

Accelerators Overview

HampsonRussell accelerators are designed to significantly improve workflow efficiency by reducing processing time and optimizing the utilization of computing resources. HampsonRussell software suite accelerators include:

Multi-Node Processing (MNP): MNP enables the execution of complex processes across single or multiple machines seamlessly, whether on-premises or in the cloud. This distributed computing capability ensures rapid results and efficient resource allocation.

Batch Processing: Our Batch Processing feature offers the convenience of portable batch scheduling, allowing you to efficiently schedule tasks to optimize machine usage.

Multi-threading and GPU support: We also provide the option to harness an additional 12 threads, further enhancing parallel processing capabilities to tackle resource-intensive tasks. GPU support is available in GeoAI and Strata.

Multi-Node Processing (MNP)

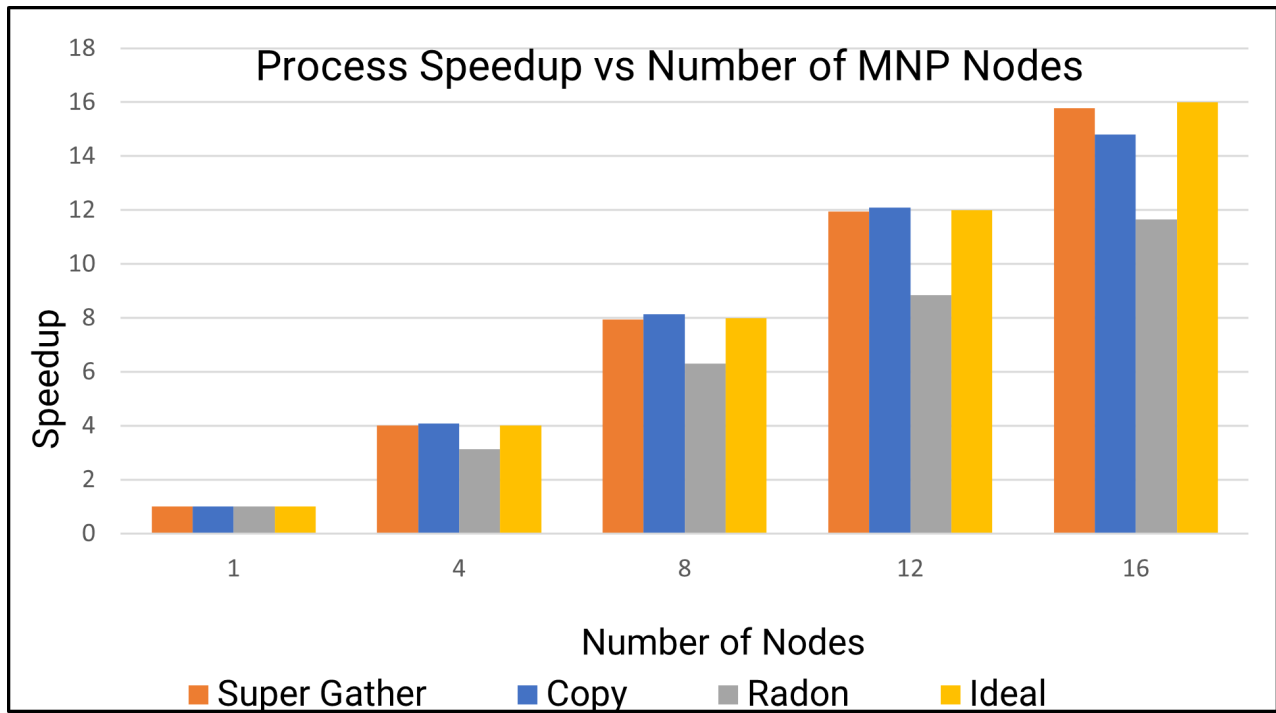
MNP is a resource framework that handles computationally intensive operations. This adaptable framework can be deployed across various network configurations, spanning from on-premises clusters to remote cloud environments such as AWS or Azure.

