

Just-in-time workflow management for DUNE

Andrew McNab, University of Manchester, for the DUNE Collaboration

DUNE is a future liquid argon TPC neutrino oscillation and astrophysical neutrino experiment that will take data at a rate of 30 PB/year.

Computing development

The DUNE computing system has evolved from the heritage of neutrino and collider experiments based at Fermilab. To achieve the increase in scale required by DUNE it has been necessary to modernise the framework that was largely built around SAM ("Serial Access to Metadata"), dating back to the Tevatron collider experiments. Where possible, DUNE is reusing tools developed for other HEP experiments and for WLCG.

DUNE has chosen to adopt Rucio as the basis for its data management strategy, and has worked with Fermilab to develop a new metadata database, MetaCat. Together these systems allow for the discovery of which files match particular criteria and where the replicas of those files are located.

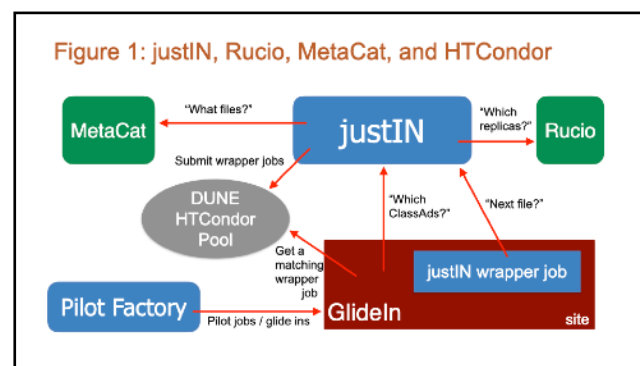
The arrival of large amounts of storage at DUNE sites other than Fermilab has also exposed limitations with the SAM-based framework when matching work to sites. To address this, the UK has led the DUNE workflow system design and development, which aims to run work at or near sites where the relevant data is stored where necessary.

justIN workflow system

One of the strengths of the SAM model was that the choice of which file a job should process next was made at the last minute based on which files were still unprocessed. Experiments were started in 2021 to extend this model to the matching of workflows to available job slots. This led to a complete just-in-time workflow system named justIN, in which the decision about which workflow to run in a job slot is based on a complex query over all the running workflows, the unprocessed files of those workflows, and the distances from the execution site in question to replicas of those files.

Services and Agents

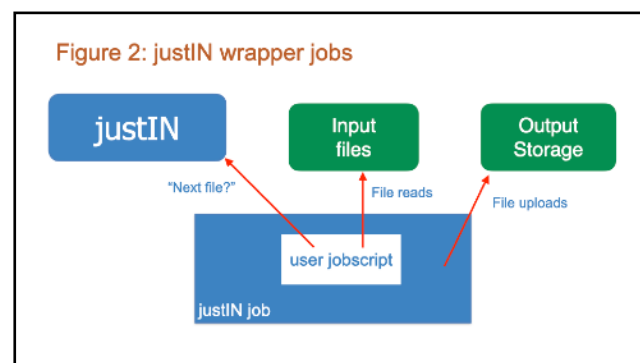
The justIN job factory agent submits wrapper jobs to the DUNE HTCondor Global Pool, with ClassAds specifying the workflow's memory and processor requirements, and the number of the workflow.



glidelnWMS pilots are submitted by the OSG pilot factory, and a custom DUNE script contacts the justIN Allocator Service to obtain the correct ClassAd requirements for this job slot to match jobs from workflows, based on the proximity of unprocessed files. Since the final matching is still done by the ClassAd system, HTCondor's accounting system is still able to prioritize justIN vs non-justIN jobs, and between different justIN users based on their recent usage history.

Other agents gather information about sites and storages from the OSG pilot factories and Rucio, and manage the lifecycle of workflows including collecting relevant file and replica location information from MetaCat and Rucio, detecting errors and retrying the processing of files if jobs fail.

Wrapper jobs



The wrapper jobs themselves execute the jobscript provided by the user in an Aptainer container to enforce privilege separation. Users' X.509 and VOMS credentials are not used, but a low privilege DUNE X.509 proxy is made

available to jobscripts which can only read files from storage. Once the jobscript is finished, the wrapper job uses a higher level X.509 proxy or token to write the output files to storage and register them in Rucio and MetaCat. This model allows DUNE to enforce rules about how the storage is used by different groups of users without the storage systems needing to enforce these distinctions.

Authentication and Authorization

Users authenticate to justIN via the justIN web dashboard whether using the dashboard itself or working at the command line.



Authentication is currently via CILogon and the Fermilab Single Sign-on service, and other eduGAIN compatible sign-on systems will be available in the future. CILogon releases the user's DUNE SciTokens / WLCG Tokens to justIN, which justIN is able to use on the user's behalf and to parse to discover the user's group memberships.

At the command line, the justin command invites the user to visit a one time URL on the dashboard website to authorise the command line environment to act as the user and the web based authentication proceeds as normal. The simplicity of this approach allows for a self-contained Python3 client script which can be installed by simply copying it in place on Linux or macOS.

justIN in use

Following its successful validation during Data Challenge 24 in 2024, justIN has been the basis for all DUNE central productions and is being rolled out for general users and working groups. Other posters and talks at CHEP describe these developments.