Monitoring Testbed Experiments with *MonEx*

Abdulqawi Saif[1,2] Alexandre Merlin[1] Lucas Nussbaum[1] Ye-Qiong Song[1]

[1]Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France
[2]Qwant Entreprise, Épinal, France

April 04, 2018
1st *Grid'5000-FIT* School - Sophia Antipolis, France

Context

Computer system experiments

- Diversity of environments
- Variety of contexts
- Many kinds of outputs

Context

Computer system experiments

- Diversity of environments
- Variety of contexts
- Many kinds of outputs

Experiments' data is often collected via:

- Ad-hoc solutions (*e.g. combining multiple scripts*)
- Manual actions (*e.g. manipulating some missed values*)
- Portability and reproducibility aware-less methods

Context

Computer system experiments

- Diversity of environments
- Variety of contexts
- Many kinds of outputs

Experiments' data is often collected via:

- Ad-hoc solutions (*e.g. combining multiple scripts*)
- Manual actions (*e.g. manipulating some missed values*)
- Portability and reproducibility aware-less methods

How to overcome these challenges?

Context

Our aim: having an Experiment Monitoring Framework

- Monitoring?, Yes!, but from experimenters perspective
- Easily collection of experiments' data
- Reducing experimenters effort during experimentation
  e.g. facilitating analysis by drawing figures
- Permit to repeat experiment analysis and comparisons

Open questions:

- What should be targeted by EMFs?
- What needed to build an EMF?
- Do we need to build EMFs from scratch?

Context

Our aim: having an Experiment Monitoring Framework

- Monitoring?, Yes!, but from experimenters perspective
- Easily collection of experiments' data
- Reducing experimenters effort during experimentation
  e.g. facilitating analysis by drawing figures
- Permit to repeat experiment analysis and comparisons

Open questions:

- What should be targeted by EMFs?
- What needed to build an EMF?
- Do we need to build EMFs from scratch?

- EMF requirements
- EMF requirements Vs related work
- MonEx EMF design
- Use case experiments of MonEx
- Grid'5000 and MonEx
- Conclusions

EMF requirements

Motivation experiment

- Target: torrent completion
- 100 peers & one seeder
- Dynamic growth of network
  $\rightarrow$ adding a peer after *n sec*
- Emulated bandwidth of peers

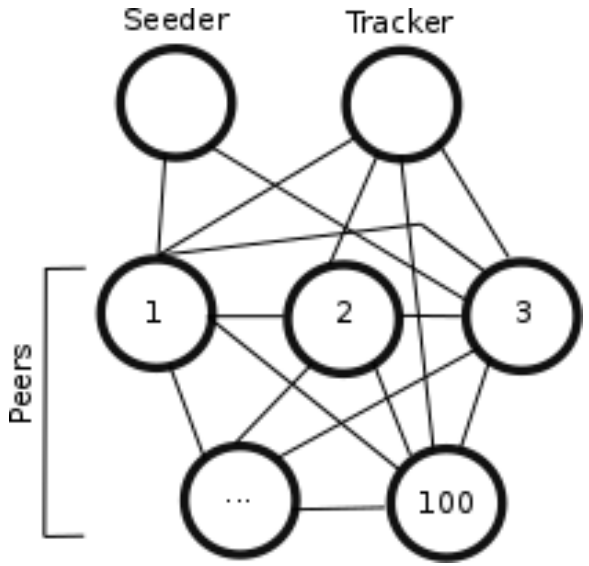

EMF requirements

Motivation experiment

- Target: torrent completion
- 100 peers & one seeder
- Dynamic growth of network
  $\rightarrow$ adding a peer after *n sec*
- Emulated bandwidth of peers

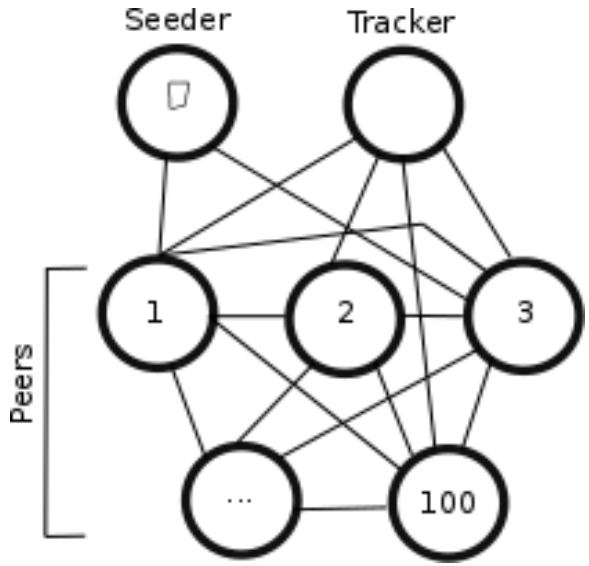

EMF requirements

Motivation experiment

- Target: torrent completion
- 100 peers & one seeder
- Dynamic growth of network
  $\rightarrow$ adding a peer after *n sec*
- Emulated bandwidth of peers

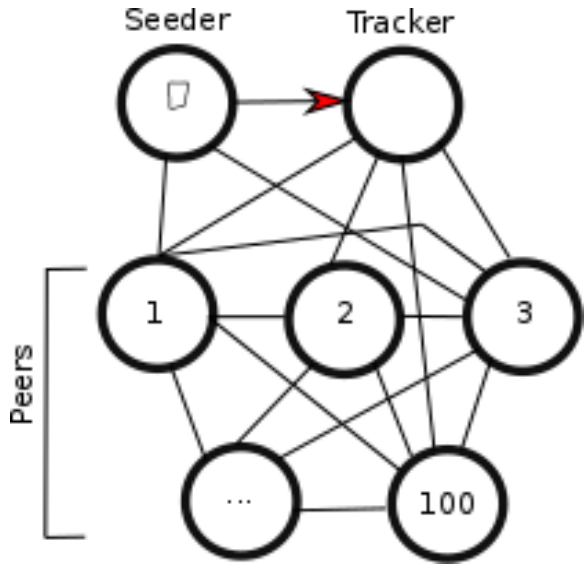

EMF requirements

Motivation experiment

- Target: torrent completion
- 100 peers & one seeder
- Dynamic growth of network
  $\rightarrow$ adding a peer after *n sec*
- Emulated bandwidth of peers

