KINEMETRICS

Advancement through Innovation

Virtualized Data Center

WWW.KINEMETRICS.COM

Vitualized Data Center

INTRODUCTION

ASPEN's Open System Solutions provide a comprehensive and integrated platform for seismic, strong-motion, geodetic GPS, and infrasound networks with both on-site archiving and telemetered data streams for central data processing and archiving.

The Aspen Data Center is highly scalable. For small research networks, a LINUX workstation is sufficient for real-time data acquisition, processing, archiving, and off-line processing. On the other hand, the Aspen Data Center can consist of enterprise class servers connected to shared and tiered storage units and a virtualized environment acquiring.

OVERVIEW

The redundant Aspen Data Center provides continuous computational resources to obtain well in excess of "2-nines" (99%) data availability.

Continuous computational resources are secured with two or more enterprise-class servers that share redundant-connected external disk space. The two virtual machines for data acquisition and processing can move between the servers in case of any necessary maintenance operation (manual failover) or failure within a server (automatic failover). The shared disk space is located on a RAID that is directly attached via two dual-controllers to each server providing failsafe connection. Moreover, employing a combination of very fast solid-state drives (SSD), fast hard disk drives, and slow large disks optimizes data access, availability of large deep storage, and cost per archived byte.

The servers communicate through two network switches, providing continuous intra-server and LAN connection.

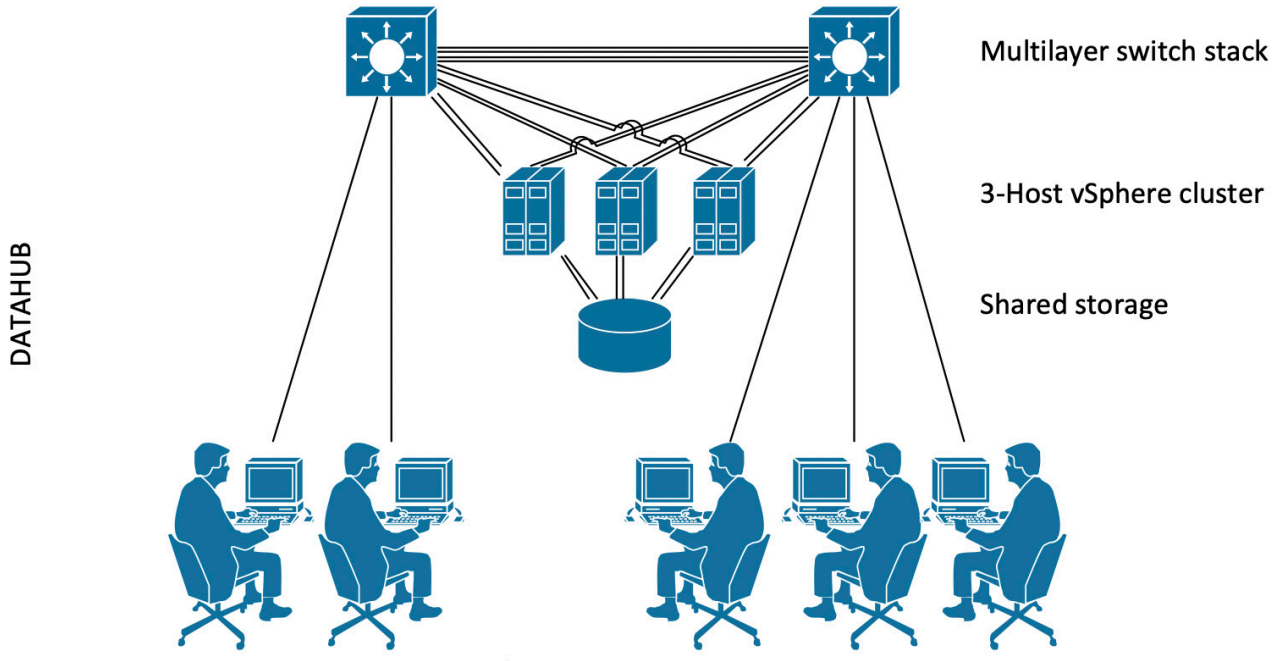

Data center HA layout

Heart of Computing Infrastructure

Heart of the computing infrastructure is a vSphere HA cluster based on the VMware ESXi hypervisor. All server-based systems (e.g., data acquisition & automatic processing, manual post-processing, data archiving, web interface, development & research) are running on virtual machines. Virtual machines are balanced over the cluster nodes to achieve the best load distribution but are prioritized based on their mission-criticality:

Priority	Description
HIGH	Real-time data acquisition & automatic processing (including dependable infrastructure)
MEDIUM	Manual post-processing and web interface
LOW	Research and development

Designated failover capacity on each server guarantees that, in case of hardware failure, mission-critical systems can automatically be restarted by vSphere HA on another server while low priority systems (development, test, research) have to be restarted manually. The vSphere service continuously monitors failover capacity and prevents virtual machines from starting that would violate the failover capacity policy. Live migration of virtual machines allows re-distribution of virtual machines without shutting them down. This is essential for planned hardware and firmware/software maintenance (without service downtime).

The shared storage employs dual active caching controllers for non-stop operation. The hosts are attached via dual 12Gbps SAS host bus adapters (HBAs), each connected to one controller (direct-attached shared storage) and I/O is balanced over all active links (ALUA). RAID6 disk groups (dual parity) allow for failure of up to (2) disks without data loss. Additional hot spares ensures that the RAID can instantly be rebuilt after disk failure (defective disks are replaced during local business hours).

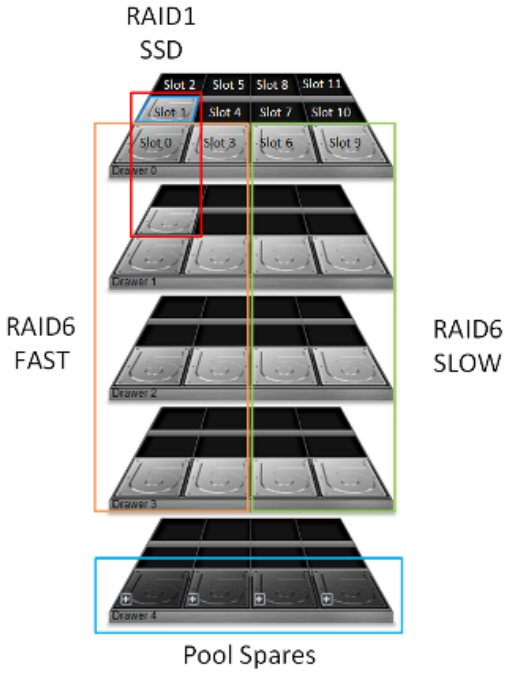

Workstations have single Ethernet connections and power supplies. At least one Operator and one Analyst workstation connected to each switch (A or B) and to different power sources. In case of switch or power failure, Operators and Analysts can continue their work on another workstation.

Due to the redundancy built into the system, maintenance and replacement of defective parts can be done during general local business hours (cost and resource effective).

Hardware & Software

A typical Aspen Data Center consists of:

- 2 (better 3) x Dell PowerEdge without local storage
- 1 x DellEMC PowerVault ME storage array
- 2 x Dell Network switches
- UPS (optional)
- Serial to USB converters for management of physical console ports (optional)

The components of the data center are fully virtualized, allowing resources to be shared or failed over, independently of the physical hardware. The main layers of virtualization are:

1. Computing
2. Storage
3. Network

Virtualization of computing resources starts with the integrated Dell Remote Access Controller (iDRAC), an independent controller that allows the hardware to be managed virtually without physical access. This includes management of power and remote access to virtual keyboard, video and mouse functions.

The Hypervisor running on the Dell PowerEdge servers provides the virtualization layer in between the bare metal and the virtual machines (VMs). The primary operating system for the virtualized Aspen System is Red Hat Enterprise Linux or alternatively CentOS). It hosts the Antelope Real-time system in two or more virtual machines.

