

UNICORE in XSEDE: Towards a large-scale scientific Environment based on Open Standards

Evolution from TeraGrid

Starting in 2001, the National Science Foundation program TeraGrid has developed into one of the world's largest and most comprehensive Grid projects, offering resources and services to more than 10,000 scientists. Its successor, the Extreme Science and Engineering Discovery Environment (XSEDE, www.xsede.org), has started in July 2011 and is expected to excel the previous program in terms of service quality while lowering technological entry barriers at the same time. These and other goals are to be achieved in the project's five year grant period with an overall budget of \$121 million. Among the partnership of 17 institutions, the Jülich Supercomputing Centre (JSC) is the only organization located outside the USA.

Open Standards-based Architecture

Since many scientific communities operate internationally, one key element of XSEDE is the use of open standards in order to promote interoperability with other distributed computing infrastructures such as PRACE in Europe. Figure 1 shows the extended reference architecture providing mandatory XSEDE Enterprise Services at every major XSEDE site as well as optionally available Community Provided Services. For many years, the JSC and several other XSEDE partners have been active

in establishing the key standards that now define the interfaces of the XSEDE Enterprise Services. Within the Grid Interoperation Now (GIN) community group of the Open Grid Forum (OGF, www.ogf.org), such key standards as BES/JSDL for running remote computations have demonstrated their impact on scientific applications. Based on these standards, scientific workflows can be executed today across different infrastructures with no less than 8 different Grid middleware technologies.

Jülich's Role in XSEDE

The JSC not only contributes its extensive Grid know-how gained from European research projects and its experience in standard-based software engineering, but also a technology called Uniform Interface to Computing Resources (UNICORE, www.unicore.eu). Being developed by partners all over Europe, UNICORE is a Grid system that provides secure and seamless access mechanisms to a variety of different computer systems and platforms. It facilitates the remote execution of scientific applications as well as sharing software, resources and data. UNICORE is fully based on Web services and open standards in order to allow seamless interoperation with other standard compliant Grid systems such as Genesis II which is developed at the University of Virginia. Being complementary to the more

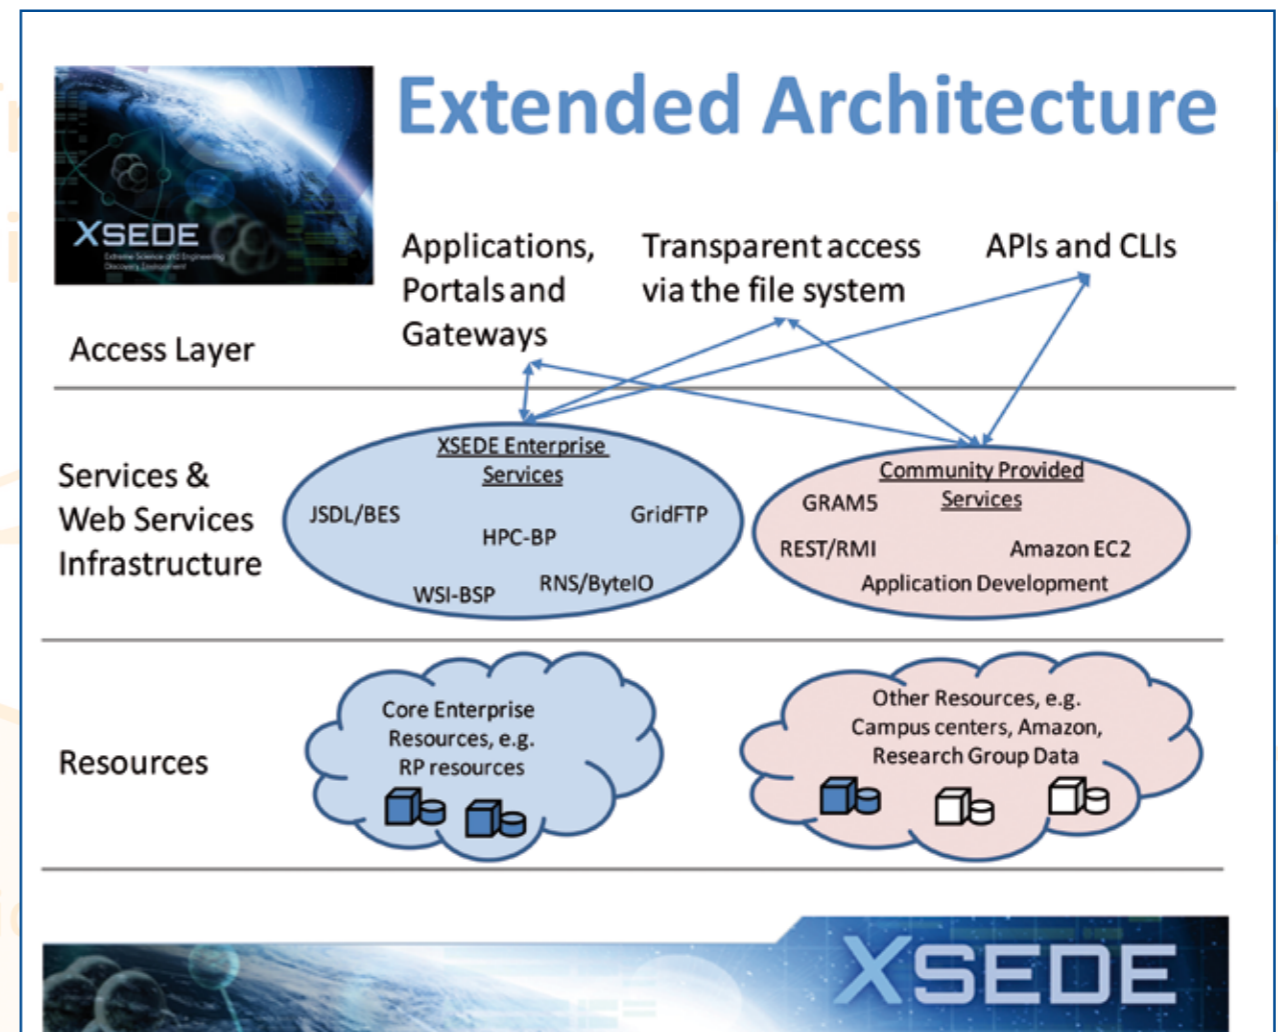


Figure 1: The current XSEDE architecture aims at providing XSEDE Enterprise Services at every major XSEDE site and optionally available Community Provided Services. The architecture will evolve over time according to end-users' needs.

• Morris Riedel
• Bastian Demuth

lightweight Genesis II services, UNICORE meets all the security requirements of modern High Performance Computing centres and provides extensive support for their highly specialized hardware as well as their varying batch systems.

Infrastructure Vision

The XSEDE architecture envisions deploying UNICORE as part of the XSEDE Enterprise Services at major US high performance centres whereas Genesis II will be used for integrating smaller computer systems such as desktop PCs in order to provide interoperability with campus Grids across the country. The

resulting infrastructure is expected to cover both high performance and high throughput computing, thus enabling innovative research and discovery requiring both types of parallel computations. Moreover, collaboration between American and European scientists will be easier than ever: UNICORE will also be deployed on the European Grid Infrastructure (EGI) and is already installed on many of the systems forming the infrastructure of the European supercomputing project PRACE.

Jülich
Supercomputing
Centre