Towards better tools for experiments on distributed systems

Lucas Nussbaum
lucas.nussbaum@loria.fr

Nancy-Université

Better?

- Push back the limits of what we can experiment on
  - Scale
  - Complexity

  $\Rightarrow$ **Answer harder questions**

- Increase the quality of our experiments
  - Credibility
  - Reproducibility

  $\Rightarrow$ **Provide more reliable answers**

Reproducibility?

Typology unclear [1]!

- **Identical replication:** (often called *replicability*)
  Exact same environment, method, parameters $\rightarrow$ same results?
- **Independent replication:**
  Same method, different {env, params, people} $\rightarrow$ same conclusions?
- **Conceptual replication:**
  Different experiment $\rightarrow$ same conclusions?

---

[1] O. S. Gómez, N. Juristo Juzgado, S. Vegas. Replications types in experimental disciplines. In proceedings of ESEM'2010.

Reproducibility?

Typology unclear [1]!

- **Identical replication:** (often called *replicability*)
  Exact same environment, method, parameters $\rightarrow$ same results?
- **Independent replication:**
  Same method, different {env, params, people} $\rightarrow$ same conclusions?
- **Conceptual replication:**
  Different experiment $\rightarrow$ same conclusions?

Similar requirements: understand, capture & describe . . .

- environment, parameters
- method, process
- raw results, reasoning

---

[1] O. S. Gómez, N. Juristo Juzgado, S. Vegas. Replications types in experimental disciplines. In proceedings of ESEM'2010.

Reproducibility?

Typology unclear [1]!

- **Identical replication:** (often called *replicability*)
  Exact same environment, method, parameters $\rightarrow$ same results?
- **Independent replication:**
  Same method, different {env, params, people} $\rightarrow$ same conclusions?
- **Conceptual replication:**
  Different experiment $\rightarrow$ same conclusions?

Similar requirements: understand, capture & describe . . .

- environment, parameters
- method, process
- raw results, reasoning

} software can help

[1] O. S. Gómez, N. Juristo Juzgado, S. Vegas. Replications types in experimental disciplines. In proceedings of ESEM'2010.

General overview

General overview

Experimental methodology:

experiment design & planning (workflow) ; description of scenarios, of experimental conditions ; definition of metrics ; laboratory journal ; analysis and visualization of results

General overview

Experimental methodology:

experiment design & planning (workflow) ; description of scenarios, of experimental conditions ; definition of metrics ; laboratory journal ; analysis and visualization of results

Experimental testbed (e.g Grid'5000):

reconfigurable hardware and network; isolation; some instrumentation and monitoring

