RUTGERS
THE STATE UNIVERSITY OF NEW JERSEY

Exploring Software Defined Federated Infrastructures for Science

Manish Parashar
NSF Cloud and Autonomic Computing Center (CAC)
Rutgers Discovery Informatics Institute (RDI[2])
Rutgers, The State University of New Jersey
http://parashar.rutgers.edu/

CometCloud

RUTGERS

Outline

- Clouds, federated computing, software defined systems, …. and Science
- Initial explorations with dynamic federation using CometCloud
- Towards a software-defined federated infrastructure for science
- Summary / Conclusion

RUTGERS

CLOUDS, FEDERATED COMPUTING, SOFTWARE DEFINED SYSTEMS

RUTGERS

The Lure of Clouds

- Cloud services provide an attractive platform for supporting the computational and data needs of academic and business application workflows
- Cloud paradigm:
  - "Rent" resources as cloud services on-demand and pay for what you use
  - Potential for scaling-up, scaling-down and scaling-out, as well as for IT outsourcing and automation
- Landscape of heterogeneous cloud services spans private clouds, public clouds, data centers, etc.
  - Heterogeneous offering with different QoS, pricing models, availability, capabilities, and capacities
  - Hybrid cloud infrastructures could integrate private clouds, public clouds, and data centers
- Novel dynamic market-places where users can take advantage of different types of resources, quality of service (QoS), geographical locations, and pricing models
- Cloud federations extend as-a-service models to virtualized data-centers federations

RUTGERS

Clouds as Enablers of Science

- Clouds are rapidly joining traditional CI as viable platforms for scientific exploration and discovery
- Possible usage modes:
  - Clouds can 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
  - Cloud support the democratization
  - Cloud abstractions can support new classes of algorithms and enable new applications formulations
  - Application driven by the science, not available resources -- Cloud abstractions for science?
- Many challenges
  - Application types and capabilities that can be supported by clouds?
  - Can the addition of clouds enable scientific applications and usage modes that are not possible otherwise?
  - What abstractions and systems are essential to support these advanced applications on different hybrid platforms?

RUTGERS

Cloud Usage Modes for Science

- ***HPC in the Cloud*** – outsource entire applications to current public and/or private Cloud platforms

- ***HPC plus Cloud*** – Clouds 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.

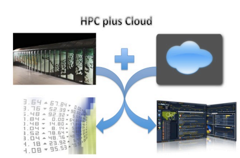

- ***HPC as a Cloud*** – expose HPC/Grid resources using elastic on-demand Cloud abstractions

RUTGERS

Federated Computing for Science (I/II)

- Scientific applications can have large and diverse compute and data requirements
- Federated computing is a viable model for effectively harnessing the power offered by distributed resources
  - Combine capacity, capabilities
- HPC Grid Computing - monolithic access to powerful resources shared by a virtual organization
  - Lacks the flexibility of aggregating resources on demand (without complex infrastructure reconfiguration)
- Volunteer Computing - harvests donated, idle cycles from numerous distributed workstations
  - Best suited for lightweight independent tasks, rather than for traditional parallel computations

RUTGERS

Federated Computing for Science (II/II)

- Current/emerging science and engineering application workflow exhibit heterogeneous and dynamic workloads, and highly dynamic demands for resources
  - Various and dynamic QoS requirements
    - Throughput, budget, time
  - Unprecedented amounts of data
    - Large size, heterogeneous nature, geographic location
- Such workloads are hard to be efficiently supported on classic federation models
  - Rigid infrastructure with fixed set of resources
- Can we combine the best features of each model to support varying application requirements and resources' dynamicity?
  - Provisioning and federating an appropriate mix of resources on-the-fly is essential and non-trivial

RUTGERS

Software Defined ….

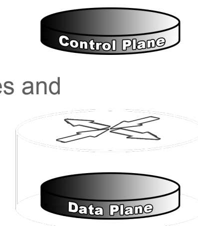

- Software Defined Networks
  - An approach to building computer networks that separates and abstracts elements of these systems (Wikipedia)
  - E.g., separation of control and data plane
- Software Defined Systems
  - Based on software defined networking (SDN) concepts
  - Allow business users to describe expectations from their IT in a systematic way to support automation
  - Enable the infrastructure to understand application's needs through defined policies that control the configuration of compute, storage, and networking, and it optimizes application execution
    - Open virtualization, Policy driven optimization and elasticity – autonomics, Application awareness
- See also software defined data centers, ….

