

## Accepted Manuscript

**Title:** Middleware Building Blocks for Workflow Systems

**Authors:** Matteo Turilli; Vivek Balasubramanian; Andre Merzky; Ioannis Paraskevakos; Shantenu Jha

**DOI:** <https://doi.org/10.1109/MCSE.2019.2920048>

**Publication:** Computing in Science & Engineering

**Citation:** Turilli, M., Balasubramanian, V., Merzky, A., Paraskevakos, I., & Jha, S. (2019). Middleware building blocks for workflow systems. *Computing in Science & Engineering*, 21(4), 62-75.

This article has been accepted for publication and undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the Version of Record. Please cite this article as doi: .....

# Middleware Building Blocks for Workflow Systems

Matteo Turilli<sup>1</sup>, Vivek Balasubramanian<sup>1</sup>, Andre Merzky<sup>1</sup>, Ioannis Paraskevagos<sup>1</sup>, Shantenu Jha<sup>1,2</sup>

<sup>1</sup> Rutgers, the State University of New Jersey, Piscataway, NJ 08854, USA

<sup>2</sup> Brookhaven National Laboratory, Upton, New York, 11973

**Abstract**—This paper describes a building blocks approach to the design of scientific workflow systems. We discuss RADICAL-Cybertools as one implementation of the building blocks concept, showing how they are designed and developed in accordance with this approach. Four case studies are presented, discussing how RADICAL-Cybertools are integrated with existing workflow, workload, and general purpose computing systems to support the execution of scientific workflows. This paper offers three main contributions: (i) showing the relevance of the design principles of self-sufficiency, interoperability, composability and extensibility for middleware to support scientific workflows on high performance computing machines; (ii) illustrating a set of building blocks that enable multiple points of integration, which results in design flexibility and functional extensibility, as well as providing a level of “unification” in the conceptual reasoning across otherwise very different tools and systems; and (iii) showing how these building blocks have been used to develop and integrate workflow systems.

## I. INTRODUCTION

Sophisticated and scalable workflows have come to epitomize advances in computational science, especially for “big science” projects, such as those in high-energy physics or astronomy. Most workflow systems for “big science” were developed in an era when the scientific distributed computing infrastructure and software ecosystem was relatively fragile, missing features and services. Thus, out of necessity, many such workflow systems adopted the end-to-end execution paradigm, and provided capabilities that enabled the end-to-end execution of workflows on distributed cyberinfrastructures.

The landscape of applications and the software infrastructure has however changed. High-throughput execution of tasks—the original driver of “big science” workflows—is still important, but is joined by other important functional and automation requirements. New application scenarios involve the time-sensitive integration of experimental data from large-scale instruments and observation systems with high-performance computing. Workflows are also becoming more pervasive across application types, scales and communities. Scientific insight typically requires computational campaigns with multiple distinct workflows, heterogeneous tasks and distinct runs. E.g., an application may involve distinct phases of parameter exploration and optimization, sensitivity analysis and uncertainty quantification.

Previously missing infrastructural capabilities that necessitated the development of end-to-end workflow systems are now relatively more reliable, better supported and more consistently available. The emergence of diverse Python-based task distribution and coordination systems, Apache data analysis tools, and container technologies provide useful examples.

Without negating end-to-end workflow systems where the socio-economic and socio-technical needs warrant them and make their use effective, an important but often overlooked fact is that many scientific applications do not use such pre-existing workflow systems. Instead, application developers tend to “roll their own”. For example, Ref. [1] enumerates in excess of 230 purported workflow systems—some partial, others closer to being end-to-end; some specific to a workload or functionality, others general-purpose; some stand-alone, others designed to be integrated with other systems. Although the full set of reasons underlying this trend defy simple reduction, some are worth highlighting: increasingly diverse, sophisticated and specific application requirements, coupled with the proverbial last mile customization challenge of traditional workflow systems. The enhanced infrastructure capabilities mentioned are an additional driver of the proliferation of “roll your own” workflow systems.

Self-evidently, this proliferation can not be reversed or even restrained, but has to be managed. It has many implications for users and developers of workflow systems, and raises important questions. Foremost, is it possible to implement workflow systems in an agile fashion to provide flexibility and sharing of capabilities while not constraining functionality, performance, or sustainability? Given the rich ecosystem of capabilities, how can the barrier to the integration of workflow systems with these capabilities be lowered? These questions are set against trends that suggest increasing functional richness and sophistication of workflows-based applications, and consequent demands on workflow systems.

We suggest that there is a need for a sustainable ecosystem of both existing and new software components from which tailored workflow systems can be composed. This entails the support of agile development and composition of workflow systems that can be responsive to the wide range of workflow requirements while leveraging the rich ecosystem of existing software capabilities. Further, this renders obsolete the historical focus of the workflow system community of competing by developing a workflow system that purports to be “better” than existing workflow systems, or the elusive need to interoperate with all other workflow systems, and emphasizes if not incentivizes the community towards development of collective capabilities.

This paper advocates for an integrative perspective to the design and development of scientific workflow systems by making the case for a building blocks approach as a first-order property. We postulate the building blocks leverages the emerging trends in software and distributed computing infras-

structure, and thus the approach supports a sustainable ecosystem of both existing and new software components from which tailored workflow systems can be composed. Building blocks enable expert contributions while lowering the breadth of expertise required of workflow system developers. Additional factors motivating an integrative approach to the design of workflow systems, include the ability to facilitate low-cost and sustainable solutions.

After a brief description of the building blocks approach and its four design principles of self-sufficiency, interoperability, composability, and extensibility, Sec. IV discusses how we used the building blocks approach to develop RADICAL-Cybertools. These are a set of software systems that can be used independently and integrated into middleware, among themselves and with third-party systems. RADICAL-Cybertools target High Performance Computing (HPC) machines to enable the execution of workflows from diverse scientific domains. We introduce a four-layered view of high-performance and distributed systems and we describe how each system implements distinctive functionalities for each layer.

