Exploration seismology and the return of the supercomputer

Exploring scalability for speed in development and delivery

Sverre Brandsberg-Dahl
Chief Geophysicist, Imaging and Engineering

Marine seismic acquisition

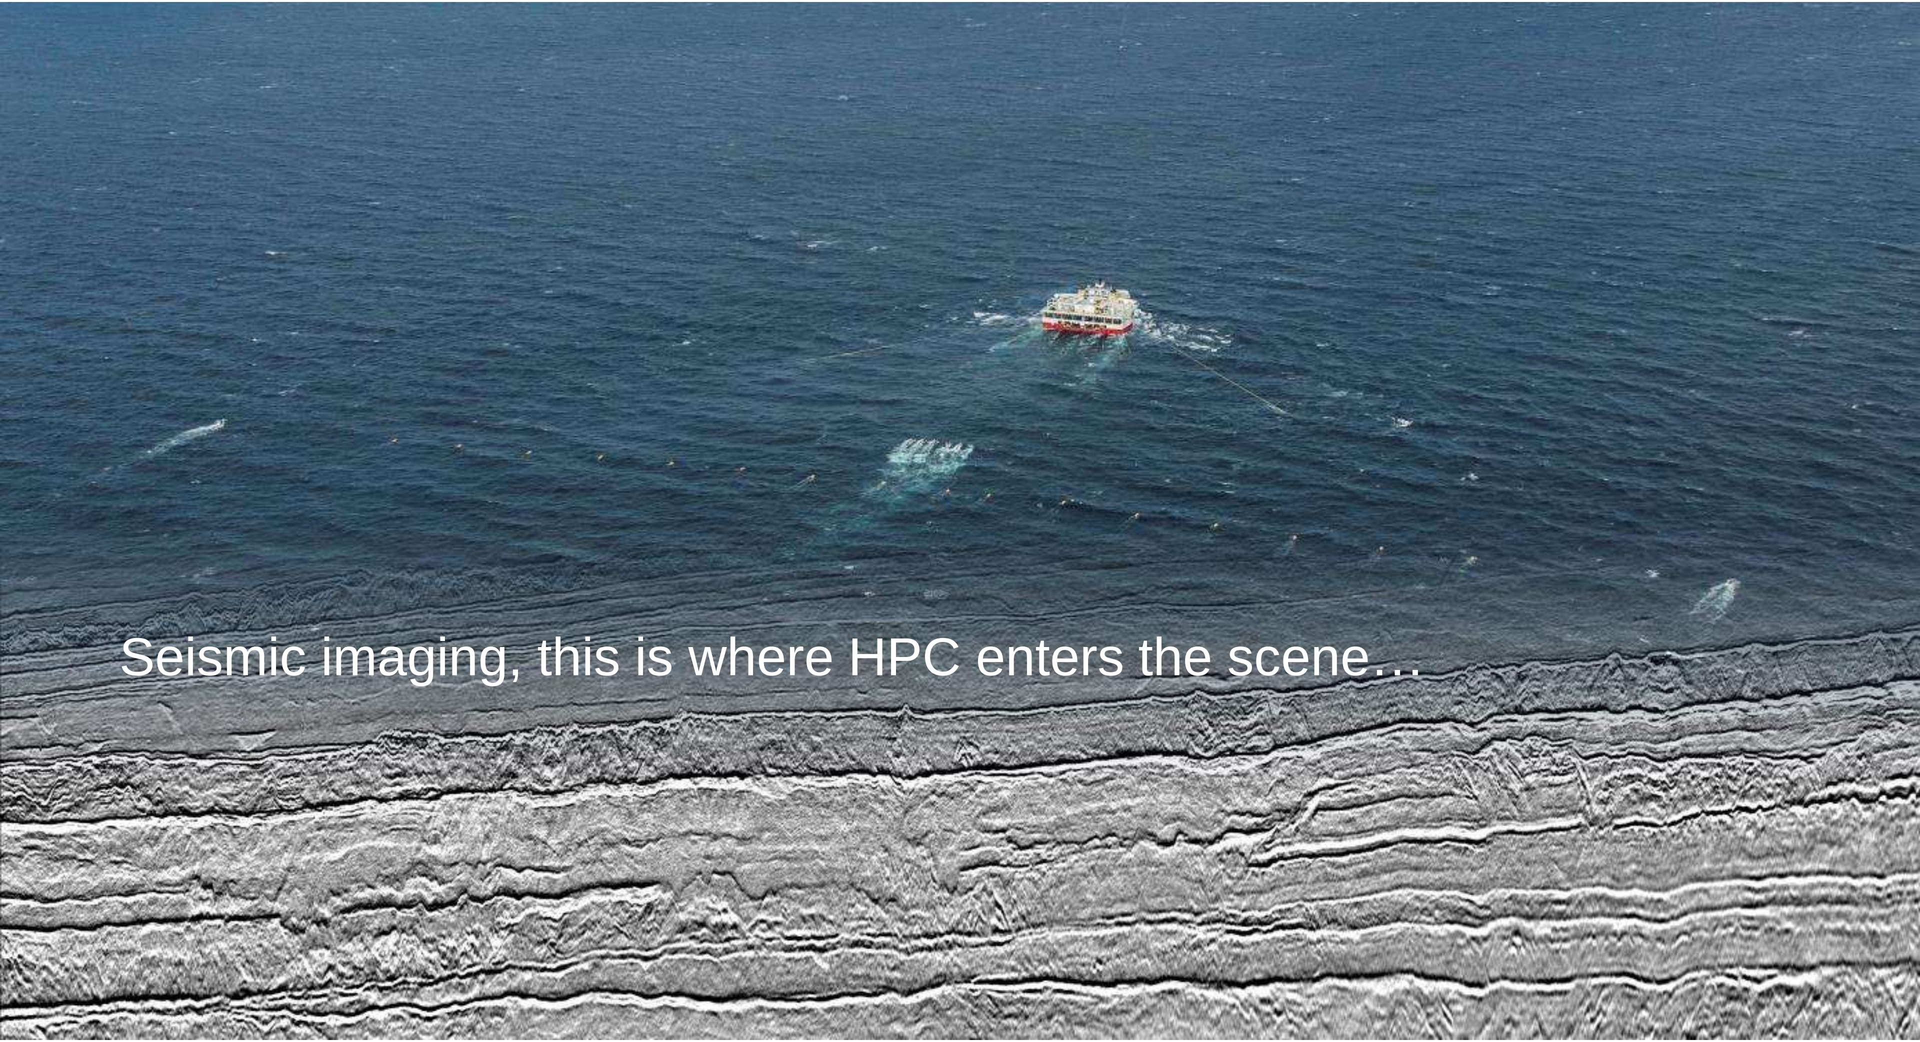
Seismic imaging, this is where HPC enters the scene…

Seismic imaging: Create an image of the subsurface from surface measurements

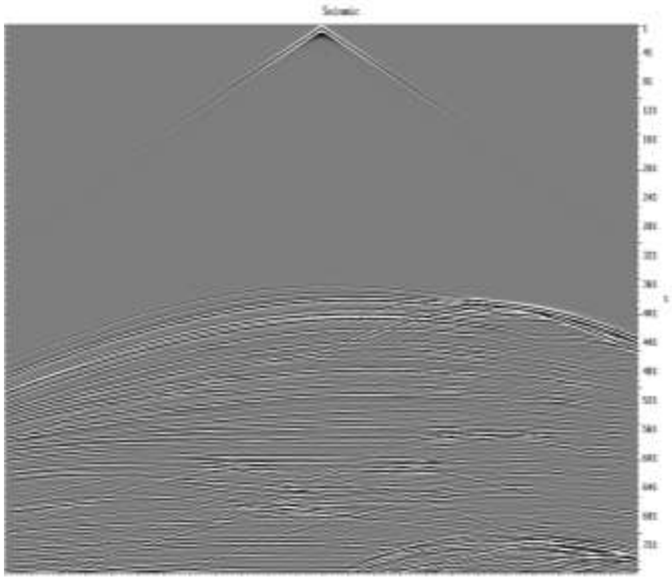

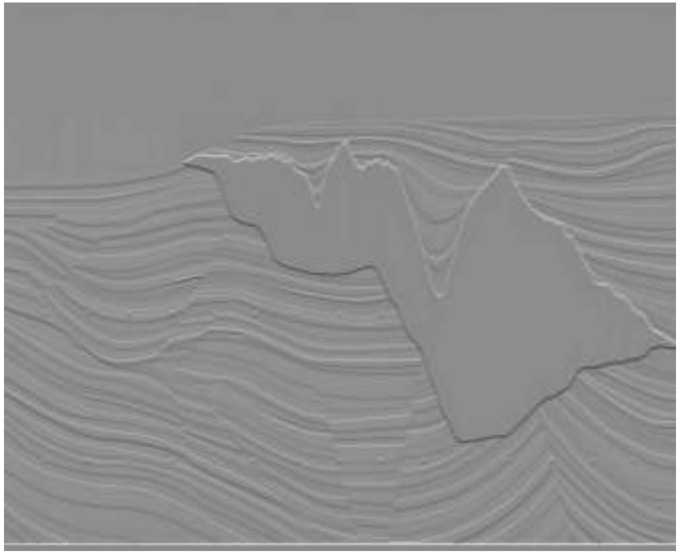

$$data(\mathbf{s}, \mathbf{r}, t) \underset{vel}{\longrightarrow} image(x, y, z)$$

How exploration seismology has been carried out for decades…

Firing a source and recording the reflection from the sub-surface:

- This forms a seismic shot record

- Treating each shot as an individual wave equation realization lead to the "embarrassingly parallel regime" of seismic imaging

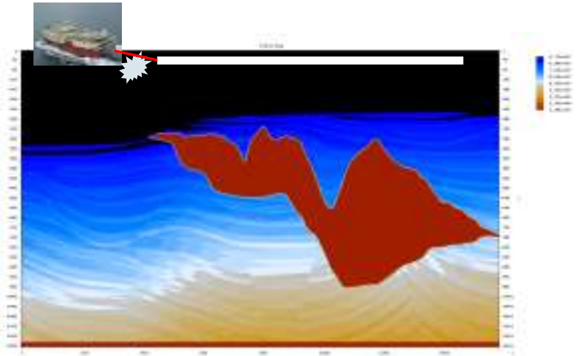

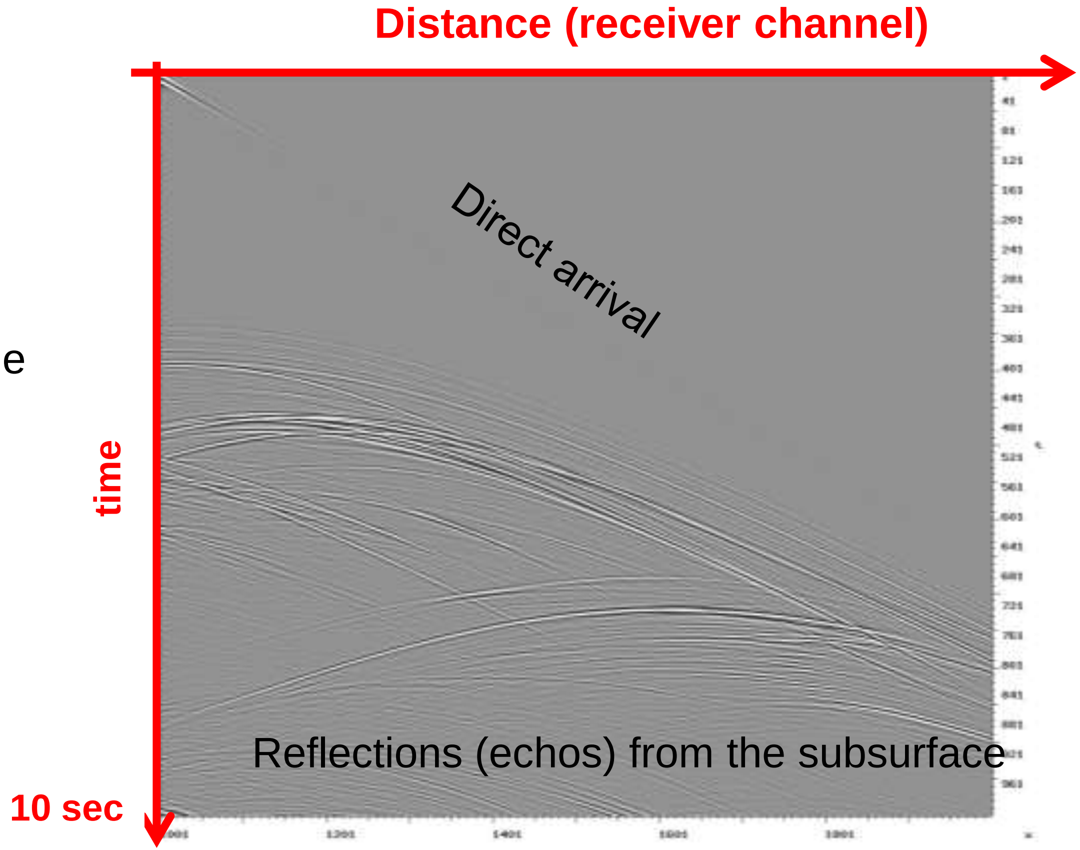

So how do we image, we solve the wave equation

$$\frac{\partial^2}{\partial x_i^2} U(x_i, t) = \frac{1}{v^2} \frac{\partial^2}{\partial t^2} U(x_i, t)$$

It is used (in approximate ways) to model seismic waves as they propagate in the earth

Some issues we face when imaging seismic data

- The seismic waves propagate through an ‘unknown’ medium, we can only measure on the (sea) surface

- We need to recover the medium properties from the data itself

- The subsurface can have large velocity and density contrasts, making the above a very complex task…

- Exponential growth in compute cost as function of frequency
  - wave equation has 3 spatial dimension and 1 temporal, so solution schemes scale as f^4

Seismic image (migration) compared to reflectivity

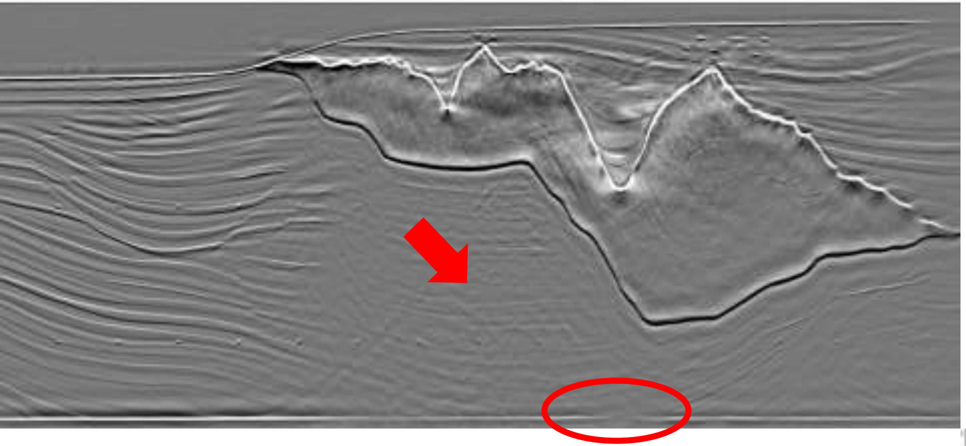

A typical RTM image

The true reflectivity

Image formed by applying adjoint of modeling operator.

No correction for data aperture, (incomplete) data geometry or illumination.

