

Scalability by Nuxeo



Copyright ©2018 Nuxeo. All rights reserved.



Introduction

Your DAM shouldn't hold you back

Nuxeo has the privilege of being a partner to some of the most forward-thinking—and demanding—organizations on the planet. These organizations operate on the leading edge of feasibility. Working with these customers has given us a unique window into the future of the market, and into the needs that leading-edge companies have today... and the rest of us will have tomorrow.

One of our biggest takeaways: Scalability matters. To propel enterprises forward instead of holding them back, a **DAM application should be able to handle huge volumes, enormous media files, and maintain responsive, fast performance for thousands of simultaneous users worldwide.**

**It's not a vision
of the future.
It's already here.
And we can
prove it.**

Customers have told us that their **legacy DAM applications have problems keeping up**: lagging searches (some taking 10 seconds or more), inconsistent naming conventions making content impossible to find, version control spinning out of control.

Why? It all comes down to scalability.

Twenty years ago, a state-of-the-art digital camera could shoot 1-megapixel photos. With 8 MB of memory, it could hold fewer than 100 images at 1024x768 resolution.¹ Today's enterprise needs to store ultra-high-definition 4K (and even 8K) video, uncompressed high-resolution image and video files, and 3-d renderings—formats that are orders of magnitude larger than anything that could be imagined in 1998.

It's no wonder that legacy digital asset management (DAM) products from 20th century companies can't scale for the needs of 21st century businesses. Too many businesses feel stuck with legacy software that yields mediocre performance even on a day with minimal load—and that can cause serious slowdowns when many assets need to be uploaded and indexed by the system simultaneously.

Nuxeo is built differently. Using the latest database technology, the Nuxeo platform gives enterprises access to **truly limitless scalability**. Searches at lightning-fast speeds, across billions of assets, high user loads, and large file sizes?

Putting Scalability to the Test

When representatives from one of the world's largest and most technologically advanced companies approached Nuxeo, they were skeptical.

We made big promises: we told them that we could handle massive quantities of large assets, concurrent users, downloads, and searches. As many as they'd ever use, and then some.


To show them we meant business, we did three benchmarking tests to gauge the performance of the Nuxeo Platform at scale. Our tests were designed to evaluate Nuxeo Platform performance not only during "a day in the life" of a worldwide corporation, but also at the time when users need it most: spikes in upload, download, search, and indexing.

In our first test—**The Billion Asset Behemoth**—we were asked to show how the system would perform under a scenario typical of daily use at a massive scale.

The second test, **Performance Under Pressure**, measured platform performance during a simulated massive spike in uploads and downloads.

Finally, we took a closer look into **Searching At Speed**, testing indexing and search responsiveness with an extreme volume of assets and searches.

Nuxeo... To Scale



1. <http://www.cnn.com/TECH/computing/9807/24/megapixel.idg/index.html>

Test #1: The Billion-Asset Behemoth

Most DAM systems struggle under the weight of a high number of assets, especially when those assets are large files with complicated metadata.

Too many teams have resigned themselves to the status quo: DAMs that take ages to perform searches (and then don't return relevant results), have slow file transfer times, and are laggy or unresponsive to user input.

When asset management underperforms, users stop working with the DAM and develop processes to avoid interacting with it as much as possible. As a result, companies are spending money on systems that were doomed from the start — all because of the inability of legacy DAM architecture to scale at the level needed by today's biggest enterprises.

24-48 — Number of 640x480 images that could be stored on a digital camera in 1998

74 — Minutes of audio storage on a CD



Setup

We wanted to evaluate Nuxeo's capabilities, applied to "a day in the life" at one of the biggest companies in the world. So we set up a system that held a truly staggering number of rich media assets: one billion.

Under our test conditions for Test #1, 100-600 concurrent users utilized the Nuxeo Platform in the following three ways:

- 40 percent of users each uploaded 5-50 assets through the DAM user interface, with individual asset sizes of up to 22GB.
- 30 percent of users did a mix of simple and advanced searches.
- 30 percent of users each downloaded 5-50 assets of up to 22GB each.


Even if you throw a ton of expensive hardware their way, legacy DAM applications struggle when assets get into the millions—just 20 million assets can slow each search to 10 seconds or more. How did Nuxeo perform?



Results

Users received search results nearly instantly throughout the test: **over 95% of searches were completed in less than 1/50th of a second**—we'd say that's "in the blink of an eye," but it's actually ten times faster than the human eye blinks.

With 100 concurrent users (and no file transfer accelerators) each transferring up to 41GB of total data, downloads took under 3 minutes and uploads under 10. Scaling up to 600 concurrent users had only a mild impact: transfers of up to 41GB of data (over 10 hours of 1080p video, or 10,000 shots from an 8-megapixel camera) took 10 minutes to download and less than an hour to upload.



Test #2: Performance Under Pressure



File transfer volumes are not constant in most enterprises. Even if most days only involve a moderate load on DAM systems, many companies experience sudden volume spikes several times a year.

