

We can make your
ProMAX[®]
run
10 times faster

with our brand-new

SeisJet Seismic Data Server software solution.

The **SeisJet Seismic Data Server** software solution drastically speeds up conventional processing flows executed by a seismic processing system (currently ProMAX[®], but integration with other processing systems is possible). The effect is especially prominent in case of processing of large volumes of data on cluster architectures. Benchmark tests show approximately 10 times speed up of typical processing batches on the same hardware.

The **SeisJet Seismic Data Server** software operates with native processing system datasets and its output is easily accessible by standard I/O facilities of the processing system.

How does it work?

The two keys to the significant speed up of the processing are the following:

- Quick data read and quick sorting on read.
- Optimized data distribution to parallelized flows.

Quick data read and quick sorting:


When the speed of sorting of the seismic data is recognized as a bottleneck in processing time, the typical solution is to buy more expensive disk array system. To double the performance, you normally need several times more expensive RAID array. Using the **SeisJet Seismic Data Server** software with the same hardware gives you 10 times the performance gain. Why? Because it reads the data in an optimum way, properly utilizing big amounts (gigabytes) of RAM available on modern computers.

It is well known that reading the data from hard discs (RAID array) in an arbitrary order (so-called *random access* to the data) is much slower than if the data is read sequentially, in the same order as it is stored (so-called *sequential access*). However, the data can be accessed in any order very rapidly when loaded into RAM (which stands for Random-Access Memory).

For this reason, the key to significant acceleration of the data input is (1) reading as much data as possible from disk into RAM in its original order, and (2) then making random-access operations with the data (e.g. sorting) in RAM.

Of course, the access to the data on disk cannot be made truly sequential if there is a need to resort big data volume on input, unless the whole data fits into RAM. However, it is possible to maximize the amount of data that is read from the disk sequentially by analysing the original order of the data on disk and the required order on input of the processing flow. The **SeisJet Seismic Data Server** builds optimized strategy of data reading and, as a result, gets the data from the disk much more sequentially than it is typically done when straightforward seismic data input approaches are used.

Optimized data distribution to parallelized flows:

Straightforward approaches to data distribution between several parallelized copies of a flow are poorly scaled. It means that after a certain (typically small) number of nodes/processes 

executing the flow in parallel, further increase of the number of nodes/processes does not speed up the processing any more: though all the computing resources (hard disks, network, processors) are not fully loaded, the execution is still slow. Typical solution is to buy expensive hardware: low latency network over expensive disk arrays. This leads to some moderate performance gain but it is not proportional to the expenses. The **SeisJet Seismic Data Server** completely solves the network latency problems by means of optimized software design.

In fact, the main reasons for the poor scaling of the parallel processing typically can be as following:

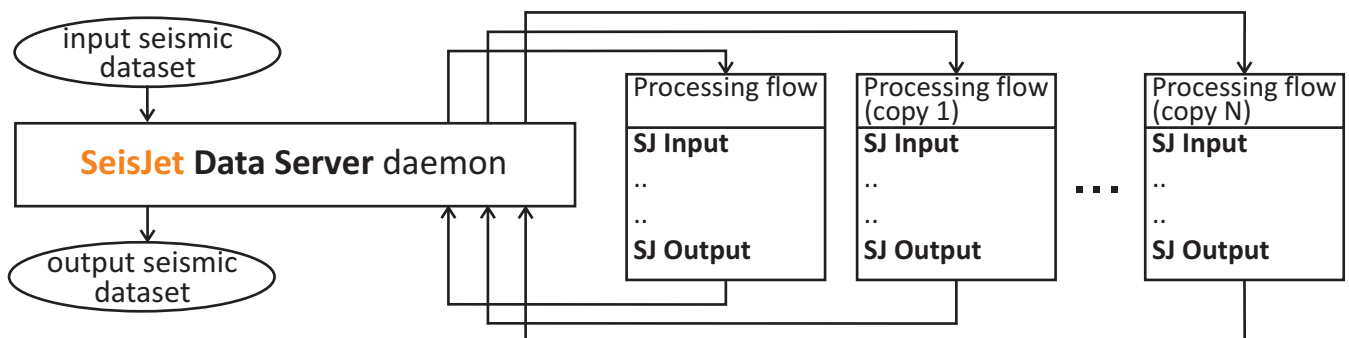
- While flow copies are executed in parallel, data reading and distribution between the copies is still made within one and the same sequence. When the processing becomes really fast, each copy of the flow manages to complete with one portion of the data earlier than the following portion is ready for its input. As a result, most of the time the parallelized batch copies are just waiting for the data and doing nothing.
- Generally, when the data is read in parallel by several processes/nodes from one and the same shared storage, the randomness of the data access increases and the whole operation slows down. When access to the shared storage is arranged via general-purpose system solutions (e.g. NFS, cluster FS, etc.) they additionally introduce overheads, as compared to direct data transfer from a data server to clients.

Obviously, this type of problems can be solved by means of software. The **SeisJet Seismic Data Server** reads the data only once and takes care of truly parallel data distribution between different nodes as quick as possible.

How does it look like?

The **SeisJet Seismic Data Server** solution includes the following components:

- **Data Server** daemon.
- Input/Output modules (**SJ Input** and **SJ Output**) integrated into the processing system.



The **Data Server** is a standalone program (daemon) that works solely with the seismic datasets of the processing system. It is capable of quick seismic data reading and sorting, and transfers the data to a number of copies of a processing flow in a high performance manner.

The data from the **Data Server** is introduced to the processing system batch through the input module (**SJ Input**) that comes along with the **Data Server**. The module initializes a new job in the **Data Server** and feeds the processing flow with the data. To provide a parallelization of the processing flow, the **SJ Input** module can start a specified number of copies of the same processing batch ("secondary" flows) on the different nodes of the computational cluster (or on the same node for a SMP machine). Each copy of the **SJ Input** module in the secondary flows will also subscribe to the **Data Server** and will feed each copy of the flow with the data. The output of the flows are transferred back to the **Data Server** by the **SJ Output** module. The **Data Server** cares of parallel processing of the requests and consistent data distribution between the flows, as well as collects the result of the processing to an output dataset.

The **SeisJet Seismic Data Server** solution does not interfere with the processing flow internals: processing modules are executed within the flows just as if the standard input and output were used. Lots of processing flows executing the same job are running in parallel just as independent flows without interfering with each other and thus achieving maximum scalability.

ProMAX® is registered trademark of Landmark Graphics Corporation.