



RSC Tornado – hyper-converged solution for future computing

Alexander Moskovsky
CEO, RSC Group

GRID 2018, JINR, Dubna (Russia)



HPC innovations since 2009

**Development of innovative ultra-high dense
energy efficient HPC solutions
delivering unique features and addressing
specific end-user needs**

**Cutting-edge supercomputers and
data centres for demanding customers**

Development directions


Computing density

Power density

Energy efficiency

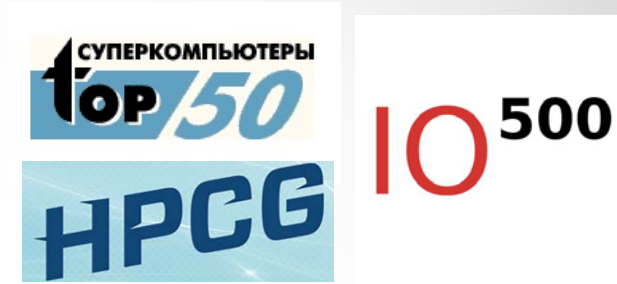
Reliability

Ease to manage and maintain

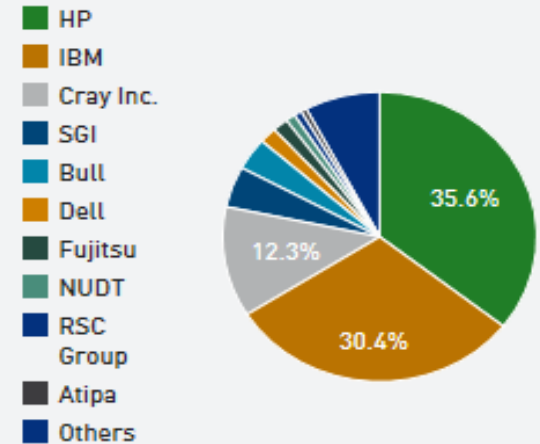


Strong market position

- Leading innovative HPC solution provider in Russia/CIS
- **24% share** in local **Top50** rating (Russia/CIS)
- **4 supercomputers** in **Top10** there
- **#9** position in **IO500** list (hyper-converged system at JINR)
- Over **70%** of all **Russian systems** in **HPCG** rating
- **National Champion rank** by the Russia's Ministry of Economic Development
- One of the leading HPC solution providers in EMEA
- Single Russian company ranked in **Top10 HPC vendors by Top500**
- Over **4.5 PFLOPS** total performance of installed base
- Over 9 years of successful deployments
- **Intel HPC Data Center Specialist elite status**
- **Intel® Select Solution for Simulation and Modeling**




Vendors System Share



Source: Top500.org (Nov'14)





Track of world records

RSC Tornado cluster solution



Highest computing density per rack* – **1.41 PFLOPS**

Leading performance density – **490 TFLOPS/m³**

Highest power density per rack – **200 kW**



Leading computing density per rack* – **685.44 TFLOPS****

Highest performance density – **535 TFLOPS/m³**

Leading power density per rack – **100 kW**


Leading PUE*** = **1.027** (measured at the customer site)



* 42U rack 80x80 cm

** for Intel® Xeon® based solutions

*** Power Usage Efficiency is less than 6%



RSC TORNADO FLEXIBLE SOFTWARE DEFINED SOLUTIONS

UNIFIED BUILDING UNITS



RSC Tornado Intel® Xeon® based node

WORLD RECORDS



Inlet Water



Water Cooling



~220V = 400V



Density

- Two Intel® Xeon® Scalable (incl. top-bin) and Intel® Xeon® E5-2600 v4 (incl. top-bin)
- Up to 256GiB DDR4-2400 RAM
- Intel® Omni-Path, EDR IB, 10/40/100 GigE
- 2x Intel® SSDs SATA and 1x Intel® SSD with NVMe incl. Intel® Optane™ SSD DC P4800X



RSC Tornado Phi Intel® Xeon Phi™ based node

- Intel® Xeon Phi™ 7200(F) (incl. top-bin)
- Up to 192GiB DDR4-2400 RAM + MCDRAM
- Intel® Omni-Path Fabric, EDR IB, 10/40/100 GigE
- 2x Intel® SSDs SATA and 1x Intel® SSD with NVMe incl. Intel® Optane™



