

Nil Solutions Configuration Guideline

Document Number: 941-07020

Product Name(s): NilRead

Product Version(s): 3.X, 4.X, 5.X

Date: July 2020

1 Executive Summary

Hyland’s Nil family of zero footprint viewers can be used by clinical personal to view any type of DICOM or non DICOM medical images inside and outside the hospital network. Nil is available as both a clinical reference viewer (“NilRead Clinical”), and a full diagnostic sophisticated viewer (“NilRead Interpretation”).

This document is intended as a guideline to design and size a Nil deployment. It provides general guideline and a set of example for common configurations.

The high-level structure of NIL deployment on a customer site in presented on the diagram below.

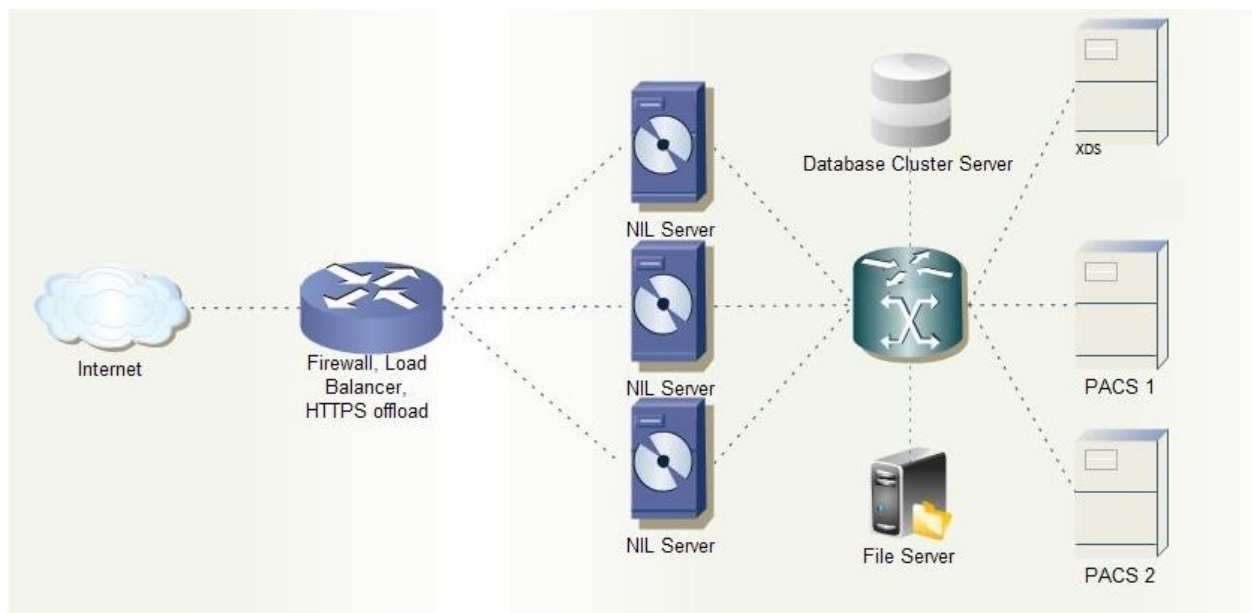


Figure 1 NIL Deployment Structure

The site will host one or more NIL web servers which provide user access to NIL imaging and reporting services. If more than one server is required to serve the number of concurrent users necessary at a site, then the access to the web servers is controlled by a load balancer (e.g. F5 BIG-IP). The load balancer is configured with session affinity by server, which ensures that user session requests are always routed to the same server after the initial assignment.

In some user deployment scenarios each NIL server uses its own NIL database to maintain user information and references to the cached data. In other scenarios Nil can pull directly from an outside DICOM archive.

Document: 941-07020, Rev 7.0

Each NIL server can be configured to interface with PACS, VNA and Enterprise archives, XDS and other data sources. DICOM data from the configured sources can be “pushed” to NIL, which caches data on a file share accessible by each NIL server. It is also possible to query/retrieve directly from multiple sources, and also setup prefetch rules to optimize the availability of priors.

