Center for Cloud and Autonomic Computing (CAC)

A CISE-funded Center

University of Florida, Jose Fortes, 352.392.9265, fortes@ufl.edu

Rutgers University, Manish Parashar, 732.445.4388, parashar@cac.rutgers.edu

University of Arizona, Salim Hariri, 520.621.4378, hariri@ece.arizona.edu

Mississippi State University, Ioana Banicescu, 662.325.7508, ioana@cse.msstate.edu

Center website: http://www.nsfcac.org/

CometCloud Manages Business Workflows on Federated Cloud Infrastructure

Public clouds have emerged as an important solution that enables the renting of resources on-demand and supports a pay-as-you-go pricing policy. Furthermore, private clouds or data centers, which cater to a restricted set of users within an organizational domain, are exploring the possibility of scaling out to public clouds to respond to unanticipated resource requirements.

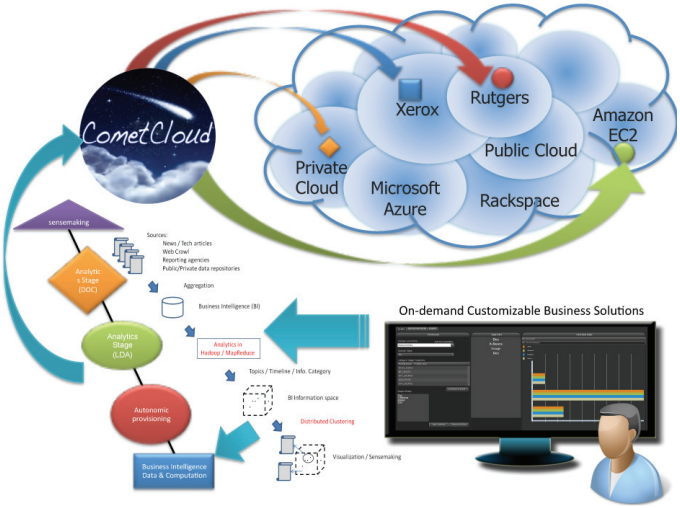

Test bed based on Supervisory Control and Data Acquisition (SCADA) systems provides a platform for testing both hardware and software.

Rutgers and Xerox collaborated to develop and deploy an innovative framework for executing business workflows using a dynamically federated cloud infrastructure. The framework builds on CometCloud, an autonomic cloud engine developed at Rutgers, and dynamically federates hybrid infrastructure, such as private clouds, enterprise data centers, grids, and public clouds, on demand, to meet the requirements and constraints of the business workflow. An enterprise workflow typically consists of an ordered set of heterogeneous applications (stages), each of which may have specific resource requirements and constrains in terms of performance, completion time, cost, data privacy, etc. Using CometCloud, different workflow stages can be deployed on the appropriate mix of resources to satisfy application requirements and constraints.

The breakthrough enabled Rutgers and Xerox to demonstrate for the first time a number of key capabilities in a single framework: 1) Dynamic cloud federation - managing resources across multiple private and public clouds in order to dynamically scale the execution of application workflows up, down, and/or out and/or to compose appropriate capabilities, according to high-level policies; 2) Programming management - resource scheduling and provisioning within the federated cloud infrastructure; and 3) Workflow deployment - deployment of real-world application workflows running on the federated cloud infrastructure. Specifically, the hybrid infrastructure used in the demonstration dynamically integrated private clouds at Rutgers and ACS with the Amazon EC2 public cloud.

Such an autonomic workflow framework can dynamically select an optimal mix of resource classes (clouds or grids provider, types of nodes, the number of nodes, etc.) based on application QoS and resources requirements (e.g., performance, latencies), user policies (e.g., budget, deadline), and constraints (e.g., security/privacy). The workflow framework can also monitor the execution of the applications services within the workflow, and can adapt both the resource provisioning as well as the services to ensure that the application requirements and user constraints continue to be satisfied. Adaptations may involve scaling resources up, down, or out within the federated cloud infrastructure and can allow the system to handle unanticipated situations such as workload bursts, system performance degradation, or resource failures.

> **Economic Impact:** 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. Enabling on-demand provisioning and construction of hybrid federated cloud infrastructures has the potential to reduce computational costs and improve efficiency of cloud computing service centers. These structures can support heterogeneous and dynamic workloads and on-demand cloud bridging. Federated cloud infrastructures also provide opportunities to improve application quality of service and lower cost by mapping applications of scientific or business workflows to appropriate resource providers.

For more information, contact Manish Parashar, 732.445.5388, parashar@rutgers.edu.

