



## Technology Overview

- Video Surveillance Storage
- Serverless Computing
- Pivot3 vSTAC Watch

February 2012

# Table of Contents

- Introduction..... 4
  - Video Surveillance Storage Requirements..... 4
  - Pivot3 Serverless Computing™ ..... 5
  - Pivot3 vSTAC Watch Appliances ..... 5
- 1. Virtual Server Environments..... 7
  - Virtual Server Resources..... 7
  - Virtual Machine I/O ..... 7
  - One VMS per vSTAC Watch Appliance..... 7
  - Integration with Existing NVRs ..... 7
  - Integration with Existing DVRs ..... 7
- 2. The Pivot3 vSTAC OS..... 8
  - Logical Volume Management ..... 8
  - Initiator Management ..... 9
  - Distributed Data Protection..... 9
  - RAID Algorithms..... 9
  - RAID Protection ..... 10
  - RAID Efficiency..... 12
  - Disk Groups..... 12
  - Changing and Mixing RAID Protection..... 12
  - Dynamic Expansion: Adding Physical Capacity to an Existing vSTAC Array ..... 12
  - Instant Volume Initialization..... 12
  - Video Performance With Unified Flash Cache..... 12
  - Write Process..... 13
  - Virtual Global Sparing ..... 13
  - Parallel Rebuild Innovation..... 13
  - Priority Rebuilds by Volume ..... 13
  - Background Verification ..... 14
  - Predictive Sparing ..... 14



Automatic Load Balancing ..... 14

SNMP Support ..... 15

Patented DDA Pipelining ..... 16

Performance Monitoring ..... 16

4. VMS Failover Support ..... 16

    Pivot3 VM Failover™ ..... 16

    Failover Software Packages from VMS Providers ..... 17

    Virtual Machine Failover using VMware HA™ ..... 17

5. Pivot3 vSTAC Appliances..... 17

    vSTAC Watch Appliances ..... 17

    vSTAC Watch Platform Components ..... 18

    vSTAC Data Appliances ..... 18

Summary..... 18

## Introduction

Video surveillance is a powerful tool for applications ranging from international border security to retail buying pattern analysis. Dramatic innovations in camera technologies and compression algorithms make it possible to economically capture high-quality content but new storage technologies are required to meet the performance, reliability and budget needs of the market.

Pivot3 has introduced an innovative line of storage products that are purpose-built for the surveillance market. This paper explains the technical features inherent in the Pivot3 design for technical architects, system integrators and end customers looking to better store, protect and quickly access the growing amount of vital video data.

### Video Surveillance Storage Requirements

Conventional IT storage systems are designed to meet conventional IT workloads. Yet video workloads differ in nearly every dimension from the data assumptions implicit in IT storage systems. This disparity provides an opportunity to re-invent storage systems that are better suited to the tasks at hand.

***Design Objective of Conventional IT Storage:*** The design goal of a conventional IT storage system is to optimize low-latency performance for small, random blocks of data that are 80% read and 20% write activity. Performance is measured in optimization of small block random IO's per second and every feature follows from this initial starting assumption:

- Head-end system with a two fixed RAID controllers and fixed SAN connections
- Multi-path connections are primarily used for failover
- Storage tiers are effective because capacity needs are small
- Primary data is never overwritten and backup to disk and tape are easily supported
- System reliability is measured in component terms like "Five 9's"
- Data is backed up and replicated offsite for disaster recovery
- Storage requirements are secondary to server considerations.

***New Design Objectives of Video Surveillance Storage Systems:*** A video surveillance storage system is optimized for high-bandwidth ingest of 99% write activity comprised of large block random data. Aggregate bandwidth is the key performance metric and the feature set is markedly different:

- Scale-out implementation with parallel, scalable controllers and SAN connections
- Multi-path connections are primarily used to aggregate bandwidth
- Parallel cache designs are required to keep up with the near-sequential data flow
- Massive capacity and bandwidth needs render snapshot, backup, and replication technologies useless
- Reliability requires no single point of failure at the system-level
- Data is too large to back up and too bandwidth intensive to replicate off-site for disaster recovery
- Storage requirements dominate server needs

The following pages describe how Pivot3 uniquely meets the requirements of this growing and vibrant market.

### Pivot3 Serverless Computing™

Pivot3 Serverless Computing was introduced in 2007 and within three years gained the top market share for IP SANs used in video surveillance storage. In 2012, Serverless Computing again led the IP SAN category that is projected to grow 39.3% annually according to IMS Research<sup>1</sup>.

Pivot3 Serverless Computing is made possible by a new software storage operating system from Pivot3 called the vSTAC OS. The vSTAC OS provides a scale-out storage system designed for the performance and reliability needs of video surveillance storage systems. The vSTAC OS integrates server virtualization technology from VMware® to eliminate the need to purchase, power, cool or cluster physical archive servers. The vSTAC OS is designed for:

- **High-bandwidth Performance:** The vSTAC OS is a massively parallel storage system that aggregates bandwidth across storage network ports, RAID controllers, flash write caches and spinning drives.
- **System-level Reliability:** In addition to redundant components and RAID protection for drives, the vSTAC OS provides complete appliance level protection and is the first and only storage system to provide fail-over protection for virtual servers as well as storage.
- **Distributed Load-balancing:** Advanced motion-based compression algorithms such as H.264 can generate large capacity and performance spikes. The vSTAC OS normalizes data ingest rate spikes across all available resources to simplify configuration, minimize over-provisioning, and eliminate operator intervention.

