Connected workflows: The new standard in marine seismic data delivery

Ed Hodges[1]*, Cerys James[1] and John Brittan[1] discuss how Low Earth Orbit (LEO) satellite constellations have revolutionised connectivity for marine seismic operations, offering high bandwidth and low latency at significantly reduced costs, and enabling near-real-time data transfer and remote processing.

Introduction

For many years, towed-streamer seismic data has been delivered from the vessel to the processing centre or to a data storage facility on physical media, typically enterprise tape cartridges such as IBM 3592 JC. This method, long considered the industry standard, often introduced significant delays between acquisition and commencing processing. Some contractors then switched to utilising physical disks (NAS drives) instead of tapes. However, this does not reduce the shipment delays but can mitigate against corrupted or bad tapes. Historically, this led to many seismic contractors placing significant amounts of High-Performance Computing (HPC) onboard their vessels along with the required infrastructure, spares and corresponding support structure and of course the highly qualified, experienced processing geophysicists to perform the processing, typically working shift patterns on an offshore rotation of 5 or 6 weeks. Often, this led to a fast-track product being either partially or fully generated offshore, onboard the vessel, because this was seen as a timely way to create an interpretable volume. Often weeks later, the field data would be delivered to an onshore processing centre where a processing team would probably take more than a year to achieve full-integrity final volumes for interpretation, which offered uplift over the similar products created offshore. Whilst the improvements may have been dramatic and certainly significant, there are numerous cases where the full integrity volumes came too late to influence initial drilling decisions.

Alternative strategies included decimating and compressing subsets of field data for satellite transmission to onshore processing centres, enabling preliminary processing. While effective to some extent, these methods involved compromises in data quality and completeness. Recent technological advancements have transformed this landscape.

The advent of Low Earth Orbit (LEO) satellite constellations, such as those provided by Starlink, has revolutionised connectivity for marine seismic operations. These systems offer high bandwidth and low latency at significantly reduced costs, eliminating the need for both data decimation and lossy compression. This technology makes near-real-time data transfer and remote processing a practical reality. In this paper, we expand on the potential applications proposed by Ewig *et al* (2024) and provide case studies detailing the experience gained actively using this approach.

LEO satellite constellations

Whilst there are numerous LEO constellations in the planning or deployment phase today, there are only really two offering fully commercial services. Starlink from SpaceX has the largest operational constellation by far, with >6000 active satellites currently. OneWeb from Eutelsat is the other fully operational LEO network. However, in 2023 the only real provider was Starlink.

Back in 2023, PGS (now TGS) initiated field trials of the Starlink Maritime Service by installing four antennas on a vessel conducting an ultra-high-resolution (UHR) survey in the Irish Sea and six antennas on another vessel acquiring 4D GeoStreamer surveys in the North Sea. The results were highly encouraging, with system uptime proving exceptional and latency significantly lower than the ~750–800 ms, typical of geostationary satellite solutions. Download speeds were also impressive. However, for marine seismic operations, the primary requirement is high-capacity data upload from vessel to shore.

The trials confirmed that the preferred use case was bulk data transfer rather than remote offshore processing. At the time, upload performance was constrained more by the limited number of antennas than by the service itself. Based on these findings, TGS implemented a solution equipping each T-Class Ramform vessel — the most capable in the fleet — with 12 antennas: 10 dedicated to seismic data transfer and two reserved for general connectivity.

Devising an agile solution

Initial tests indicated that LEO satellite systems could feasibly transfer a complete, full-integrity dataset without the need for lossy compression or decimation. During the late 2010s, the industry transitioned from recording discrete shot records to continuous recording, which enabled the widespread adoption of triple-source simultaneous shooting (Langhammer, 2015). This shift redefined 'record length' from an acquisition parameter to a processing parameter — a practice that remains standard