RUTGERS

EXPLORING FEDERATED INFRASTRUCTURE FOR SCIENCE USING COMETCLOUD

RUTGERS

CometCloud

- Enable applications on dynamically federated, hybrid infrastructure exposed using Cloud abstractions
  - **Services:** discovery, associative object store, messaging, coordination
  - **Cloud-bursting:** dynamic application scale-out/ up to address dynamic workloads, spikes in demand, and extreme requirements
  - **Cloud-bridging:** on-the-fly integration of different resource classes (public & private clouds, data-centers and HPC Grids)
- High-level programming abstractions & autonomic mechanisms
  - Cross-layer Autonomics: Application layer; Service layer; Infrastructure layer
- Diverse applications
  - Business intelligence, financial analytics, oil reservoir simulations, medical informatics, document management, etc.

http://cometcloud.org

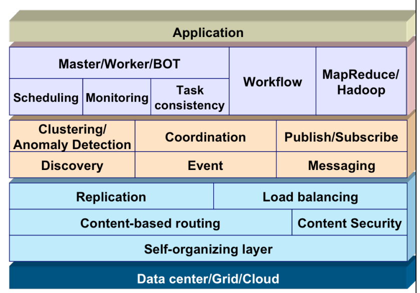

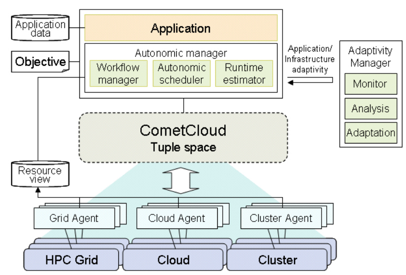

Federated (hybrid) computing infrastructure

RUTGERS

Autonomics in CometCloud

- **Autonomic manager** manages workflows, benchmarks application and provision resources.
- **Adaptivity manager** monitors application performance and adjusts resource provisioning.
- **Resource agent** manages local cloud resources, accesses task tuples from CometCloud and gathers results from local workers so as to send them to the workflow (or application) manager.

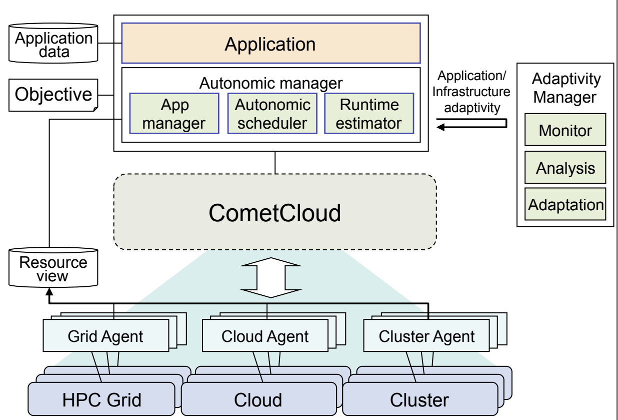

RUTGERS

On-Demand Elastic Federation using CometCloud

- Autonomic cross-layer federation management
  - Resources specified based on availability, capabilities, cost/performance constraints, etc.
  - Dynamically assimilated (or removed)
  - Resources coordinate to:
    - Identify themselves / verify identity
    - Advertise their resources capabilities, availabilities, constraints
    - Discover available resources
- Federation coordinated using Comet spaces
- Autonomic resource provisioning, scheduling and runtime adaptations
- Business/social models for resource sharing

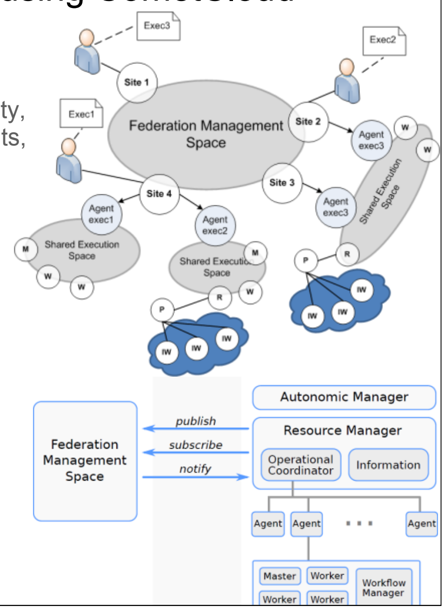

RUTGERS

An Initial Experiment: Fluid Flow in Microchannel

- Controlling fluid streams at microscale is of great importance for biological processing, creating structured materials, etc.
- Placing pillars of different dimensions, and at different offsets, allows “sculpting” the fluid flow in microchannels
- Four parameters affect the flow:
  - Microchannel height
  - Pillar location
  - Pillar diameter
  - Reynolds number

