

Scientific Workflow Repeatability through Cloud-Aware Provenance

Khawar Hasham, Kamran Munir, Jetendre Shamdasani,
Richard McClatchey

University of the West of England (UWE)

Introduction

- Scientific community is using workflows to process and analyse large amounts of data
- Recent studies [1] are evaluating workflow execution on Cloud
- Cloud provides resources on-demand and nature of resources are dynamic
 - They are acquired or released on-demand
- Provenance is an important data;
 - Help in verifying and debugging the execution
 - Verify inputs and results
 - Repeating an experiment

Motivation

- Around 80% of the workflows cannot be reproduced, and 12% of them are due to the lack of information about the execution environment [2]
- Workflow Provenance capturing in Cloud
 - Challenging due to the dynamic and geographically distributed nature of Cloud computing [3]
- Workflow Repeatability in Cloud
 - Exact resource configuration should be known
 - Resources are acquired by passing on configuration information
 - Helps in achieving similar performance with similar resources
- Workflows executed on Cloud need to provide existing provenance information along with the Cloud resource information

Literature Survey

- Provenance applications
 - Enables a researcher to reproduce to verify experimental results
- Provenance approaches or systems in Grid
 - Chimera [4] stores physical location of machine
 - Vistrails [5] and Pegasus/Wings [6] tackle the workflow evolution provenance
 - However, They lack computational environment or infrastructure information in the collected provenance
- Provenance approaches in Cloud
 - Hypervisor level information [7] (physical to virtual resource); mainly to detect data flow
 - File access level information [8]
 - These approaches don't consider workflow executions, which is the focus of this paper

Literature Survey

- Literature survey on Reproducibility
 - Systems like ReproZip [8] or CARE [9]
 - Work on individual job or execution, thus not applicable to the workflow level.
 - They don't consider the resource on which execution is taking place
 - Takes a lot of storage space.
 - Recently published (August 2014) [10], Semantic based approach uses annotations exhaustively
 - Depends on user provided data for resource provisioning and configuration
- Outcome: We can collect workflow provenance as well as virtualization level information, but we need to combine them

Re-provision an execution infrastructure on Cloud

- Focus of this paper is on re-provisioning the similar resources on Cloud, thus enabling repeatability of workflow execution on similar resources
- Information required for provisioning Cloud resources
 - Parameters require to request a new resource on the Cloud, to have similar execution infrastructure
 - CPU, RAM, Hard disk, OS Image used
 - Is this information sufficient to re-provision a resource on the Cloud? Yes
- In this paper, it is assumed that OS image contains all the required software libraries

