transtec

HPC Storage Solutions at transtec

Parallel NFS with Panasas ActiveStor

HIGH PERFORMANCE COMPUTING AT TRANSTEC

More than 30 Years of Experience in Scientific Computing

1980: **transtec** founded, a reseller for DEC equipment
1987: SUN/SPARC workstations in portfolio
1991: IBM RS/6000 systems in portfolio
2000: **„Kepler Cluster“** in Tübingen installed (#215 in **TOP500**)
2003: transtec HPC cluster at **RRZ Erlangen** no. 317 in **TOP 500**
2005: transtec is **key supplier** for **CERN** in Geneva
2007: transtec HPC cluster at **KIT Karlsruhe** no. 104 in **TOP 500**
2010: transtec focusses on HPC as a **strategic business unit**
up to now: around **500 HPC installations** in Europe

HIGH PERFORMANCE COMPUTING AT TRANSTEC

High Performance Computing in Europe

Matthew Prew
Country Manager UK & Ireland
Unit 5, 29-30 Horse Fair
Banbury, Oxon OX16 OBW
Tel. +44 1295 814501
matthew.prew@transtec.co.uk

Marc van Schijndel
Country Manager Netherlands & Belgium
Postbus 38040
NL-6503 AA Nijmegen
Tel. +31 24 34 34 210
mvanschi@ttec.nl

Vincent Pfleger
Country Manager France
Parc d'Innovation
Immeuble le Pythagore
11 Rue Jean Lapidus
F-67400 Strasbourg-Illkirch
Tel. +33 3.88.55.16.27
vincent.pfleger@transtec.fr

Maritta Hartl
Country Manager Germany
Waldhörnlestr. 18
D-72072 Tübingen
Tel. +49 7071/703-101
maritta.hartl@transtec.de

Rainer Scherf
Country Manager Switzerland
Riedmattstr. 9
CH-8153 Rümlang
Tel. +41 44/8184-840
rainer.scherf@transtec.ch

HPC Competence Center

HPC Lab:
- evaluating new technology
- configuration/preinstallation of customer systems

HPC Benchmark Center:
- running and evaluating **customers' applications**
- **remote** customer **access**

12 HPC Specialists in **Sales**, **Presales**, and **Solution Engineering**

HIGH PERFORMANCE COMPUTING AT TRANSTEC

Solutions for Getting Results Faster and Easier

transtec HPC solutions for scientific and technical **simulations** and **Big Data Analytics**:

- Life Sciences & Pharmaceutical
- Engineering, Automotive & Aerospace
- Climate & Environmental Research, Geophysics
- Finance & Insurance

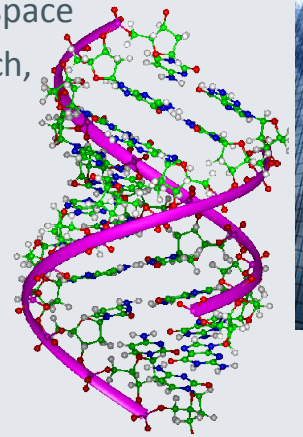

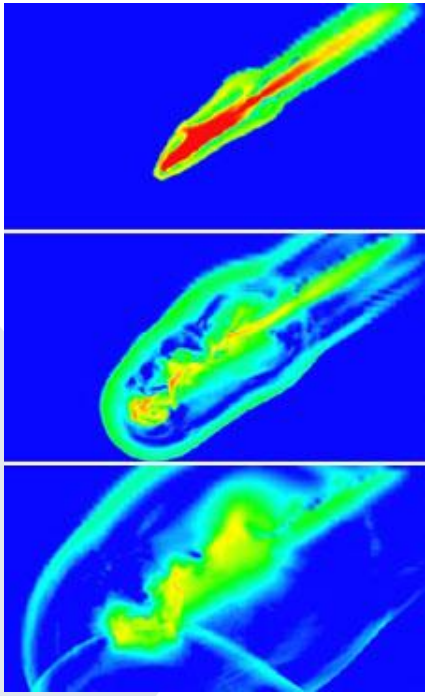

HIGH PERFORMANCE COMPUTING AT TRANSTEC

Customer Care From A to Z

comprehensive solutions for High Performance Computing (HPC)

- **turnkey**
- caring for the customer during the whole **solution lifecycle**:
  - from **individual consulting** to **managed services**
  - customers have access to **transtec benchmark center**