[1] TGS

* Corresponding author, E-mail: edwin.hodges@tgs.com

DOI: 10.3997/1365-2397.fb2025094

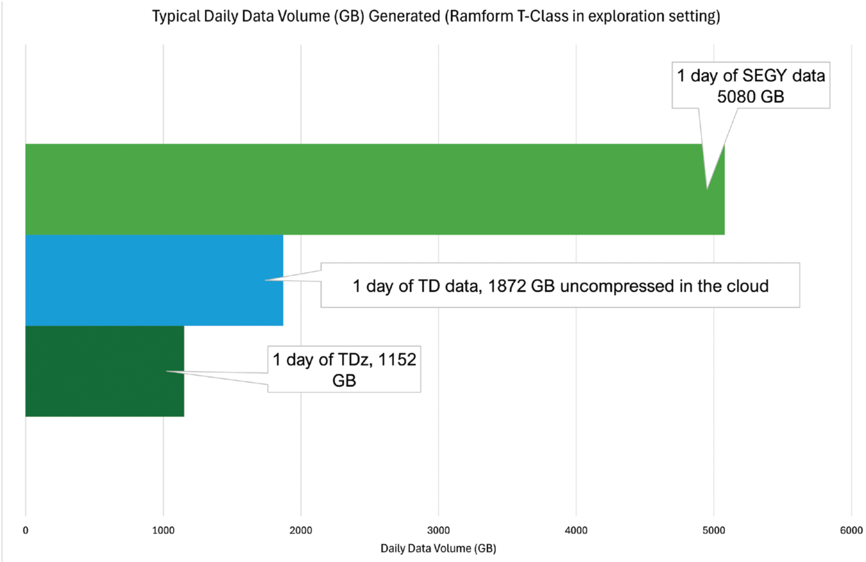

Figure 1 Illustration of typical daily data volumes generated by a large streamer vessel.

today. Consequently, the overlap between shot records can be so extensive that many samples are effectively delivered twice within discrete shot records. The resulting increase in data volume is significant, as illustrated by the progression from raw field measurements to final processed volumes, further compounded by the requirement for multiple archive copies.

The raw data from the streamer's analog-to-digital converters is stored in 24-bit integer format. This was identified early as the optimal starting point for transmitting data to the cloud. To improve efficiency, simple lossless compression algorithms were evaluated to balance computational cost against compression ratio. From the outset, it was decided that any compression applied must be fully reversible to preserve data integrity. This requirement was driven by the need for a standardised solution applicable across the entire fleet and all survey types, without requiring case-by-case testing or data-dependent decisions. The selected solution is an open-source lossless compression algorithm that typically achieves a packing ratio of approximately 1.6:1.

The vessel's recording system generates continuous time-domain (TD) data in 15-minute segments. Each segment is compressed using a multi-threaded, lossless algorithm to achieve efficiency without compromising data integrity. The compressed file is then divided into smaller chunks, which are uploaded to the cloud using proprietary software. This software employs cloud-based commands to enable parallel transfers through the firewall, which manages load balancing to optimise upload performance. The fully automated process has proven highly robust, reducing satellite transmission costs and transfer times while providing a cost-effective solution for long-term cloud-based archiving of field data.

Once the files are transferred to the cloud bucket, the proprietary software reassembles and decompresses each 15-minute segment. A single sail line typically consists of multiple such segments. After all the required segments are available in the cloud, the workflow pauses until ancillary files and necessary approvals are in place. During the initial development phase, many of these decision gates were handled manually as part of a minimum viable product (MVP). Over the course of the proof-of-concept phase and subsequent fleetwide rollout, numerous updates were implemented with the primary goal of automating steps where practical, thereby reducing manual intervention. This approach has resulted in a robust, long-term solution that remains adaptable without requiring frequent changes as survey configurations and parameters evolve.

A survey using a large streamer configuration typically generates just under 2 TB of raw data per day. By applying lossless compression, this volume can be reduced to approximately 1.2 TB for the cloud upload step. Standard practice is to deliver streamer data as discrete shot records in SEGY format, resulting in datasets of around 5 TB representing this full day's acquisition, primarily influenced by the chosen record length. This increase is driven by the switch from 24-bit integer to 32-bit IEEE format and the overlap inherent in triple-source shooting (Figure 1). Creating multiple copies of data on physical media remains common practice, for delivery to partners, government entities and archives, etc. Table 1 illustrates data volumes from a recent proprietary survey, which required five complete sets delivered on a mix of tapes and NAS drives. Transitioning this task from the field to a centralised location — enabled by cloud transfer — offers clear benefits in terms of time savings and cost reduction.