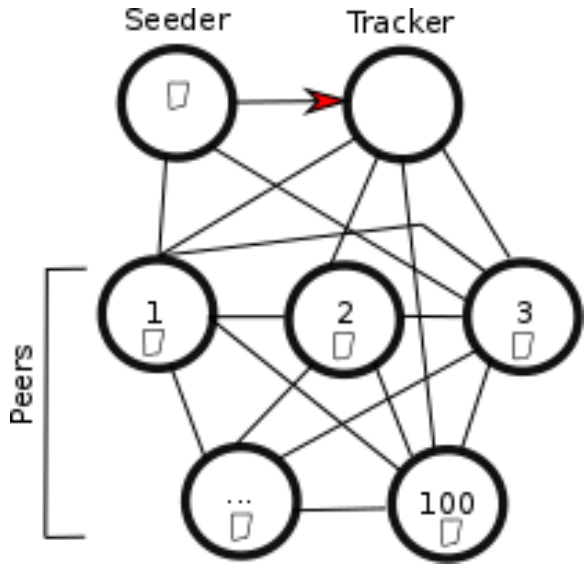

EMF requirements

Motivation experiment

- Target: torrent completion
- 100 peers & one seeder
- Dynamic growth of network
  $\rightarrow$ adding a peer after *n sec*
- Emulated bandwidth of peers

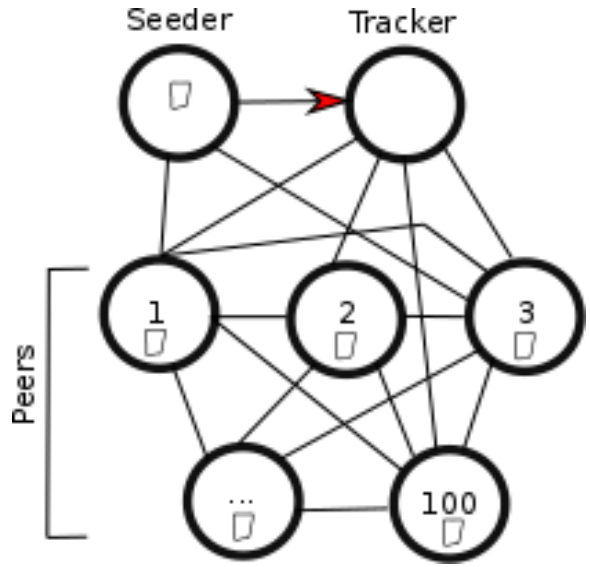

Experiment Monitoring Framework (EMF) requirements

- Experiment-focused
- Independent of experiments
- Independent of testbed services and man. frameworks
- Scalable
- Low impact
- Controllable by users
- Real-time monitoring
- Publishable figures
- Archival of experiments data

EMF requirements vs related works

	Infrastructure monitoring tools e.g. Munin	Testbed-provided measurement services	Vendetta[1]	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]	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 $\rightarrow$ code modification is necessary
OML$\rightarrow$ def. of measurement points, compilation with OML code,

	Infrastructure monitoring tools e.g. Munin	Testbed-provided measurement services	Vendetta[1]	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 $\rightarrow$ also a management tool

	Infrastructure monitoring tools e.g. Munin	Testbed-provided measurement services	Vendetta[1]	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

Inf. mon. → typical for 5-10 m measurements
Vendetta→ no info. about the measurement frequency

	Infrastructure monitoring tools e.g. Munin	Testbed-provided measurement services	Vendetta[1]	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

Inf. mon. → some with heavy nature (threads, ...)
Vendetta→ parsing events in-place

	Infrastructure monitoring tools e.g. Munin	Testbed-provided measurement services	Vendetta[1]	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 → specific infrastructure

	Infrastructure monitoring tools e.g. Munin	Testbed-provided measurement services	Vendetta[1]	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

Inf. mon. & Testbed-provided services → only by infrastructure operators

	Infrastructure monitoring tools e.g. Munin	Testbed-provided measurement services	Vendetta[1]	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]	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

Inf. mon. & Testbed-provided services → RRDtool

	Infrastructure monitoring tools e.g. Munin	Testbed-provided measurement services	Vendetta[1]	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]	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 (*Mon*itoring *Ex*periments)

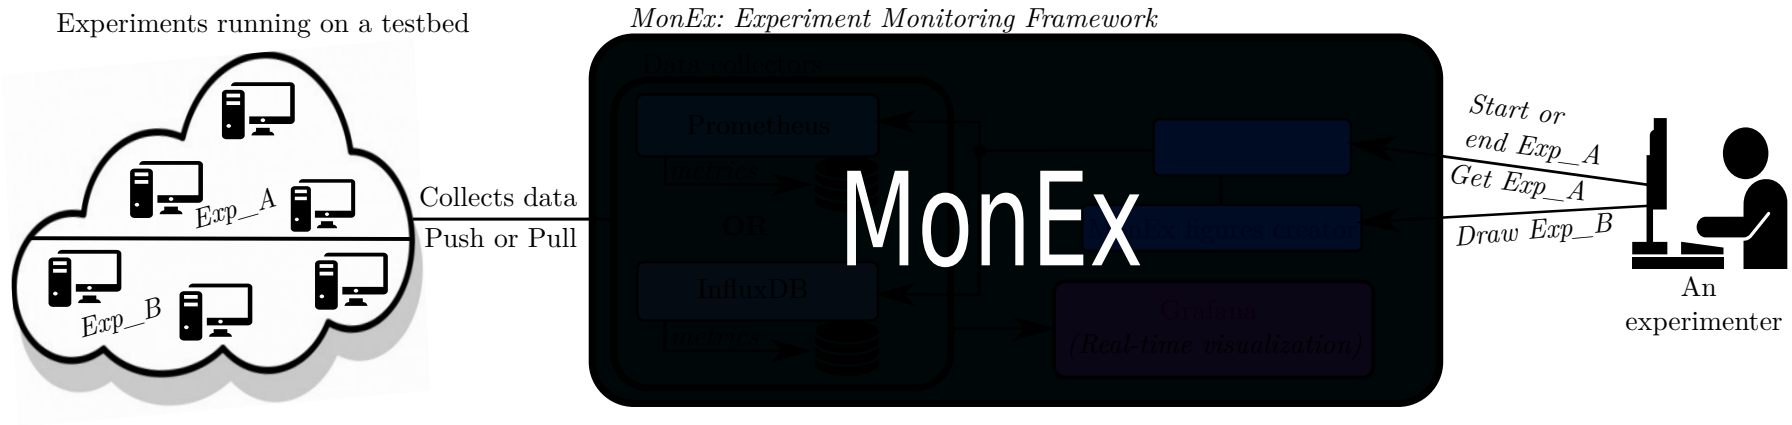

	MonEx
Experiment-focused	+
Independent of experiments	+
Independent of testbeds services	+
Scalability	+
Low impact	+
Easy deployment	+
Controllable	+
Real-time monitoring	+
Producing publication-quality figures	+
Archival of data	+

