

SPRUCE

A Special Priority and Urgent Computing Environment for supporting Urgent High-Performance Computing

Ivan Beschastnikh
Computation Institute
The University of Chicago

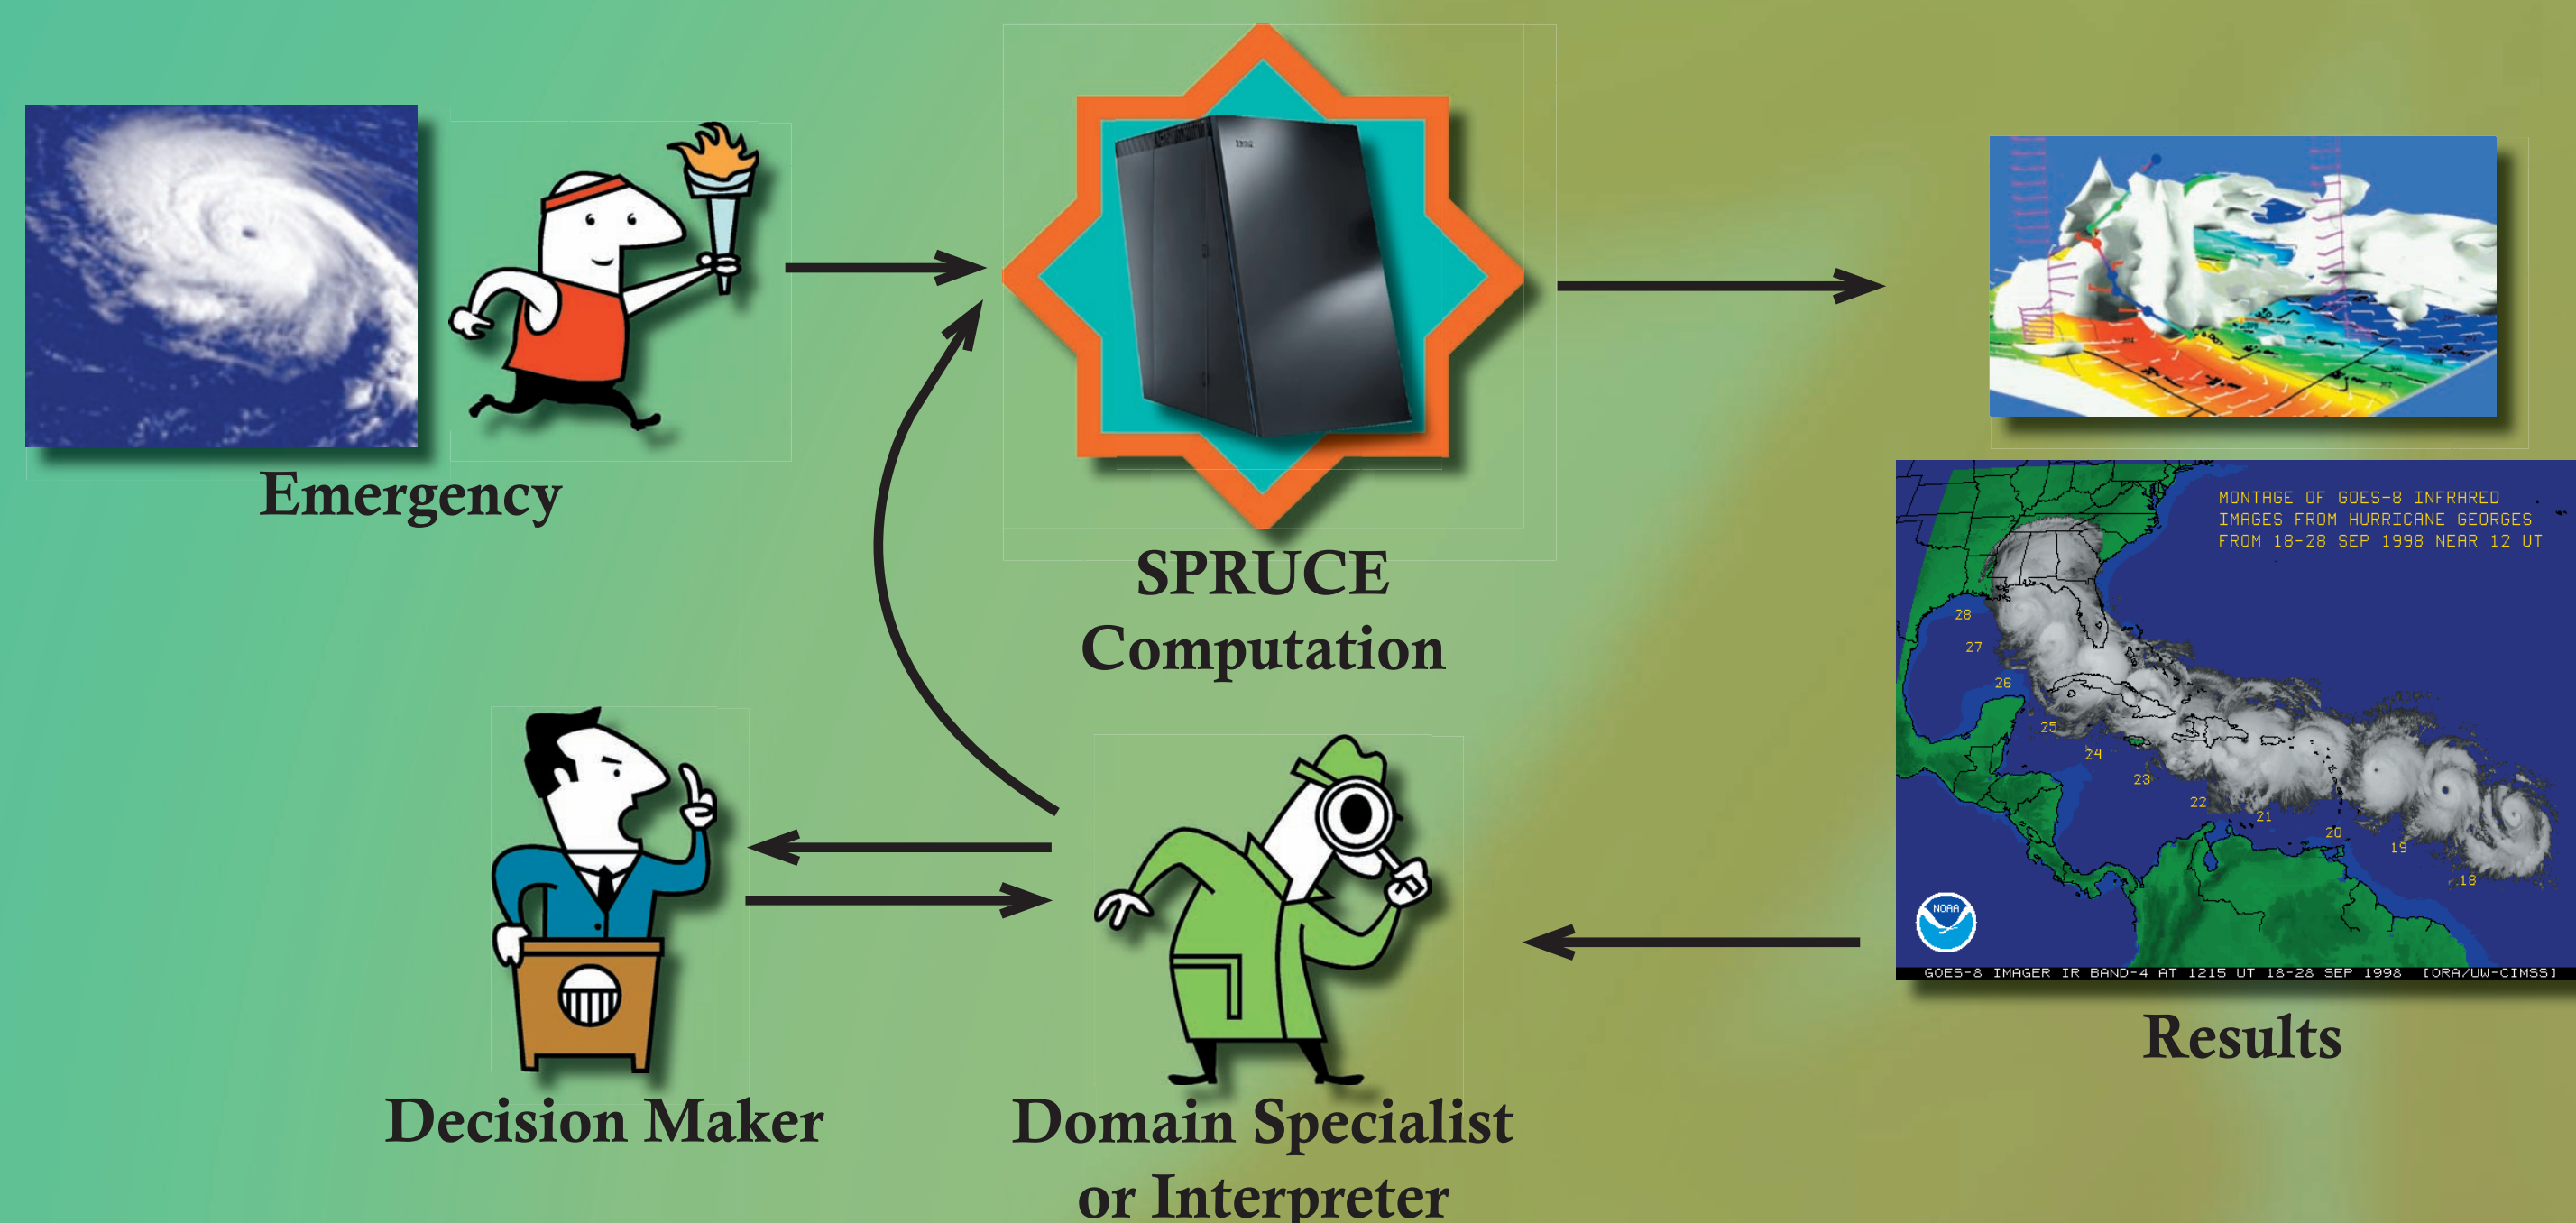


Motivation

Modeling and simulation using high-performance computing are playing an increasingly important role in decision-making and prediction. For time-critical emergency decision support applications, such as severe weather prediction, flood and influenza modeling, late results may be useless. A specialized infrastructure is needed to provide computational resources quickly, automatically and reliably. SPRUCE is a system for supporting urgent computing on both traditional supercomputers and distributed Grids.

- 1 • Should leverage existing resources with minimal dependency
• Must integrate into existing dynamic and automated workflows

SPRUCE uses existing resources to provide high priority for emergency applications, and requires minimal changes to existing workflows. Both Grid based (e.g. Globus) as well as direct submission environments are supported. Work is in progress to provide a web services interface to the SPRUCE gateway.

