Hewlett Packard Enterprise

Case Study

HUDSONALPHA
INSTITUTE FOR BIOTECHNOLOGY

Objective
Implement a new platform to cope with increased workloads

Approach
Leveraged existing Hewlett Packard Enterprise relationship

IT Matters
- Delivered as a fully integrated system that can be up and running in 20 minutes
- Based on highly configurable HPE ProLiant DL380 Gen9 servers
- Increases number of storage, network, compute and graphics options

Business Matters
- Eliminates production bottlenecks caused by 'almost-ready' code
- Reduces initial CAPEX costs with easy scale out options
- Delivers ease-of-use through an intuitive interface

HudsonAlpha streamlines IT processes to meet growing workload

HPE Hyper Converged 380 clears development bottlenecks for leading research institute

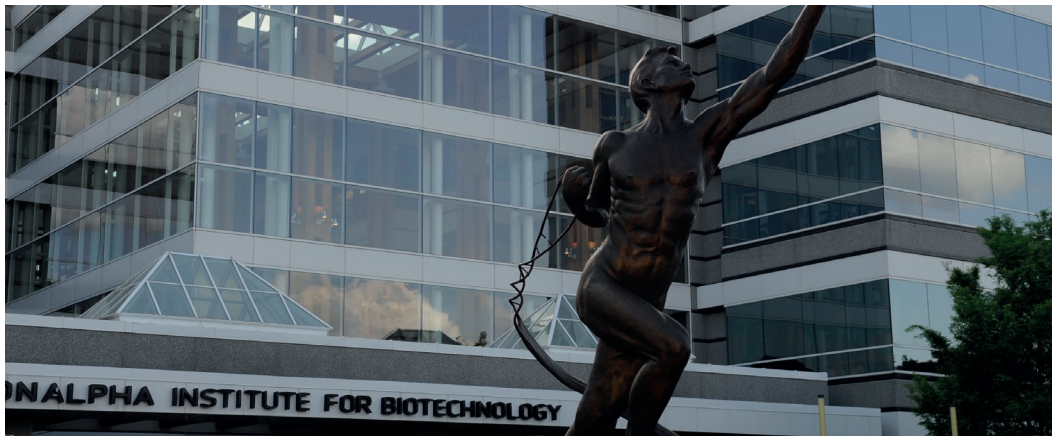

The data-intensive work carried out by one of America's leading genome research institutes, HudsonAlpha, is constantly growing. It needed a powerful and easily managed platform that would deliver the flexibility and scalability to cope with increased demands. Hyper convergence was the key with the HPE Hyper Converged 380 pre-integrated, configurable hyper-converged server.

Challenge

Growing workload

In today's organizations, growth constantly demands more from IT. Infrastructure has to keep up with fast moving environments but deploying and managing traditional platforms uses too many tools that are complex and hard to maintain, placing even more burdens on IT staff.

One organization that knows only too well about the pressures of coping with large and growing workloads is HudsonAlpha Institute for Biotechnology, a leading American genome research institute. For the past six years, its data production has doubled every six months and the amount of data is now growing at 1PB a month.

"We have recently opened a clinic for children born with conditions that can't be diagnosed and often we can save the child's life."

– Jim Hudson, co-founder and chairman of the board, HudsonAlpha Institute for Biotechnology

"The scale of our growth has become larger than most enterprise systems are built to cope with because at HudsonAlpha, we not only have growth in the number of people but also growth in the amount of work each of them is doing," says HudsonAlpha chief information officer, Peyton McNully. "Every single one of our 100+ researchers are doing something that we generalize as genomics but is very specific to their research, so there is a nuance difference in every one of these workloads. That means we have to push our IT tools to the furthest edges of the organization so that staff are able to be more engaged in their own work and we have to provision our resources accordingly. Scale is the challenge and an even greater challenge is the variability of resources we need."

HudsonAlpha Institute for Biotechnology is based on a 155-acre campus in Huntsville, Alabama and is a leading center for genomic research in the USA. Its work in this area of biotechnology has made life-saving discoveries that enable the diagnosis and treatment of human diseases, as well as leveraging agricultural genomics to feed and fuel the world. In addition, the institute runs an education outreach program that educates 100,000 children a year from Alabama and nationally in genomics and genetics.

"The research that happens at HudsonAlpha results in some of the latest and greatest findings that are going to translate into better care, better crops and a better quality of life for humans all over the world," says McNully.

With 200 individuals working on the non-profit side and 600 in 30+ associate life science companies housed on the HudsonAlpha campus, huge amounts of sequencing data are produced.