MonEx EMF (Monitoring Experiments)

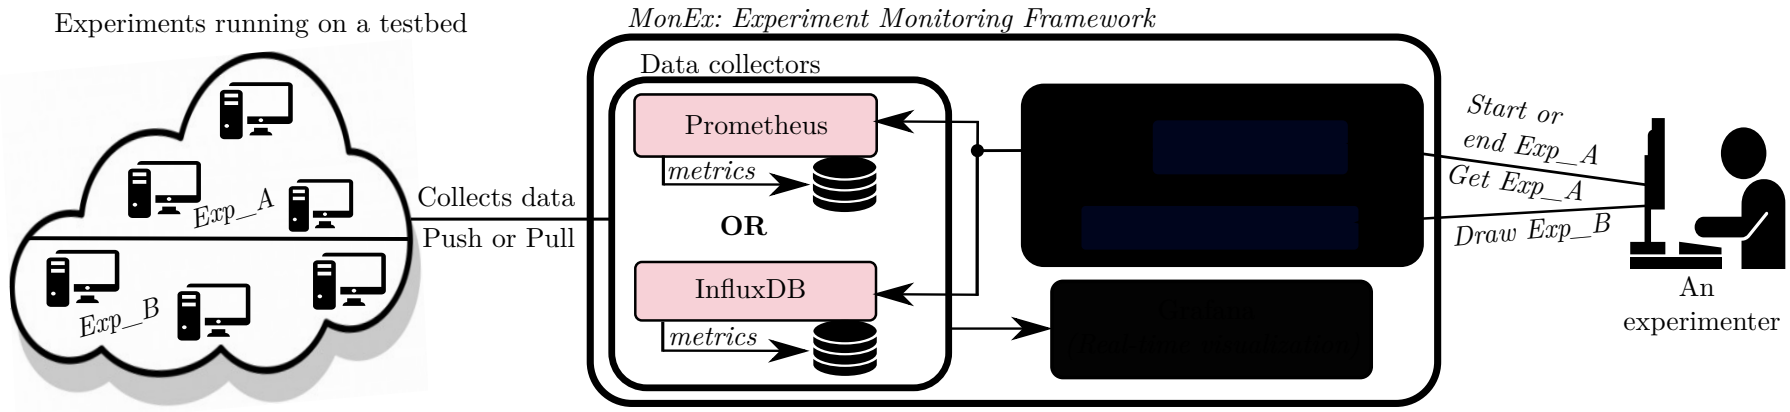

	MonEx
Experiment-focused	+
Independent of experiments	+
Independent of testbeds services	+
✓ Scalability	+
✓ Low impact	+
Easy deployment	+
Controllable	+
Real-time monitoring	+
Producing publication-quality figures	+
✓ Archival of data	+

- Prometheus
  - Up-to 1 sec as a polling interval ☺
  - A lot of data exporters
    e.g. node spec, specific app.
  - Fixed-interval polling ☹
- InfluxDB
  - Suitable for pushing data ☺
  - Variated interval measurements ☺

MonEx EMF (Monitoring Experiments)

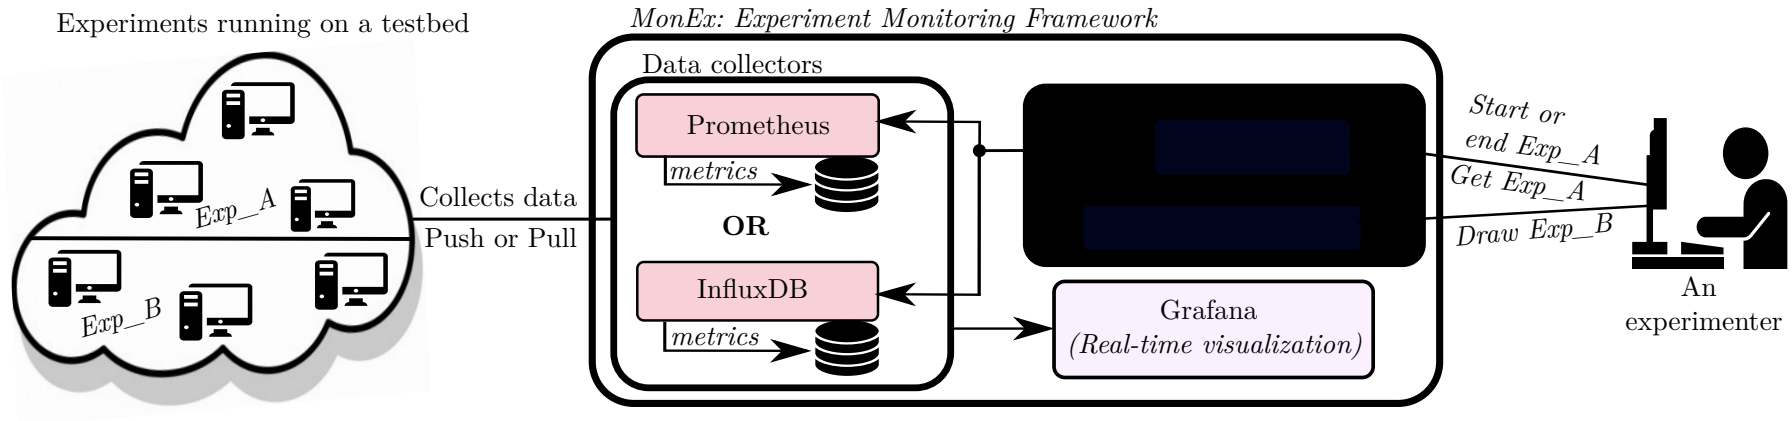

	MonEx
Experiment-focused	+
Independent of experiments	+
Independent of testbeds services	+
✓ **Scalability**	+
✓ **Low impact**	+
Easy deployment	+
Controllable	+
✓ **Real-time monitoring**	+
Producing publication-quality figures	+
✓ **Archival of data**	+

MonEx EMF (Monitoring Experiments)