Sec. V discusses how RADICAL-Cybertools complement and contribute to existing workflow systems and middleware, enabling the specification and execution of scientific workflows. We present four case studies of integrating RADICAL-Cybertools with end-to-end workflow systems (Swift), workload management systems (PanDA), general purpose computing frameworks (Spark and Hadoop) and domain-specific workflow systems (ExTASY, RepEx, HTBAC and ICEBERG). These integrations enable diverse scientific applications, including high-throughput jobs, multi-protocol simulations, adaptive workflows, data-intensive simulations, and image processing.

We conclude with a discussion of the practical impact of the case studies as well as the lessons learnt by testing the validity and feasibility of the building blocks approach. We highlight the benefits of implementing new capabilities into existing workflow systems by integrating the RADICAL-Cybertools. We also outline the limitations of our contributions as well as some open questions.

In this paper, we adopt the following definitions: A multi-task application can be represented as a **workflow**, i.e., a set of tasks with dependencies that determine the order of their execution. Subsets of these tasks can be **workloads**, i.e., tasks whose dependencies have been satisfied and that may be executed concurrently. In this way, a workflow provides a description of the application execution process while a workload identifies the tasks that are ready to be executed. We maintain that these characteristics are independent of the scale of the application, and the number of users, developers, and types of workflow.

Although the focus of this paper is on using the building blocks approach to design, develop, and integrate workflow systems, the approach we propose is in principle applicable to every type of distributed software system.

## II. RELATED WORK

We classify existing workflow systems into three categories, focusing only on those with the highest adoption and ongoing development. All-inclusive workflow systems such as Kepler, Swift, Fireworks, and Pegasus that provide full-featured, end-to-end capabilities that includes application creation, execution, monitoring and provenance. General-purpose workflow systems such as Ruffes, COSMOS, and GXPMake that enable end-to-end execution but prioritize the simplicity of their interfaces, limiting the range of capabilities. Finally, domain-specific workflow systems such as Galaxy, Taverna, BioPipe, and Copernicus that provide interfaces tailored to the requirements of specific domain scientists.

The decomposition of workflow systems into systems with high cohesion and low dependency supports decoupling of independent software development efforts and promotes the use of standardized interfaces. These systems are implemented in monolithic or modular fashion to support specific capabilities, and have been used to develop multiple workflow systems by integration. For example, Spark, Hadoop, and MapReduce can be integrated—with or without pipelining tools like Luigi, Toil, Airflow, Azkaban or Oozie—to create special-purpose workflows systems [18], [20]. Nonetheless, these tools are specifically tailored to data-oriented workflows, face several performance bottlenecks when ported to HPC machines, and require dedicated deployment [4]. Research in interoperability of HPC systems with data-parallel frameworks is ongoing and provides and extends middleware to efficiently support data-oriented workflows on HPC. A few examples are Pilot-Hadoop and Pilot-Spark, Twister or Pilot-Streaming.

Modularity, in software deployment, has evolved from chroot, jails and Solaris zones and, more in general, to what is called the “UNIX Philosophy” into modern day Service Oriented Architecture (SOA) and its Microservice variants [6]. These approaches evolve from the concepts of Component Based Software Engineering [9] (CBSE) where computational and compositional elements are explicitly separated [3], [7], [15].

We build upon CBSE and SOA concepts, investigating modularity at the level of stand-alone software systems and not at the level of modules or routines of a single system. In this context, we underline the benefits of CBSE-like concepts when applied to workflow systems for scientific computing executed on HPC resources. AirFlow, Oozie, Azkaban, Spark Streaming, Storm, or Kafka are examples of tools that have a design consistent with the proposed approach.

## III. BUILDING BLOCKS APPROACH

Each building block has a set of entities, a set of functionalities that operate on these entities, and a set of states, events and errors for each entity. Architecturally, the building blocks design requires: (i) a well-defined and stable interface for input and output that enables clean separation between computational and compositional features; (ii) one or more conversion layers capable of translating across diverse representations of

the same type of entity; (iii) one or more modules that implement the functionalities to operate on these entities and expose higher-level abstractions for their composition.

In our adaptation, the building blocks approach is based on four design principles: self-sufficiency, interoperability, composability, and extensibility. Self-sufficiency and interoperability depend upon the choice of both entities and functionalities. Entities have to be general enough so that specific instances of that type of entity can be reduced to a unique abstract representation. Accordingly, the scope of the functionalities of each building block has to be limited exclusively to its entities. In this way, interfaces can be designed to receive and send diverse codifications of the same type of entity, while functionalities can be codified to consistently translate those representations and operate on them.

Composability depends on whether the interfaces of each building block enables communication and coordination. Blocks communicate information about the states, events and errors of their entities, enabling the coordination of their functionalities. Due to the requirement of self-sufficiency, the coordination among blocks cannot be assumed to happen implicitly “by design”. Thus, coordination has to be codified on the base of an explicit model of the entities’ states. The sets of entities and functionalities of a block need to be extensible to enable the coordination among states of multiple and diverse blocks. Note that extensibility remains bound by both interoperability and self-sufficiency.

Each design principle of the building blocks approach poses unique challenges when applied to software systems used standalone and integrated with third-party systems. Choosing entities and scoping functionalities to enable self-sufficiency requires expanding the design phase and therefore longer development iterations. Further, interoperability requires system-level interfaces to become a first order concern and to be based on well defined, general purpose abstractions. The coordination protocols that enable composability require generalization of variable access, dataflow and procedure calls. Extensibility also requires shared coding convention and documentation.