### Pivot3 vSTAC Watch Appliances

The Pivot3 vSTAC Watch appliances are the third generation of Pivot3 appliances shipped with the vSTAC OS. The appliance model of storage introduces some unique advantages of its own including:

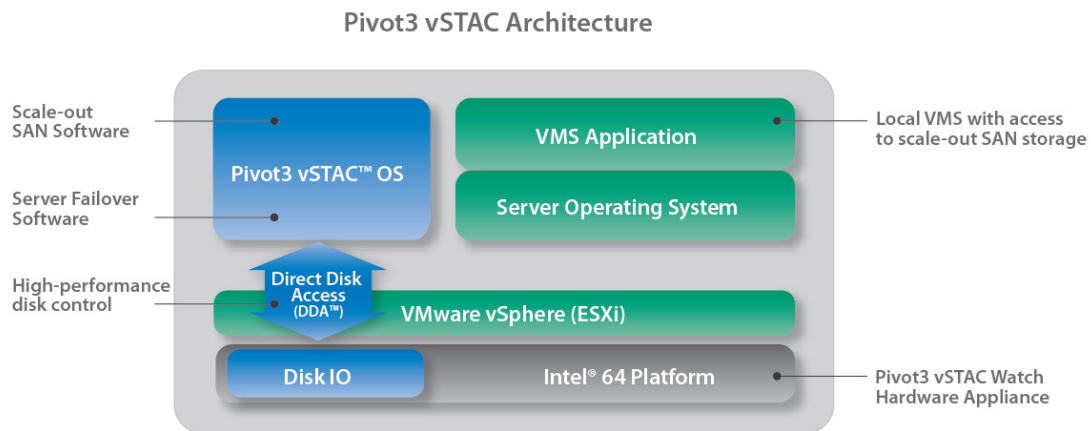
- **Scalability** refers to linearly scaling performance along with capacity and cost. It is typical to add cameras or lengthen retention time to an existing system so the ability to scale capacity and performance quickly and without disrupting ongoing video capture is essential. Recording tens, hundreds or even thousands of camera streams onto disk drives is a core requirement of a video surveillance storage system.
- **Simplicity** is essential for surveillance administrators who are experts in video security rather than complex computing systems. An appliance model offering both server and storage functionality is ideally suited to the surveillance user familiar with server, Ethernet and Microsoft standards. Surveillance directors similarly prefer self-healing systems to the complex, costly clustering and replication techniques used in conventional IT products.
- **Savings** since video surveillance systems are generally cost centers and the high cost of storage often leads to cutbacks in frame rates, resolution, or retention times to meet budget constraints. This paper describes how Serverless Computing replaces physical servers with virtual servers to save on acquisition costs, power, cooling, and rack space.

<sup>1</sup> IMS Research; Enterprise and IP Storage Used in Video Surveillance, 2012

## Pivot3 vSTAC Watch Architecture

The Pivot3 vSTAC architecture combines high performance compute and storage resources in scale-out appliances that are “stacked” together. It is the first and only scale-out architecture that deploys a storage area network (SAN) first and then layers in local virtual servers to run VMS applications. The combined server/storage platform is both a high-availability SAN storage solution with no single point of failure and a high-performance, high-availability server solution that takes advantage of commodity hardware economics.

There are five major components in the Pivot3 vSTAC Watch architecture:



Each of these elements is described in the following chapters of this report:

1. Virtual Server Environment
2. vSTAC Operating Environment
3. Direct Disk Access
4. VMS Failover™
5. vSTAC Appliances

These details are provided to help you understand and appreciate the elements of the design that uniquely solve the need of the surveillance market.

# 1. Virtual Server Environments

Each Pivot3 vSTAC Watch provides local server resources for one VMS application running on a Windows or Linux operating system running as a VMware vSphere virtual server.

## Virtual Server Resources

Each VMS application running on a vSTAC Watch has dedicated local server resources from the vSTAC Watch including dedicated CPU cores, RAM resources, and LAN ports. From a storage perspective, the virtual server acts as an iSCSI initiator with access to the shared storage across vSTAC Watch appliances. The shared storage appears as a multi-ported iSCSI target.

## Virtual Machine I/O

iSCSI I/O begins with a single iSCSI session between the iSCSI initiator in the host and each Pivot3 logical volume accessible by the host. The iSCSI sessions utilize a purely virtual network interface card (NIC) that interfaces directly to the vSTAC OS running on the Pivot3 appliance hosting the VM. This network path is immune to failure points common in physical networks (cables, switches).

A host I/O destined for a Pivot3 logical volume is sent by the iSCSI initiator to the iSCSI target for that logical volume through the virtual NIC. The vSTAC OS examines the logical block address (LBA) associated with the command to determine which Pivot3 vSTAC Watch appliance the data should be written to or read from. If the I/O can be processed on the local appliance, the request is serviced immediately and the data never traverses a physical network cable. Write or read operations destined for another vSTAC Watch appliance in the vSTAC Array are intelligently routed to the appropriate Pivot3 appliance. This intelligent I/O direction provides a fault-tolerant and load-balanced storage network over Ethernet that forms the backbone of the Pivot3 vSTAC. Once the data transfer is complete, status for the I/O is returned to the host over the virtual NIC.

