	-	-

[1] Rensfelt, O. *et al.* Vendetta-a tool for flexible monitoring and management of distributed testbeds. In TridentCom 2007.

[2] Singh, M. *et al.* ORBIT Measurements framework and library (OML) : motivations, implementation and features. In TridentCom 2005.

# EMF requirements vs related works

	Infrastructure monitoring tools e.g. Munin	Testbed-provided measurement services	Vendetta[1]	OpenStack's <i>Gnocchi</i>	OML[2]
Experiment-focused	-	-	-	-	-
Independent of experiments	+	+	-	+	-
Independent of testbeds services	+	+	-	+	+
Scalability	+ -	+	+ -	+	+
Low impact	-	+	-	+	+
Easy deployment	+	-	+	+	+
Controllable	-	-	+	+	+
Real-time monitoring	+	+	-	-	-
Producing publication-quality figures	-	-	-	-	-
Archival of data	+	+	-	+	+

[1] Rensfelt, O. *et al.* Vendetta-a tool for flexible monitoring and management of distributed testbeds. In TridentCom 2007.

[2] Singh, M. *et al.* ORBIT Measurements framework and library (OML) : motivations, implementation and features. In TridentCom 2005.

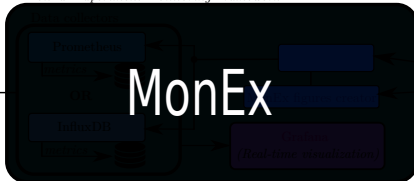
# MonEx EMF (Monitoring Experiments)

Experiments running on a testbed




Collects data  
Push or Pull

MonEx: Experiment Monitoring Framework



Start or  
end Exp\_A  
Get Exp\_A  
Draw Exp\_B  
An  
experimenter



---

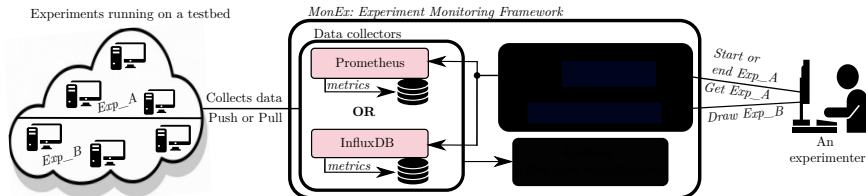
## MonEx

---

Experiment-focused	+
Independent of experiments	+
Independent of testbeds services	+
Scalability	+
Low impact	+
Easy deployment	+
Controllable	+
Real-time monitoring	+
Producing publication-quality figures	+
Archival of data	+

---

# MonEx EMF (Monitoring Experiments)



MonEx		
	Experiment-focused	+
	Independent of experiments	+
	Independent of testbeds services	+
✓	Scalability	+
✓	Low impact	+
	Easy deployment	+
	Controllable	+
	Real-time monitoring	+
	Producing publication-quality figures	+
✓	Archival of data	+

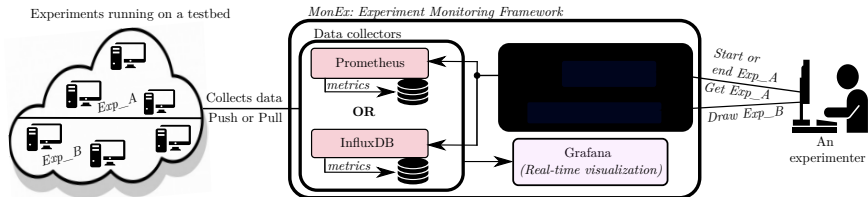
## ■ Prometheus

- 1 sec polling interval ☺
- A lot of **data exporters** ☺
- Fixed-interval polling ☹

## ■ InfluxDB

- Suitable for **pushing data** ☺
- Dynamic pushing ☺

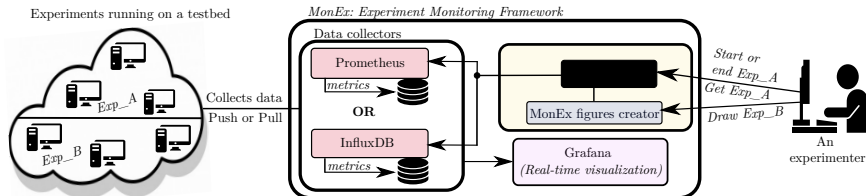
# MonEx EMF (Monitoring Experiments)



## MonEx

	Experiment-focused	+
	Independent of experiments	+
	Independent of testbeds services	+
✓	Scalability	+
✓	Low impact	+
	Easy deployment	+
	Controllable	+
✓	Real-time monitoring	+
	Producing publication-quality figures	+
✓	Archival of data	+