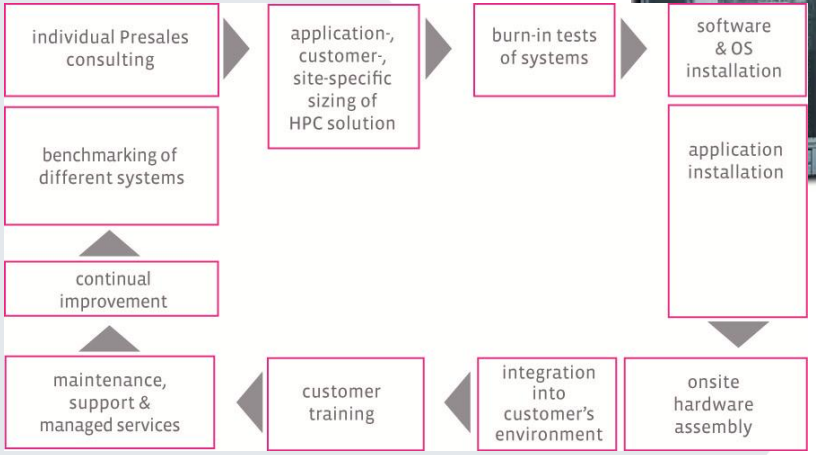

CASE STUDY 2011

Manufacturer of Engines and Generators – several European sites

Deployment of HPC cluster solutions at **2 sites** connected by a **dedicated WAN** link.

- 40 nodes, 480 cores, total of 1,920 TB RAM, QDR InfiniBand interconnect
- 200 TB parallel HPC storage
- **Moab Grid Suite** as middleware layer for unified management of 2 sites
- Installation and configuration of approx. **13 productive applications**, integration into Moab **workload management** system
- very high demands on energy-efficiency and user- and admin-friendly application management

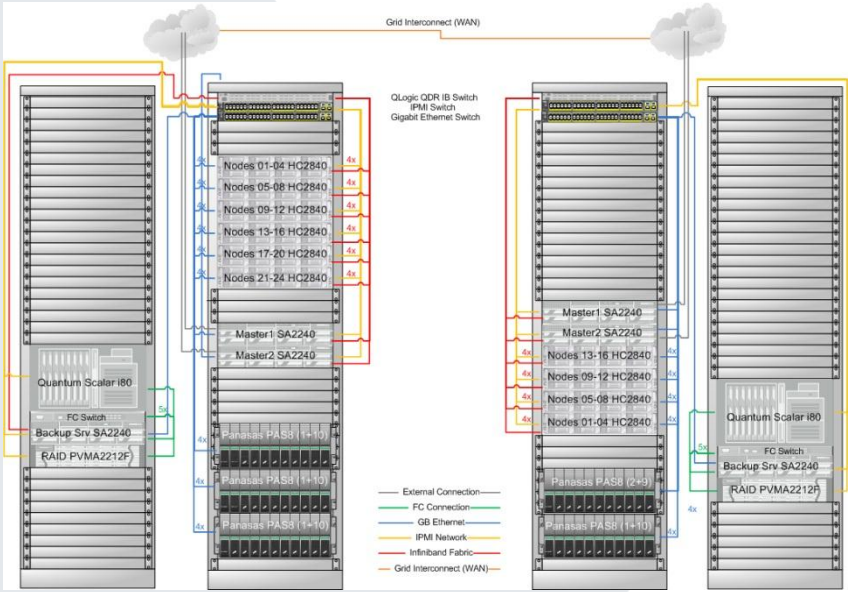

STORAGE DEMANDS IN HPC

- need for **computing power**
  - due to need to run larger and more accurate models
  - more CPUs, more cores, more nodes, more RAM
- need for **network performance**
  - more highly paralellized jobs
  - high-speed interconnects (10GbE, InfiniBand,...)

