

Recommended VMware-based Server Specification

Enterprise Imaging Workflow Unifier

Dicom Systems Unifier Platform VMWARE or HYPER-V Specifications

RAM: 16 GB min. or 32 GB recommended

INTEL CORES: 4 min. or 8 recommended

DISK I/O: NO RAID 5 configuration for this VMware setup. Speed must be equivalent to single or mirrored disks with at least 20-25 Mbytes/s min., or 40-45 Mbytes/s recommended random write speed. The root partition is 70 GB. A short-term image cache drive should be provisioned between 300 GB - 500 GB depending on customer requirements, and actual imaging volumes. Additionally, a 100 GB drive should be provisioned for logging.

High Availability configuration is recommended.

FAQs

Can we load balance the routers?

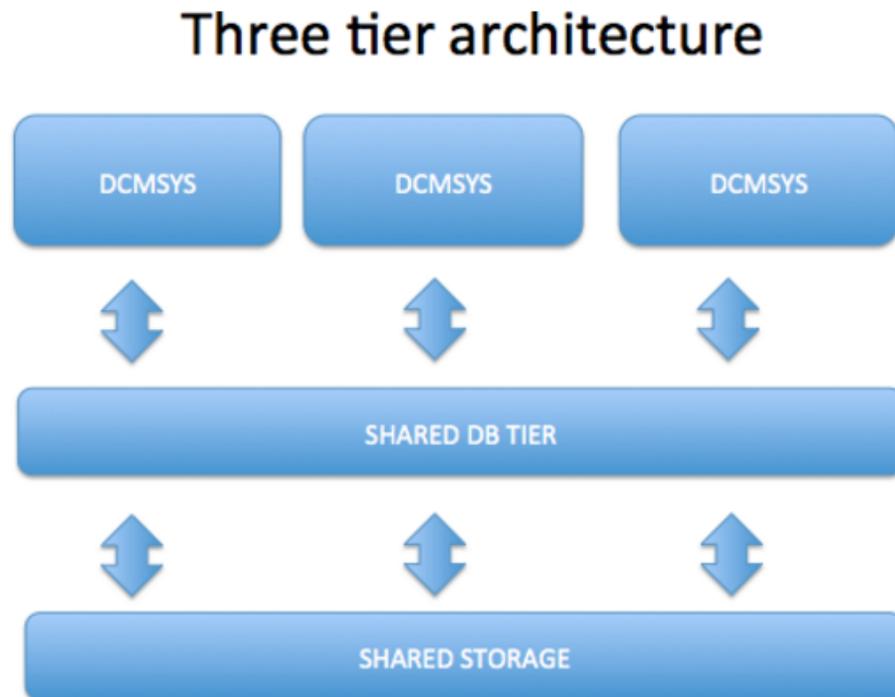
If you load balance traffic between Unifier devices you will end up with split data across all the devices. Load Balancing doesn't work on the DICOM level.

Some of our clients use our devices to load balance traffic in certain scenarios. For example, if you install one routing cluster in front with no transformation rules, or relevant priors rules with high I/O, it will be able to handle more traffic, and LOAD BALANCE to your other MAIN UNIFIERS. More information can be obtained on our website at dcmsys.com.

Another approach is to load balance inbound image traffic coming from a spoke site Unifier by sending it to alternate HUB devices if you have more than one. You can use the failover routing feature, or simply a standard routing rule to implement.

Does Unifier support an Active/Active configuration?

You can implement an ACTIVE / ACTIVE configuration if a three-tier architecture is used as shown here:



Although this configuration is supported, the client is responsible for maintaining the SHARED DATABASE TIER and SHARED STORAGE, and assuring that they have adequate speed to support the expected performance. This configuration typically requires the client to allocate substantial IT resources to maintain the SHARED DATABASE TIER, SHARED STORAGE, and connectivity between those tiers. In this architecture, if the second or third tier goes down, the whole cluster is down. This environment is susceptible to a single point of failure unless you invest in an expensive HARDWARE replication plan with multiple levels of redundancy. This method works best for customers who maintain their entire IT infrastructure in-house (software, hardware, network, IT staff), and have already invested heavily in their DATABASE TIER and STORAGE TIER architectures.

Being that we do not control on the customer's network and IT infrastructure, we design our product to be fully self-recoverable. Our watchdog services restart cluster nodes and redirect traffic to the secondary node with minimal interruption or downtime.

Can we add additional Unifiers to our cluster?

Yes, you can add multiple passive nodes to your cluster. You can also implement multiple clusters, and designate a site to distribute the traffic between multiple clusters using our Routing function, or failover routing mechanism.

How do we scale out our Unifier to accommodate larger study volumes? Would adding additional CPU cores or RAM affect our supported study volume?

The most effective way to scale in the future, and improve throughput, is to improve your I/O. For example, by implementing faster storage by using SSD disks. The complexity of the unique workflow being performed (number of transformations, routing rules, scripts, etc.) in your environment has an impact on throughput. If you are already using SSD disks, we can perform an analysis on your Unifier load and determine if allocating additional CPUs or memory is necessary to accommodate your anticipated image traffic growth.

We can horizontally scale by introducing another cluster(s) and split traffic using the logic from the spoke-side Unifier, or other devices.