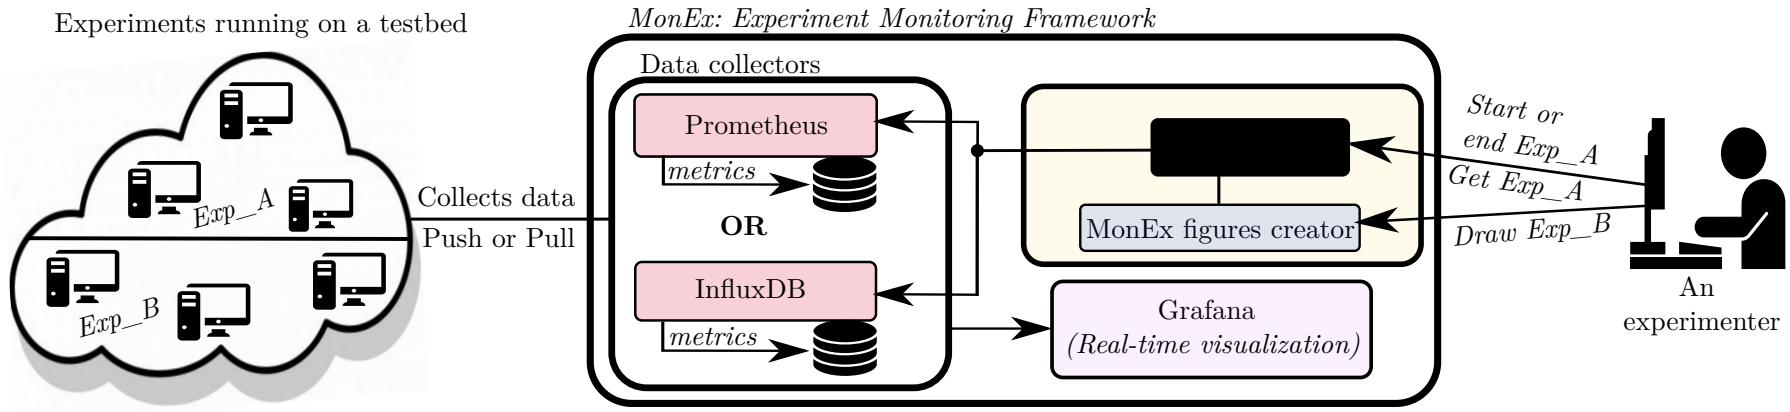

	MonEx
Experiment-focused	+
Independent of experiments	+
Independent of testbeds services	+
✓ Scalability	+
✓ Low impact	+
Easy deployment	+
Controllable	+
✓ Real-time monitoring	+
✓ Producing publication-quality figures	+
✓ Archival of data	+

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

MonEx EMF (Monitoring Experiments)

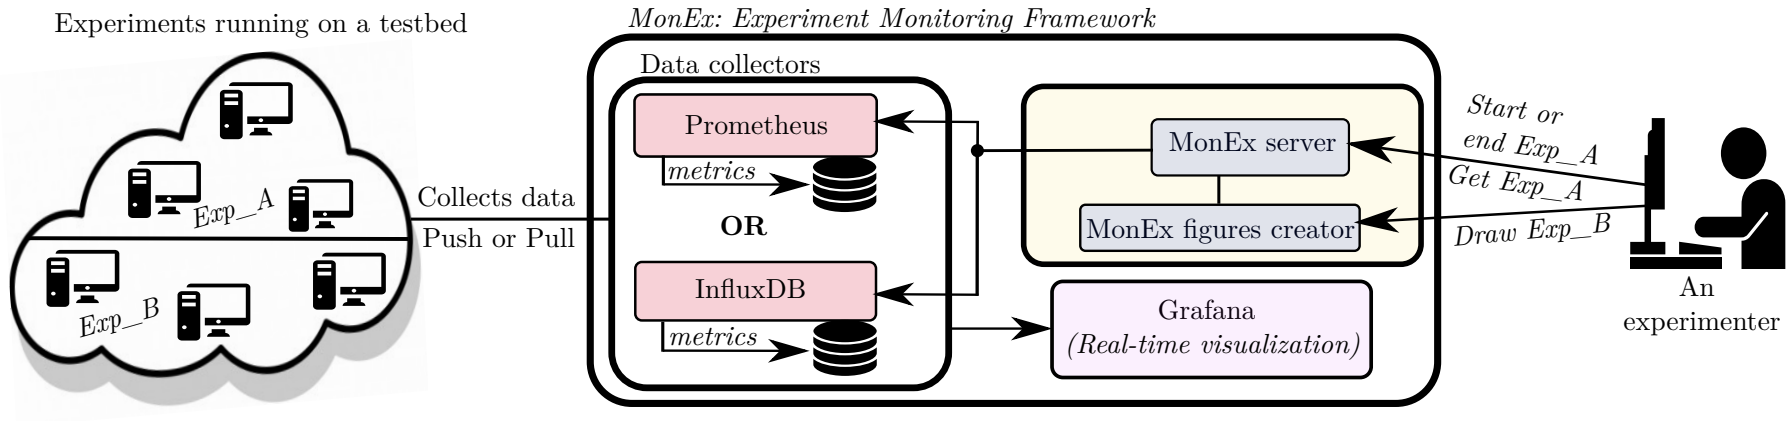

	MonEx
✓ Experiment-focused	┼
Independent of experiments	┼
Independent of testbeds services	┼
✓ Scalability	┼
✓ Low impact	┼
Easy deployment	┼
✓ Controllable	┼
✓ Real-time monitoring	┼
✓ Producing publication-quality figures	┼
✓ Archival of data	┼

- Handle experiments by name
- Specify experiment time boundaries
- Allow interacting with experiments metrics

MonEx EMF (Monitoring Experiments)