Augmenting Provenance information with Cloud resource information

- **Capture**

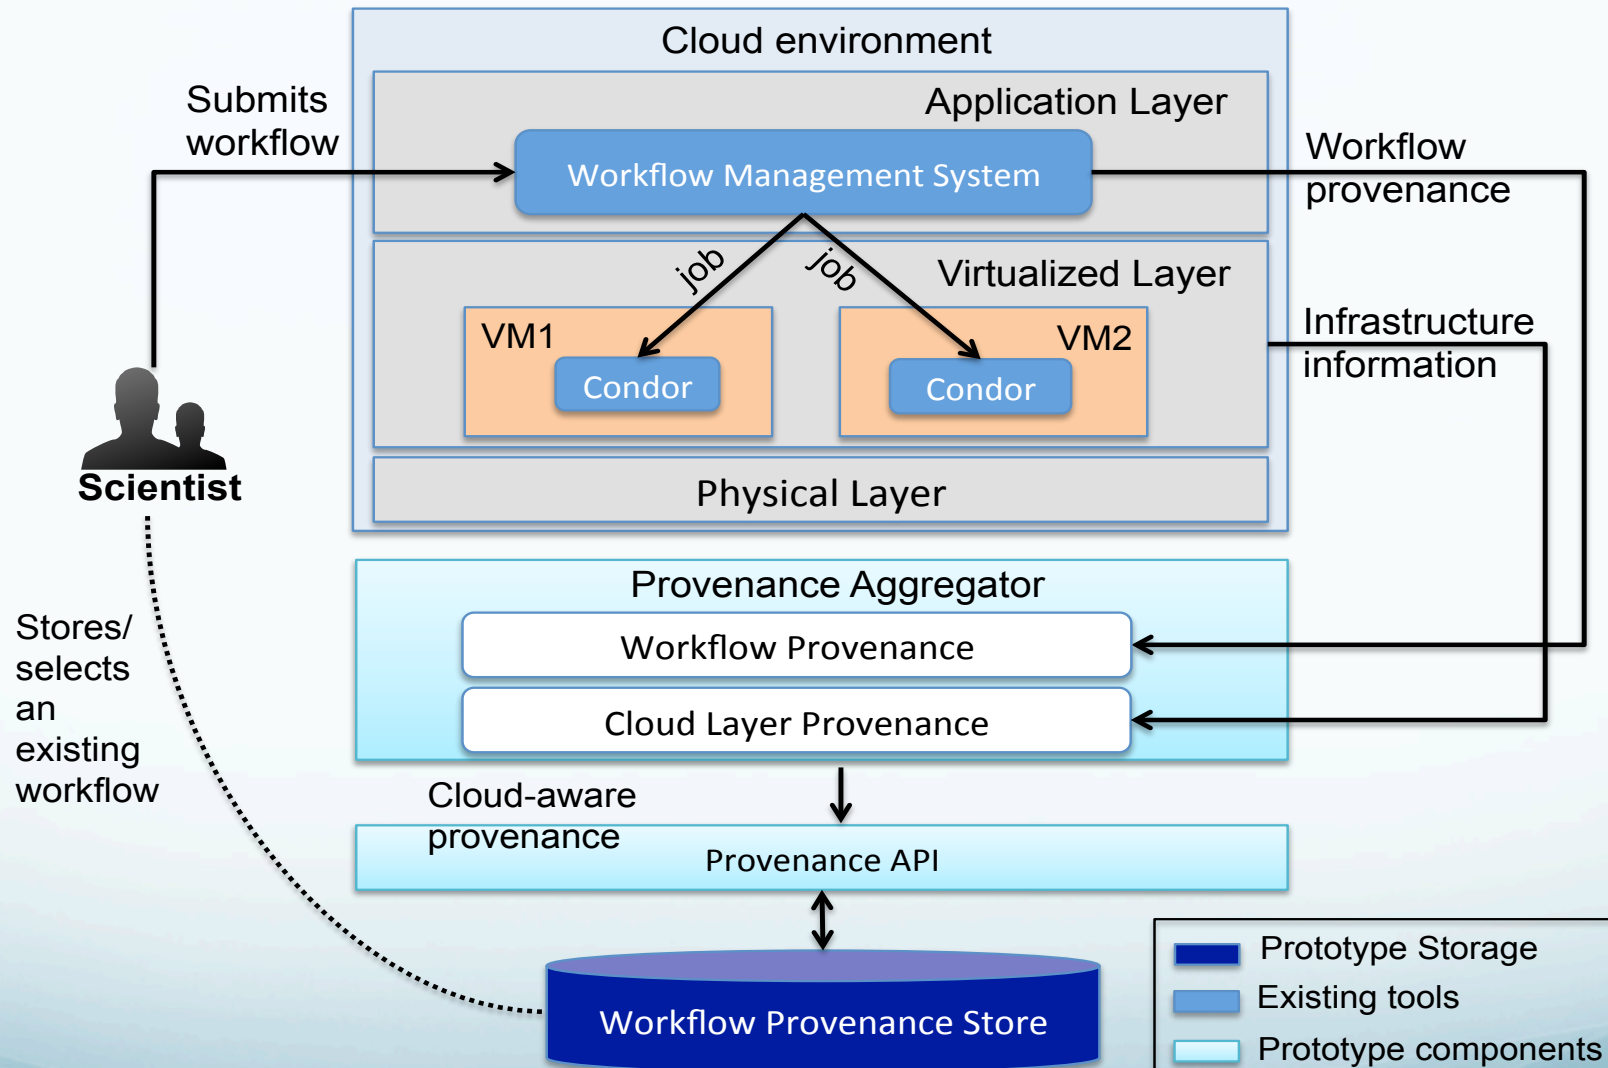
- Collect provenance data from the underlined Workflow Management System using APIs or database
 - Pegasus database
- Find APIs to interact with Cloud middleware to collect infrastructure (Cloud resource) information.
 - Apache libcloud¹

