



Global Leader in Software Defined Storage

# Nexenta Technical Sales Professional (NTSP)

COURSE CONTENT

---

Nexenta Technical Sales Professional (NTSP) Course

# **USE CASE: VIRTUALIZATION/HYPERVISOR**

# Use Case

---

## Virtualization & Hypervisors

Objectives:

- General characteristics
- System Design
- Disk Layout
- Optimization
- Hypervisor Settings
- Examples

# General Characteristics

---

- VM IOPS
  - This will vary on every system. Most Hypervisors provide a way to gather basic statistics on system performance.\*
  - Collect information on existing environment as much as possible
  - Read/Write
    - This can vary from one extreme to another
    - High Read will benefit from Memory & Disk Layout
    - High Write will benefit from ZIL/SLOG
- Rule of Thumbs
  - General VM
    - 50 IOPS
  - DB VM
    - 100 IOPS

*\*Note: See "Hypervisor Settings" pages for individual Hypervisor settings*

# System Design

---

- CPU
  - The faster the better
  - Prefer faster core speed versus number of cores
- Memory
  - Minimum 96GB
  - Larger memory for Read intensive environments
- Network
  - 4 x 10Gbe – Determine if SFP+ or BT
    - 2 x LACP across cards
    - IPMP between the two LACP pairs
- HBA
  - One JBOD per HBA if possible

# Disk Layout

---

- Pools
  - Use 2 or more pools whenever possible (Active / Active)
    - This will prevent a stall across the whole virtualized environment in case of a failover
- Redundancy Type
  - Mirror
    - This provides for the best backend disk performance (IOPS) in both Reads and Writes
    - The more vDevs the better
- ZIL/SLOG
  - Based on IOPS
    - Lower latency matters
  - Mirrored Pairs
    - Since 2 pools make sure you have at least 4
- L2ARC
  - 250 or less VMs – 400GB
  - 250-500 – 600GB
  - 500-1000 – 800GB
  - 1000 or more – 1TB

# Optimization

---

- Record Size
  - 32K is recommended
- Protocol
  - NFS with 6 folders per Pool minimum
- All I/O is random
  - All the different VM's will create random I/O to the storage
- Failover times
  - Modify protocol timeouts to 120 seconds
    - iSCSI
    - NFS
- Disable atime on the pool/s
  - `zfs set atime=off <pool_name>`

Note: See “Hypervisor Settings” pages for individual Hypervisor settings

# Hypervisor Settings - VMware

- Performance
  - VMware – Click on host → Click performance Tab → Click Chart Options.
    - There are a number of areas that you can use to get the information you are looking for:
      - Datastore
      - Disk
      - Storage Adapter
      - Storage Path
    - Select the following Counters:
      - Average Read requests per second – READ IOPS
      - Read Rate – This divided by above is avg I/O size
      - Average Write requests per second - WRITE IOPS
      - Write Rate - This divided by above is avg I/O size
- HA – Protocol Timeouts
  - iSCSI –
    - RecoveryTimeout = 120
  - NFS
    - NFS.HeartbeatFrequency = 12
    - NFS.HeartbeatTimeout = 5
    - NFS.HeartbeatMaxFailures = 10



# Hypervisor Settings - XenServer

---

- Performance
- HA – Protocol Timeouts
  - iSCSI
    - Default is 144 seconds
  - NFS
    - Modify `/opt/xensource/sm/nfs.py`
      - Adjust the NFS mount to use Hard mounts instead of soft and add the `intr` option
      - In the `soft_mount` definition modify the following line:
        - » `options = "soft,timeo=%d,retrans=%d,%s,noac"`
        - » `options = "hard,intr,timeo=%d,retrans=%d,%s,noac"`