General overview

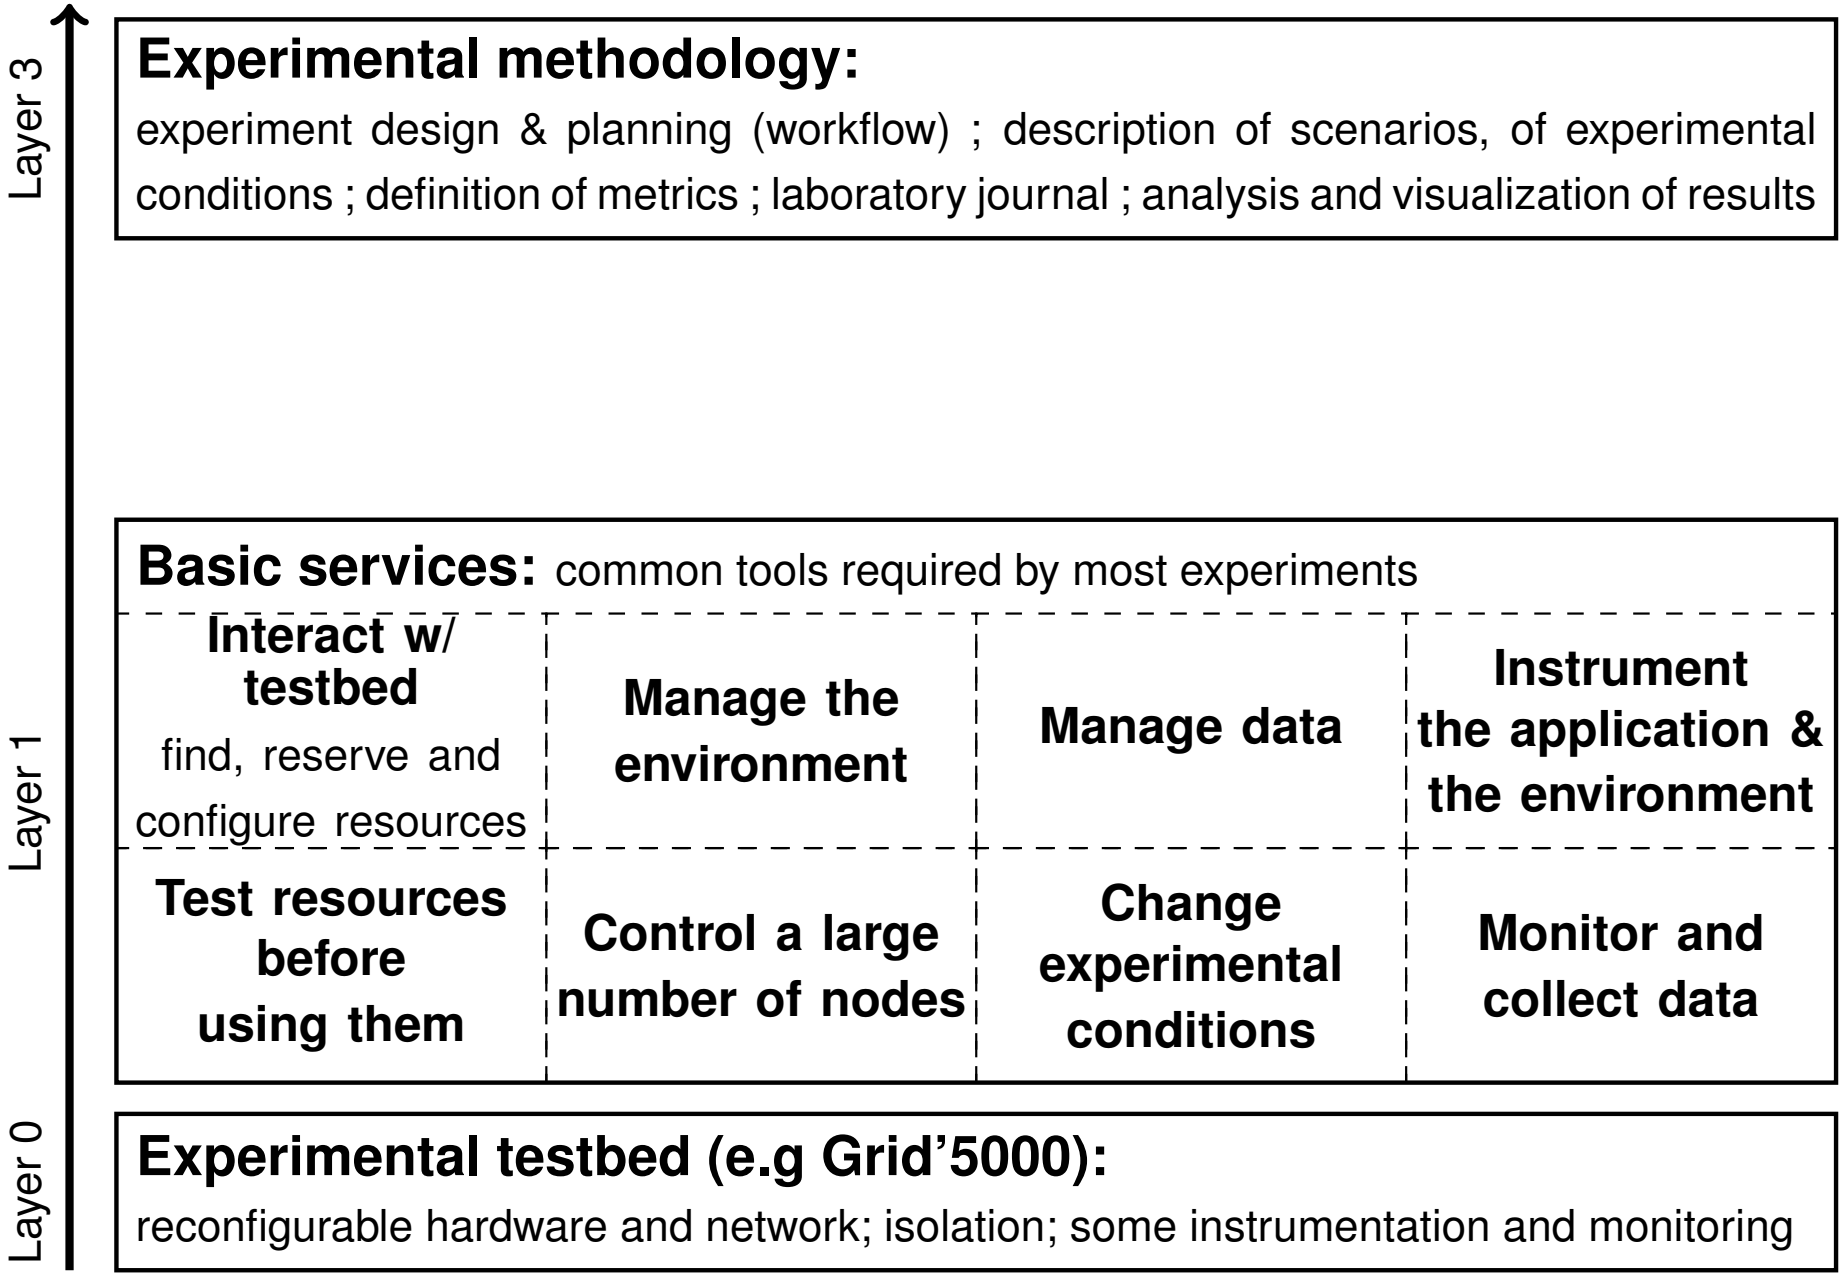

General overview

Layer 3

Experimental methodology:

experiment design & planning (workflow) ; description of scenarios, of experimental conditions ; definition of metrics ; laboratory journal ; analysis and visualization of results

Layer 2

Orchestration of experiments:

organize the execution of complex and large-scale experiments (workflow) ; run experiments unattended and efficiently ; handles failures ; compose experiments

Layer 1

Basic services: common tools required by most experiments

Interact w/ testbed find, reserve and configure resources	**Manage the environment**	**Manage data**	**Instrument the application & the environment**
Test resources before using them	**Control a large number of nodes**	**Change experimental conditions**	**Monitor and collect data**

Layer 0

Experimental testbed (e.g Grid'5000):

reconfigurable hardware and network; isolation; some instrumentation and monitoring

General overview

Layer 3

Experimental methodology:
experiment design & planning (workflow) ; description of scenarios, of experimental conditions ; definition of metrics ; laboratory journal ; analysis and visualization of results

Layer 2

Orchestration of experiments:
organize the execution of complex and large-scale experiments (workflow) ; run experiments unattended and efficiently ; handles failures ; compose experiments

Layer 1

Basic services: common tools required by most experiments

