

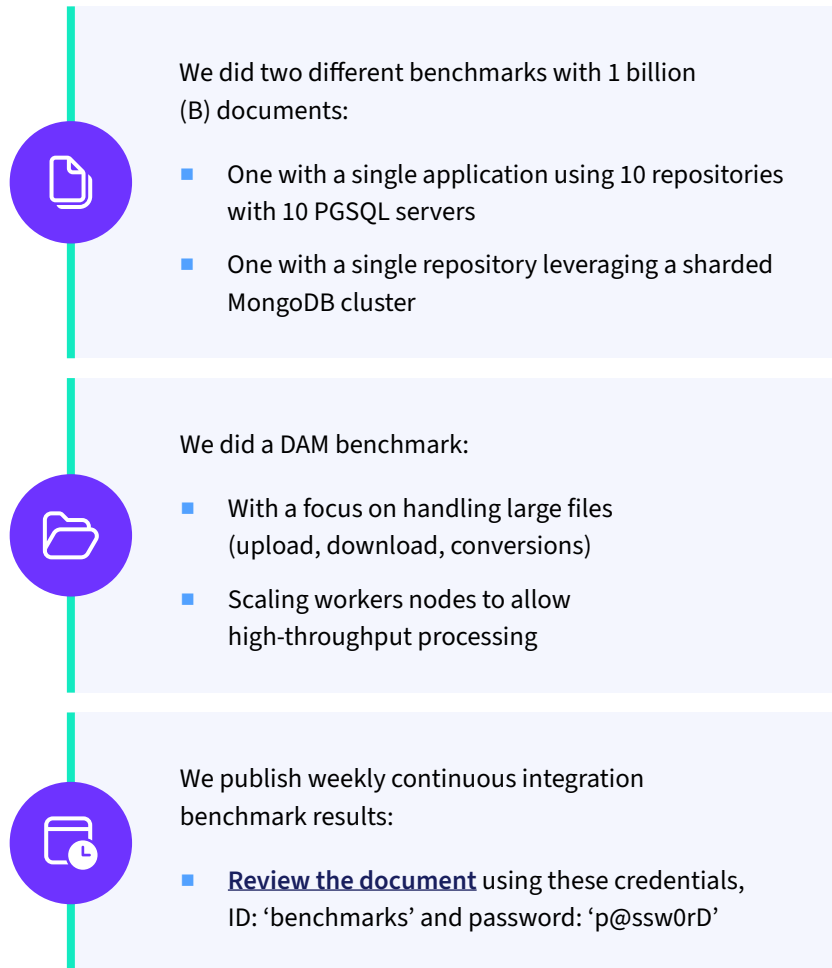



# The 11 billion documents benchmark story




# Another chapter in our benchmark story

In the past, we published several large-scale benchmarks:




 We did two different benchmarks with 1 billion (B) documents:

- One with a single application using 10 repositories with 10 PGSQL servers
- One with a single repository leveraging a sharded MongoDB cluster

 We did a DAM benchmark:

- With a focus on handling large files (upload, download, conversions)
- Scaling workers nodes to allow high-throughput processing

 We publish weekly continuous integration benchmark results:

- [Review the document](#) using these credentials, ID: 'benchmarks' and password: 'p@ssw0rD'

But this time, we wanted to do things a little differently. Here's what we decided:

- Go beyond the 1B documents mark and reach over 10B
- Make the benchmark as realistic as possible

We wanted the benchmark to be as realistic as possible to provide our customers and our ops team real answers on what a multibillion-document repository architecture looked like and the different scalability steps between 1B and 10B.

To achieve this goal, there were some additional constraints:

- Use a real production infrastructure: We can't merely rely on a benchmark "Test Lab" where we would have more tweaking capabilities.
- Test actual use cases with a meaningful dataset: We want to test a real application so we can get actionable insights.



# Production-like, not a lab

## Nuxeo Cloud Infrastructure

The goal is to leverage our Nuxeo Cloud automation to deploy the Nuxeo Application precisely the way we do for our customers.

