

White paper

FUJITSU Software HPC Gateway

Fujitsu brings HPC to the Desktop.

Simplicity of end-use enables more users – both new and experienced – to work with HPC. While in-built expertise on HPC applications allows your business to deploy robust and proven simulation methods on more projects. HPC really is now more accessible than you might think.



Content

Introduction	2
HPC as a mainstream business process	2
HPC Simplicity	3
Intuitive user interface	3
Structuring your HPC project	3
An intelligent HPC workplace	3
Building-in application expertise	4
Fujitsu Application Catalogue	4
Application On-Boarding	5
Using workflow to automate HPC methods	5
Working with the HPC Gateway	6
Application Explorer	6
Application monitoring	7
Data handling	7
Security mapped to business entities	8
Conclusion	8

Introduction

Latest advances in simulation software and hardware technologies are making High Performance Computing (HPC) ever more attractive and relevant to businesses looking to increase, or restore, competitiveness. This could mean lowering prototyping costs, shortening design or validation timeframes, or improving quality through optimisation. Equally, HPC contributes to drive innovation through design exploration, and research using virtual experiments that cannot be practically built.

Most industrial organisations would naturally acknowledge the benefits of HPC, and many larger companies have integrated such activities into their main design and production processes. HPC is helping them expand exploration of a products design space and to increase product performance, quality and reliability – whether to meet further regulations or increase competitiveness.

This whitepaper describes how a more intuitive interface and packaged expertise offer a better way to bring HPC to the desktop, allowing more organisations to incorporate HPC as a core business process.

HPC as a mainstream business process

With any technology it is the ability to embed its use into a coherent and repeatable process that determines whether that capability becomes critical to the business or remains peripheral. And here is the contradiction. HPC is often treated as a process or technique that requires some knowledge of the lower-level IT infrastructure to use effectively. For example, in many cases HPC workloads are still run manually from a direct Linux command line on the shared cluster, involving editing scripts and preparing parallel batch jobs. Simplification and standardisation of this approach will help to consolidate HPC as a fundamental business process, and one that contributes consistently to the performance of the company.

This document describes Fujitsu's HPC Gateway web-based user workplace for Linux clusters. It is available within each FUJITSU Integrated System PRIMEFLEX for HPC solution furnished with our HPC Cluster Suite (HCS) – a dedicated software stack for HPC workloads.



HPC Simplicity

Fujitsu's HPC Gateway represents a coherent and overarching user workplace for HPC projects. It bridges low-level infrastructure layer – operating system, batch resource manager, file system, networks – to provide an inclusive way of working with HPC, from novice to expert users. Companies making their first move into scalable HPC can be reassured that this investment will be utilised efficiently and confidently by the whole team. Enterprises already working with HPC will find ways to enhance the coordination of the team, optimise workloads, and broaden scalable HPC usage by more groups and on more projects. What this really means is lower risk and overall cost for your organisation.

