

The Benefits of Hyperconvergence in K12

Utilizing Virtualization More Efficiently

A District Administration Web Seminar Digest • Originally presented on February 1, 2018

The next generation of IT infrastructure, hyperconvergence is a framework that combines computing, storage and networking into a single, simplified, automated and easy-to-use system. It's a perfect fit for school districts, which have limited budgets and few staff members to maintain and operate their IT systems. In fact, hyperconverged IT infrastructures—even in large and complex environments—can be run and managed by users with no certifications or training beyond basic computer skills.

In this web seminar, presenters discussed the concept of hyperconvergence in K12, how this type of system can benefit a district, and the keys to successful implementation and deployment—which can take just a few minutes.

Craig Theriac: The common things we hear about are low budgets, a small staff and increased responsibility, including with security and privacy. There is a unique user need in this environment. You have teachers, administrators, faculty and students all accessing the student information systems and other applications, along with all the different levels of restrictions to those applications and different requirements around the use of those applications—but that all require high availability.

You don't necessarily have to understand all of the technical ins and outs to know that it's pretty evident what kind of complexity is involved in an environment like that. The problem comes down to multi-vendor support. Having to try to figure out who owns what little piece of the stack to get something supported is a nightmare, and it's a massive time sink just to keep the infrastructure up and running. HC3 can eliminate that.

Ken Munford: In our original IT infrastructure, we had about seven VMware ESX hosts, with a lot more

throughout the entire organization doing various things. We had 23 bare metal servers running with different roles and different tasks. The facility that had initially put in a lot of this equipment wasn't suited to become a data center, so we thought that at some point we were going to be bursting at the seams.

We faced non-unique challenges. We had aging equipment. We had huge power consumption issues. Every year we were adding new power to the data center to support it. We found that we had a lack of expertise, and we found that when third-party contractors or other vendors would take a look, they'd all wonder why we would do things certain ways and they'd point fingers.

Now we're running a lot of VMs for a district our size. A lot of those were moved from another system into ours—we migrated a lot of the bare metal servers we had into our new HC3 system. We thought that process was going to take weeks and weeks, but we were able to do it within 24 hours, which was great.

Craig Theriac: Scale Computing has been around



CRAIG THERIAC
Director of Product Management
Scale Computing



KEN MUNFORD
Network and Security Manager
Iron County School District (Utah)



The common things we hear about are low budgets, a small staff and increased responsibility, including with security and privacy.

since 2008, and we launched our HC3 product in 2012. The trend of hyperconvergence has been around since about then. The result is very happy customers, with a Net Promoter Score of 91. We've had a lot of success in education, but also in health care, financial services, manufacturing and government, with over 2,000 customers and over 5,500 systems out in the field.

We asked ourselves, "If the goal is high availability, how do we solve for that in a unique way that will provide simplicity for the administrator?" Instead of having a separate storage device somewhere that's using specific storage protocols to connect into, we've moved that storage into the local servers running virtual machines. Then, we've architected a storage layer so that it's redundant so that data that's sitting on this drive is also represented on one of the other drives elsewhere in the cluster. Then you can still have high availability; if this server fails, virtual machines still running on the server will pop up.

You set it up, and initially you'll probably monitor it carefully for a month, just to gain comfort with it.

Then all of a sudden it's been six months since you logged into the system to do anything other than maybe create a new virtual machine.

Ken Munford: We scaled up pretty quickly. We started out with five nodes the first year and about 40 to 50 different VMs. We had some developers on staff who create dozens of VMs for experimenting with the products they're producing. We also created the DR site, which is located in a high school about 20 to 30 miles away from us—that's our primary data center. We have gigabit Ethernet between all of our schools, and that was adequate for us to replicate VMs from our central data center.

We got a lot of benefit out of this. We have a lot fewer physical servers to deal with. Our power consumption dropped by about 90 percent, and the ambient temperature went down, so we didn't have to worry too much about environmental controls. And we can build VMs on the fly—different people from different departments inside of IT can perform their tasks very quickly, very easily.

View, comment, share this story online at DAmag.me/scalehc3

To watch this web seminar in its entirety, please visit: www.districtadministration.com/ws020118