→ **massive explosion of data sets**

→ demand for

- **large storage capacity**
- **high bandwidth**
- **low latency**

DEFICIENCIES OF TODAY'S SOLUTIONS

- most widespread solution: **single NFS server**
  - does not scale: NFS head is bottleneck
  - „high-speed" NFS server will be bottleneck by tomorrow
- **„clustered NFS":** problematic
  - either head-to-head synchronization limits scalability
  - or manual partitioning of global namespace is cumbersome
  - NFS is not suitable for dynamical load balancing (inherent state)

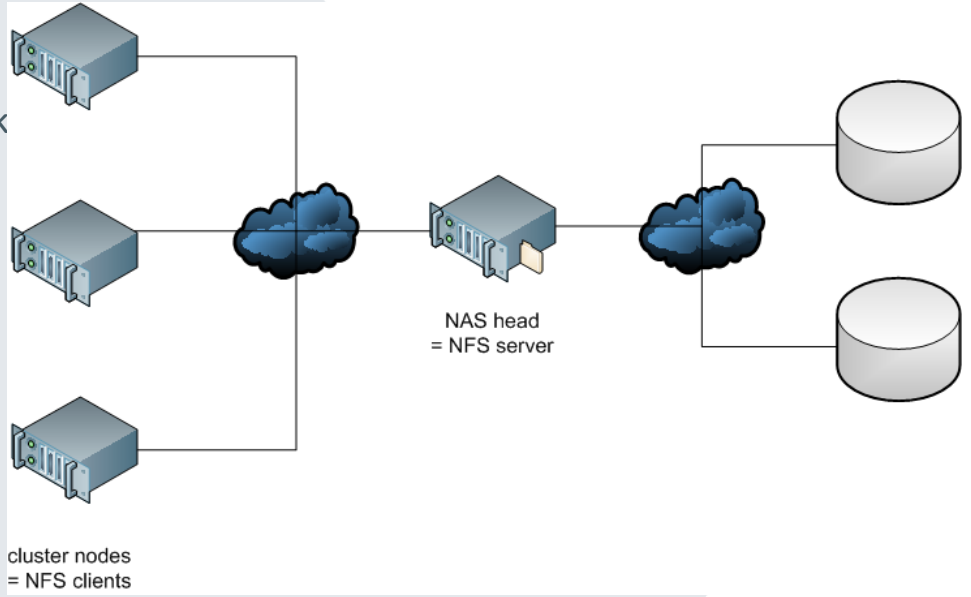

PARALLEL NFS (PNFS): GENERIC ARCHITECTURE

- **separation** of **metadata** path and **data** path (**out-of-band** global namespace)
- built for **interoperability** and **backwards-compatibility**
- **flexible** design allows for different storage implementations (**layouts**)

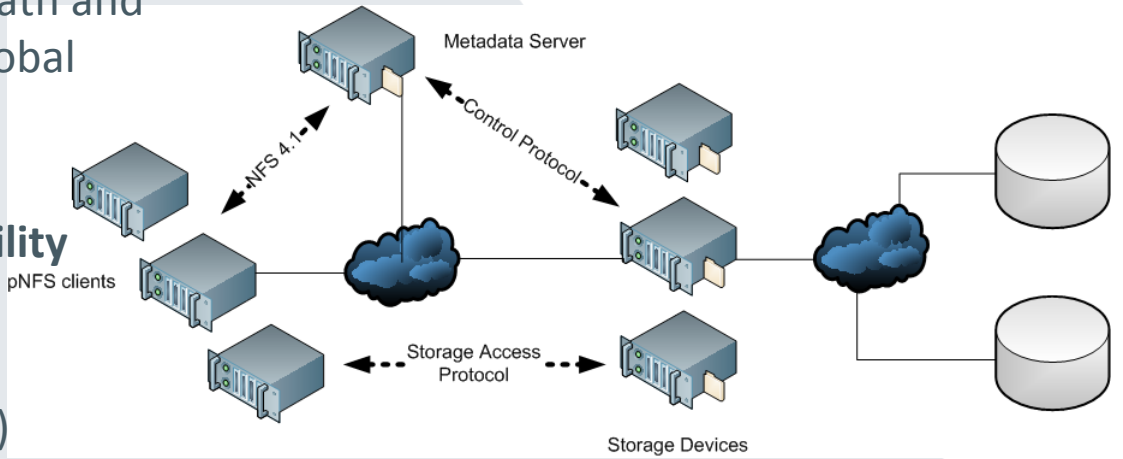

PANASAS HIGH PERFORMANCE STORAGE FROM TRANSTEC

- Blade-based parallel NFS appliance
- 11 blades per shelf = approx. 60 TB in 4 U
- Approx. 1.5 Gbyte/sec aggregate bandwidth per shelf

