



## TCW - A Migration Failure Use Case and a WIN Meet Lucene - My New Girlfriend!

TCW is an asset-based carrier performing Intermodal, Truckload, and LTL services with 12 locations in the Southeast. In the mid-1980s, TCW began writing its internal transportation management application in BBx. Over the years, the software grew in functionality. However, in early 2008, management decided to replace the “green screens” and adopt an LTL (less than truckload) software solution that had recently been developed in the transportation space. The hopes were to implement the new LTL version and retrofit the Truckload and Intermodal Drayage business later. After spending a lot of time and money, it was determined that the software was not capable of meeting the needs of all three trucking disciplines.

In 2014, TCW began to look at another out-of-the-box software solution that could offer a single pane of glass for the three differing modes of transportation, but this time it was with a market leader in the transportation software space. More time and more money were spent again trying to build a “niche” solution from a canned application. Again, it was determined that there was too much of a gap between what the application could do versus what TCW’s needs were.

In 2019, TCW looked back to an old faithful friend, BASIS. With the release of BASIS’ BBJ and Barista and their “future-proofing” model, TCW made the decision to migrate all of the legacy codebase and spend our time utilizing a new development framework while simultaneously allowing for the years of accumulated business logic to progress. Because the legacy application had been developed over many years, several key data files were non-normalized and extremely “wide.” The largest (and most utilized) having 340 fields and 625,000 rows at 2.5 GB.

### ***Let’s hear directly from Steve Thompson, Vice President of Information Technology at TCW, on the process and the results...***

As part of the new adoption strategy, we also had to displace a legacy third-party imaging application and make use of a new BBJ and Barista-based imaging solution for billing. All of our scanned pdf files are stored as BLOB fields in records within our largest VKEYED file in the imaging system, which is 621 GB on disk! BASIS’ Lucene implementation means that we’re able to search through any data within those pdfs in 3 secs or less! We scan work orders, Bill of Ladings, Gate Interchanges, Proof of Deliveries, Scale Tickets, Invoices, Accessorial Authorizations, etc. In the old days, prior to Barista/Lucene, a customer might question a charge, and we would have to determine who they emailed, search the email (hoping the employee was still employed and their email was still active and accessible) to find the email conversation that approved the charges. Nowadays, everything is scanned and is easily searchable within seconds by any employee with rights to do so within the Barista security model. The ability for our staff to find key information quickly has been one of the most applauded enhancements.

For example, the non-normalized tables mentioned earlier were impossible to efficiently find data. Work orders that we receive would have to be looked up and found in email, which might take as long as 10 minutes to scour. Now, we can search any word contained in every work order document stored in our archive in less than 2 seconds! Our end users now also have the ability to search across extraneous Note fields (in the same 2 seconds) for mention of a verbal rate confirmation, etc., where before we had no ability to locate the data easily. Lucene has given us the power and performance to MINE our data more thoroughly and give us search privileges against data that wouldn't normally be practically indexable using traditional keys.

Another example of efficiency gains is sifting through the volumes of EDI files (about 135,000) that we receive each month to locate a reference number or booking number. It used to take about 10 minutes of IT involvement to assist in locating an EDI file. Now, our end users can search in less than 2 seconds to find the same data—making it easier for our end users and freeing up valuable IT resources.

Once we realized the awesome power of BASIS' Lucene implementation in our BBJ imaging applications, we quickly made use of it on several of our other BASIS tables, including those very large non-normalized tables mentioned earlier. By quickly identifying all of the important fields in those tables that we wanted to search against, we built full-text search capabilities to allow end users to quickly search through any of those fields, including Notes fields, without adding individual indexes that would grow the base tables. We also added most of our master tables like customers, shipper/consignee, drivers, dispatch, rating, mileage, etc. We've achieved similar acclaim from our users and management team for the improved system working efficiencies and ease with which the users can now access the information in their day-to-day tasks. A measurable improvement in operational efficiencies has been achieved. All with very low risk, no changes to the data tables themselves, and delivering Google-like sub-second search results.

If you haven't given a BASIS DBMS Lucene full-text index a test drive, you're missing out! Once you test one, you'll see they multiply like rabbits!

**Steve Thompson**, VP-IT  
"Lucene's Boyfriend"

