Cloud Paradigms and Practices for CDS&E

Manish Parashar, Moustafa AbdelBaky, Ivan Rodero, Aditya Devarakonda
Cloud and Autonomic Computing Center (CAC)
Rutgers Discovery Informatics Institute (RDI^2)
Rutgers University, Piscataway, NJ
{parashar, deesha, irodero, aditya08}@cac.rutgers.edu

Cloud computing has emerged as a dominant paradigm that has been widely adopted by enterprises. Clouds provide on-demand access to computing utilities, an abstraction of unlimited computing resources, and support for on-demand scale up, scale down and scale out. Furthermore, dynamically federated "Cloud-of-Clouds" infrastructure can support heterogeneous and highly dynamic applications requirements by composing appropriate (public and/or private) Cloud services and capabilities. Clouds are also rapidly joining high performance computing system, clusters and Grids as viable platforms for scientific exploration and discovery. As a result, understanding application formulations and usage modes that are meaningful in such a hybrid infrastructure, and how application workflows can effectively utilize it, is critical.

In this whitepaper, we explore how a hybrid HPC/Grid + Cloud cyber infrastructure can be effectively used to support real-world science and engineering applications. Specifically we explore three key usage modes: ***HPC in the Cloud***, ***HPC plus Cloud***, and ***HPC as a Cloud***. This whitepaper discusses these usage modes in detail presenting illustrative scenarios in each case and outlining benefits, limitations and research challenges.

Table of Contents

This work has been submitted to the IEEE for possible publication. Copyright may be transferred without notice, after which this version may no longer be accessible.

1. Introduction

Cloud computing is revolutionizing the enterprise world, much as the Internet did not so long ago. Clouds are fundamentally changing how enterprises think about IT infrastructure, both internally and externally, by providing on-demand access to always-on computing utilities, an abstraction of unlimited resources, a potential for scale-up, scale-down and scale-out as needed, and for IT outsourcing and automation. Finally Clouds provide a usage-based payment model where users essentially "rent" virtual resources and pay for what they use. Underlying these Cloud services are typically consolidated and virtualized data centers that exploit economies of scale to provide attractive cost-benefit ratios. In spite of being in its early stages, Cloud computing is already reshaping the IT world – in fact, according to The Wall Street Journal, four out of five businesses are moving or planning to move some of their business functions to Cloud services. A recent report by Gartner estimates that Cloud services will be a $150 Billion industry by 2015 [4].

At the same time that Cloud Computing is redefining IT, extreme data and compute scales are transforming science and engineering research by enabling new paradigms and practices – those that are fundamentally information/data-driven and collaborative. Complex Computational and Data-enabled Science and Engineering (CDS&E) applications are providing unprecedented opportunities for understanding and managing natural and engineered systems, and providing unique insights into complex problems. Recognizing this data and compute driven transformation of science and engineering in the 21st Century, the NSF is rethinking the national Cyber infrastructure as part of CIF21 [2].

It is expected that Cloud services will join more traditional research cyber infrastructure components, such as high performance computing system, clusters and Grids as part of CIF21, in supporting scientific exploration and discovery. Analogous to their role in enterprise IT, Clouds can enable the outsourcing of many of the mundane and tedious aspects research and education, such as deploying, configuring and managing infrastructure, and enable scientists to focus on the science. Cloud services and the associated standardization can also improve productivity, facilitate the sharing of research results, and enable the reproducibility of associated computations. Furthermore, and once again analogous to the enterprise space, Clouds can democratize access to computational and data resources (by providing access to researchers who don't have adequate local infrastructure), which has been shown to significantly impact research productivity [5]. In fact, a recent survey of DoE users conducted by the Magellan team found that the top motivations for users' movement to the Cloud was ease of access to computing resources (cited by 79%), the ability to control software environments (59%), and the ability to share the setup of software and experiments with peers (52%) [3].

However, it is also critical to look beyond the benefits of outsourcing and understand application formulations and usage modes that are meaningful in a hybrid HPC/Grid + Cloud cyber infrastructure – for example, how emerging data and compute intensive application workflows can effectively utilize these resources, and also how this hybrid cyber infrastructure can enable new practices in science and engineering. This whitepaper explores such usage modes, and how a hybrid HPC/Grid + Cloud cyber infrastructure can be effectively used to support real-world science and engineering applications. We also discuss our initial exploration of different usage modes presenting illustrative scenarios in each case, and outlining benefits, limitations and research challenges.

The rest of this whitepaper is organized as follows. Section 2 introduces Cloud computing, its benefits, and current challenges. Section 3 discusses the different usage modes for Cloud computing in CDS&E. Illustrative scenarios are then presented in Section 4. Finally the whitepaper concludes in Section 5, and outlines a research agenda.

2. A Brief Introduction to Cloud Computing

Cloud computing aims to transform the vision of computing as a utility into a reality. In general, a Cloud can be defined as a set of network enabled on-demand IT services, scalable and with associated QoS guarantees, which can be accessed in a simple and pervasive way as a utility [32]. Cloud computing offers elasticity to applications by providing on-demand access to resources over the Internet through a pay-per-use pricing model. Cloud computing provides a layered set of service abstractions (software, middleware and infrastructure) as shown in Figure 1.

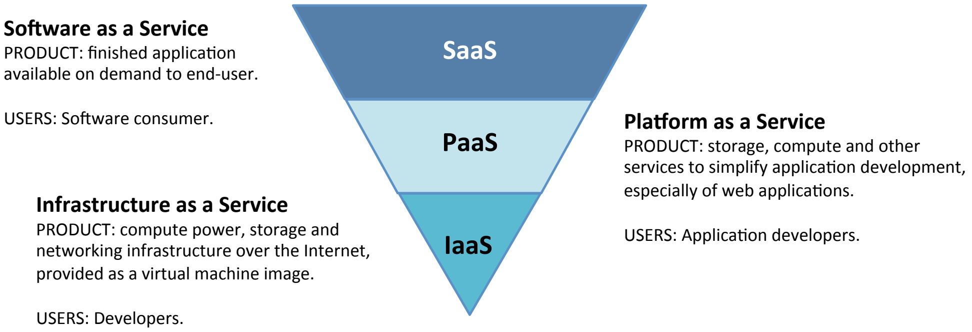

Figure 1: Cloud layered architecture (Source: An introduction to Cloud computing by Robert Grossman [18])

At the lowest level, Infrastructure as a Service (IaaS) provides raw infrastructure as a Service, where users can rent bare virtual machines to increase capacity, and customize them as needed. Platform as a Service (PaaS) provides the next abstraction layer of Cloud computing, where users can rent Cloud stacks that they can use to develop and deploy their customized Cloud services. Finally, Software as a Service (SaaS) provides the highest level Cloud abstraction, where users can rent and use complete software solutions, which are hosted in the Cloud.