- **Combine**

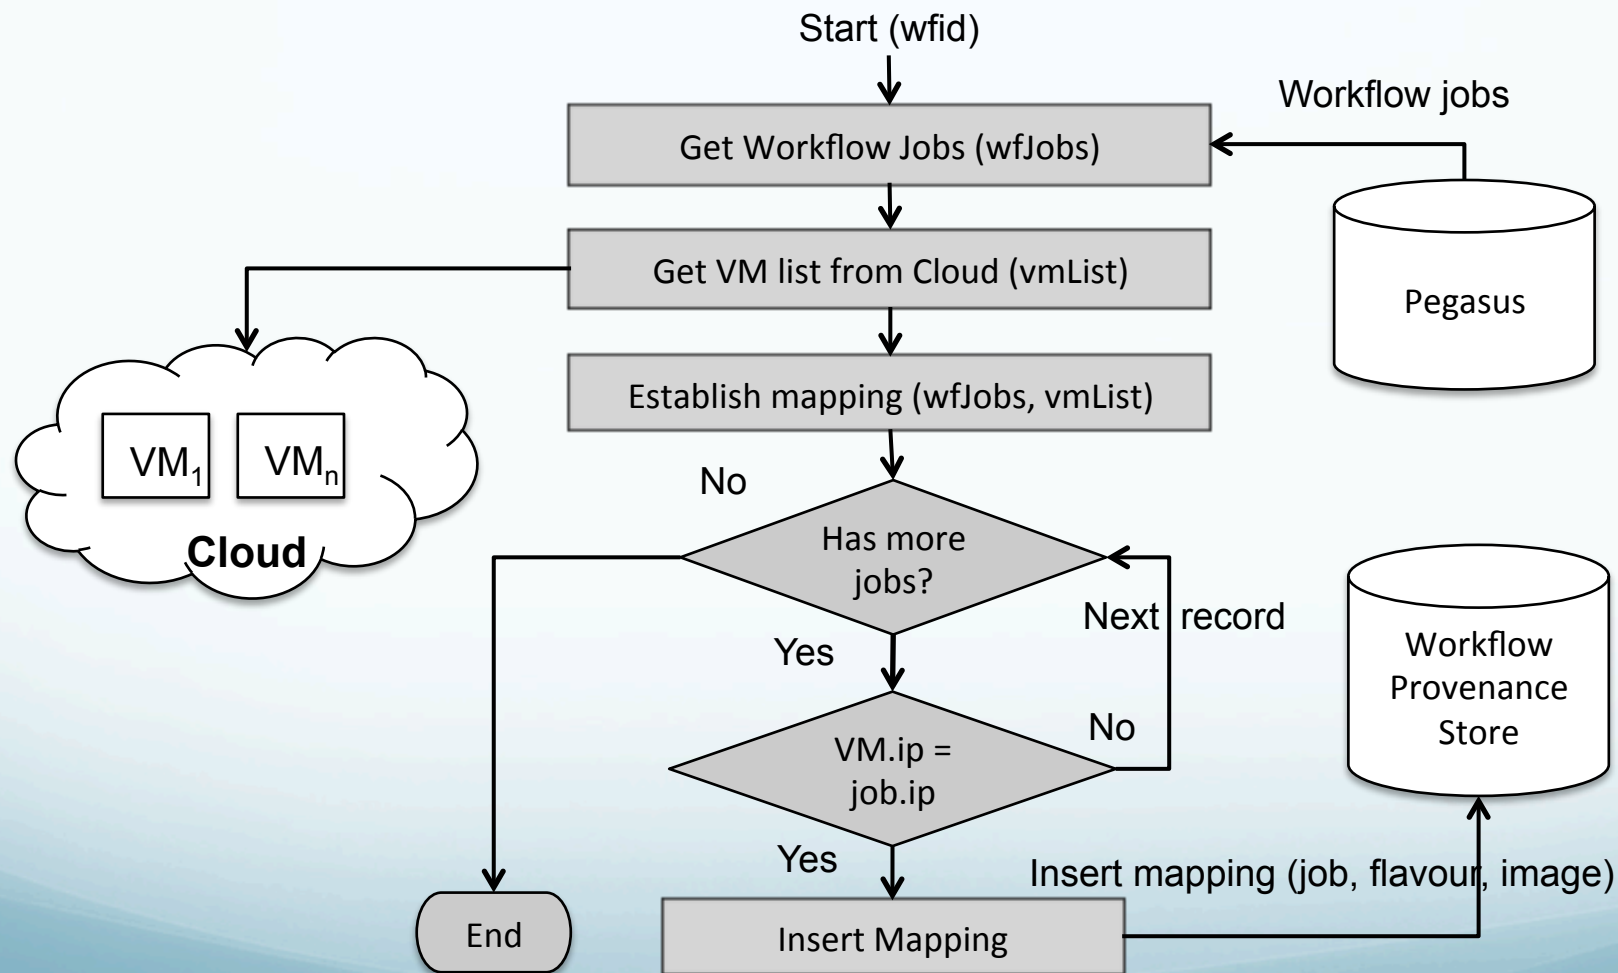
- How to combine them together
 - Identify the mapping points between two sets of information.
 - In current prototype, it is based on IP address of the resource.
- This provides detailed information about resources used for a workflow execution on Cloud

