

@GridPP

@twhyntie

Supporting a new VO: A case study

T. Whyntie*, †

** Queen Mary University of London; † Institute for Research in Schools*

#GridPP36, Pitlochry, Scotland

Wednesday 13th April 2016



Outline of the talk

- The New User Engagement Programme in GridPP4+;
- Some general observations;
- “First Contact” and the initial requirements capture;
- The case study – GalDyn;
- Summary and conclusions

The NUEP in GridPP4+ (Apr 15-Mar 16)

- **The New User Engagement Programme (NUEP):**
 - *Concerted effort by GridPP to make 10% of its resources available to non-LHC communities, including those from particle physics, physics, other sciences, other fields, and industry;*
 - *Response to talks at [GridPP32](#) and [GridPP34](#);*
- New user communities in this period – see Jeremy C’s talk:
 - *GalDyn, LSST, PRaVDA, LZ, LIGO, SuperNEMO, et al.*
- Suite of tools and software available – see the “Tools” talk:
 - *GridPP DIRAC, CernVM, CVMFS, Ganga, etc.*
- Also: the new [GridPP website](#) and [UserGuide](#).

Some general observations

- Tracking/monitoring progress:
 - *Monitoring problems is easy...*
 - *...monitoring success requires a surprising amount of effort!*
 - *“It’s all working but we haven’t had time to do anything with it yet.”*
- There is no “one-size-fits-all” solution – but should there be?
 - *All cases have used a different combination of tools and approaches.*
 - *More effort to support and document – but is this a strength of GridPP?*
 - *“..the main differentiator seems to be the level of hand-holding [GridPP] provide to users, specifically scientific and data analysis work loads. Any other cloud provider (Amazon, Microsoft et al.) could provide the same if not a greater compute footprint, but do not provide application specific support.”*

Feedback via Colin Hayhurst, Uni. Sussex Innovation Fellow

Some general observations

- We should define the different stages of engagement:
 - *First contact:*
 - One or two users testing grid technologies (on behalf of an organisation?);
 - Minimal infrastructure requirements: incubator VO, cluster/VM grid UI;
 - Copious hand-holding from a GridPP champion and GRIDPP-SUPPORT;
 - Aiming for “quick wins” that demonstrate that GridPP *is the way to go*.
 - *Development:*
 - Creation of a new VO (or integration/resurrection of an existing VO);
 - Requests to the “User Board”/PMB for site support and initial resource allocation;
 - An understanding that development and testing will be required from the users.
 - *Production:*
 - Fully-functioning system supported by resource from the community in question.
- ...and set and manage expectations accordingly.

Some general observations

- We can offer the best solution (and avoid problems) if we:
 - *Properly profile/understand the user community being engaged;*
 - *Understand what they're currently doing and how they're doing it;*
 - *Establish the feasibility of employing a distributed computing solution.*
- Suggestion: the “First Contact” phase should therefore incorporate an “initial requirements capture” exercise to do this:
 - *Carried out by the nominated GridPP Champion;*
 - *Questions answered through discussion with the new user(s) - no forms to complete, no MoUs – this is for our benefit;*
 - *Draft version presented here for discussion and feedback.*
- We will look at a case study with the GalDyn group.
 - *Completed document (plus draft template) uploaded to Indico.*

Initial requirements capture - outline

- Overview:
 - *Name, points of contact, and a brief summary of “first contact” scope.*
- The user community:
 - *List the organisations and users, their locations, prior grid involvement.*
- The current workflow(s):
 - *Describe the current workflow(s) – user interactions, software, data.*
- Moving to the grid:
 - *Establish how the current workflow could be moved to the grid (if at all);*
 - *Characterise existing infrastructure - potential benefits and barriers;*
 - *Establish User Interface, networking, software, and data requirements.*

Overview

- Basic details:
 - *Community name: GalDyn*
 - *Point of contact: Adam Clarke (UCLan)*
 - *GridPP point of contact: Tom Whyntie (QMUL)*
- Summary of the proposed activity:
 - *The Galaxy Dynamics (GalDyn) group at the University of Central Lancashire (UCLan) is part of the Jeremiah Horrocks Institute. One area of their research concerns simulating the orbits of galactic matter (stars, dark matter, etc.) in order to better understand how galaxies form. One researcher from the GalDyn group would like to investigate how the grid could be used to perform tens of thousands of simulations to study the effect of varying different orbital parameters.*

The user community

- List the organisations involved:
 - *The **Galaxy Dynamics group** is a university research group (Uni. RG) within the Jeremiah Horrocks Institute at the University of Central Lancashire (UCLan).*
- Establish if there has been any previous grid activity (including other grids e.g. OSG):
 - *None with GalDyn – a clean slate.*

Include individual institutions and any collaborations to which they belong – establish who the stakeholders are.