Cloud services thus represent a new paradigm for computing that is based on an easy-to-use as a Service abstraction, on-demand access to computing utilities, on-demand scale-up/down/out, and a usage-based payment model. Cloud computing is drastically impacting price/performance behavior and trade-offs for a wide range of applications and IT services. From a provider perspective, Cloud abstraction support building larger data centers that exploit unprecedented economies of scale [6] and lead to consolidation, better utilization and efficiencies. From a users' perspective, Clouds largely eliminate the startup costs associated with setting up a data center, provide cost associativity benefits (i.e., 1,000 computers for 1 hour has same price as 1 computer for 1,000 hours), and provide a resilient infrastructure that can be easily configured to the user's preference. Furthermore, hybrid Clouds that federate different Clouds services can match complex applications requirements, allowing such applications to be deployed easily and effectively. However, in spite of the enormous potential impact of Cloud computing, there are still several challenges that are preventing a full-scale adoption of the Cloud model. Privacy, security, service level agreements, regulatory models, and reliability are a few examples of these issues that need to be addressed immediately and simultaneously. We would like to note that given the project size of the Cloud market, industry is investing tremendous amounts of effort and resources to address many of these challenges, and academic research agendas should consider complementing these efforts. Nonetheless, there remain many such challenges.

3. Science & Engineering and Clouds

As noted earlier, Clouds are rapidly joining high-performance computing system, clusters and Grids as viable platforms for scientific exploration and discovery, and understanding application formulations and usage modes for science and engineering applications is important.

***Understanding CDS&E applications and their requirements*:** The wide range of CDS&E applications have been broadly classified, based on their requirements and execution behaviors, into high performance computing (HPC), high throughput computing (HTC), and many task computing (MTC) [30]. HPC applications are tightly coupled with large amounts of inter-processor communication, and typically require large amounts of computing power for short periods of time. Thus, fast interconnects among processors is needed in addition to the large scale needed by HTC applications. On the other hand, HTC applications are usually loosely coupled, where communication between processors is limited or non-existent. HTC applications also require large amounts of computing, but for much longer times (months or years, rather than hours or days). Finally, MTC applications are a hybrid of the above two classes. MTC applications can consist of loosely coupled tasks, where each of these tasks is a tightly coupled application that runs for a short period of time (i.e. seconds or minutes). MTC applications require fast interconnects, can potentially scale to extremely large scales (Peta & Exa scale), and often involve dynamic workflows that require elastic resources.

***CDS&E applications & Clouds*:** Cloud computing services can support CDS&E applications in multiples ways. It can provide a platform for applications, for example when local infrastructure is not available. It can also supplement existing platforms to provide additional capacity or complementary capabilities to meet heterogeneous or dynamic needs. For example, Clouds can service as accelerators, or provide resilience to scientific workflows by moving the execution of the workflow on alternative or fewer resources when a failure occurs. The simplicity of the Cloud abstraction can alleviate some of the problems scientific applications face in the current HPC environment. For example, the increasingly important and growing MTC applications can benefit from the ease of use, an abstraction of elastic and/or readily accessible resources and the ability to easily scale up, down or out. Finally, Cloud computing can not only help scientists address today's problems more effectively, but also allow them to explore new ways of formulating their application using the abstraction of on-demand access to elastic resources.

In the remainder of this whitepaper we explore these ideas in more details and investigate how the Cloud abstraction can be effectively used to support real-world scientific & engineering applications. Specifically we explore three key usage modes as illustrated in Figure 2. Each of these approaches has its benefits and limitations, and we believe a clear understanding of the usage modes and associated application scenarios is essential. In particular, we explore how using Cloud models and abstractions (i.e., elasticity) can be used to build *HPC plus Cloud* and *HPC as a Cloud* that are complementary to the *HPC in the Cloud* approach. ***Note that in this whitepaper we use HPC to refer to more traditional computing resources, such as the production resources that are part of XSEDE, OSG, and other academic computing centers (including clusters, grids, tightly-coupled systems, supercomputers, etc.) as well as their usage modes and applications. Also note that while we focus on computing Clouds, the discussion also translates to data Clouds, such as the Open Science Data Cloud project [7], which aim to provide scientific and engineering data and data analytics as a service.***

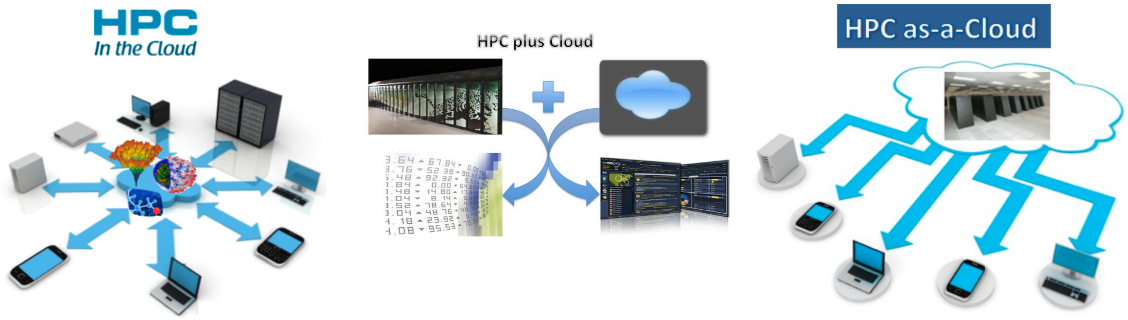

Figure 2: Hybrid HPC/Grid + Cloud usage modes for supporting real-world science and engineering applications

1. *HPC in the Cloud*, where researchers outsource entire applications to current public and/or private Cloud platforms.
2. *HPC plus Cloud*, focused on exploring scenarios where Clouds can complement HPC/Grid resources with Cloud services to support science and engineering application workflows, for example, to support heterogeneous requirements, unexpected spikes in demand, etc.
3. *HPC as a Cloud*, focused on exposing HPC/Grid resources using elastic on-demand Cloud abstractions, aiming to combine the flexibility of Cloud models with the performance of HPC systems.

3.1 HPC in the Cloud

Current Cloud platforms can provide effective platforms for certain classes for CDS&E applications, for example HTC applications. There have been several early projects that have reported successful deployments of applications on existing Clouds [13][14][20][21]. Running these applications typically involved using virtualized commodity-based hardware, which is provisioned on-demand by commercial Cloud providers such as Amazon EC2 or Microsoft Azure.