¹: <https://libcloud.apache.org>

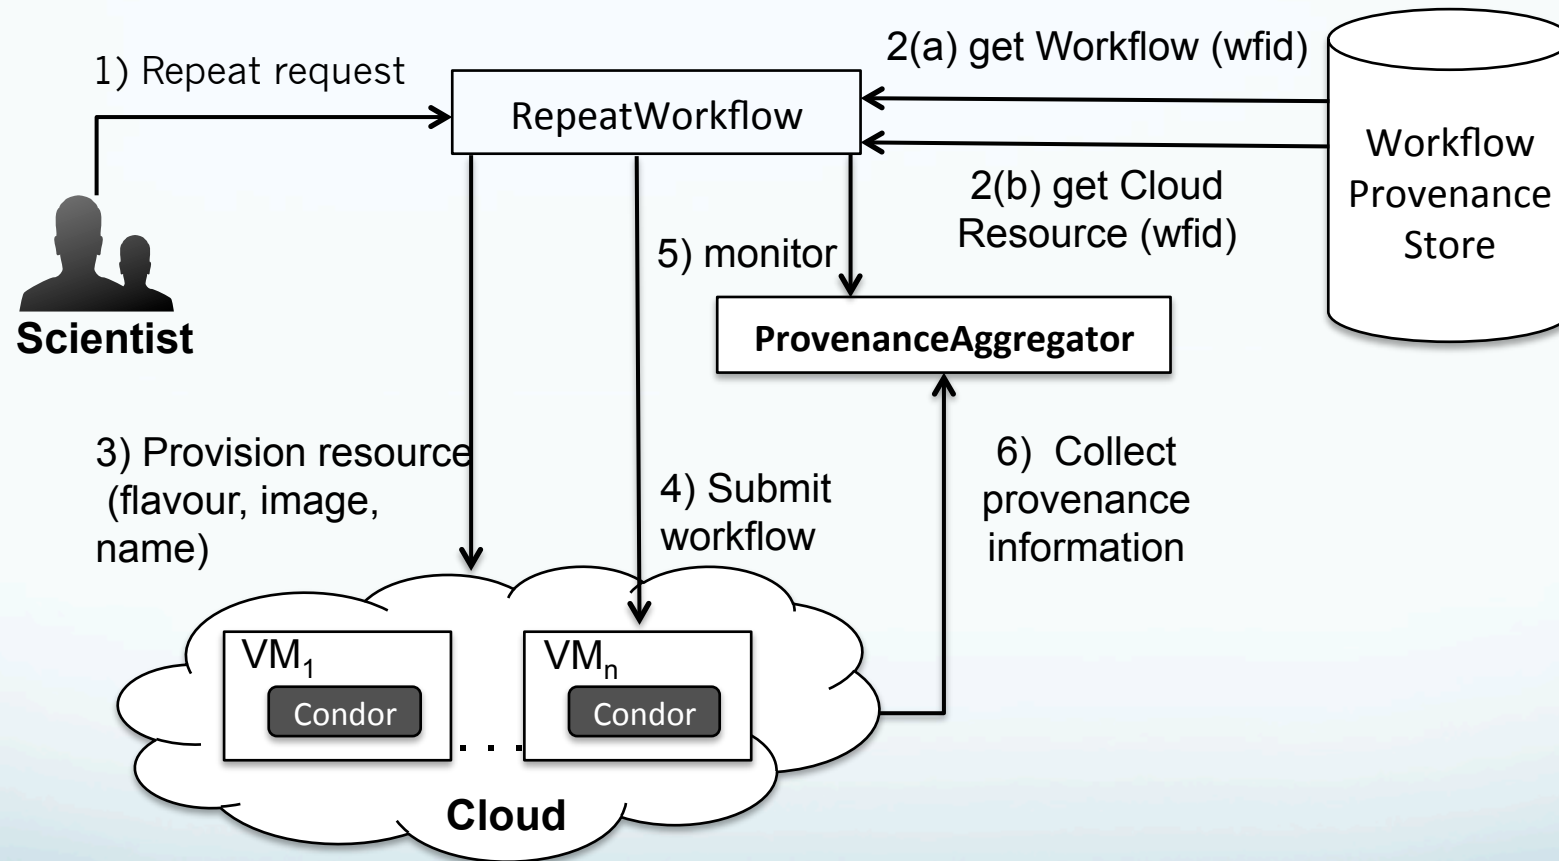
Proposed architecture



Provenance Aggregator: combining workflow provenance with Cloud information



Repeat a workflow execution using Cloud-aware Provenance



Prototype Implementation

- A proof of concept prototype has been developed using Apache libcloud API to interact with Openstack Cloud middleware.
- Pegasus WMS along with condor cluster on virtual machines is used to submit and execute workflows.
- A simple WordCount workflow with 4 jobs is executed
- Mapping between Pegasus Provenance and Cloud Virtual layer information is established
 - The virtual layer information is mapped with the workflow provenance retrieved from Pegasus
- Workflow re-execution after re-provisioning the Cloud resources using Cloud-aware provenance

Initial results – similar infrastructure

Original workflow execution – infrastructure information on Cloud

wf_id	hostip	nodename	middleware	flavorid	minRAM	minHD	vCPU	image_name	image_id
114	172.16.1.49	osdc-vm3.novalocal	OpenStack	2	2048	20	1	wf_peg_repeat	f102960c-557c-4253-8277-2df5ffe3c169
114	172.16.1.98	mynode.novalocal	OpenStack	2	2048	20	1	wf_peg_repeat	f102960c-557c-4253-8277-2df5ffe3c169