Autonomic Critical Infrastructure Protection System (ACIP)

Numerous serious cybersecurity exposures occur because of widespread use of Supervisory Control and Data Acquisition (SCADA) systems that were never designed with security in mind. Consequently, SCADA systems become a prime target for cyberattacks due to the profound and catastrophic impacts they can have on our economy and on all aspects of our life. In fact, critical infrastructures have expanded to include not only the energy critical infrastructures, but also many process control systems, networks, and infrastructures of which approximately 85% are privately owned.

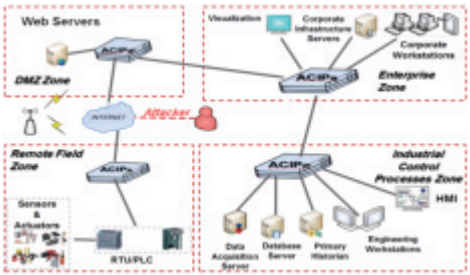

Test bed based on Supervisory Control and Data Acquisition (SCADA) systems provides a platform for testing both hardware and software.

A recent Forrester survey reported that 75% of organizations experienced Distributed Denial of Service Attacks (DDoS) even though they implemented cybersecurity solutions. One third of attacked organizations experienced service disruptions.

The problem is that many of the systems are ineffective against novel and well-organized attacks. The nation's critical energy infrastructures (power, water, gas and oil) are moving to modernize their industrial control systems to build what is referred to as "Smart Grids" that use advanced computing and communications technologies to bring knowledge so they can operate far more robustly and more efficiently.

Motivated by these exposures, researchers at the Center for Autonomic Computing (CAC) and industrial center members (Raytheon and AVIRTEK) are collaborating to develop an innovative cybersecurity approach based on autonomic computing technology. It is analogous to the human nervous system where computing systems and applications can be self-configured, self-optimized, self-healed, and self-protected with little involvement from the users and/or system administrators. CAC has successfully developed and implemented an Autonomic Critical Infrastructure Protection (ACIP) appliance and currently being tested and evaluated. This involves evaluating the appliance's self-protection capabilities using an industrial process control test bed that offers multiple capabilities for both hardware and software experimentation.

This breakthrough technology is validating the thesis that autonomic paradigms have the potential to detect and mitigate cyber threats launched against industrial control systems. Responding faster than a human operator, SCADA and their associated control elements can effectively immunize against cyber malware and mitigate the effects of control.

Economic Impact: Enhancing the ability of the nation to provide undisrupted service of electric power, clean potable water, transportation, and other necessary societal support services, saves lives, preserves the domestic tranquility, and can help protect industry's and the nation's economic vitality. The economic impacts of avoiding such cyber attacks are difficulty to estimate. That said, a study conducted by the same group, Forrester Consulting, indicated that organiza-

tions that provide online services as their core business stand to lose millions of dollars per hour when their services are down. The ACIP technology when fully matured can be exploited by western world societies to immunize critical infrastructures against being targeted by malcontents and terrorists.

For more information, contact Salim Hariri, 520.621.4378, hariri@ece.arizona.edu.

Demand-driven Service and Power Management in Data Center

Power consumption represents an increasingly significant percentage of the cost of operating large data centers. One approach to reduce power consumption in the data centers is to keep the computers in standby or off modes except when the data center workload requires them to be fully on. This approach depends on being able to effectively monitor performance, workloads, and to anticipate the need for resources required to meet the users' service-level agreements of that generate the workload.

Data centers are used by banks, investment firms, research organizations, major municipalities and by homeland security.

This project has generated: a) mechanisms to monitor, model and predict workloads associated with individual services; b) models for prediction of global resource demands; c) data management methods based on control theory and market-based approaches; d) mechanisms to minimize the cost of providing individual services while globally minimizing power consumption; and e) development and evaluation of software that implements all of the aforementioned methods. Ongoing experimental evaluations on an IBM BladeCenter have shown that the proposed approach can efficiently and stably reduce thermal hotspots, power consumption and performance degradation caused by virtual machine consolidation, while balancing conflicting objectives.

Economic Impact: Annual energy and administration costs associated with today's data centers amount to billions of dollars; power and cooling rates are increasing by an alarming 8 fold every year and are becoming the dominant part of IT budgets. The high-energy consumption of modern data centers causes excessive heat dissipation, which, in turn, increases cooling costs and server failure rates. One of CAC's main research thrusts aims to address these problems because doing so lowers data center ownership costs in all sectors of today's economy. Doing so can also increase the reliability of the infrastructure that provides critical services.

For more information, contact Jose Fortes, 352.392.9265, fortes@ufl.edu.