Archer: a shared cyberinfrastructure for computer architecture research and education

The Archer project is building a distributed, open community resource for simulation-based computer architecture research and education. Computer architecture researchers and students world-wide are invited to participate in this cyber-infrastructure, gaining access to high-throughput computing pools with easy-to-deploy architecture tools, and sharing their own idle computer cycles, applications, data sets and educational material.

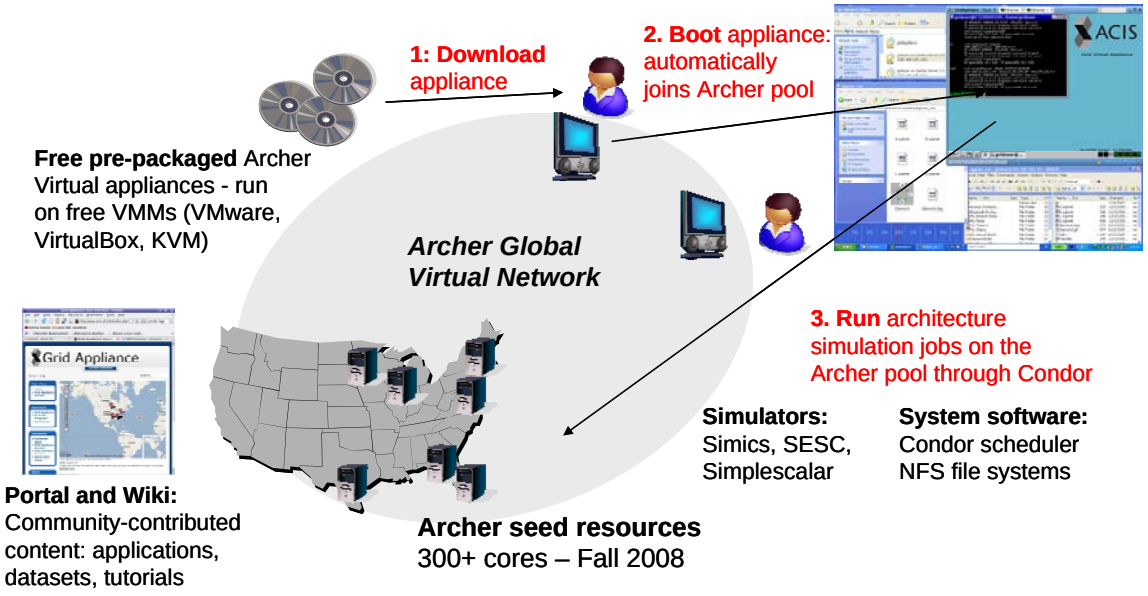

VIRTUAL CLUSTERS

Archer builds upon virtualized resources – virtual machines and overlay virtual networks – to create scalable virtual clusters over wide-area Internet resources. VM appliances and self-configuring IP-over-P2P virtual network software enables Archer nodes to be easily deployed by end users.

CONDOR

Archer virtual clusters run the Condor batch job scheduler, a robust Grid middleware that supports typical computer architecture simulation workloads – high-throughput execution of independent jobs, such as parametric sweeps.

COLLABORATION TOOLS

Archer provides a Wiki to share documents and tutorials, auto-configured NFS file systems to share datasets and tools, and user groups for questions and support. Contribute to the community by sharing your own research and education content such as tools, scripts, tutorials, class assignments and datasets.

archer-project.org

Flexible deployment options:

- ***Archer Express***: For first time users. *Quickly* gain access to a small-scale pool of resources within ~ 30 minutes to test-drive Archer capabilities on your own.
- ***Archer Global***: For academic research and education use. Gain access to a *large-scale* world-wide pool of CPU resources, shared tools and datasets. Effortlessly contribute with your own resources to the community when idle.
- ***Archer Local***: For deployment of *private* Archer pools on your own local-area resources (clusters, servers, desktops).