Interact w/ testbed find, reserve and configure resources	**Manage the environment**	**Manage data**	**Instrument the application & the environment**
Test resources before using them	**Control a large number of nodes** Taktuk talk	**Change experimental conditions**	**Monitor and collect data**

Layer 0

Experimental testbed (e.g Grid'5000):
reconfigurable hardware and network; isolation; some instrumentation and monitoring

General overview

Layer 3

Experimental methodology:

experiment design & planning (workflow) ; description of scenarios, of experimental conditions ; definition of metrics ; laboratory journal ; analysis and visualization of results

Layer 2

Orchestration of experiments:

organize the execution of complex and large-scale experiments (workflow) ; run experiments unattended and efficiently ; handles failures ; compose experiments

2

Layer 1

Basic services: common tools required by most experiments

Interact w/ testbed find, reserve and configure resources	**Manage the environment**	**Manage data**	**Instrument the application & the environment**
Test resources before using them	**Control a large number of nodes**	**Change experimental conditions**	**Monitor and collect data**

Taktuk talk

1

Layer 0

Experimental testbed (e.g Grid'5000):

reconfigurable hardware and network; isolation; some instrumentation and monitoring

Changing experimental conditions

Changing experimental conditions

Changing experimental conditions




Changing experimental conditions

When the testbed is too perfect. . . change the weather!

