



A&E Specification

COLDSTORE 3U | STORAGE SYSTEM

Version 4.4

Contents

1	General	4
2	System Overview & Specification Summary	4
2.1	System Type	4
2.2	System Application	4
2.3	System Technology	4
2.4	Basic System Specification Summary	5
2.5	Disk Type	5
2.6	Disk Reliability	5
2.7	Power Consumption	5
2.8	Additional Specification Summary	6
3	Disk Writing Pattern & Operation	6
3.1	General	6
3.2	LAID & SFS	7
3.3	Mirrored Overlapping-Pair Writing Pattern	7
3.4	Optional Full Mirrored-Pair Writing Pattern (Sequential RAID1 with LAID & SFS)	8
3.5	Zero Data Loss On Writing Disk Failure	8
3.6	Low MTBF with Any Type of SATA Hard Disk Drive	8
3.7	Simple Disk Management	9
3.8	Failed Disk Replacement	9
3.9	Identification of Time-Spans	9
3.10	Indication of Disk Failure	9
4	Power Specifications	10
4.1	Power Consumption	10
4.2	Power Supply	10
4.3	Disk Power-Up	10

5	Alarm Interface	11
5.1	Alarm Inputs	11
5.2	Alarm Outputs	11
5.3	Alarm I/O Connector	11
6	Software Management System	11
6.1	System Scan	12
6.2	System Configuration	12
7	Status Monitor System	13
7.1	System Monitoring	13
7.2	VIEWSTORE System Summary	13
7.3	Monitoring of Multiple Storage Systems	13
8	Electrical, Electronic & Environmental Specifications	14
8.1	Network	14
8.2	System	14
8.3	Capacity	14
8.4	Alarms	14
8.5	Power	15
8.6	Environmental	15
9	Mechanical Specifications	15
9.1	Rack Mount System	15
9.2	Physical Disk Access	15
9.3	Mechanical Construction & Connections	15
9.4	Cooling	16
9.5	Quiet Operation	16
9.6	Key System Dimensions	16
10	Certifications & Environmental Directives	18
11	Copyright & Trademarks	18

1 General

All equipment and materials used shall be standard components that are regularly manufactured and used in the manufacturer's system.

All systems and components shall have been thoroughly tested and proven in actual use.

All systems and components including disks shall be provided with a three-year warranty. Site-specific extended warranty and support packages shall be available from the manufacturer through authorised system integrators.

2 System Overview & Specification Summary

2.1 System Type

The System shall be a network attached hard disk array specifically designed for video surveillance storage. The System shall be designed specifically to increase the reliability of hard disks by mitigating the factors which typically cause them to fail and shall thus enable the use of very low-cost drives. Further, the System shall be designed specifically to be ideal for very high capacity (using the largest disks available) and for very long file retention times (switching off disks not being accessed) and for extremely low power consumption.

2.2 System Application

The combination of these factors (use of low-cost disk types, very high capacity disks and very low power consumption) shall reduce initial capital costs and shall reduce running costs. The System shall enable long-term video, audio and meta-data file retention in a highly practical and low-cost manner.

The System shall require no specific IT storage system skills to configure and operate and shall require little or no maintenance in normal operation.

2.3 System Technology

The System shall be based on LAID (Linear Array of Idle Disks) and SFS (Sequential Filing System) technology with sequential recording, noting that this is an entirely different approach from conventional and RAID5 or RAID6 systems. The System shall normally operate in a sequential mirrored overlapping-pair disk writing pattern but shall also be capable of encompassing LAID/SFS and sequential RAID1 simultaneously, using a sequential full mirrored-pair recording method as a configuration option.

[The difference between mirrored overlapping-pair and mirrored non-overlapping-pair disk writing is specified in more detail in the following Section 3 – Disk Handling and Operation].