If there has been, there could be legacy/external infrastructure to integrate into any solution.

The user community

- List the users involved in the “First Contact” stage:

Identifying the type of organisation can flag up potential facilitating factors and barriers in advance

Check for non-UK users

Having a local UKCA RA removes a travel barrier

User	Title	Role	Org. type	UK?	T1/2?	RA?	Grid?
AU	Researcher	Developer	Uni. RG	Yes	No	No	No
VD	Group leader	End user	Uni. RG	Yes	No	No	No

For reference

Establish who the developers, testers, and end users are

Users at Tier-1 or Tier-2 sites will have on-site support

Is the user already on the grid?

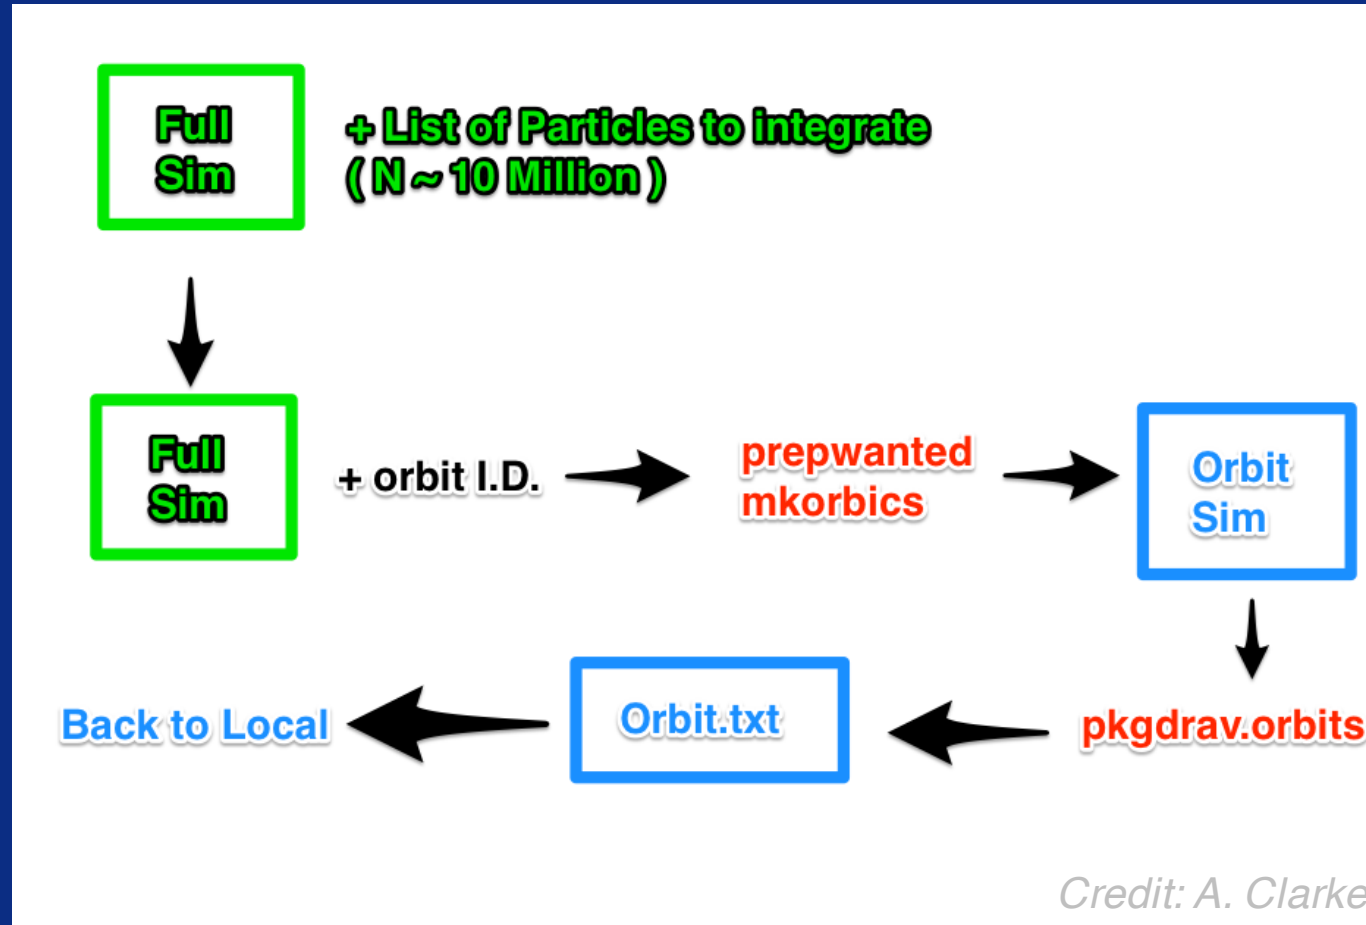
The current workflow

- Provide a brief description of the current workflow(s):
 - *A base simulation file (~1GB) is configured for a given set of parameters by two bash scripts, producing an intermediary “orbit simulation” file (~1GB). Another package then runs on this file to produce the required output containing the results (a ~2MB text file).*
- Where does it run?
 - *A single machine – user’s own or on cluster.*
- How does the user interact with it?
 - *Commands in the terminal (ssh to cluster).*

This is probably the hardest – but most important – part of the whole process. Use diagrams if you can (see next slide).

These questions are more about establishing what infrastructure is available, what the user is used to, and what could be incorporated when running on the grid.

The current workflow



The current workflow - software

- What software is required?
 - *Two bash scripts, prepwanted and mkorbits, are needed to prepare the orbit simulation file for a given set of parameters. A C++ executable, pkgdrav.orbits, is used to extract the results from the orbit simulation file to the output text file.*
- Does any of it require a license?
 - *No.*
- Is the source code available?
 - *Yes – private BitBucket repo.*

Identify the scripts, executables and libraries that feature in the workflow as it stands.

Licensing is probably a showstopper unless it can be factored out of the workflow – CVMFS certainly ruled out.

The ability to (re)compile for SL6 running makes life easier.

The current workflow - data

- Input datasets:
 - *A base simulation file (~1GB) which is generated once (independent of the workflow), and; a orbital parameter configuration file (a small text file). Both are stored locally to wherever the workflow is being run.*
- Output datasets:
 - *A ~2MB text file for each parameter set. This is stored locally.*
- Intermediary datasets:
 - *The base simulation file is configured for a given set of orbital parameters to produce an orbital simulation file (~1GB). This does not necessarily need to be stored (though this may be useful).*
- Data management:
 - *Input and output files managed in a named directory structure. No cataloguing or database usage required at this scale.*

Legacy infrastructure

- Are there any legacy infrastructure requirements to consider? Have distributed computing systems been used to execute the workflow or manage datasets in the past?
 - *No – the workflow so far has only been conducted locally on a single machine.*

While this was not an issue for GalDyn – a user community starting from scratch – this is a question with huge implications for user communities for whom the answer is “yes”.

Moving to the grid

- This section is about establishing what the most appropriate way of getting the user onto the grid would be, looking at:
 - *Which User Interface (UI) would be most appropriate;*
 - *The networking capabilities (and restrictions) of the user(s);*
 - *How the workflow could be adapted for grid running;*
 - *The software requirements for the adapted workflow;*
 - *The data requirements for the adapted workflow.*
- We will now look at the answers for GalDyn and which tools/solutions were adopted as a result of these.

The User Interface

- Grid cluster access?
 - *No.*
- Terminal with ssh access?
 - *Yes.*
- Hardware capable of running a VM?
 - *Yes.*

If the users are local to a site with a grid cluster, given them an account vastly reduces setup, configuration and networking issues. It may be worth giving non-local users access to a grid cluster, particularly if they are behind a firewall where VM networking may prove tricky.

A user's local machine should have enough memory to comfortably run a VM for their User Interface if a grid cluster isn't available.

As the GalDyn user had the hardware required to run a CernVM, we used a GridPP CernVM for their grid UI.

Networking

- How does the user connect to the Internet?
 - *The user connects to the UCLan network via an Ethernet connection on their personal laptop. They may also conduct work from home where the connection is provided by a commercial ISP.*
- Bandwidth?
 - *Typical for a university network.*
- Firewall?
 - *Yes, but not an issue for the CernVM.*

Experience to date has shown that when users are connecting via a university network it is easier to offer them some kind of cluster account. Some grid tools (e.g. those used when dealing with storage) have issues with how networking is configured on the VM.