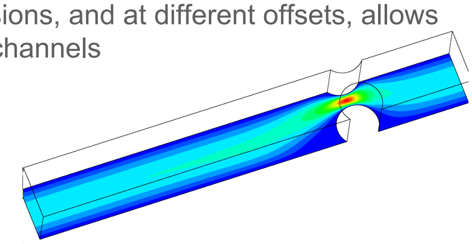

- Each point in the parameter space represents simulation using the Navier-Stokes equation (MPI-based software)
- Highly heterogeneous and computational cost is hard to predict a priori
- Global view of the parameter space requires 12,400 simulations (three categories)

RUTGERS

Experiment Summary

- 10 different HPC resources from 3 countries federated using CometCloud
- 16 days, 12 hours, 59 minutes and 28 seconds of continuous execution (in spite of failures, etc.)
- 12,845 tasks processed, 2,897,390 CPU-hours consumed, 400 GB of data generated

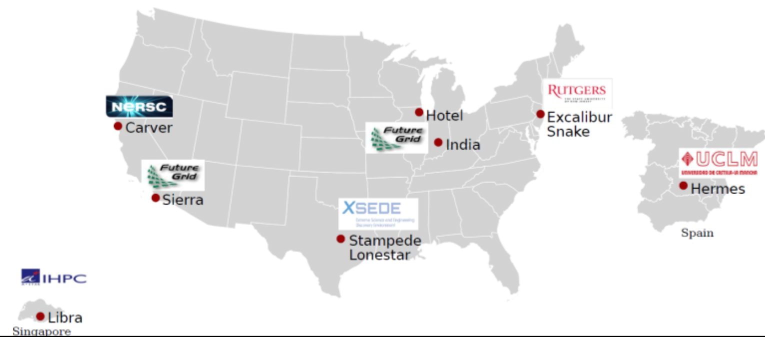

RUTGERS

Experiment Summary (II)