- Alter the platform so that it matches the experimental conditions you need
  - Introduce heterogeneity in an homogeneous cluster
  - Emulate a complex network topology
- Inject load & faults

distem – distributed systems emulator

- Uses modern Linux technology to steal resources from applications
  - Linux Traffic Control & netem
  - Linux Containers
- Easy to install and to use
  - Command-line interface
  - REST API for scripting
- Scalable: 10 000-vnodes in a single experiment

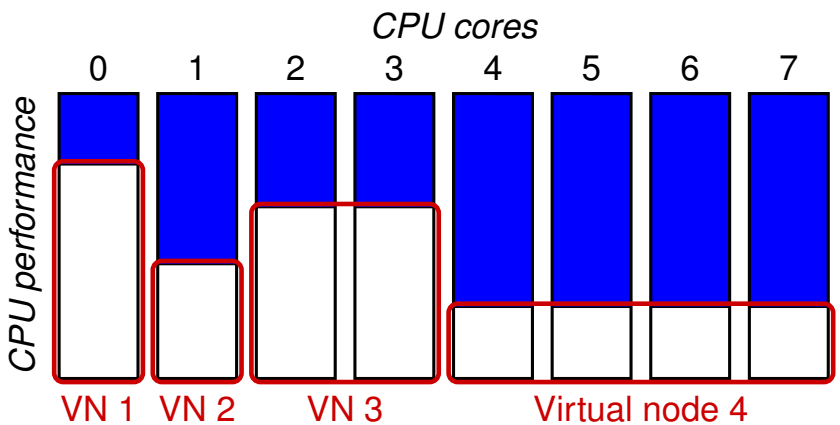

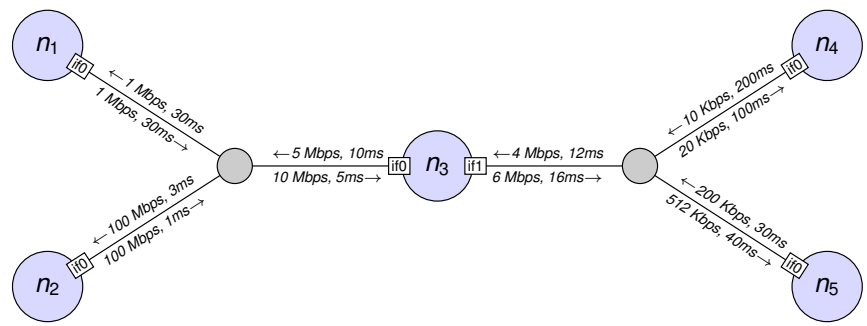

http://distem.gforge.inria.fr/

Orchestrating experiments

- Typical experiment on Grid'5000 today:
  - Quick and dirty shell scripts
    Best case: quick and dirty Ruby scripts
  - Many manual steps
- Work on how to organize experiments and combine all those services
- Requirement to increase scale, complexity, reproducibility
- Not a new problem:
  - Emulab
  - PlanetLab & GENI
  - Computational sciences
  - Other sciences
  - some Grid'5000 attempts

Emulab

Experiment management integrated in the Emulab framework

- Eric Eide, Leigh Stoller, Tim Stack, Juliana Freire, and Jay Lepreau. Integrated scientific workflow management for the emulab network testbed. USENIX'06
- Eric Eide, Leigh Stoller, and Jay Lepreau. An experimentation workbench for replayable networking research. NSDI'2007

`http://www.cs.utah.edu/flux/workbench/`

PlanetLab & GENI

- Jeannie Albrecht, Christopher Tuttle, Alex C. Snoeren, and Amin Vahdat. PlanetLab Application Management Using **Plush**. ACM Operating Systems Review (SIGOPS-OSR), 40(1), January 2006
- **Gush**: GENI User Shell – `http://gush.cs.williams.edu/trac/gush` Jeannie Albrecht and Danny Yuxing Huang. Managing Distributed Applications using Gush. TridentCom 2010.
- GENI WG: *GENI Experiment Workflow and Services*
  - Scope: What do experimenter-users need from GENI? Consider planning, scheduling, running, debugging, analyzing experiments; long running experiments & how they grow; archiving data.

Computational sciences

- Many scientific workflow management systems for computational sciences
- Kepler, Taverna, Triana, VisTrails, . . .