The building blocks approach does not reinvent modularity, it applies it at system level to enable composability among independent software systems. As an abstraction, modularity enables separation of concerns by encapsulating discrete functions into semantic units exposed by means of a dedicated interface. As such, modularity can be used at function or method level, depending on the programming paradigm and the facilities offered by programming languages, or at system-level, depending on the interface exposed by each system.

Traditionally, components of software systems independently designed by third party organizations have been difficult to integrate outside the well-defined scope of an operating system like, for example, Unix. While interfaces can hide implementation details, working as implementation-independent specifications of capabilities, integration still requires semantic uniformity across interfaces. Obtaining such uniformity is challenging and largely unsupported by specific constructs both at specification and language level. Further, integrating

independent systems poses challenges in language heterogeneity, error handling, input/output validation, effective documentation and comprehensive testing.

Building blocks approach contributes to address integration challenges across systems by specifying state, event and error models for each block. Following best practices in API design, entities are explicitly specified and implemented in the block’s interface and used as input for each exposed functionality. Each entity has a set of associated states, events and errors. The order of the state is guaranteed by the implementation (e.g., a task cannot be executed before being scheduled and scheduled before being bound to a resource) while events are unordered but always contained within two defined states. Errors are always associated to an entity, state and event. Communication is decoupled from coordination and independent from the implementation of communication channels.

Even when applied at system level, modularity, and therefore the building blocks approach, presents at least two major trade offs. Building systems as blocks increases design and implementation effort, making unfeasible an unstructured but rapid development approach. While unstructured approaches are counter productive for long-term maintenance, short-term solutions would pay an unpractical overhead to the building blocks approach. Further, integration of systems that are independently developed imposes sharing responsibility of software reliability across multiple stakeholders. Often, this can be undesirable as user attributes all the responsibility to the stakeholder of their immediate interface. This problem can be mitigated by system-level fault tolerance but it remains an element to carefully evaluate when considering the building blocks approach.

#### IV. RADICAL CYBERTOOLS

RADICAL-Cybertools are software systems designed and implemented in accordance with the building blocks approach. Each system is independently designed with well-defined entities, functionalities, states, events and errors. Fig. 1 shows three existing RADICAL-Cybertools systems: RADICAL Ensemble Toolkit (hereafter simply referred as EnTK), RADICAL-Pilot and RADICAL-SAGA. RADICAL-WMS is a workload management system still under development.

Individual RADICAL-Cybertools are designed to be consistent with a four-layered view of distributed systems for the execution of scientific workloads and workflows on HPC resources. Each layer has a well-defined functionality and an associated “entity”. The entities are **workflows** (or applications) at the top layer and resource specific **jobs** at the bottom layer, with **workloads** and **tasks** as intervening transitional entities in the middle layers. The diagram of Fig. 2 provides a reference example for the integration among entities across layers that is independent of the specifics of applications, RADICAL-Cybertools and resources.

*Workflow and Application Description Level (L4):* Requirements and semantics of an application described in terms of a workflow.

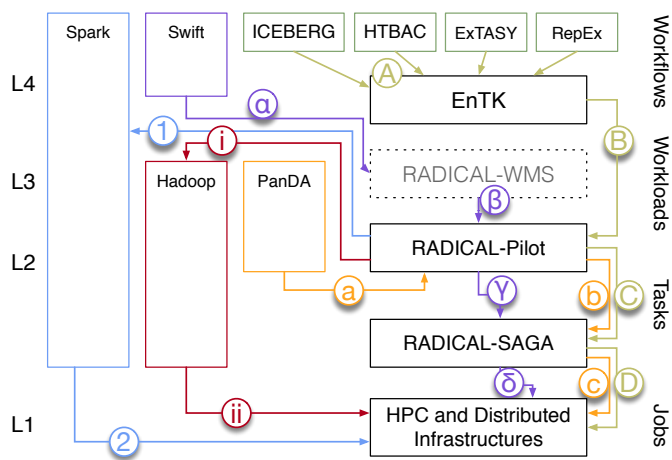


Fig. 1. Composition of RADICAL-Cybertools (black) with domain specific workflow systems (green, A–D), workflow system (purple,  $\alpha$ – $\delta$ ), workload management system (orange, a–c), framework for distributed data processing (red, i–ii), and a unified analytics engine (blue, 1–2). Numbered layers on the left; names of entities on the right. Solid colored lines indicate different integrations points with RADICAL-Cybertools; Dashed boxes indicate tools still under development.

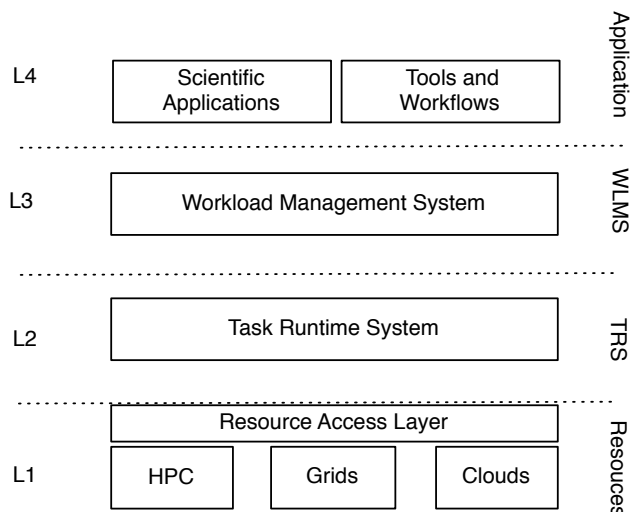


Fig. 2. Primary functional levels. The diagram supports an analysis of the functional requirements for workflow systems, and the primary entities at each level, agnostic of the applications and resources.

**Workload Management Level (L3):** Applications devoid of semantic context are expressed as workloads which are a set of tasks that can be executed concurrently. The Workload Management layer is responsible for: (i) the selection and configuration of available resources for the given workload; (ii) partitioning the workload over the selection of suitable resources; (iii) binding of tasks to resources.