NIL can be deployed in either single- or multi-tenant configurations. In the single-tenant configuration there is only one ‘operation’ context maintained which includes user accounts, data lifecycle rules, hanging protocols, DICOM connectivity setup, and cached DICOM data. In the multi-tenant configuration multiple operation contexts are supported – one per each ‘tenant’. Each tenant is provided with its own "Virtual NilRead Site". These "virtual NilRead Clinical Sites" share the same 'physical' NilRead Clinical infrastructure - front end and back end servers. Tenants are provided with a dedicated database, by default tenants do not have access to each other data, unless “trust” relationships are established

2 General Guidelines

Depending on deployment scale all NIL system components might reside on a single or multiple physical or virtual hosts. “Mission critical” scenarios would require some level of hardware redundancy, which can be implemented on both on physical and virtual level.

Nil can be run on either physical or virtual server; the difference in performance is negligible (as long as hardware supports virtualization). For scalable deployment and high availability it is recommended to use a virtualized environment. Nil supports scaling up the number of servers in conjunction with a front end load balancer to address large numbers of concurrent users.

2.1 System Resource Requirements

The system resources required by Nil depend on several factors including the usage scenario (e.g. 2D, 3D, etc.), resolution of monitors and expected performance (e.g interpretation vs referring physician).

Nil, as any other the zero footprint solutions, is a server side rendering solution. The server side does all the heavy lifting and should be dimensioned properly. On the other end the client devices can be significantly less powerful, and can range from workstations driving multiple high-resolution monitors to low end desktop and all the way to tablets and smart phones.

NilRead Clinical

Resource	Minimum requirement based on single monitor 2.3 Mpx
CPU type	Xenon E5 @ 2.2 GHZ, recommended 2.66 GHZ

# users per CPU core	Modality	#users/core
	Mixed	5
	Mostly X-ray	7-8
	Mostly CT/MR/US	4
	Mostly MPR/3D	2
Memory	Modality	
	Mixed	1 GB
	Mostly X-ray	0.5
	Mostly CT/MR/US	1 GB
	Mostly MPR/3D	1.5 GB
Application disk space	20 GB	
Bandwidth	2 Mb/sec in 'local' mode, 1 Mb/sec in internet mode.	

Nil Read/Interpretation

Resource	Minimum requirement based on 2x 5 Mpx grayscale monitor	
CPU type	Xenon E5 @ 3.0 GHZ, recommended 3.2 GHZ or better	
# users per CPU core	Modality	#users/core
	Mixed	2
	Mostly X-ray	3
	Mostly CT/MR/US	1.5
	Mostly MPR/3D	1
Memory	Modality	Assuming 1-2 priors
	Mixed	1.5 GB
	Mostly X-ray	1.0 GB
	Mostly CT/MR/US	2.0 GB
	Mostly MPR/3D	2.5 GB
Application disk space	20 GB	
Bandwidth	5-10 Mb/sec in 'local' mode.	

For interpretation, client workstations using mix of high resolution monitors and multiple video cards should have a compatible drivers. When multiple video cards are installed, Nil will use OS choice in terms of which video cards are used for driving the monitors. For multiple high resolution monitors either IE 9+ / Firefox/Safari can be used. IE8 will not use hardware acceleration and will be have a slower rendering speed when driving 10Mpx. Chrome is blocking the application's ability to set the position of

views across multiple monitors but it will work with single 'seamless' dual pane monitor (such as Barco 6 MPx color etc)

NilRead interpretation deployments should be deployed using a 'cache' approach for fast access to image pixels and metadata. To reduce the cache size, the VNA or PACS (or other pre/post fetch mechanism) should be configured to push not only the current but associated relevant priors as well, at least those used in the typical hanging protocol with prior. A 500 GB cache will hold about 7-8K studies which is fine for this scenario. NilRead is able to send a notification to the Worklist/Ris when the studies are arriving in the cache so the worklist can mark them as available for reading.