transtec

PANASAS HIGH PERFORMANCE STORAGE FROM TRANSTEC

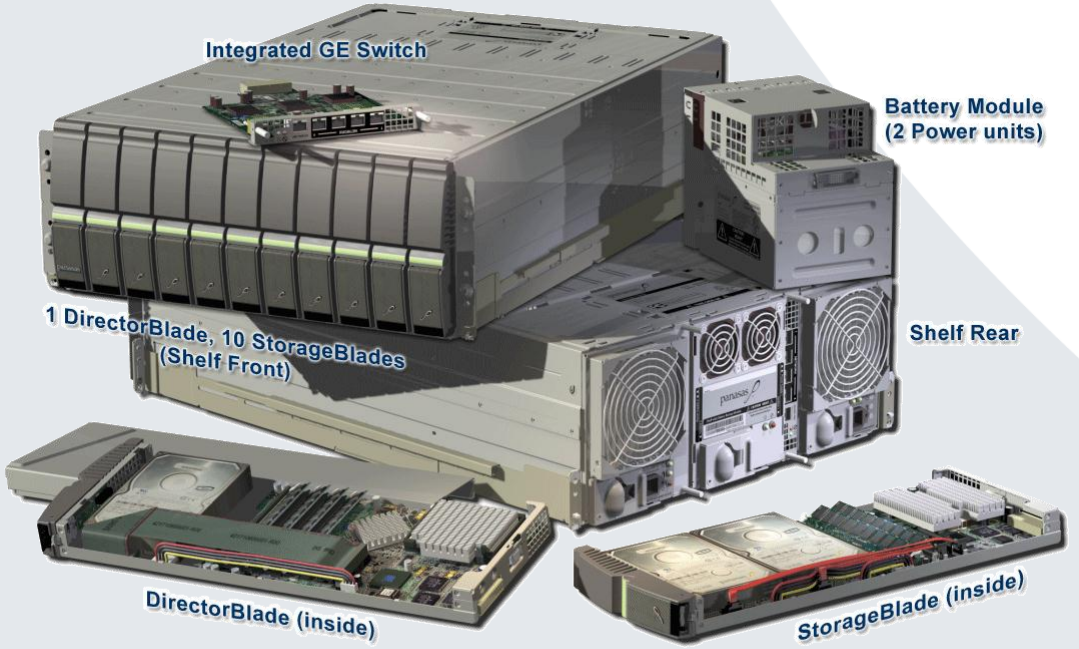

PANASAS ACTIVESTOR 11 AND 12 SPECS

	ActiveStor 11	ActiveStor 12
Product Focus	Balanced Capacity & Performance	Highest Performance
Read Throughput (MB/sec)	1,150	1,500
Write Throughput (MB/sec)	950	1,600
File Creates/Sec. per Director Blade (Metadata Performance)	4,260	6,250
Capacity (TB)	40/60	40 / 60
Cache (GB)	40 + 8	80 + 12
Architecture	64-bit	64-bit
High Availability Network Failover	Optional	Standard
Link Aggregation	No	Yes

PANACTIVE MANAGER

- Single Point of Management
- Simple out-of-box experience
- Seamlessly deploy new blades
- Capacity & load balancing
- Snapshots
- 1-touch reporting capabilities
- Scriptable CLI

