

# Intel and XENON Help Oil Search Dig Deeper Into Sub-Surface Oil and Gas Analysis

*Unique oil sector technology project returns strong cost to benefit ratio*

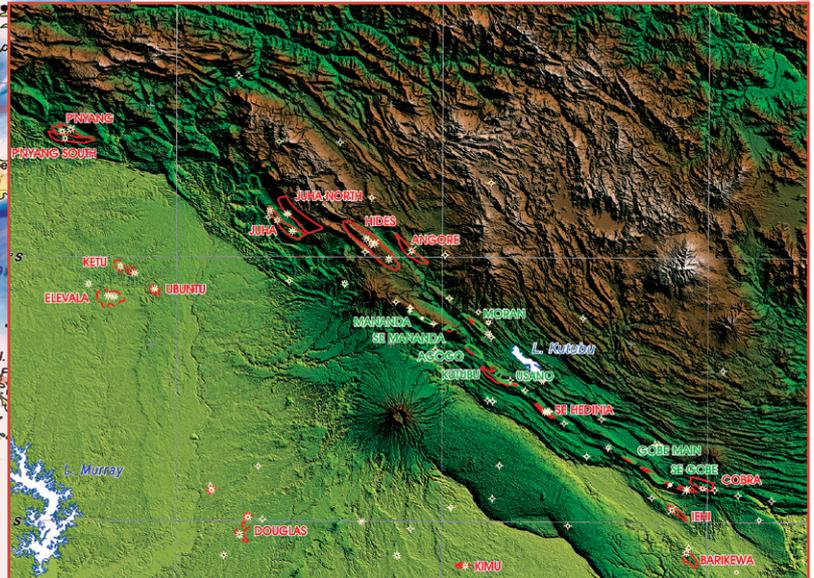
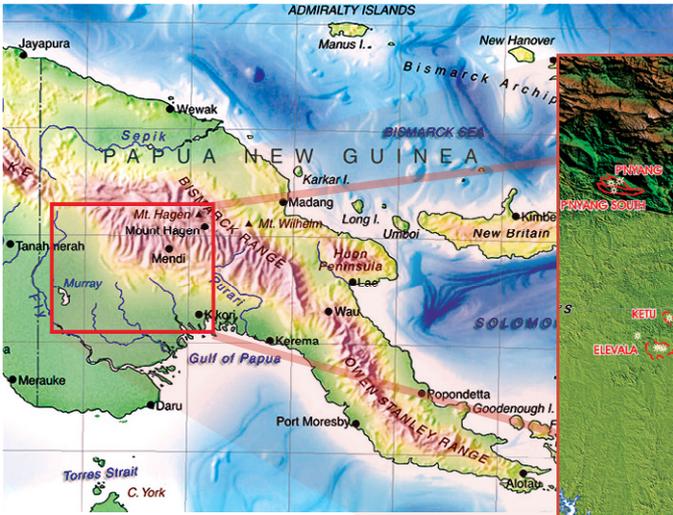
## BACKGROUND

Oil Search was established in Papua New Guinea (PNG) in 1929. More than 98% of its assets are in PNG, where it operates all of the country's producing oil fields, holds an extensive appraisal and exploration portfolio and has a 29% interest in the PNG LNG Project, operated by ExxonMobil PNG Limited. This world-class liquefied natural gas (LNG) development, which came on-stream in 2014, has transformed Oil Search into a regionally significant oil and gas producer

with a long-term, low cost, high quality LNG revenue stream. The Company also holds a significant interest in the proposed Papua LNG Project, operated by Total, which has the potential to become PNG's next major LNG development.

Oil Search is listed on the Australian and Port Moresby stock exchanges (OSH) and its ADRs trade on the US Over the Counter market (OISHY).





### CHALLENGE

Geoff Balemans, Telecommunications and High Performance Computing Manager at Oil Search works with the upstream technical computing groups. He explains that with low oil prices, it's a much less forgiving market for error and so detailed modelling of assets, using sophisticated technology, becomes even more of a priority.

Technology has recently proven an area of innovation for one very specific but critical group, the Sub-Surface team, and one whose benefits are likely to cascade across the company. Essentially, the Sub-Surface team's role is to ascertain as accurately as possible how much oil is available in a certain field, how deep a store might be and the potential flow rate – and that relies on some heavy duty modelling and simulations.