Special VMware appliances, running a customized version of Linux, provide other vSphere management functions (vCenter, Backup). The figure on the next page shows the virtualization layers within a data center.

Virtualization of network resources is implemented using Rapid Spanning Tree (RSTP), virtual LANs (VLANs) and virtual router IP addresses (VRRP) allows the automatic activation of alternative paths and router interfaces to overcome failures. Reconfiguration of network resources generally does not require any physical intervention.

The physical servers of the Aspen System do not have any local storage other than 1GB redundant SD flash cards to boot the hypervisor. All storage resources are provided by a highly redundant DellEMC PowerVault storage array with redundant controllers and fully redundant SAS connections to the two Aspen hosts. The ME4024 base enclosure has five (5) disk drawers, initially populated with 22 disks:

- SSD = 2 x 480GB SAS SSD configured as RAID1
- FAST = 8 x 1.8TB 10k RPM SAS HDD configured as RAID6
- SLOW = 8 x 4TB 7.2k RPM NL-SAS HDD configured as RAID6
- Hot Spares = 4 x 4TB 7.2k RPM NL-SAS HDD configured as pool spares

Information & Storage

The following table summarizes the information and storage capacities and gives an overview of the allocated and spare (free) virtual disk capacity:

Disk group	SSD	FAST	SLOW
RAID type	RAID1	RAID6	RAID6
Hotspares	0	4	4
Disk group capacity	441.630 GB	10.027 TB	21.799 TB
Virtual disk	0	1	2
Virtual disk capacity	441.630 GB	4096 GB	10 TB
Free capacity	-	5931.288 GB	11.799 TB
Datastore	SSD	FAST	SLOW

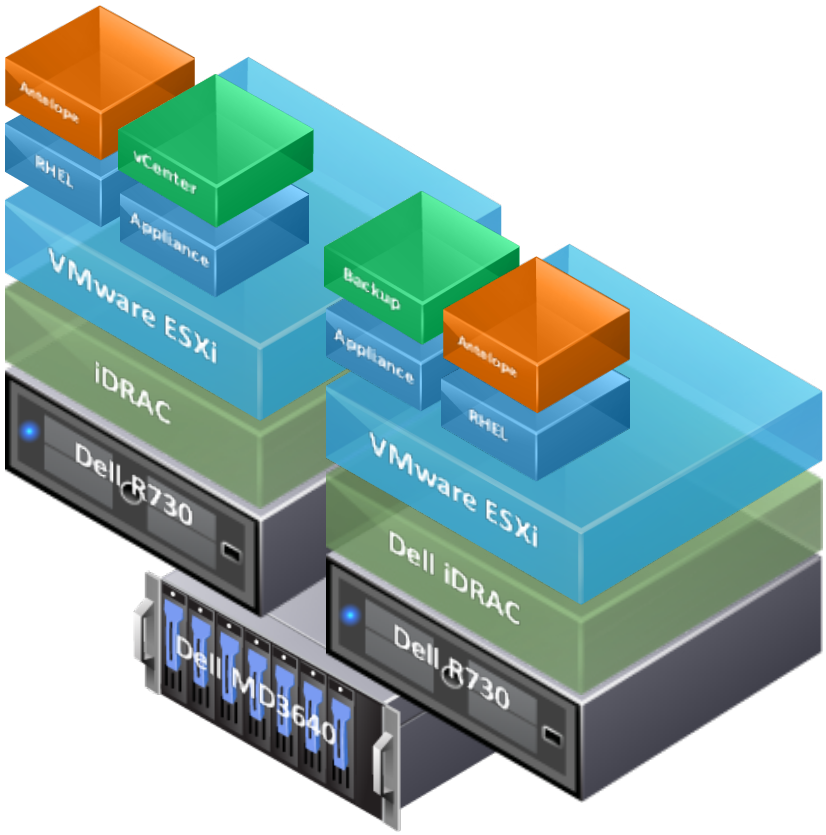

Virtualization Layers

Disks and disk groups have been arranged in a way that a failure of a whole disk drawer will not lead to any data loss. Each vertical stack of disks has at least one (1) parity disk and no RAID group spans more than two vertical stacks (RAID6 uses two parity disks). Sufficient hot spares will automatically cover any hard disk failure in the system to allow any of the two RAID6 arrays to be rebuilt without waiting for the defective disk to be replaced. The figure below shows all layers of storage virtualization from physical disk groups to the file systems of virtual machines.

