Think Global, Act Local!

Case study: Caringo Swarm empowers NEP CDN, showing performance and reliability combined

Prepared for Caringo
December 2016

TABLE OF CONTENTS

INTRODUCTION

Recently, I had the pleasure of talking to Gerbrand de Ridder, Head of R&D and Lead System Architect of NEP The Netherlands. From the Oscars to the Olympics, NEP provides the technology, the people and know-how to enable its clients to produce the world's biggest live and broadcast events. NEP's unparalleled worldwide network gives them the largest and deepest pool of outsourced broadcast engineering and production talent in the industry, providing solutions around the globe for Remote and Cloud Production, Studio Production, Video Display, Host Broadcasting, Post Production, Media Asset Management, Video on Demand, Digital Media Services, Playout and Uplink Communication.

NEP's solutions are developed by award-winning in-house integration teams such as Gerbrand de Ridder's and supported on-site and in real time by highly trained technical engineers who deliver customized solutions, superior service (24/7) and flawless execution to clients at every stage of the production cycle.

Gerbrand de Ridder's team designed and deployed NEP's CDN (Content Distribution Network) in the Netherlands, leveraging a large scale-out object storage platform, Caringo Swarm. The CDN back-end is a common use case for object storage but, in this particular case, I believe NEP The Netherlands has found a clever way to use object storage and the key to delivering a great balance between performance and reliability of the system.

THE SCENARIO

NEP is a worldwide organization with more than 2,000 employees that can manage any aspect of studio and live event video production, as well as the management and distribution of high-value media content. With an end-to-end service (glass-to-glass) portfolio, NEP has been chosen by leading broadcasters in over 65 countries on all seven continents.

> *"In 2010 we chose Caringo Swarm because it was the most flexible platform for us and had a good per TB licensing, but we are sticking with it because of its flexibility and the continuous improvement in existing and new features"*
>
> *Gerbrand de Ridder, head of R&D*

In recent years, due to the major shift by the Media & Entertainment industry towards digital content production and consumption, NEP has developed the most important Content Distribution Network in the Netherlands backed by Caringo Swarm object and cloud storage platform. The platform was born in 2010 with less than 100TB of available capacity and has been growing since then at an impressive pace of 20% to 30% per year—projected to reach 1.25 petabytes and more than 200 nodes before the end of 2016.

With a total capacity of 1.25 PB and more than 200 active nodes, Caringo Swarm is capable of delivering up to 45GB/sec of throughput, making it a key component to deliver that high-end quality service that NEP has always aimed for.

NEP's CDN serves more than 650,000 user accounts accessed with different consumption methods, ranging from unlimited monthly subscriptions to pay-per-view for programs from organizations such as FOX Sports, National Public Broadcaster (NPO) and RTL Videoland.

A large Swarm cluster configuration is the key and active component of NEP's CDN platform. In fact, at the beginning, it was only used to store all available content, given its durability and reliability. Now, the least-viewed videos are usually streamed directly from Swarm before being cached on the edge network appliances if demand increases. Internal benchmarks at NEP have shown that the Caringo Swarm cluster is capable of delivering up to 45Gb/sec of throughput, making it a key component to deliver the high-end quality service that NEP has always aimed for.

WHY OBJECT STORAGE?

Object storage brings several benefits when compared to other storage technologies. For example, the erasure coding available in Caringo Swarm is used extensively by NEP to obtain the best level of space efficiency and robust data protection at the same time. Most video content is already compressed and further compression techniques like deduplication or LZW compression, usually available in NAS systems, do not significantly improve space utilization. Additionally, traditional RAID represents a significant threat of data loss in case of multiple failures and long rebuild times due to the large size of modern disk drives.

With more than 200 nodes from two different hardware vendors, the flexibility and scalability of Swarm object storage, together with the software-defined approach, has allowed NEP to build a platform that is perfectly tailored to its requirements in terms of capacity, cost, throughput, growth and reliability, coupled with full hardware independence. Old and new hardware systems can live together in the same cluster. This allows the end user to choose the best technology available and to purchase software and servers separately, taking advantage of different depreciation cycles as well. Such a system, even though large and critical for NEP's business, does not require a full-time

"We don't have a full-time Sys Admin working on our object storage system. The only issue that usually happens to the system is that if a node or a disk fails, we need to replace or repair it."