**Task Execution Runtime Level (L2):** L3 delivers tasks to L2 which is responsible for their execution on the selected resources. L2 is a passive recipient of tasks from L3 but includes functionalities to acquire the indicated HPC resources,

schedule the given tasks over available resources, and execute these tasks with the indicated data and number of cores.

**Resource Layer (L1):** The resources used to execute tasks are characterized by their capabilities, availability and interfaces. Different resources present inconsistency in the way capabilities are provisioned but advances in syntactically uniform resource access layers enable task execution across resources.

Currently, in RADICAL-Cybertools each task defines an executable, e.g., `sleep`, `stress`, `python`, `GROMACS` or any other executable program. Each task description contains the arguments to pass to the executable, the type of parallelism required (e.g., MPI, OpenMP), the type and number of processing unit (CPU or GPU), the amount of memory, and the data staging requirements. RADICAL-Cybertools implement full task isolation, enabling the concurrent or sequential execution of heterogeneous, dependent or independent executables on a given set of acquired resources. RADICAL-Cybertools are agnostic of the operations performed by each executable: for each task, RADICAL-Cybertools satisfy its dependences, set up its environment and spawn its execution waiting for the executable to return in a final state. Therefore, RADICAL-Cybertools do not have access to the operations performed by each executable.

RADICAL-Cybertools conform to the principles of self-sufficiency, interoperability, composability and extensibility. EnTK exposes an API for the description of scientific applications as static or dynamic sets or sequences of pipelines. Pipelines are sequences of stages that, in turn, are sets of tasks. Sequences and sets formally define the relationship of priority among task executions: the tasks of a stage execute concurrently, tasks of different stages of the same pipeline execute sequentially, and pipelines execute concurrently. Resources are acquired and managed via a third-party runtime system that executes tasks on the acquired resources.

EnTK is self-sufficient as it enables necessary and sufficient functionalities for its set of entities, independently from third-party software systems. EnTK is interoperable because different representations of a workflow (e.g., DAG) can be converted to pipelines of stages of tasks, and because it is agnostic towards runtime systems and the type of resources on which they execute tasks. EnTK is also composable because it enables arbitrary coordination protocols (e.g., push/pull or master/worker) by explicitly defining the state model of its entities. Finally, EnTK is also extensible as new capabilities can be implemented for its entities, e.g., adaptivity of both workflows structure and task specifications at runtime.

RADICAL-Pilot is a pilot system that exposes an API to enable the acquisition of resources on which to schedule tasks for execution. The design of RADICAL-Pilot includes pilot and compute unit as entities. Capabilities are made available to describe, schedule, manage and execute entities. Pilots, units and their functionalities abstract the specificities of diverse types of resource, enabling the use of pilots mainly on single and multiple HPC machines, but also on HTC and cloud infrastructures. A pilot can span single or multiple compute nodes, resource pools, or virtual machines. Units of various size and duration

can be executed, supporting MPI and non-MPI executables, with a wide range of execution environment requirements.

The design of RADICAL-Pilot is: self-sufficient because, as with EnTK, it independently implements the necessary and sufficient set of functionalities for its entities; interoperable in terms of type of task, resource, and execution paradigm; and extensible as new properties can be added to the pilot, unit and resource descriptions, and more functionalities can be implemented for these entities. Currently, composability is partially designed and implemented: while the API can be used by both users and other systems to describe generic tasks for execution, RADICAL-Pilot requires RADICAL-SAGA to interface to HPC resources. A prototype interface to cloud resources based on LibCloud is available and a general-purpose resource connector component is under development.

RADICAL-SAGA exposes a homogeneous programming interface to the queuing systems of HPC resources. SAGA—an OGF standard—abstracts away the specificity of each queue system, offering a consistent representation of jobs and of the capabilities required to submit them to the resources. The design of RADICAL-SAGA is based on the job entity and its functionalities enable job submission and jobs’ requirements handling (self-sufficiency). Both entities and functionalities can be extended to support, for example, new queue systems or new type of jobs (extensibility). The SAGA API resolves the differences of each queue system into a general and sufficient representation (interoperability), exposing a stable set of capabilities to both users and/or other software elements (composability).

## V. BUILDING BLOCKS, RADICAL-CYBERTOOLS AND WORKFLOWS SYSTEMS

RADICAL-Cybertools as a whole are not an end-to-end workflow system. Each cybertool is an independent system that can also be integrated with other systems (RADICAL-Cybertools or otherwise) to form tailored middleware solutions. For example, several independent communities directly utilize RADICAL-SAGA alone, with RADICAL-Pilot or other pilot systems like, for example, PanDA Pilot. Other communities integrates all RADICAL-Cybertools with or without third-party systems to support the execution of diverse types of scientific workflows. Thus, RADICAL-Cybertools are not posed to replace existing workflow systems: RADICAL-Cybertools’ novelty is to enable the integration across systems independently developed and not necessarily designed to integrate. Crucially, this includes existing workflow, workload and computing frameworks, alongside their components.

We believe an ecosystem in which end-to-end workflow systems and building blocks coexist and, when useful, are integrated helps to avoid both lock in and fragmentation. Such an ecosystem would allow scientists with specific and stable requirements to use an end-to-end system while others to aggregate existing capabilities into tailored solutions. Inversely, a multitude of non-integrable systems built with slightly different capabilities fragments the user experience, forcing scien-

tists to learn to use multiple systems, depending on the context in which they have to operate.

As building blocks, RADICAL-Cybertools offer several benefits when used to describe and execute scientific workflows. Among these benefits, the most relevant is isolating scientists from job management (L1), task management (L2), and workload management (L3). These capabilities are further abstracted away in L4, letting scientists to exclusively focus on workflow description and application logic. Note that while this isolation is offered by other systems, RADICAL-Cybertools is agnostic towards what software tools and systems are integrated at each layer L1–4.