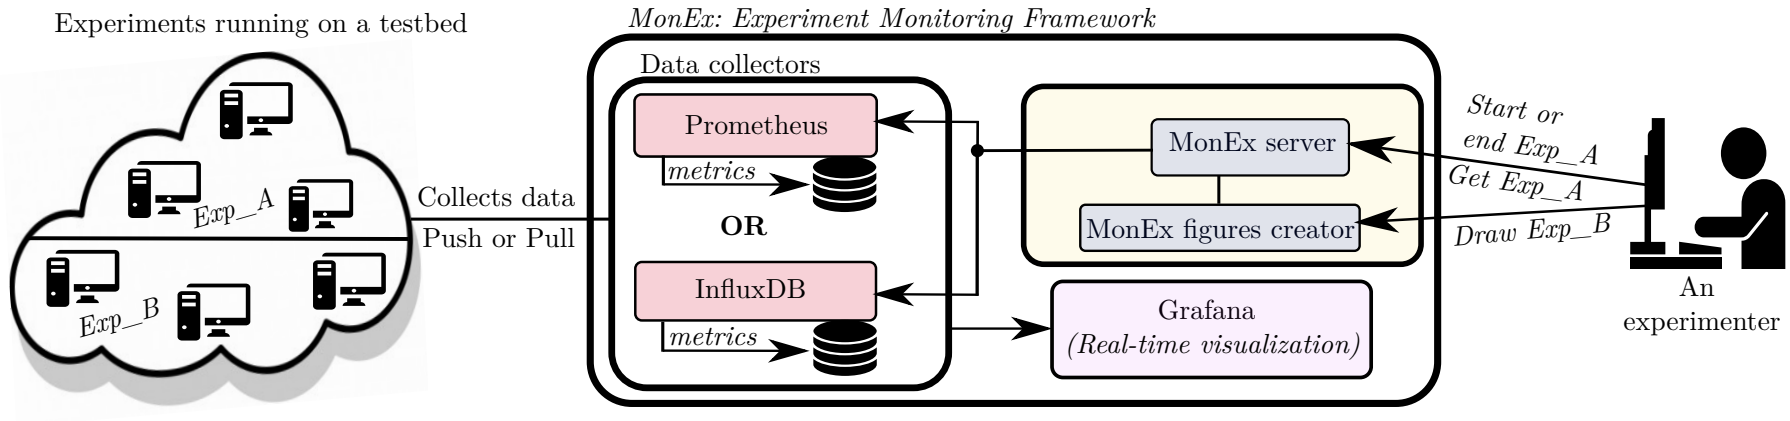

	MonEx
✓ Experiment-focused	+
✓ Independent of experiments	+
✓ Independent of testbeds services	+
✓ Scalability	+
✓ Low impact	+
✓ Easy deployment	+
✓ Controllable	+
✓ Real-time monitoring	+
✓ Producing publication-quality figures	+
✓ Archival of data	+

MonEx EMF (Monitoring Experiments)

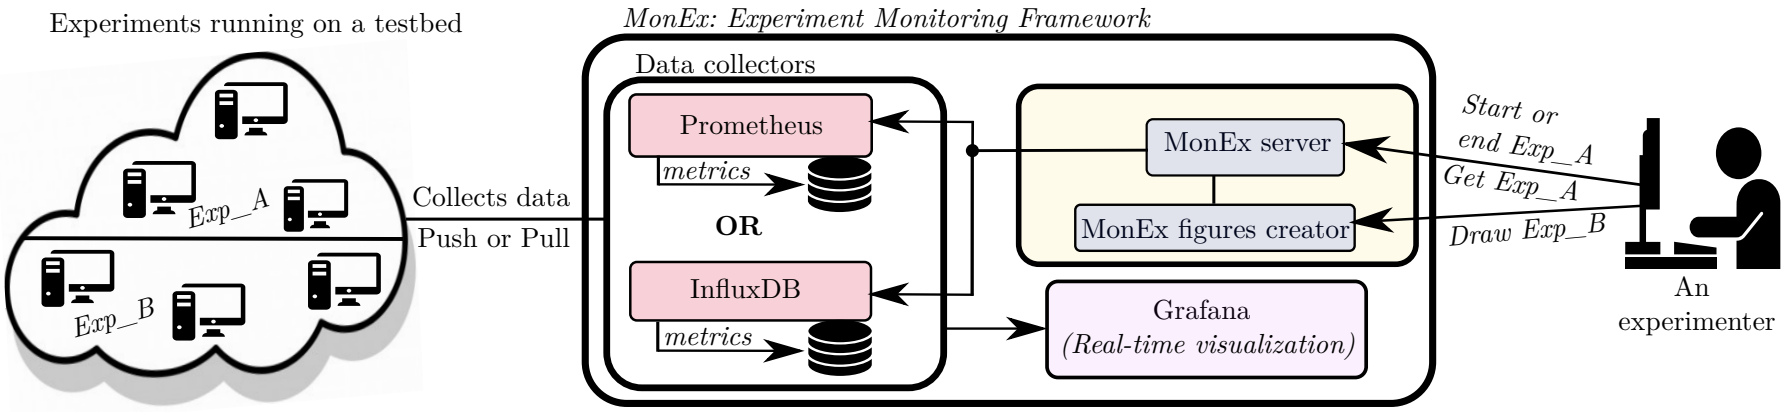

	MonEx
✓ Experiment-focused	┼
✓ Independent of experiments	┼
✓ Independent of testbeds services	┼
✓ Scalability	┼
✓ Low impact	┼
✓ Easy deployment	┼
✓ Controllable	┼
✓ Real-time monitoring	┼
✓ Producing publication-quality figures	┼
✓ Archival of data	┼

Experimenter responsibilities:

1. Make a usual experiment deployment
2. Communicate data with MonEx
   → *Expose it for Prom. or push it into InfluxDB*
   → Use or write your own exporter(s)

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*[1] (tens LOC)

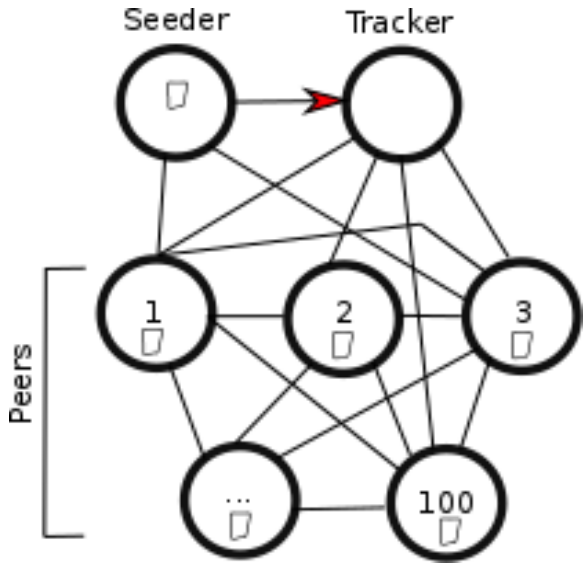

[1] https://github.com/madynes/monex/blob/master/Artifacts_and_datasets/many-nodes_bittorrent/exporter.py