Case studies – 4D surveys in the North Sea

During the summer of 2025, TGS acquired multiple proprietary 4D surveys in the Norwegian sector of the North Sea through to the Barents Sea, along with several 3D surveys. All three vessels operating in the North Sea were equipped for direct satellite transfer of data to the cloud. Several experienced offshore geophysicists transitioned from full-time offshore roles to onshore positions, working in the cloud to deliver products efficiently.

Real Example based on Hyperion Proprietary Survey	GB
Compressed Volume uploaded from vessel	73
Uncompressed Volume in the cloud	116
Set 1 Field Deliverables	206
Sum total of all sets (5) delivered	1030

Table 1 Example of delivered data volumes from a recent proprietary survey.

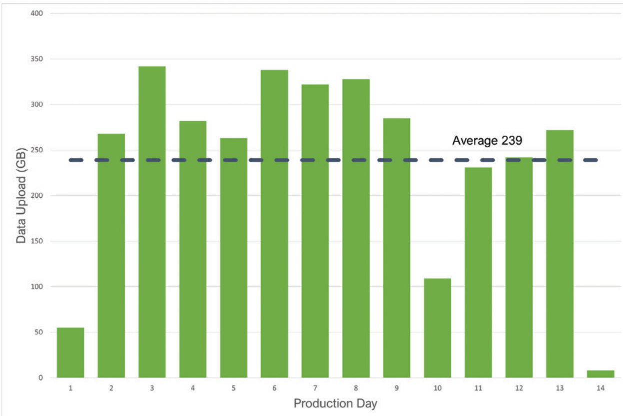

Figure 2 Raw data uploads from a small 4D survey along two weeks of production.

These 4D surveys provided an excellent opportunity to test the new connected workflows, as rapid data delivery offers a particularly high value to 4D projects.

For one of the 4D surveys, the workflow included generating rapid turnaround 4D quality control (QC) products. This involved computing repeatability attributes using various 4D binning strategies in the cloud and comparing them against both a baseline and a previous monitor survey. Additionally, and arguably more valuable, were 4D difference displays derived from 2.5D Kirchhoff pre-stack depth migrations of 4D-binned individual sail lines. The scalability and accessibility of the cloud enabled these products to be generated for every monitor survey line, typically within three days of acquisition, often sooner. This capability, however, required significant preparation of the baseline and monitor datasets by TGS Imaging before the 2025 monitor acquisition. While these extended 4D QC products could technically be generated offshore, the resources required — in terms of both computational capacity and experienced personnel — are far more efficiently and economically utilised by leveraging the cloud for this work.

One of the 4D surveys was acquired well inside the Arctic Circle. Concerns regarding bandwidth availability at such high latitudes had been addressed before the project through testing, during which upload rates exceeding 3 TB per day were achieved near the same field. This relatively small survey utilised a moderate streamer configuration, with just under 60 km of GeoStreamer deployed. The accompanying bar chart in Figure 2 illustrates daily raw data uploads, which averaged approximately 239 GB per day.

The survey acquisition lasted 14 days, with the final shot point recorded late on a Thursday night. By the following Tuesday, all deliverables had been confirmed in the third-party processing company's cloud environment via cloud-to-cloud transfer. This five-day turnaround included navigation processing, product generation, quality control, and the weekend, as well as making the data available in a shared cloud bucket for third-party access and confirmation.

Similar workflows were repeated for two additional 4D surveys acquired during the same season. In cases where the onshore imaging team handled onward processing, data delivery was seamless, involving a simple transfer from an operations bucket to an imaging bucket within the cloud — a process executed multiple times per week.

Connected workflows, combined with standardisation and event-driven services within the cloud ecosystem, enable the creation of dynamic dashboards that update automatically. Using these dashboards to track project progress across all stages provides stakeholders with a single, clear interface — improving knowledge sharing while reducing confusion and unnecessary communication. Figure 3 illustrates a high-level overview of a completed project, showing detailed daily uploads alongside an overall view of all stages through to the generation of SEGY deliverables.