2.4 Basic System Specification Summary

The storage system unit shall be a 15-bay, 3U, 19" rack-mountable, network-attached hard disk storage array with dual-gigabit Ethernet ports for network communication, and a front-mounted control panel with LCD display allowing optional local control of the system in the rack, in addition to network-based control. The System shall use an embedded Linux Operating System and COLDSTORE system firmware stored on a removable solid-state memory medium (SD Card). The Linux OS and COLDSTORE firmware shall be upgradeable via the network connection. In addition, the OS and firmware shall be upgradeable or replaceable by removing and replacing the SD Card. The SD card shall be accessible from the rear without removing the system from the rack.

2.5 Disk Type

The System shall be capable of using any type of 3.5" SATA hard drive of any capacity, including the latest 14TB drives and also including drive-managed archive-style SMR (Shingled Magnetic Recording) drives, without compromising performance.

The System shall also be capable of accepting higher capacities of SATA drives (to 20TB and beyond) without modification, as these larger drives become available.

2.6 Disk Reliability

The System shall provide a very high degree of data integrity through a number of methods which are described in more detail in Section 3 – Disk Handling & Operation. These methods result in very high disk reliability, through lower disk operating temperatures, lower disk vibration, lower disk duty cycles, limited requirement for forced cooling (and therefore lower vibration from cooling fans), low-power processors and use of sequential storage techniques to exploit the natural sequential nature of real-time video data.

Only writing disks and disks which are being accessed for playback shall be powered up. All other disks shall be switched off, isolating them from power spikes, surges and "brown outs", all which can cause damage to hard disks if not properly filtered in the power system of the building or facility.

2.7 Power Consumption

In normal operation, the System power consumption shall be less than 60 watts, irrespective of the number and type of hard disk drives fitted. The system power supply shall be a dual-module, dual 300 watt hot-swappable PSU type and shall be capable of running all of the disks simultaneously, as and when required. Even in this mode, the power consumption shall be less than 120 watts. Detailed specifications are provided in Section 4 – Power Specifications.

Upon start-up, to avoid an initial high spike in power supply current, disks will be started up and checked in series, rather than in parallel. The same start-up-in-turn process will happen in normal operation when the system receives a command to switch on all disks for search purposes. Whilst the power supply shall be capable of powering all disks at once, this technique shall avoid stress on the PSU and increase the lifetime of the PSU components.

2.8 Additional Specification Summary

The storage system shall have a physical alarm interface (input and output) in addition to alarm and status monitoring through the network. Detailed specifications are provided in Section 5 – Alarm Interface.

The System shall have a Windows client software application for system configuration and management. Detailed specifications are provided in Section 6 – Software Management System.

The storage system shall support authenticated remote status monitoring, allowing the status of multiple storage systems to be monitored in a very simple, easy to understand and user-friendly manner. Further details are provided in Section 7 – Status Monitor System.

Sections 8 and 9 cover the Electrical and Electronic Specifications and Mechanical Specifications in detail respectively.

The System shall be RoHS, WEEE, CE and FCC compliant and shall conform to all appropriate safety and environmental standards as fully defined in Section 10 – Certifications and Environmental Directives.

3 Disk Writing Pattern & Operation

3.1 General

The Storage System shall be designed to be ideal for long file retention times and shall be able to use the highest capacity disks available in an optimum manner. The System shall not lose any data due to disk failure during the writing process. There shall be no data rebuilding process as it shall not be necessary.

The storage system shall be specifically designed to minimise the factors which cause hard disks to fail. These factors are generally acknowledged to be: temperature, vibration and wear.

Thus, the system shall operate the disks at lower temperatures than normal, shall minimise vibration generation and transmission (from other disks and fans) and shall reduce the normal wear rate by significantly reducing the workload when the disk is in normal operation and by significantly reducing the

duty cycle of each disk. This handling of the disk drives shall greatly increase individual disk reliability and considerably extend the expected working life of the disks and the array.

The storage system shall be able to achieve very high levels of disk reliability, even from the low-cost desktop-type hard disk drives, through the following processes:

- 1) Reduction in disk operating temperature
- 2) Elimination of vibration
- 3) Reduction of wear

by using the disk operation and writing pattern methods described in the subsections below.