# MonEx EMF (Monitoring Experiments)

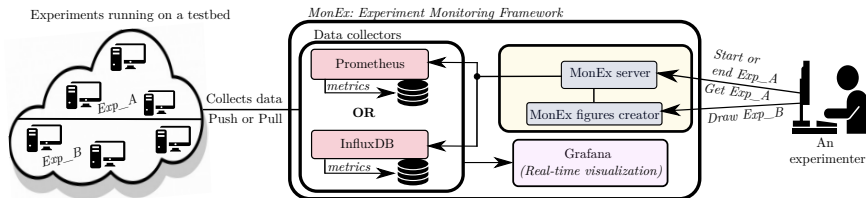


## MonEx

	Experiment-focused	+
	Independent of experiments	+
	Independent of testbeds services	+
✓	Scalability	+
✓	Low impact	+
	Easy deployment	+
	Controllable	+
✓	Real-time monitoring	+
✓	Producing publication-quality figures	+
✓	Archival of data	+

- Create figures that **respect publication aspects**  
e.g. X-Y figures, Y-figures, multiple Y-figures, stack figures, ...

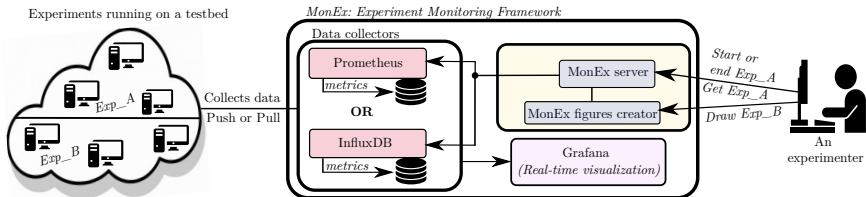
# MonEx EMF (Monitoring Experiments)



MonEx		
✓	Experiment-focused	+
	Independent of experiments	+
	Independent of testbeds services	+
✓	Scalability	+
✓	Low impact	+
	Easy deployment	+
✓	Controllable	+
✓	Real-time monitoring	+
✓	Producing publication-quality figures	+
✓	Archival of data	+

- The **main interface** for experimenters to interact
- Allow **interacting with experiments** metrics
- Specify experiment **time boundaries**
- Handle experiments **by name**

# MonEx EMF (Monitoring Experiments)



## MonEx

✓	Experiment-focused	+
✓	Independent of experiments	+
✓	Independent of testbeds services	+
✓	Scalability	+
✓	Low impact	+
✓	Easy deployment	+
✓	Controllable	+
✓	Real-time monitoring	+
✓	Producing publication-quality figures	+
✓	Archival of data	+

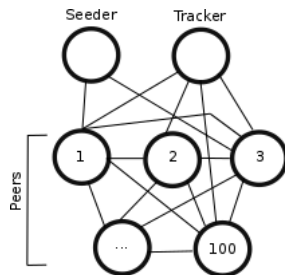
# Use case experiments

## Three different experiments

- Highlighting MonEx ability to cover EMF requirements
- Performed on the *Grid'5000* testbed

## First Experiment: many-nodes Bittorrent Download

- One seeder & 100 peers
- Transmission as torrent client
- *Distem* for emulating slow peers (30KB/s)
- Dedicated VLAN for monitoring traffic
- Prometheus as a MonEx data collector
- We use our *own exporter*<sup>1</sup> ( 15 LOC)



[1] [https://github.com/madynes/monex/blob/master/Artifacts\\_and\\_datasets/many-nodes\\_bittorrent/exporter.py](https://github.com/madynes/monex/blob/master/Artifacts_and_datasets/many-nodes_bittorrent/exporter.py)

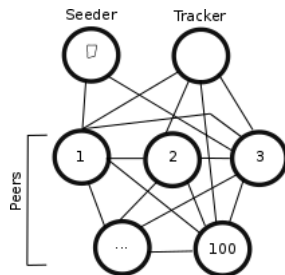
# Use case experiments

## Three different experiments

- Highlighting MonEx ability to cover EMF requirements
- Performed on the *Grid'5000* testbed

## First Experiment: many-nodes Bittorrent Download

- One seeder & 100 peers
- Transmission as torrent client
- *Distem* for emulating slow peers (30KB/s)
- Dedicated VLAN for monitoring traffic
- Prometheus as a MonEx data collector
- We use our *own exporter*<sup>1</sup> ( 15 LOC)



[1] [https://github.com/madynes/monex/blob/master/Artifacts\\_and\\_datasets/many-nodes\\_bittorrent/exporter.py](https://github.com/madynes/monex/blob/master/Artifacts_and_datasets/many-nodes_bittorrent/exporter.py)