## One VMS per vSTAC Watch Appliance

Pivot3 recommends running one VMS application per vSTAC Watch appliance because of the unique requirements of video cameras and VMS software. VMS applications are typically both CPU and IO bound, so that dedicated resources are preferred for each VMS instance. In addition, newer cameras with motion-sensitive compression techniques such as H.264 often generate large variations in data rates that are best served by allocating dedicated server resources. Pivot3 Serverless Computing provides these dedicated resources for IO and CPU loads during both steady state and motion-intensive times by deploying just one VMS application per vSTAC Watch appliance to ensure consistent performance.

## Integration with Existing NVRs

Pivot3 SAN storage capacity is not only for hosted servers within the vSTAC Array. Capacity is presented as a standard iSCSI target and can be accessed by any external physical servers acting as iSCSI initiators. For video surveillance environments with existing network video recorders (NVRs), this allows a smooth upgrade path over time.

## Integration with Existing DVRs

A virtual machine on a vSTAC Watch can also be configured as a network attached storage (NAS) head-end using, for example, the Microsoft Windows Storage Server software. This allows part or all of the capacity of the system to be accessed as a mapped network drive to expand and protect capacity of installed desktop video recorders (DVRs) in the field.

## 2. The Pivot3 vSTAC OS

The Pivot3 vSTAC (virtual Storage and Compute) Operating System runs on each Pivot3 vSTAC Watch to deliver scale-out storage across the Pivot3 appliances. The vSTAC OS provides logical volume management, distributed data protection and automatic load balancing across appliances for ease of management, high availability and high performance.

### Logical Volume Management

The vSTAC OS virtualizes physical disks and appliances in a Pivot3 vSTAC Array so capacity can be managed logically beyond the physical limits of each appliance. Reliability also improves because volume access is not disrupted by physical hardware component failures that can include an entire appliance.

Pivot3 appliances are discovered and managed by vSTAC Manager Software that runs on any PC or server connected to the vSTAC Array. Physical appliances located on the same local subnet can be selected and assigned to one or more Pivot3 vSTAC Arrays. Multiple Pivot3 vSTAC Arrays can be managed with one instance of vSTAC Manager.

**vSTAC Watch Appliances are assigned to a vSTAC Array using vSTAC Manager**

Name	Status	OS Version
VHF2-65R8HM1	Normal	pv00.5.1.1
VHF2-65R8HM2	Normal	pv00.5.1.1
VHF2-65R8HM3	Normal	pv00.5.1.1
VHF2-65R8HM4	Normal	pv00.5.1.1
VHF2-65R8HM5	Normal	pv00.5.1.1
VHF2-65R8HM6	Normal	pv00.5.1.1
VHF2-65R8HM7	Normal	pv00.5.1.1
VHF2-65R8HM8	Normal	pv00.5.1.1
VHF2-65R8HM9	Normal	pv00.5.1.1
VHF2-65R8HMA	Normal	pv00.5.1.1
VHF2-65R8HMB	Normal	pv00.5.1.1
VHF2-65R8HMC	Normal	pv00.5.1.1

Modify vSTAC Array membership Completed.

### Capacity Management

The aggregate capacity of the Watch appliances in vSTAC Array can then be parceled into logical volumes using vSTAC Manager. Attributes for each logical volume, such as RAID protection, name, rebuild priority, and access

control are set by volume without requiring knowledge of the underlying physical hardware. Capacity can be physically and logically added to both the vSTAC Array and to existing logical volumes. Capacity expansion is dynamic and does not interfere with data storage or retrieval processes. The aggregate capacity of the Pivot3 appliances is presented as a multi-ported iSCSI target.

The screenshot displays the Pivot3 vSTAC Manager web interface. On the left, a tree view shows the hierarchy: vSTAC Domain > Pivot3-Array > SATA > Volume-1 and Volume-2. The main panel is divided into several sections:
 

- Status:** Shows 'Normal'.
- Quick Links:** Includes a 'Create Logical Volume' button.
- Settings:** Lists 'Total Capacity' as 322.338 TB, 'Available Capacity' as 201.53 TB, and 'Drive Sparing' as 'Supported'.
- Capacity Detail:** A 3D pie chart shows 'Used' at 37% and 'Available' at 63%.
- Volumes:** A table titled 'Logical Volumes' provides details for Volume-1 and Volume-2.

 A status bar at the bottom indicates 'Create Logical Volume completed.'

Name	Status	Usable Capacity	RAID	Sparing	Allocated
Volume-1	Normal	14.648 TB	3.107 TB	10.52 GB	17.765 TB
Volume-2	Normal	14.648 TB	3.107 TB	10.52 GB	17.765 TB
<b>Total Allocated</b>					<b>35.531 TB</b>

**Logical Volumes are created from the aggregate capacity in the vSTAC Array**

### Initiator Management

Volume access is controlled by a host identifier. This can be an initiator name, an initiator alias, or a CHAP name. Administrator can also set attributes by initiator for Read-Write, Read-Only or no access to volumes.

Initiator names/logins have an Access Control List (ACL) that allows the administrator to define Read-Write, Read-Only or no access to volumes.

### Distributed Data Protection

The key elements of the Pivot3 vSTAC OS are distributed RAID algorithms, parallel disk write cache, Unified Flash Cache, virtual global sparing, parallel rebuilding of failed drives, priority rebuilding of volumes, continuous background verification and predictive sparing.