Geologists talk about models in terms of cells – one cell could represent 100m x 100m x a certain thickness (e.g. 5 meters), – so when Oil Search set out to model all of their PNG assets in a regional model they were looking at a 3D geological model in excess of a billion cells.

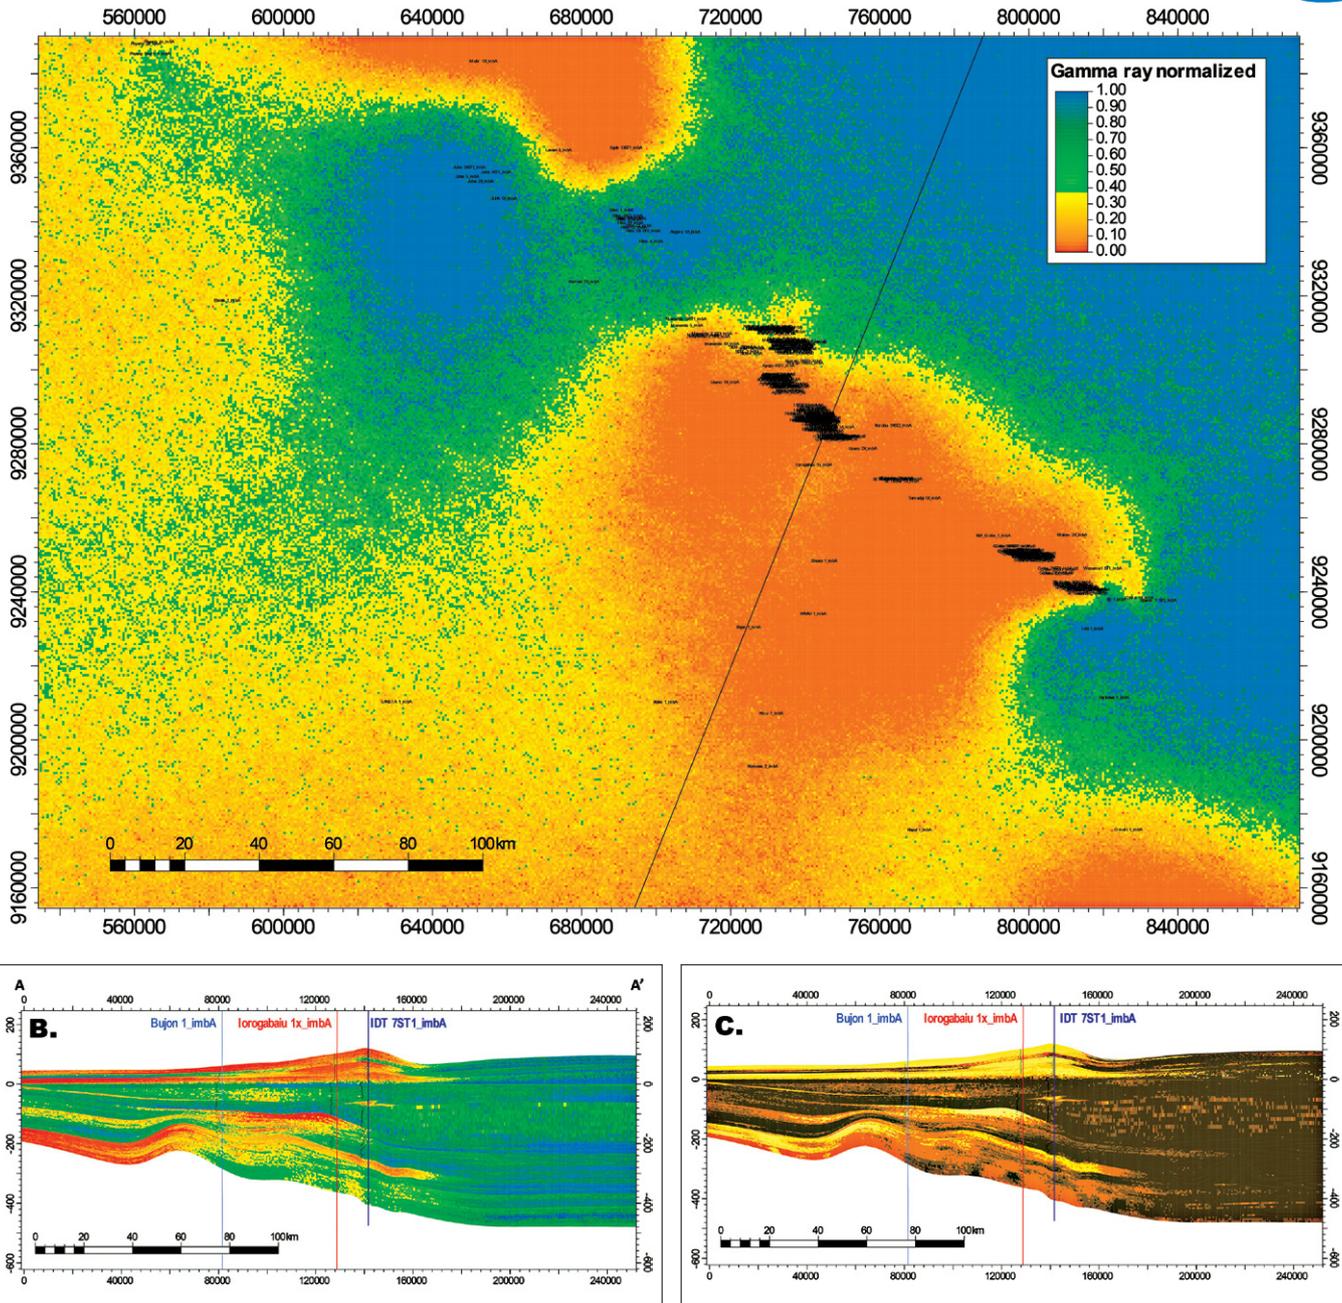
The project to help improve this process was conducted in collaboration with high performance computing specialist, XENON, who recommended NVMe based Intel® SSD DC P3700 Series and new Intel® Cache Acceleration Software for Windows\* Workstations (Intel® CAS-W) technology for faster rendering and transmission of modelling data.