3.2 LAID & SFS

The system shall operate on the Linear Array of Idle Disks (LAID) principle, where disks are written to sequentially across the array and where all non-writing, non-reading disks are switched off. The system shall use a Sequential Filing System (SFS) to effectively eliminate disk vibration from rapid read-write head movements during the write cycle (although shall not restrict such movements for random read access).

The uses of SFS shall increase the sustainable writing throughput performance for video and other sequentially-generated data and reduce the heat dissipation within the disk in normal operation.

The use of the LAID principle shall result in all but two drives in the array being normally powered off. Drives shall be automatically powered up when required for reading purposes without any special commands from the recording system (host). The use of LAID, with only two disks normally in operation at a time (ignoring playback) shall result the power-on time for each disk being less than 15% of the normal array operation time (assuming a fully-populated 15-bay array).

All data shall be written strictly sequentially to the disk via the Sequential Filing System. The use of SFS shall result in the absolute minimum of movement of the read-write heads of the disk, thus minimising or eliminating vibration from this source. SFS shall allow the use of the largest drive capacities available, without modification, including drive-managed SMR (Shingled-Magnetic Recording) disk drives, normally designed for archive purposes. (Host-managed SMR drives shall not be supported).

3.3 Mirrored Overlapping-Pair Writing Pattern

The system shall write to disks sequentially in a mirrored overlapping-pair writing pattern to provide full redundancy during the writing process without requiring twice the disk capacity overall. Specifically, the disk slot writing pattern shall follow the following mirrored pair sequence: 1+2, then 2+3, then 3+4, then 4+5 and so on. Upon reaching the end of the array, the sequence shall wrap around the array as follows: 14+15, then 15+1, then 1+2 and so on.

Any disk bays which are not populated, or which contain a failed disk shall be ignored and the next available disk in the sequence used instead.

3.4 Optional Full Mirrored-Pair Writing Pattern (Sequential RAID1 with LAID & SFS)

The system shall also be capable of writing to disk pairs in a full (non-overlapping) mirrored pair pattern. This non-overlapping writing pattern shall be functionally equivalent to the sequential use of multiple RAID1 mirrored pairs and can be considered as a sequential RAID1 system. In this writing mode the System shall still operate the LAID and SFS features to achieve the required very high levels of disk reliability features as previously described. In this mode the writing sequence shall be 1+2, 3+4, 5+6 and so on, leaving fully mirrored data in place at all times. In this mode the 15th disk slot shall be ignored.

3.5 Zero Data Loss On Writing Disk Failure

No data shall be lost during the writing process in the event of a disk failure. All writing processes shall be mirrored; thus, the System shall have data redundancy at all times during the writing process. The System shall be configurable to respond to the event of a writing disk failure in one of two ways:

- a) To immediately step to the next pair of disks in the sequence and re-commence writing (buffering data until writing re-commences on the new pair)
- b) To complete writing to the remaining space on the current drive, and only then to step to the next pair of disks in the sequence.

Option (a) shall provide mirrored-writing cover at all times, whilst option (b) shall make use of the remaining disk space on the surviving disk of the pair. It shall be acceptable that the use of option (b) shall incur a temporary and extremely low risk of loss of the second drive of the pair, resulting in a partial loss of data (on that disk only).

The LAID and mirrored-writing processes shall result in a robust system that, unlike RAID5 and RAID6, there shall be no mechanism by which the entire data array can be lost. In the event of multiple drive failures, the archive capacity of the System shall necessarily be reduced, but the recorded data shall naturally be sacrificed in order of age. That is, the oldest data shall be deleted first though the resultant process of earlier-than-normal over-writing of disks with new data, should the number of functional disks in the array be reduced.

See section 3.8 for details of failed disk replacement.

3.6 Low MTBF with Any Type of SATA Hard Disk Drive

The lower-temperature, zero vibration and minimal duty cycle of the LAID and SFS system described in above shall enable the use of standard SATA hard drives of any make, model or capacity, including those designed for light desktop use. The inherent low-stress approach of the LAID and SFS system combined shall allow even the lowest cost SATA desktop standard disk drives to attain MTTF or MTBF levels to

exceed the disk manufacturer's quoted numbers by a factor of 3 or greater. In particular, the MTBF shall exceed the average real-world reliability figures of normal RAID-level SATA or SAS drives in typical RAID5 or RAID6 storage arrays by a factor of 5 or greater.