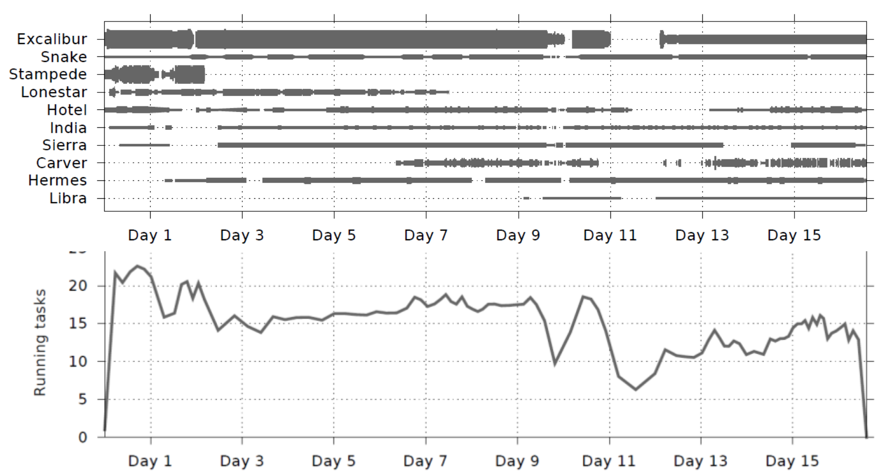

RUTGERS

Throughput and Queue Waiting Time

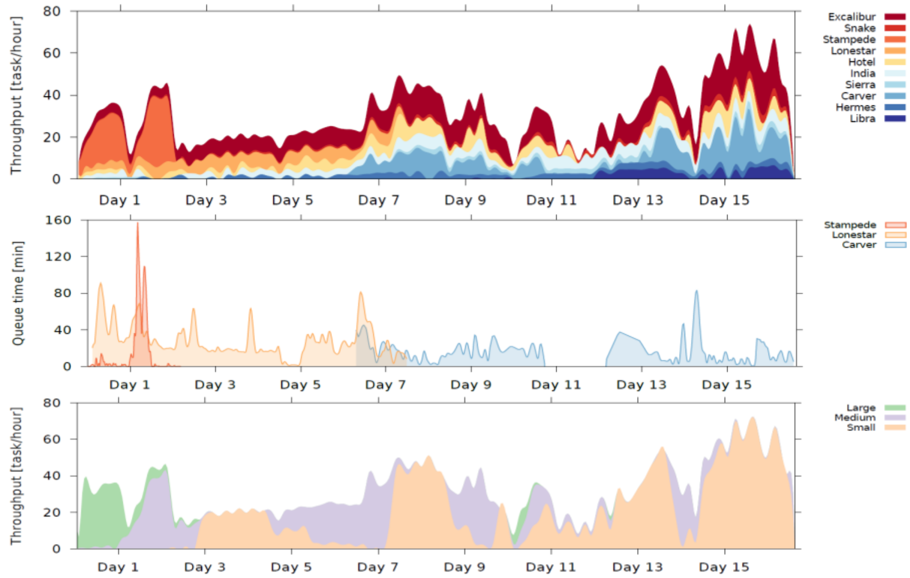

RUTGERS

Science Outcomes

- The most comprehensive data on the effect of pillars on microfluid channel flow
- Library of flow transformations
- Arranging pillars is possible to perform basic flow transformation
- What is the optimal pillar arrangement to achieve a desired flow output?
- Useful for medical diagnostics, smart materials engineering, and guiding chemical reactions

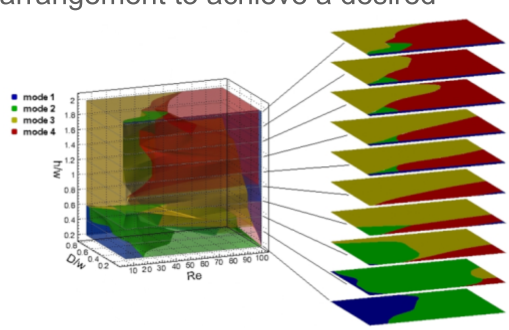

RUTGERS

HPC as a Service (Winner SCALE'11)

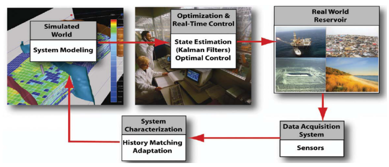

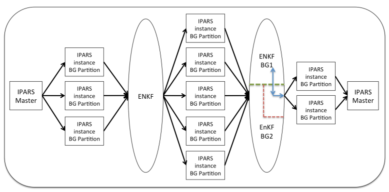

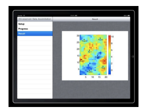

Demonstrated how the cloud abstraction can be effectively used to support ensemble geo-system management applications on a geographically distributed federation of supercomputing systems using a pervasive portal running on an iPad

http://nsfcac.rutgers.edu/icode/scale

RUTGERS

HPC as a Service (IEEE Computer 10/12)

- HPC as a Service using federation of IBM Blue Gene/P systems
- Elastically scale up to 22K processors

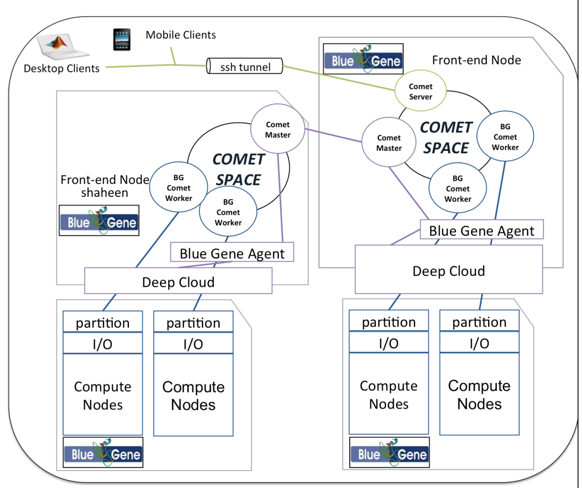

Accelerating Protein Folding using Advanced Computational Infrastructure (Rutgers + BMS)

Individual trajectories
- Parallel NAMD trajectories
- Asynchronous communication in cometCloud

Science
- Be smart about using resources
- Commodity hardware versus high end resources
- Terminate or restart resources

Infrastructure
- Federated clouds

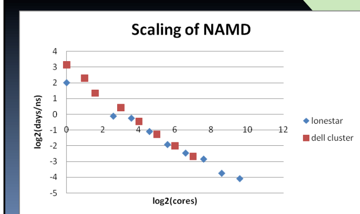

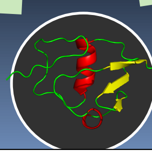

XSEDE (TACC)

Rutgers Cluster

Excalibur

RepEx App

EC²

FutureGrid

GUI

Autonomic Master (Rutgers Cluster)

Based on replica progress, Autonomic Master stops commodity trajectory and starts replica set on high performance resources.