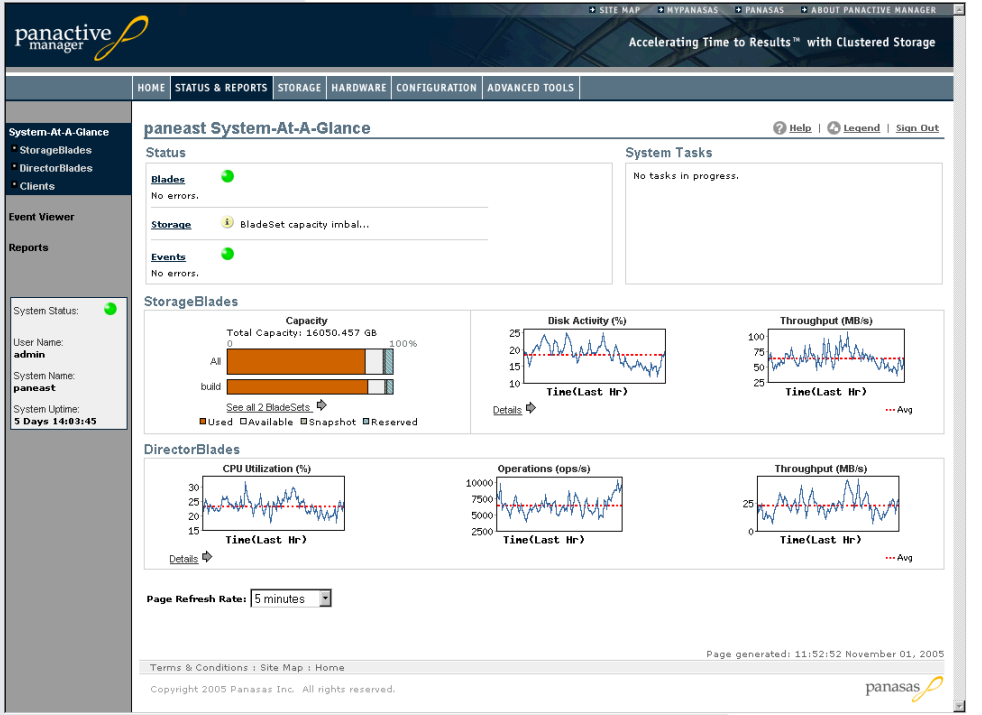

EASY TO MANAGE

- Scriptable CLI
- Easy web-based or CLI-based setup

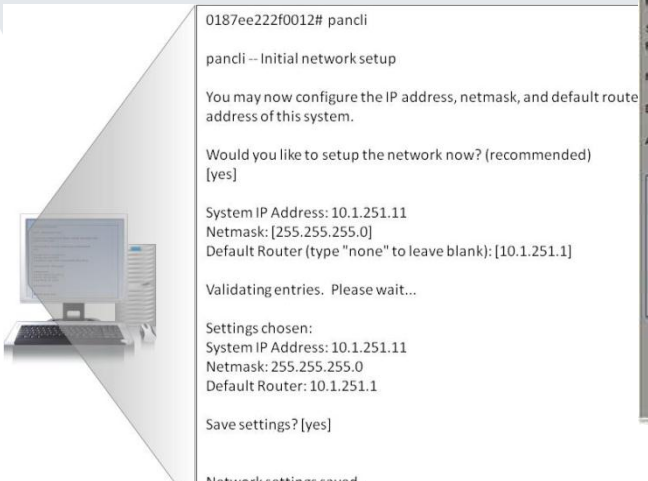

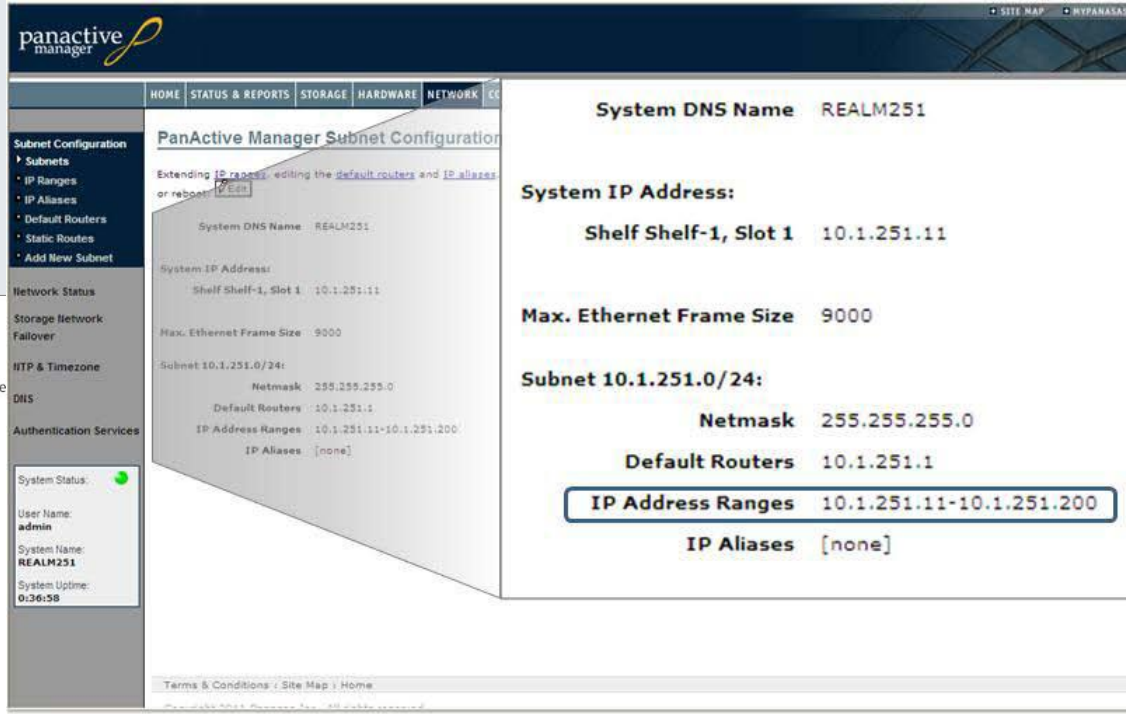

transtec

EXTENSIVE MANAGEMENT AND MONITORING CAPABILITIES

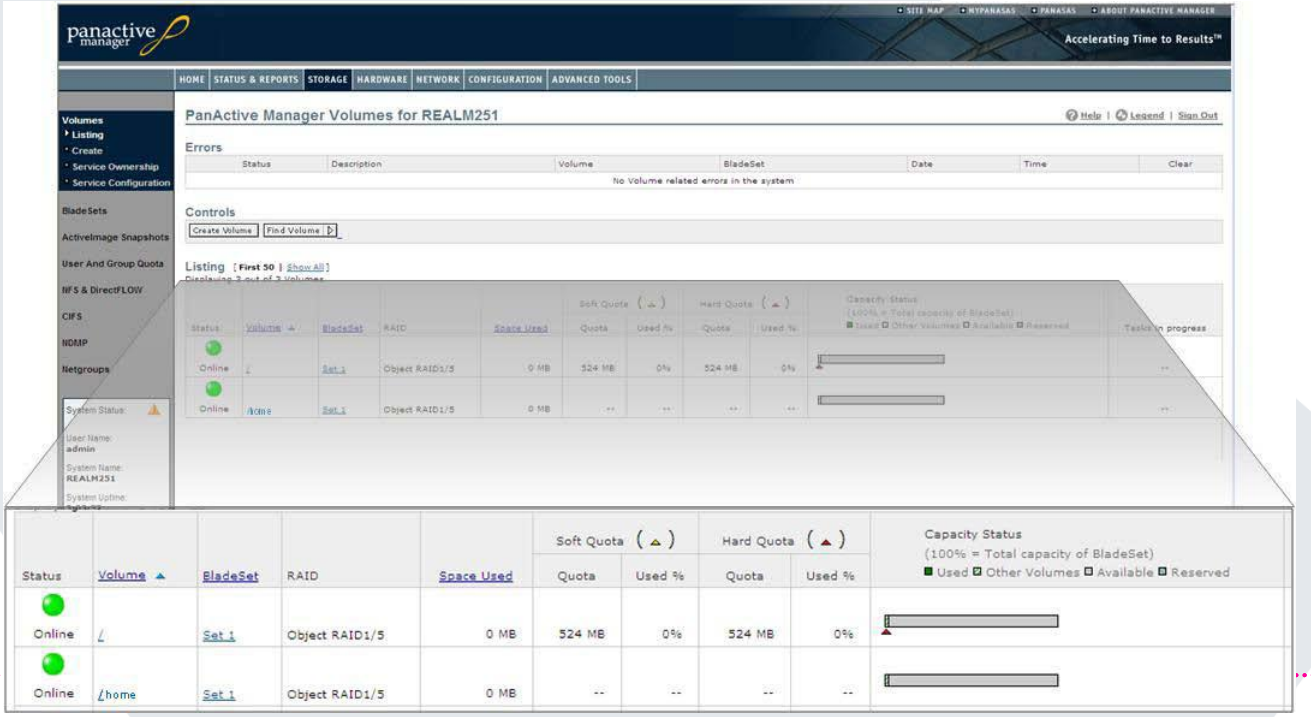

