NOAO SDM Document PL023

The Marriage of Mario (NHPPS) and Luigi (OGCE)

F. Valdes[1], S. Marru[2]

National Optical Astronomy Observatories
Science Data Management

December 15, 2010

[1]NOAO Data Products Program, P.O. Box 26732, Tucson, AZ 85732
[2]Pervasive Technology Institue, Indiana University, Bloomington, In

Table of Contents

1 Introduction **2**

2 NOAO High Performance Pipeline System (NHPPS) **2**

3 Open Grid Computing Environment (OGCE) **3**

4 The Marriage **3**

5 Conclusion **4**

Abstract

Pipeline systems, like the plumbing characters of video game fame, have their strengths and weaknesses. This is generally due to the different niches they occupy in the landscape of pipeline systems. This paper describes a marriage between two pipeline systems, NHPPS and OGCE, which capitalizes on combining their strengths.

People frequently discuss interchange of applications elements between pipeline systems. In other words, extracting some pipeline pieces (modules and algorithms) orchestrated by one pipeline system for orchestration by another. However, pipeline systems themselves are applications in their own right. This paper describes an interesting architectural marriage where one pipeline system orchestrates another pipeline system to take advantage of the strengths and niches of both. We point out this concept in general terms since there may be other pipeline marriages possible, but describe a specific example that forms the basis of a major new pipeline application for the WIYN One-Degree Imager (ODI).

The NOAO High Performance Pipeline System (NHPPS), in use for many years at NOAO, is good at orchestrating legacy software, IRAF specifically, into multi-process workflows that effectively make use of a dedicated cluster of multi-core machines without needing to write explicit threaded applications. The Open Grid Computing Environment (OGCE) is designed to wrap and manage applications into Teragrid workflows. For the ODI project we will wrap single-node NHPPS pipelines into OGCE services for orchestration into distributed Teragrid workflows. The many advantages of this married architecture will be described.

Keywords: pipelines, NHPPS, OGCE

1 Introduction

The WIYN Observatory is building a forefront mosaic camera with a one-degree field of view, hence the name One-Degree Imager (ODI), for its 3.5 meter telescope. The imager provides a high resolution of 0.1 arcsec per pixel which makes it a gigapixel camera. See Jacoby [REF] for details.

The size of the format means observers will obtain a large volume of data which most will be unable to store, calibrate, and analyze with their own resources. Therefore, Indiana University (IU) and NOAO, WIYN partner institutions, have partnered to develop a Pipeline, Portal, and Archive (PPA). In this contribution we describe the plans and synergies within this partnership for a standard calibration pipeline known as the Tier 1 pipeline. This pipeline will be a marriage of two pipeline or workflow system,

Before describing this marriage we mention a few things about the IU/NOAO partnership. The collaborating groups are the Pervasive Technology Institute (PTI) at IU and the Science Data Management (SDM) group at NOAO. PTI is an information technology group while SDM is a traditional astronomical data management group. The differences in expertise between these groups is one of the features of the partnership. SDM brings experience with pipeline processing of astronomical image data similar to ODI and managing a PI and community archive. PTI brings experience with the IT problems of data transport, storage, computations on a larger scale as well multidisciplenary science gateways. The blending of expertise is needed to address the larger scale of ODI data than other current NOAO instruments with the astronomical needs of the ODI users.

For the Tier 1 pipeline the partnership will bring the experience of NOAO/SDM with production CCD pipelines for their mosaic optical and infrared cameras together with PTI's experience with the TeraGrid and helping enable applications on this resource through middleware and portals.

Though we don't have space to describe other aspects of the PPA project one additional and pipeline related goal is to provide higher level workflows, driven by the observers, on the standard data products. This concept is called Tier 2 and Tier 3 and will again make use of the data processing expertise of NOAO/SDM and the IT expertise of PTI.

So on to the Tier 1 pipeline. The marriage we describe is between the NOAO High Performance Pipeline System (NHPPS) and the Open Grid Computing Environment (OGCE). At NOAO the NHPPS is a product of the "Mario" project, named for the video game plumber. To complement this tongue-in-check description will associate OGCE with Luigi, another video game plumber. Ihe following sections we describe the personalities of Mario and Luigi and then how their strengths are combined in a complementary manner, a marriage, for an ODI Tier 1 pipeline.

2 NOAO High Performance Pipeline System (NHPPS)

The NOAO High Performance Pipeline System (NHPPS) has been described in previous ADASS contributions [REF] and in a more detailed paper [REF]. In summary, NHPPS is an event-driven executor of host commands (aka modules or stages) on dataset objects. The pipelines, a sequence of possibly parallel stages running on one type of dataset object, are described by an XML Pipeline Description Language (PDL).