Reporting secondary structure and radius of gyration progress

Replica Set	Replica Set	Replica Set	Replica Set
(FutureGrid)	(Rutgers Cluster)	(Amazon EC²)	(Excalibur)
64 cores	128 cores	8 cores	4096 cores
2 ensembles	2 ensembles	1 ensemble	1 ensembles
4 cores/replica	8 cores/replica	1 cores/replica	512 cores/replica

*Could run multiple replicas per temperature to improve likelihood of asynchronous exchange on heterogeneous hardware.

*8 temperatures = 1 ensemble

http://youtu.be/sg2C7N7g5CU

RUTGERS

Other experiments

- Data-Driven Workflows on Federated Clouds [Cloud'14]
- Federating Resources using Social Models [IC2E'14]
- Elastic Federations for Large-scale Scientific Workflows [MTAGS'13]
- HPC plus Cloud Federations [e-Science'10]
- …. [See cometcloud.org]

- Testbed using resource in US (RU, FutureGrid, XSEDE, IBM), UK (Cardiff), Amazon EC2

- Experiments successful…. but can the model be generalized?

RUTGERS

Software Defined Cyberinfrastructure Federations for Science ?

- Combine cloud abstractions with ideas from software-defined environments
- Independent control over application and resources
- Living federation that autonomously adapting itself to:
  - Changes in the environment
  - Application requirements