### RAID Algorithms

The vSTAC OS distributes data and parity across Pivot3 vSTAC Watch appliances so that data is efficiently protected against component failures. There is no need to create physical disk or RAID sets as you would with a traditional

RAID system. Rather, disks are treated as raw capacity and the RAID function is implemented at the volume level. The normally burdensome management tasks of defining RAID groups and partitioning volumes, which are associated with traditional RAID devices, are not necessary with a Pivot3 vSTAC Array.


### RAID Protection

RAID protection levels are provided to meet the data protection goals of each application. Each RAID level is set by volume and precise usable capacity is shown in the vSTAC Manager user interface. Usable capacity is also shown in the technical specification sheets available at [www.pivot3.com](http://www.pivot3.com).

### RAID 5e

RAID 5e provides more protection than traditional RAID 5 systems since it protects data even in the case that an entire appliance fails and all its associated drives are off-line. It is the most cost effective RAID level.

- Enhanced distributed parity protects against either one disk drive failure or one appliance failure. The “e” indicates enhanced RAID 5 for vSTAC Arrays of three or more appliances where data is also protected even if an entire appliance fails.
- Data is striped across each appliance in the vSTAC Array and protected by one level of parity. Distributed parity is interleaved across each appliance in the vSTAC Array so that data stored in each appliance is protected by parity in another appliance.
- Capacity required for RAID 5e protection is roughly one appliance per vSTAC Array since one copy of parity is distributed in the vSTAC Array.



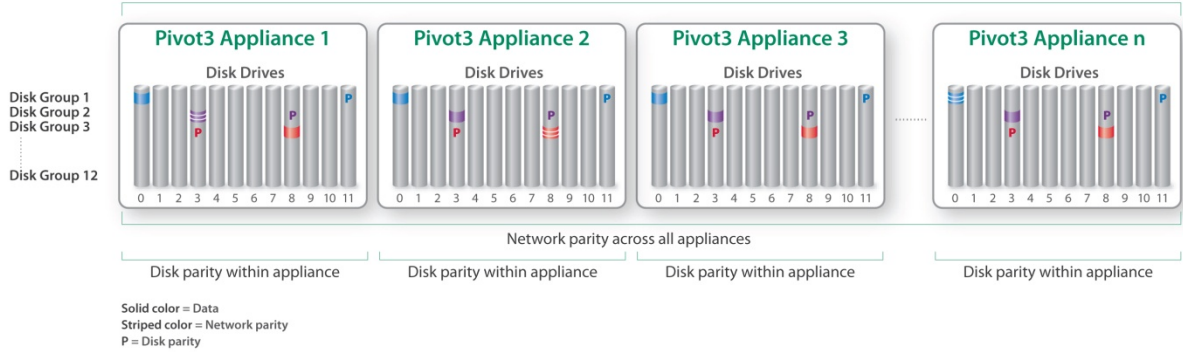
### RAID 6e

RAID 6e is recommended for most video surveillance environments. It provides cost-effective protection for more drive failures than RAID 5 systems.

- Enhanced distributed dual parity protects against either three drive failures or one appliance failure plus one drive failure. The “e” indicates enhanced RAID 6 for vSTAC Arrays of three or more appliances where data is also protected even if an entire appliance fails.
- Data is striped across each appliance in the vSTAC Array and protected by two levels of parity. The first level of parity, network parity, is striped across each appliance in the vSTAC Array. The next level of parity, disk parity, is striped across the disks within an appliance.
- Capacity required for RAID 6e protection is roughly one appliance per vSTAC Array and one drive per appliance to store two different parities in the vSTAC Array.

### RAID 6e

Protection against 3 Drive Failures or 1 Appliance and 1 Drive Failure



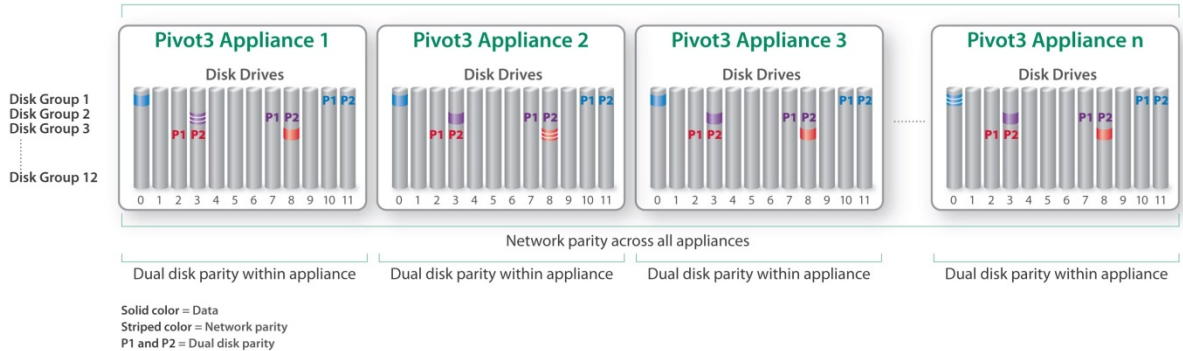
### RAID 6x

RAID 6x protects against even more drive failures and is recommended for the most sensitive surveillance applications or for applications such as border control where systems may be difficult to maintain.