In our case, it means:

- Deploy Nuxeo using Docker containers on EC2
- Leverage PaaS services
  - AWS Elasticsearch
  - AWS MSK for Kafka
  - AWS S3 for blob storage
  - MongoDB Atlas for the database
- Deploy everything using the existing automation
  - Terraform for AWS automation
  - Rely on Nuxeo Configuration template system
  - Leverage build pipeline of Nuxeo Cloud to bake the custom Docker images

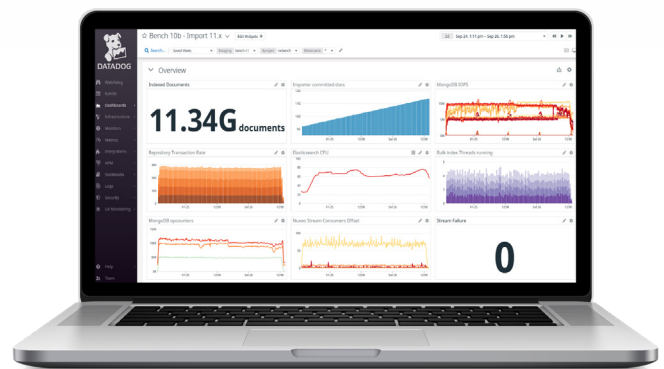


*Dashboard for Phase 1*

## Production infrastructure

Not only is the AWS infrastructure provisioned automatically using Terraform, but we also comply with all the security and production rules:

- Everything secured and encrypted
  - All communications are TLS encrypted
  - Encryption at REST for all storage
- All security systems are on
  - Including Antivirus and IDS
- Rely on DataDog for monitoring
  - Improve existing Dashboard and add metrics as needed



*Dashboard for Phase 2*

The goal is to define architecture blueprints of what infrastructure we need to deploy a Nuxeo application inside Nuxeo Cloud for 1B, 2B, 3B or 10B documents.



## Testing an application — not a storage service

To extract meaningful insights, we want to test an application that reliably translates the challenges our customers will face in the future. We want to test a content-centric application with real use cases, not just benchmarking a database.

Our starting point was to interview some customers to better understand their requirements and challenges regarding future scalability. The main takeaway is that while the existing benchmarks cover the standard document management and digital asset management (DAM) use cases, some use cases in the financial services industry were not covered.

A typical use case would be building a large, consolidated document repository, collecting all the data on customers for a bank or an insurance company.

In addition to the ability to scale to multiple billions of documents, the requirements were:

- A flexible and cost-efficient solution: Cost cannot be linear with volume
- A content application with all the expected features:
  - Attached files associated with the metadata
  - Indexing, including Fulltext
  - Security
- The resulting application to be usable in real life:
  - Can run bulk import without killing the application (support for mixed workloads)
  - Can run a full reindex
  - Can apply bulk operations on a large number of documents