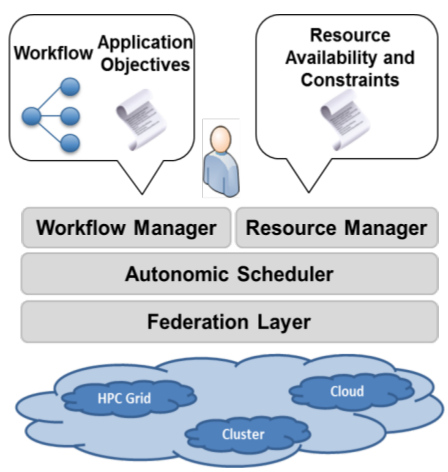

RUTGERS

Software Defined Cyberinfrastructure Federations for Science

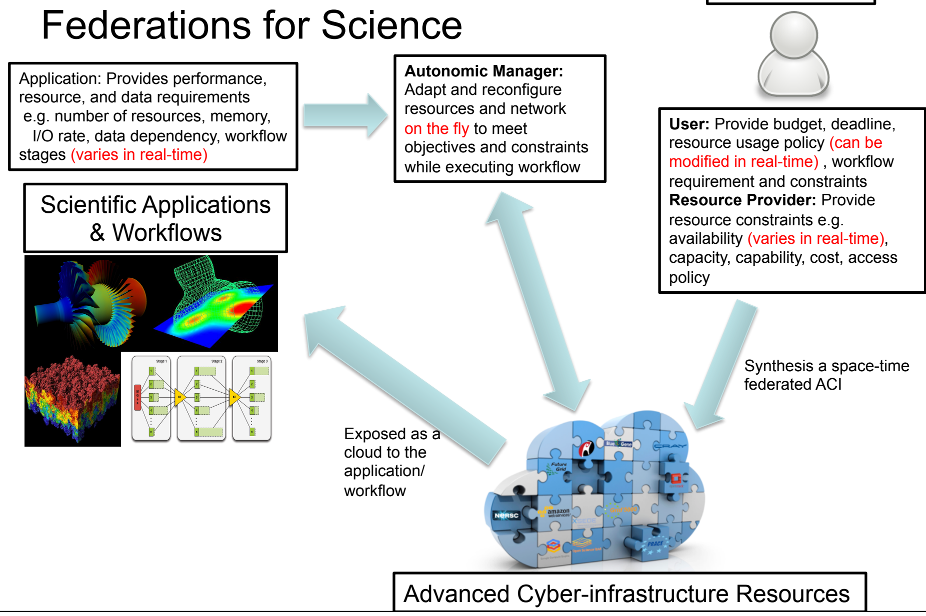

RUTGERS

Software-defined ACI: ACI-as-a-Cloud

- Software defined ACI federations exposed using elastic on-demand Cloud abstractions
- Declaratively specified to define availability as well as policies and constraints to regulate their use
  - Use of a resources may only be allowed at certain times of the day, or when they are lightly loaded, or when they have sufficient connectivity, etc.
  - Prefer certain type of resources over others (e.g., HPC versus clouds or “free" HPC systems versus the allocation-based ones)
  - Specify how to react to unexpected changes in the resource availability or performance
  - Use resources only within the US or Europe due to the laws regulating data movement across borders
- Evolve in time and space -- the evaluation of these constraints provides a set of available resources at evaluation time
- Leverage software-defined networks to customize and optimize the communication channels or software-defined storage to improve data access

RUTGERS

Software-defined ACI: Platform as a Science

- Platform as a Service to decouple applications from the underlying ACI Cloud
- Key components
    1. An API for building new applications or application workflows
    2. Mechanisms for specifying and synthesizing a customized views of the ACI federation that satisfies users' preferences and resource constraints
    3. Scalable middleware services that expose resources using Cloud abstractions
    4. Elasticity exposed in a semantically meaningful way
    5. Autonomics management is critical
- CometCloud provides some of these; currently focusing on 2

RUTGERS

Many technical issues

- **Deployability:** Must be easy to deploy by a regular user without special privileges
- **Standardization/Interoperability:** Interact with heterogeneous resources
- **Self-discovery:** Discovery mechanisms to provide a realistic view of the federation
- **Scalability and extended capacity:** Scale across geographically distributed resources
- **Elasticity:** Ability to scale up, down or out on-demand
- **Security, Authentication, Authorization, Accounting……**
- ….

RUTGERS

Some Related Efforts

- FED4FIRE (European Union FP7)
  - A common federation framework for developing, adapting or adopting tools that support experiment lifecycle management, monitoring and trustworthiness
- InterCloud (Univ. of Melbourne, Australia)
  - Utility-oriented federation of cloud computing environments for scaling of application services
- Business Oriented Cloud Federation (Univ. of South Hampton, UK)
  - Cloud federation model via computation migration for real time applications; targets real-time online interactive applications, online games
- ….

RUTGERS

Summary

- Emerging CDS&E workflows have dynamic and non-trivial computational/data requirements
  - Necessitate dynamically federated platforms that integrate heterogeneous resources / services
  - Provisioning and federating an appropriate mix of resources on-the-fly is essential and non-trivial
- Software-defined Advanced Cyber-Infrastructure for Science
  - Software defined ACI federations exposed using elastic on-demand Cloud abstractions
  - Application access using established programming abstraction/platforms for science
  - Autonomic management is critical
- Many challenges at multiple layers
  - Application formulation, programming systems, middleware services, standardization & interoperability, autonomic engines, etc.

RUTGERS

The CometCloud Team

Moustafa AbdelBaky
PhD student, Dept. of Electrical & Computer Engr.
Rutgers University
Email: moustafa@cac.rutgers.edu

Mengsong Zou
PhD Student, Dept. of Computer Science
Rutgers University
Email: mz228@cac.rutgers.edu

Javier Diaz-Montes, Ph.D.
Assistant Research Professor,
Dept. of Electrical & Computer Engr.
Rutgers University
Email: javidiaz@rdi2.rutgers.edu

Manish Parashar, Ph.D.
Rutgers Discovery Informatics Institute (RDI²)
Rutgers University
Email: parashar@rutgers.edu

And many collaborators….

CometCloud: http://cometcloud.org

36

RUTGERS

Thank You!

Manish Parashar, Ph.D.
Rutgers Discovery Informatics Institute (RDI²)
Rutgers, The State University of New Jersey

Email: parashar@rutgers.edu
WWW: http://parashar.rutgers.edu/
CometCloud: http://cometcloud.org

Acknowledgments

RUTGERS

IOWA STATE UNIVERSITY

NERSC

UCLM
UNIVERSIDAD DE CASTILLA-LA MANCHA

IHPC
A*STAR

XSEDE
Extreme Science and Engineering Discovery Environment

IBM

Future Grid

NSF