The storage array shall be able to cope with mixed disk types (manufacturers, models, speeds and capacities) without error. For any pair of disk drives, the system shall operate to the "lowest common denominator" of speed and capacity. It shall be recommended that disks of equal capacity be used for most efficient use of disk space and that disks of equal speed are used for optimum performance, but this is not essential for correct operation.

3.7 Simple Disk Management

The system shall provide simple and straightforward disk management, including the extraction of recorded disks, insertion of new disks, replacement of recorded disks, and expansion of capacity both by insertion of additional disks if free slots are available and by upgrade of individual disk capacities (swapping lower capacity disks for higher-capacity disks). All of these operations shall not require disk configuration, rebuild or anything other than simple disk extraction and insertion.

3.8 Failed Disk Replacement

The replacement of a failed disk shall be extremely simple and easily carried out by non-skilled personnel. A failed disk shall be indicated by a red LED above that disk slot. Replacement of the disk shall be as follows: The disk cradle release mechanism shall be unlocked with the key provided (if it has been locked) and then the disk extracted from the array by pulling down on the cradle release handle. A spare disk and cradle shall be inserted into the array and the cradle handle pushed to close. If required, the cradle lock can be locked in place with the key provided.

This completes the disk replacement procedure. Note that no disk configuration, system management or other processes are required. The COLDSTORE system shall automatically recognise the new disk, verify the capacity and incorporate the new disk into the array. No formatting is required as the SFS shall format the disk on-the-fly as it writes new data to the disk during the next full write cycle.

3.9 Identification of Time-Spans

It shall be possible, through the front panel controls and display or via the software management tool, to ascertain the time-span of the full array and the time-span of each and every individual disk. Conversely, it shall be possible via these same methods to identify the specific physical location of any specified recording, as defined by the time at which it was recorded.

3.10 Indication of Disk Failure

In the event of a disk failure, a red LED shall illuminate over the drive bay containing the failed disk, allowing simple identification of the disk to be replaced. As a user-configurable option, one of the four relay outputs shall be triggered in the event of a disk failure. Disk failure and other health monitoring

information shall be provided by the VIEWSTORE storage monitoring system, described in a separate A&E specification document.

4 Power Specifications

4.1 Power Consumption

The System shall consume less than 60 watts in normal operation (recording and playback), irrespective of disk types, capacities and number fitted. The achievable average power consumption with low-power disks shall be as low as 45 watts.

The System peak power consumption shall be on initial system power-up from a cold start only and shall peak at less than 200 watts for less than 60 seconds.

4.2 Power Supply

The system shall have a dual, redundant, hot-swappable power supply module with independent power connections. Each PSU module shall be rated at 300 watts independently. Each PSU module shall be capable of delivering sufficient power to run all the hard disks simultaneously (typically only 120 watts maximum).

Failure of either power supply module shall result in an audible alarm and a visible red warning LED indicating which of the two modules have failed. Removal of the failed module shall cancel the alarm and the system shall continue to operate normally, without interruption, on the remaining PSU module. A replacement for the failed module may be inserted at any time and shall automatically restore the PSU system redundancy.

4.3 Disk Power-Up

The System shall, through the normal cold-start boot-up process, query all disk slots to ascertain disk presence, size and status. In order to avoid an initial high spike in power supply current, disks shall be started up and checked in series, rather than in parallel. (All spinning disk types require a higher start-up current through the drive motor in order to spin-up the disk platters. Motor current becomes stable and constant after the platters reach their normal operating RPM).

During normal operation, with only two disks active for the writing process, the system shall be capable of receiving a command to activate all disks for broad search purposes with fast responses. In this case, the same start-up-in-turn process shall occur for all switched-off disks when this command is received.