Intuitive user interface

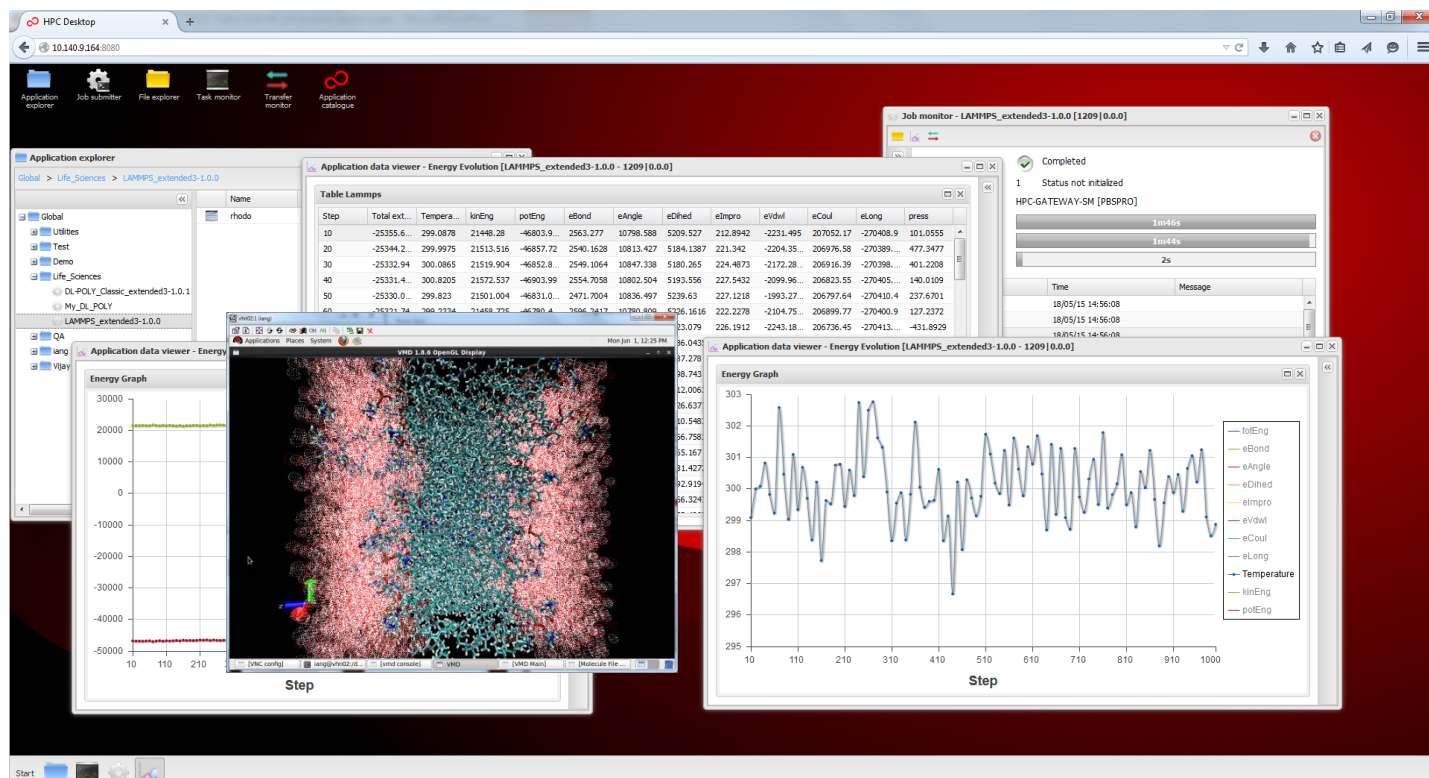
The HPC Gateway Application Desktop is web-based user interface for running and organising HPC workloads. It adopts the familiar “desktop” screen as an intuitive and adaptable workplace for organising and controlling applications running on servers. With this interface alone the HPC Gateway improves productivity across the team by providing tools adapted for different types of usage, whether running job scripts directly or, more commonly, using guided structures for preparing and submitting application jobs. One obvious reason for the general success of the desktop concept as a IT interface is the ability to retain multiple dialogues active concurrently. This advantage distinguishes the HPC Gateway from other HPC web interfaces in being more flexible, adaptable and reactive. Equally, the Gateway works entirely within the native browser, and does not require any other software components to be installed in the client device.

Structuring your HPC project

Naturally the first step in simplifying HPC usage is to decouple the primitive script edit and submission approach when using the lower-level Linux command-line interaction. In the HPC Gateway Application Desktop jobs are prepared and submitted from parameter based forms. Required and optional parameters are clearly identified, and drop-down lists offer fixed or suggested values where possible. Each form instance can be separately saved as a job profile. In this way a series of jobs can be recorded as part of a campaign of experiments. At a minimum such stored profiles can be reused as drafts for subsequent jobs, or rerun for validation purposes. More formally, the set of stored profiles may become part of an auditable trace. Individuals can further customise their own HPC Gateway setup through shortcuts on the web desktop, giving fast access to data folders, applications and job profiles.

An intelligent HPC workplace

Only simplifying the user interface is not the complete answer to broadening the reach of HPC. Without the expert-based means to drive applications and HPC workloads effectively, the range of users will always remain limited and narrow, regardless how easy the desktop interface appears. And expertise that can only be re-used by the expert or near-expert doesn't widen the appeal and opportunity of HPC. Simplicity and Expertise work in conjunction, and HPC Gateway uniquely provides a medium to capture and transfer expertise around applications and the methods to use them.