- Enhanced distributed triple parity protects against either five drive failures or one appliance failure plus two drive failures. Like RAID 5e and 6e, RAID 6x provides fault tolerance even if an entire appliance fails for vSTAC Arrays of three or more appliances.
- Data is striped across each appliance in the vSTAC Array and protected by three levels of parity. The first level of parity, network parity, is striped across each appliance in the vSTAC Array, much like the parity in RAID 5e. The next two levels of parity, dual disk parity, are striped across the disks within an appliance.
- Capacity required for RAID 6x protection is roughly one appliance per vSTAC Array and two drives per appliance.

### RAID 6x

Protection against 5 Drive Failures or 1 Appliance and 2 Drive Failures



### RAID 1e

RAID 1e is supported but is rarely used in the surveillance market because it requires double the amount of primary capacity and is therefore rarely cost effective.

- RAID 1e protection is roughly one half of the appliances in a vSTAC Array. Precise usable capacity is included in the product specification sheets.
- Enhanced network mirroring protects against either one disk drive failure or one appliance failure.
- An exact copy of data written to one appliance is mirrored to other drives in the vSTAC Array.

## RAID Efficiency

One important factor to consider when selecting an appropriate RAID level is balancing between fault tolerance and usable capacity in a vSTAC Array. Simply put, higher protection levels result in lower usable capacity. Pivot3's innovative implementation of distributed RAID delivers up to 90% usable capacity and protects against more drive failures than scale-out approaches that offer simple mirroring. The higher usable capacity similarly delivers improved power, cooling and space efficiency that is especially important for large-scale deployments.

## Disk Groups

The vSTAC OS maintains logical Disk Groups that further minimize the effect of drive failures on the overall vSTAC Array. Disk Groups consist of one drive per appliance and are automatically created and maintained by the vSTAC OS. By organizing the placement of parity and mirror data within one Disk Group, the impact of a drive failure is limited to its Disk Group.

Disk Groups effectively increase the number of simultaneous drive failures that each Pivot3 vSTAC Array can sustain without data loss since drive failures outside of a Disk Group do not affect other Disk Groups.

## Changing and Mixing RAID Protection

Because RAID protection is set by volume, different RAID levels may co-exist on a vSTAC Array and RAID levels can be changed while the system is running, without disruption to data reads or writes. For example, a RAID 6e volume can be changed to a RAID 5e volume to free up space while access to the volume continues uninterrupted.

## Dynamic Expansion: Adding Physical Capacity to an Existing vSTAC Array

Capacity can be added to an existing Pivot3 vSTAC Array by connecting new physical appliances to the SAN and then configuring them into the existing Pivot3 vSTAC Array using the vSTAC Manager. The additional physical capacity is dynamically added to the Pivot3 vSTAC Array and data is restriped across the appliances so that capacity is automatically provisioned and performance is automatically load balanced across all controllers and iSCSI connections. The complexity normally associated with managing meta-LUNs for volumes larger than the domain of one physical conventional RAID controller is consequently eliminated.

## Instant Volume Initialization

The vSTAC OS relies on an allocate-on-write technique to immediately create volumes for use. This allows new volumes to be used instantly while conventional arrays require lengthy initialization times that for large volumes can take over 24 hours.

## Video Performance With Unified Flash Cache

Pivot3 vSTAC Watch appliances feature a patented solid-state (SSD) Unified Flash Cache (UFC) that is designed to accelerate write performance for latency-sensitive database operations in modern VMS applications.

The vSTAC OS aggregates and protects SSD resources in each vSTAC Watch appliance much like it aggregates and protects spinning disk resources without requiring expensive batteries. Volumes benefit from the aggregated write cache resources and the UFC scales as appliances are added to a vSTAC Array. Disk writes are automatically load-balanced across all available cache resources in a vSTAC Array.

The Unified Flash Cache is automatically configured by the vSTAC Manager and requires no user intervention.

## **Write Process**

The following steps describe how writes are allocated across a vSTAC Array:

- The vSTAC OS creates logical volume using a logical block address (LBA) map with a unique LBA owner for each block that is being written. When the vSTAC OS receives a write request, the request is routed to the correct LBA owner who becomes the master for that write.
- The master for the write simultaneously places the data in the Unified Flash Cache and sends the data to another appliance in the vSTAC Array to be placed in the second cache. Only after the data is safely stored in both caches will the vSTAC OS send an acknowledgement for the write request to the application.
- As a background task, data is moved from cache to its final LBA location where it remains and can be read or overwritten at any time. The vSTAC OS also calculates network and disk parities and writes them across appliances.
- For RAID 1e and 5e, each appliance writes to one disk. For RAID 6e, the data and/or parities are written to two disks in each appliance and for RAID 6x, three disks in each appliance. These multiple writes protect the data from power loss or disk failure.

As additional appliances are added to a vSTAC Array, or if an appliance fails, the LBA map is updated and the new map holds the new unique locations of each data block.

## **Virtual Global Sparing**

Virtual drive sparing is used to automate and speed drive rebuilding if a drive fails in a Pivot3 vSTAC Array. The capacity of one logical spare drive is reserved across all of the drives in the vSTAC Array and removed from usable capacity. In the event of a drive failure, the rebuild process begins immediately using the previously reserved capacity.

Unlike spare drives in conventional systems that standby during normal operation, virtual global spare drives in a Pivot3 vSTAC Array contribute to the overall performance of the RAID system during normal operation.

Virtual RAID controller sparing is also supported since data is protected in the case of an appliance failure.

