Cereus: CyberInfrastructure Environments for Resource Exchange and Utility Services

Duke University, Department of Computer Science

http://issg.cs.duke.edu/cereus

Objective: manage shared cyber infrastructure as an on-demand utility in which computing resources are as easy to access as electricity, with access policies that are fair, flexible, effective, and robust.

Goals and Rationale

- Manage server and storage infrastructure as a shared resource, like the rest of the network.
  - Distributed investment costs, economies of scale
  - Distributed risk of load surges and resource failures
  - Harness surplus resources, multiplex bursty demand
- Protect shareholders from the risks of resource sharing.
  - Contractual performance assurances (leases)
  - Secure isolation from competing users (or attackers)
  - User control over configuration and access
- Adapt to dynamic conditions.
  - Continuous, reliable service: reconfigure around failures, load services, and other conditions (e.g., thermal hot spots).
  - Fast response with low operational cost: self-managing or "autonomic" system

Providers, Guests, and Brokers

- Architecture and player roles defined by SHARP framework
  - Secure Highly Available Resource Peering (SOAP/XML)
- Resource providers are cluster/storage sites (e.g., a cell of a data center), each controlled by a site authority service.
- Resource consumers are application services that run as guests on the resources of hosting sites.
  - E.g., hosted Web services or Grid/PlanetLab services
- Each guest is controlled by a service manager.
  - Negotiates for resources on behalf of the guest.
  - Monitors execution of the guest, responds to events.
- Brokers match consumers and suppliers, arbitrate demands.
  - Resource discovery and resource management

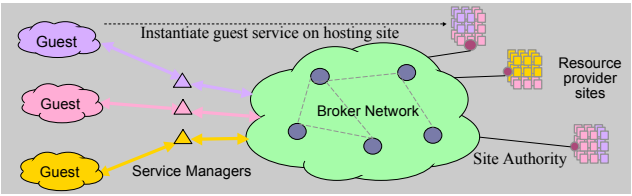

- Long-term research: policy for brokers and service managers

Resource Contracts: Leases

- SHARP protocol uses asymmetric cryptography as basis for secure and accountable resource lease contracts.
  - Digitally signed XML assertions (future: SAML)

Shorter leases have higher overhead, but are more agile with changing demand.

Longer leases give stronger assurance of stable resource control by the holder.

```
<lease>
  <issuer> Site A's public key </issuer>
  <signed_part>
    <holder> SM's public key </holder>
    <rset> 20 type 4 servers </rset>
    <start> Monday 12:01 AM </start>
    <end> Tuesday 11:59 PM </end>
    <sn> unique ID at Site A </sn>
  </signed_part>
  <signature> A's signature </signature>
</lease>
```

Cluster-on-Demand (COD)

- Example: Cluster-on-Demand (COD) is a SHARP authority service for LDAP-administered cluster sites.
  - NSF National Middleware Infrastructure (NMI) funded project
  - Duke/CSEM deployment in progress, seeded by IBM
- COD guests may run within isolated virtual clusters.
  - Supports virtual machines (e.g., Xen, VMware).
- Default COD service manager is a web/email user interface.
  - Pluggable automated service managers for feedback control, e.g., meet objectives in Service Level Agreements (SLAs).
  - Batch schedulers for computational clusters (e.g., Globus GRAM with PBS/SGE batch job management).
- Flexible and reliable configuration management with user-specified XML event handlers (based on Apache Ant).

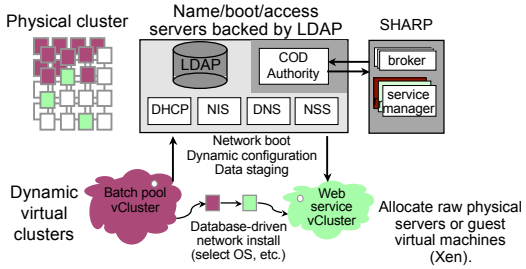

A Leasing Market for Cyberinfrastructure

- **Research focus**: market-based control for resource access based on cash, credits, or a virtual currency.
  - Balance resource supply and demand
  - Allocate resources to their "highest and best use"
- Self-organizing, self-optimizing emergent global behavior
  - *Sustainability*: create incentives for sites to contribute resources
  - *Arbitration*: create incentives to schedule usage under constraint
- Brokers conduct regular auctions for resource contracts.
  - Service managers bid from a currency budget.
  - Incentive-compatible auction protocols with posted price signals.
  - Brokers trade resource rights: resources diffuse through the network of brokers by following the price gradients.

Accountable Auctions and Virtual Currency

- Accountable cooperation: participants are community members securely bound to identities and subject to sanction.
  - All SOAP/XML message exchanges are digitally signed and timestamped: actions are tamper-evident and non-repudiable.
  - Untrusted auditors inspect signed action records to detect lying, cheating, or stealing, and prove it to a third party.
- Self-recharging virtual currency
  - Budgets expire/replenish automatically: limits instability due to hoarding and predatory behavior in money economies.
  - Currency is issued by a central bank, but participants trade directly; received currency is redeemable at the bank.
  - Participants are accountable for their currency management.