Building-in application expertise

There are three main ways in which an application can be enabled for use within the HPC Gateway:

- Pre-built package import
- Application On-Boarding
- Workflow development

Import is a function to allow transfer of application-related expertise. Fujitsu maintains an online Application Catalogue of pre-built packages that comprise encoded methods and monitors for selected codes. The HPC Gateway also has an Export function that creates the pre-built packages for given application methods. Expert service providers can also develop such pre-built packages to convey to customer HPC Gateway sites through this export capability.

Fujitsu Application Catalogue

The Fujitsu Application catalogue contains pre-built packages for many widely-used HPC applications. Each package encodes an automated and intelligent method to run the application that captures our expertise and insights into optimal ways to run, monitor and control applications on HPC clusters. This knowledge is continually enriched from interactions with ISVs, end-users and both internal and external research activities. New and extended methods are regularly added to the catalogue with enhancements aimed at higher robustness, more detailed tracking and a richer monitoring and reporting experience. Currently the Application Catalogue mainly addresses codes in the CAE and Computational Chemistry disciplines.

HPC Gateway users are encouraged to regularly consult our Application Catalogue to check the latest additions, whether for new applications in their field, emerging disciplines or methods with additional features and options. This exchange ensures that the methods used at the client site are constantly refreshed, with the aim of sustaining high productivity and broadening the accessibility of HPC to more users through such standardised and expert-driven run-time methods.

Using the Fujitsu Application Catalogue

Importing a pre-built package

Clicking an icon on your HPC Gateway desktop will open a link to the Fujitsu Application Catalogue web site. As an authenticated HPC Gateway user you can then select and download a pre-built package – a single packed archive file. Once this file is placed on the cluster filesystem run the import utility through the HPC Gateway to install the new application method.

The diagram illustrates the process of importing a pre-built package from the Fujitsu Application Catalogue to the HPC Gateway user site. It consists of two main components: the HPC Gateway at the user site and the Fujitsu Application Catalogue site.

HPC Gateway at user site: This component shows a web browser window with the HPC Gateway interface. The 'Application Catalogue' link is highlighted in the top navigation bar. Below the navigation bar, the 'Application Catalogue' link is also highlighted in the left sidebar. The main content area displays a table of application packages. The table has columns for 'Name', 'Status', 'Revision', 'Owner', 'Creation time', and 'Description'. The table lists various application packages, including 'DL_POLY_4solve', 'BLAST_21_extended2-1.0.0', 'CPK_21_extended2-1.0.0', 'NWCHEM_61_solve', 'NWCHEM_61_solve', 'GROMACS_46_solve', 'LAMMPS_solve', 'DL_POLY_Classic_solve', 'CPK_21_solve', 'CASTEP_60_extended2-1.0.0', 'DL_POLY_Classic_solve', 'LAMMPS_solve', 'NWCHEM_61_solve', and 'QUANTUM_EXPRESSO_50'. The 'DL_POLY_4solve' package is highlighted in the table.