## **Parallel Rebuild Innovation**

Conventional RAID systems are constrained by the physical relationship between RAID groups and their member disk drives. As a result, sparing and rebuilds are similarly constrained to physical drives. This becomes an important limitation as rebuild times for large drives increases.

Pivot3 vSTAC Arrays provide extremely fast parallel rebuilds of failed drives because of the distributed nature of data allocation and sparing. Many drives contribute to the rebuild process and the recovered data is written to all drives resulting in a massively parallel activity. Only sectors of a failed disk that actually have data allocated and written need to be rebuilt which further speeds rebuild times in lesser utilized vSTAC Arrays.

## **Priority Rebuilds by Volume**

Rebuilds are performed by volume. All volumes in the vSTAC Array are allocated to all of the drives in the vSTAC Array. An added benefit is the ability to designate a priority level for each volume so that higher priority volumes are rebuilt first. Rebuilding any specific volume may require rebuilding only a small portion of a drive.

With conventional systems, rebuilding happens at the disk level, which generally means volumes are only protected once the entire disk is fully rebuilt.

### Background Verification

The Pivot3 vSTAC Array continuously performs background disk verification. Each disk is completely scanned to identify disks that are beginning to fail and to detect and repair bad blocks on the media. This is another process that benefits from the massive available bandwidth of the vSTAC Array and the processing power available in the Pivot3 appliances.

### Predictive Sparing

Pivot3 Predictive Sparing is a background routine that continuously monitors disk drives to identify drives that could negatively impact overall system performance or that have a high likelihood of failure. Since drive deterioration is often characterized by gradual performance degradation, predictive sparing is an important method of gracefully removing suspect drives from the vSTAC Array, both to optimize performance and to keep the Pivot3 vSTAC Array in a fully protected state.

Predictive Sparing includes two factory-set thresholds for each drive. If the performance parameters for a drive exceed the first threshold, the low performing drive is automatically spared and the drive is subsequently and gracefully removed from the vSTAC Array. If the second higher threshold is exceeded Predictive Sparing fails the drive immediately and automatically initiates drive sparing, except in the case where failing the drive would cause data loss.

### Automatic Load Balancing

Dynamic load balancing of bandwidth and capacity across initiators, network ports, appliance controllers, disk drives and UFC is managed by the vSTAC OS with no administrative intervention. Since data is equally distributed across the Pivot3 vSTAC Array, changes to either the physical infrastructure or logical entities can be quickly accommodated to eliminate disk, controller and network hot spots.

For write operations, the dedicated IA-64 processors in each appliance have ample processing power for both RAID operations and TCP offload processing. For read operations, Pivot3 vSTAC Watch appliances return data to the application servers in parallel, providing load-balanced performance across all appliances and all drives.

### Physical Pivot3 Appliance Connections



The parallel architecture and load balancing of the vSTAC OS allows Pivot3 vSTAC Arrays to effectively aggregate many 1Gbps Ethernet ports and quickly surpass the bandwidth in proprietary Fibre Channel systems.

Load balancing of capacity and performance extends to physical reconfiguration of each Pivot3 vSTAC Array. Following additions or removals of physical appliances to an existing vSTAC Array, the vSTAC OS restripes data across the new physical appliance count and automatically optimizes the load across the new physical network connections.

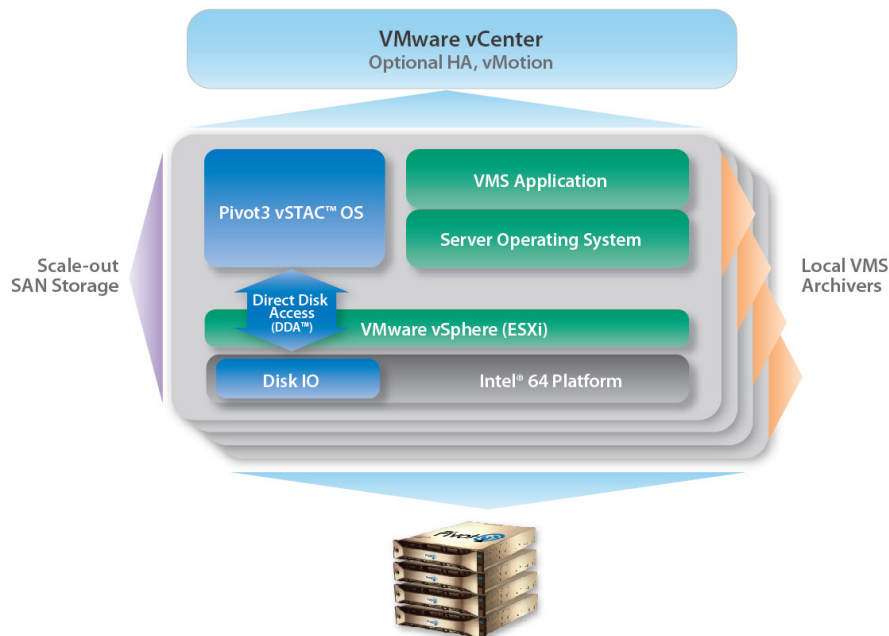
### SNMP Support

Pivot3 appliances can be monitored using Simple Network Management Protocol (SNMP). Community strings for SNMP are configured through the vSTAC Manager and the SNMP MIB (management information base) is provided with the Pivot3 software. Because appliances cooperate within a vSTAC Array, SNMP agents can be set once at a vSTAC Array level and do not need to be set for each appliance.