Why PGS did computers a bit differently

In a good position to be different

- Relatively small market share in data processing, but growing
- Focused company: marine streamer seismic
- In-house software platform used only by PGS

This allowed us to do things differently, and not copy what was current standards or adopted trends in the industry

- "Sticking to the plan", business and technology strategy aligned and in place since 2009

PGS Imaging

Vision

- To deliver the highest quality images of the subsurface in the shortest possible time

Mission

- To become the market leader in high-end imaging
- To achieve a market share similar to that of PGS marine acquisition

Strategy

- Improve quality and consistency of services
- Grow behind our fleet
- Increase productivity (more automated processes)
- Develop new interaction with customers

Guiding principles for PGS Imaging

- Acquire “complete data”: spatial sampling and bandwidth
- 3D operators throughout the processing flow, from wavefield separation through de-multiple
- Use the complete data in VMB and imaging: reflections, refractions and multiply-scattered waves
- Move towards full wavefield, and full-survey processing and imaging (ensemble-based processing)
- Leverage angle-domain image space for analysis, residual corrections and conditioning
- Inversion algorithms in processing and imaging

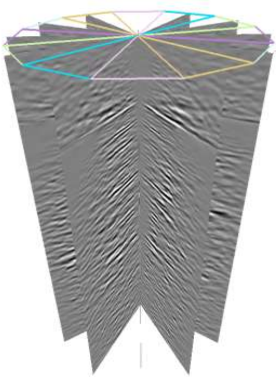

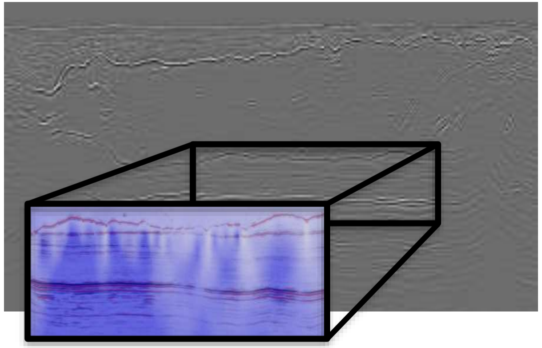

Late to the game, catching up to the data explosion…

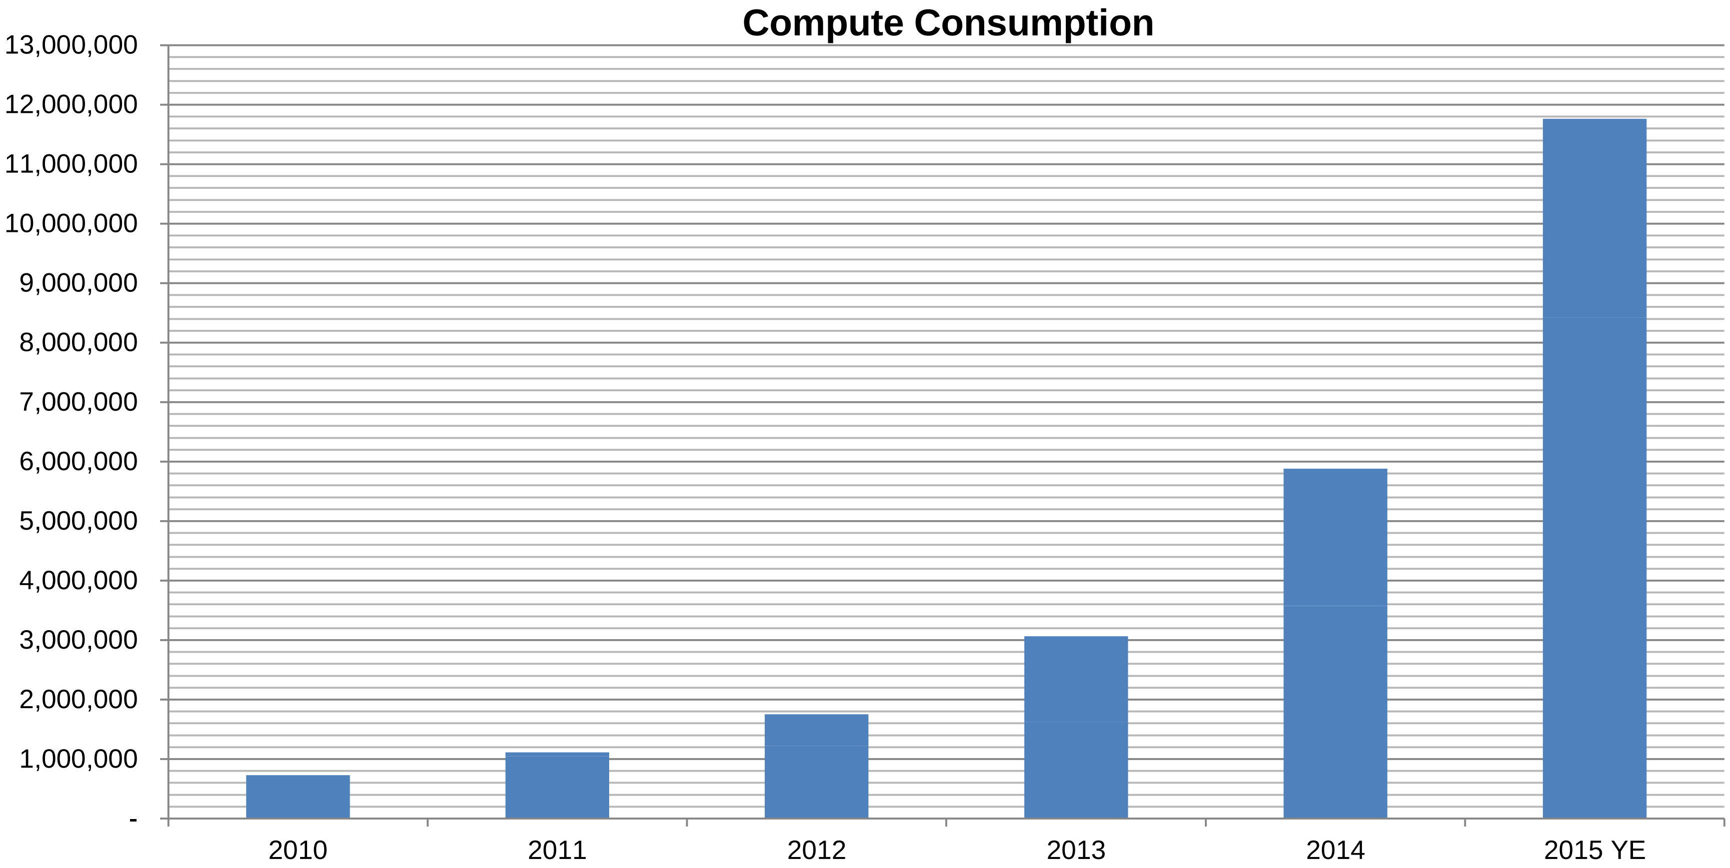

But also facing up to new realities…

FOTO: Colourbox

Relying on Linux clusters, but not at scale

Compute Consumption

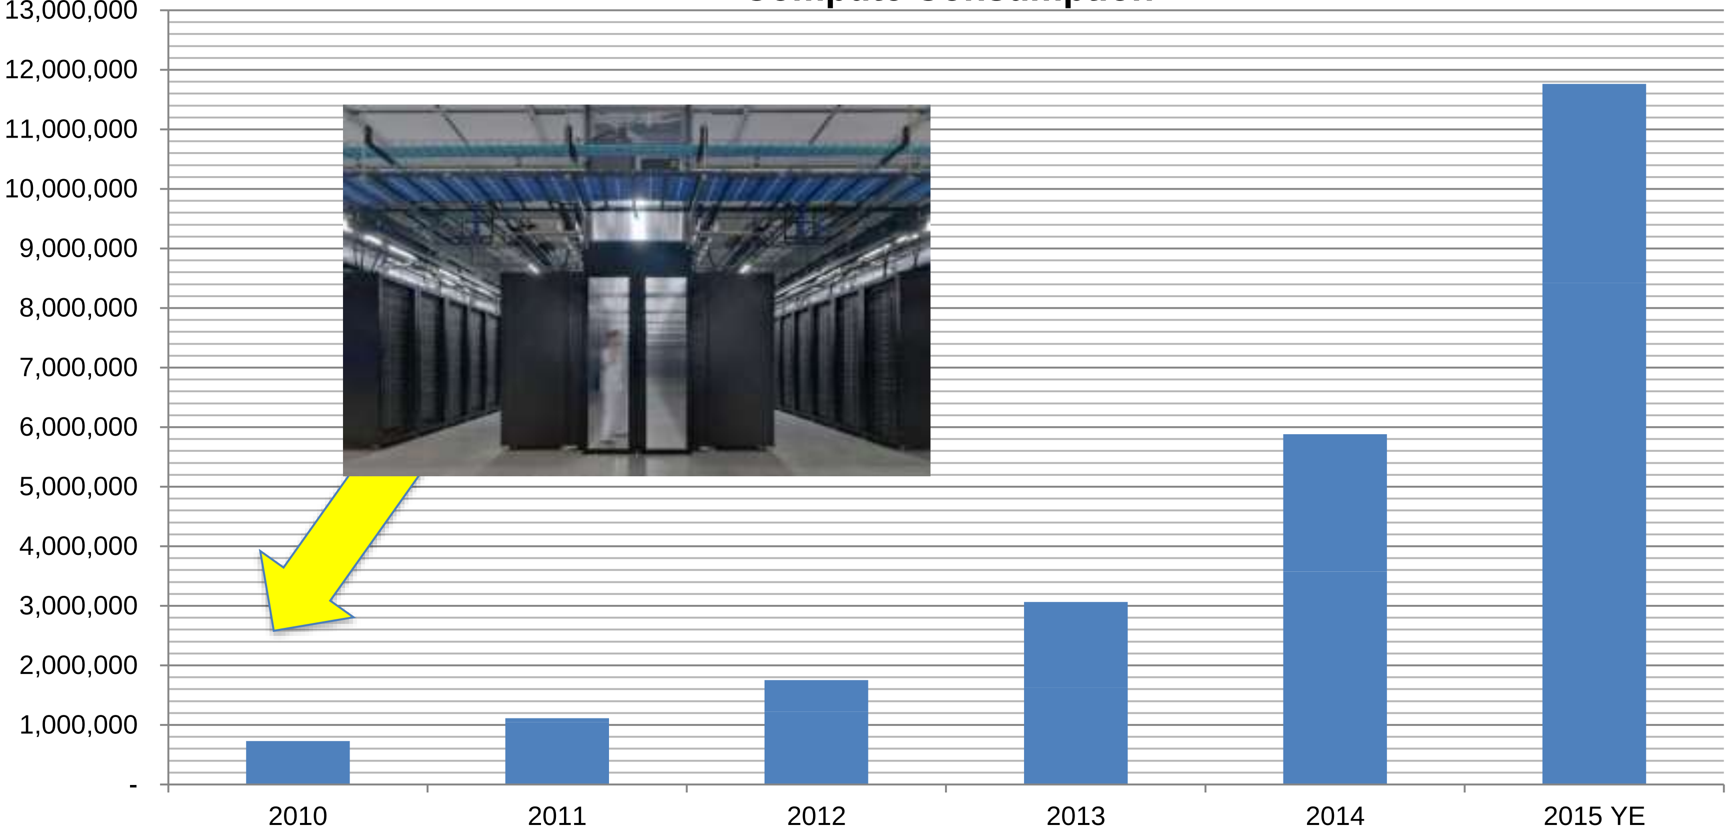

Missed the wave on GPU clusters…

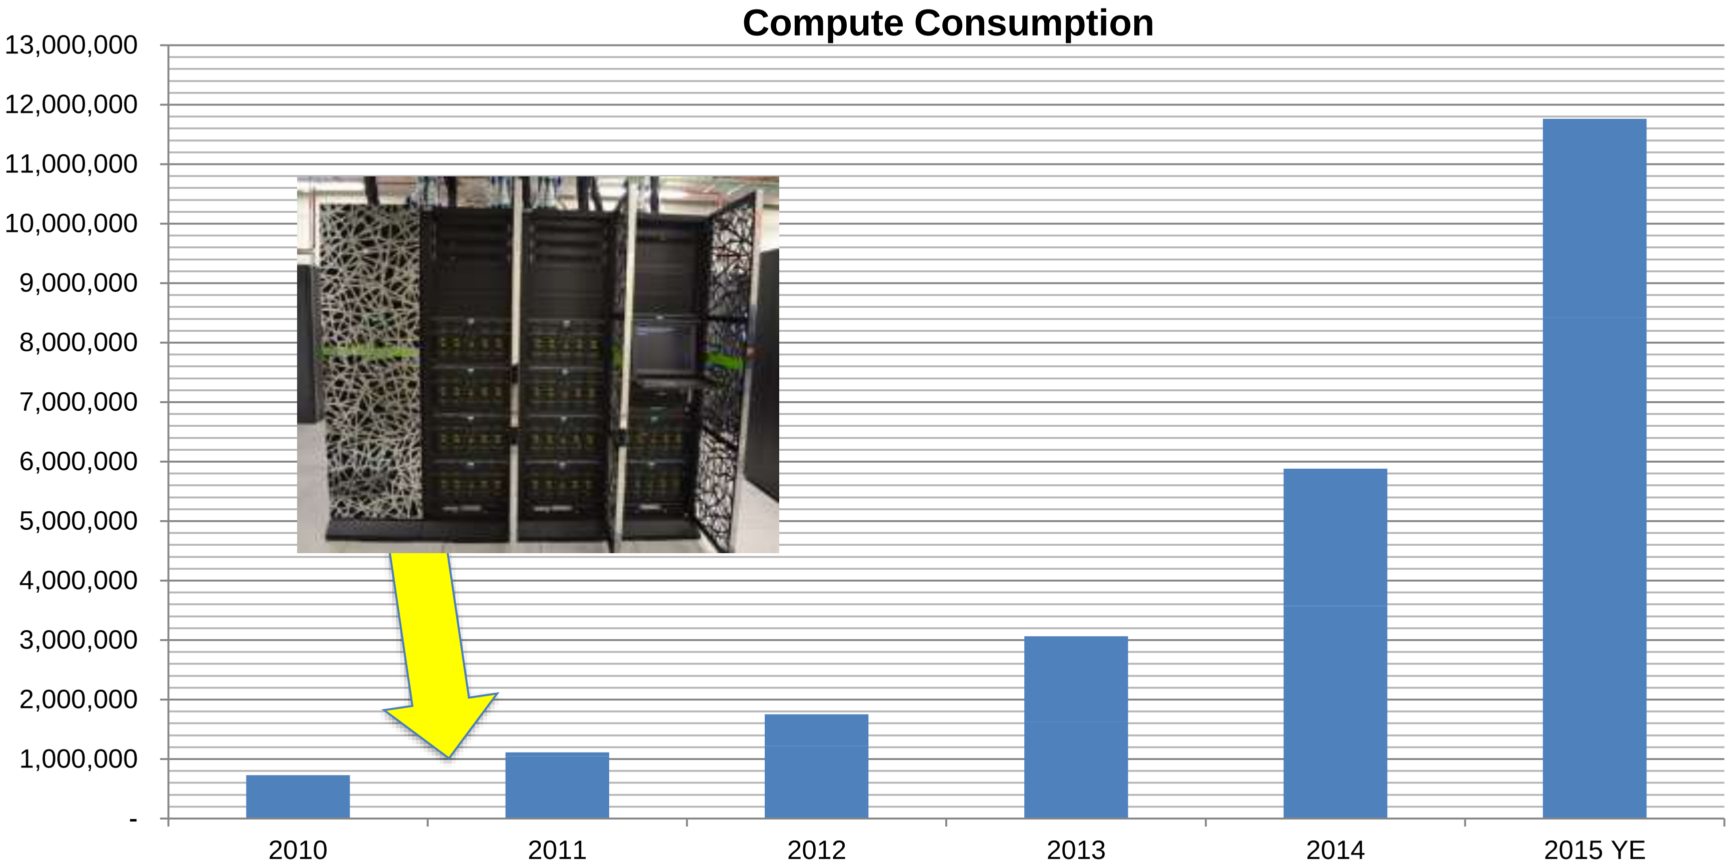

Rather than trying to implement what others were already doing…

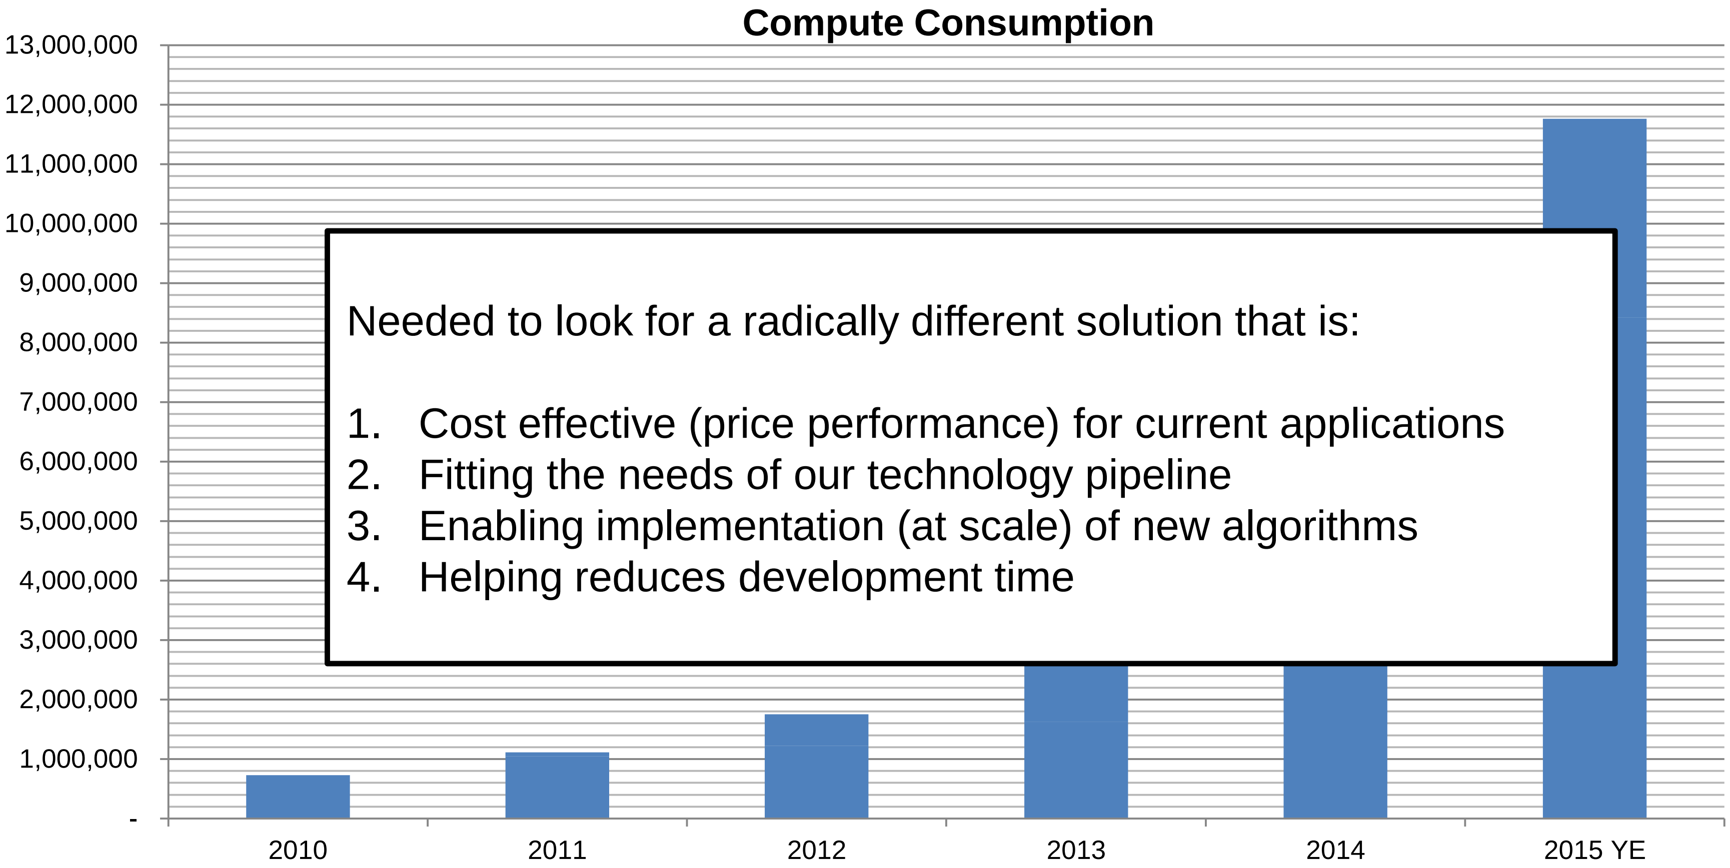

Needed to look for a radically different solution that is:

1. Cost effective (price performance) for current applications
2. Fitting the needs of our technology pipeline
3. Enabling implementation (at scale) of new algorithms
4. Helping reduces development time

Small XC30 system for R&D

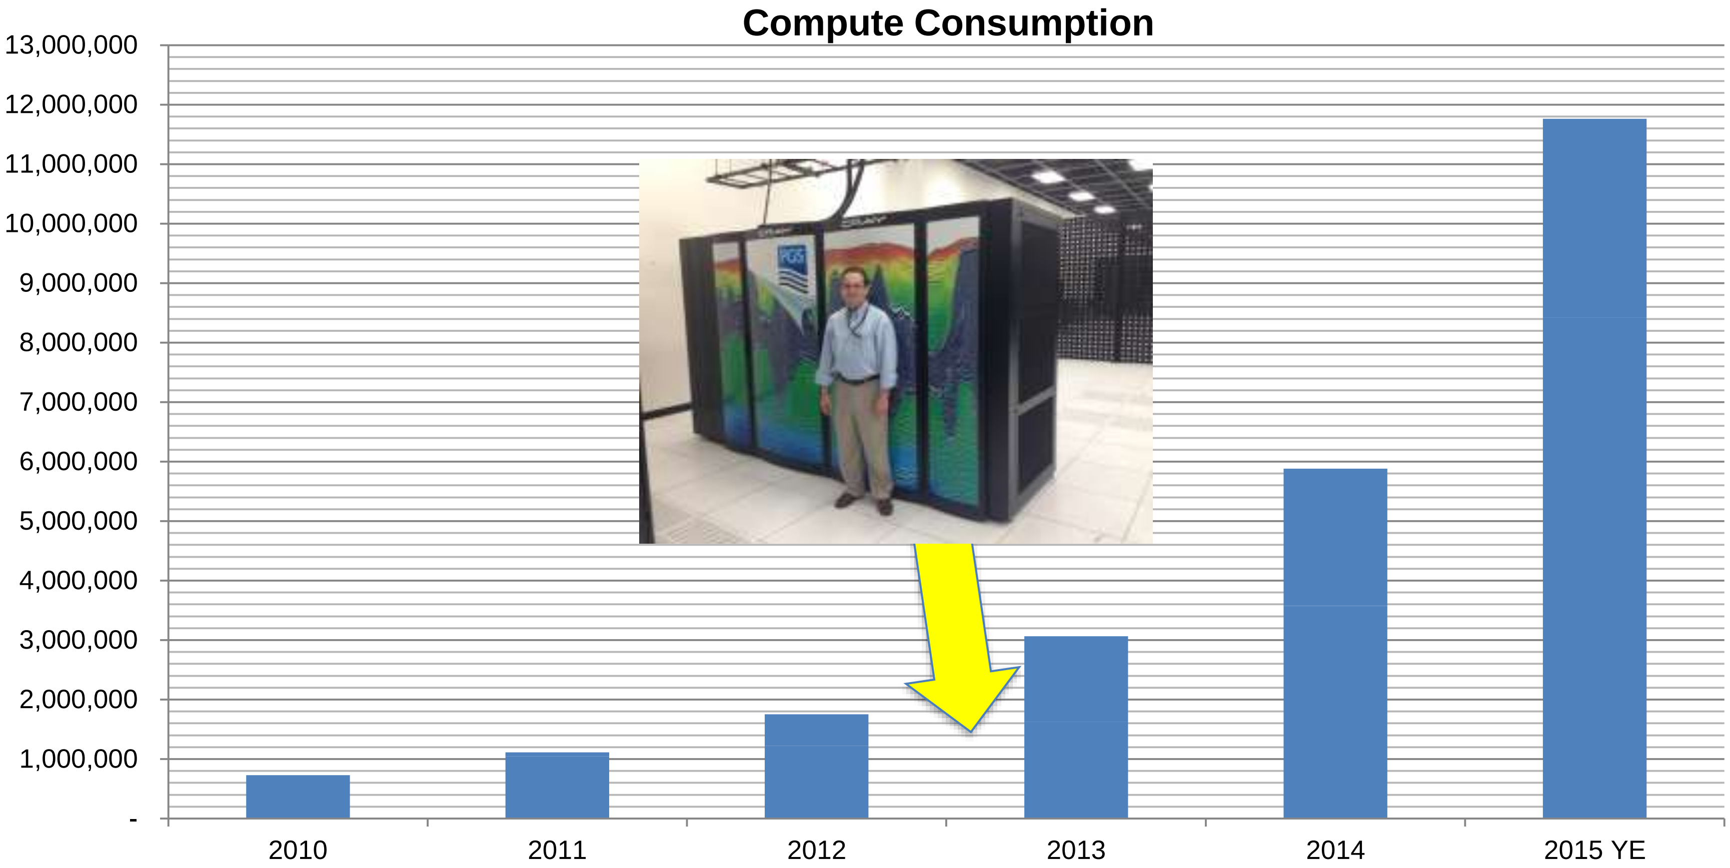

Moving to the next level; spurred by growth in business- and data-volumes, and readiness in the organization

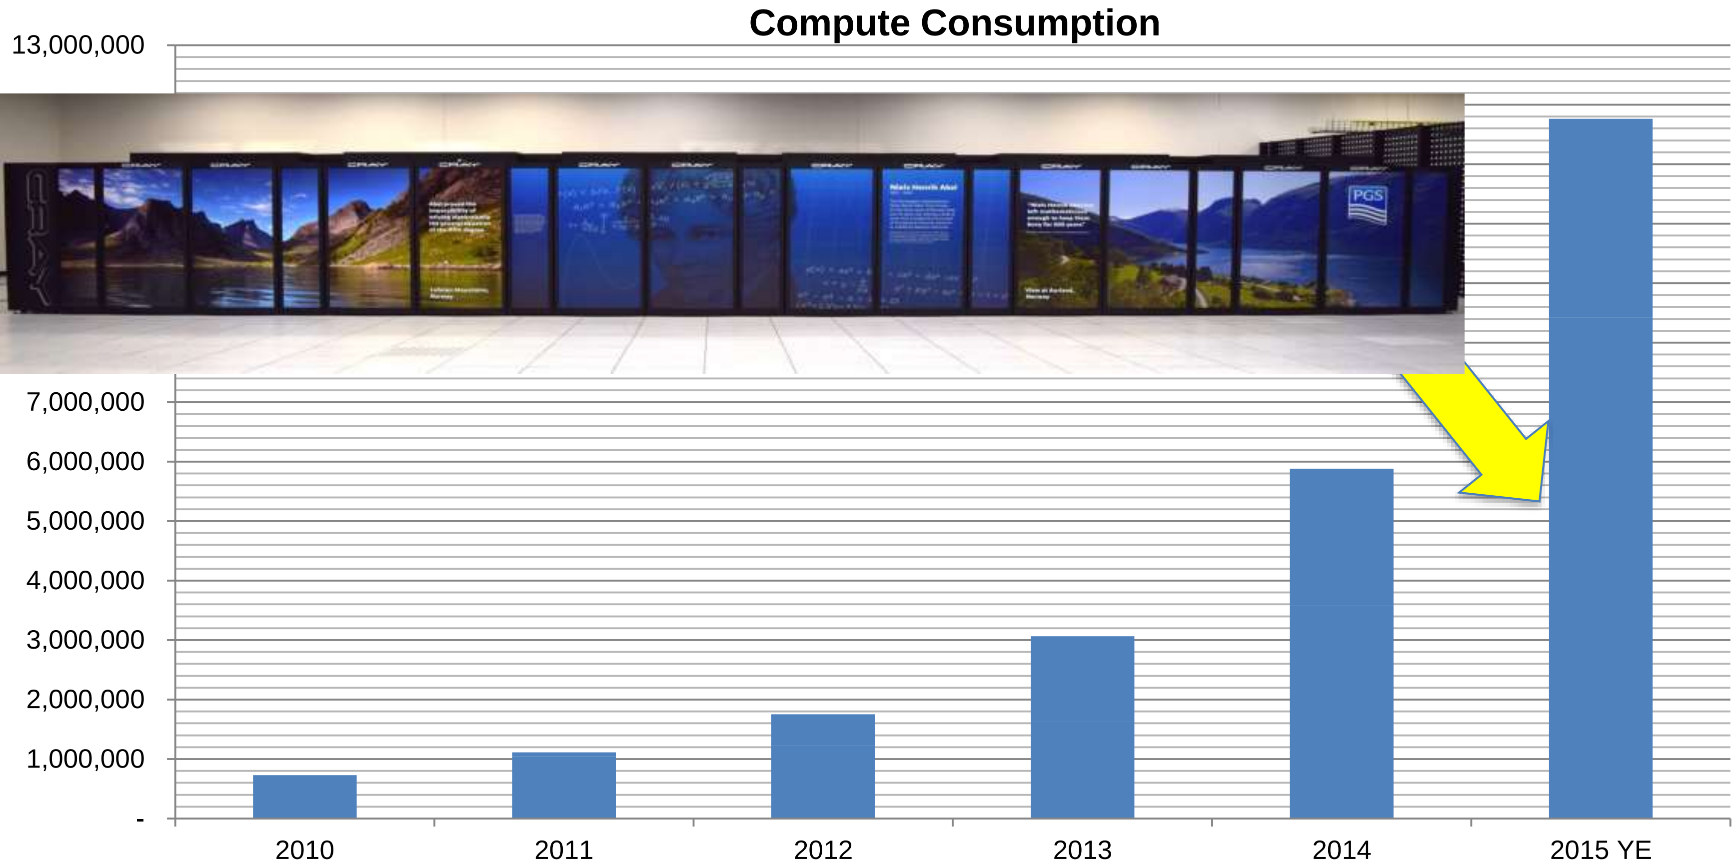

Grow behind our fleet, and keep up with a growing data flow

The big data challenge… in the seismic business

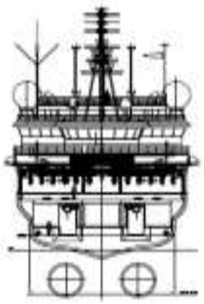

1995
6 streamers

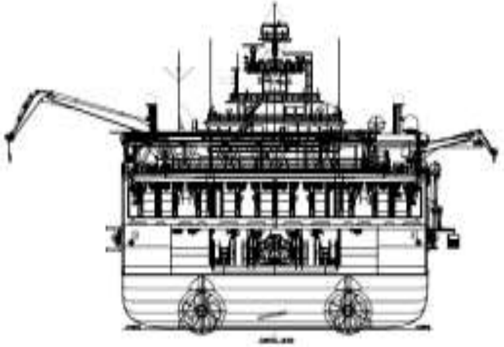

2005
16 streamers

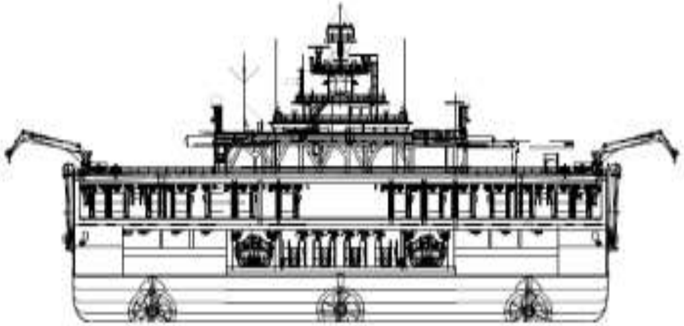

2015
24 streamers

Titan Class vessel

Seismic acquisition data volumes

- A typical streamer is 8000 meters long and contains 1280 receivers
- Data is recorded in time chunks or as continuous series, 2ms sample interval, generating 500 samples per second per receiver
- A streamer (single sensor) generates 640,000 samples per second
- One streamer spread (10 streamers) generates 6,400,000 samples per second
- Big spread, 20 streamers, dual sensor will generates 25,600,000 samples per second
- Typical acquisition can generate multi-TBs of data per day

Outline

- Why PGS did computers a bit differently
- Grow behind our fleet, and keep up with a growing data flow
- Triton, imaging a 660TB dataset using distributed memory and fast interconnect
- New geophysical science on a Cray XC40 supercomputer
  - WEI: Wave Equation Inversion
  - Deconvolution imaging condition
- Future challenges; software, people and machines

Triton, imaging a 660TB dataset using distributed memory and fast interconnect

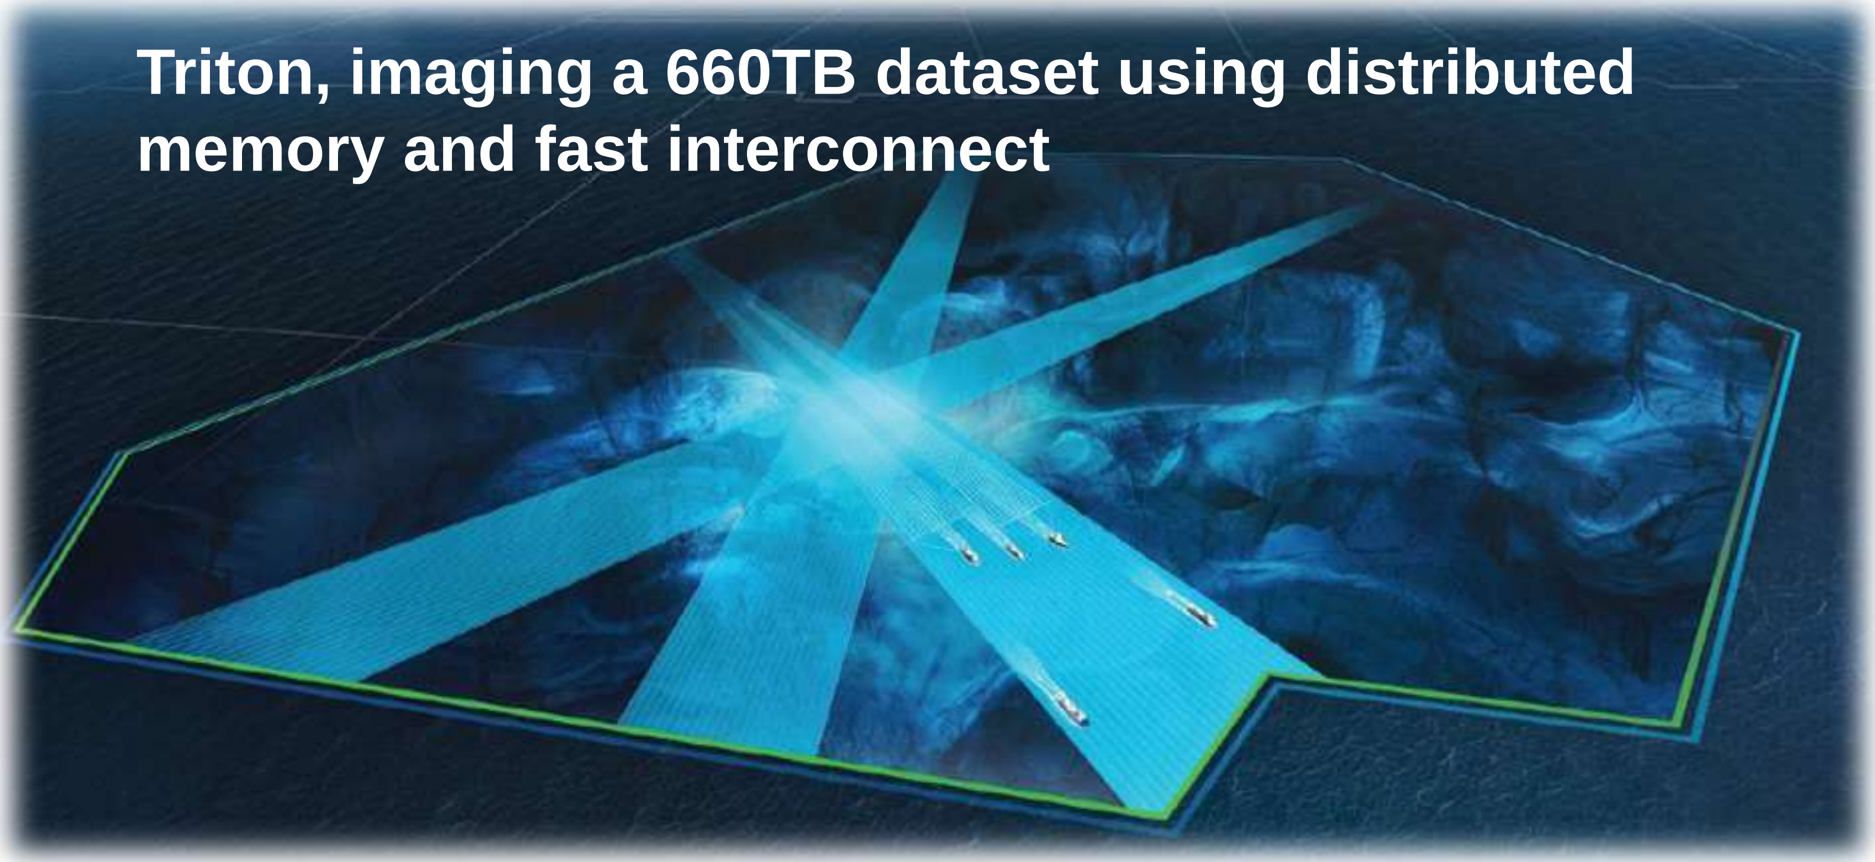

Triton acquisition

281 sail lines over 3 azimuths

2.6 million CMP km acquired

enough to go to the moon…
…and back
almost 3.5 times

Triton pushes us to the next level

Reverse Time Migration (RTM) 16km aperture, 16km depth

6 Azimuths, 46 Angles

RTM Problem size approaches 200 GB per shot

Hardware requirement -- 6000 nodes with 256GB memory.