TIERED-PARITY RAID

- 3-level parity: network, horizontal, and vertical

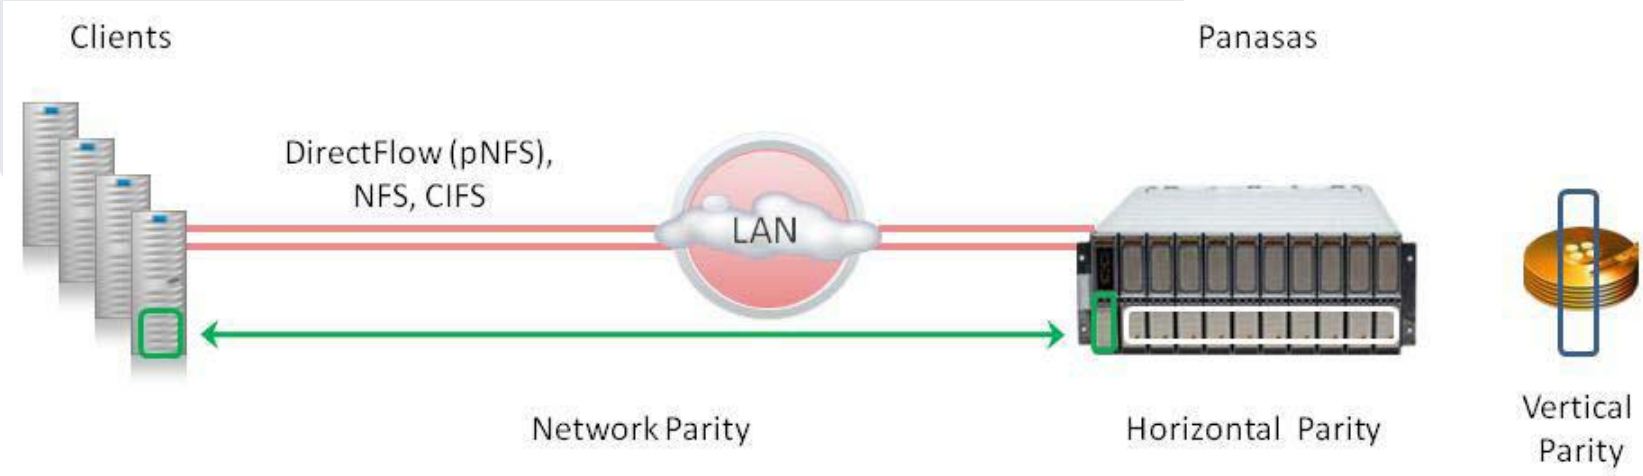

SCALABLE PERFORMANCE

- ESG Labs Performance Tests with ActiveStor12 chassis

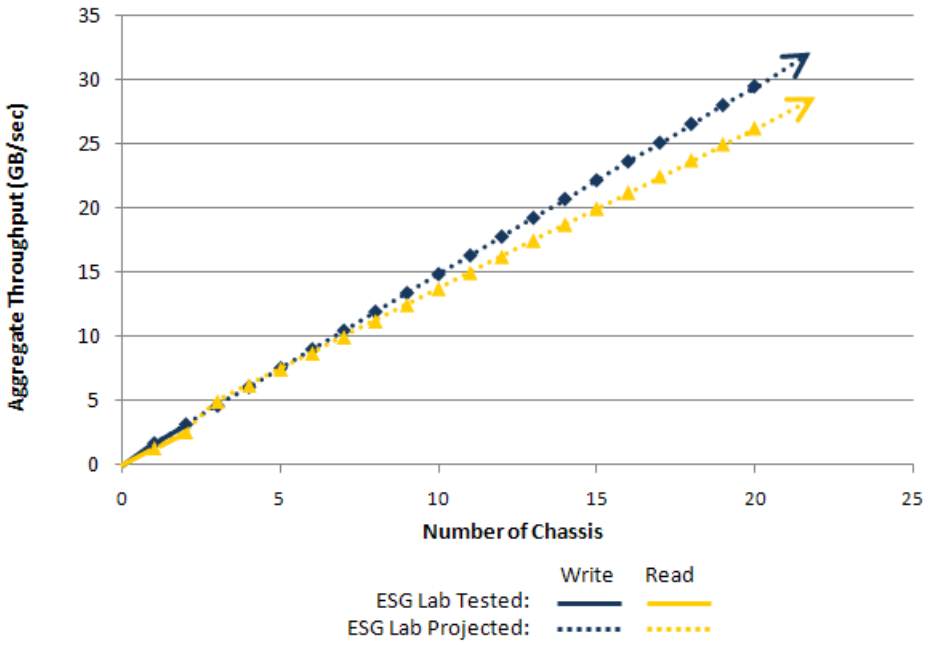

NEW: PANASAS ACTIVESTOR 14

- additional SSD drives for increasing IOPs and metadata performance
- 30-50% faster RAID reconstruction
- enhanced management GUI with PanFS 5
- approx. 80 TB per shelf raw capacity
- 1.6/1.5 Gbyte/s write/read throughput per shelf
- Approx 14,000 IOPs per shelf
- More efficient double-parity algorithm

transtec

accelerate productivity