

GENI Experiment Control Using Gush

Jeannie Albrecht and Amin Vahdat
Williams College and UC San Diego



~~GEC 3~~
GEC 4




Overview

- How do we actually use GENI?
- Goal: Develop abstractions and tools for addressing the challenges of managing distributed applications
 - Make it easy for a range of users to run a variety of experiments on GENI
- Talk outline
 - Summarize existing work and design of **Gush**
 - Describe implementation/integration plan and timeline


Running an Experiment

- Suppose we have written our software, we have created our slice, and we are ready to **deploy** on PlanetLab for the first time
- We could...
 1. Connect to each of the 800+ PlanetLab machines
 2. Download software (no common file system)
 3. Install software
 4. Run application and analyze performance
 5. Check for errors on each machine
 6. When we find an error, we start all over...
- Or we could use **Gush**



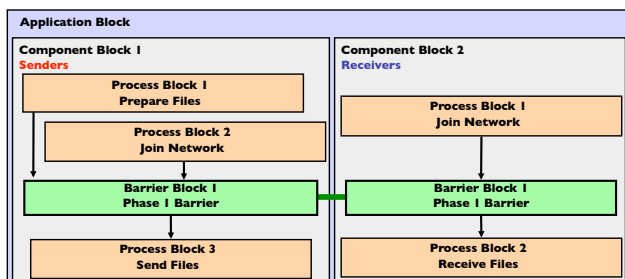
Gush

- A distributed application management infrastructure
 - Extends functionality of **Plush** to support experiment control on GENI
 - Designed to simplify deployment of distributed applications
 - Provides abstractions for configuration and management
 - Allows users to “remotely control” computers running distributed applications worldwide




Step 1: Describe Application

- Describe experiment using application “building blocks”
- Create customized control flow for distributed applications
- **Application specification** blocks are described using XML




Step 2: Acquire Resources




- How can we find “good” machines?
 - We want machines with specific characteristics
 - PlanetLab services perform *resource discovery* to find machines that satisfy our requirements (e.g., SWORD)
- Gush interfaces directly with these services
 - Eventually Gush will talk directly to GENI Clearinghouses to find resources


Step 3: Configure Resources




- Connect to and configure selected resources
 - Optionally create a tree for achieving scalability in communication
 - **Controller** “remotely controls” the **clients** on our behalf
 - Install software on clients (some are **senders**, some are **receivers**)




Step 4: Start Application




- Controller issues commands to clients telling them to start running our application
 - **Senders** begin running sender processes
 - **Receivers** begin running receiver processes




Step 5: Monitor Application




- We want to make sure the processes keep running
- Gush clients monitor experiment processes for failures
 - If a failure is detected, client notifies controller
 - Controller decides to tell client to restart failed program or process



Step 6: Cleanup

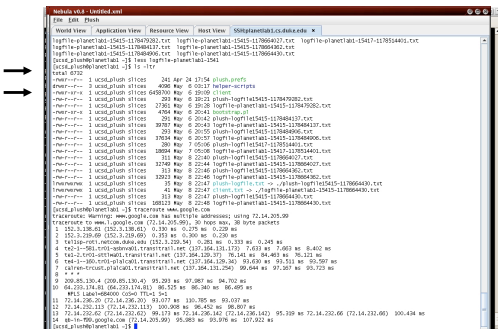


- Gush clients make sure all programs exited cleanly
- Remove logs and software from remote machines
- Disconnect clients from controller



Gush User Interfaces

- Command-line interface used to interact with applications (see DEMO)
 - Provides single point of control for remotely controlling resources
- Nebula (GUI) allows users to describe, run, monitor, & visualize applications
- XML-RPC interface for managing applications programmatically



Current Status & Recent Activities

- Gush works with current PlanetLab XML-RPC API
 - Have written some code for geniwrapper integration (but it has not been tested)
- Preliminary user study completed
- Gush web site is “live” now (<http://gush.cs.williams.edu>)
 - Instructions for obtaining and using Gush
- Undergraduate student involvement at Williams
 - Gush web interface (GoogleMaps) [summer 2008]
 - Batch scheduler that uses Gush XML-RPC interface [summer 2008]
 - Gush and Nebula development [summer 2009]
- Gush in the classroom
 - “Bringing Big Systems to Small Schools: Distributed Systems for Undergraduates.” SIGCSE 2009.

Summary & Next Steps

- Gush provides abstractions for managing a range of distributed applications on GENI
 - Provides three different user interfaces to meet the needs of a variety of researchers with varying levels of expertise
- Next steps
 - Continue obtaining user feedback to enhance usability and provide additional functionality
 - Develop better user and developer documentation
 - Continue to integrate with other Cluster B services (e.g., Raven)
- Long term goals
 - Enhance GUI (Nebula) functionality
 - Extend Gush to support experiments on sensor/mobile networks

Thanks!

For more info, visit

<http://gush.cs.williams.edu>

Email:

jeannie@cs.williams.edu