When integrated, RADICAL-Cybertools simplify the codification of workflows, lowering the barrier to adoption, maintenance and reuse. When using EnTK, workflows are codified as pipelines in a general purpose language (Python) and application-specific constructs (Task, Stage, Pipeline and App-Manager). As programs, workflows can be maintained following diverse approaches: from keeping a simple script on a scientist’s workstation to sharing a more complex application among multiple scientists via a collaborative version control system. Codifying workflows as code but without a dedicated domain language, offers the opportunity of reusing portion of code in the form of methods, classes and modules. Further, scientists have the option to grow the code as needed, typically starting from a small script and growing it into an application as the research advances, alone or with the help of other scientists and software engineers.

Building logically self-contained software blocks, lowering the technical barriers to their composability while enabling their interoperability and extensibility allow designing workflows as domain-specific applications, developed to solve classes of scientific problems, not issues of resource and execution management. These applications, alongside the blocks they use, become sustainable because they can be understood and maintained by diverse, invested communities. This is the sustainability model of successful open source software, including some of the existing solutions for certain types of workflows and resources, e.g., the Apache Hadoop ecosystem.

Supporting the development and maintenance of domain-specific applications is becoming increasingly important to enable scientific workflows. Alongside large communities in which the same workflow is used for many years (e.g., the LHC community), many research fields increasingly require running rapidly-evolving workflows with relatively short computation campaigns. These workflows depend on simulations and analysis procedures that evolve during the campaign, integrating new models and methodologies. As such they require a software ecosystem with independent systems that can be easily integrated and extended, depending on evolving scientific requirements.

### A. Integrating End-to-end Workflow Systems

As an example of how the building blocks approach can be utilized in other systems, we map the primary functional levels described in Fig. 2 to Pegasus [5], one of the most

adopted end-to-end workflow system. Scientific applications are described as abstract workflows using the HubZero API, or workflow composition tools such as Wings and Airvata. These interfaces correspond to the application layer (L4) of Fig. 2.

The abstract workflow is transformed to a concrete workflow by the Mapper component. The transformation takes into account the availability of software, data, and computational resources required for execution, and can restructure the workflow to optimize performance. A concrete workflow with several interdependent jobs, each consisting of several interdependent tasks, is passed to a workflow engine. Pegasus utilizes different engines, depending on the target resource: (1) lightweight execution engine for local resources; (2) HTCondor DAGMan and HTCondor Schedd for clusters and HPC machines; and (3) HTCondor with Glide-in WMS for grids. Functionally, Mapper, workflow engine, and local scheduler correspond together to the workload management layer (L3) of Fig. 2.

Pegasus supports three modes of job execution, depending on the execution environment and architecture of the remote machine: (1) PegasusCluster, a single-threaded engine that submits one task at a time; (2) PegasusLite, for handling tasks input and output data on resources with no shared filesystem; and (3) Pegasus MPICluster, for systems with a shared filesystem where MPI is used to implement a master-slave layout for task binding and execution. Collectively, these three remote execution engines correspond to the task runtime layer (L2) of Fig. 2.

Pegasus uses GlobusGRAM and CREAM-CE to submit jobs directly to remote resource managers such as SLURM, PBS, LSF, and SGE. These tools correspond to the resource access layer (L1) of Fig. 2.

Following this mapping, end-to-end workflows systems or some of their component can be integrated with RADICAL-Cybertools as building blocks to enable new capabilities. We used this approach to integrate Swift [19] with RADICAL-Pilot and RADICAL-SAGA. Swift has a long development history, with several versions that supported diverse case studies. Swift also integrated pilot systems of which Coasters is actively supported. The design of Swift is modular and it relies on connectors to interface with third-party systems.

In Swift, the language interpreter and the workflow engine are tightly coupled but connectors can be developed to stream the tasks of workflows to other systems for their execution. As seen in Sec. IV, RADICAL-Pilot can get streams of tasks as an input and submit these tasks to pilots for execution.

We integrated Swift with RADICAL-Pilot to enable the distributed and concurrent execution of Swift workflows on multiple HPC machines and HTC infrastructures (Fig. 1, purple  $\alpha$ - $\delta$ ). The distributed scheduling capabilities of RADICAL-Pilot offered the possibility to minimize the time to completion of tasks execution, obtaining both qualitative and quantitative improvements [16]. Qualitatively, RADICAL-Cybertools enabled Swift to execute workflows concurrently on both HPC and HTC resources via late binding of both tasks to pilots

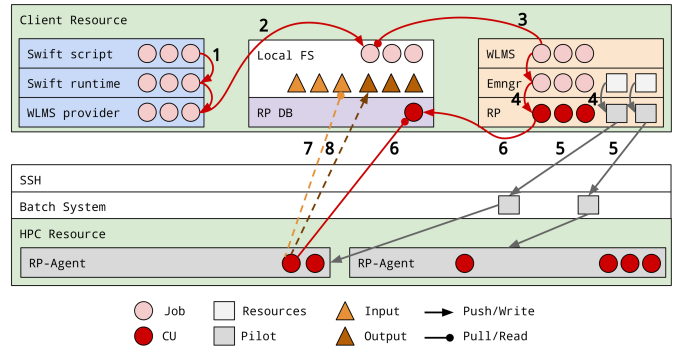


Fig. 3. Integration between Swift and RADICAL-Pilot. The two systems exchange task descriptions via a local filesystem. RADICAL-WMS derives the size and duration of the pilots from the task requirements, independently from Swift.

and pilots to resources. Quantitatively, the time to completion of workflows was improved by leveraging the shortest queue time among all the target resources.