Notes

1. Power savings features of the hosts should be disabled. By default Xenon E3/E5/E7 are not set in 'performance' mode and the CPU frequency is lowered causing very slow image interaction speed. This can be changed wither in BIOS or VMware ESX.
2. The number of cores is based on 'real' cores no vCores. If hyper-threading is on, each real core is exposed as 2 cores. In general, Nil will be faster by 10-15% if HT is disabled, for the same number of real cores.
3. Nil requires at least 4 cores to run smoothly. For a clinical proof of concept this can be lowered to 2-3 real cores (4-6 with HT on).
4. The difference in rendering speed between a VM and physical hardware is about 15%.
5. If is a 'cache' solution and the Dicom/Cache server is running on the same system as the rendering servers, it will need an extra 1-2 cores and 2-6 GB of RAM depending on the study volume. For heavy sites (> 250K studies/year) we recommend deploying the Dicom/Cache server on a separate VM (2-4 cores, 6-10 GB of RAM including OS), see details in "Repository I/O Capability" section for recommended and minimum I/O benchmarks.
6. Large number of concurrent users will typically use 30-50% of the total calculated bandwidth value due to normal interaction with the image patterns (x-rays will use very little bandwidth while scrolling modalities such as CT/MR and especially US/cardiology will use more)
7. Nil stores by default images in RLE lossless format, which uses about 65% of original. If Nil is deployed in a no-cache solution (reading directly from VNA) we strongly recommend using Jpeg Lossless instead of Jpeg 2000 storage syntax as the VNA store format. Jpeg 2000 is a very computationally intensive format for pixel decoding and will cause slow pixel access (scrolling). In a cache setup for interpretation, Nil should be configured to not accept Jpeg 2000 through Dicom as by default Nil stores incoming 'compressed' data in the original format.

The repository (cache) I/O capability also has a significant impact on overall performance, the following is the recommended and minimum I/O benchmark requirements:

Repository I/O Capability (NilRead Clinical/Interpretation)

Benchmarks		
Recommended: (SSD or similar):	#users/core	
	IOPs	> 25,000/s
	Throughput	> 150MB/s
	Average latency	< 3ms
	Latency distribution	90% < 6.0ms
Minimum: (10K RPM, Raid 5 SAN, or similar)	#users/core	
	IOPs	> 7,000/s
	Throughput	> 50MB/s
	Average latency	< 9ms
	Latency distribution	90% < 12.0ms

Note

The I/O capability should be verified with the Microsoft storage Testing Tool, Diskspd Utility: <https://gallery.technet.microsoft.com/DiskSpd-a-robust-storage-6cd2f223>, with the following measuring parameters (the cache is located in C:\temp in the example):

```
diskspd.exe /w10 /d300 /W30 /b8K /t1 /o8 /h /L /Z1G /c1G c:\temp\test1.dat c:\temp\test2.dat
c:\temp\test3.dat c:\temp\test4.dat c:\temp\test5.dat c:\temp\test6.dat c:\temp\test7.dat
c:\temp\test8.dat
```

The command simulates the Nil I/O access pattern, it will run 30 seconds warm-up before measuring, followed by 300 seconds 8K block sequential I/O tests against 8 1GB test files on the target drive (C: drive in the example), using a 30% write and 70% read ratio with 8 worker threads that each leverages 8 overlapped IOs and a write entropy value seed of 1GB, 1 thread per target. Software and hardware write caching are disabled.

Software

NIL can be deployed on all English editions of Windows Server x64 starting with 2012 R2. Nil will need as prerequisites: IIS 8.5 or newer .Net 48 or newer, Microsoft SQL Server Express 2008 SP2 or newer.

For 'no-cache' deployments or when the cache size is under 5 million images an SQL Express database can be used. For DR/HA deployments, the NIL database is expected to be deployed on a clustered Microsoft SQL Server Enterprise edition and the databases synced between primary and DR data centers.

High availability and Disaster Recovery

Large scale NIL deployments would require usage of hardware load balancers (e.g. F5 BIG-IP) providing cookie based user session persistence, as well as HTTPS offloading. Another key factor for large scale

Document: 941-07020, Rev 7.0

deployments is performance of file storage, which can be achieved by using multi-channel 10GbE or FDDI based SAN/NAS/DAS storage array.

The typical medium-large scale deployment network topology is presented on Figure 2. It includes a redundant firewall/routing tier. The load balancer tier would offload SSL and redirect client requests to a server associated with the request authentication token (implementing as “sticky” session load balancing). If the load balancer uses a ‘client IP’ sticky session management, the SSL offload can be done on the Nil server. In this mode a valid wildcard SSL certificate should be installed on each of the rendering servers.

Nil servers run on either physical or virtual servers connected to high-throughput storage system using either 10GbE or FDDI channels.

