Building a unified grid, Part 4: Managing the workflow

Level: Intermediate

Abel W. Lin (awlin@ncmir.ucsd.edu), Software Architect, National Center for Microscopy and Imaging Research

02 May 2006

The "Building a unified grid" series has highlighted the component implementations needed to grid-enable an end-to-end process. Part 1 discusses what it means to grid-enable an end-to-end process and described the grid-based system architecture developed in the Telescience Project at the National Center for Microscopy and Imaging Research. Part 2 shows why the grid interface is critical to the end-to-end process and how a unification of resources in an on-demand fashion can collect better data faster. Part 3 describes the unified authentication and authorization bundle of the Telescience ATOMIC package. Now, we discuss the use of workflows within the Telescience architecture.

Why workflow? What are workflows?

Use models for first-generation grids and their supercomputing-center predecessors were akin to the hub-and-spoke model used by the airline industry. User data environments were treated as the *hub*. At every step, the user was required to log in, and data was passed across the country to one of the few virtual organizations, or *spokes*, to execute their jobs (see Figure 1).

As mentioned in Part 1, data in an on-demand world is automatically maintained and flows freely from instrument to computation to analysis. The results of that analysis interface directly to the instrument without going through a central hub, providing automated feedback that constantly refines data-collection parameters and techniques. Faced with these requirements, grid applications are no longer built as monolithic all-encompassing programs, but rather are built from often pre-existing program components. Workflow tools are generally defined as the tools necessary to manage the activities (for example, authentication and data I/O) between program components. The use of workflow tools enables a point-to-point process necessary for an on-demand unified grid.

Figure 1. Hub-and-spoke vs. point-to-point processes (solid lines indicate data transfer; hashed lines indicate user-monitoring tools)