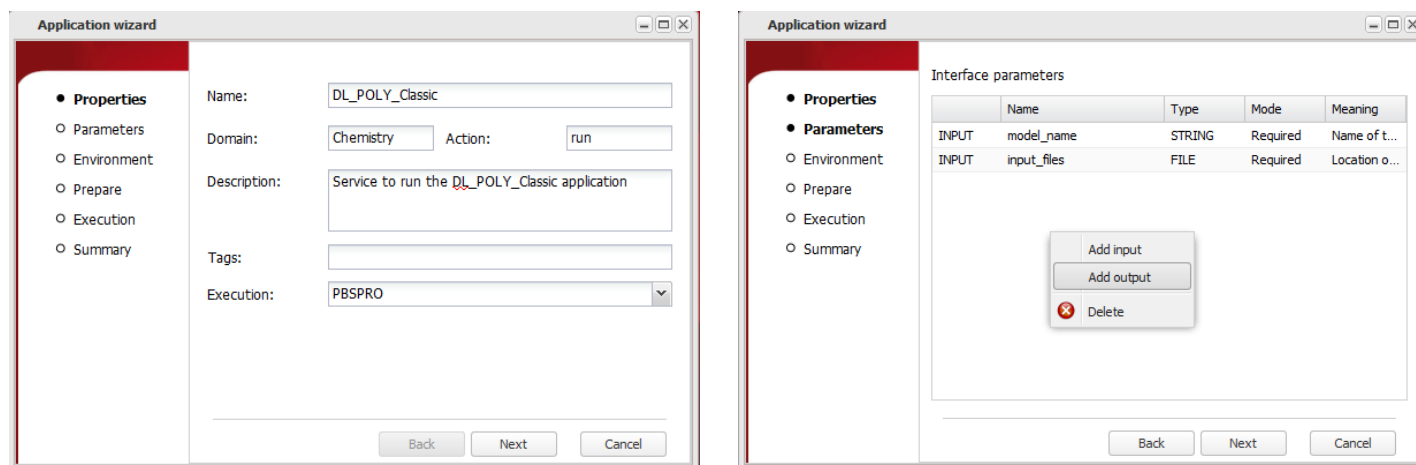
Fujitsu Application Catalogue site: This component shows a web browser window with the Fujitsu Application Catalogue site. The site displays a table of application packages. The table has columns for 'Category', 'Application name', 'Package description', 'User guide', 'Instruction guide', and 'Download'. The table lists various application packages, including 'Life Sciences', 'NCBI BLAST v2.2', 'NCBI BLAST v2.2', 'CASTEP v6.0', 'CASTEP v6.0', 'CASTEP v6.0', 'CPK v2.3', 'CPK v2.3', 'DL-POLY Classic (based on DL-POLY version 1.9)', 'DL-POLY Classic (based on DL-POLY version 1.9)', 'DL-POLY Classic (based on DL-POLY version 1.9)', 'GROMACS v4.6', 'GROMACS v4.6', 'GROMACS v4.6', 'GROMACS v4.6', and 'GROMACS v4.6'. The 'DL-POLY Classic (based on DL-POLY version 1.9)' package is highlighted in the table.

Consult online Fujitsu Application Catalogue: A red arrow points from the 'Application Catalogue' link in the HPC Gateway user site to the Fujitsu Application Catalogue site.

Download and import selected method: A red arrow points from the 'DL-POLY_4solve' package in the HPC Gateway user site to the 'DL-POLY Classic (based on DL-POLY version 1.9)' package in the Fujitsu Application Catalogue site.

Application On-Boarding

Benefits from deploying the HPC Gateway grow as more applications are enabled. For specific solutions the packages in our Application Catalogue may cover all applications used on the cluster. Bringing the complete set of codes used for all projects into your HPC Gateway system allows individual users to remain within the one consistent environment. It also ensures that users new to HPC, or even temporary consultants, can work from a validated template – sustaining a higher aggregate productivity and lowering support efforts. For these reasons the HPC Gateway provides an Application On-Boarding utility, allowing you to define and on-board a new interface. No web coding is required to enable an application to be used through the HPC Gateway. This definition is created through a wizard dialogue that guides you through the setup of each parameter. From an existing command-line script an application can be enabled in the Gateway in a few minutes.



The use of expert-programmed methods is part of a broader imperative within HPC. A presentation from the NCSA (National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign) at the HPC User Forum meeting in July 2014 noted that although commercial applications are now central to all manufacturers 90% remain on PCs. To continue to shorten design time and increase realism much more of this usage will need to move to HPC; that is, run their models on scalable clusters. Some of the inertia can be due to limited parallelism in the code itself. But a larger resistance comes from the barriers perceived by end-users themselves in running their workloads on HPC systems. When HPC experts encode the methods they use into shareable and consistently reusable objects then many of these barriers can be overcome. Less experienced users can then work with a simple interface to sophisticated and intelligent processes, delivering more valuable results and maximising the HPC potential.

