

CUSTOMER WHITE PAPER

Three companies. One choice.

How a Fortune 500 Retailer, Its Global Outsourcer, and Its Systems Integrator All Chose PropelZ to Move Mainframe Data to Google Cloud

PropelZ™ by VirtualZ Computing



In brief

A Fortune 500 retailer needed to move its largest mainframe datasets, some over 200 million records, from IBM Z to Google Cloud to support a strategic business priority with a fixed deadline. It already owned the legacy mainframe data movement tools, including enterprise managed file transfer. So did its global infrastructure outsourcer and its global systems integrator. Those tools could not move the data fast enough.

The retailer tested PropelZ, a no-code mainframe data movement product from VirtualZ Computing, in a structured proof of concept against its existing managed file transfer approach. PropelZ moved the data at roughly twice the observed speed, with every transfer validated for accuracy and completeness. In the final 24 hours before cutover, the integrator measured file transfers that would have taken 12 to 14 hours finishing in 2 to 4. The retailer, the outsourcer, and the integrator each tested it and each approved it. From first planning call to a validated decision took about six months.

2x observed speed vs. the retailer's managed file transfer approach	2–4 hrs final-day transfers before cutover, down from 12–14 hours
~95% of data moved ahead of cutover, with CDC and other measures	3 yeses retailer, outsourcer, and integrator, in about six months

Why we wrote this down

Most mainframe teams we talk to are in the same spot this retailer was. The tools are already in the building, the contracts are signed, and the cloud program is still waiting on data. We wanted to lay out plainly what this customer tried, how they tested, what they found, and why three large organizations agreed. If you are weighing the same decision, the details should be more useful than a summary slide.

01 / THE CHALLENGE

Why weren't the existing tools enough?

The short answer is speed at volume. The retailer's data movement process had several steps, and at hundreds of millions of records each step cost time the batch window didn't have.

Like most large mainframe shops, the retailer moved data to the cloud in four hops. Data was unloaded from VSAM or sequential datasets into flat files. The files were sent through managed file transfer. They landed in cloud storage, and then they were loaded into their target. Every hop adds elapsed time, and every hop is another place a job can fail and need a rerun.

For small files this works fine. For a dataset with 200 million records it can mean the difference between finishing overnight and running into the next business day. The retailer ran into three specific limits:

- **Elapsed time.** The largest transfers were pushing past the batch window.
- **File size.** Staging steps put a practical ceiling on how large a single transfer could be.
- **Effort per dataset.** Getting each new dataset moving meant custom scripting and hand-tuning.

The outsourcer that runs the mainframe had its own tools. The integrator, which picks up the data once it lands in the cloud, had its own file transfer technology and data mover. For file data at this volume, though, its fallback would have been building something custom. The team wasn't short on software. It was short on a way to hit the program's dates.

01 / THE CHALLENGE — CUSTOMER PERSPECTIVES

What the teams said

The retailer's head of engineering for the program put it this way:

FORTUNE 500 RETAILER

”

“We already owned most of the legacy mainframe data movement tools. None of them could move data to Google Cloud at the scale, and with the simplicity, this program requires. PropelZ did.”

Senior Director of Engineering

The integrator's managing director described what the alternative would have looked like:

GLOBAL SYSTEMS INTEGRATOR

”

“Without PropelZ, we would have had to come up with this type of solution on our own. We would have had to go custom, looking at compression technology to accelerate the transfers and at network options. Being able to buy it meant we didn't have to solve this problem ourselves. PropelZ solved it for us.”

Managing Director

02 / THE EVALUATION

How was PropelZ tested?

The retailer ran a proof of concept on its own mainframe, with its own data, and wrote down what success meant before any testing began. PropelZ had to do four things:

1. Move large volumes of mainframe data to the cloud reliably
2. Perform better than the current approach
3. Support more than one way of moving data
4. Fit cleanly into the retailer's Google Cloud architecture

Operating environment

The outsourcer installed PropelZ on a dedicated test LPAR under the existing RACF security, after the retailer's security review. That was deliberate. The test had to show how the product would behave inside the real operating environment, including the enterprise job scheduler, standard JCL, and existing change control.