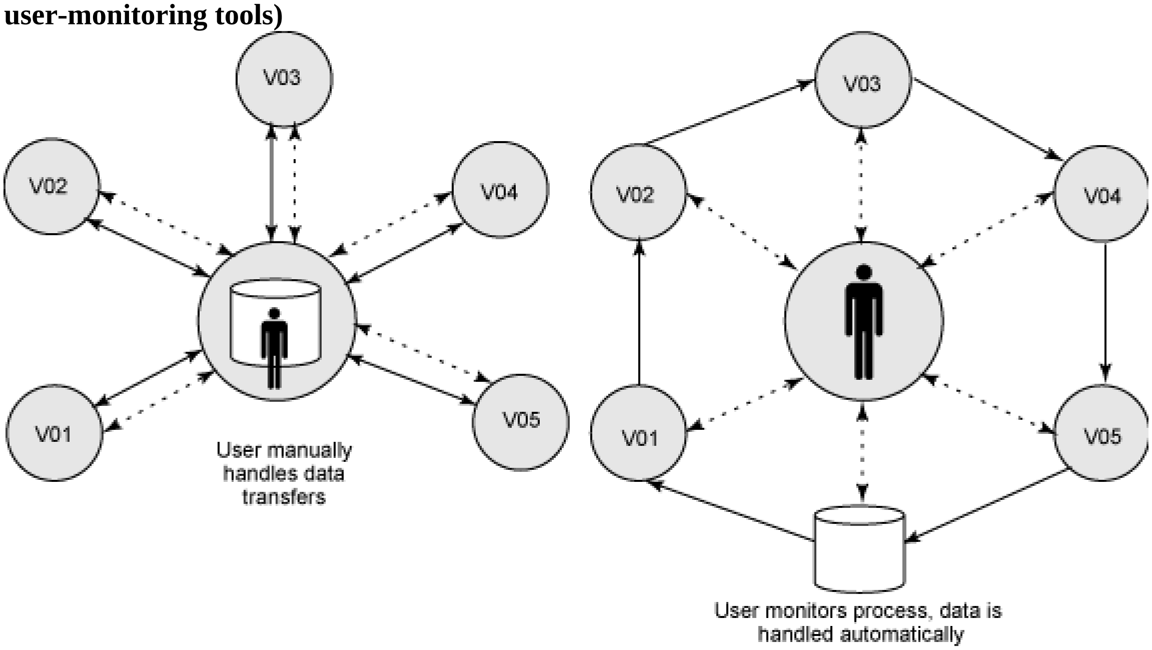

The challenges of using workflows

Like the core grid tools (now bundled by Application to Middleware Interaction Component (ATOMIC) into thematic Web services API), there are numerous workflow tools in the Telescience Project that provide complementary, yet overlapping, functionality. As workflow tools are only now beginning to gain adoption, they are in even more flux than the core grid middleware components.

Those workflow tools fall into the following hierarchal classes:

Workflow tools

Type of workflow	Description
Process management	Frames the highest-level scientific (laboratory) processes and provides policy, process, state management, and administrative tools that coordinate and manage lower-level workflows and pipelines that comprise a scientific study (or an instance within that study)
Interapplication	Pipeline or plan-building tools that streamline computational operations
Intra-application	Planners and execution engines that optimize the execution of these plans on heterogeneous physical resources

The Telescience workflow approach is to facilitate coordination and sharing of stateful information between these workflow layers. Each layer has unique abilities and requirements. Process and state management tools (typically portal-based) are necessary to preserve and delegate the contextual information with regard to the user. This information includes process management, management of the scientific process, authentication and authorization, and high-level state information. Within the Telescience Project, much of this information (authentication and authorization, for example) is delivered to the lower-level workflows via the ATOMIC tool kit API (see Figure 1 in Part 3).

Interapplication tools create process pipelines, which are subcomponents of the highest-level experimental process-management workflow. These tools are typically user-driven GUI environments that are ordered within the process-management workflow or presented as a general tool to serve the process-management workflow as needed. The lowest-level intra-application workflows are composed of the executable plans that have been mapped to heterogeneous pools of physical resources. In the context of the Telescience Project, we describe the use of portals as *workflow controllers* and the use of the intra-application class workflow tools to bring scientific algorithms to the grid.

Revisiting portals as workflow controllers

The end-to-end process a scientist embarks upon is defined as all the steps between the conception of an experiment and the final discoveries made based on experimental findings, including initial planning, information gathering, data collection, analysis, and potentially many iterations of this process at one or many decision points. The laboratory process isn't simply a linear stove-pipe process; it's a dynamic and highly iterative process with multiple points of user interaction, data visualization, and feedback. Enabling this feedback is a goal of a unified grid.

In the context of workflows, the laboratory process is the first-order workflow in the hierarchy of workflow tools, and it's the first workflow level that directly interacts with the end user. As described in Part 2, the Telescience infrastructure allows the team to move beyond interfaces with singular actions and integrate them into a richer user environment that is automated and dictated by the process and not by the grid middleware (see Figure 2 in Part 2). In this role, the portal is used more for process workflow management, where more emphasis is placed on managing state and persistence information of the components. Less emphasis is placed on the mechanics of launching application components than first-generation grid portal implementations for hub-and-spoke processing. While not traditionally thought of as workflow tools, the team has found portals to be critical to applications and workflow information delivery.

Intra-application workflows

The goal of the portal workflow controller is to coordinate the activities of the applications that "do the work." These applications might be stand-alone applications that simply require access to grid-managed data, or they might be complex parallel applications that use myriad heterogeneous computational or data resources.

Unlike first-generation grid codes that were large and monolithic, these component modules are small and dynamic. And unlike early codes, which tended to be "pleasantly parallel," these modern codes are heterogeneously parallel, often requiring more than one precursor component to be completed before computation can begin. This parallel heterogeneity, mixed with resource heterogeneity, requires sophisticated workflow planning and execution tools to abstractly plan and execute the workflow. As these requirements are shared across many scientific communities, several tools exist.