… while waiting for the next bigger survey that would push us over 256GB of memory.

So it was decided to do something radically different

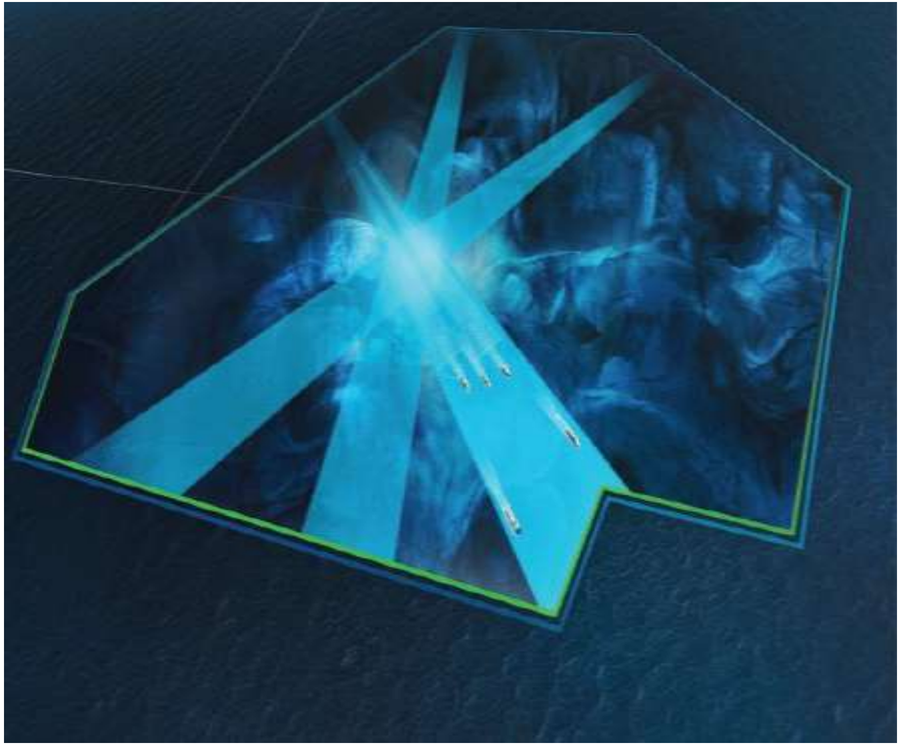

- Move RTM to a massively-parallel computer with distributed memory
- In the process the “embarrassingly parallel paradigm” was busted
- RTM with angle/azimuth gather output was implemented, and Triton production imaging was run on a Cray XC40 system

PGS Abel supercomputer

- 24 cabinet Cray XC40
- All CPU
- Almost 600TB of memory
- One-step all to all communication

This is a key technology differentiator for PGS

- Will help with delivery of current projects
- Will enable development of new imaging technology

Implementing RTM on the Cray

- Our Cray has a fantastic network, but nodes have no scratch disk.
- Legacy RTM implementation uses lots of /scratch to store wavefield snapshots
- RTM with pseudo-spectral extrapolator (FFT-based)

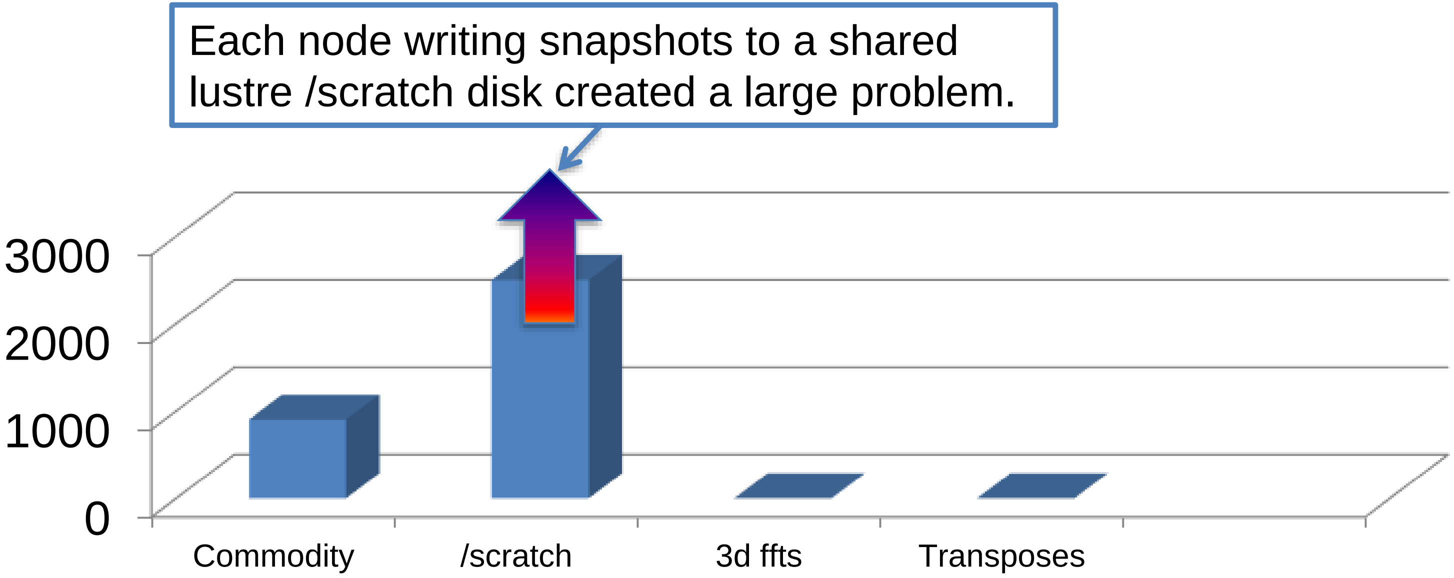

Implementing RTM on the Cray

Parallelizing one shot over multiple nodes and storing snapshots in memory is better, but distributed memory 3d FFTs slow things down.

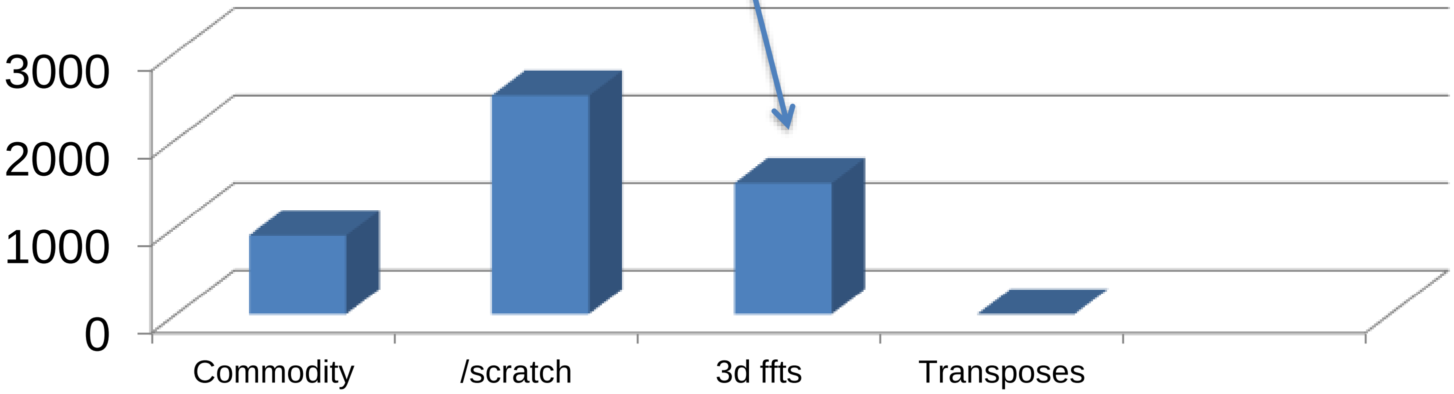

Implementing RTM on the Cray

- Each shot migrated using tens of compute nodes: no more limits on frequency and aperture
- Transcending the standard RTM implementations: densely sampled gathers, scalable
- Improved runtime over commodity: one node – one shot

Reorganize math to reuse transposes and work two dimensions at a time. Total aggregate runtime is faster than single node

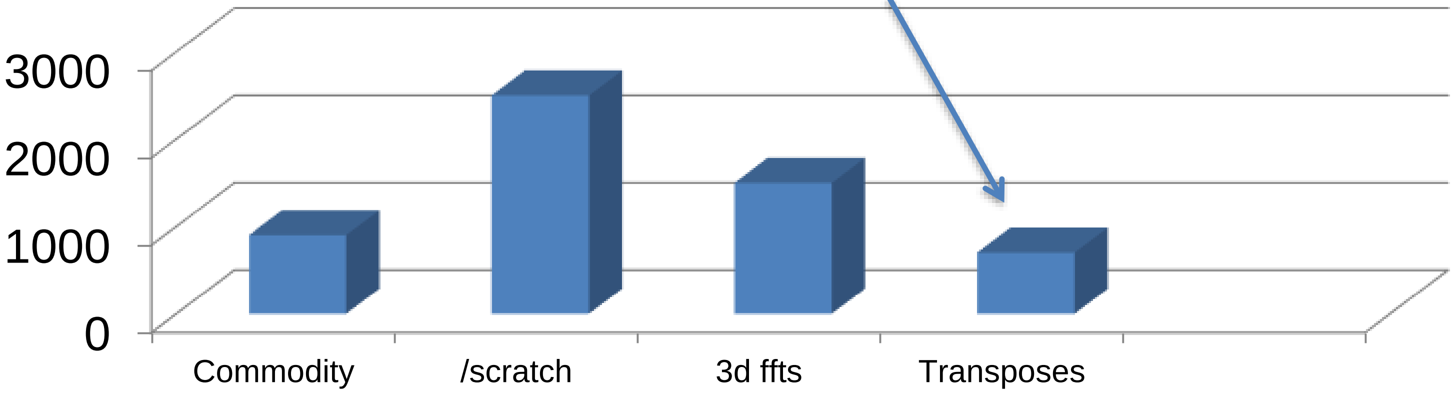

RTM imaging – fun facts

Summed over the full output space of 11,847km^2 (red polygon)…

…6 azimuths x 11,847km^2 = 71,082km^2

46 angles x 71,082km^2 = 3,269,772km^2

with a trace spacing of 30m x 30m…

…or for each surface location there is 276 angle traces i.e. images

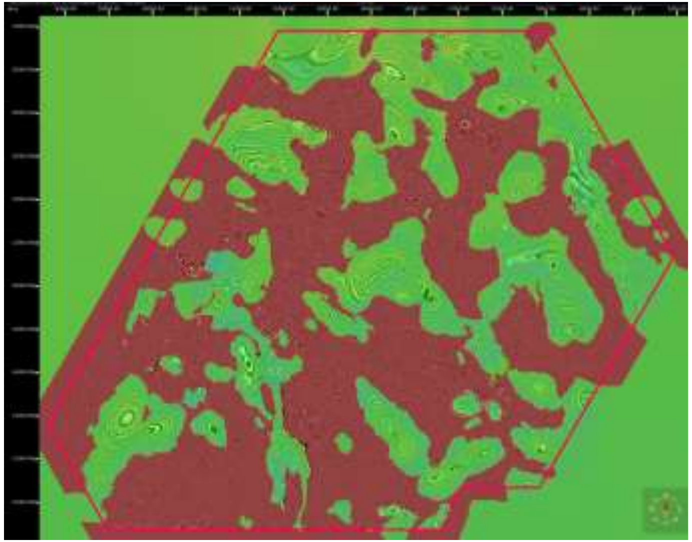

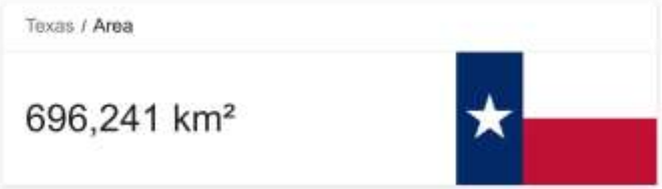

Or stated differently: a 3D image with a pixel size of 30x30x10 meters, that covers an area about 5 times the size of Texas, to a depth of 16km.

So why? The sub-salt imaging challenge

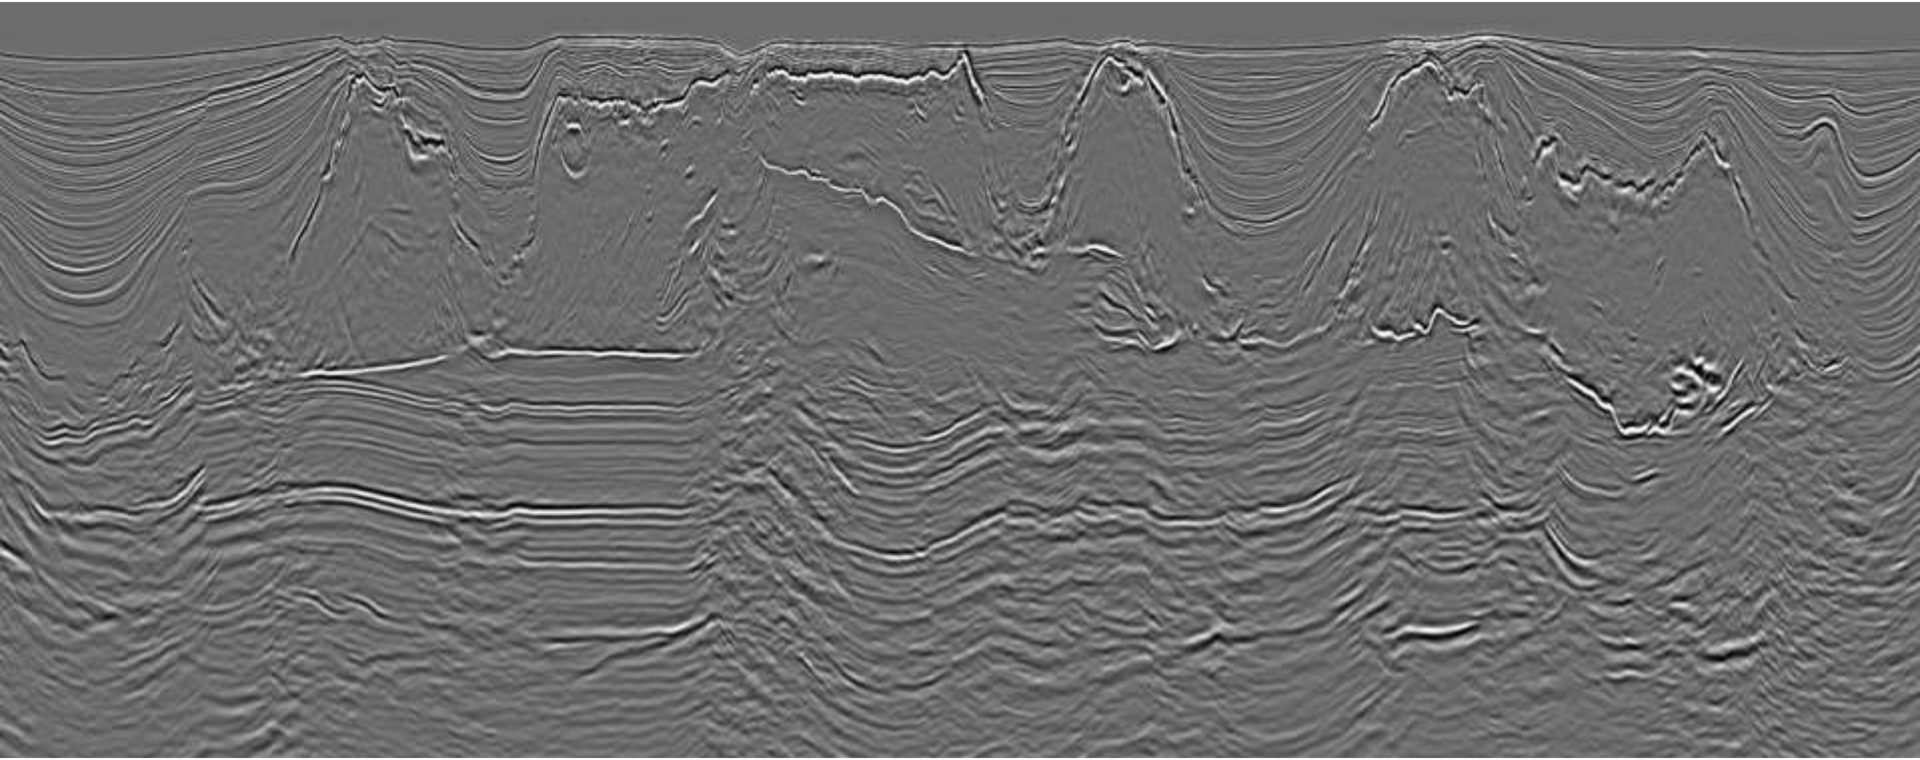

Leveraging a supercomputer in 3D image optimization

The 5D output data volume is about 33TB, we now want to analyze and edit this to improve the image further i.e. post processing, and generate a final 3D image

- When imaging beneath complex geology, it is common that the signal of interest is confined to just a few partial images (angle-range and azimuth).