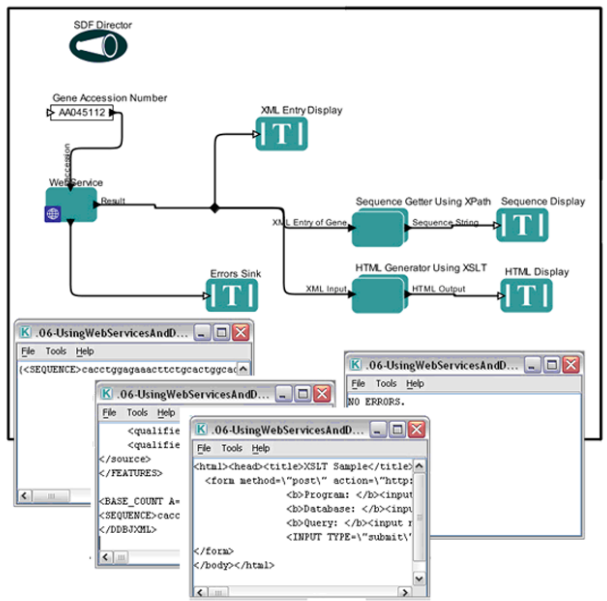

`https://kepler-project.org/`

Computational sciences

- Many scientific workflow management systems for computational sciences
- Kepler, Taverna, Triana, VisTrails, . . .