Because any host command may be used a variety of possible workflow styles can be implemented. NOAO uses NHPPS to produce pipelines that are connected in a map/reduce style of computation. In other words pipelines decompose dataset objects into smaller dataset objects which can be processed in parallel. These datasets are sent to pipelines designed for those types of objects and the results are collected for further use.

NHPPS was designed to efficiently run as many datasets and stages in parallel as desired. This happens both on a single node and distributed across peer nodels in a dedicated cluster.

The framework executes any host command whether threaded or not. NOAO developed and uses this framework because it effectively with noan-threaded or MPI commands. In particular, NOAO capitalizes on the nunderds of man-years of work in IRAF [REF] so that most NOAO/SDM pipeline modules are IRAF-shell scripts using standard IRAF commands.

To understand the marriage a key point to understand is that the NHPPS framework is basically a server, called the Node Manager, which runs on each compute node and responds to events by executing modules on particular dataset objects. The efficiency occurs because it manages many processes at the same time so that operating system can keep the CPU cores which it operates busy.

When NHPPS runs as a distributed pipeline system the Node Manager communcate with each other and manage multiple instances of the same pipeline or different pipelines on different nodes depending on the configuration design.

3 Open Grid Computing Environment (OGCE)

The Open Grid Computing Environment (OGCE) is essentially described as tools to wrap host commands as web services which are deployed on the TeraGrid (or potentially any grid) to run a workflow application. There are a variety of components and aspects of OGCE which are described by XXXXX [REF].

A key point is that OGCE does not do any workflow orchestration on a node. So primary it launches a single host command on a grid node. What the command does may be multi-threaded, MPI, simple shell script, etc. So the utilization of the node resources resides in the application. It also does not facilitate any communcation between applications and nodes except for parameters (of various types).

4 The Marriage

The technical question for te NOAO/IU partners with respect to the Tier 1 pipeline was how NOAO developed science pipeline software would be integrated with the IU Grid workflow system. After technical exchanges and prototyping

The NOAO pipeline team began with a data flow design [REF]. From this design, which built on knowledge of the NOAO Mosiac Pipeline, the number of modules is estimated to be on the order of 150 and the number of logical pipelines to be on the order of 20.

NOAO desired to maintain the ability to run the Pipeline as an NHPPS pipeline on a dedicated cluster. The reasons for this are to allow a development and test pipeline for the developers with

which they have experience, to maximize the use of pipeline modules from the NOAO production pipelines, and to maintain a fallback operations mode. A consequence of this desire is that the modules require no major API and interface changes. One could achieve this with module wrappers but the number of wrappers would

From these considerations and the realization that the NHPPS framework could be encapsulated as part of a single host application, the architecture of wrapping entire NHPPS pipelines as single OGCE service applications was identified.

After this ”epiphany” the technical requirements and additional benefits became evident.

- the main requirement for the NHPPS framework that multiple instances can coexist on the same node by using independent socket addresses - the NHPPS pipelines as implemented by NOAO have a very regular structure an interface - NHPPS pipelines make effective use of a node’s resources by keeping a large number of modules active at the same time

The benefits of OGCE are the separation of the applications, the NHPPS pipelines, from the complexities of scheduling, job management, security, etc. of the grid.

- needs to provide the map/reduce capability in the workflow - needs to provide some basic conditional workflow constructs tied to the output of the modules

So here is the core of this contribution. The marriage of NHPPS with OGCE consists of wrapping NHPPS pipelines as single applications. Because these pipelines have a standard structure a single basic wrapper is all that is needed. This wrapper accepts input parameters and produces output parameters (which include status) that OGCE services expect. The wrapper also takes care of starting the NHPPS framework, supplying the triggers that start the pipeline, wait for the pipeline to complete, and shutdown the framework.

For OGCE, one web service wrapper is created for each of the pipelines. These web services are then orchestrated by OGCE as a Grid workflow. OGCE provides the workflow monitoring and collecting and managing of the pipeline input and output in addition to the typical submission and security aspects of a Grid application.

An organizational benefit of this design is that the interface, and resulting division of work, between NOAO and IU is simple an clear. NOAO can concentrate on the NHPPS pipelines knowing the wrapper strategy is easy and clear and IU can concentrate on the map/reduce workflow with little need to know about the modules and algorithms.

- only little compiled coding required; i.e. no special multithreaded or MPI development

which normally

5 Conclusion

The partnership of NOAO and IU/PTI for providing an ODI Tier 1 pipeline has many benefits. The architectural marriage described in this paper has a clean interface in software responsibililities. The NOAO developers do not need expertise in all the details of Grid applications and IU developers do not need expertise in the algorithms of the pipeline modules.

The other key beneit of the NHPPS/OGCE marriage is the ability to make efficient use of multicore Grid nodes using the large body of existing CCD reduction software in IRAF.