The integration with RADICAL-Pilot required the development of a dedicated connector for Swift by iterating on the already available shell connector. We used the opportunity to prototype a distributed workload management (RADICAL-WMS) as a means of research. The connector enabled saving task descriptions on the local filesystem from where RADICAL-WMS was able to load and parse these descriptions without needing any added functionality. This type of integration was made possible by sharing the task entity semantics between the two systems and by isolating distinct functionalities operating on that entity in two distinct software systems. Note how these two systems were not designed to be integrated and were developed by independent teams.

### B. Integrating a Workload Management System

PanDA is a workload management system designed to support execution of independent tasks on Grid computing infrastructures like WLCG [12] and leadership-class HPC machines. The use of leadership HPC machines for executing large amounts of small jobs presents two main challenges: using a queue system that privileges large MPI jobs; and accessing untapped resources without disrupting the overall utilization of the machine. Pilots can address both challenges but pilot systems are difficult to deploy on HPCs. The main problem is efficiently managing the concurrent and sequential execution of small heterogeneous jobs at scale.

We developed a single-point solution to enable a workload management system designed for HTC to execute workflows on leadership-class HPC machines (Fig. 1, orange a-c). The PanDA team developed a job broker to support the execution of part of the ATLAS Monte Carlo workflow on Titan, while RADICAL-Pilot was used to enable pilot capabilities on Titan via an interface to RADICAL-Pilot we called Next Generation Executer (NGE).

PanDA Broker uses NGE to exchange information about tasks descriptions and resources requirements, while

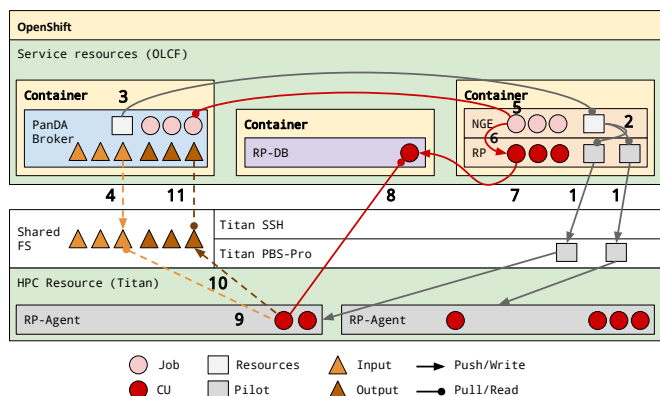


Fig. 4. Integration between PanDA and RADICAL-Pilot via the Next Generation Executor (NGE) REST interface. All systems execute on OLCF service resources within containers. Pilots are exposed to PanDA as an aggregation of available resources (steps 2 and 3).

RADICAL-Pilot behaves like a resource queue for PanDA Broker (Fig. 4). Both systems require no modifications to be integrated but the development of a coordination protocol to pull/push information about entities and their states. As with Swift, PanDA Broker and RADICAL-Pilot are independently developed and their integration was performed when the two stacks were already in production.

### C. Integrating General Purpose Computing Frameworks

Hadoop exposes an API for users and other software systems (composability) to describe distributed applications, mostly in terms of the MapReduce programming model, supporting distributed filesystem capabilities. Accordingly, Hadoop implements the necessary and sufficient functionalities of a workload management system (self-sufficiency). Hadoop can aggregate and manage diverse storage resources via a master/worker subsystem composed of multiple Namenode and DataNode instances (interoperability), and supports diverse runtime systems like Mesos, YARN, and others, to schedule and execute tasks on computing resources (extensibility).

Spark can be considered as a self-sufficient implementation of a Workflow system. It enables necessary and sufficient workflow functionalities, for machine learning, iterative analytics and streaming, independent of the underlying Task Runtime System. Spark enables interoperability by supporting different execution engines, such as Hadoop, MPI and others. Spark exposes an API that can be used to develop distributed applications (composability). Spark can be extended to support different types of workflows.

As building blocks, Hadoop, Spark and RADICAL-Pilot can be integrated into dedicated frameworks for HPC machines [11] (Fig. 1 red i–ii and blue 1–2). The integration avoids the need for dedicated deployment and customizations of Hadoop and Spark while retaining the full functionalities of both systems on HPC resources. RADICAL-Pilot configures, starts and manages an Hadoop/Spark cluster, and then executes a user’s Hadoop/Spark application on that cluster. Once

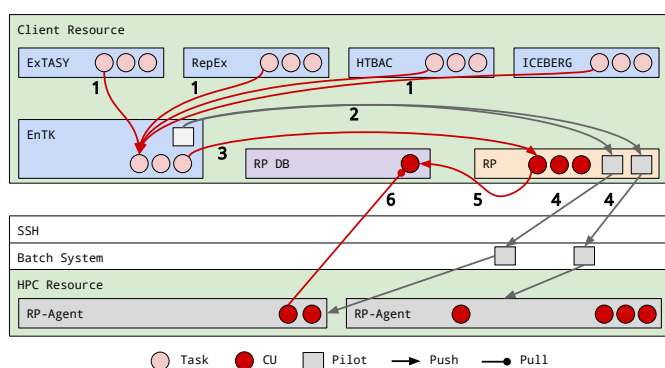


Fig. 5. Integration between four domain-specific workflow (DSW) systems—ExTASY, RepEx, HTBAC, ICEBERG—and EnTK. Numbers indicate the execution flow. RADICAL-Pilot (RP) database (DB) can be deployed on any host reachable from the resources.

done, RADICAL-Pilot shuts down the cluster and cleans up the environment.

We used RADICAL-Pilot’s integration with Spark to parallelize MDAnalysis and characterize its performance. MDAnalysis is a Python library that provides a comprehensive environment for filtering, transforming and analyzing molecular dynamics simulation trajectories [13]. Currently, we use RADICAL-Pilot’s integration with Spark to support diverse imagery analysis algorithms for geological and polar sciences.