A recent technical report by G. Fox and D. Gannon [15] has extensively studied running HPC applications in the Cloud. According to this study, commodity Clouds work effectively only for certain classes of HPC applications. Examples of these applications are embarrassingly parallel applications that analyze independent data or spawn independent simulations that integrate distributed sensor data, science gateways and portals, or data analytics that can use MapReduce-like applications. In addition, research work by Fox et al. [16] and Juve et al. [19] show that different variants of MapReduce computations do well on current Cloud platforms such as iterative MapReduce. In general, HPC applications with minimal synchronization and minimal communications requirements, small I/O requirements and modest scales are well suited for current Cloud platforms and can be successfully outsourced to Clouds. In cases where existing application formulations are not directly suited for available Cloud platforms, alternate formulations need to be explored before these applications can benefit from Cloud services. For example, the asynchronous replica exchange [36] formulation is a novel, decentralized, asynchronous and resilient formulation of the "replica exchange" algorithm for simulating the structure, function, folding, and dynamics of proteins, and has been successfully deployed on existing Clouds.

<u>Research Challenges:</u> Several aspects of current commercial Clouds continue to limit the more general deployment of applications onto Cloud platforms. These include the capabilities and heterogeneity of the typical underlying hardware, and the lack of high-speed interconnects to support data exchanges required by these applications. In these cases, if Clouds are the platform of choice (possibly due to availability), then alternate formulations must be explored, such as in the case of the replica exchange example above. Furthermore, even in case of the suitable CDS&E applications, the cost and relative performance can become a concern [13]. For example, once again the Magellan report [33] cites that Cloud services were found to be 7 to 13 times more expensive. Finally, HPC applications based on loosely or bulk synchronizations cannot take advantage of the scalability, elasticity and fault-tolerance due to the high latency in commercial Clouds. For example, Iosup et al. [19] evaluated the performance of Amazon EC2 for scientific workloads and found that the reliability and the performance of EC2 are low. However, the performance and reliability of virtualized resources, depends heavily on the underlying hardware and network. For example, Younge et al. [34] compared the performance of various virtualization technologies against bare-metal on a high performance cluster and InfiniBand network provided by FutureGrid, and found that the performance on a virtualized HPC environment is close to that of bare-metal in this case.

3.2 HPC plus Cloud

Running CDS&E applications on more traditional HPC resources and bursting onto a Cloud when there is a spike in the demand for computing capacity, or a specific capability is needed to meet heterogeneous applications requirements (also known as Cloud bursting), is an attractive approach for heterogeneous and dynamic application workflows. This holds true especially as CDS&E application are increasingly becoming end-to-end workflows consisting of coupled simulations and integrated I/O and data analytics pipelines. A hybrid Cloud plus traditional HPC infrastructure can also support new and potentially more effective usage modes and enable new application formulations that use Clouds to achieve acceleration, resilience, or more appropriate cost/power/ performance tradeoffs.

Existing efforts that have explored such hybrid infrastructure include InterGrid [12] and meta-schedulers [10], which interconnect different Grids. There are also efforts to include Clouds, such as Amazon EC2, into integrated computing infrastructures. Buyya et al. [11] described an approach of extending a local cluster to Cloud resources using different scheduling strategies, while Ostermann et al. [29] extended a Grid workflow application development and computing infrastructure to include Cloud resources, and experimented with Austrian Grid and an academic Cloud installation of Eucalyptus using a scientific workflow application. Similarly, Vazquez et al. [31] proposed architecture for an elastic Grid infrastructure using the GridWay meta-scheduler, and extended Grid resources to Globus Nimbus. Our previous work has also explored such a hybrid-computing infrastructure, which integrated clusters with Clouds [26] and Grids with Clouds [23] and enabled on demand autonomic Cloudbursts using CometCloud [1][27].

Research Challenges: The federation of traditional HPC/Grid/Cluster resources with elastic Cloud resources provides a hybrid infrastructure with the ability to dynamically add (or remove) capacity as well as capability. While such hybrid infrastructure can support new and potentially more effective usage modes, and even enable new application formulations, its inherent heterogeneity and dynamism presents significant challenges. For example, the different resource classes available in such a hybrid infrastructure can vary significantly in their usage costs, configuration, performance, availability, runtime behaviors and the guarantees for quality of service provided. On the other hand, on the application side, we are often faced with dynamic requirements and constraints. As a result, provisioning and scheduling an appropriate mix of resources for these applications requires considering appropriate cost/performance tradeoffs. Furthermore, these requirements/constraints may change, due to a failure or an application event for example, requiring dynamic system and/or application adaptations. As manually monitoring these dynamic behaviors and requirements and enforcing adaptation can quickly become unfeasible, and autonomic management approaches become crucial for such workflows. Finally, the challenges also include programming the framework required for specifying workflow structures as well as requirements and constraints.

3.3 HPC as a Cloud

Due to the limitations of commodity Clouds in serving general HPC applications, Cloud providers realized the need to provide Cloud solutions that are built specifically for HPC applications (i.e., hardware with faster processors and interconnects). Some providers have even provided non-virtualized hardware, in order to provide bare to the bone performance that these applications require. This is commonly referred to as "HPC as a Cloud" (i.e., running HPC applications on HPC resources that are exposed as on-demand resources using Cloud computing abstractions, in order to take advantage of the Cloud model without sacrificing the HPC performance that scientific applications require). There are two main approaches used for providing HPC as a Cloud summarized in Table 1. The first approach uses large HPC systems that can be provisioned as Clouds. AbdelBaky et al. [7] illustrate this approach using the IBM Blue Gene/P supercomputer. The second approach uses small HPC clusters that can be connected together to form a large Cloud. These HPC clusters can be virtualized or non-virtualized to provide better performance. Amazon EC2 Compute Cluster, Azure, Adaptive Systems, and Platform HPC Computing are commercial providers that use the virtualized approach, while SGI and Penguin Computing on Demand are commercial providers using the physical approach. In addition, Goscinski et al. [17] proposed and validated a similar academic approach for providing HPC as a Cloud using small clusters, known as HPCynergy. While HPCynergy is a more effective approach than commercial offerings, the scale of HPCynergy is relatively small (224 cores) especially for realistic science and engineering applications.