MILLION YEARS (Myr ca. 1927)	AGE <small>(Data in brackets from Harland et al. 1992)</small>	BIOSTRATIGRAPHY			LITHOLOGY <small>SW LANDWARD NE BASINWARD</small>
		Robertson	BP	AUSTRALIAN	
125	Berriasian <small>(137.5)</small>	E.torynum	EK10	P10	TORO SST. TORO A TORO B TORO C TORO D
		L.pinosum	EK11		
		P.apiculatum	EK12	P11	
		P.mirabilis	EK13	P11A	
130	Tithonian <small>(144)</small>	R.serrata	LJ8	P12	IMBURU FM. IMBURU A IMBURU B IMBURU C
		B.simplex	LJ4	P13A	
		N.similis	LJ5	P13B	
135	LATE JURASSIC <small>(151)</small>				JAGIEU MBRA A & B

Oil and gas modelling applications are very different to generic technology applications. Engineers and geologists in the Sub-Surface modelling group at Oil Search use highly sophisticated modelling applications to manipulate large data files, and so the environment requires a high performance (and specialised) technology environment to manage these massive datasets.

By way of example, a modeller might need to load a project to perform modelling and calculations on an active oil field. This would require extremely large files to be sent over the network to a workstation. The modellers simply couldn't load the model from the regional project because the network protocols to deliver these files weren't fast enough because of the file size. Additionally the applications need to be able to write back results from these calculations causing a bottle neck, time-outs and crashes. This severely slowed down the modellers' workflows.

▲ Figure 1: Location, study area and reservoir stratigraphy.



▲ **Figure 2:** A) Normalised property map. B and C) Cross-section (A-A') of reservoir properties used to characterise the subsurface.

Each pixel represents a cell.

As a result, the modellers would split up and load the model in three parts and when they then performed the simulations, they could often only do so once per day without crashing. Typically it would take around twenty hours to process the model because of the application constantly reading and writing over the network.

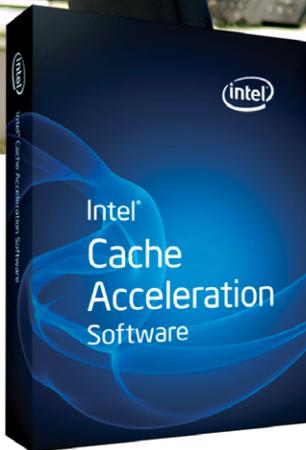
This meant that everyone in the team was reliant on this initial process taking place to progress each next step of the project.

“To positively impact the business, it was all about how fast we could process these models. That’s when the modellers came to us with the ‘billion cell model challenge,’” said Geoff Balemans.