### D. Domain Specific Workflow Systems

We call a workflow system that provides a specific higher-level functionality a Domain Specific Workflow system (DSW). We aggregated RADICAL-Cybertools, developing four DSW: ExTASY, RepEx, HTBAC and ICEBERG (Fig. 1, green A–D). Although, driven by specific application needs, each DSW is characterized by a unique execution and coordination pattern and can serve multiple applications.

Our DSW systems use EnTK to support ensemble-based workflows. EnTK is agnostic to the details of both the specific executables run by the ensemble members, and the system used to manage their execution. Fig. 5 shows how EnTK is coupled with RADICAL-Pilot to execute the ensembles via pilots on HPC resources. Note that, in principle, EnTK could use a different runtime system and type of infrastructure.

ExTASY [2] supports several advanced sampling methods (e.g., DM-d-MD and CoCo) in biomolecular simulations, using the EnTK API to provide the simulation-analysis execution pattern. RepEx enables multiple replica-exchange methods which vary in the coordination patterns across replicas, e.g., global synchronization barrier, or pair-wise synchronization etc. RepEx supports multi-dimensional exchange schemes, both synchronous and asynchronous [14], separating performance and functional layers while providing simple and easy methods to extend interfaces. Finally, HTBAC implements multiple pipelines of heterogeneous tasks, wherein both pipelines and tasks within a pipeline can change at runtime. HTBAC runs on several HPC machines, including ORNL’s and NCSA’s leadership-class machines. ICEBERG supports

scalable image analysis applications using multiple concurrent pipelines.

ExTASY, RepEx, HTBAC and ICEBERG benefit from integrating RADICAL-Cybertools by not having to reimplement workflow processing, efficient task management and interoperable task execution capabilities on distinct and heterogeneous platforms. This, in turn, enables both a focus on and ease of “last mile customization” for the DSW.

## VI. DISCUSSION AND CONCLUSIONS

Traditionally, assumptions about types of applications or resources have led to software systems that, while modular, have not allowed reuse outside their original requirements. We believe this is why functionalities pertaining to entities like tasks or pilots are often reimplemented. Each system serves well the single research group or the large scientific project but not each other.

As argued in Sec. III, software systems should be self-sufficient, interoperable, composable, and extensible so to be able to serve arbitrary requirements for a well-defined set of entities. For example, a workflow manager should provide methods for DAG traversing independent of how and when the DAG is specified or where the tasks of the workflow will be executed. Analogously, a pilot manager should provide multi-staging and task execution capabilities independent of the task scheduler or the compute resources on which tasks will be executed.

Modularity is not a design principle strong enough to realize this type of software systems. Modularity needs to be augmented by API and coordination agnosticism alongside an explicit understanding of the entities that define the domain of utilization of the software system. Each system developed following this approach, implements a well-defined set of functionalities specific to a set of entities, with minimal assumptions about the system that will use these functionalities or the environment in which they will be used.

Systems like Celery, Dask, Kafka, or Docker are early examples of software designed by implicitly following the proposed building blocks approach. These tools implement specific capabilities like queuing, scheduling, streaming, or virtualization for the domain of distributed computing. Consistently, they assume a set of core entities like workloads, tasks, pipelines, or messages, each with well-defined properties like concurrency and states. Their integration in multiple domains shows the potential of their underline design approach.

This paper offers four main contributions: (i) showing the relevance of the building blocks approach for supporting the workflows of various scientific domains on HPC machines; (ii) illustrating building blocks that enable multiple points of integration, resulting in design flexibility and functional extensibility, and providing a level of “unification” in the conceptual reasoning (e.g., execution paradigm) across otherwise different tools and systems; and (iii) showing how these building blocks have been used to develop and integrate workflow systems for HPC machines.

Sec. V highlights the practical impact of the building blocks approach. All the integrations required minimal development, mainly focused on translation layers and glue interfaces. Importantly, no refactoring was required within the systems we integrated. Explicit and agreed upon engineering processes was necessary to enable the integration among systems developed by independent teams and institutions. GitHub proved to be fundamental to enable these processes and to manage the engineering process. Explicit agreement on written use case and software requirements specifications greatly increase development coordination and, ultimately, efficiency. Lastly, weekly meeting among the lead developers helped the coding process and establishing a shared development culture.

It is important to outline what this paper does not attempt to achieve. This paper presents a preliminary study focused on one approach to building blocks for workflows systems, without a quantitative analysis of its benefits. Our work also does not attempt to distinguish (or identify) either the set of applications or systems where a building blocks approach will surpass alternative approaches. Finally, our paper does not analyze the wider implications for the middleware ecosystem for scientific computing. Although preliminary, this work is not premature: Conceptual formalisms that are too far ahead of proof-of-concepts and demonstrable advantages are unlikely to yield practical advances. Thus, even though the building blocks approach is still a work in progress, we believe early reports of success are necessary.

The building blocks approach spawns many new questions. A prominent one pertains to the issue of how we might model workflows systems and tools, so as to provide a common vocabulary, reasoning and comparative framework. Ref. [17] provided the architectural paradigm for pilot systems, however it is still unclear how an analogous paradigm would complement the work done on reference architectures for workflow systems [10], [8], and whether, given the very broad diversity of workflow systems and tools, we can even formulate a single architectural paradigm. This paradigm has been elusive so far, but it might be more fruitful to formulate system-level paradigms that have the properties of building blocks.

An end-goal and intended outcome of this paper is to begin a discussion on how the scientific workflows community—end-users, workflow designers and workflow systems developers—can better coordinate, cooperate, and reduce redundant and unsustainable efforts. We believe the building blocks approach contributes towards an examination and investigation of design principles and architectural patterns for workflow systems that may facilitate this discussion.

