Tintri VMstore™ T7200 Series

NVMe Platform

VMstore

Workload Intelligence Drives Tintri's Most Powerful and Flexible Auto Adaptive Platform

The Tintri VMstore T7200 series is the latest addition to the storage industry's first and only purpose-built, workload-aware portfolio. The T7200 series delivers NVMe-driven performance and efficiencies, supporting up to 10,000 managed workload objects in just two rack units. It offers the kind of scalability, availability and security that you'd expect from an enterprise-grade system that powers numerous Fortune 100 firms. Our customers love the outstanding AI-driven efficiency and optimizations that set VMstore apart from traditional approaches. VMstore systems deploy in minutes, self-optimize and dynamically adapt to accommodate the most demanding workloads in the data center maintaining quality of service (Auto-QoS) for each managed workload object. Autonomous operations and advanced real-time and predictive analytics consistently drive down storage management activities and costs – by as much as 95%.

The T7200 series leverages award-winning NVMe technology and provides 40% faster performance than previous generation VMstore systems. VMstore T7200 series customers benefit from using a storage hardware controller that is deployed in the most demanding AI, analytics, deep learning and high IOPS workloads on the planet.

Scale up to 2.58PB[a] with a single VMstore T7200 system and grow beyond 164PB in a scale-out configuration. All your VMstore systems – supporting up to hundreds of thousands of VMs, container persistent volumes, or databases – can be managed from a single console with Tintri Global Center™. Enterprises choosing NVMe as their preferred platform can now gain the advantages of Intelligent Infrastructure to consolidate their enterprise applications, manage storage with generalist IT skills and re-focus on their business.

Features

Consistent Performance
Resource isolation for each and every managed workload, either VM, container or database, all the time

Data Services
Real-time deduplication and compression, public cloud connector, copy data management, and more

Real-Time Analytics
Granular observability across storage, network, and hosts on a per-VM or per-workload basis

Actionable Analytics
Get to root cause in one click. Fix and see results instantly

Predictive Analytics
Accurately model and forecast both capacity and performance needs over the next 18 months by application type

Developer Flexibility
Select Tintri's native REST APIs, PowerShell toolkit, Python SDK, or our many different hypervisor plugins

Per-Application Data Management
Manage snapshots, clones, replication, and QoS policies on a per-app, per-VM, container persistent volume, or per-database basis

High Availability
Dual-controllers, hot spares, inline checksums and referential integrity with real-time error detection

Choice of Application Environments
Supports VMs, container persistent volumes, and databases

Benefits

- Guaranteed high performance for every managed workload without manual intervention
- Concurrent multi-hypervisor and container support enables you to operate with an ever growing list of hypervisors, database systems and container management systems, simultaneously on a single system without partitioning
- Remote management of both initial system configuration and power simplifies administration of your distributed environment
- Share analytics data with vRealize Operations, Microsoft System Center Operations Manager, and other platforms to gain valuable holistic insights
- Easy configuration enables you to go from box to production workloads in under an hour in most cases
- Open APIs deliver flexibility while making scripting simple, standardized, and powerful
- Autonomous operation eliminates the vast majority of manual tasks, saving time and reducing errors, integrates with your existing tools for DevOps, CI/CD or others
- Dependability when and where it's needed most
- Native CSI driver allows for integration directly with multiple container platforms, managing the persistent volume with granular observability, including Kubernetes and OpenShift

Product Specifications		T7290	T7270	T7250
Application Density	Managed Objects	10,000	7,500	5,000
Flash	Effective capacity[ab]	Up to 2580TB		
	Raw capacity	38 to 737TB		
	Data protected as DP/DR target[ac]	Up to 5160TB		
Onboard Network ports per controller	Admin ports	2 x 1/10GBASE-T		
	Data/Replication ports	2 x 1/10GBASE-T		
	Management port - IPMI	1 x 1GBASE-T		
Optional Network ports per controller	Data ports	2 x 100GbE or 2 x 40GbE or 2 x 25GbE[d] or 4 x 10GbE[ef]		
	Replication ports	2 x 100GbE or 2 x 40GbE or 2 x 25GbE[d] or 4 x 10GbE[ef]		
Physical Specifications	Dimensions (HxWxD)	2RU, 3.5" x 19.0" x 34.5" (89mm x 483mm x 850mm) without bezel		
	Weight 10x NVMe drives	84.2 lbs (38.19kg)		
	Weight 24x NVMe drives	90.6 lbs (41.09kg)		
	Power supplies	Dual fully redundant hot swappable power supplies; IEC plug types		
Operational Specifications	Voltage	200-240 VAC / 50-60 Hz		
	Watts (avg./max)	870 / 2400		
	BTUs (avg./max)	2969 / 8184		
	Operating temp.	0°C to +50 °C		
	Operating humidity	8% to 90% (non-condensing)		
	Non-oper. temp.	-40°C to +70°C		
	Non-oper. humidity	5% to 95% (non-condensing)		
System	Type	Fully redundant all-flash hot swappable dual controllers		
Software	Tintri OS	Software Tintri OS TintriOS 6.0 or higher		
Workloads	Supported Hypervisors, Container Platforms, and Databases	VMware vSphere®, Kubernetes, Microsoft® Hyper-V, Xen Hypervisor, Red Hat® Enterprise Virtualization (RHEV), OpenStack, Platform9, OpenShift, and Microsoft® SQL Server		
Additional Software	Management	Tintri Global Center™ Standard (included)		
	Analytics	Tintri Analytics[g]:Up to 3 years of detailed operational historical metrics		
	Tintri Software Suite	ReplicateVM™: Asynchronous Replication Synchronous Replication VMstore Cloud Connector™: Public Cloud Connector Tintri Global Center™ Advanced: VM Scale-out[a] SyncVM™: Copy Data Management: SecureVM™: Data-at-rest Encryption SQL Integrated Storage		
	Other Optional Licenses	Additional features available soon		
Product Support	Administration	Revolutionizing and Minimizing Data Management with autonomous operations for Managed Workloads		
	Support	Proactive support with automated phone home and case creation		
Regulatory	UL/CSA/EN/IEC 60950-1, EMC Emissions Class A, FCC, IC, CE, VCCI, RCM, BSMI, EAC, KC, ROHS, REACH, WEEE			

a. Effective capacity refers to usable space. It is calculated by removing data protection overhead from RAW capacity, and then a space savings multiplier is applied. Data protection overhead includes double parity, hot spare and internal reserves for metadata. Space savings is derived from inline deduplication, compression and clone savings, but does not include thin provisioning. Data reduction typically provides 2.5-5x capacity savings; 5x was used for the value shown.

b. One MB is equal to one million bytes, one GB is equal to one billion bytes and one TB equals 1,000GB (one trillion bytes) when referring to storage capacity. Accessible capacity will vary from the stated capacity due to formatting and partitioning of SSDs, the operating system, and other factors.

c. Assumes minimum policy of 8 hourly snapshots, 7 daily snapshots, and 4 weekly snapshots. All snapshots are logically represented as full recoveries.

d. Supports Negotiation down to 10 GbE via SFP module swap

e. Supports Auto-Negotiation down to 1 GbE.

f. Copper 10GBASE-T only for 4 port card option.

g. Included with an active VMstore maintenance contract.

Tintri