Electronic data delivery

Using dedicated shared cloud buckets has proven to be a reliable and robust method for transferring data between companies. Electronic data delivery using this approach offers significant advantages over traditional tape shipments or NAS drive transfers, beyond obvious time savings. These time savings are particularly valuable for international transfers, where customs delays often compound physical shipping times. From a risk management perspective, the profile of electronic transfer differs markedly from that of physical media. Cloud-based workflows typically incorporate built-in integrity checks, providing a high degree of confidence in data accuracy. In the rare event of an error or incomplete dataset, remediation is generally faster and less complex than with traditional methods. While checksums remain standard practice, errors are exceedingly

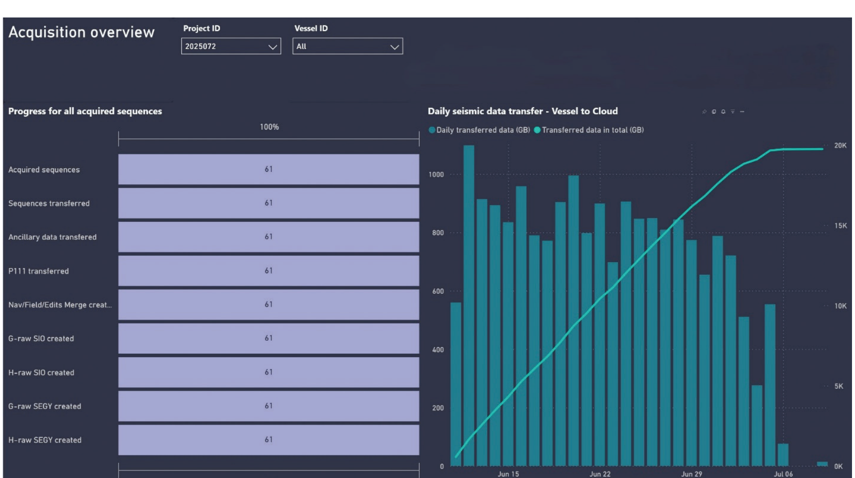

Figure 3 High-level dashboard providing the basic overview of an acquisition project and the progress of the deliverables generation.

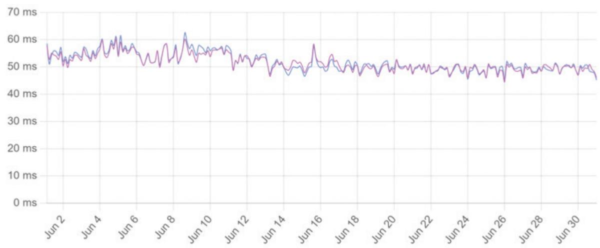

Figure 4 Ping latency for two of the Starlink antennas onboard Ramform Tethys during June 2025.

rare because cloud-native transfer commands usually include checksum validation and failover logic.

Electronic delivery is almost certainly the most cost-effective option and aligns with the industry's demand for timely data access. While physical media for long-term archiving will likely remain for some time, this approach decouples archive creation from immediate delivery, enabling centralised resources for more cost-efficient archive generation.

At first glance, the benefits of digital delivery appear obvious, yet industry adoption has been surprisingly slow. Perhaps the situation is analogous to the persistence of physical books despite the convenience and lower cost of e-readers — physical books still outsell e-books by nearly 4:1.

Reasonable expectations for the LEO satellite constellation

The LEO design is very favourable for low latencies. Over the three years since Starlink was first installed in the fleet, we have observed further improvements in the ping latency. Today, in Q3 2025, we observe values in the 20-30 ms range when the vessel is operating within the most densely populated part of the constellation between ±60° N/S. During the summer months, a vessel at ~71°N, well inside the Arctic Circle, observed values in the 50-60 ms range (Figure 4).

Experience has shown that the Signal Quality remains largely unaffected by all but the most intense electrical storms. The flat high-performance marine antenna appears unaffected by any vessel motion. The current solution for the offshore data transfer makes use of 10 dedicated antennas to upload the data. We routinely see a single Starlink subscription using a bonded pair of antennas having a throughput of >500 GB in 24 hours; similarly, a lone antenna uploading 250-300 GB within the same period. The antennas are located on the upper deck with a clear view of the sky, as shown in Figure 5.