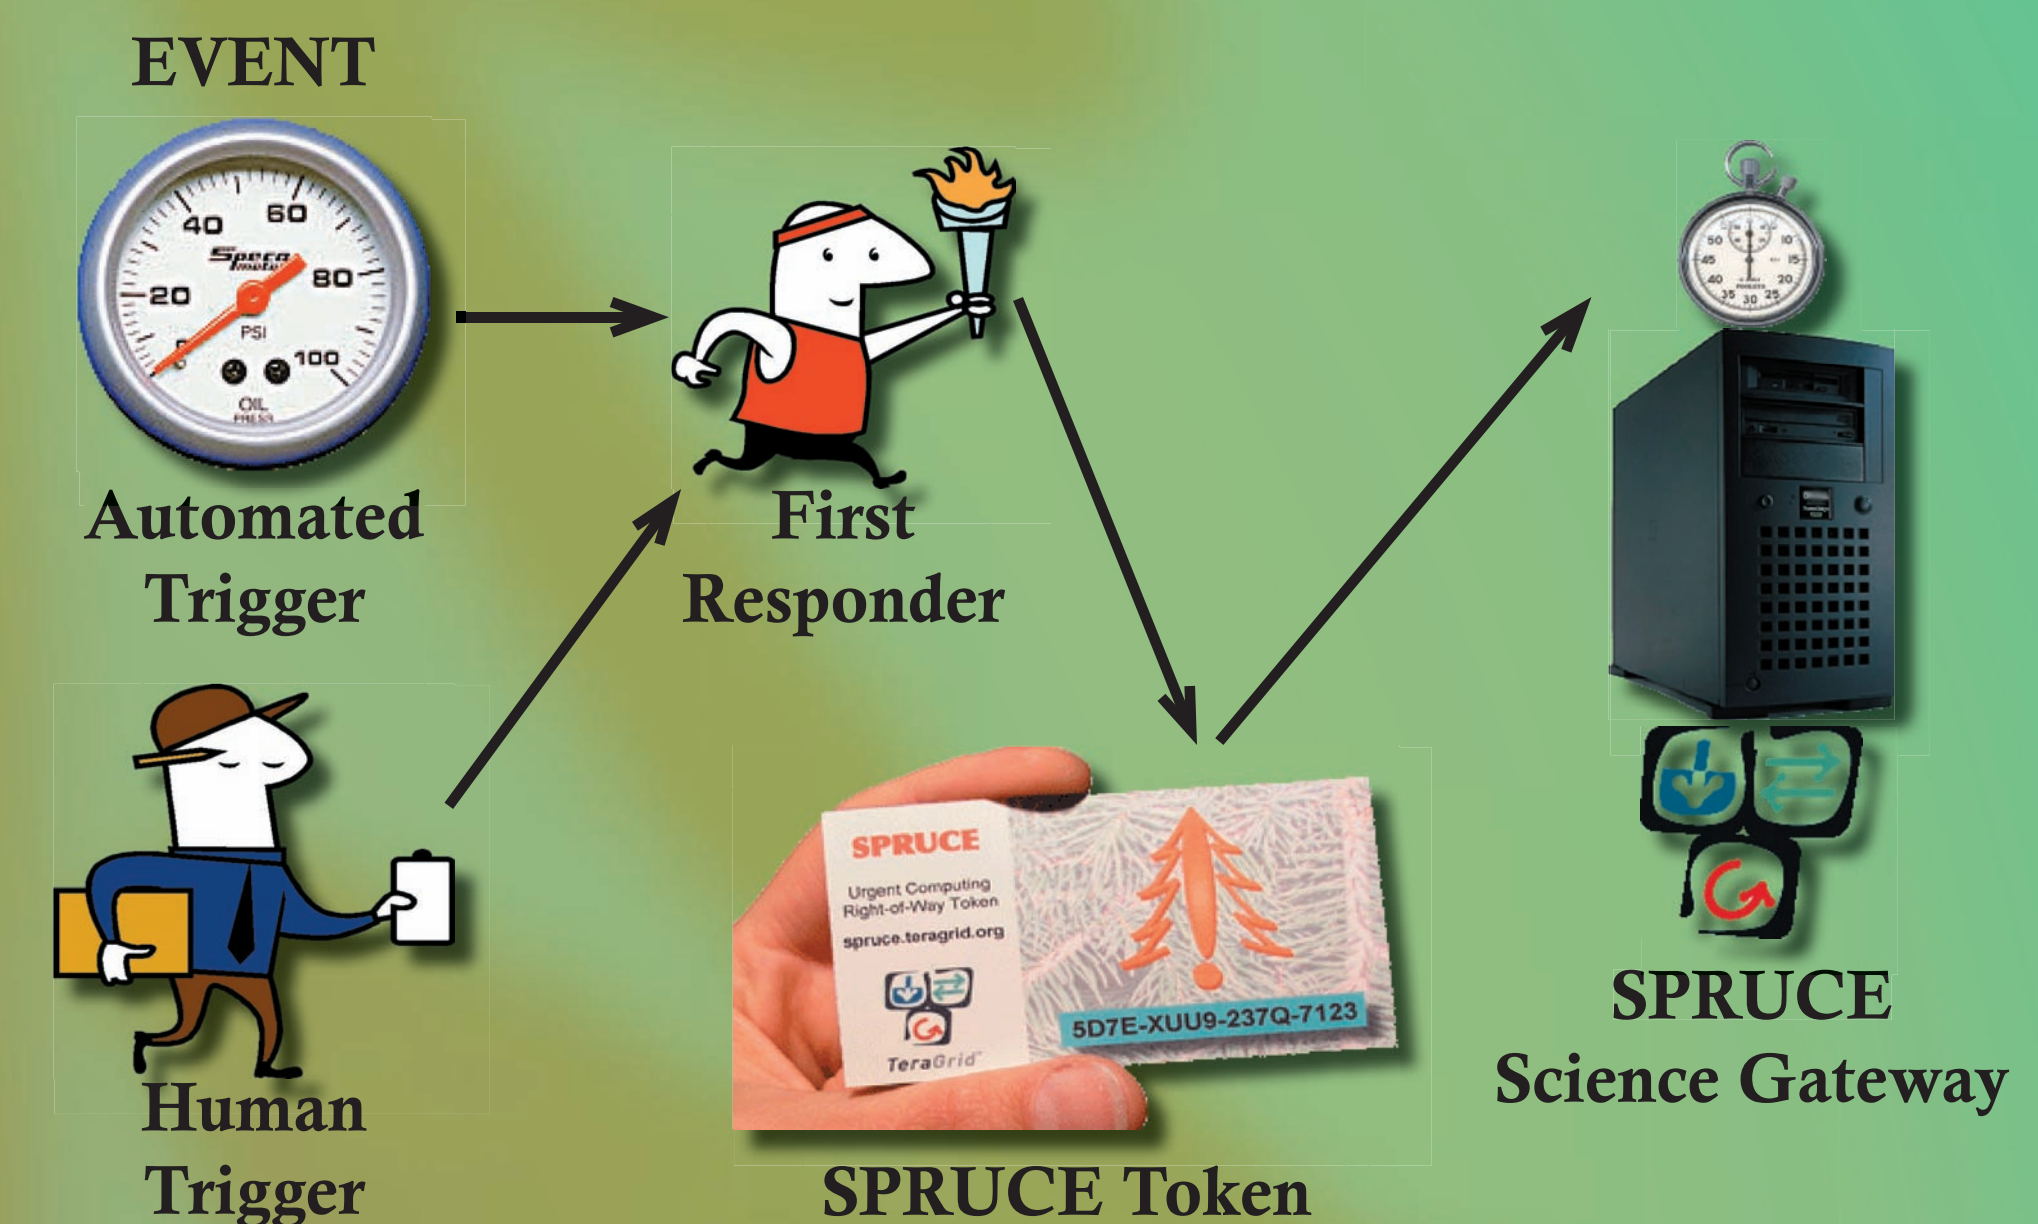


Deployment Status

SPRUCE is a TeraGrid Science Gateway, currently deployed at the UC/ANL, Purdue and TACC resource providers. Work is in progress to extend the system to include the NCSA, Indiana and SDSC resources. As of now, we support Torque and LSF, and are working on adding support for LoadLeveler and PBS Pro. This poster presents requirements for a system supporting urgent computing and illustrates how the SPRUCE architecture meets these requirements.

