

How GEN3i Tripled Performance Speeds on Their Custom Application



Relativity developer partner GEN3i needed to migrate 500 million documents from ViewPoint to Relativity—but they didn't want to do the traditional “dump and load,” where a ton of unnecessary duplicates are also transferred over. Plus, the more automated they could make the process, the better.

So using the Relativity Import API, they spun up a custom solution that connects Relativity directly to ViewPoint's SQL databases to automatically migrate all document metadata, natives, and extracted text.

But there was a problem.

“As soon as we started migrating data, we noticed a slower than expected transfer rate of approximately two million documents per day,” says Driss Mechiche, senior solutions architect at GEN3i. “At this rate, it would take more than six months to migrate all the documents, so we started looking at ways to increase our throughput.”

They called in Relativity's developer experience (DevEx) team for backup—but DevEx didn't find anything obviously wrong. GEN3i had followed all the [best practices](#) when building the application and there wasn't any issues with infrastructure that would cause performance problems.

Given the circumstances, DevEx offered up Nate Noonan, a senior architect, to help.

Nate shared some internal benchmarks for throughput with GEN3i and, compared against what they were seeing, it was clear there was room for improvement. The team worked together and identified two problems: 1) the batch size was too large; and 2) the indexing interval was too frequent.

“Nate recognized the challenge we were dealing with and following his suggestions, we reduced our batch size to 1,000 documents per batch and increased our indexing interval to every two million documents,” says Driss.

Driss and team also modified their workflow to import the documents in two separate steps:

- Step 1: Import document metadata and natives via the Import API
- Step 2: Import extracted text separately via SQL

“Importing documents and natives without extracted text was extremely fast and allowed case managers to begin QC on the migrated documents much sooner, which led to a more efficient workflow overall,” says Driss.

With those adjustments, Driss has seen a huge improvement: *Performance is now three- to five-times faster than before, resulting in a transfer rate of six to ten million documents per day.*

“This is a huge win for us and our client,” says Driss. “It was great having a team to work with that is specifically devoted to making external developers successful. It shows a true commitment to the developer community that DevEx is willing to give us access to a top architect like Nate.”

Ask Nate: Relativity Import API vs. Relativity Integration Points

If you’re familiar with building integrations on top of Relativity, you may be wondering if and when you should use the Import API instead of [Integration Points](#). The short answer? Always start with Integration Points.

“Since introducing Integration Points, we’ve found it’s the best route for most people, as it’s optimized to run multi-threaded imports and leverage optimal Import API settings for most use cases,” says Nate. “If your data volumes or speed are very high—as was the case for GEN3i—or if the data requirements are highly customized, then we would recommend building an integration from scratch with the Import API.”

Need help with one of your applications? [Talk to the developer experience team](#). And don’t forget to check out our resources, including our open source [Relativity Test Helpers](#) project for help with testing strategies.