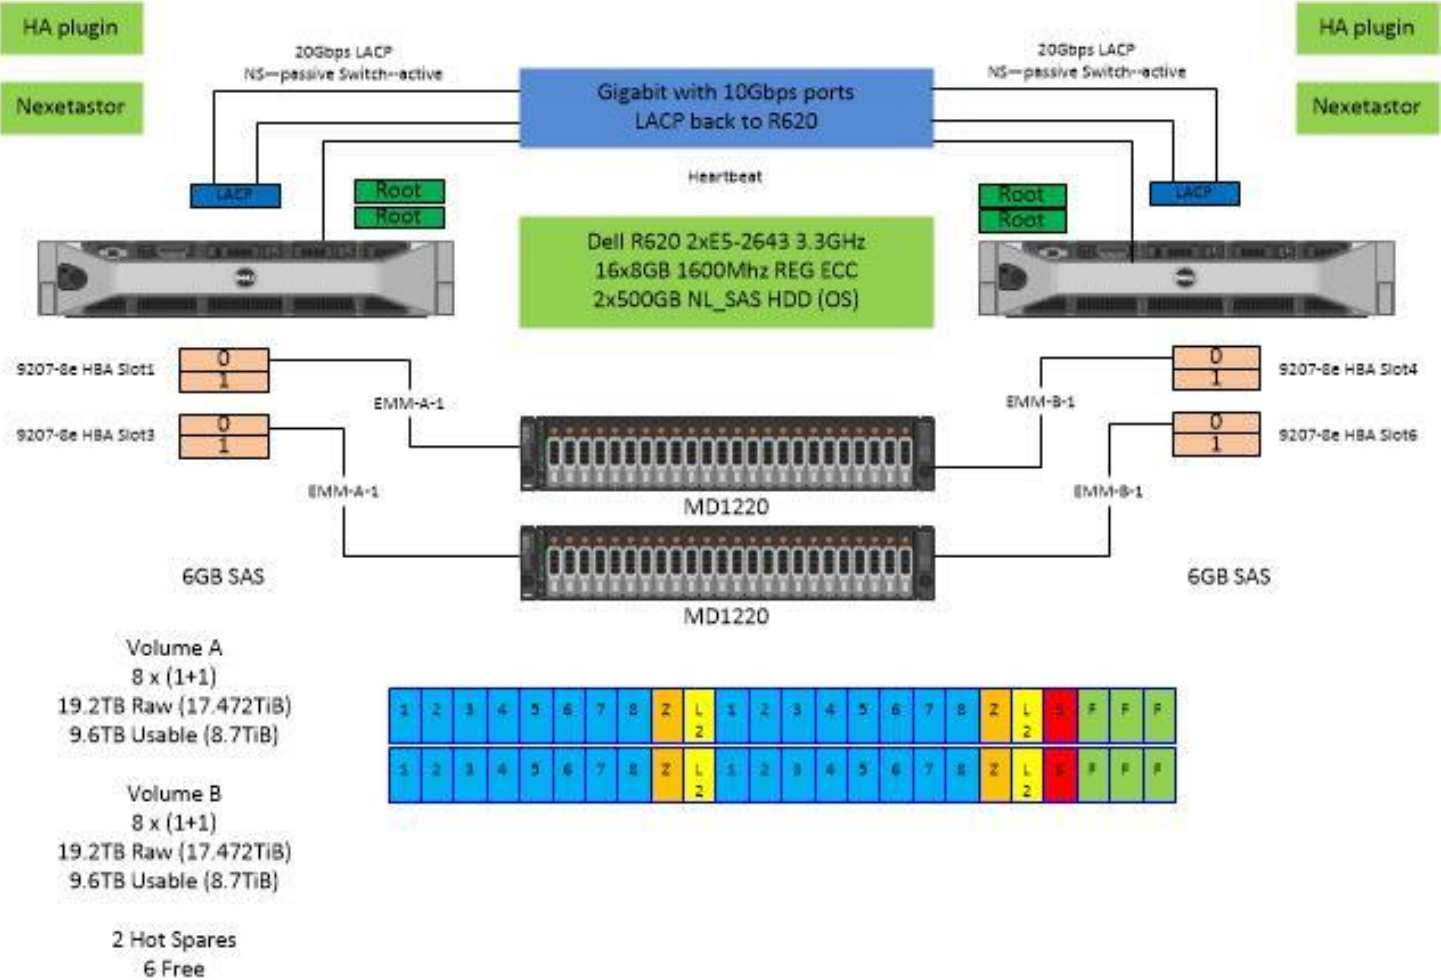
# Hypervisor Settings – HyperV

---

- Performance
- HA – Protocol Timeouts
  - iSCSI
    - Windows 2000, 2003, 2008, Vista, & Windows 7
    - Open the registry editor (regedit) and navigate to:
      - HKEY\_LOCAL\_MACHINE / System / CurrentControlSet / Services / Disk
        - » Look for 'TimeOutValue'. If it exists, edit it, and if it does not exist, right-click and choose 'Edit/Add Value' and create it. The type is REG\_DWORD, and the value should be set in decimal to the timeout in seconds that you desire (minimum of 120).

# Small Virtualized Environment Example

## 2 Volume @ 8TB



# Module Quiz Questions

---

- What is the minimum memory?
  - 96GB
- What is the recommended Block Size?
  - 32K Block Size
- What is the preferred Redundancy Type?
  - Mirroring
- VM environments are all sequential I/O
  - T/F – False, Virtual environments create completely random I/O
- What are the preferred number of pools?
  - 2 – One per head to reduce disruption in the event of a failover