

WARP 41001-M "MemoryMatrix" All-flash Parallel Storage

The **WARP 41001-M** is an ultra-dense scale-out storage appliance for HPC environments that deliver the power, scale, and performance needed to support Big Data for organizations of every size without breaking the bank.

Using Lustre, HPC architects have built parallel file systems support the world's fastest, most scalable storage systems in the world. Unfortunately, teams of PhD-holding administrators were required to implement and maintain such environments. WARP Mechanics has simplified Lustre deployment centered on a scalable storage *appliance*, while retaining the flexibility of a customized solution. The appliance combines storage servers and targets into a tightly-integrated package making it easy to deploy in a modular fashion. The WARP 41001-M can be rolled out in a turnkey manner, and then tailored as needed to meet customer-specific requirements.

Integrating the advanced WARP*fs* file system on enterprise class hardware, a WARP 41001-M is ideally suited for high performance applications, which combine read- or write-intensive data and low-latency onto the same system. Available in several hardware configurations to optimize the appliance for diverse customer workloads, when configured for maximum scalability, it can hold as much as 3.0PB raw in a single rack. That's over 2.2PB usable space, with double parity protection.

The system can also *scale out* to add both performance and capacity as needed by simply deploying more nodes to the same storage namespace.

Truly turnkey, all software is pre-loaded and pre-configured at the factory, so only site-specific parameters need to be configured by the customer. Simply rack the appliance, cable it, and go. Customers purchasing large configurations can order the system pre-racked and cabled in a WARP Mechanics-supplied cabinet.

Don't trust your data to a mere commodity white-box chassis. WARP Mechanics uses reliable OEM-grade storage hardware and SAS drive modules, with best-in-class open storage software in production at the largest HPC shops in the world. The entire system is supported by WARP Mechanics and its resellers for a single support contact.

Software on the WARP 41001-M is not merely a collection of open-source components. WARP*fs* integrates advanced features such as thin provisioning, de-duplication, block-level checksums, copy-on-write, snapshots, replication, and more, that traditional Lustre systems lack, finally bringing data integrity and reliability to the Lustre file system.

- Optimized for high-performance workloads
- Pay as you grow starting from 20TB
- Typical 600TB+ usable with RAID6 (~768TB raw) per rack
- Maximum 2.2PB+ usable (~3.0PB raw) with high-cap SSDs
- Turnkey deployment & Enterprise support and service
- Simple but powerful GUI and CLI management
- Advanced data integrity protections
- Replication with WAN optimization

Whether you deploy it as a stand-alone storage system, or as part of a comprehensive WARP Mechanics datacenter architecture, the WARP 41001-M will provide unmatched density and performance.



Superior Metadata Services

Central to the appliance is the **WARP 41111-M** system to provide Lustre metadata services. Built on the AP-2214 platform and containing up to 24x SSD modules and redundant MDS/MGS controllers, the 41111-M serves as the point of contact for all components of the Lustre environment to obtain information about the landscape. It runs commercially-supported WARPfs built on Lustre + ZFS, developed in partnership with **Lawrence Livermore National Lab**. This code is powering tens of petabytes at LLNL and massive amounts of storage at other Top500 sites. When configured for maximum performance, each 41111-M can deliver over 1M IOPS and over 100GBps.




The MDS/MGS appliance contains two controllers (HA), each with dual high-speed Intel Broadwell processors and up to 1.5TB RAM and multiple SAS-3 connections to the MDT. That's 3TB of RAM and 384Gb total SAS bandwidth in a maximal HA configuration. Front- and back-end interfaces are available for each controller in 10Gb/40Gb Ethernet, 56Gb/100Gb InfiniBand or 100Gb OmniPath ports to deliver redundant, load-balanced access for WARPfs.

Designed for high availability, with no single point of failure for data, cooling, processing, or power, both controllers access to all disks in the enclosure, allowing for multiple MDTs per file system using Distributed Namespace feature of WARPfs. Should one node fail, the other will assume operation transparently with no downtime.

Performance-Intensive Storage

IT architects require much larger building blocks to meet modern data growth rates, and to access giant data sets, storage must also have high IOPS and throughput. The **WARP 41021-M** fills these needs in an efficient, high-performance package by combining the WARP Mechanics® AP-2214 controller with the WDS-2224 JBOD. The AP-2214 controller is configured as an Object Storage Server, providing redundant Lustre OSS nodes in an HA appliance.



In the simplest configuration, the WARP 41001-M uses the WDS-2224 JBOD to provide capacity provided in a three-chassis "Pod" configuration: 2x OSS heads and one OST JBOD with enterprise SSD modules to achieve maximum write and read performance. Each SSD exceeds 100,000 IOPS per module, making the WARP 41001-M easily capable of 192TB and 1 Million IOPS per OST shelf.

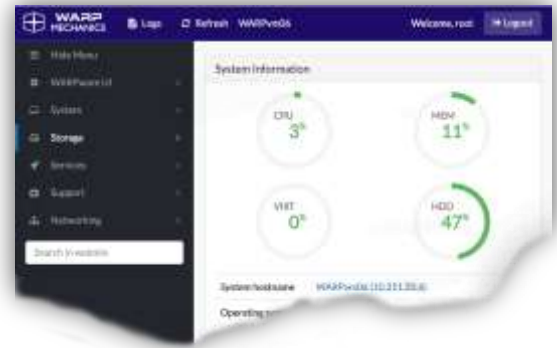
To expand the WARPfs file system, customers can just add another WARP 41021-M to simultaneously increase capacity and performance. A few simple commands expand the file system within the existing namespace with zero downtime.

