

## Supporting 1,400 End Users with a Staff of 2 at Bloomer School District

by Chris Whitman

As the IT Director for Bloomer School District, my primary job is to support student learning. In that way, my team isn't dissimilar to the teachers. Before every decision we make, every asset we purchase, every tool we provide, we need to think about what's best for the student.

Unfortunately, our old SAN setup just wasn't what was best for *anyone*, let alone the students. Our traditional SAN infrastructure was beginning to show its age. They were older pieces of Dell equipment that, with age, brought reliability issues. On top of it all, they had finally fallen out of support.

We were quickly approaching the time we needed a change, particularly with our storage infrastructure. We even had one instance where our entire system went down—including our wireless and our file server—and I needed to bring in an outside consultant to help get us back up. It was a huge deal.

Part of our issue ultimately came down to our size. We're a growing K-12 district, which is rare since there aren't many districts that are growing in rural Wisconsin. I think that says a lot of things about not only our district, but the community in general. People are choosing to move here when there are much bigger districts down the road

At around 1,200 students, we're not a small district, but we're not so large that everyone is just a number. We're right in that sweet spot, where the staff care about the kids and they have the time necessary to devote the maximum amount of attention to each one.

But while our district itself may be growing, my team certainly isn't—and that's another one of the reasons why aging hardware quickly became a compounding problem. Even today, we have just two people supporting 1,400 end users. As the latter number climbs over the next few years, the former is likely to stay about the same. At that point, there's no such thing as a "small issue." Every performance problem we run into or asset that goes offline represents time we can't devote to supporting student learning.

**As a small staff supporting a growing number of end users, make the flexibility and reliability of your infrastructure a top priority.**

So as we began our hunt for a better option, we knew a new solution needed two qualities above all others: flexibility and reliability. We needed the flexibility to continue to support our district as it grew, and the reliability to free up our team's time.

## **SAN vs Hyperconvergence**

Over the course of our research, two main options presented themselves. We could go for an updated version of traditional SAN, which is essentially the newest model of what we already had, or we could go with hyperconvergence.

On the one hand, we knew what a traditional SAN could do—but it just didn't give us the right level of flexibility without the overhead. The setup was static, which meant it might not fit our growing needs. With flexibility being a high priority, we knew we needed to go with hyperconvergence. But we also had to find the right solution for our small team.

We looked into a few potential solutions and performed a cost analysis. Once we saw that many of the options were in the same ballpark, we began to look at their functionality. One option quickly stood out from the rest: [Nutanix](#).

## **Working with mature platforms means you'll never feel like a beta tester.**

We didn't want to pay for future software subscription costs—if we could avoid it. That made Nutanix's license-free option very attractive. Also, when it came to full-stack support, Nutanix was the most mature solution out there. And as we looked to future developments, Nutanix's disaster recovery (DR) offerings were incredibly exciting.

For disaster recovery, I loved that I could execute DR on a virtual machine level. That level of cloud failover isn't natively part of many providers' tech stack. Or, if it is, it's very complex.

Prior to this, we didn't really have a traditional disaster recovery solution. We had what I refer to as "limited disaster recovery." In other words, if a tornado came through and took out our secondary failover cluster, we would be totally down. I wouldn't consider that to be true disaster recovery, so it was exciting to see that Nutanix could help there as well.

I'm not saying that other companies didn't offer these things. But when I looked across the market to see the extent of what everyone was offering, Nutanix rose to the top in every area that was important. It became clear that Nutanix offered the reliability and flexibility we needed.

## **An Implementation Begins**

Any seasoned IT veteran will tell you, however, that picking a partner is just one step of a larger journey. At that point, you know they talk the talk. Now, you get to see whether or not they walk the walk.

In terms of the actual roll out, I couldn't have asked for a more painless process with Nutanix. We began by moving over VMs that had the least impact on production: supporting VMs and supporting infrastructure. Thankfully, we had no issues, even with older VMs.

The new system ran all of our VMs incredibly well and I've been moving the rest of them over ever since. We still have a few left on that old infrastructure, but once those are done and everything is migrated to Nutanix, we'll turn those old SANs into targets for backup.

Initially, we worked with a Nutanix partner to set up the first cluster. They helped me with some initial migrations and configurations. Beyond that, everything has fallen to me. Amazingly, I haven't had any issues to speak of—I think I may have called support once, possibly twice. Working with the Nutanix partner was helpful because they had done all this before. They knew what to expect when those virtual machines were migrated. After everything was handed off to me, however, things continued to be just as easy.

### **A Single Pane of Glass**

Once everything was up and running, there were a few key things I noticed almost immediately. To start, I loved the quality of the snapshot functionality that Nutanix had to offer. The file-level restoration that was built right into the solution was something I had been clamoring for. We've always had a backup solution, but we've never been able to do it natively (read: easily) before.

Today, I can natively find a file without having to initiate my backup appliance. This saves me a tremendous amount of time on something that's supposed to be simple.

When working with VMware, you always have to worry about what potential problems due to old snapshots. Now, I don't have to worry about snapshots. They even clean and take care of themselves.

Especially when I'm doing server work, I can just take a snapshot in a matter of seconds and it's a fast, effortless process. It makes my life easier because I can easily create a snapshot and work with data exports and imports and file-level restoration.

This is a night-and-day difference from before. Our old process was slow and, often, snapshots would fail for various reasons. Our old infrastructure wasn't redundant, so I couldn't do updates to the cluster because I'd have to check for the replication. It wasn't overly difficult, but it wasn't easy, either.

### **A single pane of glass view means your small team doesn't need to hunt for a solution.**

As someone who doesn't have a massive team, all of this time saving has been a tremendous help. If I want to check on the health of a virtual machine, I now have that one pane of glass. I'm not looking at different managers to troubleshoot an issue. If you have the time to do that,

that's great. But I don't. I want to be able to look at one spot, understand the health of the cluster, the health of our VMs, and immediately look for any bottlenecks.

### **Empowering Our Students**

Overall, we've seen a roughly 75% reduction in our data center by going from 10–12 racks down to 4. On the heat and power side alone, we've cut our footprint down dramatically.

Now, full disclosure: we're still using those remaining racks as a backup target. But I view it as a saving because we're not using it like we would have if we went the traditional SAN route. It was always the plan to use them as backup targets.

With our new setup, the project deployment process has become significantly easier, all because of snapshotting and templates. I can quickly spin up new VMs and if it doesn't work, I can revert to a checkpoint. I even have scripts built up now in Nutanix I can use to pre-configure templated VMs.

To see the impact on deployments, let's look at our statewide testing. Each server runs caching software that caches the test from students' Chromebooks. This software was very resource heavy and led to a laggy experience for students. On something as important as state testing, this was a big disruption.

Previously, I would have to load all of our testing services onto one VM. Now, I have four to work with. Because of data deduplication, I'm not using up any additional storage space. But from a memory and a CPU standpoint—which is where this testing impacts—I can have four VMs spun up. Because of how quick the cluster is, the testing software is lightning fast for students—and I don't have to worry about those CPU or memory overheads like I used to.

Our students' performance on state testing is vitally important. As the IT director, I never want to feel like I'm standing in the way of that. Today, I'm not.

For my small team, our new hyperconverged infrastructure means we can do more with less. With everything in a single pane of glass, I have one go-to dashboard to troubleshoot and maintain our most vital services. Looking to the future, I'm very excited about our options for DR.

And as our district grows, I know my small team can support our needs. Instead of putting out fires, I look forward to helping light a fire within our students.