




Trantor®

HIGH PERFORMANCE COMPUTING AND AI · QUANTUM COMPUTING

Trantor® for HPC

**Supercomputing delivered turnkey,
with quantum included.**

Supernodes with 1, 4 or 8 GPUs, TVirt® serving the compute and AI machines, Altair/Siemens HPCWorks PBS Professional and Navops for scheduling, a catalog of open source AI models already tuned and optimized, and one or two 3-qubit quantum computers, integrated into the same console. One bundle, already configured, hardware included.



GPU NODES

SUPERCOMPUTING VMS

ALTAIR/SIEMENS HPCWORKS

OPTIONAL QUANTUM MACHINES

WHERE THE OTHERS DON'T GO

High performance computing, in your datacenter or on the field. Your data stays under your control and your rules, always.

trantor.it · info@trantor.it

10 August 2026

The contents of this brochure (text, graphics, logos) are the property of PetaBit S.r.l. and are protected by copyright and industrial property law; they may not be reproduced without written permission. All other trademarks cited belong to their respective owners. The information published is for information purposes only and does not constitute a contractual offer, and PetaBit does not warrant the absence of inaccuracies. All prices shown are subject to change and do not constitute a contractual offer.

What this solution solves

A compute cluster is traditionally a slow and expensive integration project. Meanwhile the cloud meter keeps running, and prying eyes keep watching.

Long, costly integrations

Months spent making hypervisor, GPUs, scheduler and storage talk to each other, using skills that are rare, contested and expensive. Every upgrade reopens the building site.

AI still waiting

AI is becoming part of everyday work, but you don't want to hand your confidential data to third parties, nor be at the mercy of their political decisions.

Eyes everywhere

Your company's confidential data is precious material for cloud providers and for many companies overseas. Moving supercomputing to the cloud, or using private AI models as a service, exposes your data to unwanted eyes.

Cloud on the meter

Pay-as-you-go GPUs feel nimble until the project grows: then costs explode, and the price list changes overnight without your consent.

Quantum postponed

It looks like science fiction, so it gets postponed. But hybrid classical-quantum workflows are learned now on real hardware, not on slides.

Field teams left out

Your company may often work with people or IoT devices deployed in remote or constantly changing locations, where poor connectivity prevents continuous, full use of supercomputing and AI.

THE PHILOSOPHY

«Supercomputing should be within reach of every company: turnkey bundles, everything under control, and everything usable on the field too.»

Five souls, one dashboard

GPU supernodes

Compute nodes with 1, 4 or 8 GPUs, sized on your workload, configured and installed by us. No bills of materials to chase: one supplier, one signature.

TVirt®, the hypervisor

Supercomputing and AI machines are VMs served by Trantor® TVirt®: isolation between projects and groups, snapshots, and compute that becomes an internal service, not a privilege for whoever holds the keys to the rack.

Altair/Siemens scheduling

Queues, priorities and job accounting speak the language of HPC: Altair/Siemens HPCWorks PBS Professional and Navops, integrated and configured by us. Users launch jobs; they don't have to care about what sits underneath.

Trantor® AI, ready to go

A catalog of open source models already tuned and optimized for the supernodes: pick them, serve them, use them. Keeping them current is our job; the data stays in your company.

Quantum, integrated

One or two 3-qubit quantum computers, built by a partner of ours, with Trantor® software integrating them into TVirt®: **in production**, they generate the random seeds (TRNG) for the cryptographic generators of TVirt® and of the VMs, and they also serve as the training ground for hybrid classical-quantum workflows, all from the same dashboard.

ONE BUNDLE, ALREADY CONFIGURED

The five components arrive together, powered up and tested: integration is part of the product. In the pages that follow, each area is covered in turn.

GPU supernodes: the specifications

The first area is the hardware: servers with 1, 4 or 8 GPUs, supplied, installed and tested by us. They are designed and built to suit your requirements.

Form factor

The servers come in standard 19" rack format, from 1U to 8U depending on the number of GPUs and on the power required.

Compute power

The servers carry latest-generation Intel or AMD CPUs (a single server can reach 384 physical cores and 6TB of DDR5-6400 ECC RAM). The graphics cards selected are typically NVIDIA RTX 6000 PRO (96GB GDDR7 VRAM, optionally ECC) or NVIDIA Blackwell Ultra (279GB HBM3e VRAM + NVLink).