Reproduced workflow execution – infrastructure information on Cloud

wf_id	hostip	nodename	middleware	flavorid	minRAM	minHD	vCPU	image_name	image_id
117	172.16.1.183	osdc-vm3-rep.novalocal	OpenStack	2	2048	20	1	wf_peg_repeat	f102960c-557c-4253-8277-2df5ffe3c169
117	172.16.1.187	mynode-rep.novalocal	OpenStack	2	2048	20	1	wf_peg_repeat	f102960c-557c-4253-8277-2df5ffe3c169
122	172.16.1.114	osdc-vm3-rep.novalocal	OpenStack	2	2048	20	1	wf_peg_repeat	f102960c-557c-4253-8277-2df5ffe3c169
122	172.16.1.221	mynode-rep.novalocal	OpenStack	2	2048	20	1	wf_peg_repeat	f102960c-557c-4253-8277-2df5ffe3c169

Future work

- Improve the initial prototype
- Detailed evaluation
 - Comparison of outputs produced by two workflow runs
 - Detailed provenance comparison of two workflow runs to evaluate the workflow reproducibility
 - Workflow structure
 - Outputs
 - Infrastructure
- Evaluate the impact of different resource configurations on workflow's overall execution performance



Thank you!

References

1. G. Juve and E. Deelman, "Scientific workflows and Cloud", Published in Crossroads - Plugging Into the Cloud Magazine, vol. 16, pp. 14-18, March, 2010
2. Zhao, J., Gomez-Perez, J.M., et al., "Why workflows break - understanding and combating decay in taverna workflows", in IEEE 8th International Conference on E-Science, pp. 1-9, 2012
3. Y. Zhao, Z. Fei, I. Raicu, and S. Lu, "Opportunities and Challenges in Running Scientific Workflows on the Cloud", ICM Proceedings of the 2011 International Conference on Cyber-Enabled Distributed Computing and Knowledge Discovery, pp. 455-462, 2011
4. I. Foster, J. Vockler, M. Eilde, and Y. Zhao, "Chimera: A virtual data system for representing, querying, and automating data derivation", In International Conference on Scientific and Statistical Database Management, pp. 37-46, July 2002
5. C. Scheidegger, D. Koop, E. Santos, H. Vo, S. Callahan, J. Freire and C. Silva, "Tackling the Provenance Challenge One Layer at a Time", In Journal of Concurrency and Computation: Practice & Experience, Vol. 20 Issue 5, pp. 473-483, April 2008
6. J. Kim et al., "Provenance trail in the Wings-Pegasus system", Concurrency and Computation: Practice and Experience, vol. 20, pp. 587-597, 2008 doi: 10.1002/cpe.1228
7. P. Macko, M. Chiarini and M. Seltzer, "Collecting Provenance via the Xen Hypervisor", 3rd Workshop on the Theory and Practice of Provenance (TaPP '11), June 2011
8. O. Q. Zhang, M. Kirchberg, R. K.L. Ko and B. Sung Lee, "How to Track Your Data: The Case for Cloud Computing Provenance", IEEE Third International Conference on Cloud Computing Technology and Science (CloudCom), pp.446-453, Nov. 29 2011-Dec. 1 2011
9. F. Chirigati, D. Shasha and J. Freire, "ReproZip: Using Provenance to Support Computational Reproducibility", 5th USENIX Workshop on the Theory and Practice of Provenance (TaPP), 2013
10. J. Yves, V. Cedric, and D. Remi, "CARE, the Comprehensive Archiver for Reproducible Execution". Proceedings of the 1st ACM SIGPLAN Workshop on Reproducible Research Methodologies and New Publication Models in Computer Engineering, TRUST'14, pp. 1- 7, 2014
11. I. Santana-Perezantana-Perez, R. Ferreira da Silva, M. Rynge, E. Deelman, M. Pérez-Henández and O. Corcho, "Leveraging Semantics to Improve Reproducibility in Scientific Workflows", reproducibility@XSEDE workshop, 2014