

Tasseography , Scrying , & Cluster Status

As cluster supercomputers have grown from a few nodes to giant systems, monitoring health and performance has become both more important and harder. Seeing cluster status is about as intuitive and effective as reading tea leaves or gazing into a crystal ball. That's the problem **nodescape** is trying to solve.

This isn't a problem unique to clusters. In fact, **nodescape** is really the grandchild of *Firescape*, a system we were creating to help firefighters find their way through burning buildings. The general concept is what we call a *Senscape*: an integrated presentation of multidimensional sensory data allowing a human to understand and use properties of the environment that might otherwise have been beyond human perception.

Seeing Is Understanding

How can one see what's going on in a cluster? How about:



That's an actual photo of one of our clusters, NAK, tinted by **nodescape** as a real-time status display. Nodes are tinted from **blue-to-green-to-red** based on attribute values. As data ages (i.e., no new sensor readings are recorded), the tint increasingly dithers with **magenta**. A node with no data is tinted solid **magenta**.

Here's the little cluster in our SC12 exhibit:

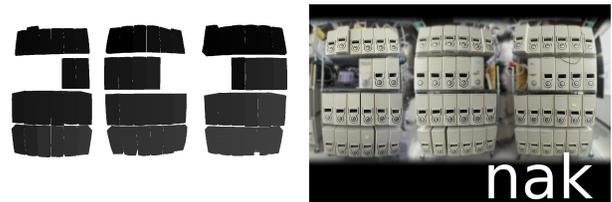


The image of NAK happens to be showing CPU sensor temperatures, but it could be any measured or derived attribute of the cluster. For example, the cluster in our SC12 exhibit is playing multi-voice music with nodes tinted according to the frequency of the note most recently played by each node. Not only can the live status display be put on a video wall (as we do in the front of our machine room), but it also can be accessible via a URL or even painted on the front of the actual machines using a video projector – as in our SC12 exhibit.

How Does nodescape Work?

It's pretty simple. Code running on each node (**epacedon**) throws sensor readings at **nodescape** via short UDP messages. These messages are logged by **nodescape** in a **MySQL** database, much as other health or performance monitoring systems log their data. However, that's not what **nodescape** is about. It's really about analysis and presentation of the data to help humans better understand what's important.

Each new message can trigger derived attribute value updates that may cause actions – such as updates of tinted status displays. The tinting is done by incrementally updating a memory-mapped image. That image is derived using a key image to identify which pixels in a base image correspond to each node:



Are We There Yet?

Although we've been using **nodescape** for over a year now, it is still very much an active research project. Watch **Aggregate.Org/NODESCAPE** for more information.

This document should be cited as:

```
@techreport{sc12nodescape,  
author={Frank Roberts and Henry Dietz},  
title={Tasseography, Scrying, and Cluster Status},  
institution={University of Kentucky},  
address={http://aggregate.org/WHITE/sc12nodescape.pdf},  
month={Nov}, year={2012}}
```