Large HPC systems	**Small HPC clusters connected together**	
Physical Solutions	Physical Solutions	Virtualized Solutions
E.g., Blue Gene/P as a Service	E.g., Penguin on Demand, Silicon Graphics, HPCynergy	E.g., EC2 Compute Cluster, Azure, Adaptive Systems, CloudCycle

Table 1: HPC as a Cloud classification

Research Challenges: HPC as a Cloud has a clear potential to alleviate some of the problems scientific applications face in the current HPC environment, and significantly improve the usability and productivity of traditional HPC overall. However, the current HPC as Cloud offerings have several limitations.

- *Scale & Performance:* The scale of current commercial HPC as a Cloud is very small (128 processors on Amazon HPC instance vs. 294,912 processors on an IBM Blue Gene/P), and therefore cannot accommodate the requirements of CDS&E applications. In addition, although the performance of these Cloud services is better than commodity Clouds, they are still outperformed by local HPC clusters that are specifically designed to support these applications [35]. Finally, spending more resources (i.e., scaling beyond a physical cluster) does not guarantee better performance, and performance can in fact get worse (i.e., due to communication overhead), contradicting one of the main premises of Cloud computing.
- *Provisioning:* In HPC as a Cloud, provisioning refers to building a virtual layer to easily configure and access HPC resources in a "Cloud-like" manner. Building this

virtual layer can be very challenging. For instance, the layer should be general enough that it can encompass different HPC systems and architectures. However it should also be specific enough to take full advantage of the underlying resources.

- *Elasticity and On-demand Access:* Elasticity by definition is the abstraction for providing on-demand access to resources and the ability to scale up, down or out as needed. However, most HPC resources currently use a batch-queue system for running applications, which in turn prevents on-demand access to resources. Much research is needed in order to provide elasticity on top of batch-queue systems, and support on-demand access to these systems as needed.
- *Programmability:* While most web applications running on commodity Clouds are agnostic of the underlying hardware, scientific applications are often customized to the underlying hardware. For instance, maintaining data locality when running scientific applications can minimize communication overhead. Therefore, building scientific applications that are agnostic to the underlying hardware without compromising performance is essential. Finally, exposing Cloud abstractions provided by HPC resources (e.g., elasticity, data locality, and ease of use) to scientific applications is also critical.

A summary of the different usage modes and their advantages is shown in Table 2.

	Advantages	**Disadvantages**	**Typical applications**	**Sample systems**
HPC in the Cloud	Easy to use, infinite resources	Low performance, expensive	E/P, L/S, A, HTC	Amazon EC2 Microsoft Azure
HPC as a Cloud	Elasticity, on demand access	Small scale, limited QoS, medium performance	S, M, HPC, MTC	EC2CC, SGI, POD
HPC plus Cloud	Acceleration Conservation Resilience	Different prices, performance, and availability of resources	HPC + Analytics /Visualization	CometCloud, InterGrid, GridWay

E/P	Embarrassingly Parallel	M	Metaproblems
L/S	Loosely Synchronous	HPC	High Performance Computing
S	Synchronous	HTC	High Throughput Computing
A	Asynchronous	MTC	Many Task Computing

Table 2: HPC Cloud landscape

4. Initial Explorations with CDS&E and Cloud Computing

We have used the CometCloud [1] framework developed at Rutgers to explore CDS&E usage modes on existing cyber infrastructure and public Clouds. Specifically, The goal of our efforts is to explore application formulations, Cloud and hybrid HPC/Grid + Cloud infrastructure usage modes that are meaningful for various classes CDS&E application workflows. Our experiments explored a broad set of federated resources spanning existing public commodity Clouds, Grids, supercomputing resources, and local clusters. We also explored a broad set of applications classes including business intelligence, financial analytics, oil reservoir simulations, medical informatics, document management, etc. A summary of the different specific applications and platforms explored in this study is summarized in Table 3. A more detailed summary of all our experiments is presented in Appendix A. The discussion in this section is focused on summarizing our experiments and experiences related to the usages modes outlined in this whitepaper using a single ensemble application as a driver one. The rest of the section is organized as follows. An overview of CometCloud is presented in Section 4.1. The driving ensemble application is described in Section 4.2, and is followed by a detailed case study of each usage modes, *HPC in the Cloud, HPC plus Cloud, and HPC as a Cloud*, in Sections 4.3, 4.4, and 4.5, respectively.

Applications	Description	Platforms	Description
Value at Risk	Simulate stock prices at time t+1 given prices at time t	**Amazon EC2**	Amazon Elastic Compute Cloud.
Medical Image Registration	Mapping of two images of similar objects or same object	**TeraGrid**	e-Science grid computing infrastructure
Heat distribution problem using Jacobi Iteration	Calculate heat distribution inside uniform metal plate	**FutureGrid Nimbus**	Distributed test bed for Cloud/grid Open source service package
Asynchronous Replica Exchange	Sampling algorithm to preserve canonical distribution	**FutureGrid Eucalyptus**	Eucalyptus (open source Cloud computing software)
Protein Data Bank	Database of crystal structures and nuclear magnetic resonance	**Rutgers Cluster I**	Microsoft HPC cluster
IPARS + EnKF	Oil reservoir history matching application with Ensemble Kalman Filter	**Rutgers Cluster II**	Dell Cluster
Decentralized Online Clustering	Proactive self-monitoring of distributed systems	**Blue Gene /P**	IBM supercomputer
Mandelbrot	Mathematical set of points in complex plane	**Rutgers Cluster II**	Dell Cluster

Table 3: Applications and platforms supported by CometCloud

4.1 Overview of CometCloud

CometCloud [1][27] is an autonomic computing engine that enables dynamic and on-demand federation of Clouds and Grids as well as the deployment and robust execution of applications on these federated environments. It supports highly heterogeneous and dynamic Cloud/Grid infrastructures, enabling the integration of public/private Clouds and autonomic Cloudbursts, i.e., dynamic scale-out to Clouds to address extreme requirements such as heterogeneous and dynamics workloads, and spikes in demands. The CometCloud programming layer provides a platform for application development and management. It supports a range of paradigms including MapReduce, Workflow, and Master/Worker/BOT. Finally the CometCloud autonomic management layer enables the autonomic management of applications workflows as well as the federated infrastructure, and ensures that application objectives and constraints are satisfied. Specifically, the autonomic management framework can provision the appropriate mix of HPC/Grid and public/private Cloud resources based on application requirements and constraints, monitor system/application state (e.g., workload, availability, delays) and adapt the application and/or the resources (e.g., change algorithms used or re-provision resources) to respond to changing applications requirements or system state. A schematic overview of the CometCloud-based autonomic application management framework for enabling hybrid HPC platform usage modes is presented in Figure 3.