When employing the MNP methodology, a master node effectively delegates a task to numerous worker nodes by partitioning the data into smaller subsets. MNP has the capacity to engage up to 64 nodes and additionally supports multi-threading and batch scheduling.

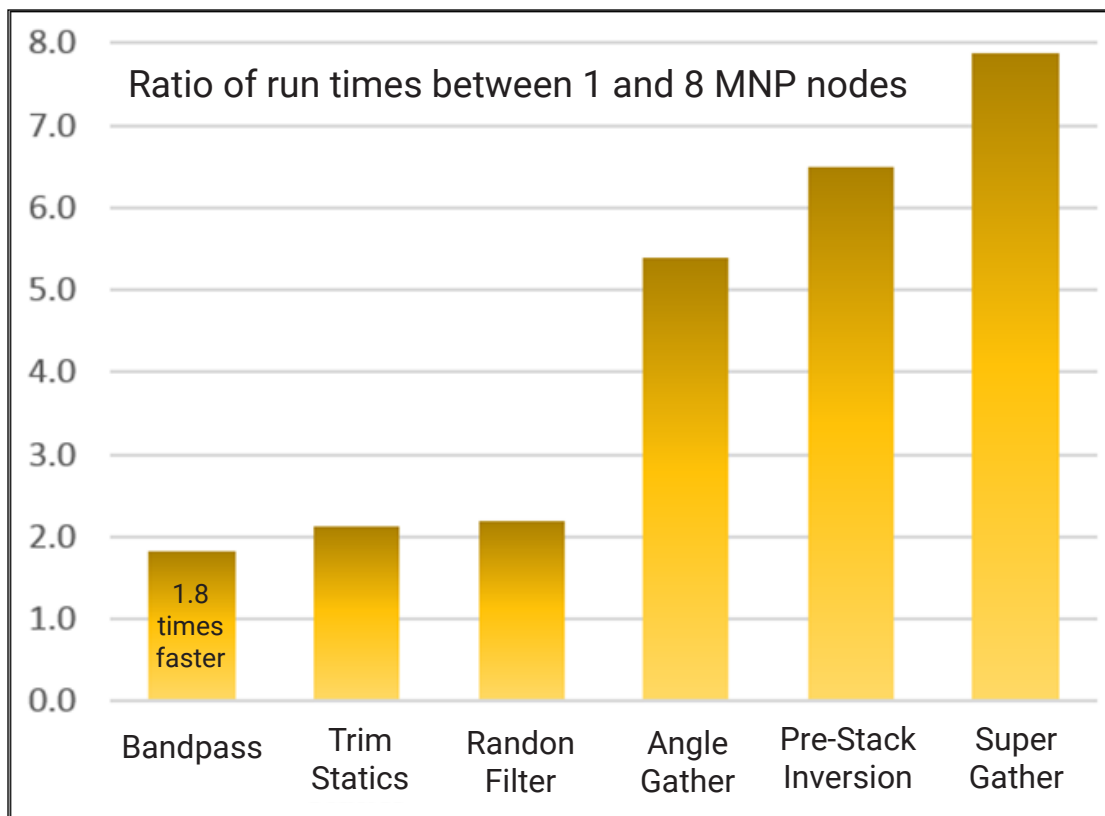
MNP Cluster: Network and Cloud

MNP provides a significant processing time uplift for CPU-intensive processes such as Pre-stack Inversion, Inverse Q, Radon Filter, and Emerge Neural Network application as compared to the performance of a single machine.

MNP also excels in accelerating I/O-bound operations and enhances the performance of data conditioning tasks such as Super Gather or Angle Gather on multi-terabyte data sets.



