



REIMAGINE YOUR STORAGE REDEFINE YOUR POTENTIAL



ABOUT NYRIAD

Nyriad, Inc. is the developer of the UltraIO storage system, an all-new system that combines the processing speed of GPUs and advanced algorithms to deliver unprecedented performance, resiliency, and efficiency. The ground-breaking design enables UltraIO systems to support block storage media and block, file, and object data types in a single system for maximum flexibility. UltraIO systems run on industry-standard hardware, use the highest capacity, most efficient storage media, and simplify storage management to achieve low total cost of ownership. Headquartered in Austin, Texas, Nyriad empowers businesses to grow and adapt their storage to stay competitive in a data-driven world.

Visit us at nyriad.io



RELIABILITY

UltraIO storage system achieves 99.999% reliability utilizing an active passive dual controller with automated failover. The hardware design is highly redundant at all levels and uses high quality components. The system remains available during planned maintenance or administration activity.

PERFORMANCE

UltraIO storage system has been designed to deliver maximum performance from the storage drives under its control. It does this primarily by utilizing all the drives for all the work all the time. This is achieved through its intelligent placement mechanisms and its use of a large persistent cache. By placing blocks intelligently, UltraIO can utilize the entire drive array to provide write operations concurrently. It places the data to optimize both writing and subsequent concurrent reading ensuring maximum performance. It also does write deduplication when high frequency block updates take place. The optimization of data operations provides good performance regardless of workload type.

RESILIENCE

The UltraIO storage system allows customers to choose the amount of resilience (the number of times data can be recreated without loss) that they require. With UltraIO individual blocks are rebuilt when a read request fails. The rebuild takes place on only the first read, and once rebuilt it is written and operates normally. No duplication of data is required to achieve resilience. So long as the number of failures is less than the resilience the customer requested, regardless of the pattern of failures, the system is resilient.

CONSISTENCY

The UltraIO storage system provides consistently high performance. In a test of a 200-array UltraIO system with 10% parity drives, all 20 parity drives were made to fail concurrently with no loss of data and a performance degradation of 5%. Intelligent data placement ensures work is balanced across drives so that throughput is not constrained by load concentration on a few drives. If a drive fails, the system continues to use all other drives for writes making the impact normally unnoticeable. It recreates any block read dynamically, once. When a drive fails it does not require urgent drive replacements and attendant rebuilds which induce extended periods of poor performance in many systems today, especially if larger drives are used. Whole drives may, but typically do not need to be replaced, and not urgently.

ULTRAIOTM ADMINISTRATION GUI

UltraIO Administration GUI is a simple but powerful browser based management tool to simplify ongoing management and administration. The entire administration software is part of the UltraIO controller and does not need an external management server.

INTEGRITY

Many applications today, such as high performance computing and media, want to ensure every bit is correct, every time the block is read. UltraIO storage system stores an integrity hash with every block written. It then integrity checks each block as it is read and ensures it is correct. If not, it recreates it (see resilience).

EFFICIENCY

UltraIO storage system makes every bit the customer buys count. It is highly efficient. It does not use copies, reserve drives for hot spares and minimizes the use of space for garbage collection and storing parity. This delivers a high percentage of usable storage. It utilizes the largest storage drives available. The cost of implementation and administration is low. Drive durability is enhanced by elimination of unnecessary system activities like patrols, writing sequentially and operating with very low write amplification. Overall, it delivers high value for low TCO.

TECHNOLOGY LEVERAGE

UltraIO storage system uses no proprietary hardware. It operates on Linux. It is designed to take advantage of new technologies as they become available. At the core is a combined CPU + GPU architecture. As more advanced CPU's, GPU's, communication technologies, persistent memory technologies, and storage media are developed they will make our systems better, faster, and larger.

SCALABILITY

Our customers can grow vertically (within the same array) from terabytes to multiple petabytes within a single array operated by a single controller. They can grow horizontally by adding additional arrays as targets within their desired clustered file system.

SIMPLIFIED STORAGE MANAGEMENT

Customers can non-disruptively connect UltraIO as a POSIX Compliant Block Target into their existing infrastructure. All of the file and other systems tested so far have successfully connected and operated. Typically, most file systems do not expect a POSIX target that is as fast as, with such large volumes or as resilient as, the UltraIO storage system so file system configurations need to be revised to take advantage of them. Tomorrow the system will support Block File and Object.

RESTful API AND SECURE SHELL

UltraIO provides a REST API (also known as RESTful API) which is an application programming interface (API or web API) that conforms to the constraints of REST architectural style and allows for interaction with UltraIO RESTful web services. The UltraIO Secure Shell is based on a secure protocol for operating storage and other services securely over a network.

ULTRAIO™ SPECIFICATIONS	H1000	H2000
USABLE CAPACITY RANGE	0.9PB to 1.6PB	1.8PB to 3.3PB
READ THROUGHPUT	10.5GB/s	18GB/s
WRITE THROUGHPUT	13.5GB/s	18GB/s
ECC MEMORY	6 x 32GB DIMMs DDR4-2933	
PERSISTENT MEMORY (NVDIMM)	6 x 16GB NVDIMM DDR4-2933	
NETWORK I/O	4x 100GbE/EDR/HDR100 IFB QSFP56 ports (2x NICs)	
MANAGEMENT NETWORK I/O	2 x RJ45 10GBase-T Management Ports 1x RJ45 Dedicated IPMI LAN port	
PROCESSORS	CPU, GPU	
FILE SERVICES SUPPORTED	All POSIX Compatible	
STORAGE PROTOCOLS SUPPORTED	iSCSI/iSER	
AVAILABILITY	Full Redundant System Components, Active/Passive HA	
SUPPORT & SERVICE	12 to 60 Months Available, 24x7	

PHYSICAL ENVIRONMENT	H1000	H2000
MAXIMUM WEIGHT (kg)	175	294
FORM FACTOR	12U	16U
OPERATING TEMPERATURE	10°C ~ 35°C (50°F ~ 95°F)	
OPERATING RELATIVE HUMIDITY	8-90% Non-Condensing	
POWER CONSUMPTION (W)	4,500	6,000
INPUT VOLTAGE (VAC)	200-240VAC 50/60Hz	
INPUT CURRENT (A)*	24	32
ROHS	ROHS Compliant	
CERTIFICATIONS	FCC 47 Part 15, Class A (EMC); UL 62368-1 (Safety)	

* Assuming a minimum power factor of 0.95