4.2 Ensemble Kalman Filter (EnKF) Applications

Ensemble applications explore large parameter spaces in order to simulate multi-scale and multiphase models and minimize uncertainty, which is achieved by running thousands of ensembles simultaneously and then assimilating their results. An ensemble filter then analyzes these results before an updated set of realizations is run. Each ensemble can be a traditional parallel HPC application that requires a varying number of processors and fast communication among processors. In addition, typically a large and varying number of ensemble members are also required to achieve acceptable accuracy, which in turn requires a very large and dynamic pool of HPC resources. The EnKF workflow presents an interesting use case due to the heterogeneous computational requirements of the individual ensemble members as well as the dynamic nature of the overall workflow. Figure 4 illustrates the variability between stages of a typical EnKF based simulation. EnKF is a recursive filter that can be used to handle large, noisy data; data in this case are the results and parameters from ensembles of reservoir models that are sent through the filter to obtain the “true state” of the data. Since the reservoir model varies from one ensemble to another, the runtime characteristics of the ensemble simulation are irregular and hard to predict. Furthermore, during execution, if real historical data is available, all the data from the different ensembles at that simulation time must be compared to the actual production data before the simulations are allowed to proceed. This translates into a global synchronization point for all ensemble-members in any given stage, which can present challenges. The end-to-end application consists of several stages, and in general at each stage the number of models generated varies in size and duration.

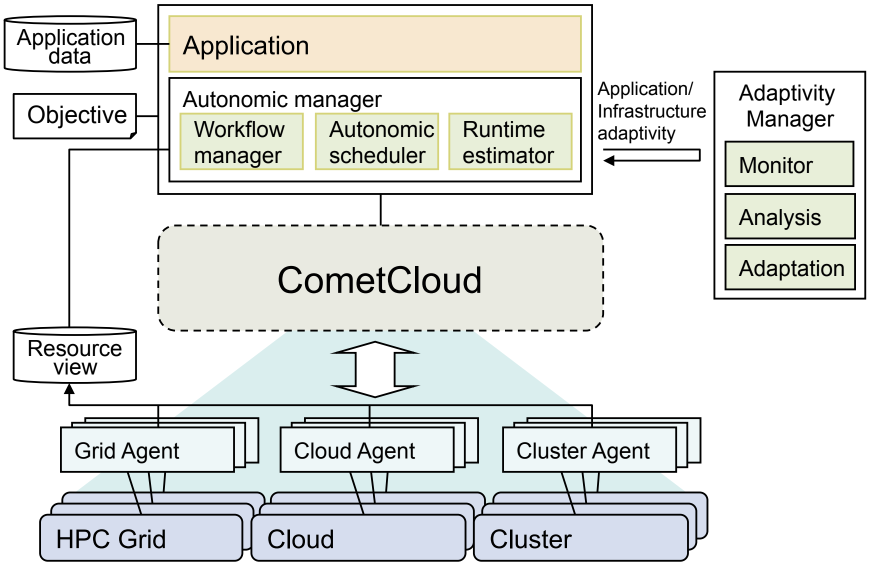

Figure 3: Architectural overview of the autonomic application management framework and supported platforms

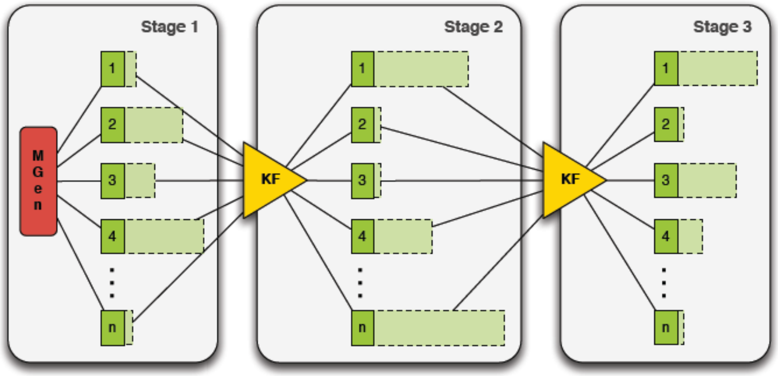

Figure 4: Stages of a typical ensemble Kalman filter based simulation

4.3 HPC plus Cloud using CometCloud

We have used CometCloud to explore meaningful usage modes for a hybrid *HPC plus Cloud* infrastructure. In particular, we used EnKF workflow complementing TeraGrid resources with Amazon EC2 public Cloud instances [23]. We explored 5 different usage modes: (1) *acceleration* – using Clouds as accelerators to reduce the application time to completion, for example given budget constraints, (2) *conservation* – using Clouds to conserve HPC allocations, within the appropriate runtime and budget constraints, and (3) *resilience* – using Clouds to handle unexpected situations such as an unanticipated HPC downtime, inadequate allocations, unanticipated queue delays or failures of working nodes, while meeting user objectives, (4) *Cloud bursting* – using Clouds to perform the actual computation in the Cloud if this is more effective than moving the data to HPC resources, and (5) *analytics/visualization* – using Clouds to perform data analytics or visualization at the same time that complex simulations are run in HPC resources.

4.4 HPC in the Cloud using CometCloud

In this investigation we studied the performance [24] and performance fluctuation [22] of the EnKF workflow on EC2 and FutureGrid. We observed that the underlying hardware impacts performance as measured by time-to-completion, as well as cost. We concluded that reducing the fluctuation on a single instance is a mandatory step to expand to large core counts and running MPI applications efficiently across several instances. In addition to exploring how CometCloud can support a range of well-suited application on public and private Cloud infrastructure (e.g., image analysis, value-at-risk computations, etc.), we also explored the Replica Exchange Molecular Dynamics application [36]. This application, which is not traditionally well suited for Clouds, can be reformulated to an asynchronous and resilient application that can be effectively run on public and private Clouds using the messaging and coordination abstractions provided by CometCloud.

4.5 HPC as a Cloud using CometCloud

In an early experiment we explored how a Cloud abstraction can be effectively used to provide a simple interface for current HPC resources and support real-world applications [7]. In particular, we experimentally validated the benefits of the Cloud paradigm, such as ease of use and dynamic allocation, and their application to supercomputers, specifically, on an IBM Blue Gene/P system. The CometCloud-based framework essentially transformed Blue Gene/P into an elastic Cloud, bridged multiple Blue Gene/P systems to create a larger HPC federated Cloud, and supported dynamic provisioning. The exercise demonstrated the ease-of-use of the elastic Cloud abstraction, and its effectiveness in improving utilization. This experiment was demonstrated at the 4th IEEE SCALE Challenge, and was awarded the first place. During the experiment, Blue Gene/P resources varied from 640 to 22,016 processors, spanning across two Blue Gene systems in two continents.

