



Big Data Meets Software-Defined Storage at EngineRoom.io - the BigData Company

EngineRoom.io
Sydney, Australia
Technology



System configuration

Storage Software: Two 100TB NexentaStor HA Clusters
CPU: Dual 8 Core
SSD: Write Cache: Fusion IO IO Accelerators
Read Cache: 12TB of SSD Cache
HDD: 64TB of 2.5" 15K SAS / 32TB of 3.5TB 7K SAS
Networking: 40GB Brocade

Business Overview

Having a large data set and deriving value from that data are two vastly different processes. For this reason, Big Data is an asset that many businesses have but don't know what to do with. That's where EngineRoom.io[®] comes in. EngineRoom.io is a leading provider of Big Data Platform as a Service for enterprise businesses.

The company, founded in 2011 by Stefan Gillard, provides cloud services for the entire Big Data pipeline. Businesses from all industries come to EngineRoom.io because the company offers a secure private cloud upon which it runs High Performance Data Analysis. The term 'Big Data' means different things to different people, but at EngineRoom.io they fundamentally believe that Big Data means deriving value and competitive advantage from within a businesses' data streams.

The company, headquartered in Sydney, Australia, focuses on providing Big Data as a service; the means to aggregate structured and unstructured data, to handle high-velocity streaming data, to process large datasets at scale, and to provide advanced visualization,

consulting, analytics and modeling expertise.

Their customers faced a build or buy decision, but ultimately didn't see the need or the economic viability of creating their own high-performance data analysis infrastructure. Not when they can turn to a reliable, enterprise-focused cloud-based Big Data service.

The service is built upon virtualized cloud infrastructure to offer not just compute cycles, but a full range of tuning and other services to enable effective and successful use of cloud for Big Data and High-Performance Computing (HPC) workloads.

Problem

As the business began to grow in early 2013, the company faced a critical build or buy decision of their own with respect to their storage needs. Given their business model, EngineRoom.io had to deal with massive amounts of data; but data that also fluctuated wildly based on platform usage. If you're going to collect terabytes, and even petabytes, of business-impacting information, it needs to be stored properly.

“We looked closely at between twenty and thirty different storage solutions, or hybrids thereof,” according to CTO Nicky Ladas.

“We weren’t satisfied with any of them for a wide variety of reasons, including TCO, vendor lock-in, scalability, and more. Our preference for a flexible, tunable, scaleout storage file system. After extensive R&D and production testing at scale, we finally decided on Nexenta’s Software-Defined Storage offering.”

Big Data and Software-Defined Storage go hand in hand. With Big Data, companies can generate breakthrough insights and analysis like never before but are hindered by how to store and extract the maximum value from vast volumes of raw, unstructured data sets. Software-Defined Storage abstracts storage software from the underlying hardware, reducing storage infrastructure costs by moving to commodity hardware and freeing users from hardware vendor lock-in.

EngineRoom.io was looking for a solution that met the following criteria:

Time to Market: EngineRoom.io needed a solution they could implement quickly. They had an opportunity create a foothold in the Big Data as a Service market, but couldn’t wait for the next generation of storage hardware to hit the market. Nexenta’s software-only model fit the bill.

Simplicity: The product had to be well designed, mature and easy to use.

Licensing: EngineRoom.io needed a flexible licensing model that allowed them to grow at their own pace.

Support: The company wanted advanced technical support; a company they could rely on rather than an outsourced call center.

“At the end of the day,” added CEO Stefan Gillard “Nexenta was the only offering that checked all the boxes. NexentaStor provides us with a solution that is high-performance, scalable, and with tremendous linear scale out IOPS capabilities.”

Solution

EngineRoom.io is now running NexentaStor with the HA plug in and managing upwards of 100TB of data in two discrete stacks.

NexentaStor is running EngineRoom.io’s VMware’s data stores, and high performance

multi-tenanted production storage. Networking is facilitated via both Brocade and VMware via 10GB trunks. The companies’ EngineBlock

storage platform utilizes large capacity RAM footprints for ARC, FusionIO as local cache accelerators, large format 800GB SSD’s for pool caching and 1.2TB 10k SAS drives.

“It makes our job much simpler to know that NexentaStor can act as a management framework and work with both physical and virtual environments simultaneously from a single console,” said Ladas. “This was a very attractive feature to us.”

In addition, NexentaStor is helping EngineRoom.io with additional capabilities such as inline de-duplication, unlimited snapshots and cloning, transparent compression, ability to leverage global name space features, hardware-agnostic solutions, and high availability support.

Benefits

“We made the decision to move ahead with Nexenta because they met our somewhat unique list of requirements we had in place,” said Gillard. “We were pleasantly surprised, however, when the implementation came in at one-twentieth the cost of competing solutions.”

“There are now hundreds of hours that we no longer have to dedicate to administrating our storage infrastructure. We’ve taken those hours and focused more on R&D and innovation,” according to Gillard.

Because EngineRoom.io works as a service, that means the company has its footprint in multiple industries – and the data particular to those industries. Media and entertainment, life sciences, financial services, oil & gas, engineering and more all have one thing in common: they are producing lots and lots of data. Some studies are claiming that 2.5 quintillion bytes of data are created each and every day. Regardless of whether the data is comprised of medical images or visual effects, and regardless of whether the data is structured or unstructured, EngineRoom.io now has it covered.

“We knew that if we wanted to succeed in the Big Data world, we had to resolve the complexity and cost of storage management,” said Gillard. “Nexenta not only helps us make those challenges disappear, the sophistication and economics of NexentaStor gives us a significant competitive advantage.”