For these companies, volume spikes often correspond to the most exciting, critical days for the business. Imagine the needs of a sports broadcaster during a major event—like the World Cup. Or a video game publisher on the day of a much-hyped product announcement. Or a leading news network during the first few hours of breaking a big story.

Customer experience and overall brand satisfaction are on the line during these volume spikes, more than at any other time. No company wants their DAM to start slowing down right when they need it most—so we stress-tested the Nuxeo Platform with extreme loads.

7,390

The highest number of consumer magazine titles ever simultaneously circulated in the United States

10,000



Setup

In order to test two dimensions of scalability and show the flexibility of the Nuxeo Platform, we analyzed file transfer performance under a heavy user load—600 simultaneous users uploading or downloading files in the DAM interface at the same time. Both upload and download performance were tested, with transfers to and from a system with 1 billion assets.

We also examined the transfer speed for extremely large files, to test the ability of the system to handle a 4K video file with a total file size of 1 TB.



Results

Whether they're uploading or downloading files, Nuxeo Platform doesn't slow users down—it allows transfers to move as fast as computing hardware permits.

Upload

With 600 simultaneous users, uploading 1-50 files ranging from 1 MB-500 GB, assets uploaded as fast as the storage hardware could allow: 9 hours for 4 TB of total data uploaded.

A massive, 1 TB video file required 3.5 hours to upload, an average of 90.1 MB/s. The video was transcoded to a 720p, 1.8 GB file in just over two hours. Similarly to the high-volume test, the file could be uploaded as fast as the storage hardware allowed.

Faster storage speeds—or using Nuxeo's native capacity to use multiple storage locations—could further accelerate file transfers.

Download

With 600 simultaneous users downloading a mix of files with a total size of 4 TB, Nuxeo Platform allowed users to sustain a total 450 MB/second download speed, an upper bound dependent on the storage hardware.

The **massive 1 TB video file was downloaded in just 73 minutes, at an average of 250 MB/second**, again limited only by the storage hardware.

100,000

400,000

Total number of scrolls contained by the Library of Alexandria before its destruction

500,000

Total number of feature-length movies produced in human history

Test #3: Searching at Speed

When very large assets are ingested into a DAM, they need to be indexed and searchable as quickly as possible. Searches that take too long have real consequences: team members engage in more “hunt and peck” menu navigation to find files—a process that can result in older versions of assets being re-used and outdated information finding its way to the customer. Or, worse yet, they stop using the system and revert to just asking team members for a recent version, or using whatever version is on their desktop.

To keep up with the speed of business, a modern DAM needs to be able to search through vast numbers of files with the kind of speed that users have come to expect from browsing on Google or Amazon. Moreover, these searches need to be just as responsive on days with high search load as they are on a normal day.

1,000,000

8,000,000

Total number of artifacts contained by the collections of the British Museum

10,000,000

Approximate number of music releases (singles and albums on CD, cassette, etc.) in English since recorded music began



Setup

Two tests were performed. In both tests, a high number of assets were ingested by the DAM and indexed. Then, a high volume of searches were performed, representing a search volume spike at scale.


We first ran this scenario with 500 million assets being ingested, then doubled the number of assets for a subsequent test, ingesting a full 1 billion assets.



Results

For both scenarios, ingestion was completed at an average rate of 40,000 assets per second. All assets were ingested in 3.5 hours for 500 million assets, and 7 hours for 1 billion assets. Indexing was complete for these massive file ingestions after 6 and 12 hours, respectively.

Nearly all searches (a mix of simple and complex) took less than 1/50th of a second in the 500 million asset system, and less than 1/30th of a second for the 1 billion asset system. On average, Nuxeo Platform performed 2,000 searches xeo's native capacity to use multiple storage locations—could further accelerate file transfers.



Beyond the Numbers: What Scalability Means for Your Business

Not every enterprise needs to index hundreds of millions of files in its DAM—at least, not yet. One thing is certain, though: the universe of data is expanding at a truly astonishing rate. **Over 90% of the data that exists today was created in just the last 2 years.**

For many companies, simply storing more data hasn't translated into an avalanche of new insights or more efficient use of resources. Few have the capabilities to quickly and easily sift relevant assets out of a widening sea of content.

What does scalability do for your company? You can stop worrying that adding additional sources of content—like assets from new agency relationships, acquired companies, or partners—will limit your capabilities. **Forget the legacy DAM nightmare of file transfers slowing to a crawl whenever activity peaks.**

With the billion asset benchmark reached (and even more on the horizon), Nuxeo is prepared for the file formats and technologies of the future. Take the long view: with Nuxeo, you can adopt new strategies and technologies without being concerned about impact on search, transfer, and indexing capabilities.

By creating the first DAM that **scales to nearly limitless proportions**, we've created what was once unimaginable: a DAM that won't get in your way, no matter how large you grow.

500,000,000

Total number of Harry Potter books sold

1,000,000,000

Assets in Nuxeo's benchmarking test... enough to handle the content needs of today and tomorrow.