"Data is our biggest challenge. We generate over 1PB of data a month that needs to be stored, manipulated, computed and so forth," says co-founder and chairman of the board, Jim Hudson. "Managing that data and being able to query it is very essential for us. We analyze it on large computers but we needed a platform that would allow us to very efficiently bring all this together and take advantage of all the resources that we have in a much better way than we've been able to do in the past."

Solution

Tight integration of resources

HudsonAlpha decided that hyper-convergence was part of the answer to its challenging growth. Hyper-convergence is a software-defined architecture that tightly integrates compute, storage and virtualization resources in a density-optimized platform that is centrally managed.

Case study
HudsonAlpha Institute for Biotechnology

Industry
Biotechnology

As users transition to a software-defined data center (SDDC), hyper-converged products serve as self-contained, modular building blocks that can handle changing workloads and accommodate new business.

Because of its existing relationship with Hewlett Packard Enterprise and following a visit to HPE Discover in Las Vegas, HudsonAlpha chose the all-in-one compute, storage and virtualization platform, HPE Hyper Converged 380.

"At Discover, I was better able to understand the strategy in what's taking Hewlett Packard Enterprise forward into transformation areas, and where those can be applied," says McNully. "The entire organization is now focused on delivering changes to an IT environment that have business value and that's hugely important because there is a lot less education that has to take place between the sales organization and the IT organization here at HudsonAlpha."

Described as a Virtual Machine (VM) vending machine, HPE Hyper Converged 380 is a configurable, scalable and agile hyper-converged virtualization system that delivers a simple solution stack with extended flexibility and manageability.

It builds on the powerful, industry standard HPE ProLiant DL380 Gen9 server platform and is combined with VMware vSphere®. Using the new HPE OneView User Experience (UX) adds full lifecycle management, VM provisioning and updates in a single pane of glass.

Designed for the software-defined data center, the HPE Hyper Converged 380 enables a standardized approach to virtual server deployment and is available in three workload configurations: General Virtualization, CloudSystem and a Virtual Desktop Infrastructure (VDI). It can be customized at the time of order with all hardware and software components pre-installed and pre-integrated at the factory. Quick customization is available using the HPE OneView UX software which is designed to do away with need for a manual.

"The fact that the HPE ProLiant DL380 is something that we are all familiar with and able to use on existing workloads takes a lot of the magic out of pushing towards a hyper-converged environment which enables us to do away with a lot of very expensive and elaborate SAN environments," explains McNully.

Case study
HudsonAlpha Institute for Biotechnology

Industry
Biotechnology

Customer at a glance

Hardware
- HPE Hyper Converged 380
- HPE ProLiant DL380 Gen9 server

Software
- HPE OneView

"We have created a pre-production sand box with HPE Hyper Converged 380 and I think that productivity across the board is far better because production is no longer clogged with almost-ready code. We are now writing something once and using it multiple times."

– Peyton McNully, chief information officer, HudsonAlpha Institute for Biotechnology

HudsonAlpha's solution incorporates HPE Synergy which brings compute, network and storage infrastructures together as a single platform with integrated management to give a single API so companies can use that one API to program their whole infrastructure in a very automated manner. It is used with the converged management tool, HPE OneView.

Benefits

Improved deployment process
HPE Hyper Converged 380 is enabling HudsonAlpha to streamline the processes for dealing with its increasing workloads. It is supporting the IT team's aim to achieve version control of both software and infrastructure in GitHub, the web-based Git repository hosting service that offers distributed revision control and source code management. The HPE OneView service portal means that every developer visiting GitHub will be able to see what is running where, when and how and will enable them to integrate with that. In this way, the developers will increase their self-service and having set up the system, IT can step back, reducing management time and effort.

"Our best use of HPE Hyper Converged 380 so far has been in ultimately using it as a development test environment because it's really similar to what a user will actually have in a production environment on HPE Synergy. It's also functionally simpler for a developer to hit the FENIX web server user interface window, spin up their instance, log in, conduct their test and then deploy their workload to a production environment whenever they see fit," says McNully.

"That's been useful because if there are 33 developers working that day, there are 33 fewer tickets plus the added benefit that if that developer has not loaded up all the dependencies or 'boot strapped' their code they are getting a vanilla OS layer and they get to drag all of their dependencies into that OS at start-up. If they can't do it in that environment, then they can't do it on the HPE Synergy instance as well so they are not ready to send it to a production workload. This means that we are not clogging up production with almost ready code."

Learn more at
hpe.com/info/hyperconverged

Sign up for updates

4AA6-8183ENW, February 2017 Rev.1