



Evolving with IT to Support Research

Modernizing the campus approach to IT infrastructure for supporting research has never been so urgent.



Damien Eversmann

*Delivery Strategist for Higher Education,
North America Public Sector, Red Hat*

WHILE INSTITUTIONS HAVE EXPRESSED continuing concern about wobbling tuition and ancillary dollars, one source of revenue remains healthy for higher education: COVID-19 research funded by federal and state programs. The full measure, from community colleges to Research 1s, are at the forefront of projects to develop vaccines; uncover the sources of coronavirus and its evolving replication patterns; create new initiatives for public health response; understand the impact of the virus on various populations; study the physical and mental health and learning effects of prolonged quarantine; and explore numerous other facets..

However, the heightened attention on campus research comes with a continuing challenge: how to keep up with IT infrastructure needs, typically assembled once the grant funding arrives. Since many of these recent grants are short-term, turnaround time can be tight. In many cases, research teams are going from near-zero infrastructure to running as quickly as possible – and not just serving applications to users, but storing, processing and sharing astronomical amounts of data.

In addition, the workloads for these research initiatives are constantly changing. It's no longer about starting up an application and letting it run for days, weeks or months. The researchers need to do their data crunching quickly and then move on to the next job – in other words, spin up huge amounts of compute as quickly as possible and then spin it back down.

“Containing” the Challenges

Researchers have pursued two routes in addressing their needs, each with its limitations:

- Going out and buying as much compute as possible to accommodate peak demand. Those resources may then sit unused during the times when researchers aren't running their applications.
- Turning to the public cloud. Even then, not all data

can be maintained in the cloud; some of it has to stay on-premise by virtue of its sensitivity, restriction or regulation.

Neither solution fully allows researchers to put the right workload in the right place at the right time and do so efficiently without wasting IT resources.

Here's where containerization – the next evolution in helping IT make more efficient use of existing technology – comes into play. The container concept can serve as either an alternative or a companion to virtualization. All of the code and its dependencies, such as configuration files and libraries, are packaged or contained. The benefit is portability: The software is abstracted away from the operating system and runs consistently and reliably on whatever infrastructure is available – on-premise, in the cloud or in a virtualized setting.

The container concept can serve as either an alternative or a companion to virtualization. All of the code and its dependencies, such as configuration files and libraries, are packaged or contained.

Hardware that was formerly running multiple virtual machines (each with its own instance of the operating system [OS] and instance of the application) now can host dozens or even hundreds of applications, all sharing the same OS kernel in a lightweight approach. As a result, universities can migrate applications from monolithic systems to containerized applications and repurpose hardware they already have to run those applications more effectively.

Open Source Management Oversight

Since nobody lives 100% in the cloud or on-premise, what's also needed is a platform that accommodates management across the board. **Red Hat OpenShift** has found a home in many schools for good reason: It includes an enterprise-grade operating system (Linux) with a heavy emphasis

on security. At the heart of OpenShift is Kubernetes (aka, “kube”), an open source platform and the de facto standard in container orchestration technology that automates the manual work of deploying, managing and scaling containerized applications. OpenShift clusters groups of hosts running Linux containers and streamlines development and administration. Moreover, multiple OpenShift clusters can be managed via Red Hat’s Advanced Cluster Management for Kubernetes, whether those exist on AWS, Azure, IBM Cloud, some other cloud, on-premise or as a mix.

The advantages of choosing open source can’t be overestimated. As numerous institutions have found, Red Hat’s products and services eliminate costly vendor lock-in and represent the best thinking of the community, which the company participates in, fosters and works with.

Modernizing the campus approach to IT infrastructure for supporting research has never been so urgent. Fortunately, that’s where Red Hat shines – in helping reduce the effort and expense of building, moving and managing workloads, to keep the insights and discoveries coming.

Damien Eversmann is a delivery strategist for higher education, North America public sector, for Red Hat.

ANSIBLE FOR ENTERPRISE AUTOMATION

As colleges and universities have sought ways to make their processes more efficient, a lot of them have turned to the use of **Ansible**, Red Hat’s automation platform, which is relevant for both IT and business users. The idea is simple: Take repetitive tasks that are a time sink and involve infrastructure, applications, networks, security or business functionality, and automate them – whether that’s on-premise, virtual, cloud-based or in a container (or some combination).

There are two ways to deploy automation with Ansible. IT staffers can create tasks for themselves or other users in YAML (Yet Another Markup Language) and run them sequentially in “Playbooks.” Or users can pick what’s needed out of the self-service Ansible Services Catalog.



The Modern Campus: Security, Automation, Hybrid Cloud

Transform the student
experience through innovation.

Download the eBook:
carah.io/TheModernCampus