## VII. ACKNOWLEDGMENTS

We thank our collaborators: Peter Coveney and Dave Wright (UCL/HTBAC); Pavlo Svirin, Ruslan Mashnitov, Danila Oleynik, Kaushik De, Jack Wells and Alexei Klimentov (ATLAS Project/PanDA); Charlie Laughton and Cecilia Clementi (Nottingham/Rice/ExTASY); Srinivas Mushnoori. We thank Daniel Smith, Levi Naden and Sam Ellis (MolSSI) for useful discussions and insight. This was

was supported primarily by NSF 1440677 and DOE ASCR DE-SC0016280. We acknowledge access to computational facilities: XSEDE resources (TG-MCB090174) and Blue Waters (NSF-1713749).

[20] Z. Zhang, K. Barbary, F. A. Nothaft, E. Sparks, O. Zahn, M. J. Franklin, D. A. Patterson, and S. Perlmutter. Scientific computing meets big data technology: An astronomy use case. In *Big Data (Big Data), 2015 IEEE International Conference on*, pages 918–927. IEEE, 2015.

## REFERENCES

- [1] Computational Data Analysis Workflow Systems. <https://s.apache.org/existing-workflow-systems>.
- [2] V. Balasubramanian, I. Bethune, A. Shkurti, E. Breitmoser, E. Hruska, C. Clementi, C. Laughton, and S. Jha. Extasy: Scalable and flexible coupling of md simulations and advanced sampling techniques, 2016. <https://arxiv.org/abs/1606.00093>.
- [3] D. Batory and S. O'Malley. The design and implementation of hierarchical software systems with reusable components. *ACM Trans. Softw. Eng. Methodol.*, 1(4):355–398, 1992.
- [4] N. Chaimov, A. Malony, S. Canon, C. Iancu, K. Z. Ibrahim, and J. Srinivasan. Scaling spark on hpc systems. In *Proceedings of the 25th ACM International Symposium on High-Performance Parallel and Distributed Computing*, pages 97–110. ACM, 2016.
- [5] E. Deelman, G. Singh, M.-H. Su, J. Blythe, Y. Gil, C. Kesselman, G. Mehta, K. Vahi, G. B. Berriman, J. Good, A. Laity, J. C. Jacob, and D. S. Katz. Pegasus: A framework for mapping complex scientific workflows onto distributed systems. *Scientific Programming*, 13(3):219–237, 2005.
- [6] N. Dragoni, S. Giallorenzo, A. L. Lafuente, M. Mazzara, F. Montesi, R. Mustafin, and L. Safina. Microservices: yesterday, today, and tomorrow. In *Present and Ulterior Software Engineering*, pages 195–216. Springer, 2017.
- [7] D. Garlan, R. Allen, and J. Ockerbloom. Architectural mismatch or why it's hard to build systems out of existing parts. In *Proceedings of the 17th International Conference on Software Engineering, ICSE '95*, pages 179–185, 1995.
- [8] P. Grefen and R. R. de Vries. A reference architecture for workflow management systems. *Data & Knowledge Engineering*, 27(1):31–57, 1998.
- [9] G. T. Heineman and W. T. Council. Component-based software engineering. *Putting the pieces together, addison-westley*, page 5, 2001.
- [10] C. Lin, S. Lu, X. Fei, A. Chebotko, D. Pai, Z. Lai, F. Fotouhi, and J. Hua. A reference architecture for scientific workflow management systems and the view soa solution. *IEEE Transactions on Services Computing*, 2(1):79–92, 2009.
- [11] A. Luckow, I. Paraskevagos, G. Chantzialexiou, and S. Jha. Hadoop on hpc: Integrating hadoop and pilot-based dynamic resource management. *2016 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 1607–1616, 2016.
- [12] T. Maeno. Panda: distributed production and distributed analysis system for atlas. In *Journal of Physics: Conference Series*, volume 119, page 062036. IOP Publishing, 2008.
- [13] N. Michaud-Agrawal, E. J. Denning, T. B. Woolf, and O. Beckstein. Mdanalysis: A toolkit for the analysis of molecular dynamics simulations. *Journal of Computational Chemistry*, 32(10):2319–2327, 2011.
- [14] B. K. Radak, M. Romanus, T.-S. Lee, H. Chen, M. Huang, A. Treikalis, V. Balasubramanian, S. Jha, and D. M. York. Characterization of the three-dimensional free energy manifold for the Uracil Ribonucleoside from asynchronous replica exchange simulations. *Journal of Chemical Theory and Computation*, 11(2):373–377, 2015.
- [15] J.-G. Schneider and O. Nierstrasz. Components, scripts and glue. In *Software Architectures*, pages 13–25. Springer, 2000.
- [16] M. Turilli, Y. N. Babuji, A. Merzky, M. T. Ha, M. Wilde, D. S. Katz, and S. Jha. Evaluating distributed execution of workloads. In *e-Science (e-Science), 2017 IEEE 13th International Conference on*, pages 276–285. IEEE, 2017.
- [17] M. Turilli, M. Santcroos, and S. Jha. A comprehensive perspective on pilot-job systems. *ACM Computing Surveys (CSUR)*, 51(2):43, 2018.
- [18] J. Vivian, A. A. Rao, F. A. Nothaft, C. Ketchum, J. Armstrong, A. Novak, J. Pfeil, J. Narkizian, A. D. Deran, A. Musselman-Brown, et al. Toil enables reproducible, open source, big biomedical data analyses. *Nature biotechnology*, 35(4):314, 2017.
- [19] M. Wilde, M. Hategan, J. M. Wozniak, B. Clifford, D. S. Katz, and I. Foster. Swift: A language for distributed parallel scripting. *Parallel Computing*, 37(9):633–652, 2011.