Easy steps to get started:

1. Install a free VM monitor (VMware Player, VirtualBox, or KVM) in your desktop - x86-based Windows, Linux, or MacOS.

2. Go to ***archer-project.org***, download and run the Archer Express Grid appliance image

3. Follow the “quick start” tutorial to submit your first simulation run

National Science Foundation
WHERE DISCOVERIES BEGIN

This work is sponsored by the National Science Foundation under CRI collaborative awards 0751112, 0750847, 0750851, 0750852, 0750860, 0750868, 0750884, and 0751091.

Archer: a shared cyberinfrastructure for computer architecture research and education

PARTICIPATING INSTITUTIONS

Archer is a collaborative effort led by University of Florida, Florida State University, Northeastern University, University of Texas at Austin, Northwestern University, Cornell University, University of Minnesota and University of Wisconsin-Madison.

APPLICATIONS

Simulation tools widely-used in the community are available in Archer, including SESC, Simics, SimpleScalar, and PTLSim. In the typical case, Linux 32-bit x86 binaries will work on Archer, unmodified.

MANAGEMENT AND USER SUPPORT

Management staff develops and maintains the Archer virtual machine appliance image so Archer users can focus on the science rather than the infrastructure. Management staff also provides support through user mailing lists.

A community cyberinfrastructure:

- ***Shared distributed computers*:** Archer makes high-throughput computing accessible to the computer architecture community at large. Seeded with hundreds of CPUs across a distributed system (to be deployed over three years), the infrastructure grows as a community effort because users can effortlessly contribute with resources of their own to the community when they are idle.

- ***Shared architecture content***: Archer facilitates the sharing of tools, datasets, self-contained experimental setups, and educational materials. Sharing of tools and experiments is facilitated by providing a consistent virtual environment to all Archer users and automatically-configured NFS file systems. Sharing and dissemination of content is facilitated by the Archer Wiki.

The Archer technology targets simple deployment and use:

- ***Pre-packaged VM appliances***: Archer packages all the software needed to run architecture simulation jobs conveniently. It uses freely available virtual machine software to create self-configuring, plug-and-play virtual appliances and the robust Condor batch job scheduler. Archer VM appliances enable functionality similar to systems such as SETI@Home, but support unmodified applications and provide sandboxing of remote jobs through the use of virtualization.

The Archer software makes it very easy for computer architects to deploy their own Linux Condor pools on their own resources (for example, instructional lab PCs). A small local pool can be brought up in a matter of hours.

Frequently Asked Questions

- ***What are the software licenses used in Archer?***

  The core of the Archer middleware is open-source (GPL and Apache2 licenses). Applications include both open-source packages (e.g. SESC) and simulators with academic licenses available to Archer users (e.g. Simics).

- ***Who can use Archer? Is it free?***

  Archer is available and free for use by computer architecture researchers and educations within the US and abroad. Usage of the Archer Global pool is limited to academic use only; Archer middleware may be used to deploy Archer Local pools for non-academic purposes.

- ***Can I use Archer to create an isolated, private Condor pool in my own resources?***

  Yes, you can use the VM appliances to deploy Archer Local pools, which are independent from the shared pool. However, Archer staff will not be able to assist with Archer Local management.

- ***What security mechanisms are provided?***

  All jobs run within a virtual machine sandbox. Access control is mediated by Condor, which means users do not have direct access (e.g. ssh) to a VM, except through Condor. Remote jobs run as UNIX user "nobody" within a VM sandbox. Archer jobs can only send packets to other virtual machines within the Archer virtual network. Packets sent to other Internet hosts are blocked by a firewall. Archer confines its traffic to the virtual network through firewalls to avoid the use of Archer to launch denial-of-service attacks. All virtual network IP traffic between Archer VMs is authenticated and encrypted.

ACIS
Advanced Computing and Information Systems Lab
http://www.acis.ufl.edu