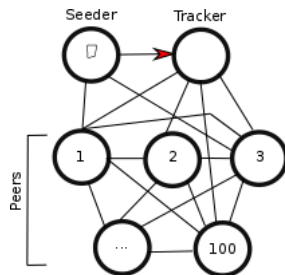
# Use case experiments

## Three different experiments

- Highlighting MonEx ability to cover EMF requirements
- Performed on the *Grid'5000* testbed

## First Experiment: many-nodes Bittorrent Download

- One seeder & 100 peers
- Transmission as torrent client
- *Distem* for emulating slow peers (30KB/s)
- Dedicated VLAN for monitoring traffic
- Prometheus as a MonEx data collector
- We use our *own exporter*<sup>1</sup> ( 15 LOC)



[1] [https://github.com/madynes/monex/blob/master/Artifacts\\_and\\_datasets/many-nodes\\_bittorrent/exporter.py](https://github.com/madynes/monex/blob/master/Artifacts_and_datasets/many-nodes_bittorrent/exporter.py)

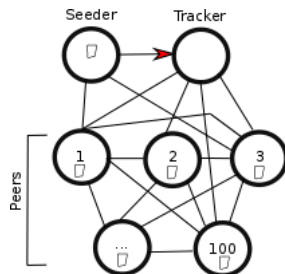
# Use case experiments

## Three different experiments

- Highlighting MonEx ability to cover EMF requirements
- Performed on the *Grid'5000* testbed

## First Experiment: many-nodes Bittorrent Download

- One seeder & 100 peers
- Transmission as torrent client
- *Distem* for emulating slow peers (30KB/s)
- Dedicated VLAN for monitoring traffic
- Prometheus as a MonEx data collector
- We use our *own exporter*<sup>1</sup> ( 15 LOC)

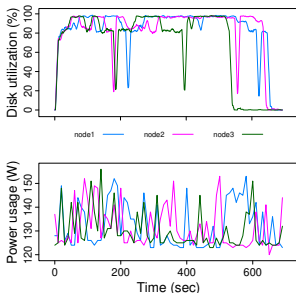


[1] [https://github.com/madynes/monex/blob/master/Artifacts\\_and\\_datasets/many-nodes\\_bittorrent/exporter.py](https://github.com/madynes/monex/blob/master/Artifacts_and_datasets/many-nodes_bittorrent/exporter.py)

# Use case experiments

## Second Experiment: metrics diversity

- Metrics: HDD Util & power of a *MongoDB* cluster
- Three-shard nodes & 80 GB of total data
- Prometheus as a data collector & Grafana for visualization
- Two default node-exporters of Prometheus
  - ⇒ node-exporter → nodes spec.
  - ⇒ SNMP-exporter → PDUs data



## Third Experiment: time-dependent & heavy-measurement

- Metric : file offsets vs sequences of I/O requests
- High-frequency measurement (thousands of I/O)
  - pushing data to InfluxDB
- Generate I/O traces using IOscope [1]
- Fio benchmark generates rand. read workload

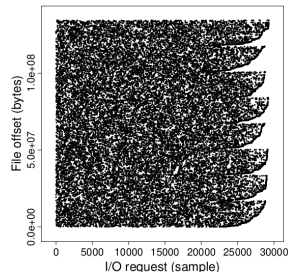
## Third Experiment: results

### ■ Get data

```
curl -X GET 127.1:5000/exp/accessPattern -H  
"Content-Type: application/json" -d  
'"measurement":"ebpf","server":"influx",  
"database":"monex","type":"sample"' > d.csv
```

### ■ Draw-it!

```
./monex-draw -F d.csv -x "I/O request  
(sample)" -y "Offset (bytes)" -p -r "random"
```



[1] A. Saif et al. IOscope : A Flexible I/O Tracer for Workloads' I/O Pattern Characterization, in WOPSSS, 2018

**Demo time!**

## Conclusions:

- We defined the **minimal requirements of EMFs**
- **Partial coverage** of requirements in the literature
- We proposed **MonEx EMF** to enhance monitoring experiments
  - ◇ Usable on various experimentations' contexts
  - ◇ Efficient for analysis repetition and metrics comparisons
  - ◇ First step to **unifying** experiments' data collection methods
  - ◇ Towards **experiments portability** and reproducibility

## Future prospects:

- MonEx maturity?
- Monitoring experiments on federated testbeds

# Any Questions ?

## Conclusions:

- We defined the **minimal requirements of EMFs**
- **Partial coverage** of requirements in the literature
- We proposed **MonEx EMF** to enhance monitoring experiments
  - ◇ Usable on various experimentations' contexts
  - ◇ Efficient for analysis repetition and metrics comparisons
  - ◇ First step to **unifying** experiments' data collection methods
  - ◇ Towards **experiments portability** and reproducibility

## Future prospects:

- MonEx maturity?
- Monitoring experiments on federated testbeds

# Any Questions ?