During Q2 of 2025, TGS tested the High Data Throughput service offered by Marlink using Starlink as the satellite service. The objectives here were to test the limit of the system both in terms of overall throughput, antenna performance, and global coverage. From the early days of testing Starlink, it had been made very clear that what is provided is a 'best effort' service without the SLAs (Service Level Agreements) that the wider business world is so accustomed to. This High Data Throughput offering was an attempt to at least provide a priority service. One of the vessels was transiting from Angola to Las Palmas; during this time, the upload capacity was constantly tested with real data being uploaded to the cloud. The tests included investigating the effect of varying the number of active antennas between 6 and 10. Over an 8-day period, the upload performance was robust. The key measurement was data volume uploaded per antenna per hour, which varied between 18 and 23 GB/antenna/hour. The rate per antenna was slightly higher with 6 active antennas rather than 10, but increasing to 10 increased the overall throughput to the stage that ~4.8 TB was uploaded within a 24-hour period.

Another test was performed on the *Ramform Tethys* as she transited from Lerwick to just north-west of Hammerfest. This relies on the limited number of satellites in the 70° inclination or the 97.6° polar shells. The throughput observed was consistently ~15 GB/antenna/hour, as shown in the figure below. Whilst this is significantly less than what was observed at the lower latitudes, it still offers a daily upload exceeding 3.5 TB, which is way more than we envisage a streamer vessel requiring (Figure 6). This figure also highlights that one antenna (vlan217) was underperforming compared to the others; the issue was later resolved by replacing the cable.

The OBN case

Predicting daily data generation from a streamer seismic survey is challenging due to factors such as weather, SIMOPS, and technical downtime. However, there is a clear upper limit on the maximum daily volume, typically around 2.5 TB, which reduces

Figure 5 Location of the Satellite antenna following the minimum separation criteria.

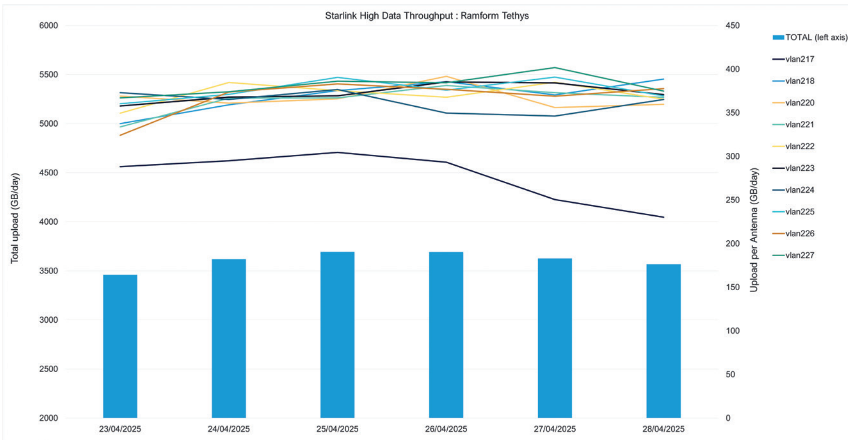

Figure 6 High Data Throughput test in high latitudes. Note that the Total upload per day is using the left axis, whilst the individual antenna uses the right axis.

to approximately 1.6 TB after lossless compression. This is well within the capacity of the onboard hardware and is now routine for large exploration surveys, such as those in the Equatorial Margin of Brazil, where a single line into the current can take more than 24 hours.

For OBN surveys, the variables differ, even though the data type remains similar, essentially continuous trace data. Daily volumes depend on the rate at which ROVs recover nodes and the duration of their recording, which varies with survey design, node density, and crossline coverage. Many surveys reach 3-5 TB per day, and short periods exceeding 5 TB per day are not uncommon. Within the OBN arena, this journey is just starting out; to date, OBN teams have used LEO satellites primarily to accelerate delivery of low-frequency hydrophone-only datasets for FWI model-building. This approach was recently applied in the Gulf of Mexico on the Laconia ultra-long-offset sparse OBN program. In that case, data volumes were relatively modest because preconditioning was performed onboard, and only resampled hydrophone SEGY files were uploaded. The satellite companies are bringing though the next generation of hardware, which is expected to be able to handle these data volumes. We expect to be able to leverage this to provide similar cloud access and rapid data turnaround.

The way ahead

How might the availability of seismic data in the cloud transform workflows beyond simply accelerating delivery? When the current solution was designed, it was accepted that some delay would exist between acquisition onboard the vessel and dataset availability in the cloud. At the time, real-time streaming from vessel to cloud was considered too risky and incompatible with the Agile development approach selected. The team has chosen incremental MVPs (minimum viable products) that could deliver immediate operational benefits whilst being aligned to the general roadmap.