## SOLUTION

Having defined the problem and identified the bottleneck, that’s when Balemans turned to long-term consulting partner, XENON. As a specialised high performance computing consultancy with deep blue-chip technology industry relationships, XENON often trials new technology ahead of the general market release and is well-positioned to advise companies when facing specialised computing challenges. On this occasion, XENON designed a solution using Intel® technologies, Intel® SSD DC P3700 Series and Intel® CAS-W technology.

Enterprise data sets are growing at a staggering rate, and applications need to access that data faster than ever before.



▲ Intel® Cache Acceleration Software combined with high performance Solid State Drives (SSDs), increases data centre performance via intelligent caching rather than extreme spending.

Hard drive speeds cannot keep up with the significant performance gains provided by Intel® Xeon® processor-based servers. Critical business applications are “data starved” as they wait for legacy mechanical hard disk drives and network I/O bottlenecks to catch up.

Intel® CAS-W offers an innovative solution to this challenge by using an Intel® SSD for data centers as a cache for data that is frequently accessed. By storing a copy of this “hot” data in the SSD, server applications can access the data and run faster. Intel® CAS-W offers up to 12 times more performance on transactional database processing and up to 36 times faster processing of I/O intensive virtualised workloads.

Oil Search and XENON started a dialogue with Intel® locally and with Intel® headquarters to trial the Intel® CAS-W software together with Intel® SSD DC P3700 Series for an initial three months in a non-production environment. Testing proved so successful that it was a straightforward business case to get approved and rolled out.

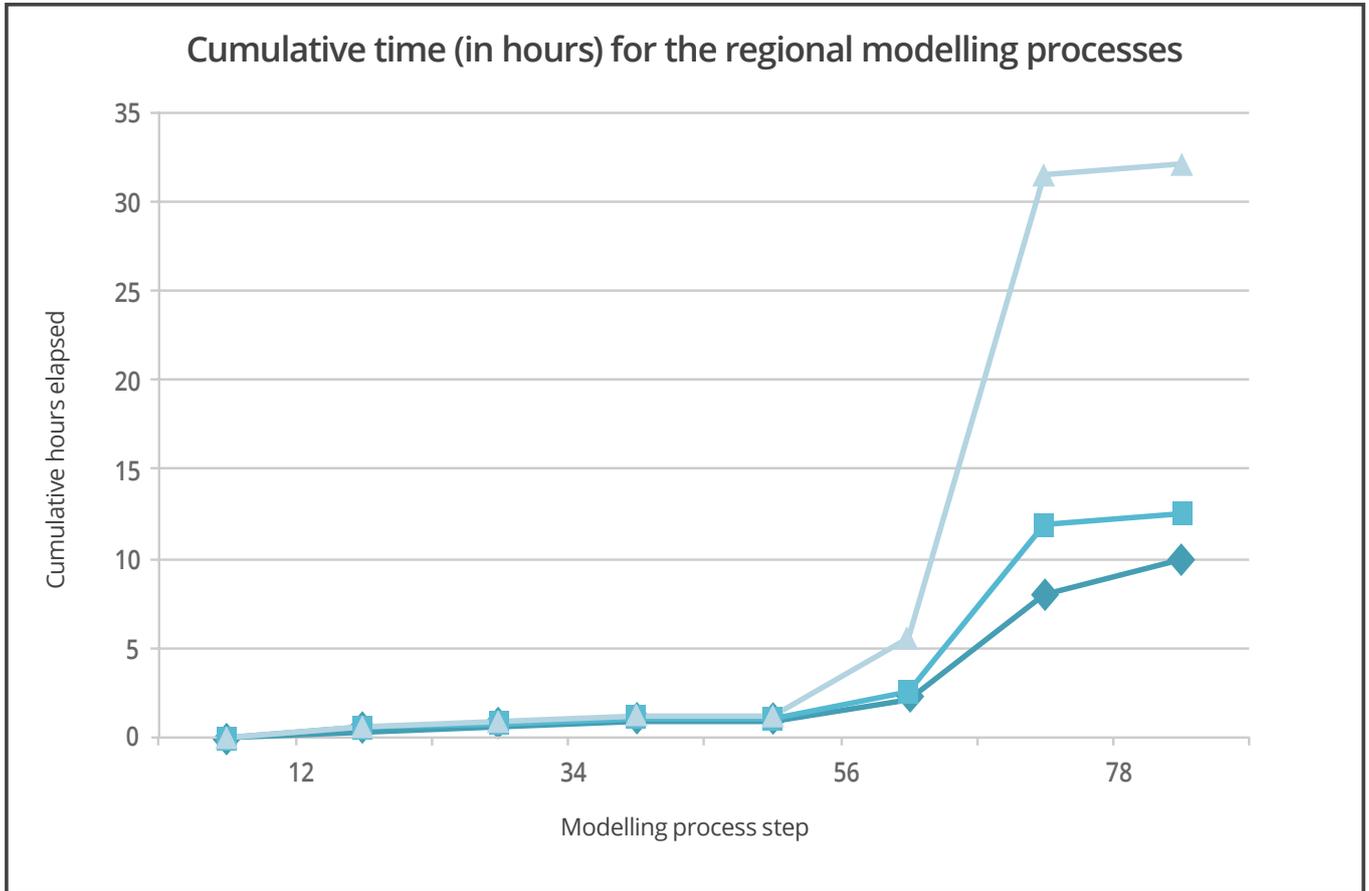
“It sounds almost trivial for such a complex set of processes to say that it was practically plug and play and worked right off the bat – we weren’t expecting that. XENON was