- Selective stacking optimizes the signal by combining the gathers that best illuminate the target while discarding noisy parts that would deteriorate the final image.

- The aim is to improve the Signal/Noise in the final image

Azimuth Sectored
Angle gathers

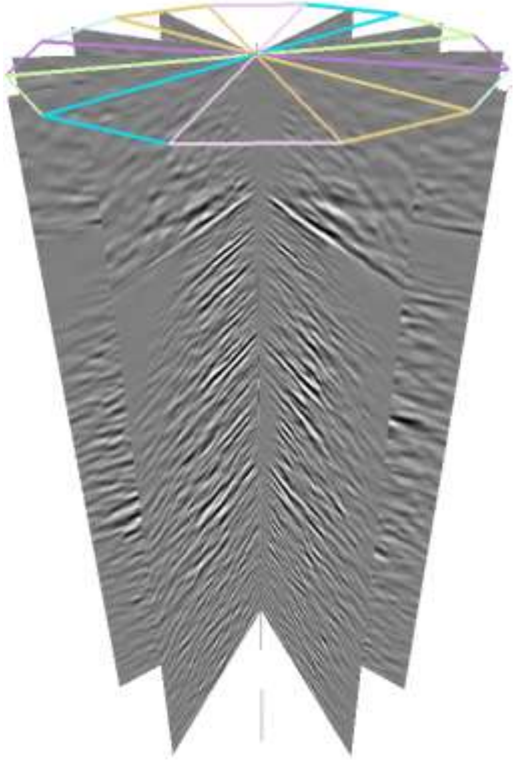

276 angle traces in a 5D hypercube giving a different view of the geology

Azimuth Sectored
Angle gathers

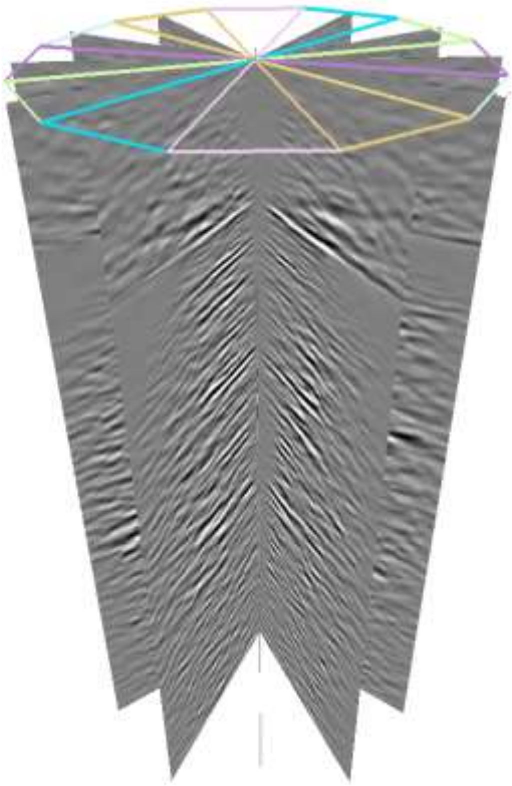

Butterfly Angle gathers

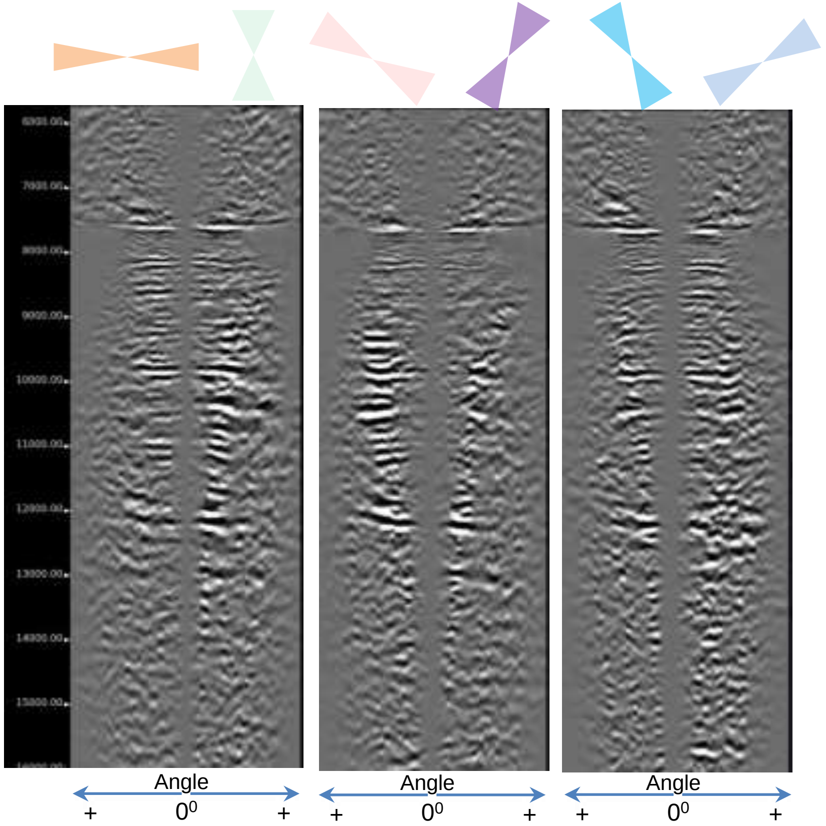

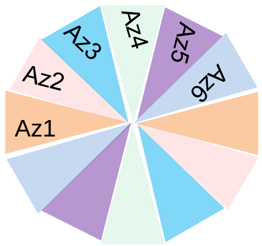

Azimuth sectored stack sections

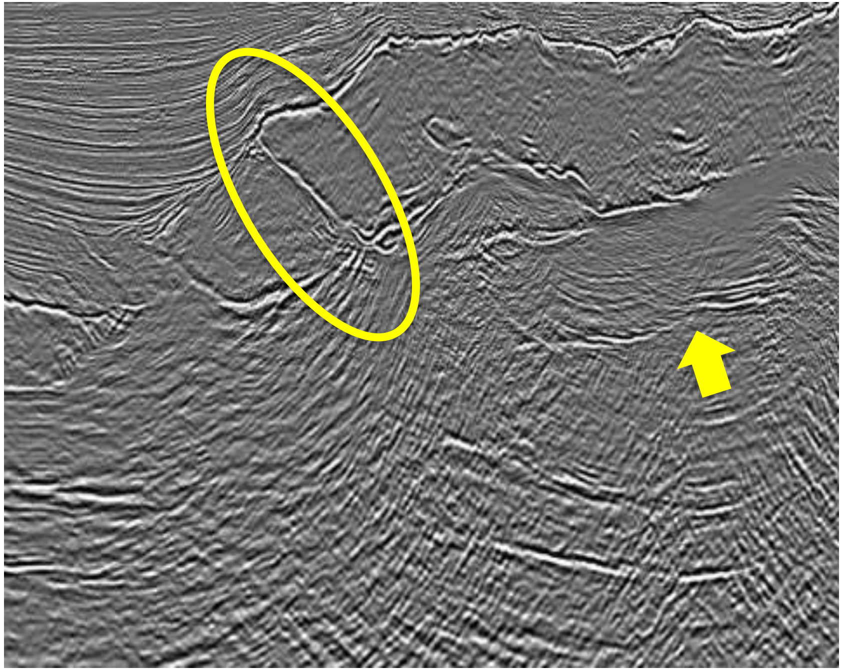

AZ 5

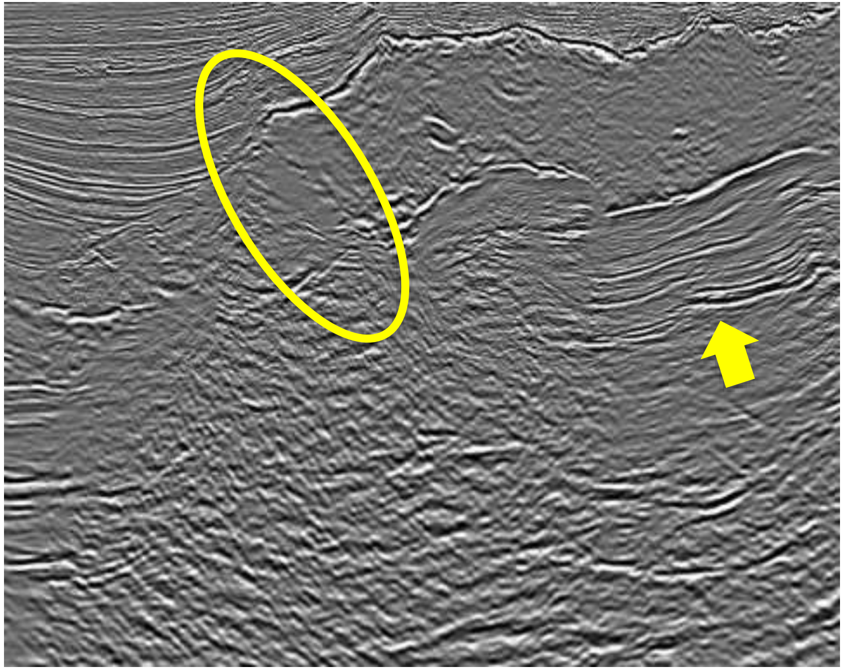

AZ 2

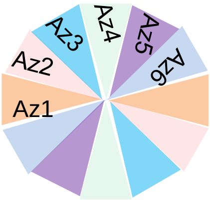

Angle domain is great for combining data (5D binning)

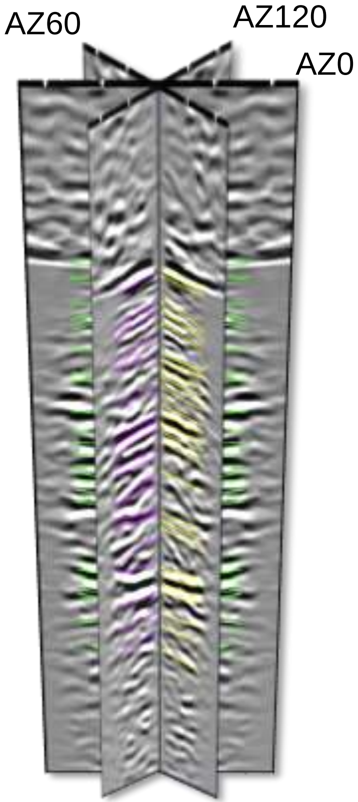

- Utilizing full azimuth RMO picks from near and far offsets.
- Enhanced illumination from offsets up to 16km.

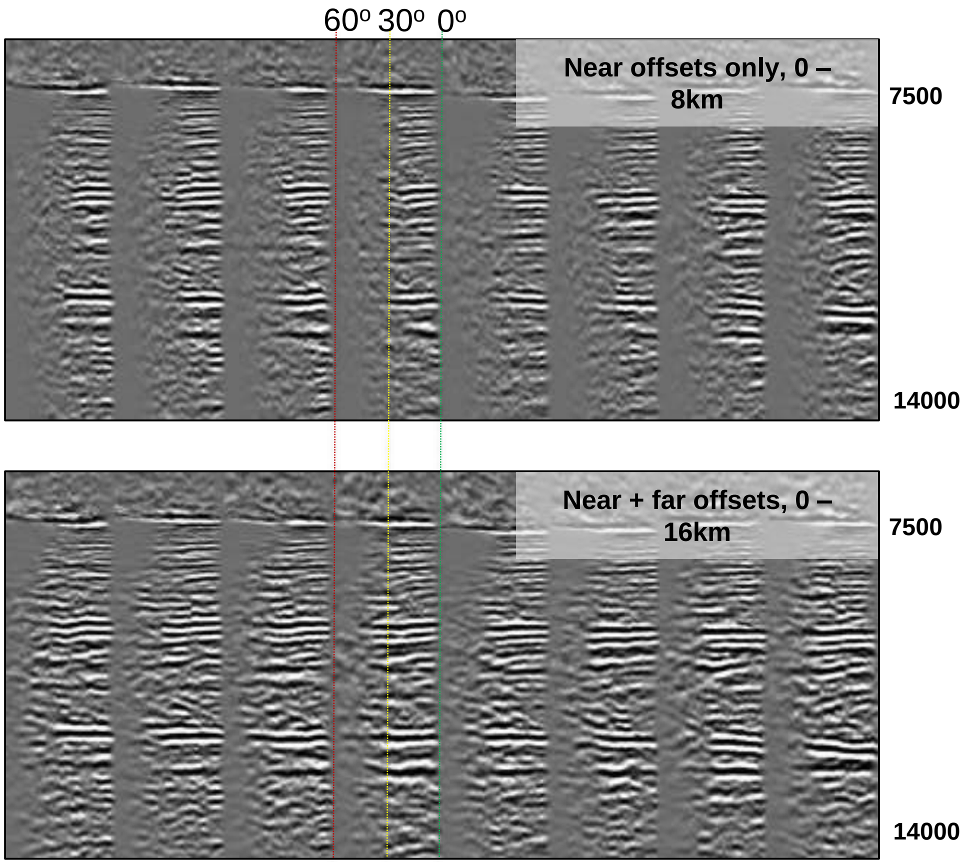

A different perspective on the RTM computational domain

Images are formed for individual supershots, each shot is processed inside a computation box with a surface area of 576km^2.

The survey contains >100,000 supershots

On the other hand, only 18 such computational boxes are needed to cover the complete survey area

…and the size of a single image pixel

So more than 100,000 images were in reality formed, each with a computational domain of 24x24x16km, and with a cell size of 30x30x10m.