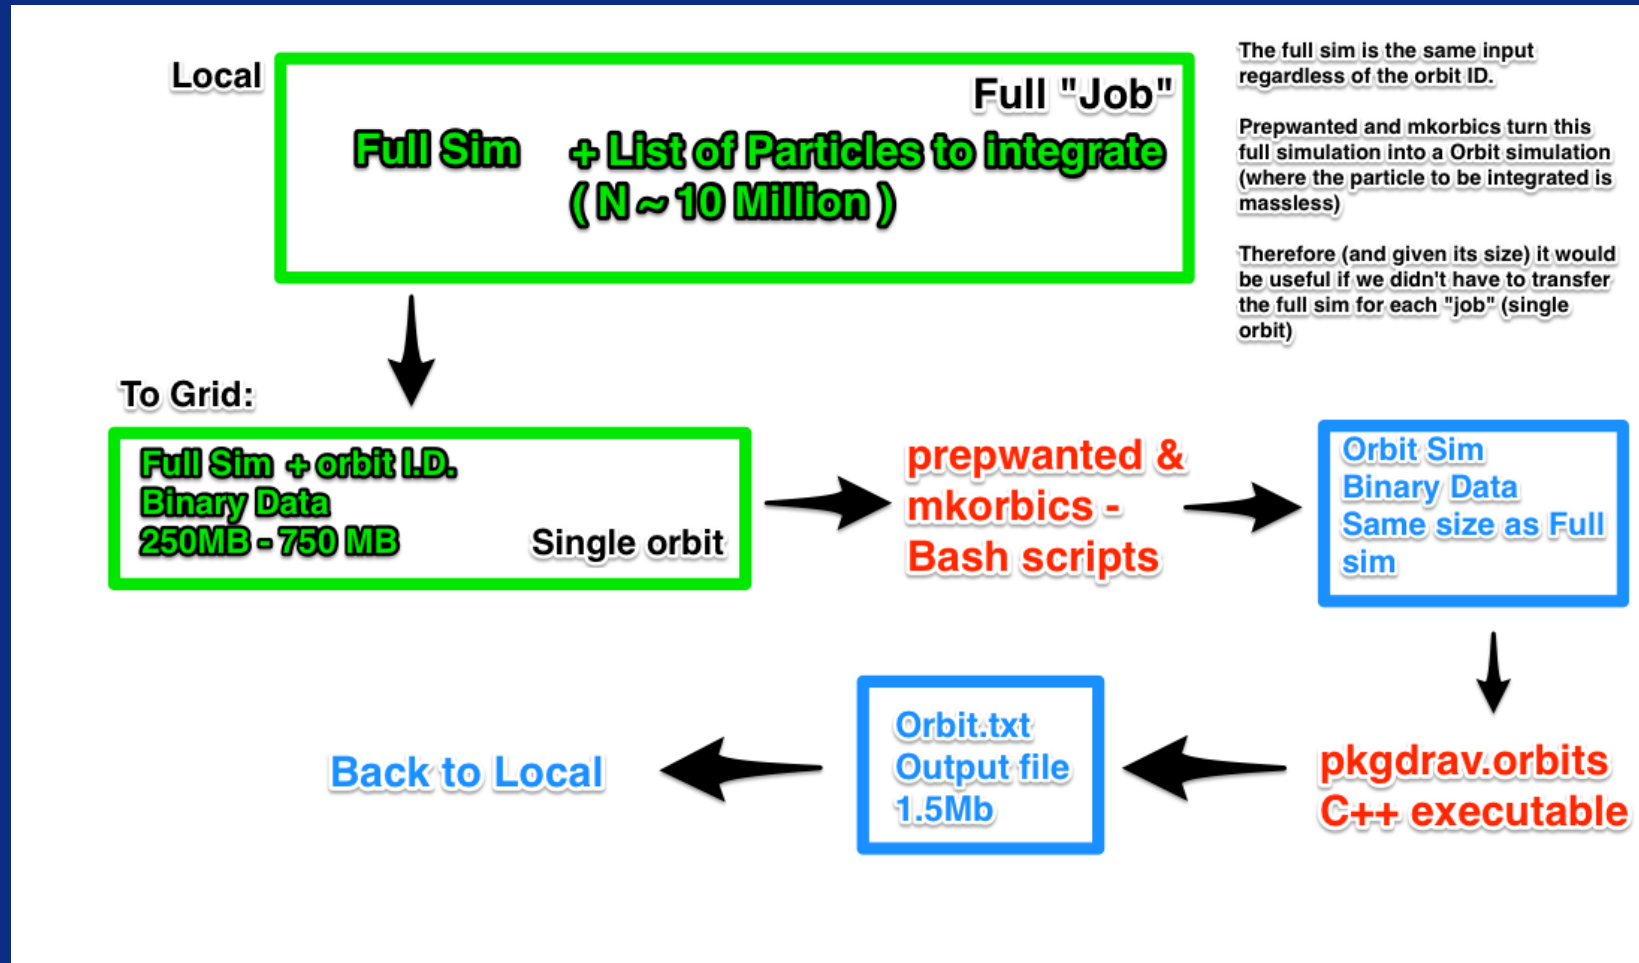
The cluster option for GalDyn has not been ruled out for future operations.