5 Alarm Interface

The System shall have four alarm inputs and four relay-controlled alarm outputs.

5.1 Alarm Inputs

The alarm inputs shall be software-configurable to allow specific functions to operate on alarm. These may include:

- a) Controlled shut-down
- b) Signal from external UPS to indicate external power loss and operation on UPS power only
- c) Other functions as may be required for specific integrations or installations

5.2 Alarm Outputs

The relay outputs shall be software-controlled and configurable to enable signalling of the following functions:

- a) Hard disk failure
- b) Power status of the System (open on power loss, closed on power present)
- c) Hard disk removed unexpectedly
- d) Target archive period not achieved
- d) Other functions as may be required for specific integrations or installations

5.3 Alarm I/O Connector

All alarm connections shall be via a multi-pin removable connector with screw terminals for simple connection of external wiring.

6 Software Management System

The System shall be provided with a separate Windows application for managing and configuring the system over the network. This Manager Application (SITESCPE) shall be Windows 10 (and later) compatible.

6.1 System Scan

The Manager Application shall automatically search the network and display a list of Storage Systems available on the network and shall list the following information:

- a) IP Address (Ethernet Port 1)
- b) MAC Address (Ethernet Port 1)
- c) System Name (if assigned)

of each system. This list shall be visible at all times when the Manager application is running. It shall also be possible to add Storage Systems to the list manually by entering the IP address of the System.

6.2 System Configuration

Upon selecting one of the Systems listed, the Windows Client application shall provide the means to configure all system parameters and access various system status and diagnostics information including, but not limited to:

- 1) Array & Disk Status
- 2) System Status
- 3) System Configuration
- 4) System Diagnostics

Access shall be password protected on a per-system basis. Access shall allow the following parameters to be read and or configured:

- a) Array name
- b) IP Address 1 (inc. netmask & gateway)
- c) IP Address 2 (inc. netmask & gateway)
- d) NTP Server IP address 1
- e) NTP Server IP address 2
- f) NTP Server IP address 3
- g) File Retention Limits (days)
- h) Time Zone
- i) Array mode (mirrored overlapping pair, full mirrored pair)
- j) Disk failure action
- k) System firmware version (including upgrade option)
- l) System array status
- m) Individual disk status (including capacity)
- n) Disk array usage

- o) Time span of individual disks
- p) Array time span

7 Status Monitor System

7.1 System Monitoring

As a modern, simpler and more user-friendly approach, the System shall not support SNMP but instead shall support remote visual real-time monitoring through the VIEWSTORE status monitoring system. See VIEWSTORE A&E Specification document for full details.

7.2 VIEWSTORE System Summary

The VIEWSTORE system is specified separately, but a short summary is provided here. VIEWSTORE is a system application running on a micro-PC and drives a full HD monitor, providing both graphical and tabular display options for multiple COLDSTORE storage system units. This dedicated display system is designed primarily for control room wall display but may be used anywhere for disk array status monitoring purposes.

The graphical display provides a clear and obvious real-time status of all System storage arrays and individual disks, highlighting any failed disks or arrays which are not responding. This monitoring system is extremely simple to use and understand, requiring no skills and or specific training to use.

7.3 Monitoring of Multiple Storage Systems

The VIEWSTORE system shall be available as an option from the manufacturer and shall provide single-screen monitoring for up to 40 COLDSTORE 3U systems whilst simultaneously incorporating monitoring of up to 600 disk drives in these 40 systems.

It shall be possible to employ multiple VIEWSTORE systems to monitor installations larger than 40 storage units. It shall also be possible to use multiple VIEWSTORE systems in order to provide monitoring of the same storage systems in multiple locations.