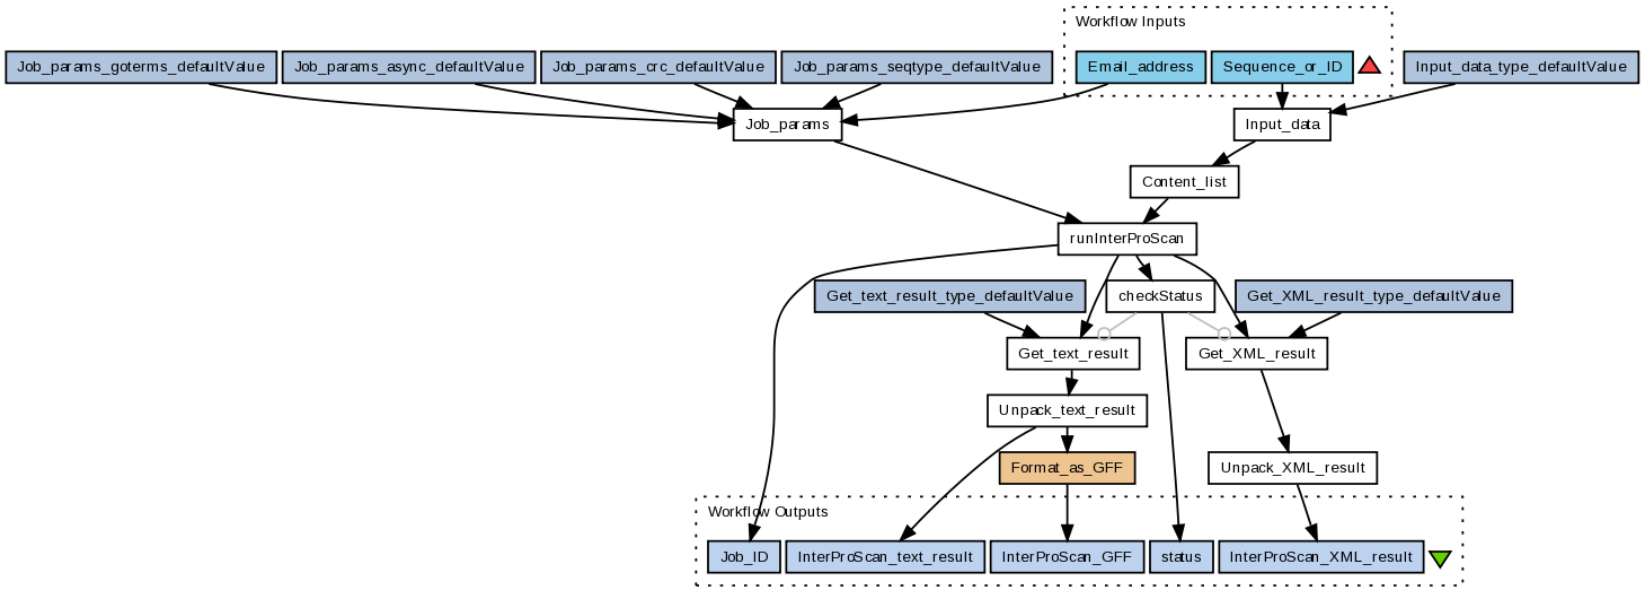

`http://www.taverna.org.uk/`
Integrated with `http://www.myexperiment.org/`

Other sciences

NATIONAL INSTRUMENTS

Cart | Help Search

Improve your ni.com experience. Login or Create a user profile.

MyNI | Contact NI | Products & Services | Solutions | Support | NI Developer Zone | Academic | Events | Company

All Products and Services > NI LabVIEW

The LabVIEW Environment

The LabVIEW Environment

Engineers and scientists can rapidly and cost-effectively interface with measurement and control hardware, analyze data, share results, and distribute systems through intuitive graphical programming.

» What is LabVIEW?

How Can I Use LabVIEW?

Applications are as varied as the engineers who create them. Fortunately, LabVIEW combines the flexibility of a programming language with the power of an advanced engineering tool so users can complete their projects regardless of their unique, custom requirements.

Browse all applications

LabVIEW Product Options

Compare LabVIEW development systems, explore add-ons, and see pricing.

Shop for LabVIEW Products

On Grid'5000

Several attempts already:

- GRUDU (Lyon – GRAAL)
  GUI focused on reservation and deployment
- NXE (Romaric Guillier, Lyon – RESO)
  automating network experiments
- Expo (Brice Videau & Olivier Richard, Grenoble)
  Ruby DSL, integrated with Taktuk
- Execo (Matthieu Imbert, Lyon)
  `https://gforge.inria.fr/projects/execo/`
  focus on large-scale command execution
- g5k-campaign (Cyril Rohr, Rennes)
  `http://g5k-campaign.gforge.inria.fr/`
  focus on reservation and deployment, uses Grid'5000 API

But:

- Not clear where we should go
- More attempts are probably needed

Workflows?