RSC Tornado Unified Cabinet

- Flexible configuration options:
 - Up to 153 RSC Tornado servers [685 TFLOPS]
 - Up to 153 RSC Tornado Phi nodes [528 TFLOPS]
 - Mixed RSC Tornado/RSC Tornado Phi nodes
- From 1 to 9 fully independent domains
- 0.64 m² / 6.9 ft² footprint, 2 m / 6.6 ft height



RSC Tornado Expansion Pack

- HPC, BigData, VDI, Security, Machine Learning Expansion Packs and others by request



RSC Tornado Power Supplies

220-400V AC/DC 12 kW
220-12V AC/DC 2.1 kW

- Direct liquid cooling & node-like design
- Thermal and power management
- Flexible redundancy (N+1 to N+N)
- Power conversion efficiency up to 96%





RSC Tornado node

based on the top-bin


Intel® Xeon® Platinum 8180 Processor \Rightarrow 4.48 TFLOPS




Easy access to components
Memory replacement windows
Broad range of available components



Protective cover



Unified power connectors
for easy upgrade/reconfiguration




Compact Node Design

based on the top-bin 72-cores



Intel® Xeon Phi™ 7290 Processor ⇒ 3.4 TFLOPS



Easy access to components
Memory replacement windows
Broad range of available components



New protective cover



Unified power connectors for easy upgrade/reconfiguration



Hot Water Cooled Intel® OPA


World's first 100% 'hot water' liquid cooled 48-port Intel® Omni-Path Edge Switch 100 Series



ISC'17 (Frankfurt, Germany)



SC'17 (Denver, USA)




New: RSC Tornado HYPeR-converged

RSC Tornado node:

- 2 x Intel® Xeon® Scalable (Skylake-SP) processors up to 205W with 28 cores each
- Intel® Server Board S2600BP with two 10GigE ports on-board and (optional) Intel Quick Assist support
- RSC Management Module with dedicated Ethernet fabric
- Up to 12 hot-swap NVMe SSDs, for example each can be:
 - Intel® SSD DC P4511 (NVMe, M.2) 1-2TB configured as disk
 - or Intel® Optane™ SSD DC 4801X (M.2) 375GB as disk or memory via IMDT
- Memory per node – up to 768GiB DDR4 Reg ECC up to 2666
- 2 x Intel® Omni-Path 100 Gb/s adapter (or EDR InfiniBand or Ethernet) providing up to 200Gbps external fabric bandwidth

Intel® SSD DC P4511 (NVMe, M.2)




Intel® Optane™ SSD DC 4801X (M.2)

NVMe-attached SSDs can provide:

- **Large and fast storage:** up to 24TB+ per node today
- **Large Memory Capacity** node via Intel Memory Drive Technology (IMDT) with up to 4.2TB of RAM today
- **Many combination of previous two options**, e.g. 3TB RAM and 8TB disk

100% 'hot water' liquid cooled solution for stable operation and high safety of components



RSC Variety of Hyper-converged node types




Nodes servicing fast scalable parallel file system (Lustre, BeeGFS etc.)



Nodes with large memory



Standard compute nodes






RSC Tornado Storage-on-Demand

Software Defined

High Performance Storage


On Demand

It provides outstanding performance at 5% of traditional parallel file system cost (TCO)



World first 100% hot water cooled and hyper-converged system at JINR





Hyper-Converged Solution at JINR



Intel Omni-Path

RSC Tornado nodes based on Intel® Xeon® Scalable:

- Intel® Xeon® Gold 6154 processors (18 cores)
- Intel® Server Board S2600BP
- Intel® SSD DC S3520 (SATA, M.2), 2 x Intel® SSD DC P45xx (NVMe, M.2) 1TB
- Memory per node – 192GB DDR4 2666 GHz
- Intel® Omni-Path 100 Gb/s adapter


100% 'hot water' liquid cooled Intel Omni-Path switch:

- 48-port Intel® Omni-Path Edge Switch 100 Series
- Intel® Omni-Path optical cables

RSC Tornado nodes based on Intel® Xeon Phi™:

- Intel® Xeon Phi™ 7190 processors (72 cores)
- Intel® Server Board S7200AP
- Intel® SSD DC S3520 (SATA, M.2)
- Memory per node – 96GB DDR4 2400 GHz
- Intel® Omni-Path 100 Gb/s adapter

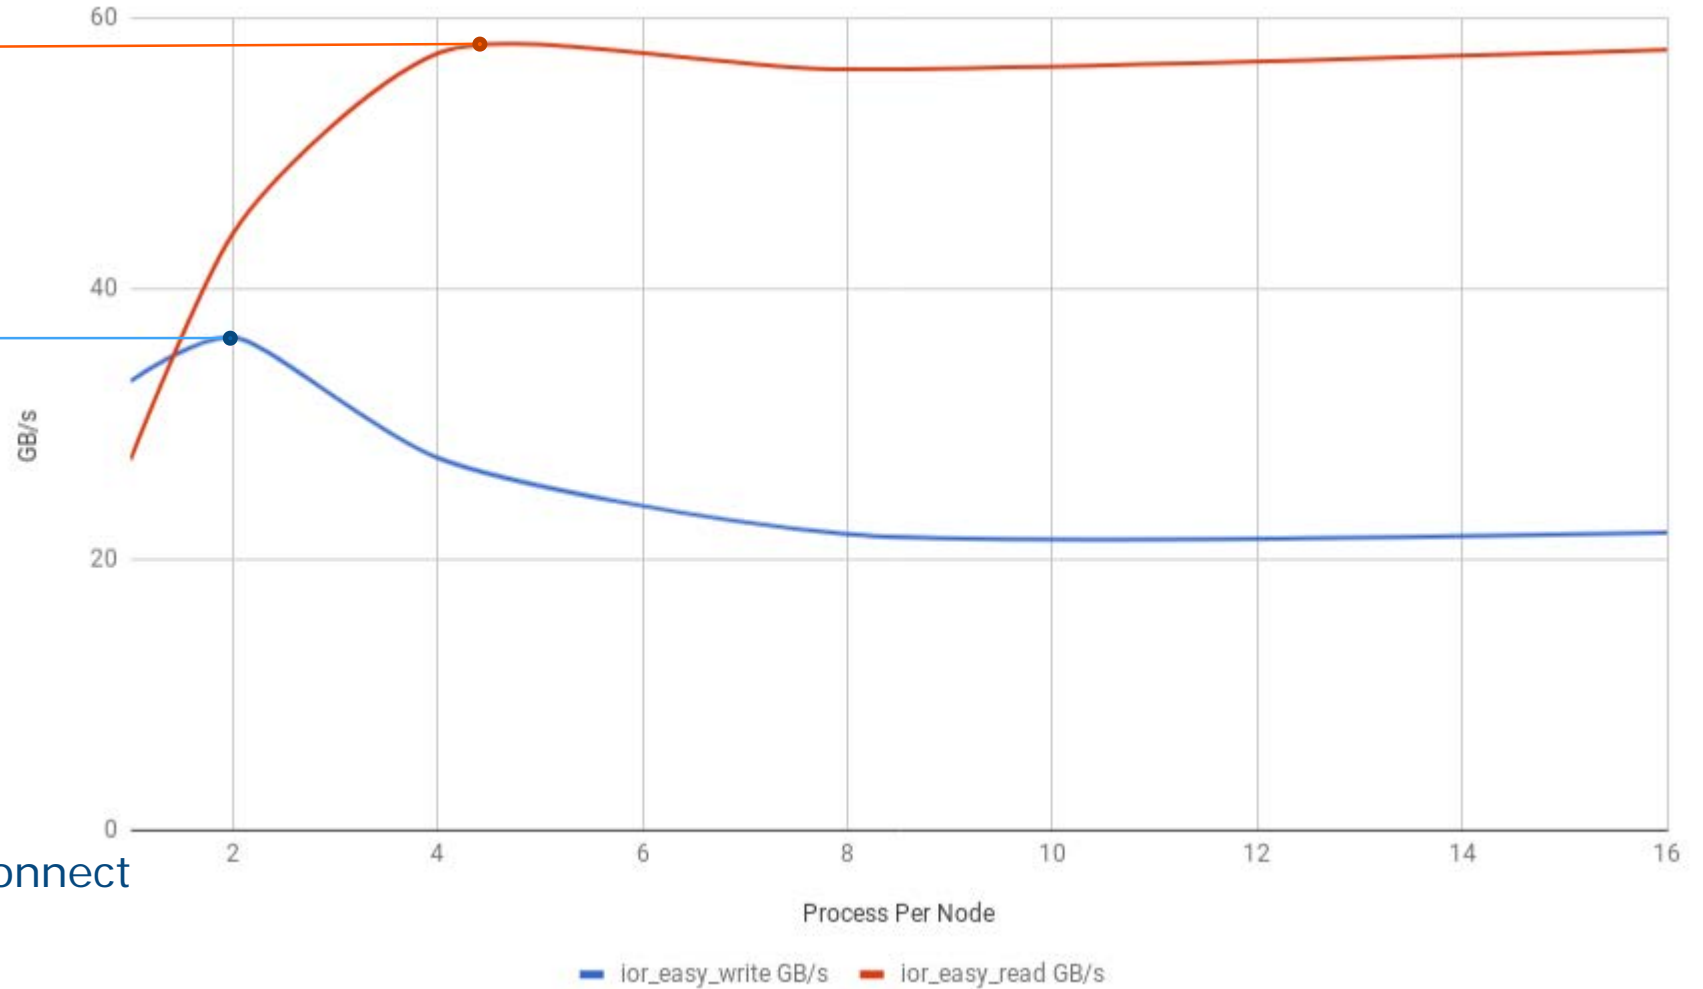
RSC BasIS integrated software stack for system monitoring and management