Adapting the workflow

- Which parts of the workflow could be moved to the grid?
 - *The GalDyn group would like to investigate the effect of varying the orbital parameters of interest on the scale of tens of thousands of parameter sets. As runs for given parameter sets are independent of each other, both the orbit simulation and results extraction stages could be performed on the grid. A suggested modification to the workflow is shown in the following diagram (next slide).*

This is probably the second-most important part of the whole process. Don't forget, it may well be that a distributed computing solution is not appropriate for the workflow in question.

Adapting the workflow



Software

- What would need to be on the grid?
 - *The prepwanted and mkorbics bash scripts and the C++ executable pkgdrav.orbit would need to be run on the worker nodes.*
- Can these run on SL6?
 - *Yes – the source code could be compiled to run on SL6.*
- Are any licenses required? *No.*

Where source code is unavailable, custom/proprietary libraries are required, or the user is unable/unwilling to use SL6, a VAC-type solution may be necessary. The overheads associated with both approaches need to be carefully considered.

Licensing software for grid use could throw up many barriers...

GalDyn's software was compiled and tested on the GridPP CernVM before being deployed via a custom CVMFS – galdyn.egi.eu

Datasets

- Input datasets:
 - *The full simulation file (~1GB) is generated once – upload to SEs and retrieved on demand. Parameter set files uploaded on a job-by-job basis.*
- Output datasets:
 - *The user(s) will need to retrieve the output text file (~2MB) generated by each parameter set.*
- Intermediary datasets:
 - *The orbit simulation files (~2GB each) could be stored for later use but this is not necessary.*

At this point, the options for grid-based data management solutions should be considered, particularly if there are legacy catalogues/databases.

Solutions used and current status

- These requirements suggested the following approach:
 - *User obtained a grid certificate, joined NorthGrid VO;*
 - *A GridPP CernVM was used as the UI;*
 - *Software compiled and tested on the CernVM, deployed with CVMFS;*
 - *GridPP DIRAC used to manage both jobs and data;*
 - *Input data (full simulation file) stored and catalogued via the DIRAC File Catalog (DFC). Parameter files uploaded via the InputSandbox. Intermediary/output data not stored on the grid (yet).*
- Workflow successfully implemented on the grid – the next stage is to scale up to the O(10k) jobs:
 - *Time for Ganga! (With the DIRAC backend.)*

Summary and conclusions

- The New User Engagement Programme is in full swing:
 - *Plenty of progress to track/monitor;*
 - *Is a “one-size-fits-all” approach useful?*
 - *We need to define phases of engagement – First Contact? Development? Production? – and manage expectations accordingly.*
- To offer the most appropriate distributed computing solution, we need to understand what a given community needs:
 - *As part of the “First Contact” phase, useful to conduct a “initial requirements capture”;*
 - *Case study featuring the GalDyn group presented;*
 - *Completed example and draft template provided (see Indico material).*

@GridPP

@twhyntie

Thanks for listening! Any questions?

T. Whyntie*, †

** Queen Mary University of London; † Institute for Research in Schools*

#GridPP36, Pitlochry, Scotland

Wednesday 13th April 2016



Types of organisation

- Independent individual (i.e. no organisation)
- University
- School
- Other academic institution
- Other research institution
- Government department
- Non-Government Organisation (NGO)
- Small-to-Medium Enterprise (SME)
- University spin-off company

Actually not very useful information in its own right – most facilitating factors/barriers dependent on an individual user's specific situation.