

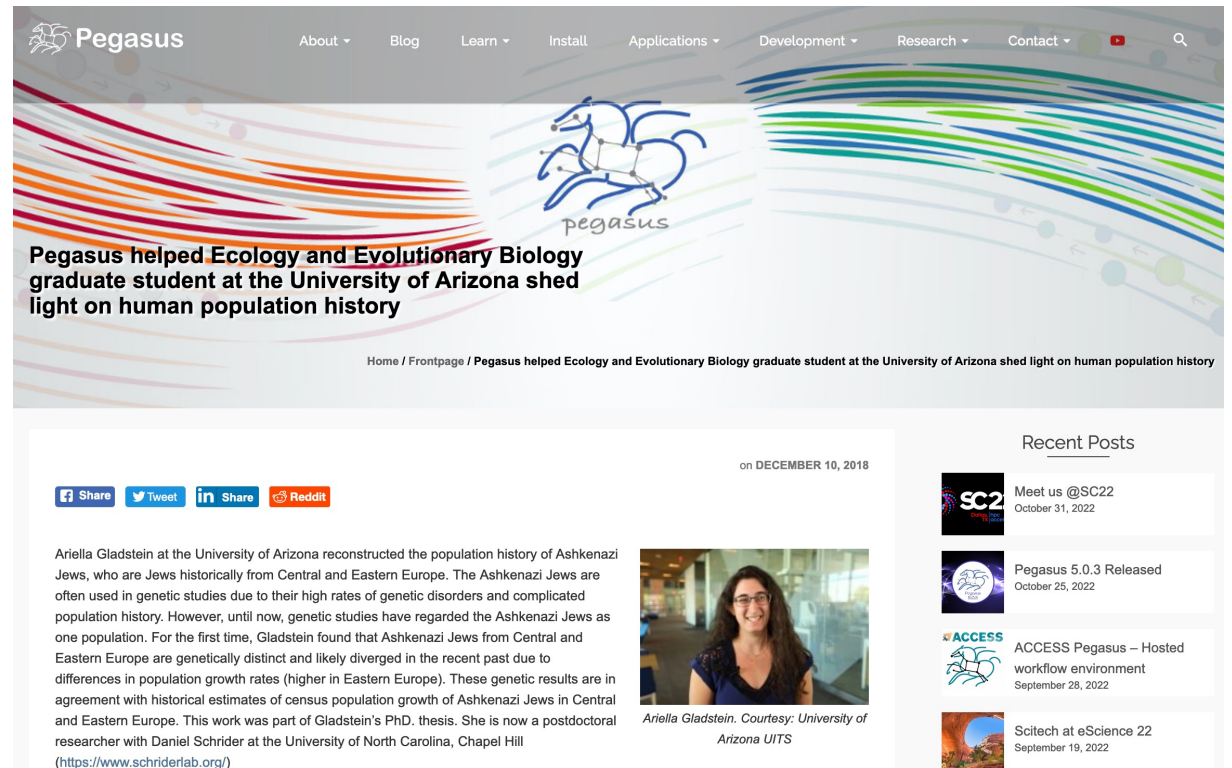
# I Am Not an RSE, but I'm a Believer. Here's Why:

By Blake Joyce  
Data Science Manager  
Research Computing, UAB IT

We run an HPC, a research cloud, and a K8s platform

```
Job ID: 11681113
Cluster: slurm_cluster
User/Group: [REDACTED]
State: CANCELLED (exit code 0)
Nodes: 1
Cores per node: 9
CPU Utilized: 14664-06:34:13
CPU Efficiency: 80184.68% of 18-06:54:54 core-walltime
Job Wall-clock time: 2-00:46:06
Memory Utilized: 71.99 GB
Memory Efficiency: 99.98% of 72.00 GB
```

- A number of jobs per day fail because of code/compute mismatch
- I suspect that our 'need' for larger systems is actually a need for more efficient/stable/reliable code
- We're working through a year's worth of SLURM efficiency data to determine the frequency of failed jobs and lost CPU/GPU hours
- Also want to identify researchers to help/notify this is happening to minimize frustration



The screenshot shows the Pegasus website with a navigation menu (About, Blog, Learn, Install, Applications, Development, Research, Contact) and a search icon. The main banner features a blue Pegasus logo and the headline: "Pegasus helped Ecology and Evolutionary Biology graduate student at the University of Arizona shed light on human population history". Below the banner is a "Recent Posts" section with three items: "Meet us @SC22" (October 31, 2022), "Pegasus 5.0.3 Released" (October 25, 2022), and "ACCESS Pegasus - Hosted workflow environment" (September 28, 2022). A "Scitech at eScience 22" post is also visible (September 19, 2022). The article text mentions Ariella Gladstein at the University of Arizona and her PhD thesis work on Ashkenazi Jews.

- Ariella was genetics PhD trying to run 10+ million simulations
- Legacy code comparing genetics for SNPs
- Code compared using A, T, G, C's
- Julian changed it to be a bitwise comparison of 0/1 for ab/presence of SNPs
- Speed up of 30,000x +