Representative datasets

The datasets were chosen to reflect real work: VSAM files copied from production candidates, sequential files large and small, generation data groups, and tape-resident data. Data with personal information was left out for security reasons. Transfers were timed against the managed file transfer process the retailer already used.

Refining the test

The proof of concept also changed as it went, which we think is how a good one should work. Midway through, the retailer's team named three priorities: show that data could land as files, not only in a database; confirm consistent performance; and prove there was no ceiling on file size. Each one was tested and closed before the evaluation ended.

03 / THE RESULTS

What did the testing show?

PropelZ met or exceeded every success criterion. The retailer's close-out summary recorded these results:

What the retailer needed to know	What the testing showed
Can it handle our volumes?	Datasets of 100 million to over 200 million records moved to Google Cloud
Is it faster than what we use now?	Throughput observed at roughly twice the existing managed file transfer approach
Does the data arrive intact?	Every transfer validated for accuracy and completeness, with stable, repeatable runs
Can we land data more than one way?	A database staging path (PostgreSQL) proven production-ready, and a direct file-to-cloud path validated for enterprise workloads
Does it fit our cloud architecture?	Integrated with Google Cloud storage and services, in line with the target design
Can our operations team run it?	No custom code; jobs set up by configuration and run as batch, JCL, or z/OS UNIX jobs under the enterprise scheduler

Three of these results mattered more than the others.

03 / THE RESULTS — SPEED

Speed that changes the cutover window

The first was speed. A 2x gain sounds like a benchmark figure, but on a job that runs for hours it decides whether the job fits in the batch window. The retailer's infrastructure architect described it that way:

FORTUNE 500 RETAILER



“At our data volumes, 2x isn't a benchmark number — it's the difference between a transfer that fits in the batch window and one that doesn't. PropelZ moved our largest mainframe datasets to Google Cloud at twice the speed of every other data movement tool we tested.”

Infrastructure Architect

The integrator saw the same effect where it mattered most, in the last hours before cutover:

GLOBAL SYSTEMS INTEGRATOR



“Moving the files in the final 24 hours before cutover was going to take 12 to 14 hours, which would have spilled over into our cutover downtime window. Our target for that window is 8 to 10 hours. With PropelZ, we measured it at more like 2 to 4 hours. That was significant.”

Managing Director

03 / THE RESULTS — DELIVERY AND CDC

Delivery and readiness

The second was having two ways to land data. The integrator did not want the end state to depend on a database alone. PropelZ could stage data in PostgreSQL, which suits incremental updates, or stream files straight to cloud storage. The file-size ceiling the team had seen earlier came from staging. The streaming path removed it, and it was tested on large files in the retailer's environment.

The outsourcer's technical lead described both paths during testing, along with early results from the optional multiplexer, which runs more transfers in parallel:

GLOBAL INFRASTRUCTURE OUTSOURCER



“We've tested the output that sends data directly to Google Cloud Storage, and it works well. We've tested sending data to a PostgreSQL database, and that works well too. With the multiplexer, it speeds up quite a bit.”

Mainframe Technical Lead

The third was change data capture (CDC). PropelZ copies a baseline of each file once, then moves only the changes after that. It delivers CDC the way it delivers ETL: through a more modern, efficient, no-code approach than traditional tools, with nothing to script or maintain. The retailer's data keeps changing right up to cutover, so this mattered. Most of the data can be in the cloud and current well before the switch, and the cutover window only has to carry the last changes.

GLOBAL SYSTEMS INTEGRATOR



“The CDC component of PropelZ lets us not only replicate the baseline files but represent change, and that moved quite a bit of our work in front of the cutover window. Combined with other things, it's allowed us to move about 95% of the data ahead of time.”