5. Research Agenda for CDS&E and Clouds

Aggressive Cloud computing technology development has resulted in many different classes of Cloud services that provide attractive solutions for different types of business applications. It is also clear from previous research and our own experiences summarized in this whitepaper, that there are real benefits of using Clouds and Cloud computing abstractions as part of a hybrid cyber infrastructure to support CDS&E, for example, to simplify the deployment of applications and the management of their execution, improve their efficiency, effectiveness and/or productivity, and provide more attractive cost/performance ratios. Furthermore, Clouds and Cloud computing abstractions can support new classes of algorithms and enable new applications formulations, which can potentially revolutionize CDS&E research and education. However, before CDS&E can fully realize the potential benefits of a hybrid cyber infrastructure that integrates Cloud services, several research issues remain as outlined in the previous sections and summarized below. Although we envision new usage modes, our research agenda leverages the existing industry standards and best practices that have been defined during the last decade.

5.1 Algorithms and Application Formulations for Clouds

A key attribute of Clouds is on-demand access to elastic resources – i.e., applications programmatically access more or less resources as they evolve, to meet changing needs. Such a capability can have a significant impact on how algorithms are developed and applications are formulated. For example, the execution of an application no longer has to constraint itself to a fixed set of resources that are available at runtime and can grow or shrink based on the demand of the science – the science can drive the scale and type of resources involved, based on for example, the levels of refinement required to resolve a solution feature, or the number of ensembles that need to be run to quantify the uncertainty in a solution, or the type of online analytics services that need to be dynamically composed into the application workflow.

In addition to the usage modes and application scenarios discussed earlier in this whitepaper, understanding how CDS&E applications can effectively utilize Clouds and Cloud abstractions as part of the a hybrid cyber infrastructure, to enable new practices and levels of scientific insights remains a research challenge. Research is needed to explore the meaningful science, engineering and business application scenarios that can take advantage of such hybrid infrastructure. For example, a meaningful HPC plus Cloud use case may consist of simulations with online data analytics/visualization needs. In such a scenario, exposing the ability to modify goals/configurations based on data analytics feedback to the user will be critical to ensuring impact on the science. For example, in data-intensive computations the use of feature tracking might allow the scientist to adjust application parameters based on the analysis of meaningful features using a public Cloud, where the analysis can be performed in a timely manner due to shorter resource provisioning times compared to a high-end HPC system. We believe that such meaningful scenarios will exist in all areas of CDS&E.

5.2 Programming Systems and Abstractions

One of the key research challenges is developing appropriate programming abstractions and language extensions that can enable CDS&E application to simply and effectively take advantage of the elastic access to resources and services during application formulation. Furthermore, it may be necessary to define constraints (for example, budgets, data privacy, performance, etc.) to regulate the elasticity, and the programming abstractions may provide support for expressing these constraints so that they can be enforced during execution. Similarly, such annotations can also define possible adaptations, which could then be used to increase performance, manageability and overall robustness of the application. For example, an application may dynamically increase assigned resources in order to increase the resolution of a simulation under certain convergence constraints, or modify convergence goals to avoid failure or guarantee completion time. The Cloud service models can also lead to interesting services specialized to CDS&E that provide entire applications or application kernels as a Service (i.e., SaaS). Furthermore, and arguably more interestingly, it can also export specialized platforms for science as a Services (i.e., PAAS), which encapsulate elasticity and abstract of the underlying hybrid cyber infrastructure. In return, the scientists are only required to provide, core kernels, meaningful parameters, and basic configurations. For example, our own research is exploring the idea of EnKF as a Service.

5.3 Middleware stacks, management policies, economic models

Middleware services will need to support new CDS&E applications formulations and services. A key research aspect will be the autonomic management and optimization of application execution through cross-layer application/infrastructure adaptations. It will be essential for the middleware services to be able to adapt to the application's behavior as well as system configuration, which can change at run time, using the notion of elasticity at the application and workflow levels. Furthermore, appropriate services are necessary to be able to provision different types of resources on demand. For example, if we target HPC as a Cloud and HPC plus Cloud approaches on national cyber-infrastructure such as XSEDE, Open Science Grid and FutureGrid along with commercial Clouds such as Amazon EC2 or Microsoft Azure, autonomic provisioning and scheduling techniques, including Cloud bursting will be necessary to support the usage modes outlined in this whitepaper. Finally, monitoring, online data analytic for proactive application/resource management and adaptation techniques will be essential as the scale and complexity of both the applications and hybrid infrastructure grows.

5.4 Security policies and mechanisms

Clouds tend to emphasis the need for quality security mechanisms due to the sharing of storage and computing. One research area investigates hybrid architectures with algorithms, including low cost but non privacy preserving mechanism running on an intrinsically secure private Clouds, and time consuming but privacy preserving ones executing on a public Cloud. In addition to basic security for computing and storage, another research area is privacy preserving search with the elegant but time consuming concept of Homomorphic Encryption, which allows encrypted data to be searched by encrypted queries. Some key research areas were highlighted at the recent NSF Workshop on Security for Cloud Computing [6].

5.5 Deployment/Transition to practice

There are many important standard activities, from those specifying the basic virtual machine structure to higher-level standards defining the PaaS/SaaS environment. Although there is some support for these standards, such as OCCI (from OGF) in OpenNebula and OpenStack, this area is still under development with NIST and IEEE are playing leadership roles. In addition, substantial progress needs to be made in order enable procurement of services through capped purchase orders, or subcontracts; administration of sub-accounts and delegation of resources and authorities; reporting and alarming/limiting; and several others. Finally, the development and/or modification of codes adapted to the Cloud environment require a unique set of skill, and appropriate educational and training structures are required

5.6 Conclusion

Clouds are becoming an important part of production computational environments and are joining high performance computing system, clusters and Grids as viable platforms for supporting Computational and Data-enabled in Science and Engineering. The goal of this whitepaper was to understand application formulations and usage modes that are meaningful in such a hybrid infrastructure, and how application workflows can effectively utilize it, is critical. In this whitepaper, we explored how a hybrid HPC/Grid + Cloud cyber infrastructure can be effectively used to support real-world science and engineering applications. Specifically we presented three key usage modes: ***HPC in the Cloud***, ***HPC plus Cloud***, and ***HPC as a Cloud***. This whitepaper discussed these usage modes, discussed initial explorations and experiences, and outlined a research agenda for CDS&E with Cloud computing.