RSC Lustre Storage-on-Demand io500 benchmark results

56 GB/s easy_read ior test

36 GB/s easy_write ior test



Configuration:

- 1 x MDS
- 12 x OSS (6 x 1TB NVMe SSD)
- 24 Clients to load Lustre
- 100GBps Intel Omni-Path interconnect

IO-500




This is the official list from [ISC-HPC 2018](#). The list shows the best result for a given combination of system/institution/filesystem.

#	information						io500		
	system	institution	filesystem	storage vendor	client nodes	data	score	bw	md
								GiB/s	kIOP/s
1	Oakforest-PACS	JCAHPC	IME	DDN	2048	zip	137.78	560.10	33.89
2	ShaheenII	KAUST	DataWarp	Cray	1024	zip	77.37	496.81	12.05
3	ShaheenII	KAUST	Lustre	Cray	1000		41.00*	54.17	31.03*
4	JURON	JSC	BeeGFS	ThinkparQ	8		35.77*	14.24	89.81*
5	Mistral	DKRZ	Lustre2	Seagate	100		32.15	22.77	45.39
6	Sonasad	IBM	Spectrum Scale	IBM	10	zip	24.24	4.57	128.61
7	Seislab	Fraunhofer	BeeGFS	ThinkparQ	24		16.96	5.13	56.14
8	Mistral	DKRZ	Lustre1	Seagate	100	zip	15.47	12.68	18.88
9	Govorun	Joint Institute for Nuclear Research	Lustre	RSC	24	zip	12.08	3.34	43.65
10	EMSL Cascade	PNNL	Lustre		126		11.12	4.88	25.33
11	Serrano	SNL	Spectrum Scale	IBM	16		4.25*	0.65	27.98*
12	Jasmin/Lotus	STFC	PanFS	Panasas	64	zip	2.33	0.26	20.93

Values with * indicate that a value for the computation was missing.

Further lists with more detail can be found on the navigation menu. For example, a list with [radar chart](#) and [configurable scoring](#).



Top projects

- **Joint Institute for Nuclear Research (JINR)**
- **Saint Petersburg Polytechnic University (SPbPU)**
- **Russian Academy of Sciences (JSCC RAS)**
- **Siberian Supercomputer Center (SSCC SB RAS)**
- **South Ural State University (SUSU)**
- **Institute of Oceanology of Russian Academy of Sciences (IO RAS)**
- **Russian Weather Forecast Agency (Roshydromet)**
- **Moscow Institute of Physics and Technology (MIPT)**
- **Aviation Industry, Energy sector, Computer Graphics, Oil&Gas**
- ... and many others

Over 4.5 PFLOPS

of installed base of innovative liquid cooled supercomputers



www.rscgroup.ru

hq@rsc-tech.ru