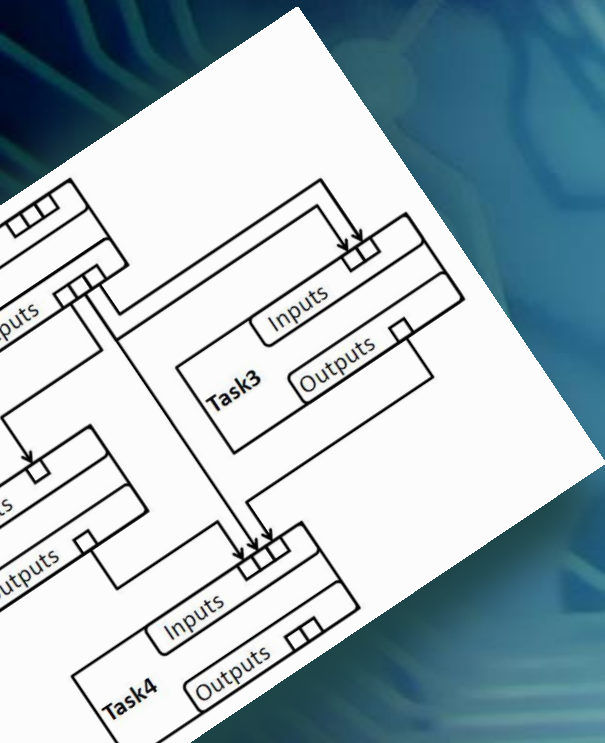
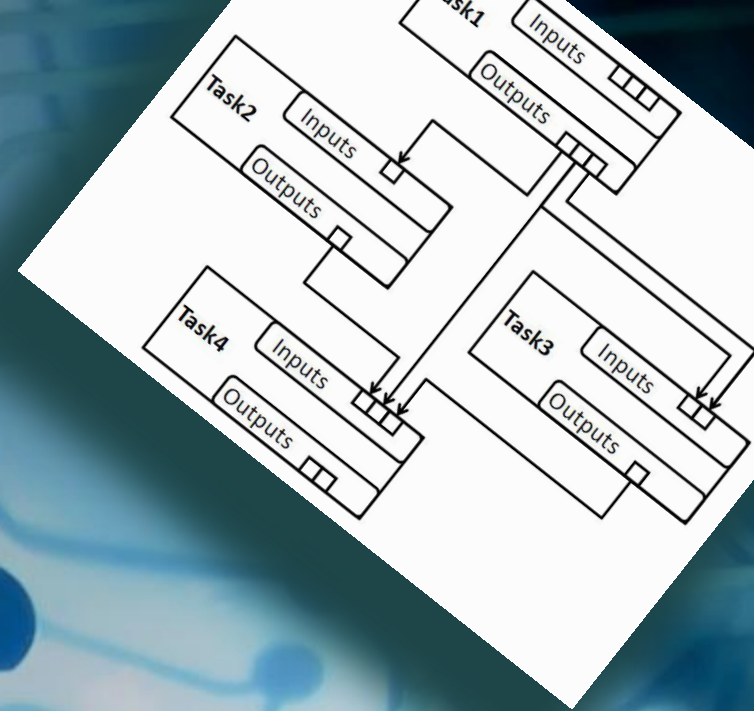


First Steps to Workload Management: Job Wrapper



Maxim Potekhin
BNL

LBNE Software Workshop
November 14, 2013

A note on WMS

What is a Workload Management System? Here's my own quick (and imperfect, but better than Wikipedia) definition based on experience in this area: a modern and capable system presents a combination of three major capabilities (we can also call them "benefits") –

- **Brokerage:** optimal routing of the payload to available resources (which may be transient and/or opportunistic) according to policies which may change with time.
- **Execution environment management:** the pilot can take care of a lot of setup steps that users otherwise have to keep in their own wrapper scripts; standard payloads can also be defined via a Web service, making their management more robust and transparent.
- **Monitoring:** comprehensive set of tools that allows the operator to keep track of the payload execution at a few abstraction levels (workflow, task, job) with desired level of detail – which is sometimes substantial. Sometimes described as "drill down" functionality. Quite often, easy access to log files and similar information is indispensable for efficient debugging and operations on the Grid.
- **Data Management:** Grid Tools that allow data distribution among computing centers according to a chosen policy, sometimes via a subscription mechanism, sometimes combined with dynamic distributed data access (e.g. via xrootd) and ties into the payload distribution based on data locality.

Things to consider

Do we need a WMS at all?

I believe that *eventually* we do. We need to hedge against future uncertainties of our computing power requirements combined with future uncertainties in what's likely to be available at a particular point in time. This includes LBNE users at remote computing centers who wish to process distributed data.

Do we need to select a WMS now, i.e. in two weeks?

No we do not. Our requirements are still evolving and so is the technology landscape. We do need it in foreseeable future, however.

What are the WMS features that we can achieve on a relatively short time scale and without commitment to a specific and complex solution?

Monitoring, accounting and data handling.

“Pilot jobs”

What is a Pilot?

Pilot's functions include validation of the environment, monitoring and logging for the payload job. In some cases, the Pilot can perform more complex functions such as interfacing data management system.

What is late binding?

The Pilot, when it starts, is not aware of its future payload (if it comes at all). The payload is deployed to validated slots, which insulates the end user from problematic sites and WNs.