Using workflow to automate HPC methods

To further support the drive towards expertise and knowledge capture, HPC Gateway includes both workflow programming and a workflow engine. In this context workflows are a set of elemental application and utility services arranged to be automatically executed in a particular sequence. Just some of the potential value from workflows includes:

- Tuned methods that link dependent and repeatedly used sequences of tasks.
- Expert knowledge that is more effectively disseminated as a workflow.
- Best-practice to be standardised across the organisation.
- Process steering based on data manipulation to control workflow path.
- Process and data provenance – the attribution of both services and data to identifiable entities (projects, suppliers, partners).

Within HPC Gateway the workflow capability is highly evolved, with support for split/join and loop structures, and in particular data flow (optionally compressed/encrypted) between tasks. Such workflows may incorporate services already imported or on-boarded into your deployed HPC Gateway, or may involve adding new application services.

The use of workflow across HPC operations is growing in importance. Basic examples include the expansion of optimisation and parameter sweep studies, as well as other design exploration methods. And on a more general level workflow when developed by the organisation workflows can thus represent defining assets of the company, methods that are proven and highly optimised for the business operations.

Working with the HPC Gateway

Fujitsu's HPC Gateway is part of the dedicated HPC Cluster Suite offered with PRIMEFLEX for HPC cluster solutions. Together, the HPC Gateway functions provide a consistent and standardised range of tools for the majority of HPC workloads. Specifics on the HPC Gateway functions are given in the HPC Cluster Suite (HCS) datasheet. Here we describe the use of the main working elements.

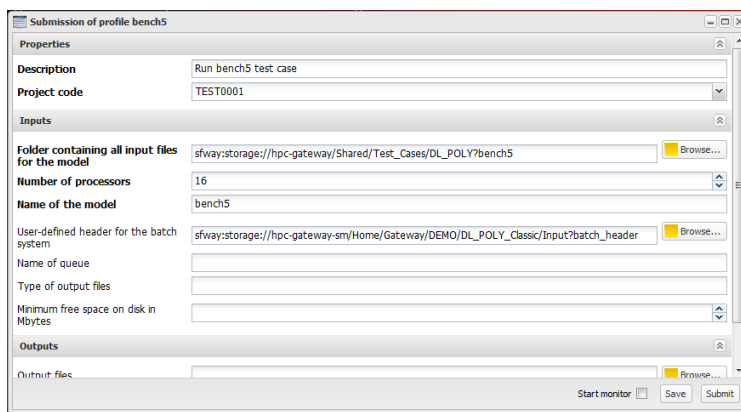

Main Gateway functions

HPC Gateway is a web-based environment that runs on Linux. It supports all common browsers such as Mozilla Firefox, Google Chrome and Microsoft Internet Explorer. The HPC Gateway requires a positive authentication to identify the end-user. This identity is then authorised to objects in the Application Desktop through an integrated role-based access control (RBAC) authorisation mechanism.

Function	Description
Application Explorer	Prepare and organize job parameters, run application jobs, run workflows
Job Submitter	Run own scripts directly, edit and manage script files
File Explorer	Navigate cluster file systems, upload/download files, edit text files, cut/paste
Task Monitor	Monitor jobs and workflows tasks, track and visualise key application results
Application On-Boarding	Mechanism to incorporate single script (single task) applications in the HPC Gateway
Application Catalogue	Link to Fujitsu site hosting downloadable pre-built application processes
Workflow Editor (HCS Advanced Edition)	Separate desktop tool to build more complex multi-task processes
Administration Dashboard (HCS Advanced Edition)	Separate desktop tool to customise and manage the HPC Gateway environment.