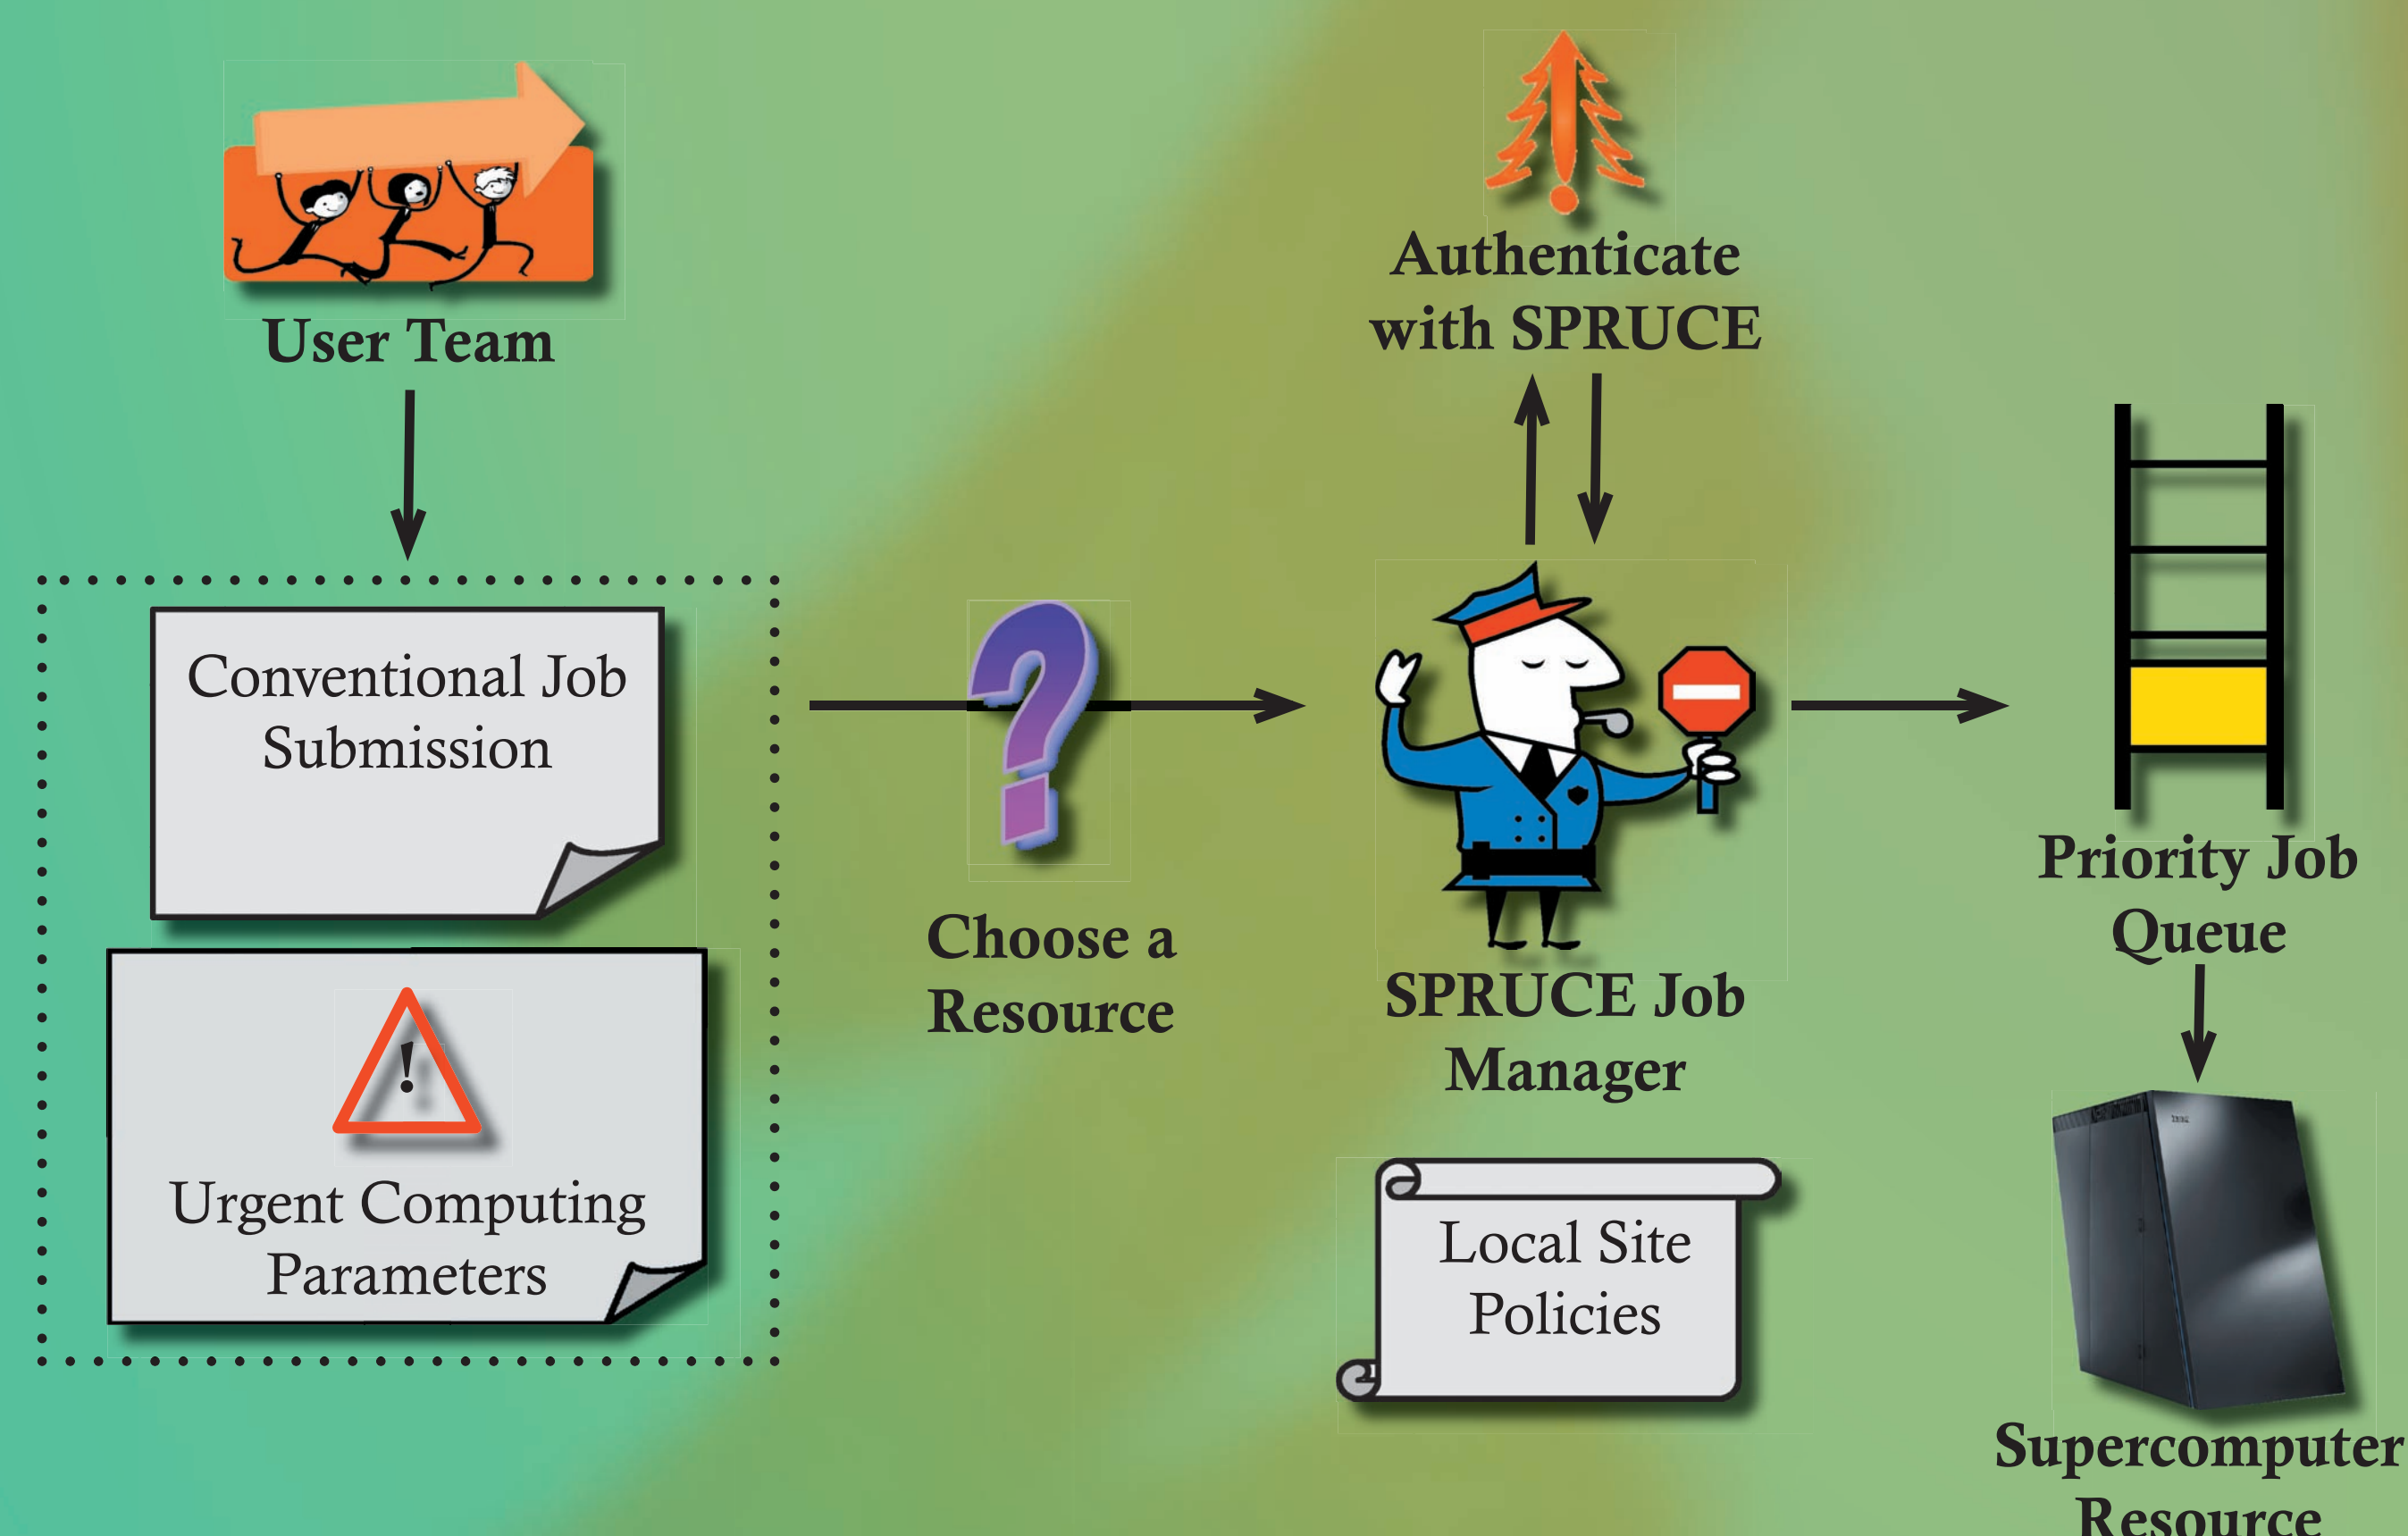
- 2 • Allow for a flexible, portable, and secure urgent access method
• Expose functionality through a Gateway for ubiquitous access