Use case experiments

First Experiment: workflow

Bittorrent experiment

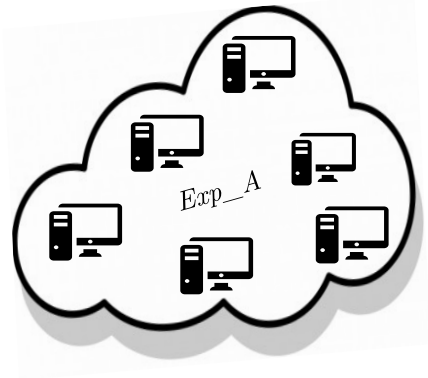

MonEx: Experiment Monitoring Framework

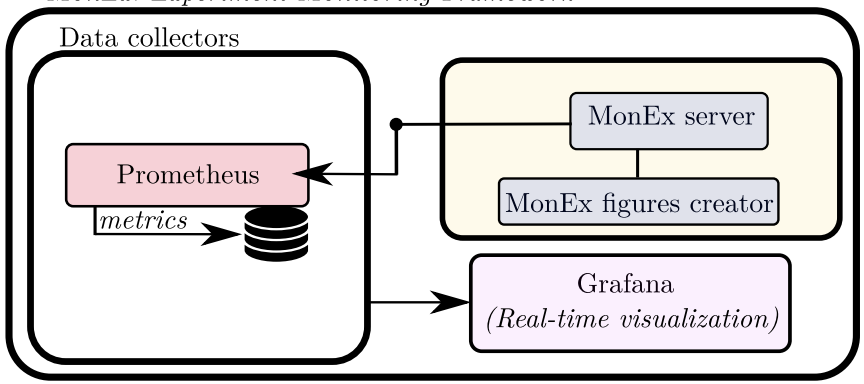

An experimenter

Use case experiments

First Experiment: workflow

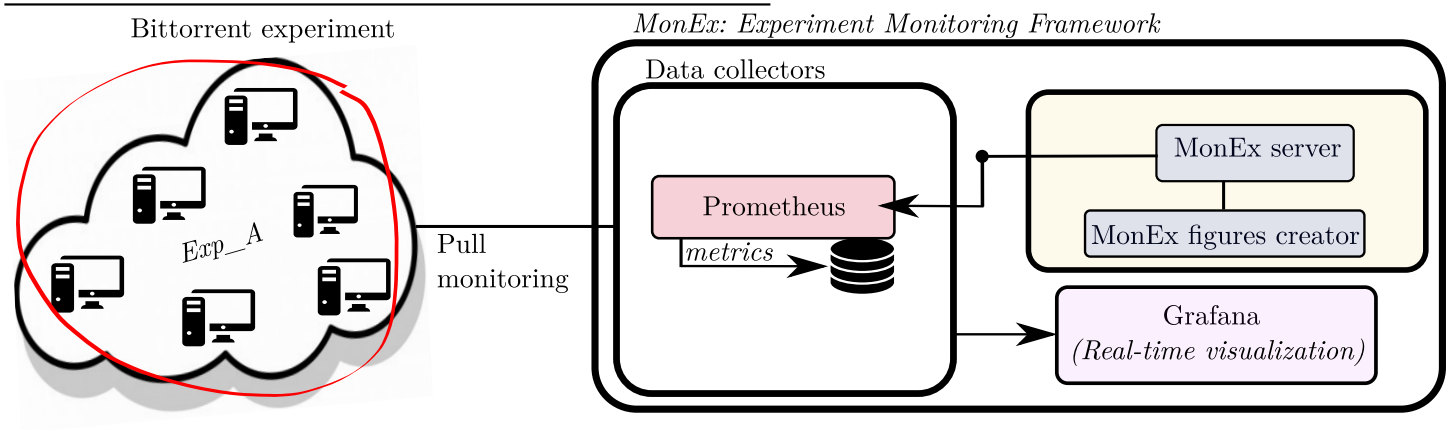

Make a usual deployment and install *MonEx*

Use case experiments

First Experiment: workflow

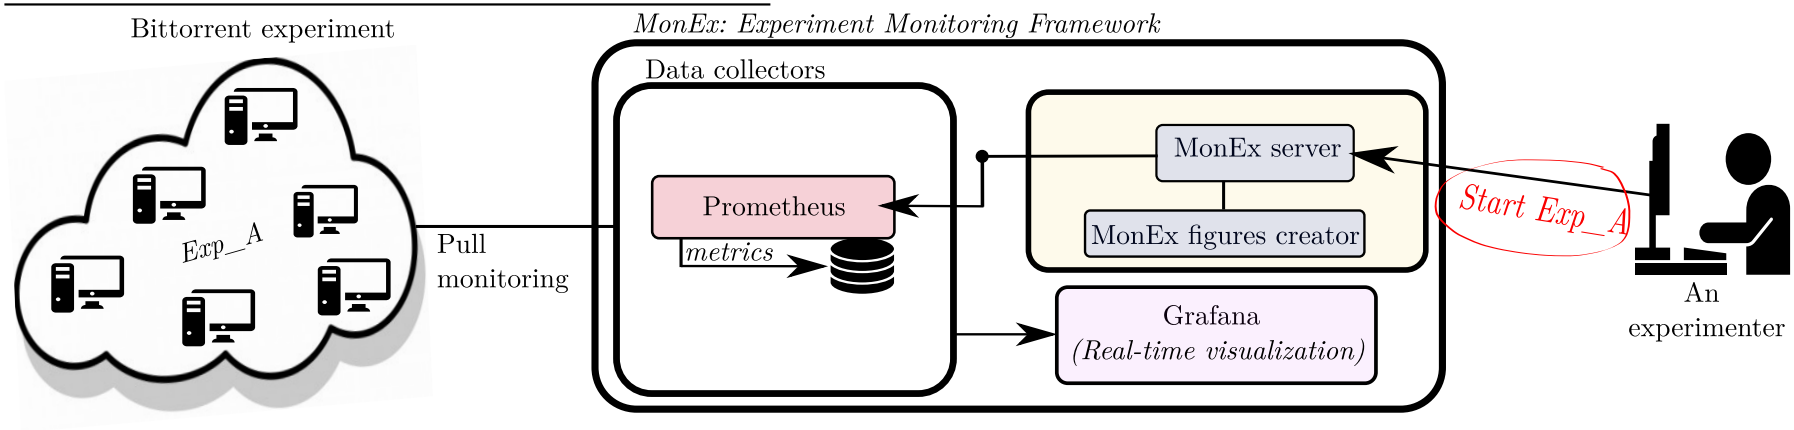


```
curl -X POST MonEx_IP:5000/exp/Exp_A
```