The Planning for Execution in Grids (Pegasus) environment is the workflow planning tool primarily used within the Telescience Project. Pegasus is a framework that maps complex scientific workflows onto distributed resources, such as the grid. Pegasus maps an abstract workflow description to its executable form, and Condor Directed Acyclic Graph Manager (DAGMan) executes the jobs specified in the executable workflow. Pegasus and DAGMan map and execute workflows on Condor pools, clusters managed by LSF or PBS, TeraGrid hosts, and individual hosts.

Pegasus operates on abstract workflow descriptions where the analysis is described in terms of application components and the data that the components use. The workflow is abstract because it does not identify the resources necessary for execution. Pegasus takes this abstract workflow description and produces an executable workflow that identifies the compute resources needed and includes data management nodes that stage the data in and out of the computations. Additional workflow nodes are added to register the newly derived data products so they can be located at a later time.

Bringing it all together

In the Telescience model, the portal environment is required to curate all the pertinent information regarding the user and session state required by lower-level workflow classes. While the portal serves as the curator of user information and state, ATOMIC serves as the delivery vehicle, providing downstream applications and workflows with access to the appropriate information necessary for a given process. This abstraction of the session information is necessary for maintaining a seamless user environment during the transition between workflow classes, and it scales to the needs of future workflow technology developments.

Figure 2 shows a high-level outline of a typical multiclass workflow initiated by the end user. From the main scientific process workflow-controller portlet, the user launches an external application (in this example, a Telemicroscopy control session). Session information curated by the portal upon login (authentication and data management parameters, for example) is passed to the application at runtime via ATOMIC tools and services. Using those parameters, the application initiates a lower-class workflow (in this example, a Pegasus-planned workflow for parallel tomographic volume reconstruction executed on a heterogeneous collection of computational resources).

Next-generation ATOMIC Web/grid service-based implementation will further allow dynamic notifications of progress at the external application level and at the main portal workflow. All of this takes place in a seamless user environment where the typical overhead of transitioning between workflow classes is passed neither to the end user nor to the application developer. For example, we anticipate the inclusion of more robust resource and network discovery tools within Pegasus without modification of current applications.