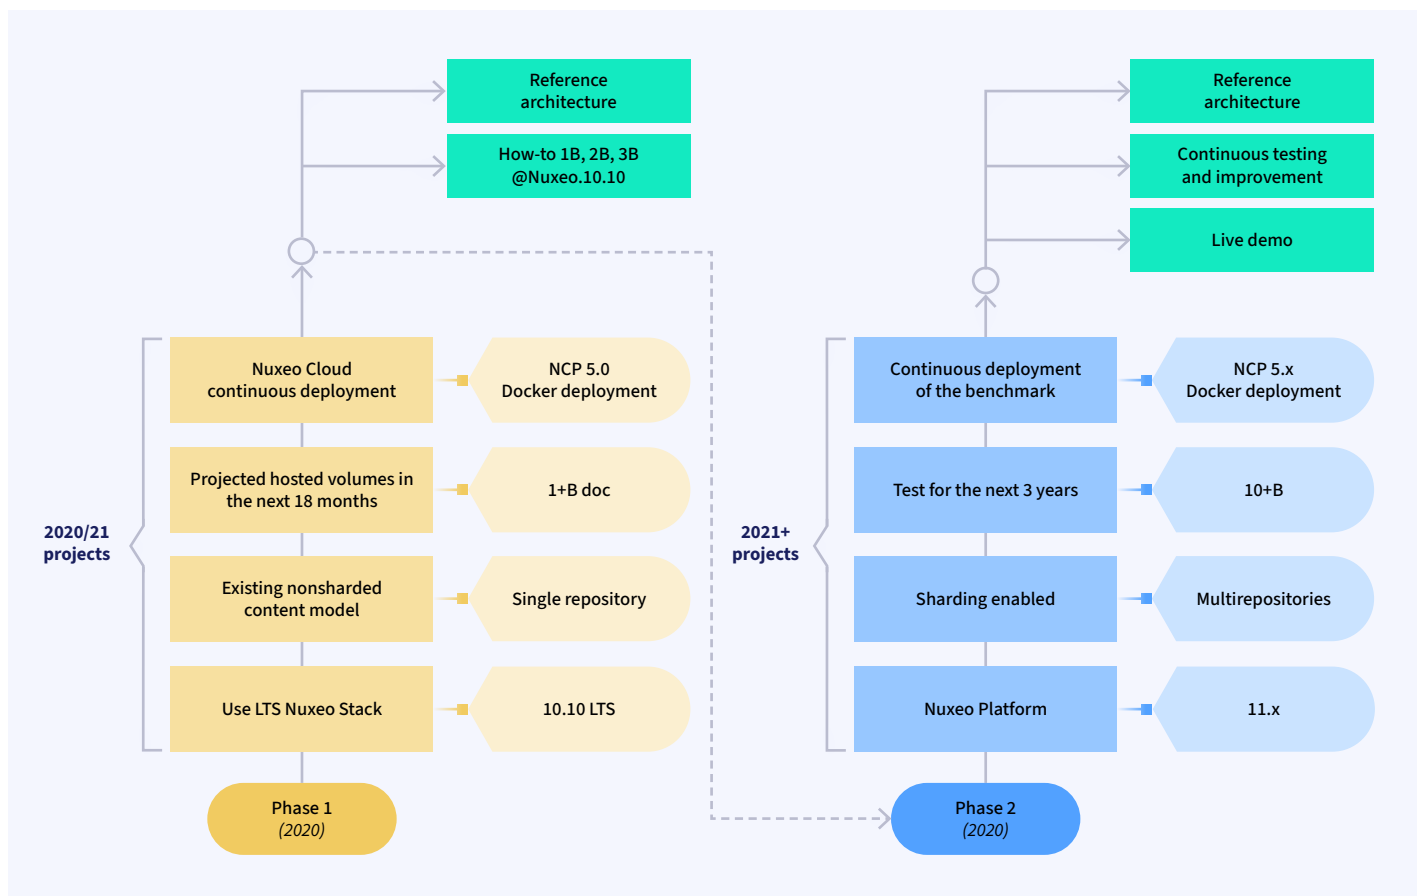


# Phases and methodology

This chapter presents how we structured our benchmarking effort.

## Two different phases

We know from our past experiences that doing a benchmark can be time-consuming, and we decided to split the effort into two phases with different goals and deliverables.



### Phase 1

This first phase aims to provide guidance to our customers who will reach 1B, 2B or more documents in the next 18 months.

Since most of these customers have already deployed their Nuxeo application, we need to apply some constraints:

- Use an already-released LTS version: Nuxeo 10.10 LTS
- Do not alter the existing file plan (sharding is not an option)

Deliverables for Phase 1:

- Tests against a 3B documents repository
- Key scalability steps
  - Know what the bottlenecks are
  - Understand what we need to monitor
  - Know when we need to scale out/up
- Reference architectures for 1B, 2B and 3B documents
  - Hardware
  - Expected performances



## Phase 2

We knew that we would host multibillion repositories within the next 18 months. We want to be ahead of these deployments by clearing the way for these future projects and having some data to be able to guide them.

For phase 2:

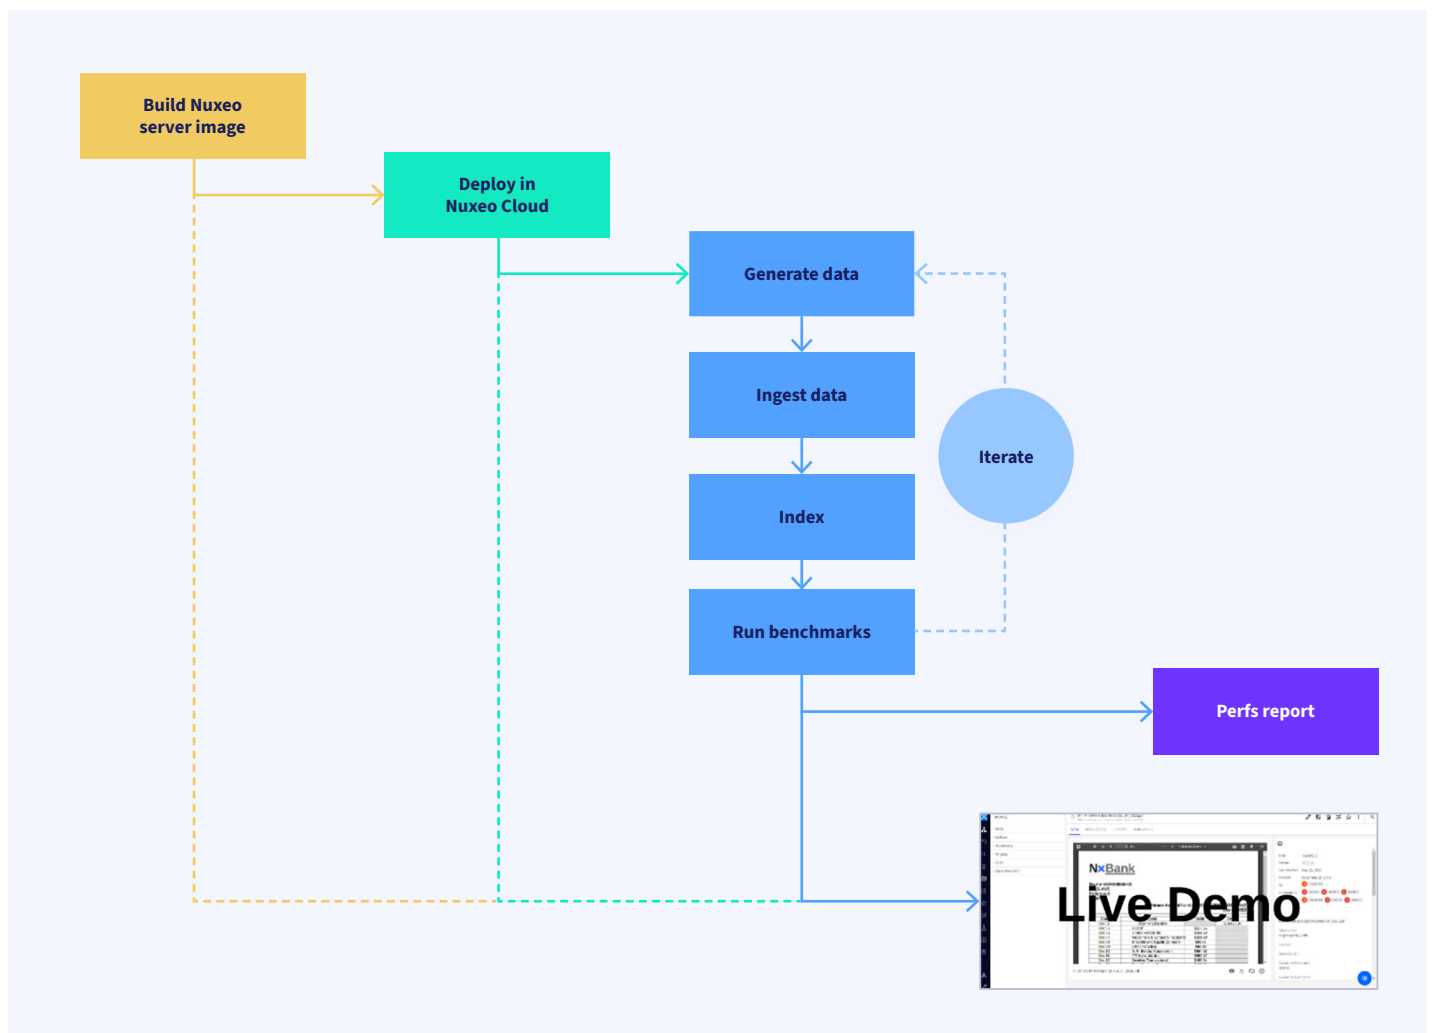
- Push the benchmark to 10B+ documents
- Leverage sharding capabilities to optimize infrastructure costs
  - Sharding at the application level (multirepository)
  - Sharding at the storage level (i.e., MongoDB and Elasticsearch sharding)
- Leverage the “Cloud version” of Nuxeo
  - Nuxeo 11.x

Deliverables for Phase 2:

- Tests against a 10B+ documents application
- Sharding architecture
- Live demo and continuous testing

This last point means we are not aiming for a “one-shot” benchmark. We want to run this benchmark continuously to verify that performance evolves with the platform in the right way.

We also want to have the capability to rebuild this 10B+ documents demo on-demand, so that people who need to see it to believe it can do it.



# Phase 2 principles: Building a reproducible 11B documents repository benchmark

This chapter explains how we designed Phase 2 of the benchmark and the underlying reasoning.

## 11B is not just about doing x3.5 on the 3B benchmark

The first difference with Phase 1 is that we want to reach 10B+, but we do not want to keep increasing the 3B architecture until it gets to the target volume.