8 Electrical, Electronic & Environmental Specifications

8.1 Network

Network Interface:	Dual Gigabit Ethernet
Data throughput:	Up to 400Mbit/sec (unlimited scalability with multiple storage array units)

8.2 System

Operating System:	Embedded Linux on solid state memory.
System integration:	Direct from 3rd party client application (SDK & network protocols available).
Configuration:	Over LAN via Veracity client application or 3rd party client application.
Management:	By front panel, or over LAN via Veracity client app. or 3rd party client app.
Time synchronization:	Via NTP (essential)

8.3 Capacity

Drive bays:	15 x 3.5" (lockable drive trays included)
Maximum capacities:	
Raw	210TB (with 14TB disks)
Effective	196TB (Normal COLDSTORE mirrored overlapping-pair mode) 98TB (Sequential full mirrored-pair mode) (14 disks)

8.4 Alarms

Alarm relays:	4 software configurable e.g. disk fail, disk inserted, disk extracted, PSU fail, etc. (See section 5 for full detail).
Alarm inputs:	4 software configurable e.g. UPS active, power down, start fans, etc. (See section 5 for full detail).

8.5 Power

Power supply:	Dual hot-swap PSU
Power consumption:	Maximum 60 watts average
Power supply rating:	300W (See actual power consumption above)

8.6 Environmental

Operating temp:	5 degrees C to 50 degrees C (14 degrees F to 122 degrees F)
Relative humidity:	Up to 85%, non-condensing.

9 Mechanical Specifications

9.1 Rack Mount System

The System shall be 19" rack mountable and 3U in height. The depth of the unit in the rack shall be no more than 450mm.

The System shall have 15 hard disk drive bays with lockable, removable, disk cradles. The disk cradles shall incorporate light pipes to allow display of LED indicators from the disk connector backplane of the system.

9.2 Physical Disk Access

There shall be a hinged lid with fastener to provide front access to the disk drive cradles. The hinged lid shall be ventilated to allow airflow into the System and shall incorporate a simple removable, washable dust filter material into the lid assembly.

9.3 Mechanical Construction & Connections

There shall be no SATA cables connecting the disks to the backplane, nor shall there be any SATA cables connecting the backplane to the main CPU board. The main CPU board shall be directly connected to the backplane via a high-quality multi-pin connector system. Both main CPU board and backplane shall

be secured bolted to the chassis. The resulting assembly shall therefore largely eliminate the use of SATA cables.

9.4 Cooling

The main CPU board shall not require any cooling fans or on the processor or other components. The heat dissipation of the main board shall be so low such that only a low-volume airflow will be required from two low-vibration fans mounted at the rear of the system.

There shall be two software-controlled low-vibration fans at the rear of the unit to provide additional airflow when required. When the System is operating below a threshold temperature, the fans shall stop, and shall switch on above the threshold temperature. The temperature shall be measured by a minimum of four temperature sensors built-into the main CPU board system.

Each power supply module shall have a single temperature-controlled variable speed fan. Immediately after a cold start, the PSU fan speed shall automatically decrease to a low level, unless the ambient operating temperature is above 25 degrees C.

The two fans in the main chassis plus those in the power supply modules shall be the only fans in the system. Thus, there shall be no more than four fans in total.

9.5 Quiet Operation

The system, including fans, shall operate with very low noise levels. In normal operation, the System shall have a sound pressure level of less than 45dB at 1 metre.

9.6 Key System Dimensions

The System shall conform to the following dimensional specifications:

Width: 19" rack mount (chassis 438mm wide, maximum width (in mounts) 482mm)




Height: 3U (133.3mm)

Depth: 491.8 mm



COLDSTORE

 SURVEILLANCE STORAGE SYSTEM



Dimensioned Drawings Dimensions in mm	Product Name COLDSTORE 3U
Veracity Product Codes CSTORE15-3U	

V3.1

10 Certifications & Environmental Directives

The Storage System shall comply to standard EMC Directives defined for industrial electrical equipment. Specifically, the System shall include CE (EEA) and FCC (USA) certifications.

The System shall also be RoHS, WEEE and REACH compliant.

The System shall be designed and manufactured under both ISO 9001:2008 and ISO 14001 standards for quality management and environmental principles.

11 Copyright & Trademarks

This document is © Veracity UK Ltd 2018

COLDSTORE, LAID, SFS VIEWSTORE and SITESCAPE are all trademarks or registered trademarks of Veracity UK Ltd.