Information & Storage

Users and operators connect from their desktop computers or any devices with a browser with Starnet's FastX-3 to their respective accounts. Because the connections are stateless (residing on the server), any session can be stopped and resumed from any devices. FastX-3 can display a Linux desktop inside a browser.

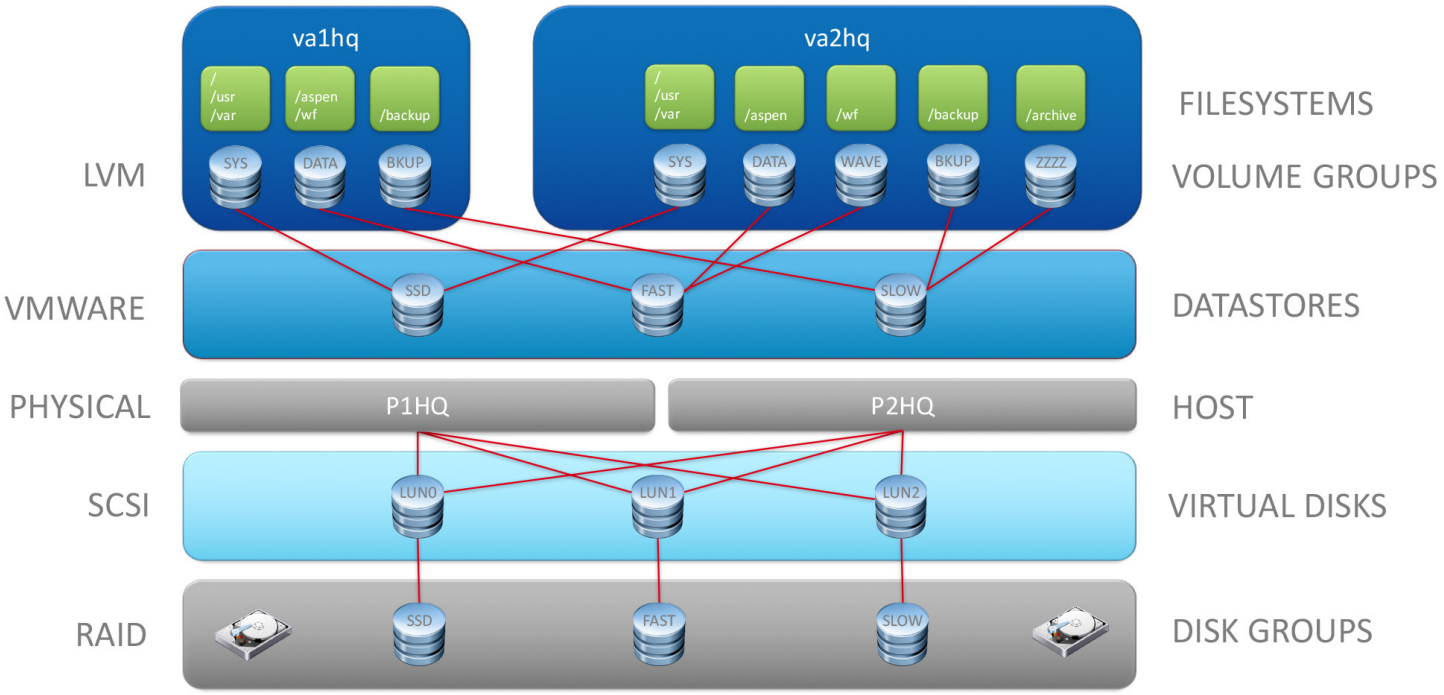

Storage Virtualization

222 Vista Ave., Pasadena, CA 91107 USA
Telephone: +1-626-795-2220
Fax: +1-626-795-0868
E-mail: sales@kmi.com