

# Ian Kelley

---

## Research Statement

My main areas of research interest are in distributed computing and network security. I am a member of the Distributed & Scientific Computing group at Cardiff University, with my expertise lying mainly in peer-to-peer networking and high-performance/Grid computing. As I have become involved in developing distributed data management solutions for volunteer computing, I have taken a renewed interest in virtual machines and Cloud technologies. I am especially curious to see how VM technologies can be used in the scientific domain as a viable computation platform for many loosely coupled and parallel applications that demand high scalability. In addition to distributed systems and networking, I have a strong interest in collaborative environment and user-interaction technologies, such as portals and mobile devices.

## Background

I started my research career in 1999 at the Max Planck Institute for Gravitational Physics (AEI) in Germany. During my time at the AEI, I worked mainly in high-performance computing. First, on the Cactus Computational Toolkit, where I developed management tools and reporting services. In 2002, I began work on a large EU Grid project entitled “Gridlab,” where I designed user-facing portal applications for the numerical relativity group and was involved in Grid infrastructure setup and new development.

In late 2003, when the head of the AEI’s numerical relativity group, Edward Seidel, accepted a new position as Director of Louisiana State University’s (LSU) Center for Computation & Technology (CCT), I was offered the opportunity to move to Louisiana to help start the Center by leading a new Collaborative Environments group. I accepted and formed a team, consisting of myself, a part-time researcher, and between 4 and 7 graduate students. I focused my work on developing new ways to bridge the gap between end-users and the services and tools that are being developed within the research community.

CCT afforded me a unique opportunity to work with many different user groups, ranging from the technically savvy numerical relativists I had worked with for many years at the AEI, to chemists, and coastal modelers. Through discussions with these varying groups about their scientific workflows and processes, it became apparent that much could be done in the computer science area to integrate their systems, and streamline processes. CS efforts in these areas benefits not only computer science research by broadening the field and application domains, but also enables new kinds of science and data analysis for other scientists. I enjoyed the aspect of being able to “enable” non-CS communities with new technologies that distanced them from the computer science details. I believe computer science can play a key role in obfuscating underlying technological details from our users, letting them focus more on their respective areas of expertise, and less on the intricacies of how the infrastructure functions.

After many years working as a researcher, I wanted to complete a PhD in computer science to further my career and pursue a faculty-track position. I moved to the UK, where I accepted a student position at Cardiff University, while maintaining a part-time role at CCT. At CU, I focused my thesis on finding new ways to distribute large datasets in wide area networks, a key problem confronting data-intensive applications wishing to leverage volunteer compute resources.

## Research Interests

My PhD thesis is narrowly focused on distributed data management for volunteer computing environments. However, I am motivated to expand my research towards wider applications that make use of resources at the “edge of the network” (e.g., P2P). I believe that peer-to-peer technologies could be further leveraged to facilitate greater levels of resource sharing, ranging from their current use as distributed file storage and volunteer computing-style processor sharing, to broader applications such as P2P proxies and distributed web hosting and routing.

In addition to further P2P research, I am curious to explore virtual machine technologies. Specifically, I am interested to see how VMs might be applicable not only as fault tolerant solutions and scalability enablers, but also for service cloning and running scientific applications such as parameter sweeps and task farming, where the system could spawn off clones to pursue alternative paths. Another VM topic that I find intriguing is whether VMs could easily be employed as both high and low-level reusable services, and then tied into complex workflows. How will this differ from the Grid research that has been undertaken in the last decade to integrate Grid systems and provide common service interfaces?

# Ian Kelley

---

## Grants and Funding

During my time at the Cardiff School of Computer Science & Informatics (Cardiff University), I have been very proactive in pursuing grant funding and developing collaborations with international researchers. I have had the opportunity to be involved in the conceptualization, writing, and subsequent management of several research projects. They are as follows:

- 2010 – 2012: “European Desktop Grid Initiative” (EDGI)
  - 2.4M€ (total) European Union Framework 7 infrastructure project.
  - 10 partners and four subcontractors. Project coordinator: MTA SZTAKI.
  - I helped write the proposal and currently serve as the co-manager/co-PI (for Cardiff U.) under Ian Taylor
  - My responsibilities include overseeing of the project’s data management tasks, CU’s deliverables, monitoring our budget, managing student activities, attending project meetings and workshops, and representing Cardiff University at EU reviews and project board meetings.
- 2008 – 2010: “Enabling Desktop Grids for e-Science” (EDGeS)
  - 2.9M€ (total) European Union Framework 7 infrastructure project.
  - 9 partners and eight subcontractors. Project coordinator: MTA SZTAKI.
  - I helped write the proposal and was the co-manager/co-PI (for Cardiff U.) under Ian Taylor.
  - I was responsible for the Data Management Joint Research Activity (JRA3), one of three research work-packages in the project. I was also in charge of Cardiff’s deliverables, monitoring CU’s budget, managing several students and researchers, and coordinating partner activities in JRA3. I was CU’s representative to project meetings and workshops, and I presented for JRA3 at both EU project reviews.
- 2007 – 2008: “WHIP: The myExperiment workflow companion.”
  - £177K Open Middleware Infrastructure Institute (OMII UK) project.
  - I co-wrote the proposal with PI Ian Taylor, and acted as co-manager and advisor.
  - The project was a fusion of Cardiff University’s workflow engine “Triana” and portal/collaborative technologies. I was experienced in portal development from my time developing *GridSphere*, and helped oversee the project and provided guidance.
- 2006 – 2007: “GGF Grid Portal Standardization: Towards providing common portal services for Grid operations.”
  - £4,500 in combined GridNet2 grants I received for travel expenses to attend GGF/OGF and contribute to standardization. Mainly in the Grid Computing Environments (GCE-RG) research group.
- 2006: Global Grid Forum 17 (GGF17) Student Travel Scholarship
  - Travel scholarship I applied for and received to attend the Global Grid Forum (GGF) in Tokyo

In addition to the list above, I also was also involved in CoreGRID, where I proposing a data management research topic and applied for several successful CoreGRID travel grants to bring fellow researchers and students to Cardiff (for several months at a time) and send one of Cardiff University’s PhD students to Portugal.

Note that the list above depicts only the successfully funded grants I have directly contributed to in a significant way. During my time at Louisiana State University and Cardiff University I have also been involved in the design, writing, and submission of several other grant proposals.