Use case experiments

First Experiment: workflow

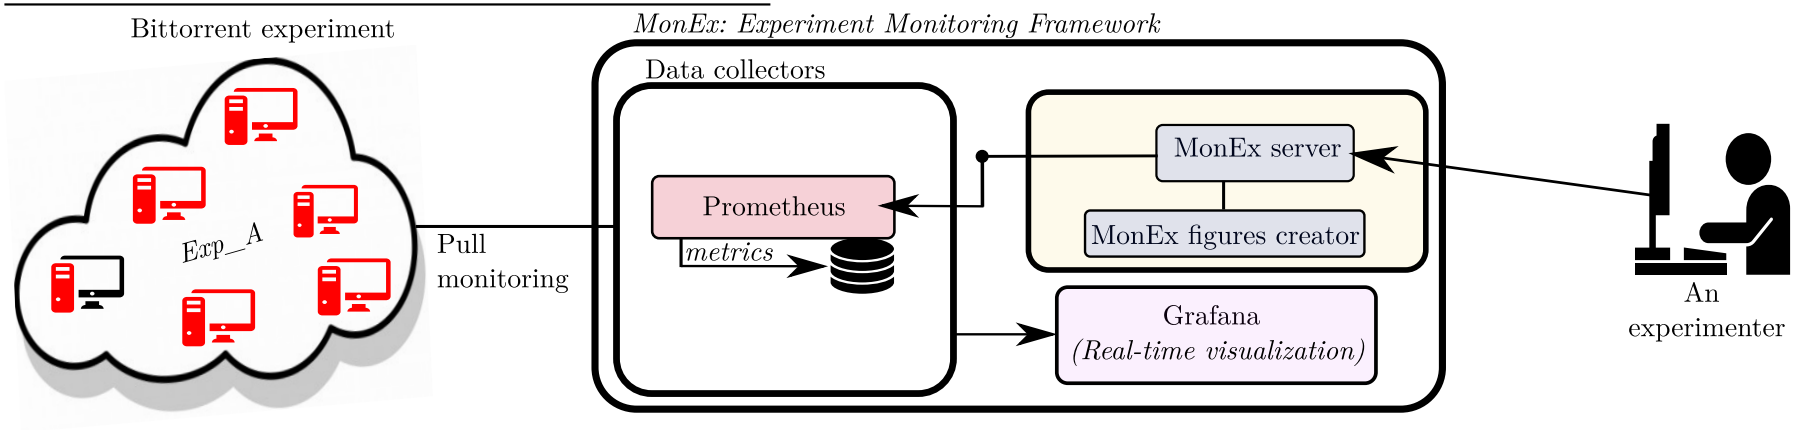

Run experiments' workload

Use case experiments

First Experiment: workflow

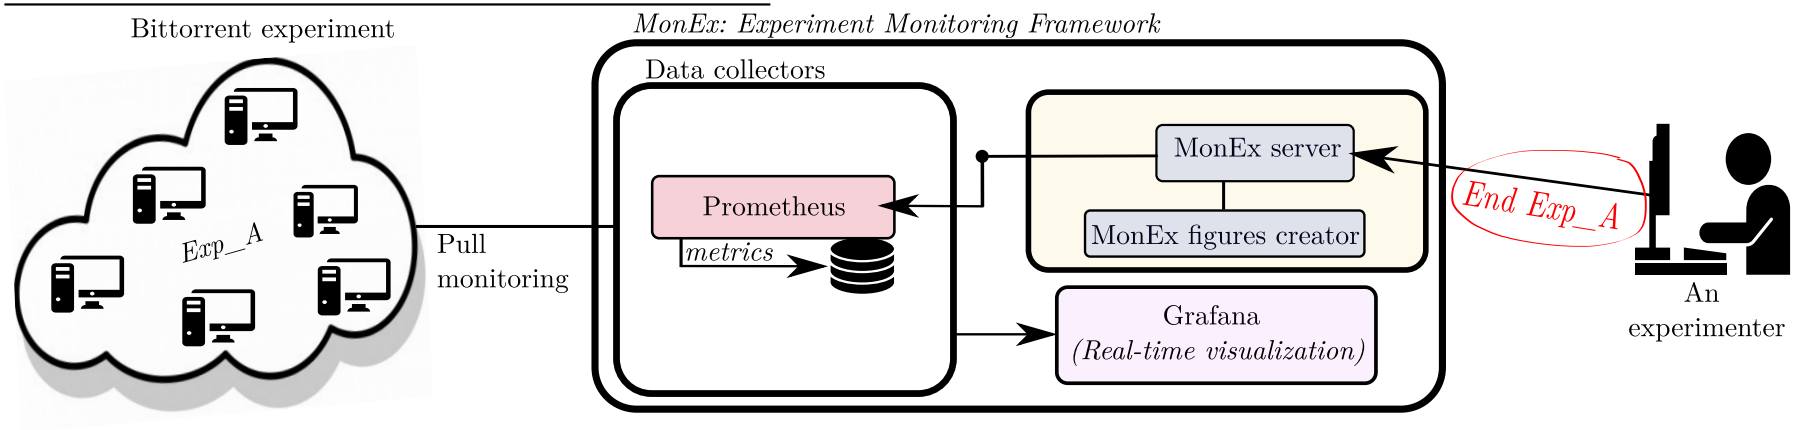


```
curl -X PUT MonEx_IP:5000/exp/Exp_A
```

Use case experiments

First Experiment: workflow

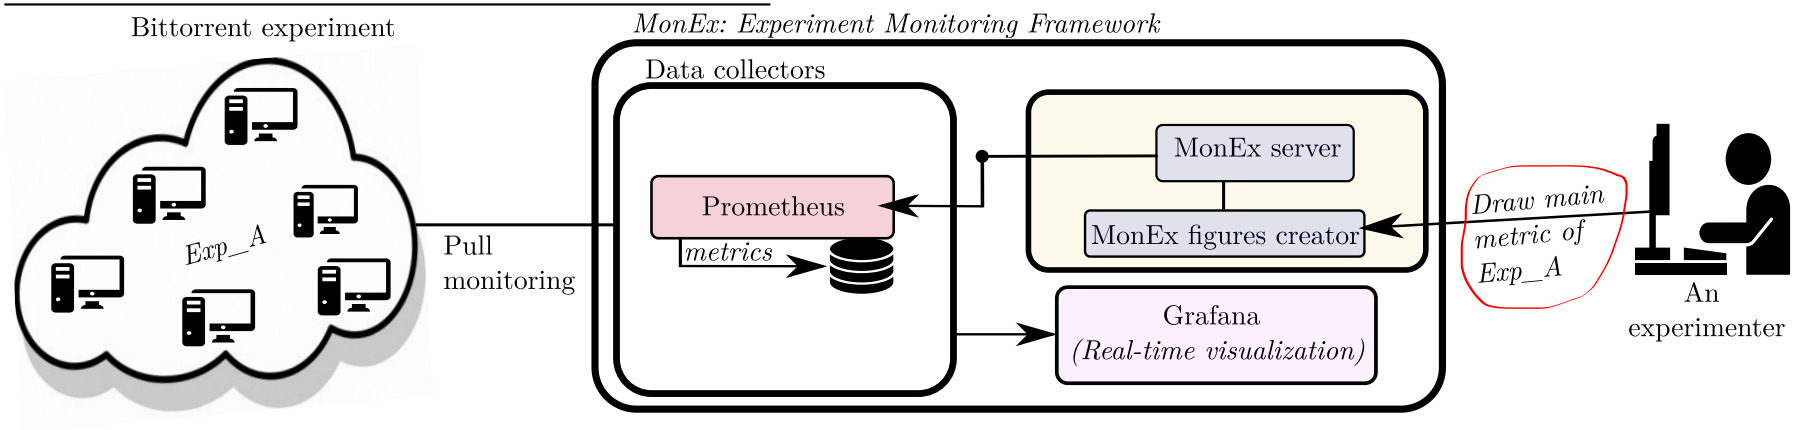


```
./monex-draw -S M_IP:5000 -exp
"Exp_A" -metric "completion" -type
"duration" -t "Torrent ..." -y
"Completion (%)" -x "Time (sec)" -n
```