Pivot3 appliances run the Pivot3 SNMP agent and send out SNMP events to third-party software applications that receive, or trap, the SNMP notifications. Since SNMP traffic is on the storage network, the server running third party application or an SNMP trap receiver needs to have access to both the storage network and the server network. Many network-management applications use SMTP to forward traps as email. The network-management application will push an SMTP message to the corporate email server, which will forward the message as an email to a specified email address.

## 3. Direct Disk Access

A critical part of the technical innovation inherent in the vSTAC architecture is in a patented technique for integrating virtual servers with a high-performance distributed RAID stack. This innovation is called Direct Disk Access (DDA) because it pipelines video data IO directly to native disk controller. This performance optimization eliminates intermediate virtual disk services to deliver performance that is ideally suited for the video surveillance market where IO-intensive servers are not easily virtualized.



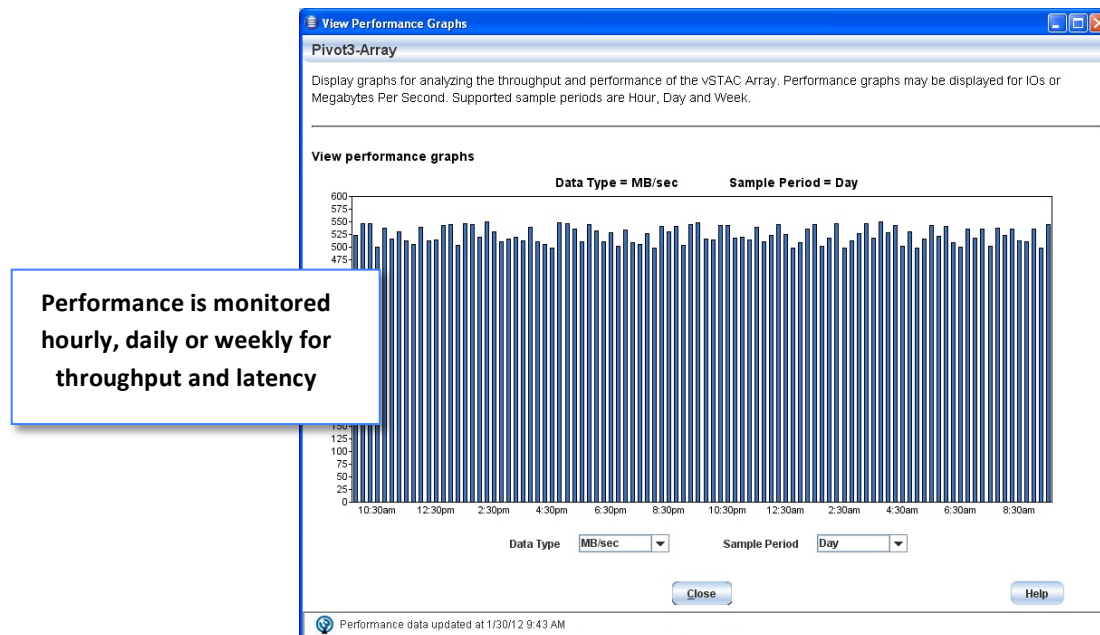
### Patented DDA Pipelining

In conventional virtualized server, the hypervisor, such as ESXi, has full control of the platform hardware and emulates I/O devices. This emulation layer enables the hypervisor to provide a consistent view of a hardware device and to isolate and contain device accesses to only those VMs that are assigned to the device by the hypervisor. However, it also slows down the I/O performance, resulting in as much as 30% overhead.

The Pivot3 DDA feature assigns the native disk controller directly to the vSTAC OS while maintaining isolation of VMs. This unique approach optimizes disk performance for the write-intensive that is typical in video surveillance environments.

### Performance Monitoring

Dynamically updated performance graphs are provided for both bandwidth and latency. The graphs can be updated on an hourly, daily or weekly basis.




## 4. VMS Failover Support

Video surveillance systems require high availability (HA) of both server and storage resources. Pivot3 supports a variety of approaches to match customer needs.

### Pivot3 VM Failover™

Conventional systems require separate high availability (HA) solutions for server and storage while the vSTAC Array uniquely solves both requirements in a single approach. Pivot3 offers a patented failover feature where a VMS application running on a vSTAC Watch appliance will automatically restart on another available vSTAC Watch appliance. If a vSTAC Watch appliance fails, the VMS running locally on the affected appliance will reboot on another appliance in the vSTAC Array without user intervention and will automatically re-establish both camera connections and storage connections to the recording volume. This self-healing restart is accomplished without separate clustering software, extra cabling, or VMS failover licenses. The feature requires that one vSTAC Watch appliance in

the vSTAC Array is available as a failover target. This is possible since the virtual machine boot volumes are protected against an appliance failure and are available from the remaining appliances. The virtual machine dynamically re-establishes virtual network and storage connections and will typically be running in less than two minutes.



#### Failover Software Packages from VMS Providers

Many VMS providers offer failover licenses for real-time continuous failover protection. These offerings are complementary to Pivot3 VM Failover and can be used in conjunction with Pivot3 VM Failover so that recording is never interrupted and users have access to video stored up until the server event.

#### Virtual Machine Failover using VMware HA™

Each VMS application can be viewed and managed as a standard VMware virtual server using VMware vCenter. VMware offers a failover mechanism called VMware HA that can be used although separate licensing is required for both VMware HA and for vCenter.