Customers can bolt on additional JBODs to existing OSSes up to 16x WDS-2224 per AP-2214 OSS. Larger options are available, with half- and full-rack configurations, for greater flexibility to expand the system. An organization can expand the system by adding disks, JBODs, or half- and full-rack bundles, depending on their required growth rates.

Simple but Powerful Management Software

Like all OEM-class product companies, WARP Mechanics leverages existing open source software where appropriate in its WARPware OS. This approach benefits customers through broad development, cost efficiency, higher reliability, and faster delivery of new features and updates.

The WARPware stack starts out as an enhanced version of CENTOS Linux: the leading supercomputing operating system deployed at Lawrence Livermore National Lab, Sandia National Lab, Los Alamos National Lab, and many other top-ranked HPC facilities.



WARP then adds a web-based GUI and a powerful CLI to the tools already included in the OS. The complete management stack is integrated with the WARPfs hardware and software layers, fully supported by WARP Mechanics.

At a lower level, this strategy gives WARPfs the benefits of ZFS, which brings scalability, performance, and enterprise storage features.

For instance, WARPfs supports effectively **unlimited file system sizes**. It is practical to scale one appliance beyond a petabyte in a single filesystem with current disk densities. Livermore, for example, is running a **~70 petabyte** filesystem using this ZFS code.

It also supports virtually **unlimited snapshots**. Other legacy solutions are limited to 255 snapshots, but WARPware can accommodate 2×10^{48} . Run hourly snapshots if desired, to support extremely granular recover points.

Of particular interest in large configurations, WARP Mechanics has features intended to stop **silent data corruption** that traditional Lustre systems lack. WARPware provides end-to-end checksums and transactional copy-on-write IO operations. CoW eliminates RAID write-holes, and checksums eliminate corruption that has plagued legacy storage solutions at this scale as individual disk sizes get larger and lower-reliability SATA moves into the Enterprise.

Further simplifying the environment, the WARP 41001-M uses the same WARPware feature set available in all WARP Mechanics products: numerous monitoring and reporting tools, iSCSI target mode, NAS protocols, mirroring and replication for disaster recovery and remote site backups, thin provisioning, pNFS, Ceph for object storage, and more.

For a complete list, or to schedule an evaluation, contact your WARP Mechanics authorized reseller.

Application-Centric Performance Optimization

The WARP Mechanics name is synonymous with ultimate power and performance, but not all data access needs to take place at those speeds. The goal of the WARP appliance portfolio is to allow customers to match the capacity to performance ratio appropriate for their application.

The **WARP Mechanics 41001-M** appliance delivers a truly massive amount of scalable solid-state storage ideal for applications which need:



- High capacity
- SSD or even RAM performance for IO
- Optimal cost per terabyte
- Turnkey deployment model
- Simple management interface

It is targeted at customers who serve large amounts of frequently accessed, high-performance data such as big data analytics, life sciences, or engineering or any organization considering a new scale-out storage rollout or adding to an existing HPC environment. It can also be paired with other WARP Mechanics appliances to create a complete solution for all levels of enterprise storage.

WARP Mechanics Scalable All-flash Parallel Storage Appliance Architecture





The WARP 41001-M appliance is constructed using WARP Mechanics hardware platforms, a variety of WARP-tuned and certified SSD modules, WARP software, and fully-embedded partner software.

Intelligent MDS, MGS, and OSS controllers leverage widely deployed, industry-standard hardware with the latest components such as CPU, memory, and IO bus. Each appliance includes up to two controllers. With each expansion enclosure attached, up to 192TB of raw SSD storage capacity is added, depending on workload needs. Each 41001-M head is outfitted with available network interfaces for the intended environment, such as 1Gb, 10Gb, or 40Gb Ethernet, or 40/56Gb InfiniBand and even 100Gbps EDR or Omnipath.

A customer can start with a single 41001-M head, using only a portion of its drive slots: just a few TB of main storage. Yet it can grow to multiple petabytes without architectural changes or even downtime. Simply scale vertically by adding more JBODs connected to the OSS head.

It is also possible to scale horizontally outward by adding OSS heads for additional performance *and* capacity. This can produce a scale-out storage solution delivering more than a Terabyte per second (>100Terabits/s).

This system can be combined with the other WARP Mechanics appliances such as the HybridMatrix and StorageMatrix systems. Using tiering software, data can be migrated between the systems as its business value changes over time.



Copyright © 2016 WARP Mechanics Ltd. All Rights Reserved

WARP Mechanics, WARPware, the WARP Mechanics logo, the WARP Mechanics icon, and SmartStorage System are trademarks of WARP Mechanics Ltd. in the United States and other countries. Other brand, product, or service names may be trademarks or service marks of, and are used to identify, products or services of their respective owners. This document is supplied "AS IS" for information only, without warranty of any kind, expressed or implied. WARP Mechanics reserves the right to change this document at any time, without notice.