What is the role of the Pilot job in monitoring?

The Pilot communicates with a network service (can be Condor, or a Web service typically Apache-based). This makes it possible to monitor the state and other parameters of a running job. It can also manage placement of the log files, error files and some other information that is useful in debugging and operations.

What is the role of the Pilot in data handling?

The Pilot is capable of interfacing distributed data management systems.

Can we use the Pilot without a WMS?

Yes we can, but let's just call it a wrapper. Or “execution environment manager”.

There is *no late binding* anymore, but some important functions can still be taken care of.

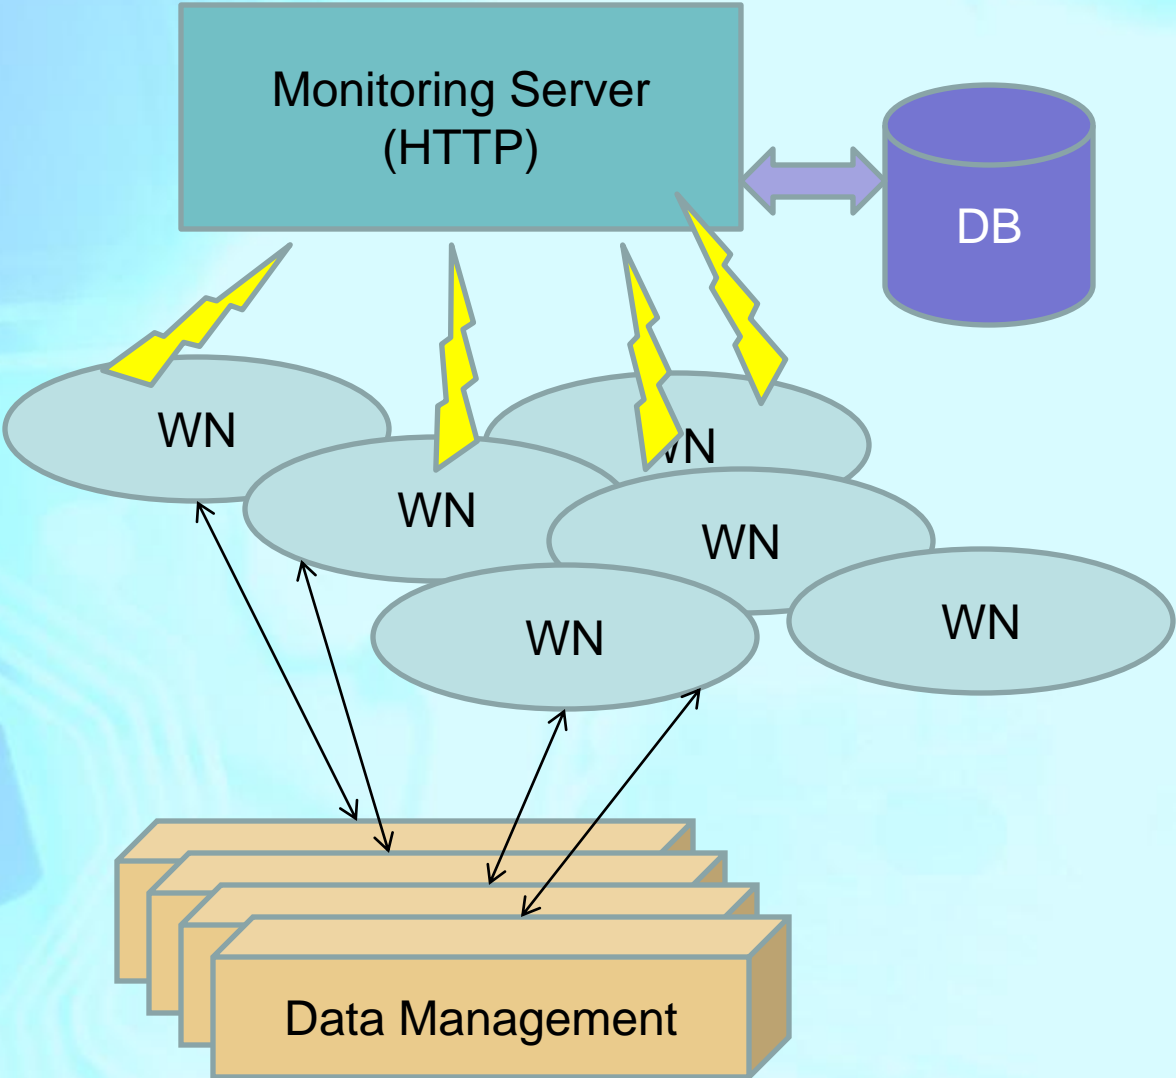
What are immediate benefits?

Monitoring – we shall be able to offer a nice and user friendly Web UI for monitoring functions

Accounting – we will be able to know who needs resources and how they are used, which is crucial for meaningful planning. Dependencies and environments can be also tracked.

Data management – there will be a way to run the payload without any coupling to, or awareness of the location and other attributes of the input and output data.

Conceptual Diagram



Let's consider this option.

Needs a thorough discussion with the Physics Tools Group.

Who is going to do it?

That's a fun project and it involves a monitoring UI. Should be possible to locate willing experts (myself included).

What do we need?

A machine, a DB but most importantly cooperation and lots of good will.

What is the first thing to do?

The monitoring capability.