Future: treat the full survey area (100km x 100km) as as one domain, rather than 100,000 separate (overlapping) experiments

Outline

- Why PGS did computers a bit differently
- Grow behind our fleet, and keep up with a growing data flow
- Triton, imaging a 660TB dataset using distributed memory and fast interconnect
- New geophysical science on a Cray XC40 supercomputer
  - WEI: Wave Equation Inversion
- Future challenges; software, people and machines

New geophysical science on a Cray XC40 supercomputer

RTM imaging conditions: leveraging shared memory and fast interconnect to implement new geophysics

- Move away from the classic cross-correlation imaging condition
- Better image conditions = more data available during imaging step
- Using the full wavefield in imaging, not just a synthesized source (snap shots), including very long, continuous records with energy from multiple sources
- This field is in constant development, it is key to move fast from concept to implementation

Classic acquisition scheme – individual shots

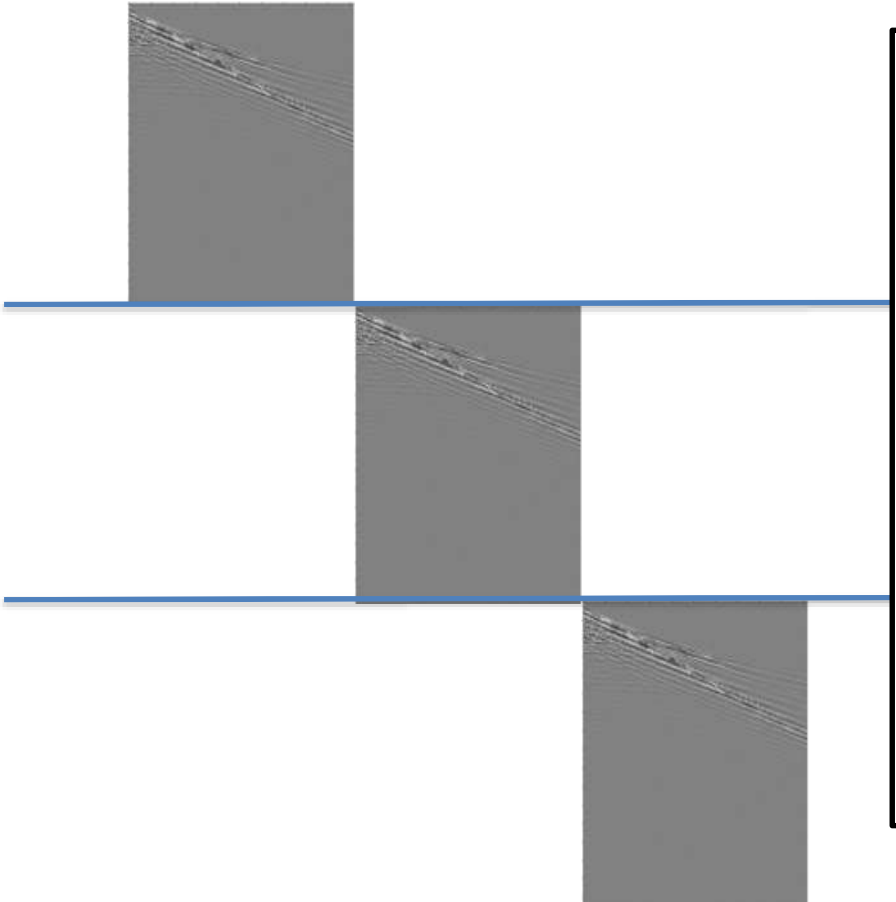

This data structure lends itself to (embarrassingly) parallel processing on clusters, one shot per machine,

then accumulate individual images to form final image. Data reduction.

Emerging – continuous shooting and recording, irregular sampling

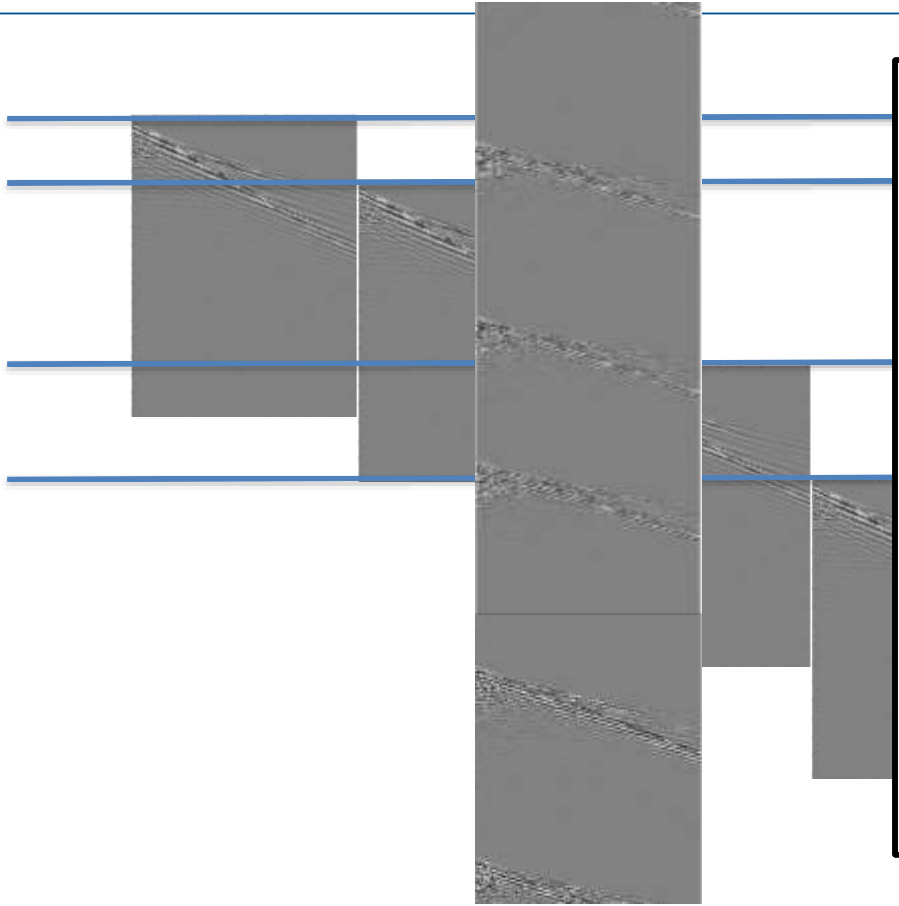

This data structure does not lend itself to individual shot processing.
New approach, massive data volumes, very large linear systems

Map data into image space, no/little data reduction, treat survey as one ensamble

Migration = Imaging with an adjoint operator

$$\mathbf{d}_{obs} = \mathbf{L}\mathbf{m}$$

$\mathbf{d}_{obs}$	recorded seismic data
$\mathbf{L}$	modeling operator
$\mathbf{m}$	reflectivity

Migration (RTM) = Imaging with an adjoint operator

$$\mathbf{m}_{mig} = \mathbf{L}'\mathbf{d}_{obs}$$

$\mathbf{d}_{obs}$ recorded seismic data

$\mathbf{L}'$ migration operator

$\mathbf{m}_{mig}$ migrated image

Least Squares Imaging - Wave Equation Inversion (WEI)

$$\mathbf{d} = \mathbf{L}\mathbf{m}$$

$$\hat{\mathbf{m}} = \left(\mathbf{L}'\mathbf{L}\right)^{-1}\mathbf{L}'\mathbf{d}_{obs}$$

Massive matrix, dimension is Data x Data

WEI: fast solvers for sparse linear systems

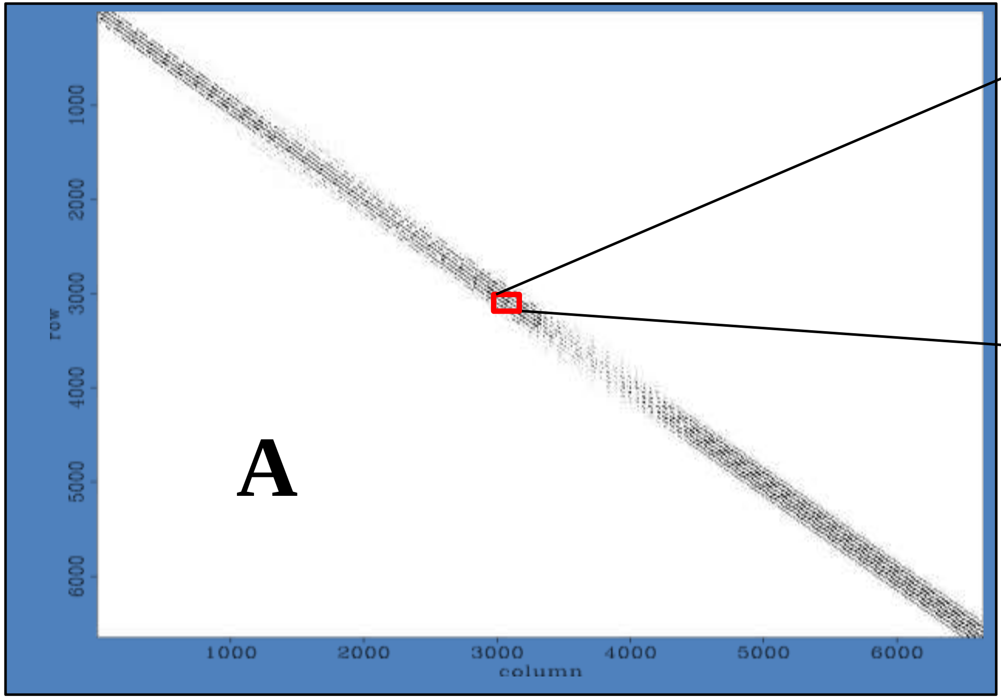

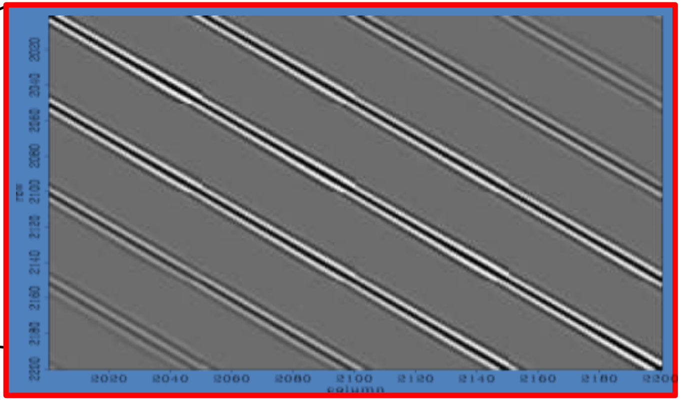

$$\hat{\mathbf{m}} = \left(\mathbf{L}'\mathbf{L}\right)^{-1}\mathbf{L}'\mathbf{d}_{obs}$$

$$\mathbf{Ax} = \mathbf{b}$$

WEI: fast solvers for sparse linear systems

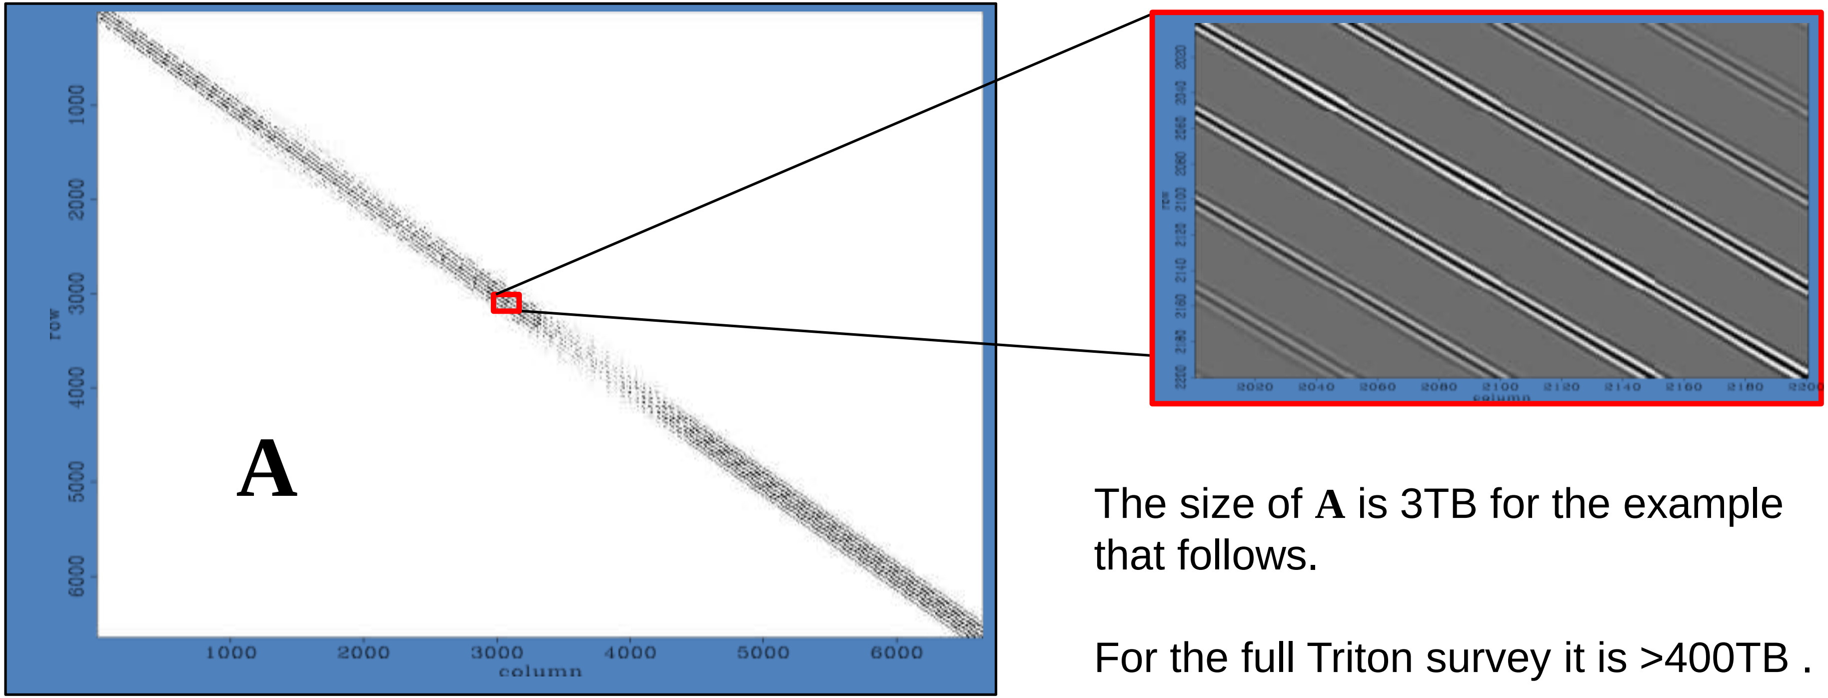

The size of **A** is 3TB for the example that follows.

For the full Triton survey it is >400TB .