creative and innovative in knowing this product and how it would perform in our environment, and the Intel® CAS-W software worked exactly as we needed,” explained Sigurdur Jonsson, senior reservoir engineer at Oil Search.

Balemans explains that the ‘smarts’ are in how the technology knows what to cache and that the only other way would be to copy data down to the local machine then copy it back. However, that’s not something his team would recommend from a back-up and security perspective.

“It’s not a process-heavy solution. Yes there’s an extra step of mapping a drive but certainly nothing hard or time-consuming about it. The team operates largely as normal and that’s where the Intel® CAS-W product really shines.”

“Intel® is excited to be working with customers, such as XENON and Oil Search, to deliver innovative solutions to take data center performance to new levels.” David Tuhy, General Manager, NSG Storage Software, Intel® Corporation. “The unique combination of Intel® Cache Acceleration Software and Intel® SSD Data Center Family for NVMe enables end-users to address real-world problems.”



▲ **Figure 3:** The cumulative time taken to run the regional modelling workflow.

Balemans has a view that if he can change something to get a 20-30% advantage then it's worthwhile investigating. This has been a proven approach in other areas too, for example, selecting a specific type of CPU to instantly run 4 times faster – saving on license fees and running faster simulations.

they have and how to model them. The project is quite unique globally in the oil sector.

“To put simply, the more simulations we can run, the more accurate we can be, so this modelling can have a really direct impact to the business,” added Balemans. For us, this project represents significant return for a modest investment, not just field by field, but also regionally. This spreads out into oil and gas exploration more generally for the company.”

## RESULTS

Now, the senior geologist can do in a day what he was doing in a week, which means he can do in a week what took a month.

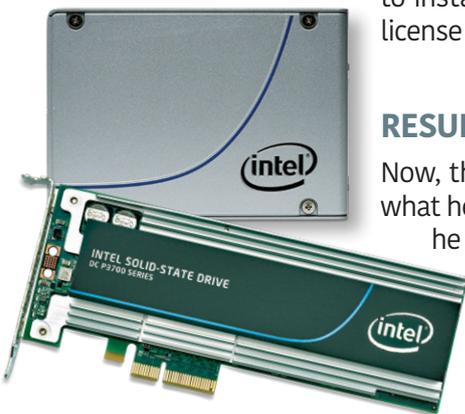
It's not just a huge time saver but it has huge ramifications for the whole team. It's essentially like having increased the team four or five-fold.

Fig. 3 shows the improvements in time taken to run the regional model and clearly shows a significant speed improvement.

With everything hanging off the regional model, Oil Search now has a much better understanding of the oil and gas assets

## FUTURE

“Now that our core need has been taken care of, we will look to utilise Intel® technologies, Intel® SSD DC P3700 Series and the Intel® CAS-W, to further improve data warehousing, databases, and production accounting. Indeed, we will continue to work with XENON and Intel to extract value from our existing technology hardware, and the sophisticated applications we use, wherever possible,” concluded Balemans.



▲ Architected with the new high performance controller interface, Non-Volatile Memory Express\*, delivering leading performance, low latency, and quality of service.