While we're confident we could scale Nuxeo, MongoDB Atlas and Elasticsearch to fit 10B+ documents inside a single nonsharded repository, it doesn't seem to be the best approach:

- Not optimal on costs
- Not optimal on operations

This is why for Phase 2 of the benchmark we want to partition the data between:

- Two live repositories that:
  - Will have adequate resources to ensure data stays performant
  - Can easily be updated
  - Are located somewhere we can perform daily operations
- One archive repository:
  - Where the data is mainly read-only
  - Where we can optimize storage for cost

The need to partition the data between “live/fresh” and “archive” was clearly expressed by most of our customers looking at a multibillion document application. In most cases, you do not have multiple billions of documents you need to work with regularly. However, you can have billions of documents that you rarely access, but still need to be accessible and searchable in case you need them.

The partitioning of the data impact:

- What kind of MongoDB cluster we use
  - Typically, we want to downscale the “Archive” MongoDB cluster to support mainly read-only
- What type of indexing we apply
  - We probably don't need to index all attributes and the full-text of the archived statements
- What kind of binary storage we use
  - We could use AWS S3 infrequent access for the blobs associated with the archive repository



## Testing the application — not the storage service

The goal for this exercise was not simply an academic one or to test the cloud and storage infrastructure. The goal was to provide meaningful insights and best practices for customers who are building large-scale information management solutions.

To that end, our objective was to test the conditions that our customers would face when building content-centric business applications. We started by interviewing several of our customers to better understand their requirements and the challenges they expected as they scale their implementations.

To that end, we constructed scenarios that covered document management use cases that are common for our enterprise customers, especially in the financial services industry.

To support these use cases, we included customer IDs (images), statements, account documents and even customer correspondence.

In addition to scaling storage to support multiple billions of documents, additional benchmark requirements included:

- **Cost-efficiency:** The solution should be flexible to ensure that costs do not scale proportionately with document volume.
- **Full support for content functionality:** To ensure the results were meaningful, the benchmark included full support for expected content services features, including metadata, search indexing, security controls and more.
- **Support real-world requirements:** To ensure the benchmark provided meaningful information for large-scale implementations, we tested for several real-world situations, including mixed workloads (e.g., bulk import simultaneous with daily usage, ability to fully reindex, bulk operations across large numbers of documents).



## Results summary

### 11B+ documents

#### Multiple repositories | Single index

##### CRUD / REST API

Response time < 220 ms  
Throughput > 1,900 Req/s

##### Bulk import (3+B)

**Throughput > 25,000 docs/s**  
(Phase 1 > 12,000 doc/s)

##### Navigation

Response time < 500 ms  
Throughput > 800 Req/s

##### Bulk indexing (3+B)

**Throughput > 57,000 doc/s**  
**3B 14h**  
(Phase 1 > 24h)

##### Search with FullText and Facets

Response time < 600 ms  
Throughput < 400 Req/s

##### Bulk import + Index (5B)

**Throughput > 15,000 doc/s**



# Key takeaways

This groundbreaking benchmark test for Hyland Nuxeo focused on scaling enterprise storage and system performance to handle **11 billion documents**.

The benchmarks were completed in **two phases**:

- 1 Phase 1** focused on scaling storage and testing for up to **3 billion documents**, delivering insights on performance optimization and infrastructure scalability.
- 2 Phase 2** pushed capabilities further, achieving **11+ billion documents** across multiple repositories, incorporating sharding for cost-efficiency and performance.

Realistic testing was conducted on **production-grade cloud infrastructure**, including AWS services like S3, Elasticsearch and MongoDB Atlas.

Key deliverables included:

- **Architectural blueprints** for enterprise scalability at 1B, 2B, and 10B+ document levels
- **Live performance testing** that validated usability, speed and reliability for real-world applications
- Enhanced capabilities for tasks like **bulk imports, reindexing** and **CRUD operations**, achieving up to **25,000 document imports per second**

The study highlighted **cost-efficient storage strategies**, partitioning active and archive repositories, and leveraging cloud elasticity to scale infrastructure dynamically.

This performance testing demonstrated how Nuxeo can manage massive repositories while maintaining operational efficiency, empowering organizations that need to scale.

---

➔ Learn more about [Hyland Nuxeo](#) and how it can transform your enterprise content management capabilities.



➤ Try [Hyland Nuxeo](#) for free and experience the power of enterprise-scale content management firsthand.