WEI: fast solvers for sparse linear systems

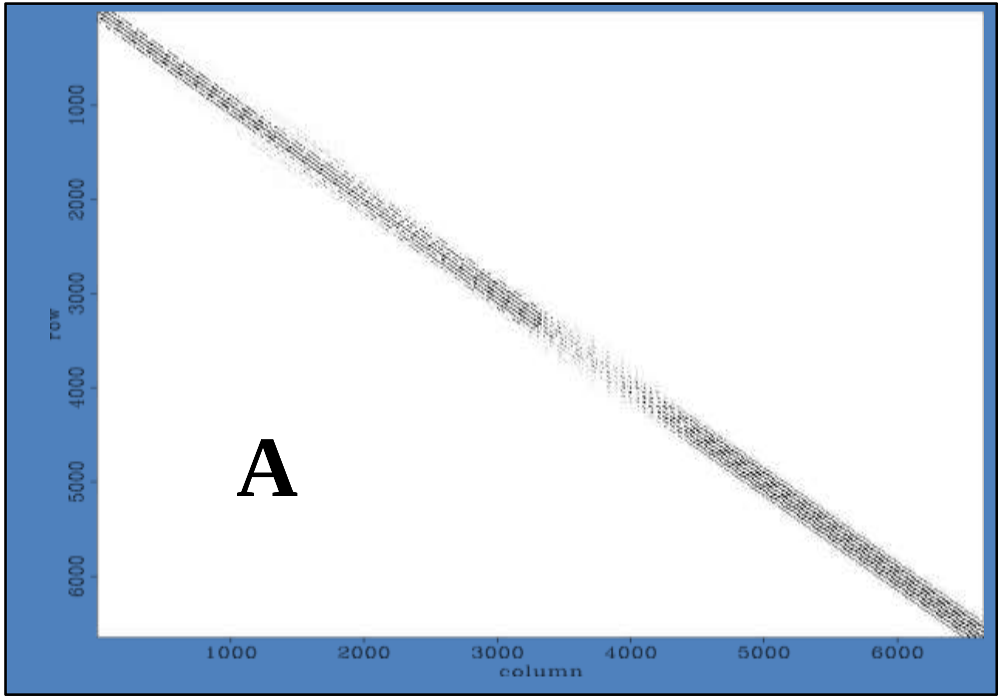

$$\mathbf{Ax} = \mathbf{b}$$

Need to solve an sparse linear system with an scalable (parallel) solver

A fast, low-latency interconnect is needed to obtain good performance

WEM

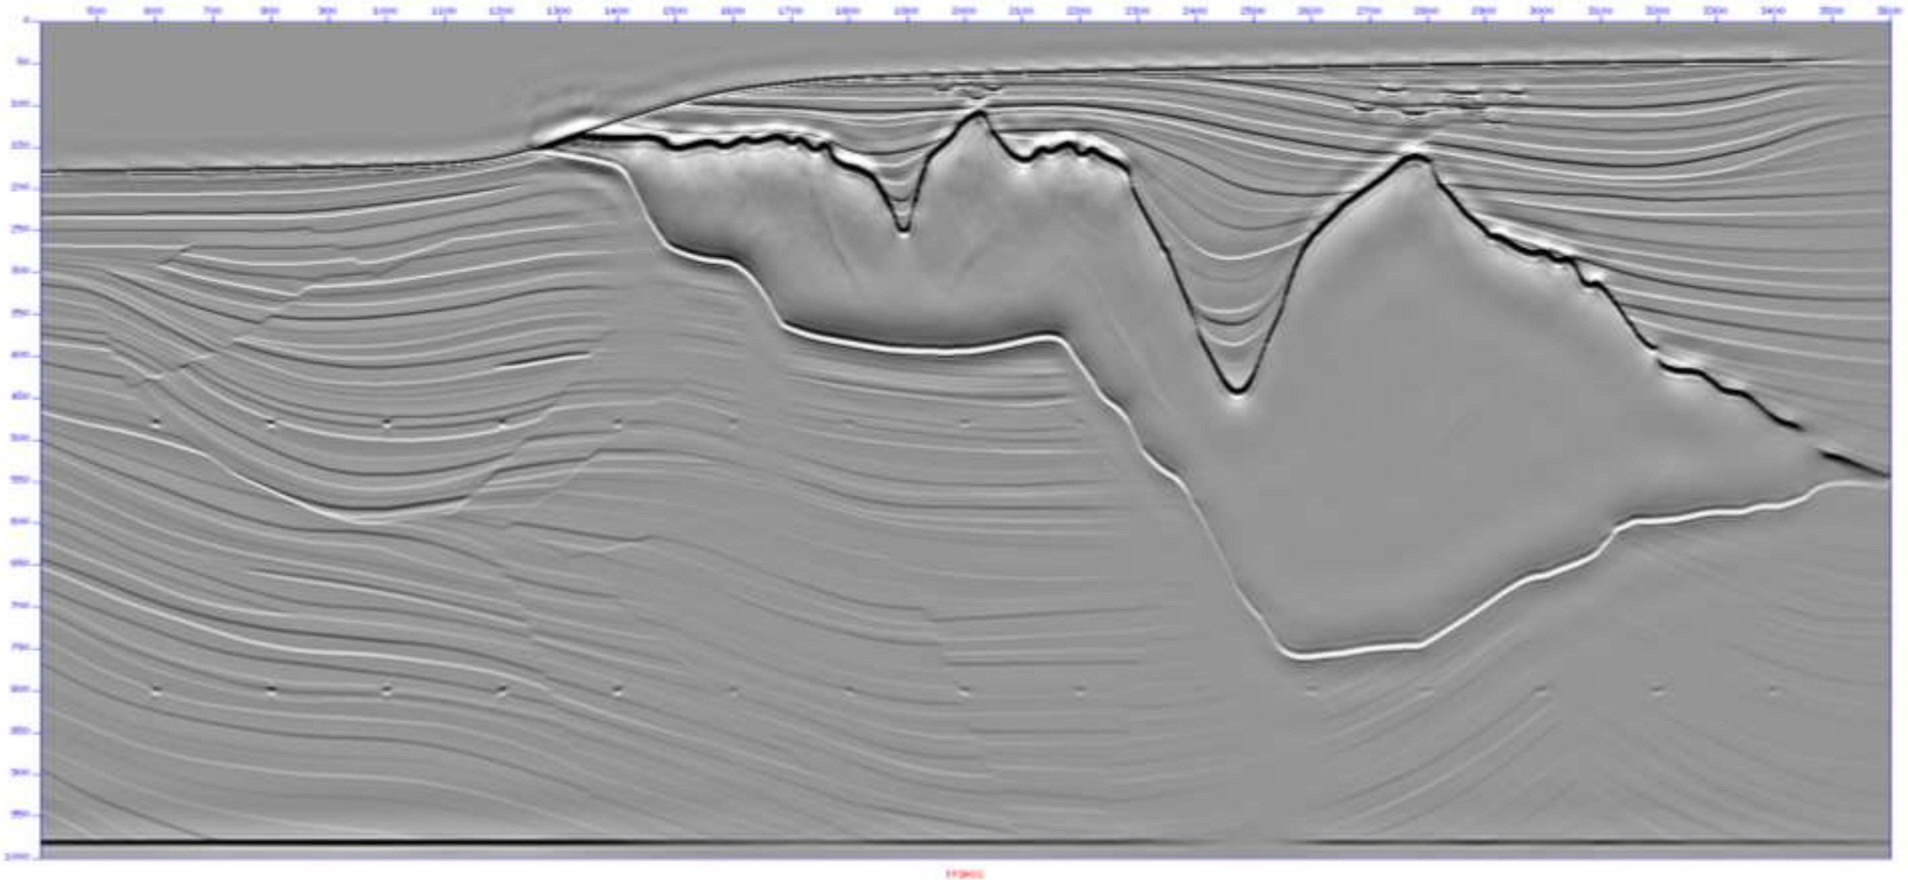

WEI

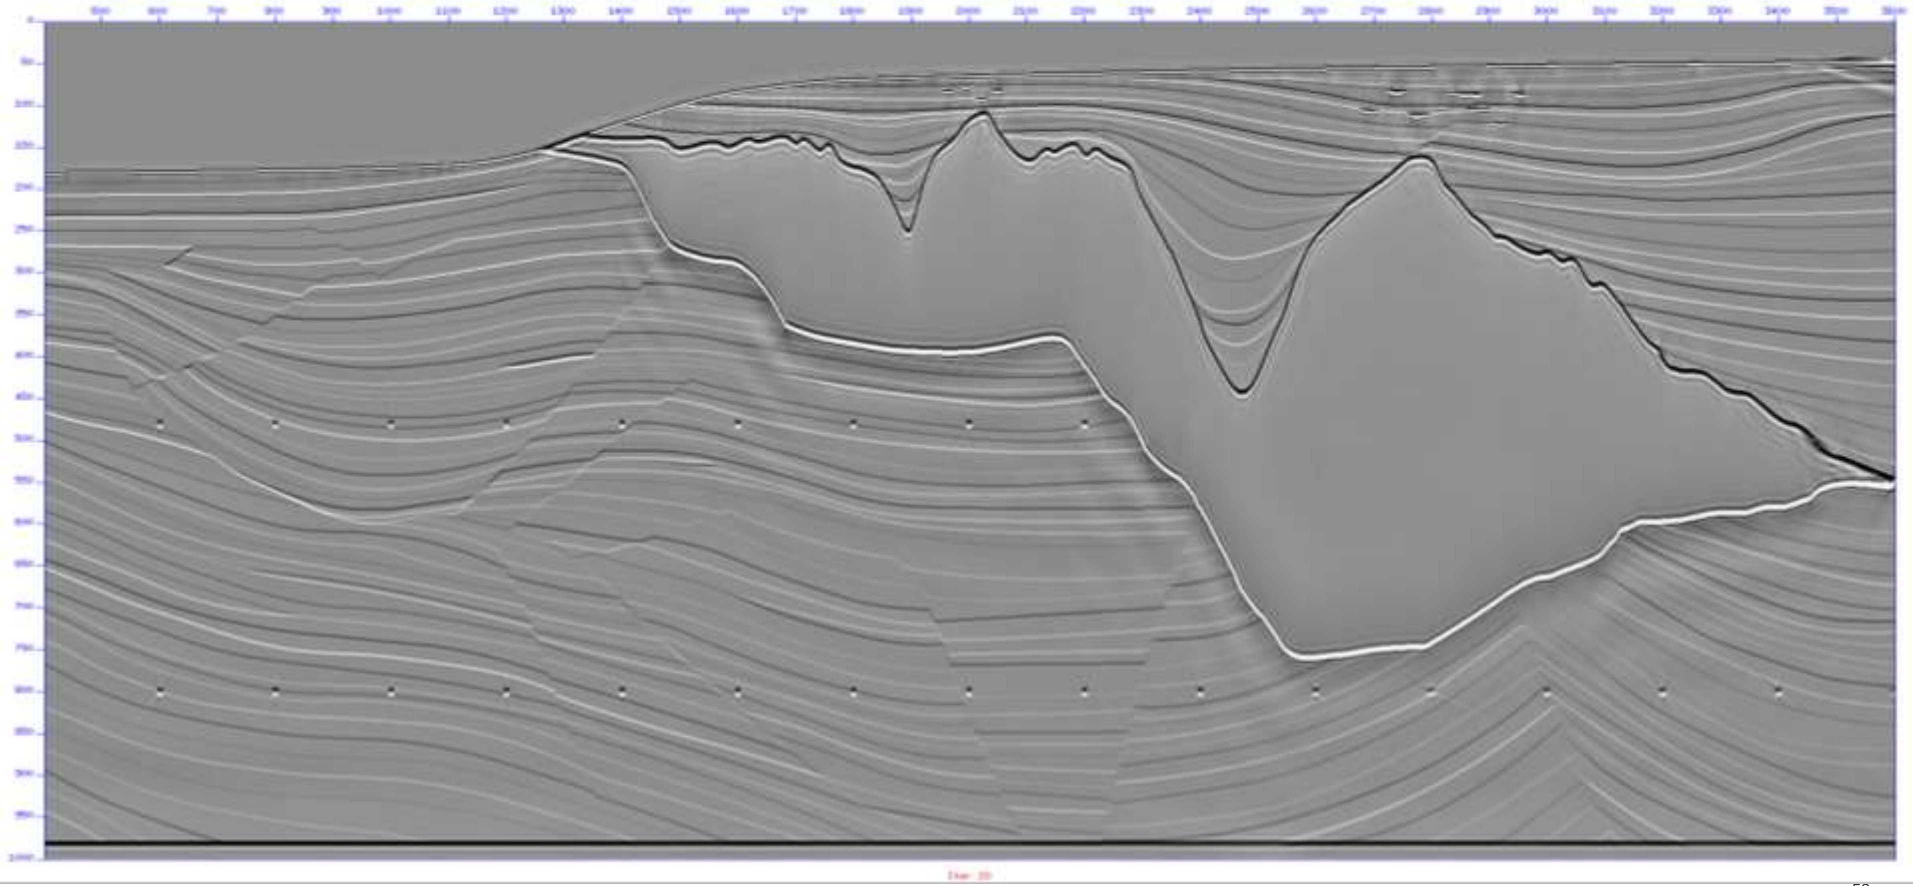

WEI

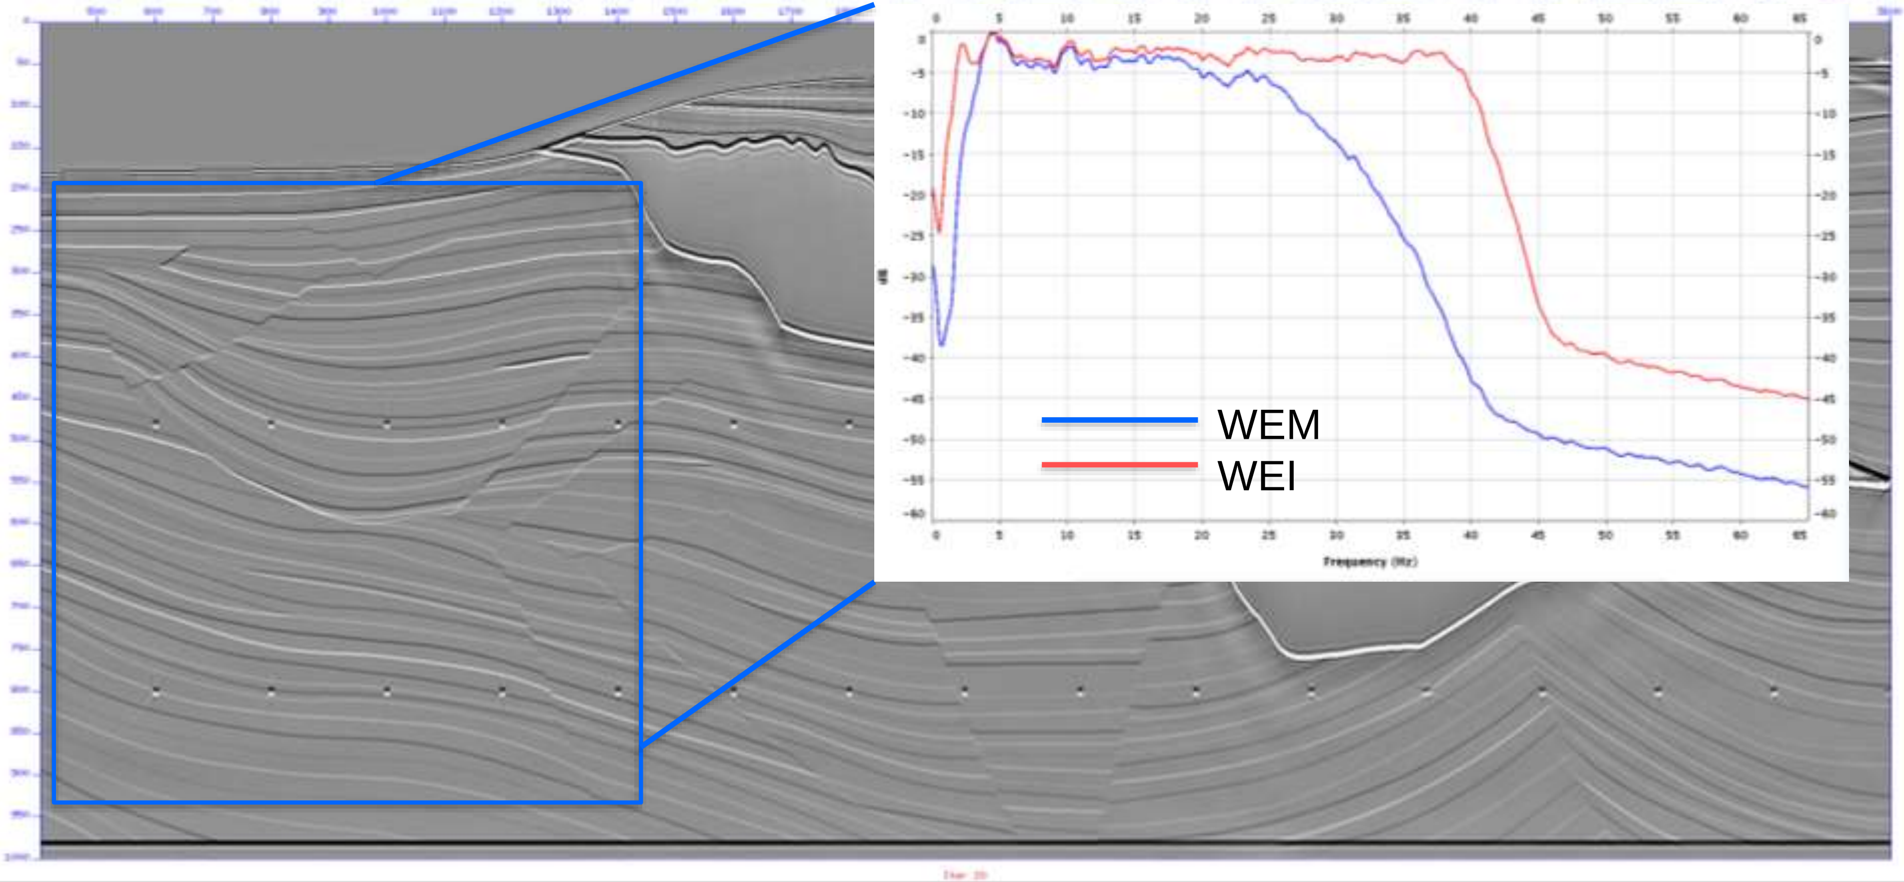

Depth Slices WEM (35Hz)