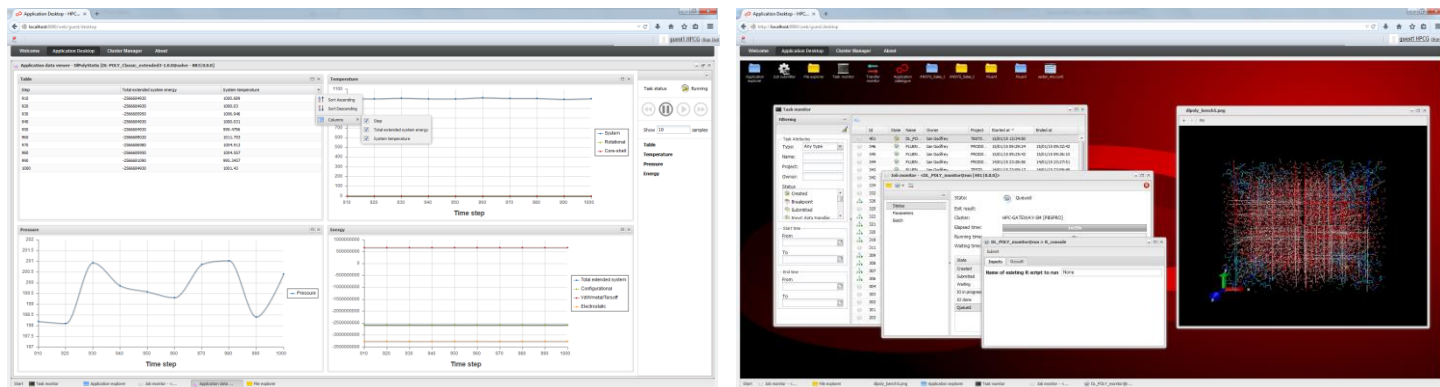
Application Explorer

This tool shows the application methods available to the individual user. For each tool the user may create one or more job profiles in which they set the parameters for the given run of the associated method. The method itself may be a single batch task, or an automated workflow of dependent tasks. Methods are arranged in folders (hence the name of Explorer for this tool), conventionally named according to discipline or project. Folders may be stacked in a hierarchy, and the same method can appear in multiple folders. Such folder organisation can be fully adapted by the customer. A profile may be completely new, or adapted from a previous setting. Having saved the profile the job can then be submitted. Users can maintain multiple profiles to establish a progression within a particular study. Or they may be retained to rerun later.