SPRUCE provides users with transferable “Right-of-Way” tokens restricted to specific resources and priorities. In case of emergency, tokens can be activated by the first responder at a Web-based SPRUCE gateway. The token holder associates the response team’s user identities with the activated token, giving them permission to run urgent jobs. Tokens carry an expiration date and a lifetime during which computations requesting priority access can run.



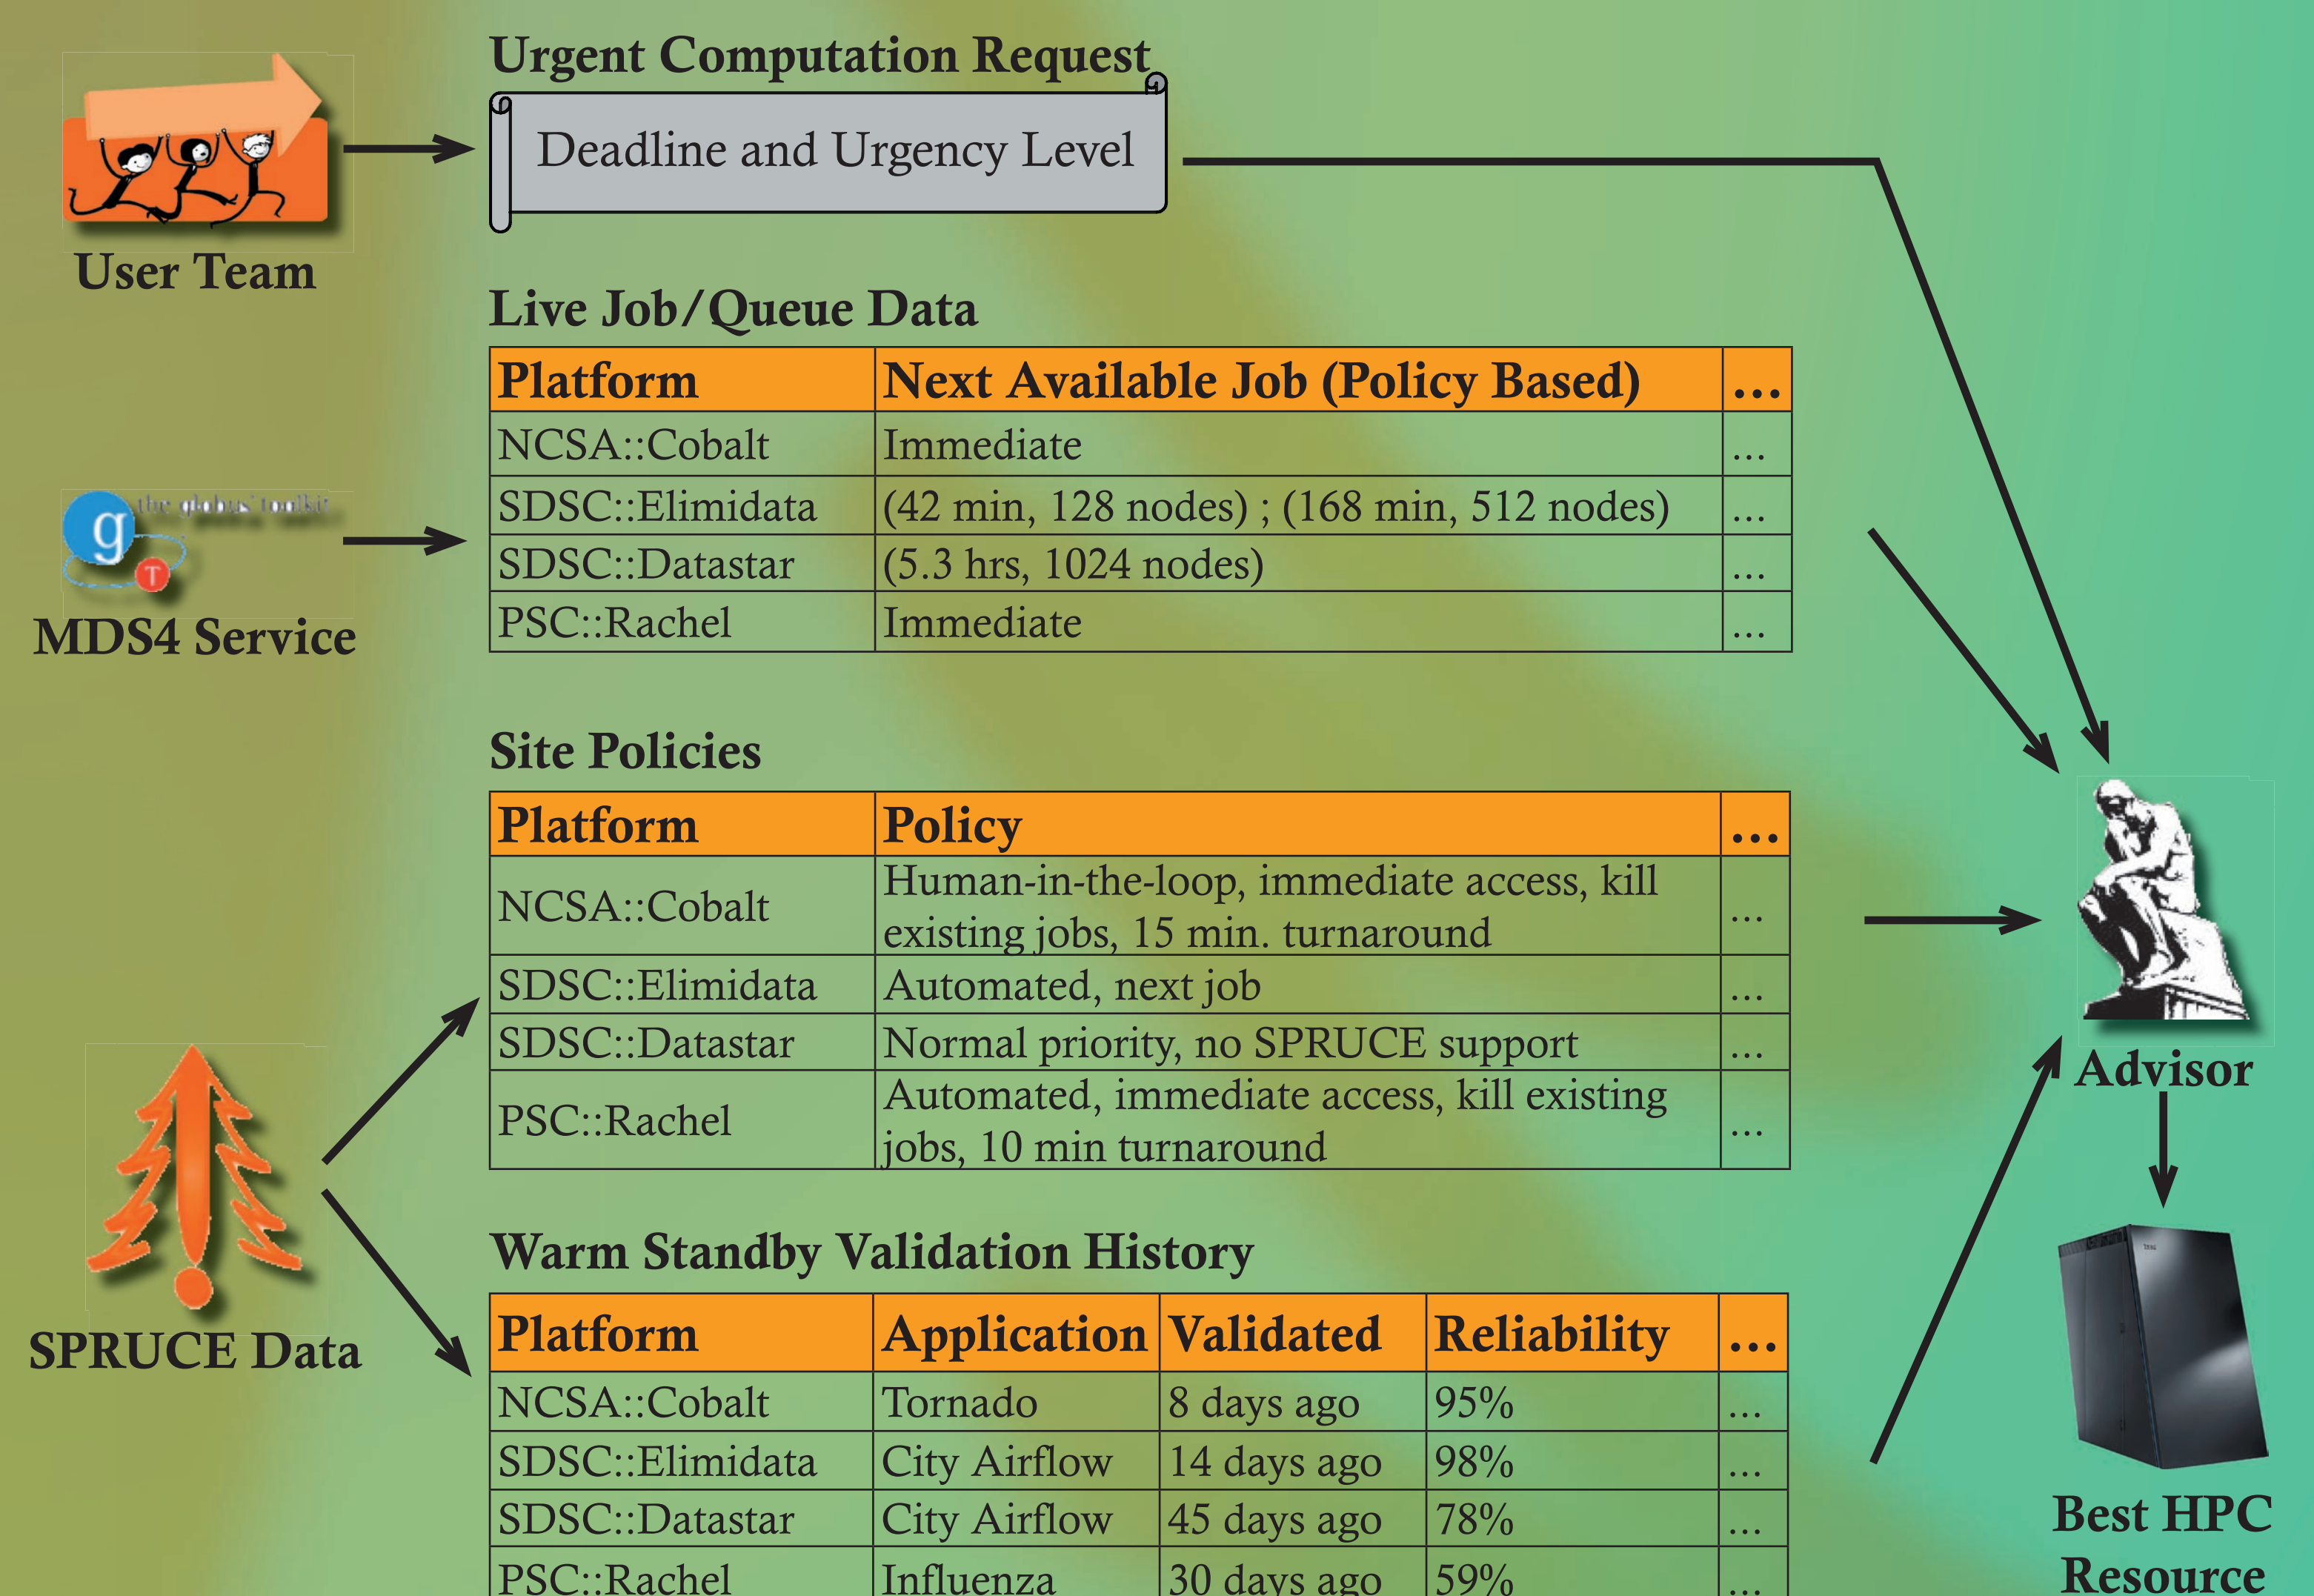
- 3 • Priority job submission must be simple
• Should support resource providers’ policies for handling urgent jobs