Acknowledgements

The research presented in this work is supported in part by National Science Foundation (NSF) via grants numbers IIP 0758566 and DMS-0835436 and through XSEDE resources under award TG-CCR110035, by the Department of Energy ExaCT Combustion Co-Design Center via subcontract number 4000110839 from UT Battelle and via the grant numbers DE-SC0007455 and DE-FG02-06ER54857, and by an IBM Faculty Award, and was conducted as part of the NSF Cloud and Autonomic Computing (CAC) Center at Rutgers University.

References

[1] CometCloud. http://www.cometCloud.org.

[2] Cyber infrastructure Framework for 21st Century Science and Engineering (CIF21), National Science Foundation, http://www.nsf.gov/funding/pgm_summ.jsp?pims_id=504730

[3] “Evaluating Cloud Computing for HPC Applications,” in NERSC Brown Bag, October 2010. http://www.nersc.gov/assets/Events/MagellanNERSCLunchTalk.pdf.

[4] Gartner report, “Forecast: Public Cloud Services, Worldwide and Regions, Industry Sectors, 2009- 2014.” The report is available on Gartner’s website at http://www.gartner.com/DisplayDocument?id=1378513

[5] IBM Research Doubles Its Productivity with Cloud Computing. Case Study QuickView. http://www-935.ibm.com/services/in/cio/pdf/oic03013usen.pdf.

[6] NSF Workshop on Security for Cloud Computing http://illinois.edu/blog/view/695/66281?count=1&ACTION=DIALOG

[7] Open Science Data Cloud. http://www.opensciencedatacloud.org/

[8] The economics of the Cloud. http://economics.uchicago.edu/pdf/Harms_110111.pdf.

[9] M. AbdelBaky, H. Kim, M. Parashar, K.E. Jordan, H. Jamjoom, V. Sachdeva, Z.Y. Shae, J. Sexton, G. Pencheva, R. Tavakoli, and M.F. Wheeler, “Scalable Ensemble based Oil Reservoir Simulations using Blue Gene/P as-a-Service,” White paper, SCALE 2011, the 4th IEEE International Scalable Computing Challenge was held on May 26th 2011 in conjunction with the 11th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing at Newport Beach CA, USA.

[10] N. Bobroff, L. Fong, S. Kalayci, Y. Liu, J. C. Martinez, I. Rodero, S. M. Sadjadi, and D. Villegas, “Enabling interoperability among meta- schedulers,” in CCGRID ’08: Proceedings of the 2008 Eighth IEEE International Symposium on Cluster Computing and the Grid. Washington, DC, USA: IEEE Computer Society, 2008, pp. 306–315.

[11] M. D. de Assuncao, A. di Costanzo, and R. Buyya, “Evaluating the cost- benefit of using Cloud computing to extend the capacity of clusters,” in HPDC ’09: Proceedings of the 18th ACM international symposium on High performance distributed computing. New York, NY, USA: ACM, 2009, pp. 141–150.

[12] M. D. de Assuncao, R. Buyya, and S. Venugopal, “Intergrid: a case for internetworking islands of grids,” Concurr. Comput. : Pract. Exper., vol. 20, no. 8, pp. 997–1024, 2008.

[13] E. Deelman, G. Singh, M. Livny, B. Berriman, and J. Good. The cost of doing science on the Cloud: the montage example. In SC ’08: Proceedings of the 2008 ACM/IEEE conference on Supercomputing, pages 1–12, Piscataway, NJ, USA, 2008. IEEE Press.

[14] C. Evangelinos and C. Hill. Cloud computing for parallel scientific hpc applications: Feasibility of running coupled atmosphere-ocean climate models on amazon’s EC2. In Cloud Computing and its Applications (CCA-08), Chicago, IL, USA, Oct 2008.

[15] G. Fox, D. Gannon, “Cloud Programming Paradigms for Technical Computing Applications,” technical report.

[16] G. Fox, “Data Intensive Applications on Clouds”, Keynote at The Second International Workshop on Data Intensive Computing in the Clouds (DataCloud-SC11) at SC11 November 14, 2011.

[17] A. Goscinski, M. Brock, P. Church, “High Performance Computing Clouds, Cloud Computing: Methodology, Systems and Applications, Wang et. al. (Ed.), CRC Press, ISBN 9781439856413

[18] R. Grossman, “An introduction to cloud computing,” August 6, 2009 in Philadelphia

[19] A. Iosup, S. Ostermann, N. Yigitbasi, R. Prodan, T. Fahringer, and D. Epema, “Performance analysis of Cloud computing services for many-tasks scientific computing”, IEEE TPDS 2011 Los Alamitos, CA, 2011.

[20] Y. Kang, G. Fox, “Performance Evaluation of MapReduce applications on Cloud Computing Environment”, FutureGrid GDC2011, 2011 Conference on Grid and Distributed Computing December 8-10, 2011 in Jeju Grand Hotel, Jeju Island, Korea.

[21] K. Keahey and T. Freeman. Science Clouds: Early experiences in Cloud computing for scientific applications. In Cloud Computing and Its Applications 2008 (CCA-08), Chicago, IL, USA, Oct 2008.

[22] Y. el-Khamra, H. Kim, S. Jha, and M. Parashar, "Exploring the Performance Fluctuations of HPC Workloads on Clouds," 2nd IEEE international conference on Cloud computing technology and science (CloudCom), Indianapolis, Nov. 30-Dec.3, 2010.

[23] H. Kim, Y. el Khamra, S. Jha, and M. Parashar, “An autonomic approach to integrated hpc grid and Cloud usage,” in e-Science, 2009. e-Science ’09. Fifth IEEE International Conference on, 2009, pp. 366–373.

[24] H. Kim, Y. el-Khamra, S. Jha, and M. Parashar, “Exploring Adaptation to Support Dynamic Applications on Hybrid Grids-Clouds Infrastructure,” 1st Workshop on Scientific Cloud Computing, in conjuction with the ACM International Symposium on High Performance Distributed Computing (HPDC), Chicago, Illinois, June 20-25, 2010.

[25] H. Kim, Y. el-Khamra, I. Rodero, S. Jha, and M. Parashar. Autonomic management of application workflows on hybrid computing infrastructure. Sci. Program. 19(2-3):75-89, 2011.

[26] H. Kim, S. Chaudhari, M. Parashar, and C. Marty, “Online risk analytics on the Cloud,” in Cluster Computing and the Grid, 2009. CCGRID ’09. 9th IEEE/ACM International Symposium on, May 2009, pp. 484–489.

[27] H. Kim and M. Parashar, “CometCloud: an autonomic Cloud engine”, Chapter 10, pp. 275-297, Cloud Computing: Principles and Paradigms, Wiley, 2011.