Torrent completion of a 500MBytes file over time on a 30 kBytes/s network

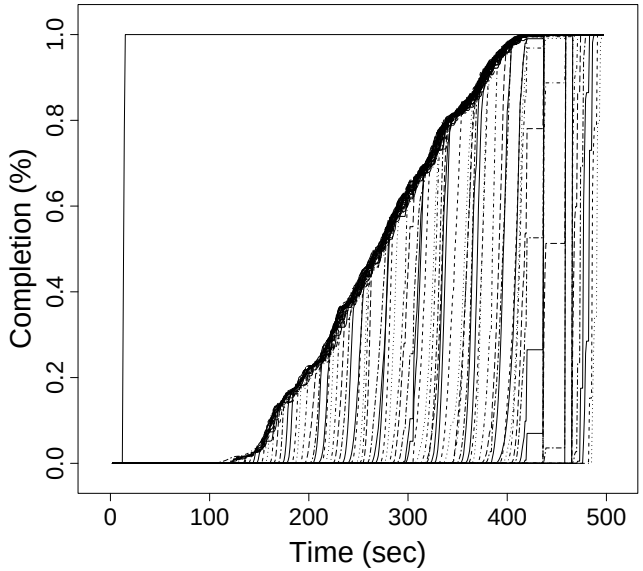

Use case experiments

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

  $\Rightarrow$ node-exporter $\rightarrow$ nodes spec.

  $\Rightarrow$ SNMP-exporter $\rightarrow$ PDUs data

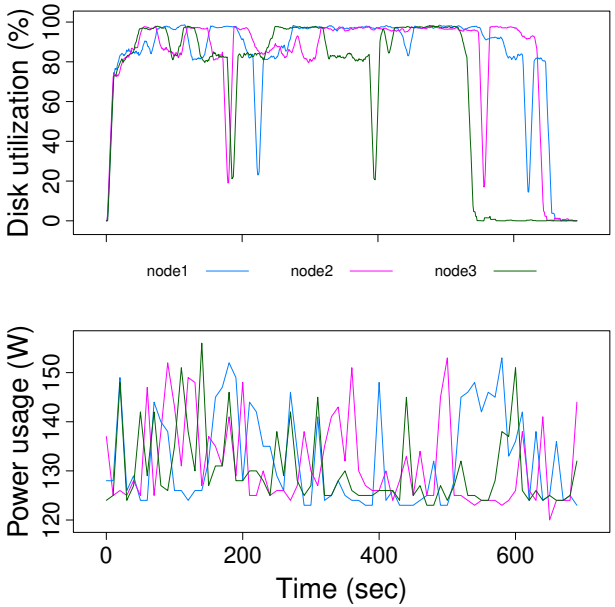

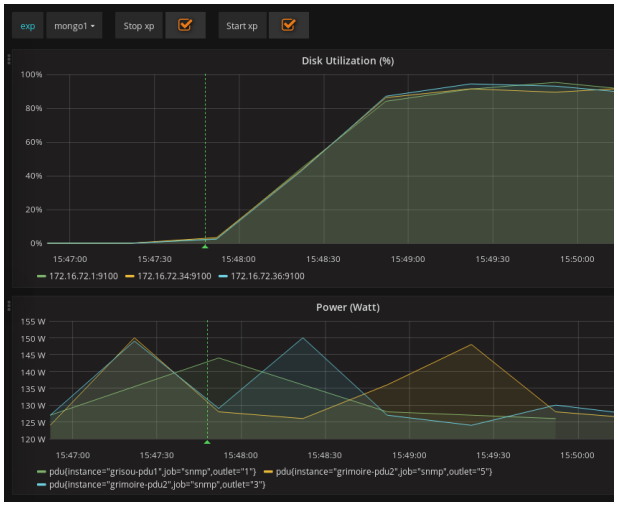

Third Experiment: time-dependent & heavy-measurement

- 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 an eBPF tool
- 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"
```

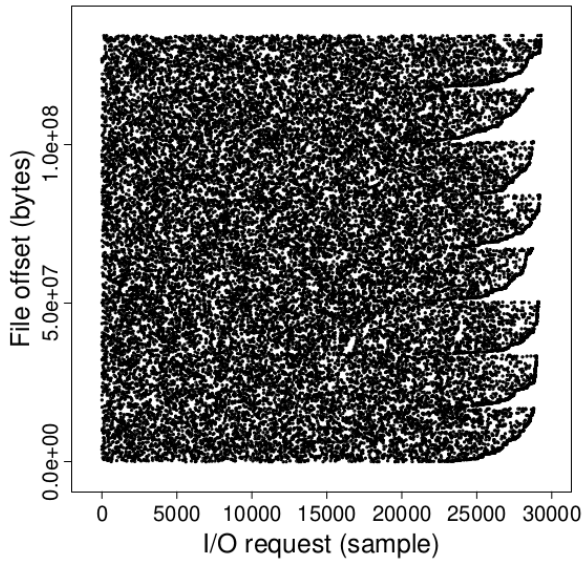

Grid'5000 and MonEx

MonEx leverages several features of Grid'5000:

- Monitoring traffic on separate LAN or another NIC
  - Targeting scalable experiments
    e.g. torrent experiment
- Up-to-date docs and rich API
  e.g. We automated the power experiment thanks to a Grid'5000 wiki page
- Co-operating easily with outside environment
  e.g. placing MonEx outside the testbed
- Deploying customized images with root privileges
  e.g. with built-in node-exporters

Conclusions

Conclusions:

- We defined the minimal requirements of EMFs
- Absence of related work with full coverage of requirements
- We designed MonEx EMF on top of Prometheus and InfluxDB monitoring systems
  - Usable on various experimentations' context
  - 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

Conclusions:

- We defined the minimal requirements of EMFs
- Absence of related work with full coverage of requirements
- We designed MonEx EMF on top of Prometheus and InfluxDB monitoring systems
  - Usable on various experimentations' context
  - 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 ?