Once a user is associated with an active SPRUCE token, an extra ‘urgency’ parameter in the job script is all that is necessary to initiate an urgent job run. On submission, the user’s identity is authenticated against the SPRUCE Gateway and the SPRUCE job manager implements the policy to grant the necessary access. SPRUCE supports a variety of site-specific policies that dictate the resource response. These may include providing “next-to-run” status or preempting other jobs, and may differ for different priority levels.



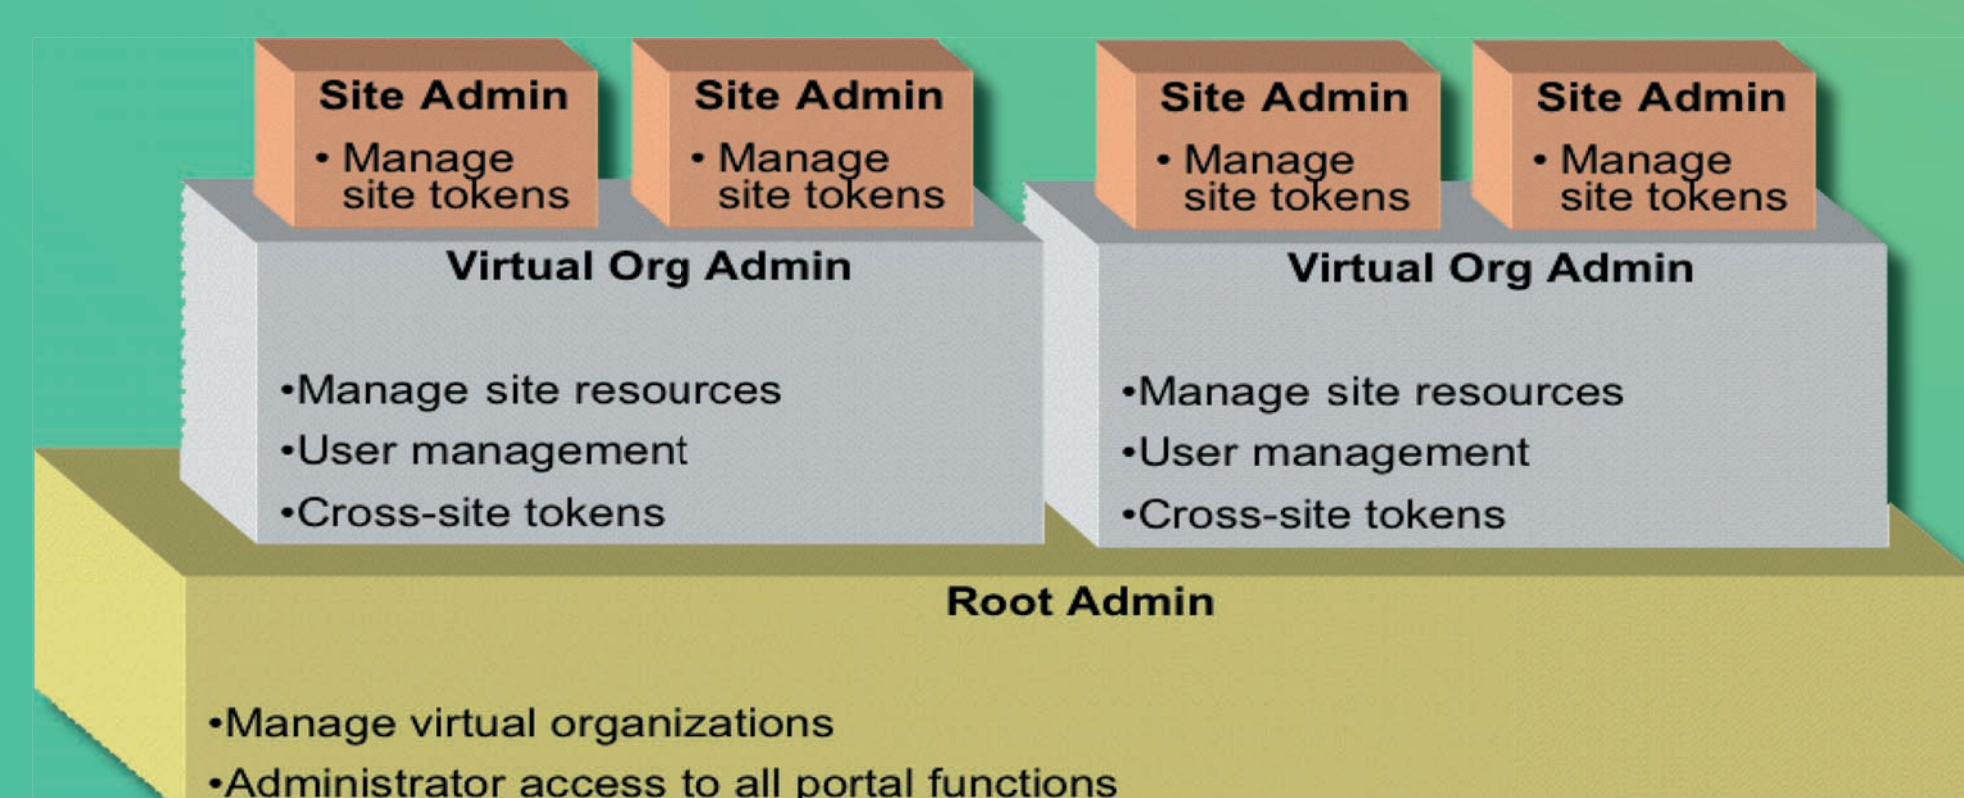
- 4 • Provide best resource selection based on deadline, previous validation history, queue status, site policy, and other information.

An automated “advisor,” currently in development, will help SPRUCE users to determine the best resource based on deadline, queue status, site policy and warm-standby validation history (see 6) at SPRUCE supporting sites.



- 5 • Allow for flexible resource, token, and user management by resource administrators
• Expose complex functionality in a service-oriented way

The SPRUCE Gateway allows for three administrative domains, each of which has token issuing and monitoring powers for their respective resources. Administrators can login and manage tokens, users and usage on their particular site. The Virtual Org (e.g. TeraGrid) administrators can issue cross-site tokens to teams that may require co-scheduling capabilities.




- 6 • Periodic code validation to maintain the user application in ready mode

In the same way that emergency equipment, personnel, and procedures are periodically tested for preparedness and flawless operation, SPRUCE applications are tested in ‘warm standby’ mode to ensure readiness. The codes are periodically validated and a grid wide catalog with results of the validation runs is published for use by the advisor (see 4). This crucial feature is still in development.

SPRUCE Team: Pete Beckman (PI), Suman Nadella, Ivan Beschastnikh, Nick Trebon
This work was supported by NSF SCI-0503697

and by the U.S. Dept. of Energy under Contract W-31-109-Eng-38

<http://spruce.teragrid.org>



THE UNIVERSITY OF
CHICAGO