Gerbrand de Ridder, head of R&D

Storage Administrator and the most common activity on the system is for repairing or replacing damaged disks and nodes. Even demanding system upgrades, like capacity expansion activities, are performed with minimal user intervention by adding new nodes to get data rebalanced across the cluster.

WHY CARINGO SWARM

After a long evaluation process that included several vendors, Caringo Swarm was selected for its infrastructure flexibility and an attractive licensing model. This has not changed over the years, but now NEP has also found a partner in Caringo who listens to needs and feature requests while developing new and innovative enhancements such as the SwarmNFS converter and FileFly, which simplify the use of multiple protocols. In fact, even though a large part of data accessed at NEP is performed via API, there are still some legacy application components that require a file-based access—a common scenario across all types of use cases.

"We can talk to Caringo's engineers and they listen to us about our needs and feature requests."

Gerbrand de Ridder, Head of R&D

Caringo Swarm has also demonstrated its ability to scale not only in capacity but also in performance. NEP has several versions of each piece of content to serve different types of devices and bandwidth. Therefore, the object store has copies of hundreds of thousands of files of different sizes which, when not yet cached, must be streamed directly to the end user device. Despite the complexity of streaming content of varied sizes, Caringo Swarm has been able to provide an impressive 45Gb/s, which is likely to improve after implementation of further infrastructure optimizations currently planned to increase the bandwidth availability.

BOTTOM LINE

One of the most common use cases for object storage is near-line archiving. Content Delivery Network infrastructures usually take advantage only of its scalability and low $/GB for their storage backend. NEP took a step forward. Thanks to the brilliant design of their infrastructure and Caringo Swarm's capabilities, it is able to provide top-notch services to its clients and end users. Streaming the least-viewed videos directly from the object store and being able to pre-fill cache servers at high speeds makes the full content available immediately to users while improving overall infrastructure efficiency.

Although developing and building a CDN infrastructure is still a job reserved only for large media companies or cloud providers, there is no doubt that the proliferation of devices capable of recording high-quality videos is producing a huge amount of data

Large enterprises, without building a complex infrastructure such as a CDN, could benefit by using an object storage solution like Caringo Swarm, to store more data safely, sustain growth and stream content to end users directly when required.

that must be stored safely and then quickly accessed as needed. Large distributed enterprises, without having to build a complex infrastructure such as a CDN, could benefit by using an object storage solution like Caringo Swarm to store large amounts of data safely and to sustain growth and stream content to end users directly when needed.

Caringo has proven that with the right configuration it is possible to provide content at high throughput from large, cost-conscious storage systems.

JUKU

WHY JUKU

Jukus are Japanese specialized cram schools and our philosophy is the same. Not to replace the traditional information channels, but to help decision makers in their IT environments, to inform and to discuss the technological side that we know better: IT infrastructure virtualization, cloud computing and storage.

Unlike the past, today those who live in the IT environment need to be aware of their surroundings: things are changing rapidly and there is a need to be constantly updated, to learn to adapt quickly and to support important decisions—but how? Through our support, our ideas, the result of our daily global interaction on the web and social networking with vendors, analysts, bloggers, journalists and consultants. But our work doesn't stop there—the comparison and the search are global, but the sharing and application of our ideas must be local and that is where our daily experience, with companies rooted in local areas, becomes essential in providing an honest and productive vision. That's why we have chosen: "think global, act local" as a payoff for Juku.

AUTHOR

Enrico Signoretti is an analyst, trusted advisor and passionate blogger (not necessarily in that order). He has been immersed in IT environments for over 20 years. His career began with Assembler in the second half of the 80′s before moving on to UNIX platforms until now when he joined the "Cloudland". During these years his job has changed from highly technical roles to management and customer relationship management. In 2012 he founded Juku consulting SRL, a new consultancy and advisory firm deeply focused on supporting end users, vendors and third parties in the development of their IT infrastructure strategies. He keeps a vigil eye on how the market evolves and is constantly on the lookout for new ideas and innovative solutions. You can find Enrico's social profiles here: http://about.me/esignoretti