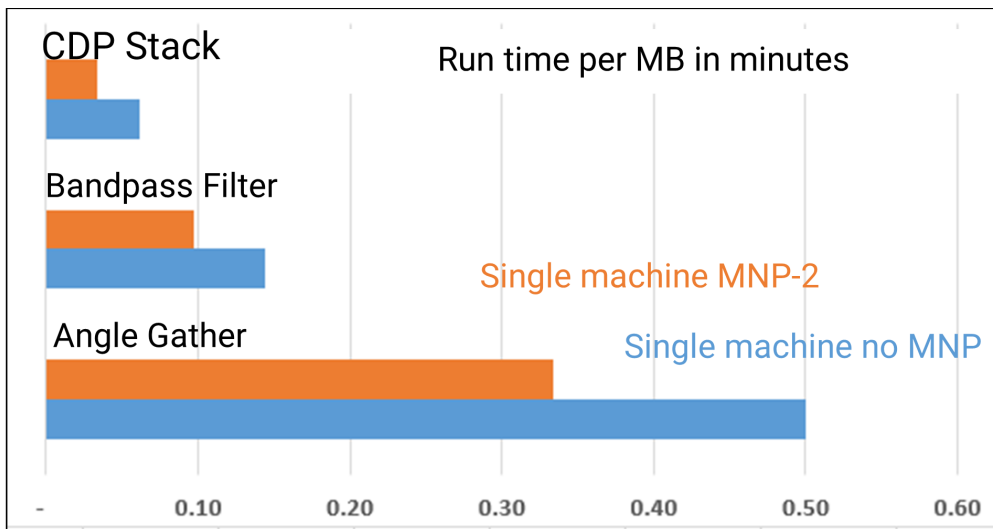
MNP configurations tested on AWS, utilizing a 48 core machine with a Lustre FSx file system processing a 10 TB volume. The processing speedup is roughly linear with the number of nodes.



Performance testing in the cloud environment, showing speed gains with MNP-cluster for both CPU and I/O intensive operations. The Bandpass process shows speeds of 1.8 times greater with 8 nodes while Super Gather has a speedup of almost 8 times with 8 nodes.

MNP Local

Local MNP, operating within a single machine, offers a boost in runtime speed for input/output bound processes. Significant time savings are achievable when processing extensive terabyte-sized input datasets, even operating on a single machine. In this configuration, each node reads and processes separate portions of data from the disk, effectively bypassing the delays associated with disk wait times during trace header operations, resulting in significant time savings.



MNP on a single machine: tested on a 4 core, 8 thread machine with local SSD. Comparison used MNP-2 option with 2 nodes, 4 threads each and a single machine with 8 threads, using the full machine capacity in both cases.

Various factors, including parameters, core count, threads, disk speed, memory, cache, and network speed, impact run times and the benefits of MNP across machines or a single machine.

Batch Processing

When handling extensive seismic datasets on your local machine, consider using batch processing for job scheduling. Note that the batch scheduler continues running even if HampsonRussell is closed.

- Plan to execute resource-intensive processes during non-working hours.
- Run processes in the background, allowing you to work on other tasks simultaneously.

Multi-threading and GPU support

Multi-threading enhances software performance by enabling concurrent task execution on multi-core processors. By default, 12-thread support is included. You can enhance this capability with an additional 12-threads license to further expedite resource-intensive processes. GeoSI and GeoAI automatically utilize an unlimited number of threads, constrained only by your machine's resources.

HampsonRussell offers a feature to estimate the optimal thread count for your hardware and parameters, useful for identifying I/O bound processes where performance is constrained by disk or network speeds rather than CPU power. In these cases, MNP processing is preferable.

NVIDIA GPU support is available in Strata Pre-stack Inversion and in GeoAI Convolutional Neural Networks (CNNs).

The HampsonRussell Advantage

HampsonRussell's advanced accelerator offerings empower you to achieve faster results while optimizing the utilization of computational resources, ultimately boosting productivity, and enabling more efficient workflows.

Contact us to learn how HampsonRussell accelerators can help to optimize your project turnaround times and take full advantage of your hardware resources.