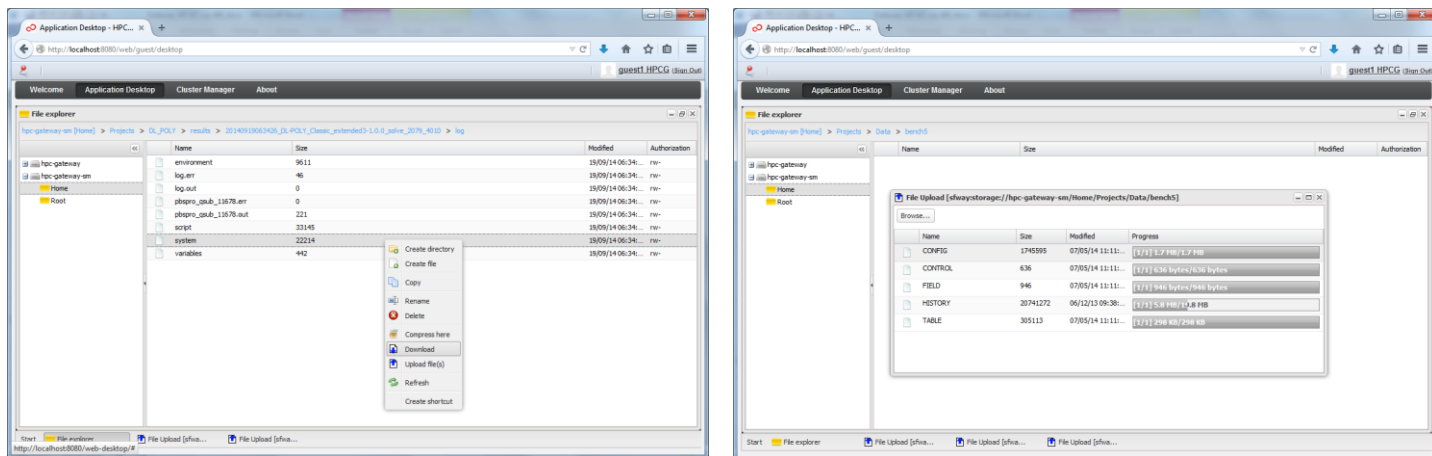
Application monitoring

The Task Monitor shows all the user jobs active or completed in the HPC Gateway system. For easier viewing these can be filtered by start/end time ranges, and by key attributes (name, project, status). Selecting a task opens the detailed Job Monitor. Here the user can track progress in the system, check input parameters, and inspect batch status. More importantly the HPC Gateway also provides support for direct application result monitoring. Using an active sensor capability relevant application metrics can be displayed dynamically by the HPC Gateway. Commonly these are data results written into output files at regular intervals, and which provide the user with reference points to track progress of the simulation or analysis. Many packages from the Fujitsu Application Catalogue are developed to extract and display such metrics in plots and tables that are updated automatically during execution. Simpler yet still effective monitors are available through HPC Gateway task sub-actions. These can be used to perform any type of programmable operation directly on the running task. For example sub-actions are widely used to inspect parts of the main result files to validate the state of the computation.



Data handling

The HPC Gateway provides a File Explorer tool to work with files/folder on the host cluster filesystem, and to upload/download between the client device. Conventional file and folder treatment is supported, and multiple files can be exchanged in a single operation. Unlike other such Linux visual file explorers, however, the HPC Gateway introduces its own set of virtual mount points to map onto the remote filesystem. First, such mount points allow for better organisation of the filesystem layout since mount points can be defined to represent project areas, application data locations and other business-related objects. HPC Gateway RBAC permissions apply directly to such mount points, enabling a more business-oriented security model to be established within the shared cluster environment.



Security mapped to business entities

The HPC Gateway incorporates its own security layer above, and complementary to, the underlying Linux structures. It uses role-based access control (RBAC) to limit the operations that can be performed by the connected identity. RBAC provides a powerful and adaptable mechanism to structure permissions for objects used through the HPC Gateway – filesystem mount points, application methods, project folders, job profiles, etc. – and allows better alignment with business entities (projects, teams, skills, responsibilities).

- **Role:** A Role is a job function in an organization. Roles can be defined in a hierarchy. The role hierarchy defines a seniority relation between roles: senior roles acquire the permissions of their juniors.
- **Identities:** An identity is usually and individual users of a named group. Groups can be created for specific business structures, typically projects or teams. User identities can be assigned to one or more groups.
- **Permissions:** A permission is an allowed *action*, on a given *resource type*, with a specific *rule*.

The combination of pre-defined application methods, web interface and strong RBAC permissions adds to an overall defence-in-depth security model that protects the server from misuse.

Conclusion

FUJITSU Software HPC Gateway encompasses an approach to broaden the accessibility to and value derived from scalable cluster computing. Organisations in a growing number of sectors are recognising the advantages of moving from physical to virtual methods of design, prototyping and testing. This means more users working on more projects are looking at how they can use such systems. While Linux remains the primary operating system, newer users have little or no familiarity in working with the command line and scripts. Instead expectations are for a more visual and intuitive workplace as the basis.

Simplicity therefore provides the entry point, allowing more users to learn quickly and work effectively with HPC. Our opinion is that simplifying the HPC interface alone is not the complete answer. Without the expert-based means to drive applications and HPC workloads effectively, the range of users will always remain limited, regardless how easy the interface appears. With the HPC Gateway we also have a means to productise and transfer expertise related to applications. Fujitsu maintain an Application Catalogue of encode HPC methods, productized for download and import into the local HPC Gateway installation. Such industrialised expertise can then be scaled to the largest number of individuals and teams, whether practiced users or HPC newcomers.

Contact

Fujitsu Technology Solutions GmbH

Website: www.fujitsu.com/global/hpc
2015-01-16 WW - EN

© 2015 Fujitsu Technology Solutions GmbH. Fujitsu and the Fujitsu logo are trademarks or registered trademarks of Fujitsu Limited in Japan and other countries. Other company, product and service names may be trademarks or registered trademarks of their respective owners. Technical data subject to modification and delivery subject to availability. Any liability that the data and illustrations are complete, actual or correct is excluded. Designations may be trademarks and/or copyrights of the respective manufacturer, the use of which by third parties for their own purposes may infringe the rights of such owner.