Managing Director

04 / THREE ORGANIZATIONS, ONE DECISION

Why did three companies have to say yes?

In a large mainframe shop, the customer is rarely the only decision-maker. This retailer's mainframe is run by a global infrastructure outsourcer, and its cloud conversion is led by a global systems integrator. Each of the three could stop a new tool, and each had its own reasons to be careful.

Organization	What it answers for	What it needed to see
The retailer	The program's results, budget, and deadlines	Real speed gains on real data volumes
The outsourcer	Keeping the mainframe stable and secure	A clean install, a passed security review, normal scheduler and JCL use, and no effect on production
The integrator	The cloud architecture and the data conversion	Complete, accurate data that downstream processes could use, without depending on one landing method

All three already had tools they knew and vendors they worked with. Choosing a newer product meant each of them taking on some risk. They each tested PropelZ against their own concerns, and each approved it.

Asked what he would tell a peer who called him about PropelZ, the integrator's managing director said:

GLOBAL SYSTEMS INTEGRATOR



“We had a very specific problem to solve, driven by the volume of data and the amount of time we had for cutover. We saw at least a 50% improvement over the other methods we were using to transfer that data, and that was very real. Overall, the experience has been very good, and it's been successful for what we needed to do.”

Managing Director

05 / INSTALLATION AND OPERATIONS

How fast is PropelZ to install and prove?

Installing PropelZ takes hours, not weeks. When a mainframe solutions provider installed PropelZ in its own environment, it went from installation to operation in 24 hours. The consultant who did the install described it this way:

MAINFRAME SOLUTIONS PROVIDER

”

“PropelZ is a very straightforward install. You create the zFS, upload it, untar it, and it's ready to go. It doesn't take much time at all, just a couple of hours to get the product laid down and installed in a testing environment.”

Senior Solutions and Innovation Consultant

The steps are ones any z/OS systems programmer already knows:

- **Create a zFS file system of about 200 MB**, upload the package, and untar it
- **Set up RACF permissions** and reserve TCP/IP ports
- **Automate startup** with standard UNIX scripts and JCL, and connect it to the automation tools already in place
- **In a sysplex**, share one read-only zFS across LPARs and keep log files local to each LPAR

At the retailer, the outsourcer installed PropelZ on a dedicated LPAR under the existing RACF security, through its standard change process, with no custom code. From the first planning call to a validated decision took about six months. That included a security review, change control, and sign-off from three large organizations. Decisions that involve a customer, an outsourcer, and an integrator usually take much longer in mainframe software.

06 / PRODUCT ARCHITECTURE

How does PropelZ work?

PropelZ isn't limited to projects like this one. It moves data between the mainframe and the cloud in either direction, for any reason a business needs to share it. That includes feeding cloud analytics and AI with mainframe data, keeping hybrid applications on both platforms in sync, and bringing cloud data back to mainframe processes. The same questions this retailer tested, volume, speed, integrity, and operational fit, apply to all of them.

PropelZ is software that runs on z/OS, the operating system of IBM Z mainframes. It moves data directly between mainframe datasets and cloud databases or cloud object storage, in either direction. For this retailer, that meant reading mainframe data where it sits and writing it straight to Google Cloud, without separate unload, managed file transfer, and reload steps.

In practice, that means:

- **It runs like any other mainframe job.**

Teams start it from their job scheduler using standard JCL, as a batch job, or under z/OS UNIX System Services, and they get the usual return codes and logs.

- **It reads mainframe data formats.**

That includes VSAM files, sequential datasets, generation data groups (including relative generations such as 0 and -1), and data on tape. Binary data can move without a COBOL copybook.

- **It writes to databases or files.**

It can load more than 100 industry-standard JDBC targets, including PostgreSQL, Oracle, MySQL, Microsoft SQL Server, and Snowflake, or stream files straight to cloud object storage without staging.

