



TRE (Tele-Rilevamento Europa)

TRE, located in Milan, Italy, was established in 2000 as the first commercial spff from the Politecnico di Milano Technical University. It is the leading global expert in data processing services derived from satellite-borne Synthetic Aperture Radar (SAR).

TRE is the world-wide exclusive licensee of the Permanent Scatters Technique (PS Technique) which was patented in 1999 by the Politecnico. The PS Technique allows unprecedented accuracy in collecting satellite radar data used for the detection and monitoring of surface deformation phenomena affecting large geographic areas (subsidence, seismic faults, landslides, etc.). In addition, PS technology can highlight the motion of individual buildings.

SUMMARY

Industry: Satellite data processing

THE CHALLENGE

To increase the amount of data being processed in less time and ensure that production capacity scales with their HPC investment.

THE SOLUTION

The integrated software/hardware solution includes the Panasas Operating Environment and the PanFS™ parallel file system with the Panasas DirectFLOW® protocol.

THE RESULT

- Halved the time required to process satellite data
- Doubled production capacity
- Implemented a scalable, future-proof storage solution
- Reduced storage system management time significantly

Working with its demanding customers TRE is able to detect, measure and monitor geophysical phenomena and verify the stability of individual buildings. Its customers include government bodies, energy companies, academic research institutions and space agencies.

The Challenge

The increasing global demand for monitoring land surfaces, analyzing potential environmental hazards and increasing public safety has contributed to increased customer demand for TRE's services. The requirement to deliver better information faster with increased processing performance and storage capacity was becoming a serious challenge. As TRE's customer demand for processed satellite data grew, they increased their compute infrastructure to meet it. However they found that their application performance and production capacity didn't increase linearly as they added more processing performance.

TRE's high-performance computing (HPC) infrastructure is based on a Linux cluster with 35 compute nodes. Each node contains two, dual-core, Opteron

64-bit processors totalling 140 cores. The cluster uses a gigabit Ethernet interconnect and the operating system is Red Hat Enterprise Linux. Their previous storage infrastructure was a fibre-channel SAN (Storage Area Network) containing four storage arrays.

"Our previous storage solution was based on a fiber-channel SAN, but it quickly became a bottleneck as we grew the number of compute nodes to process the growing number of customer jobs and the increasing sizes of the data-sets," said Alessandro Menegaz, IT & Security Manager, TRE. "Our production capacity is tightly coupled with our hardware performance, so we needed a solution that could scale linearly in terms of capacity and performance as we increased the number of jobs we were processing."

The Solution

TRE quickly realized that their storage solution was becoming a bottleneck and impeding their processing performance. With multiple compute cores processing data, but a single data path to storage, it could be equated to rush-hour traffic only

having access to a single lane road. TRE's IT group quickly ascertained that if all of the compute cores could write data simultaneously it would be similar to moving from that single lane road to a multi-lane highway.

TRE sought a storage solution that matched its requirements for performance, scalability and manageability. The data needed to be shared across all the compute nodes in the cluster. Multiple data paths between the compute nodes and the storage would allow their application performance to scale linearly and would allow TRE to process customer jobs faster.

Working closely with their Italian IT partner, E4 Computer Engineering, TRE focused on identifying the best storage solution for their needs. Given the criticality of the results TRE engaged E4 Computer Engineering to create a simulated test environment to compare the capabilities of various storage solutions. According to Simone Tinti, System Engineer, E4 Computer Engineering, Italy, "The Panasas Storage solution excelled in all the competitive tests and was the only solution to exceed the performance, scalability and manageability requirements."

Panasas Storage includes the embedded PanFS™ parallel file system enabling the input and output of data (I/O) to run in parallel across compute nodes and compute jobs. The solution provides multiple data paths and simultaneous parallel access directly between the compute nodes and the Panasas storage. The PanFS parallel file system also offers a single global namespace so that all the compute nodes can see all of the data. The solution is based on storage blades that are added to shelves and placed in industry-standard racks. As new capacity (shelves) is added, the system automatically recognizes that capacity and balances the workload. The Panasas Storage solution provided TRE with the performance, scalability and manageability they sought in order to meet their increased customer demands today and in the foreseeable future.

The Result


The Panasas Storage solution halved the time it takes TRE to read and write stored information. This has significantly increased its production capacity and allows TRE to deliver processed satellite data much faster to its customers, giving them a distinct competitive advantage in the market.

About Panasas

Panasas, Inc., the leader in high-performance scale-out NAS storage solutions, enables enterprise customers to rapidly solve complex computing problems, speed innovation and bring new products to market faster. All Panasas solutions leverage the patented PanFS™ storage operating system to deliver exceptional performance, scalability and manageability.


PREVIOUS SOLUTION

COMPUTE CLUSTER
35 Nodes, 140 Cores, GigE,
Red Hat Enterprise Server



PANASAS SOLUTION

COMPUTE CLUSTER
35 Nodes, 140 Cores, GigE,
Red Hat Enterprise Server



Alessandro Menegaz at TRE summarized their results in this way: "The Panasas solution has increased our processing capability significantly by providing parallel data paths directly between the server nodes and the storage system. We've literally halved the time that our system spends on processing all the SAR data. Within five minutes of installing the Panasas Storage cluster we could see all clients and all data".