Connectivity

The supernodes come with connectivity starting at 100 Gbit/s and reaching up to 800 Gbit/s.

Mobility

Whatever their size, our supernodes can be fitted with a wheeled trolley (on continuous tracks, if needed) to make transport into the field easier.

Physical security

Every supernode leaves the factory with an integrated TPM 2.0 module and full support for digitally signed bootloaders (with the customer's own signing key, if preferred), plus encrypted disks with automatic unlock and the recovery key held in the customer's safe.

Why buying our bundle with the hardware included pays off

- ✓ **Sized on your workload:** the right model is chosen from your requirements, not from a price list.
- ✓ **Delivered installed and tested:** the equipment arrives already configured and ready to run.
- ✓ **Grows node by node,** under the same scheduler: to add power you add a node, from any vendor.
- ✓ **One supplier, one signature:** hardware, software and support answer the same phone number.

EVERYTHING INCLUDED

No tangled configurations, no compatibility to study: the supernodes are born for the stack that governs them, and every model is already well defined and tested.

GPU supernodes: the models

This page gives you a look at some of our best-selling configurations. Prices move from week to week with the volatility of the hardware market, so a current quotation is always required.

SKU	TRNT-HPC-N8-B300	TRNT-HPC-N8-6000PRO	TRNT-HPC-N4-6000PRO	TRNT-HPC-N1-6000PRO
Form factor	19" RACK 8U	19" RACK 4U	19" RACK 2U	19" RACK 1U
CPU	2x AMD EPYC 9965 (384 physical cores total)	2x AMD EPYC 9845 (320 physical cores total)	2x Intel Xeon 6788P (172 physical cores total)	1x Intel Xeon 6745P (32 physical cores total)
RAM	6TB DDR5-6400 ECC RDIMM	3TB DDR5-6400 ECC RDIMM	2TB DDR5-6400 ECC RDIMM	1TB DDR5-6400 ECC RDIMM
GPU	8x NVIDIA Blackwell Ultra (2232GB HBM3e VRAM total, with NVLink)	8x NVIDIA RTX 6000 PRO (768GB GDDR7 ECC VRAM total)	4x NVIDIA RTX 6000 PRO (384GB GDDR7 VRAM total)	1x NVIDIA RTX 6000 PRO (96GB GDDR7 ECC VRAM total)
Network	2x NICs with one QSFP112-DD 400Gbit/s port	1x NIC with two QSFP112 200Gbit/s ports	1x NIC with two QSFP112 200Gbit/s ports + 1x NIC with four RJ45 10Gbit/s ports	1x NIC with four SFP28 25Gbit/s ports
Disks	2x 400GB M.2 NVMe SSD DWPD3 in RAID1	2x 480GB SATA SSD DWPD1.5 in RAID1	2x 400GB M.2 NVMe SSD DWPD3 in RAID1	2x 400GB M.2 NVMe SSD DWPD3 in RAID1
Power draw (peak)	19800W	6000W	2000W	860W
Software included	Trantor® TVirt®, Trantor® AI General Purpose, Altair/Siemens HPCWorks PBS Professional and Navops	Trantor® TVirt®, Trantor® AI General Purpose, Altair/Siemens HPCWorks PBS Professional and Navops	Trantor® TVirt®, Trantor® AI General Purpose	Trantor® TVirt®, Trantor® AI General Purpose
Support included	hardware 5 years on-site, software 3 years	hardware 5 years on-site, software 3 years	hardware 5 years on-site, software 3 years	hardware 5 years on-site, software 3 years
Price excl. VAT (as of July 2026)	€1,600,000.00	€640,000.00	€358,000.00	€130,000.00

FULL FLEXIBILITY

Can't find the configuration that suits you? Any change is possible: we look at your requirements together and adapt the models.

TVirt® for supercomputing and AI

Installed inside the GPU supernodes, Trantor® TVirt® turns the HPC cluster into an internal service: high performance computing and AI are served as virtual machines and containers, with the isolation, the snapshots and the benefits of a real hypervisor.

HPC goes virtual

- ✓ **VMs preconfigured for supercomputing and AI:** drivers and environments ready to use and already optimized.
- ✓ **Isolation per project and per group:** every team gets its own pool of resources.
- ✓ **Snapshots, cloning and templates:** management and maintenance are easier and faster than on bare metal.
- ✓ **No elevated privileges:** departments are given VMs and containers, not control of entire racks.