Figure 2. The Telescience Portal

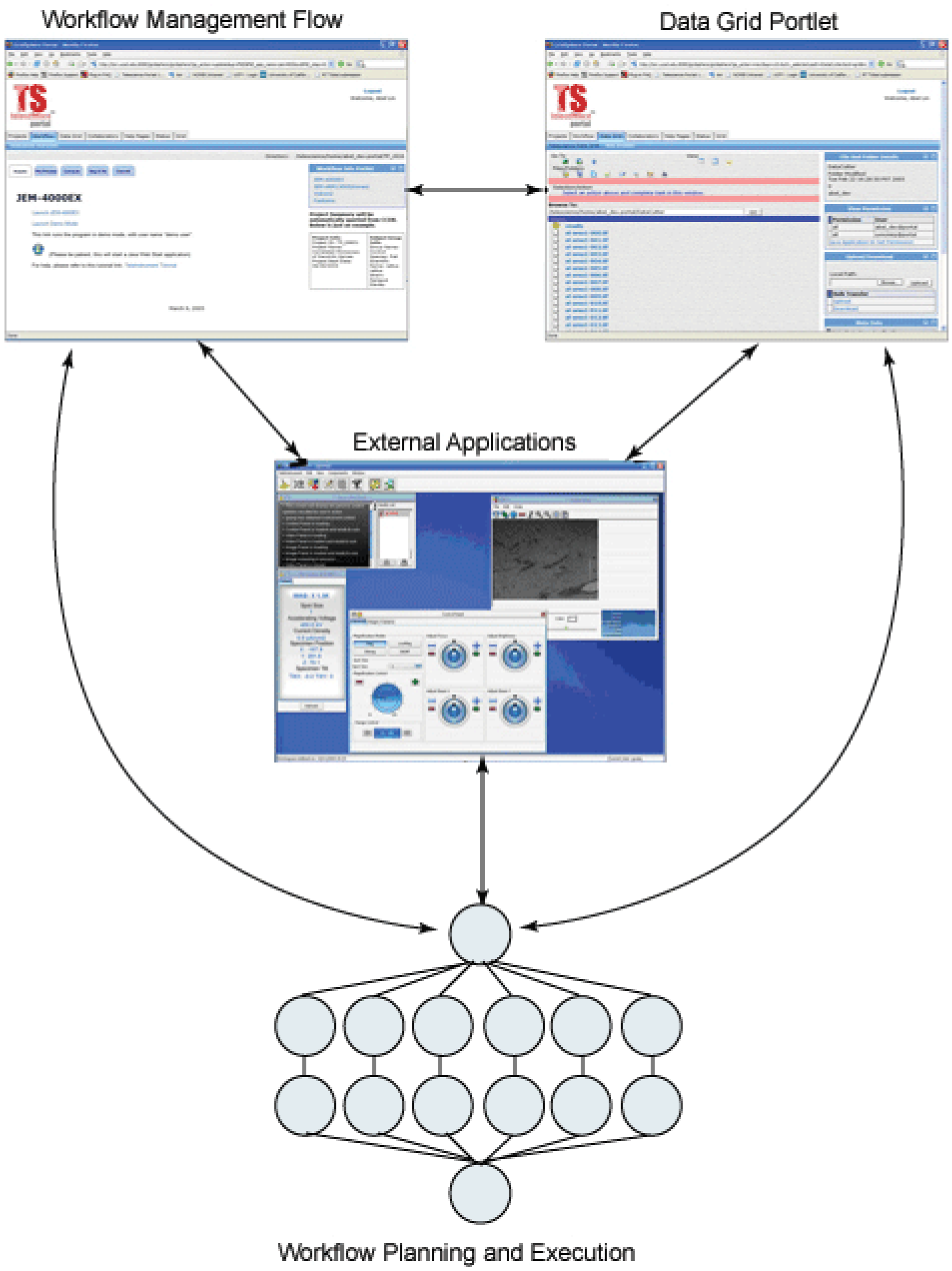

Conclusion

If we define an on-demand unified grid to encompass all the steps -- including feedback loops -- that take place between data or information acquisition and the final discoveries that take place as a result of the initial data acquisition, it's clear that no single grid middleware or workflow tool is adequate to address this need. It can be addressed only through the integration of several interoperable tools.

In this series, we've demonstrated this integration though the bundling of core grid components (ATOMIC) and the hierarchal use of workflow tools within the larger architecture. In the next (and final) article of this series, we show specific application examples of the Telescience Project using the Telescience architecture to enable an on-demand unified grid, and we'll discuss future development and challenges.

Resources

Learn

- Learn more about the Telescience Project.
- Learn more about research at the National Center for Microscopy and Imaging Research (NCMIR).
- Check out the Pegasus Project.
- Learn more about the various grid software components at the NSF Middleware Initiative.
- Learn more about workflows at the Workflow Management Research Group at the Global Grid Forum.
- Visit the developerWorks Grid computing zone for extensive information to help you develop with grid computing technologies.

Get products and technologies

- Download a beta version of the Telescience ATOMIC Toolkit.

Discuss

- Get involved in the developerWorks community by participating in developerWorks blogs.

About the author

Abel W. Lin is the architect and technical lead for the Telescience Project at the National Center for Microscopy and Imaging Research. He has more than five years' experience applying grid and other computer science technologies to scientific processes. He designed and led the implementation of the first-generation proof-of-concept systems for the Telescience Portal and ATOMIC components of the Telescience Project, and is an interdisciplinary scientist with published research in biology and computer science. His interests include distributed systems architecture, software project management, and structural biology. Outside of the office, he is an avid reader, golfer, and bodysurfer.