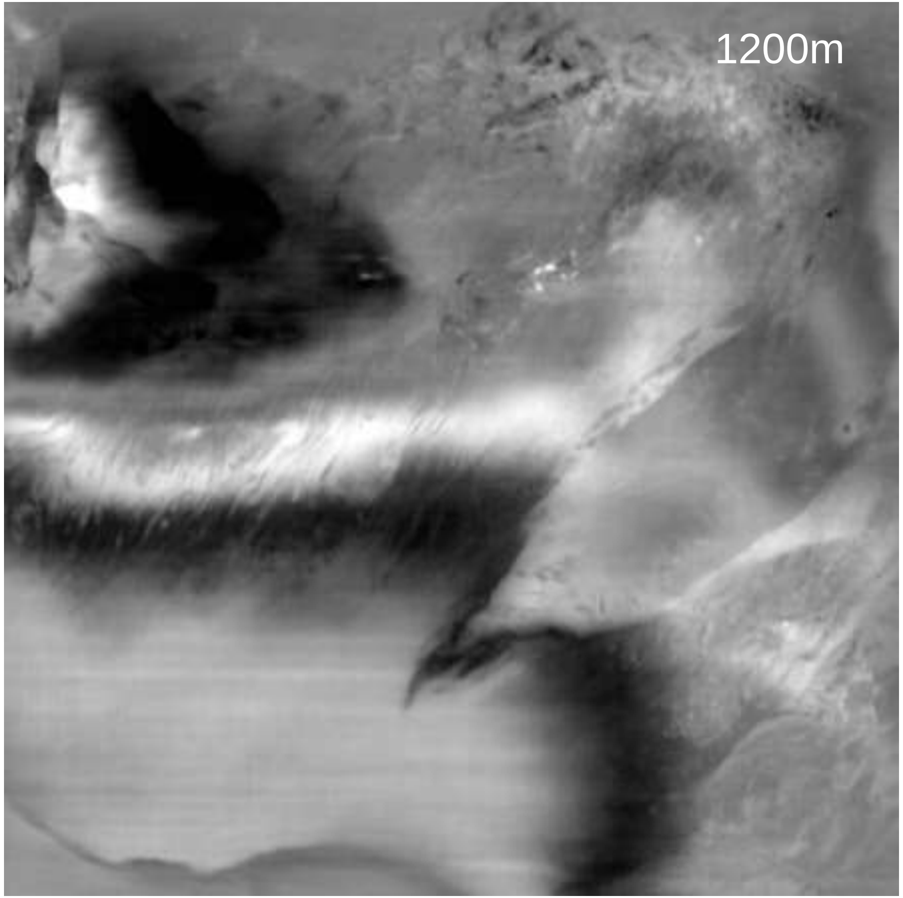

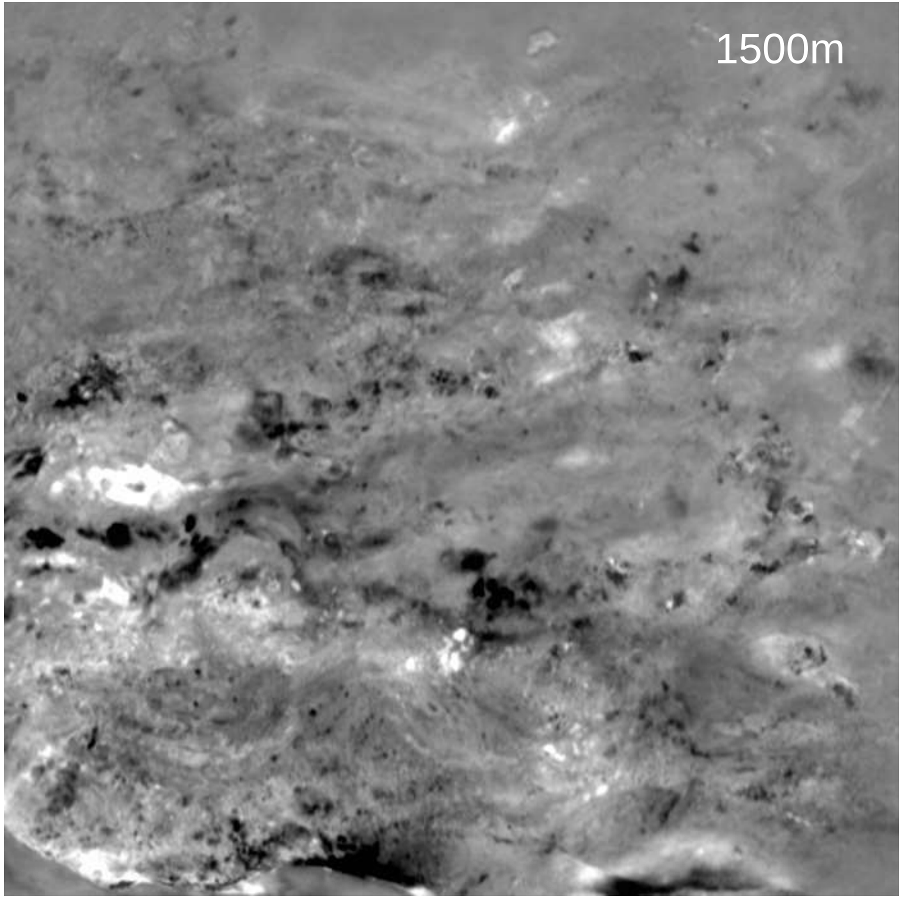

Depth Slices WEI (35Hz)

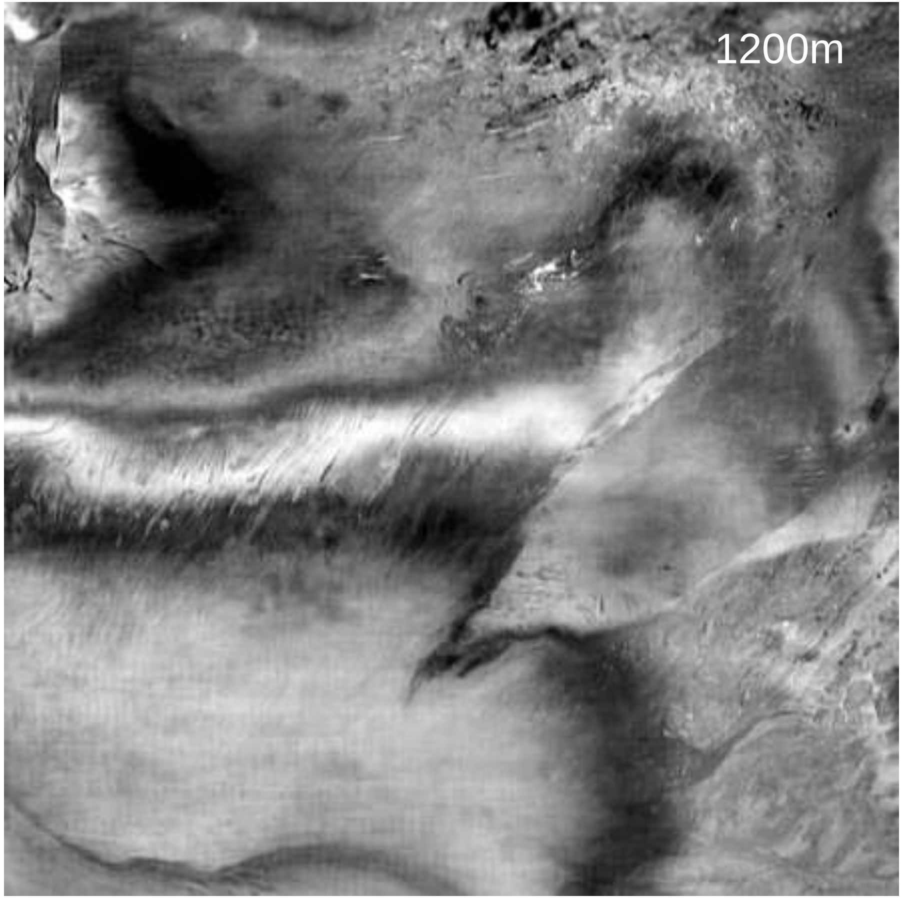

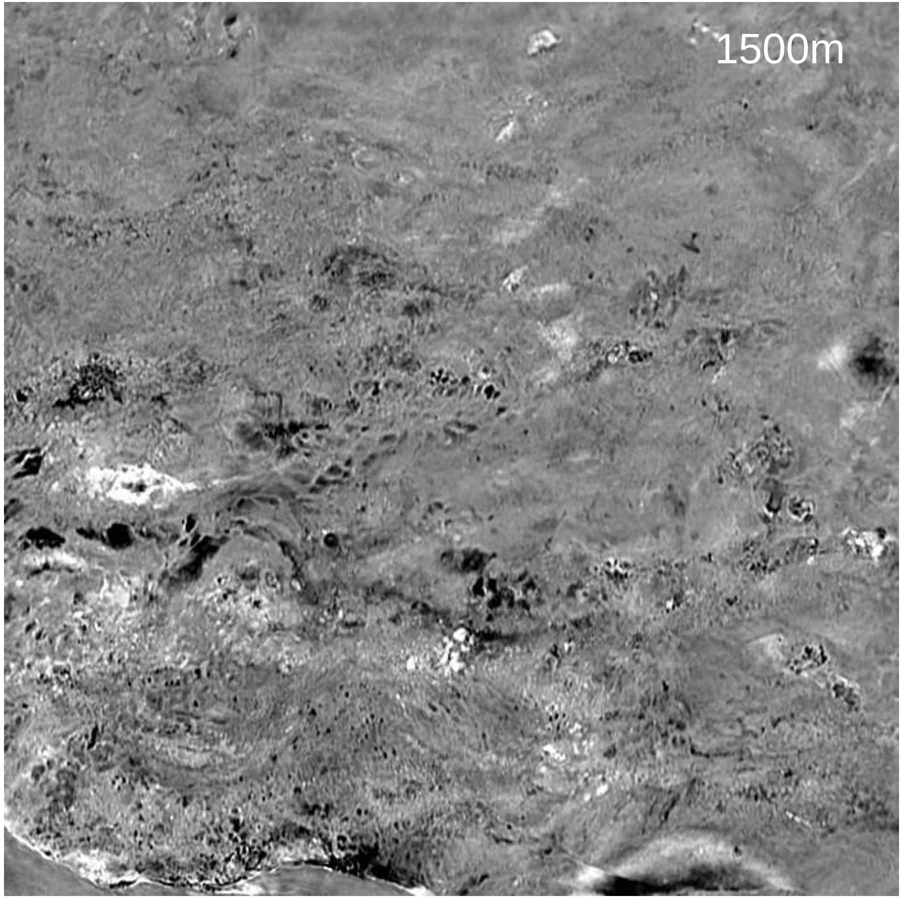

Illumination below salt and Wave Equation Inversion

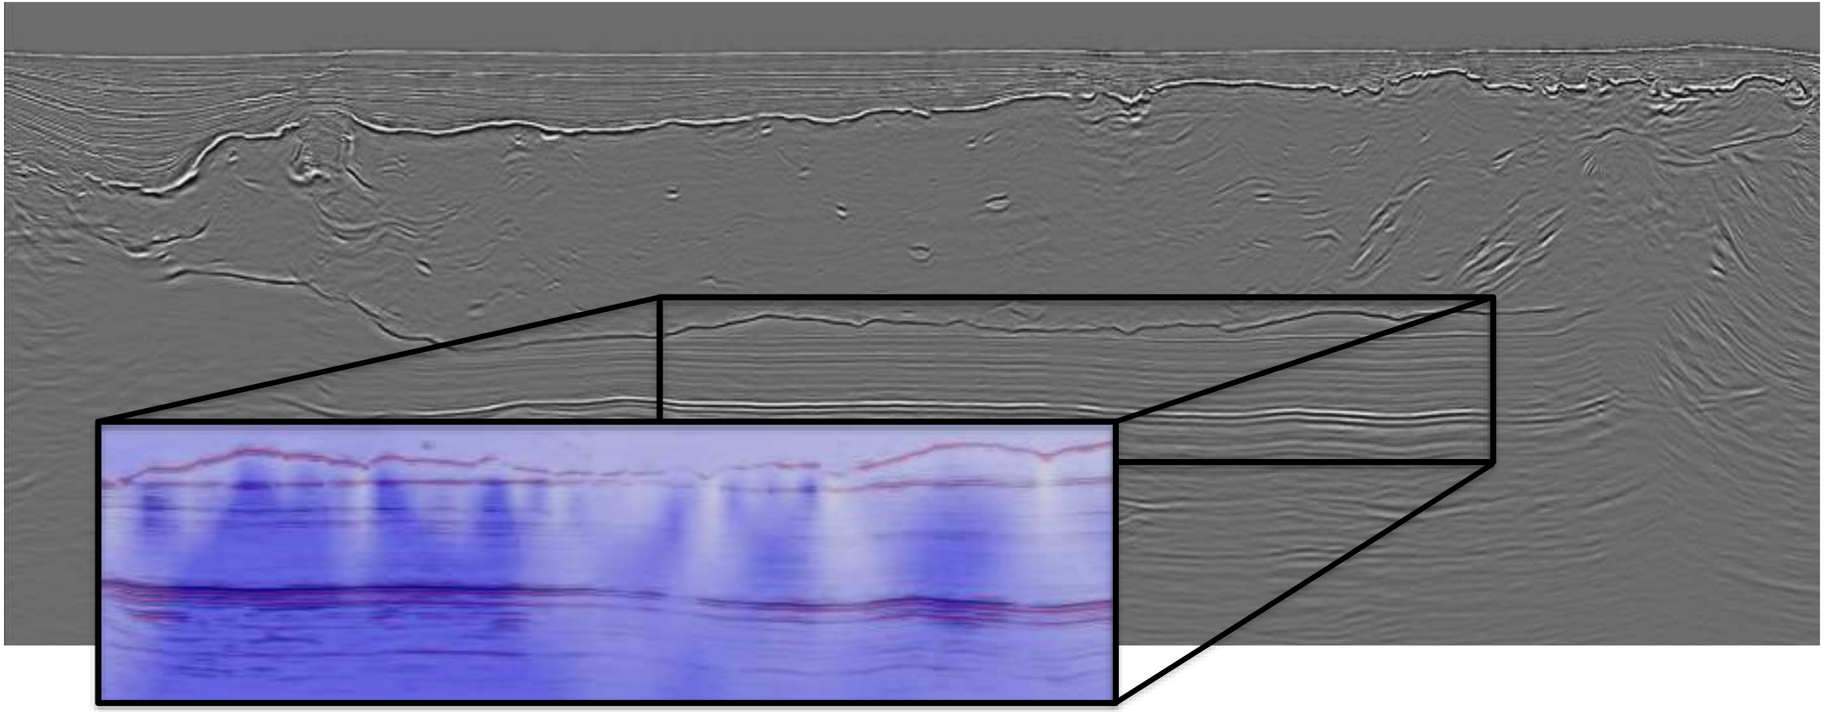

WEI applied in a small target area of 375 km^2

z	INLINE	XLINE
3 km	25 km	15 km

Amplitudes extracted along target horizon

Migration result

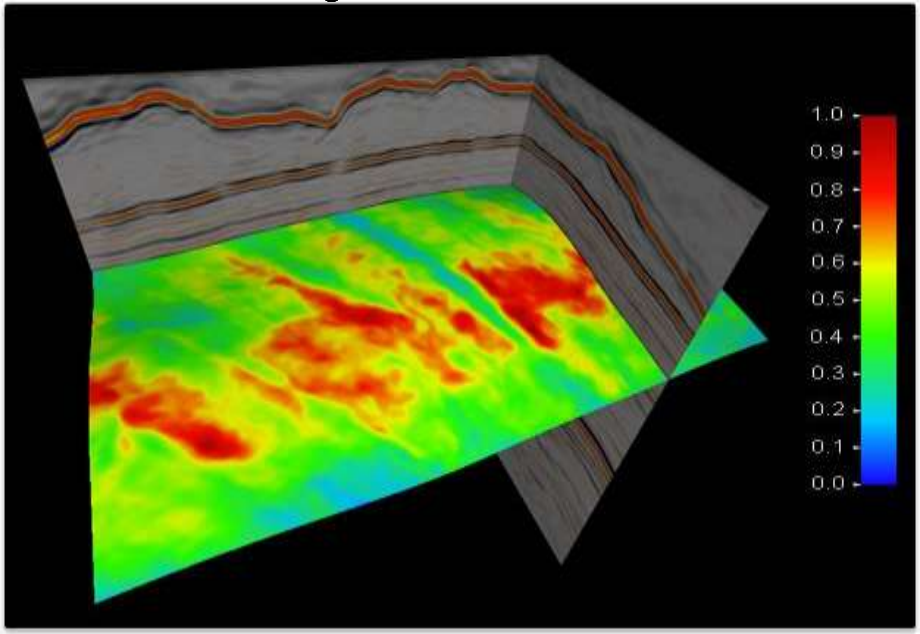

WEI result

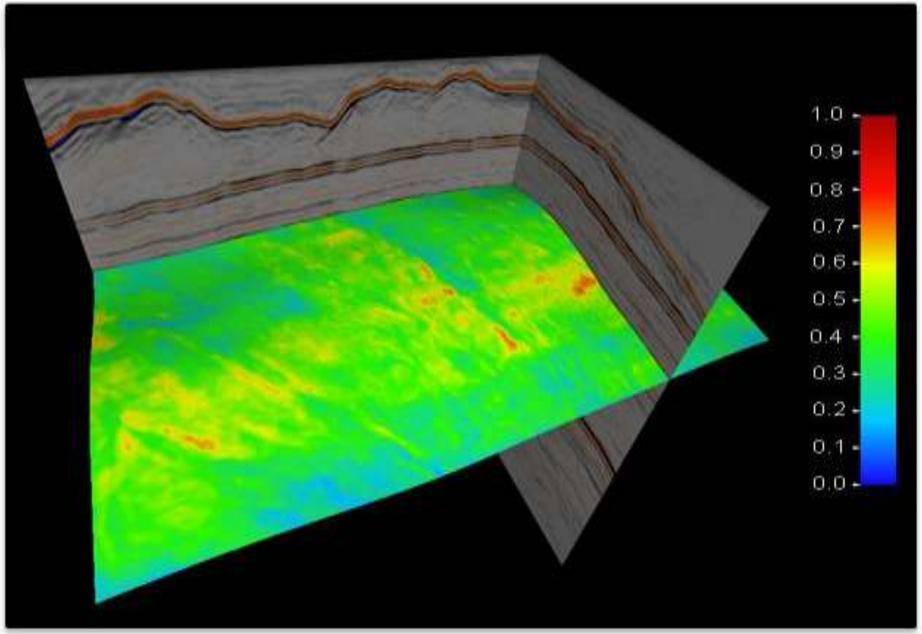

Some thoughts on what's to come

- Relentless focus on turnaround and cost
  - price-/performance is key, will pursue anything that provides improvements
  - Balanced against speed of development: how long to bring new physics to market

- Bigger problem size
  - More and more data, multi-component, continuous recording, density of sensors
  - Non-uniform data: data with less inherit parallelism
  - Inversion… very large linear systems

- Rapidly move from concept to large-scale deployment: "From Matlab to 200,000 cores in a week..."
  - Small nimble teams, stable hardware
  - Programming model is a challenge

Conclusions

- How our size and strategy for growth pushed us to look beyond typical solutions used by peers

- The challenge of adapting new HPC hardware while running a small margin service business
  - Organizational capabilities: commodity vs. integrated solution
  - RTM ported to Cray XC40

- New algorithms brought to market thanks to
  - Large memory pool
  - Fast interconnect
  - Acceptable programming model and “uniform” hardware

- Ensemble based processing, moving from single trace to full survey, then be able to dynamically focus on the parts that are of interest (anomalies)
  - Challenges the software environment and programming model
  - Will pose a challenge also for hardware