Trantor® AI

- ✓ **A catalog of open source models** already tuned and optimized for the supernodes.
- ✓ **Raise your company's productivity:** automation of repetitive tasks, intelligent search across email and knowledge base, and much more.
- ✓ **Model updates handled by us,** like the rest of the stack.
- ✓ **Your data never leaves the company:** the AI works on your data without your data working for others.

DON'T GET LEFT BEHIND

Trantor® turnkey solutions for supercomputing and AI let any company benefit from state-of-the-art technology without spending months on research and training.

Altair/Siemens HPCWorks

Queues, priorities and job accounting are handled the way any HPC engineer expects: Altair/Siemens HPCWorks PBS Professional and Navops arrive integrated and configured by us, inside the same bundle.

Scheduling, as standard

- ✓ **Altair/Siemens HPCWorks PBS Professional:** queues, priorities, reservations and job accounting.
- ✓ **Altair/Siemens HPCWorks Navops** for orchestrating workloads across the cluster, with native TVirt® integration.
- ✓ **Users worry only about their compute jobs;** nobody needs to know what happens behind the scenes: the platform underneath is an implementation detail.
- ✓ **Integration is our job:** scheduler, hypervisor and nodes arrive already working together, thanks to the partnership between Trantor® and Altair/Siemens HPCWorks.

Altair, Siemens, HPCWorks, PBS Professional and Navops are trademarks of their respective owners: they enter the bundle under their own name and with their own licences.

THE USABILITY OF A STANDARD HPC

Those coming from standard HPC environments find the tools they already know; those starting from scratch find a platform that is easy to use.

Quantum in your datacenter

Three or six qubits will not speed up your workloads, that is what the supernodes are for. But these three or six qubits work in production from day one, doing what no CPU or GPU can do. Trantor® lets companies become familiar with a technology of the future by using it in the present.

In the bundle

- ✔ **One or two 3-qubit systems** built by a partner of ours, working alongside the GPU supernodes.
- ✔ **In production from day one**, doing what no CPU or GPU can do: every quantum measurement is turned by Trantor® software into **TRNG seeds** feeding the cryptographic generators (CSPRNG) of TVirt® and of the virtual machines. VMs are famously hungry for entropy, especially at boot: here they get it from quantum mechanics.
- ✔ **To practise and learn hybrid classical-quantum computing**: algorithms, skills and integration are learned today, on real quantum machines.

The model

SKU	TRNT-QUANTUM-3QB
Form factor and size	Desktop 405x370x220mm (HxWxL), weight 23 kg
Processor	3-qubit quantum processor, NMRQC type
Quantum characteristics	Coherence time: ~6s(T ₁), ~300ms(T ₂); Grover Algorithm Fidelity: 0.83; Deutsch Algorithm Fidelity: 0.88
Power draw	60W
Software included	Trantor® Quantum Bridge
Support included	hardware 3 years, software 3 years
Price excl. VAT (as of July 2026)	€300,000.00

CPU, GPU AND QUANTUM: EACH HAS ITS ROLE

Just as GPUs do a different job from CPUs, quantum computers do a different job from both. Trantor® has invested so that companies can both exploit this new technology straight away and prepare for a future in which it will be widespread.

Computing on confidential data

Sovereignty over your computing

- ✓ **Everything on-premise:** supernodes, scheduler, models and qubits sit inside your company. The hardware is your property and no data leaves your datacenter.
- ✓ **Fully offline AI, high quality:** the AI works on your data without your data working for others.
- ✓ **No mandatory telemetry:** suitable for isolated and classified environments.
- ✓ **Clear terms:** components, versions, support and boundaries, all stated before signature.

For those who really compute

Research, engineering, life sciences, finance, universities: wherever computing is the craft and data is the asset, sovereign supercomputing is not an indulgence, it is the condition for doing the work properly.

For those who want to use AI

Companies hold millions of data points that could feed advanced analysis and research: emails, knowledge bases, chats, phone calls and much more. But taking that data outside puts the company's security and its people's privacy at risk.

THE NEXT MOVE

See for yourself what the system can do.

You will be able to run one of your own jobs or AI tasks, and judge the performance and the quality of the system first hand.

trantor.it
info@trantor.it