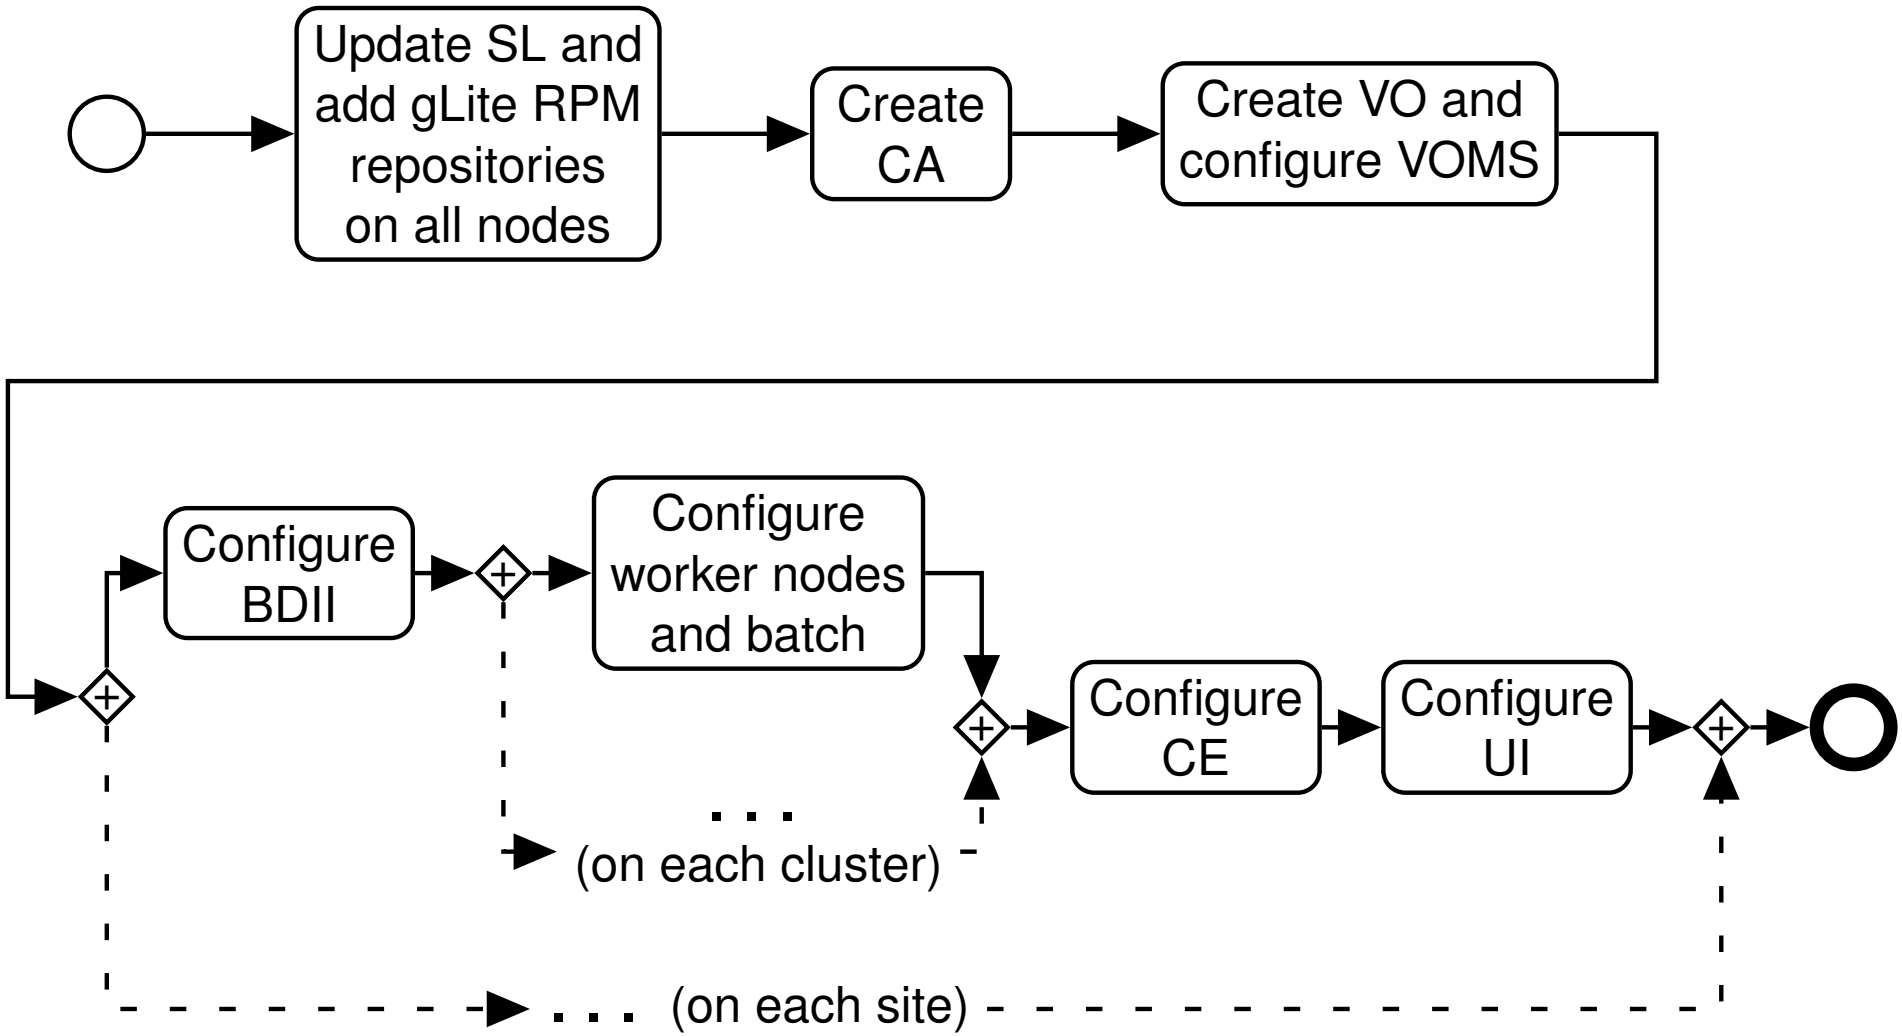

Need specialized workflow engine?

Conclusion: a long way forward