[28] H. Kim, M. Parashar, D. J. Foran and L. Yang, “Investigating the use of Cloudbursts for high-throughput medical image registration,” the 10th IEEE/ACM International Conference on Grid Computing (GRID2009), Banff, Canada, Oct. 2009.

[29] S. Ostermann, R. Prodan, and T. Fahringer, “Extending grids with Cloud resource management for scientific computing,” in Grid Computing, 2009 10th IEEE/ACM International Conference on, 2009, pp. 42 –49.

[30] I. Raicu, Z. Zhang, M. Wilde, I. Foster, P. Beckman, K. Iskra, B. Clifford. “Towards Loosely-Coupled Programming on Petascale Systems”, IEEE/ACM Supercomputing 2008.

[31] C. Vazquez, E. Huedo, R. Montero, and I. Llorente, “Dynamic provision of computing resources from grid infrastructures and Cloud providers,” in Grid and Pervasive Computing Conference, 2009. GPC ’09. Workshops at the, may 2009, pp. 113 –120.

[32] L. Wang, J. Tao, M. Kunze, A.C. Castellanos, D. Kramer, and W. Karl, 2008. "Scientific Cloud Computing: Early Definition and Experience," High Performance Computing and Communications, 2008. HPCC '08. 10th IEEE International Conference on, pp. 825-830.

[33] K. Yelick, S. Coghlan, B. Draney, R. S. Canon, et al., The Magellan Report on Cloud Computing for Science. U.S. Department of Energy Office of Advanced Scientific Computing Research (ASCR), December, 2011.

[34] J. Younge, R. Henschel, et al., "Analysis of virtualization technologies for high performance computing environments", IEEE CLOUD 2011, Washington, DC, IEEE, 2011.

[35] Y. Zhai, M. Liu, J. Zhai, X. Ma, and W. Chen, “Cloud versus in-house cluster: evaluating Amazon cluster compute instances for running MPI applications,” In State of the Practice Reports (SC '11), ACM, New York, NY, USA, Article 11, 10 pages.

[36] R. L. L. Zhang, M. Parashar and E. Gallichio, “Salsa: Scalable asynchronous replica exchange for parallel molecular dynamics applications,” in Proceedings of The 2006 International Conference on Parallel Processing, 2006, pp. 127–134.

Appendix A

Scenarios	Goals	Federated Platforms	Result and Observations
Baseline EnKF Time To Completion (TTC) on TeraGrid and all resources classes for EC2 (m1.small, etc.) Acceleration: Utilize Cloud resources to accelerate the running time. Conservation: Conserving HPC resources (i.e., CPU time usage) and using Clouds to achieve the same TTC. Resilience: Using adaptive resource provisioning to recover from failures. Use cases: Failure on TeraGrid, Failure on EC2 and Failure on both.	Running workflows (EnKF) on EC2 and TeraGrid and utilizing CometCloud to autonomically manage the workflow.	Amazon EC2 (Cloud) TeraGrid (HPC)	Showed that CometCloud can be used to manage workflows on hybrid infrastructure. Showed that Clouds can be utilized to accelerate HPC time to completion, to conserve HPC resources. Showed that CometCloud is adaptive and can recover from failures in the hybrid infrastructure. **Publications and Reports** [23][25]
Deploy the EnKF workflow using three different solvers with (without) block-Jacobi preconditioner. Generalized Minimal Residual Method (GMRES): Uses a large number of MPI reduction operations. Conjugate Gradient (CG) and Bi-Conjugate Gradient (BiCG): Uses matrix vector multiplications but does not guarantee convergence.	Understand the performance fluctuations in Cloud environments using FutureGrid and Amazon EC2 as the experimental platforms.	Amazon EC2 (Cloud) FutureGrid Eucalyptus (Cloud)	Found that IO bottleneck is not the primary cause of the performance fluctuations. There is significantly more variation in performance when 8 or more cores are used per VM. Amazon EC2 has more variation than FutureGrid. This is due to the better hardware used in FutureGrid than in EC2. **Publications and Reports** [22]
Baseline EnKF Time To Completion (TTC) on Amazon EC2 using only the m1.small and c1.medium instance types. Application Adaptivity: Deploy the optimized application on EC2 and evaluate the effects on TTC and EC2 cost. Infrastructure Adaptivity: Deploy the workflow without the instance type restrictions. Understand the effects on TTC and EC2 cost. Hybrid Adaptivity: Use both Adaptivity models to see the effects on TTC and EC2 cost for the EnKF workflow.	Establish the benefits from a hybrid HPC Grid-Cloud execution environment for a sophisticated scientific application. Investigate application infrastructure adaptivity, how it can be supported on hybrid infrastructure, and the subsequent performance advantages.	Amazon EC2 (Cloud)	Experimental results showed that time to completion of the workflow decreased with system and application-level adaptivity. Enabling both system and application-level adaptivity further decreased the time to completion. Applying system adaptivity increased the monetary cost of using EC2 primarily because it caused the application to use more expensive EC2 instance types. **Publications and Reports** [24]
Deploy the Medical Image Registration application using various combinations of private and public Cloud resources to understand the effects of autonomic scheduling on execution time. Deploy the application on the hybrid (private, public) Cloud environment with a budget constraint and use the autonomic scheduler to optimize the application's execution time and cost.	Investigates the use of Clouds and autonomic Cloudbursting to support a medical image registration application. Enable a virtual computational Cloud that integrates local computational environments and public Cloud services on-the-fly with varied computational requirements and QoS constraints.	Rutgers private Cloud (Cloud) Amazon EC2 (Cloud)	The experimental evaluation demonstrated the effectiveness of autonomic Cloudbursts and policy-based autonomic scheduling for the Medical Image Registration application. The performance capabilities of each Cloud have a corresponding impact on the execution time of the application. If there were no budget restrictions then using only EC2 instances is the best solution. However, autonomic scheduling is the best solution when budget is limited. **Publications and Reports** [28]
Deploy an oil reservoir history matching application with an EnKF filter on multiple supercomputers running as-a-Cloud	Investigate transforming supercomputers to a HPC Cloud	Multiple Blue Gene/P	The framework essentially transformed the Blue Gene/P supercomputer into an elastic Cloud, bridging multiple systems to create a HPC federated Cloud, and supporting dynamic provisioning and efficient utilization while maximizing ease-of-use through an as-a-service abstraction **Publications and Reports** [7]

Appendix A: Detailed summary of all experiments using CometCloud and Hybrid HPC/Grid + Cloud infrastructure