#### VMware vMotion

VMware vMotion can be used with the vSTAC Watch as with any virtual machine. Pivot3 vSTAC Arrays support vMotion by providing multi-initiator support so that multiple external physical servers can access the same virtual machine volume files concurrently. vMotion, which is separately licensed by VMware, allows a running VM to be dynamically migrated from one physical server to another to:

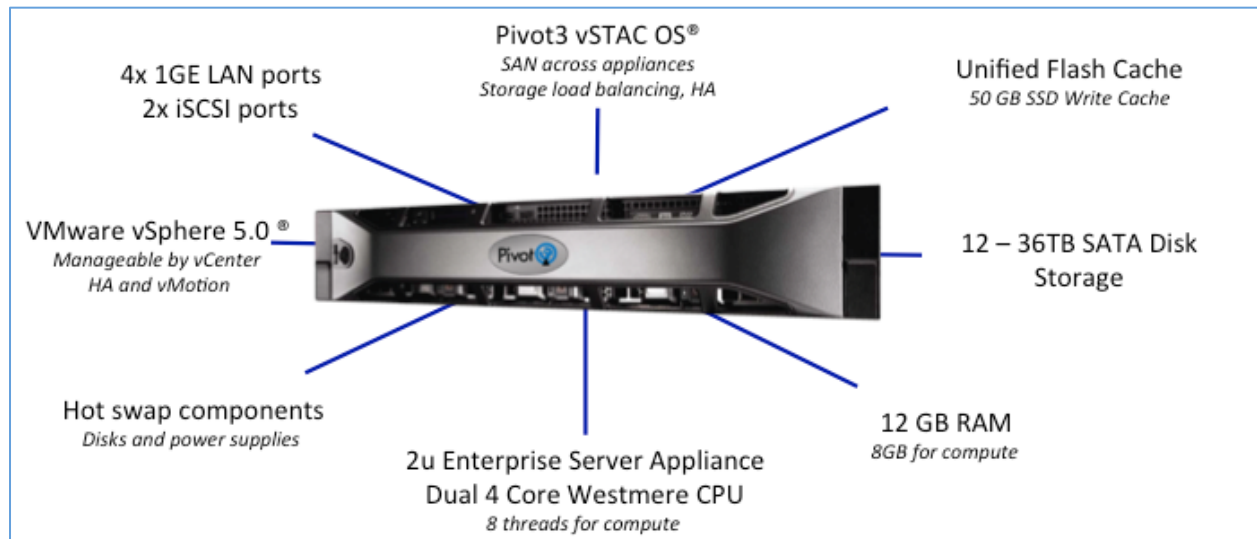
- Continuously and automatically optimize VMs within resource pools
- Perform hardware maintenance without scheduling downtime and disrupting business operations
- Proactively move VMs away from failing or underperforming servers

## 5. Pivot3 vSTAC Appliances

#### vSTAC Watch Appliances

Pivot3 vSTAC Watch appliances are the hardware building blocks of Pivot3 surveillance storage solutions. Each appliance is based on off-the-shelf server hardware.

### vSTAC Watch Platform Components



For the latest vSTAC Watch technical specifications, please visit [www.pivot3.com](http://www.pivot3.com).

### vSTAC Data Appliances

Some video surveillance installations, particularly those with small numbers of cameras and long retention times, require a higher mix of storage to compute. For these environments, the Pivot3 vSTAC Data appliances can be introduced into a vSTAC Array to optimize cost and power needs. Each vSTAC Data appliance contributes additional scale-out storage to a vSTAC Array but does not contain a second CPU or the VMware vSphere hypervisor and, as a result, cannot run a VMS application or participate as a VM Failover target. In all other aspects, the vSTAC Data appliance is identical to the vSTAC Watch appliance.

For the latest vSTAC Data technical specifications, please visit [www.pivot3.com](http://www.pivot3.com).

## Summary

The video surveillance market has unique needs that reach beyond product design. Pivot3 has invested in the following initiatives to be sure its products are easily supported in open systems environments:

**The HD Storage Alliance™** - Pivot3 works closely with leading video surveillance partners to ensure smooth integration and support of open systems in the field. Pivot3 camera simulators, automated test beds, and configuration tools are used to certify authorized partners and assist in field deployments. The latest list of certified partner products is available at [www.pivot3.com](http://www.pivot3.com). New partners can request information on certification procedures by sending email to [partners@pivot3.com](mailto:partners@pivot3.com).

**Architectural and Engineering (A&E) Specifications** - For architects and engineers, Pivot3 has developed A&E specs that can be used to quickly respond to bid requests and user inquiries. A&E specs can be found in the library on [www.pivot3.com](http://www.pivot3.com).

For other questions or inquiries, please send email to [info@pivot3.com](mailto:info@pivot3.com)

Pivot3, Inc.  
1900 Embarcadero  
Palo Alto, CA 94303

[www.pivot3.com](http://www.pivot3.com)

Tel: 855-EXABYTE



Copyright © 2012 Pivot3, Inc. All rights reserved. Specifications are subject to change without notice. Pivot3, vSTAC OS, vSTAC Watch, vSTAC Data, Serverless Computing, HD Storage Alliance, Unified Flash Cache and DDA are trademarks or registered trademarks of Pivot3. All other trademarks are owned by their respective companies.