After nearly 20 surveys using this system, data transit times observed range from approximately 45 minutes for small 4D spreads to around three hours for large exploration spreads. This performance has delivered most of the anticipated commercial benefits. It is difficult to identify significant additional value from moving to a fully real-time vessel-to-cloud system, especially given the added complexity and bandwidth redundancy requirements such a system would entail. The optimal approach for real-time interaction may instead involve leveraging onboard data in combination with low-latency connectivity for remote collaboration when necessary.

In streamer seismic acquisition, QC is traditionally divided into real-time and offline processes. Real-time QC, performed onboard by observers and QC geophysicists, relies on live attribute displays to enable immediate interventions — such as detecting an autofire or air leak to prevent costly infill or reshoots. Offline QC, typically executed at end-of-line (EOL), provides statistical checks against contractual specifications (e.g., less than 3% bad channels per streamer). Historically, rigid threshold-based acceptance has evolved into more intelligent, data-driven approaches such as Marginal Line Acceptance, where processing demonstrates that data remains fit-for-purpose despite marginal deviations, avoiding unnecessary reshoots. Machine learning has accelerated this trend; for example, RIDNet, a convolutional neural network for marine seismic denoising (Farmani et al., 2023), can often recover marginal data — though its computational demands are significant. This is where cloud scalability becomes transformative: workflows that transfer raw data to the cloud enable these advanced, compute-intensive QC processes to run quickly and cost-effectively, far beyond what is feasible onboard.

Currently, vessels often generate only a brute stack onboard, typically a 2D stack of a single subsurface line using hydrophone data only, representing less than 1.5% of available traces in a 12-streamer, triple-source configuration with dual-component streamers. With data in the cloud, there is virtually no barrier to producing a full 3D stack cube during acquisition, delivering far richer insight into the acceptability of the entire line.

Another critical factor is the growing prevalence of triple-source acquisition. Advanced deblending algorithms (e.g., Udengaard et al., 2025) remain too computationally heavy for timely onboard execution, excluding them from Marginal Line Acceptance decisions. Cloud-based workflows may change this paradigm, enabling application of advanced processing techniques during acquisition — unlocking QC and decision-making capabilities that were previously out of reach.

Conclusion and outlook

This paper has outlined how LEO satellite connectivity has rapidly transformed the acquisition and delivery of marine seismic data. The technology is proven, reliable, and operational at scale.

The remaining barriers are less about capability and more about cultural inertia — reluctance to change or hesitation to be perceived as a first mover. These tendencies are counterbalanced by an economic environment that increasingly favours Opex-based solutions over Capex-heavy alternatives.

Looking ahead, further development is inevitable. We are focused on greater standardisation and automation — automation that delivers tangible benefits while simplifying workflows. The foundation established through this work positions the industry to capitalise on these advancements, paving the way for more agile, scalable, and cost-effective seismic operations.

Acknowledgements

The development of these projects in TGS has been a collaborative effort involving many individuals whose contributions have been invaluable, even if they are too numerous to list here. Sincere appreciation is also extended to our satellite communications partner, Marlink, for its exceptional support and patience in helping to implement the vessel connectivity solutions that underpin this work.

References

Ewig, E., Brittan. J., James, C., Brodersen, J.O. and Olsen, S. [2024]. Adopting technology to revolutionise and accelerate the flow of seismic data from sensor to customer. *First Break*, **42**(2), 69-73.

Farmani, B., Pal, Y., Pedersen, M.W. and Hodges, E. [2023]. Motion sensor noise attenuation using deep learning. *First Break*, **41**(2), 45-51.

Langhammer, J. and Bennion, P. [2015]. Triple-Source Simultaneous Shooting (TS3), A Future for Higher Density Seismic? 77th EAGE Annual Conference & Exhibition, *Extended Abstracts*.

Udengaard, C., Alexander, L., Gabioli, L., Cochran, S., Cvetkovic, M. and Unger, D. [2025]. Deblending 3D NAZ Data in High Sea Surface Current Environments. IMAGE 2025, *Conference Proceedings*.