06 / PRODUCT ARCHITECTURE — CONTINUED

Built for mainframe operations

- **It captures changes, not just full copies.**

Change data capture (CDC) moves a baseline once and then only what changed, so copies stay current for a cutover or for ongoing sharing. It's a more modern, efficient approach than traditional CDC tools, with no code.

- **It transforms as it moves.**

A built-in transformation engine does the work of ETL without ETL scripting.

- **It's set up, not programmed.**

Jobs are defined with parameters, and operators control how much runs in parallel. An optional multiplexer adds parallelism to speed up large transfers. There are no scripts to write or maintain.

PropelZ needs z/OS 2.4 or later, Java 11 (preferred), and a network connection to the target. It was designed to read production data without disrupting production work and to run within existing workload management settings. It comes out of VirtualZ Computing's patented technology portfolio.

PropelZ can run on zIIP specialty engines as well as general processors. The outsourcer's technical lead described how it behaved during testing:

GLOBAL INFRASTRUCTURE OUTSOURCER



"If the resources are there, if the zIIPs and the general processor engines are available, PropelZ will take advantage of them. Any waiting we saw wasn't because of the product. It was other processes competing for the same resources in a development environment."

Mainframe Technical Lead

07 / COMMON QUESTIONS

Common questions

Is PropelZ only for moving data off the mainframe?

No. This retailer used it for a critical, deadline-driven program, but PropelZ moves data between the mainframe and the cloud in either direction. Teams use it whenever data has to be shared across the two, whether that's mainframe data feeding cloud analytics and AI, or cloud data that mainframe applications need.

Is PropelZ a replacement for managed file transfer?

For moving data between the mainframe and the cloud, it can be. PropelZ skips the unload-and-send steps that managed file transfer depends on. Many shops will keep managed file transfer for other file exchange.

How much faster is it?

The retailer observed roughly twice the speed of its managed file transfer approach, and the integrator reported at least a 50% improvement. In the final 24 hours before cutover, transfers estimated at 12 to 14 hours measured 2 to 4 hours with PropelZ. Results depend on the data, the network, and system load, so the best way to know is to test on your own datasets.

How large a dataset can it move?

This retailer moved datasets of more than 200 million records. The streaming path sends data directly to cloud storage without staging, which removes the file-size ceiling staging creates.

Does it require writing code?

No. Jobs are set up with parameters and run through the tools mainframe operations teams already use.

Will it affect production workloads?

This was the outsourcer's first concern. The testing confirmed PropelZ could read production-representative datasets without disrupting production processing.

Which clouds and targets does it support?

This customer used Google Cloud. PropelZ writes to more than 100 industry-standard JDBC targets, including cloud databases, data lakes, AI and analytics platforms, and security monitoring (SIEM) tools, and it can stream files to cloud object storage. Beyond z/OS, it runs on Windows, Linux, UNIX, IBM i, and other Java-enabled platforms.

08 / YOUR NEXT STEP

Mainframe data at the speed of innovation

The mainframe holds some of the most valuable data a business owns. What has changed is how many places that data now needs to be: cloud analytics, AI models, new applications, and back again. The question is no longer whether mainframe data should move. It is whether it can move fast enough, safely enough, and simply enough to keep up with the business.

This retailer had one critical business goal and a date it couldn't miss. It already owned the legacy tools, and so did its partners. What changed the outcome was a willingness to test something new, on real data, with everyone who had to say yes in the room from the first day. Three large organizations looked at the same evidence and reached the same conclusion.

If your transfers are pushing past the batch window, if your cloud and AI teams are waiting on mainframe data, or if every new dataset needs its own scripts, that doesn't have to be the cost of running a mainframe. PropelZ goes live in hours, and it can be proven on your data, against your current tools, with your partners at the table. It can feed AI, analytics, security, and the cloud, in either direction, for whatever your business needs next.

One solution. Any data. Every platform.

This retailer proved it works at enterprise scale.

PropelZ is a product of VirtualZ Computing. To see it move your data in your environment, [schedule a briefing](#) or try PropelZ at www.virtualzcomputing.com/propelz.