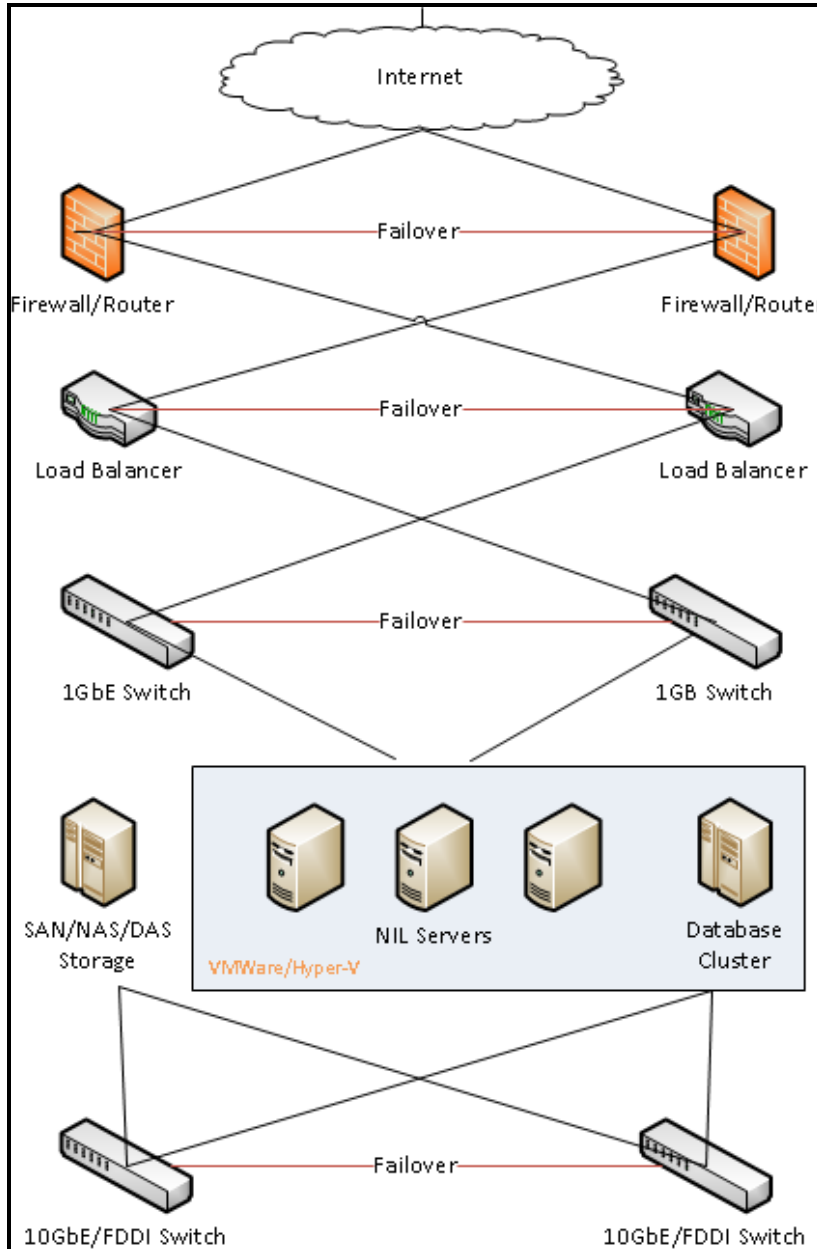


Figure 2 NIL site network topology

2.2 Cybersecurity Best Practices

It is important that the Nil Systems are included in the site's cybersecurity framework to ensure protection of confidential patient data and continuous operation.

While not exhaustive this framework should consider

Document: 941-07020, Rev 7.0

- Network defenses such as Firewalls and “denial of service” appliances. Nil will normally require just the HTTPS port to be opened on front end servers, and DICOM/communication ports 1100-1104 on backend servers.
- Wireless network security such as usage of appropriate encryption standards and deployment of intrusion detection systems.
- Ensure that OS is updated with security patches and an antivirus is deployed. Perform regular scan for vulnerabilities. Note that NilRead applications and images cache should be normally ‘excluded’ from the ‘real-time’ antivirus protection.
- Physical security to servers and network devices as well as social engineering fraud.
- Incident response plans that can mitigate the risks following a cybersecurity incident.