

NQ Storage

SMB3 Server Implementation For Storage Systems

Table of Contents

1. Executive Summary	3
2. THE ISSUE: LACK OF EFFECTIVE SMB SOLUTION FOR NON-WINDOWS TO WINDOWS INTEGRATION	4
3. THE SOLUTION: NQ™ STORAGE	5
4. ARCHITECTURE OF NQ™ STORAGE	6
5. QUALITY, PERFORMANCE AND SCALABILITY	9
6. COMPLIANCE	10
7. CONCLUSION	11

Table of Figures

FIGURE 1: NQ STORAGE ARCHITECTURE FOR A STANDALONE SERVER	6
FIGURE 2: NQ STORAGE ARCHITECTURE FOR A CLUSTERING STORAGE SYSTEM	7
Figure 4: NQ Storage capacity.....	9

1. Executive summary

The market for applications and data servers is growing fast with the emergence of new Data Centers and transformation of the existing ones. It is also a heterogeneous environment where both Windows and non-Windows platforms are concurrently used. The protocol used for connectivity between these environments is the SMB protocol, and implementing it is essential for Data Centers in a non-Windows environment. Storage providers are thus looking for an SMB solution that offers high performance, scalability, flexible integration and reliable support.

NQ Storage is a highly scalable and portable SMB Server that can be integrated into any storage solution. The objective of NQ Storage is to provide SMB connectivity to those storage vendors that develop on top of a non-Windows platform, but require connecting to a Windows-based environment to benefit from their products.

NQ™ Storage 1.01 is the latest SMB server library developed by Visuality Systems Ltd. Released by Visuality Systems in partnership with Microsoft Corp., this product features high performance, scalability and configurability. It also supports a complete list of server side file transactions, which makes it suitable for any storage system.

2. The issue: Lack of effective SMB solution for non-Windows to Windows integration

A Data Center hosts applications and data servers; the latter also known as Storage. New research from International Data Corporation (IDC) suggests growing expectations of Data Centers from storage solutions.

Today, the growing market of storage solutions is comprised of various environments from big entities, such as Microsoft, EMC and NetApp, as well as numerous small vendors. Most, if not all, of the smaller (though not small) storage vendors use non-Windows platforms such as Linux or UNIX. However, many application servers rely on Windows platforms in order to execute extremely critical applications such as HyperV and SQL Server. These platforms in turn expect the SMB protocol to be the main means of connectivity.

The SMB protocol, also known as CIFS, is a file and printer sharing protocol which serves as the basis for Microsoft's Distributed File System implementation. Thus, implementing the SMB protocol becomes essential to Data Centers using non-windows platforms since they cannot use native protocols to connect to Windows machines.

In this heterogeneous environment, the SMB file sharing protocol becomes the ultimate means of connectivity between Windows-based application servers and non-Windows Storage machines. Therefore, Storage providers look for an SMB solution that offers them the following features:

- a. Provides high performance
- b. Scalable for a wide range of usage and loads
- c. Grants seamless and flexible integration
- d. Comes with effective and reliable support

3. The solution: NQ™ Storage

NQ Storage is a highly scalable and portable SMB Server that can be integrated into any storage solution. The objective of NQ Storage is to provide SMB connectivity to those storage vendors that develop on top of a non-Windows platform, but require the connectivity to a Windows-based environment to benefit from their products.

NQ™ Storage is the latest SMB server for storage systems developed by Visuality Systems Ltd. It extends the SMB range of Visuality Systems' products meant for the storage market. NQ™ Storage inherited its technology from the embedded SMB solution, NQE™, which runs in millions of scanners, printers, home routers and other devices. NQ™ Storage development is dedicated to giving the best performance without compromising on reliability.

NQ™ Storage 1.01 was released on January 31, 2018 by Visuality Systems in partnership with Microsoft Corp. The product is an outcome of Visuality Systems' 20 years of experience in the SMB market with the last four years dedicated to the development of NQ™ Storage.

According to Thomas Pfenning, Director of Development at Microsoft, the collaboration between Microsoft with Visuality Systems “is extremely valuable as it maximizes the SMB support efficiency in the Storage market. The partnership enables storage manufactures to turn their storage system into NAS by the integration of Microsoft's SMB protocol.”

4. Architecture of NQ™ Storage

NQ Storage is a highly portable library with the simple aim to be integrated into a storage solution. It is important to distinguish between three levels of integration:

- 1) Porting
- 2) Standalone integration
- 3) VFS integration

The Porting process occurs when NQ Storage is about to be used on an additional platform, be it another operating system (OS), another CPU or other avenues. Porting NQ™ Storage involves implementing its low layer by means of the most common platform services. This process is seamless and requires minimum efforts, while for Linux/UNIX platforms an off-the-shelf solution is also available. The process of Standalone Integration occurs when NQ Storage is being incorporated into a new solution on a platform for which Porting has been already done. In most cases, this does not require any significant efforts besides fine tuning of a couple of parameters. VFS Integration refers to connecting NQ™ Storage to a storage solution on a higher (semantics) level. This is an advanced process which is only needed for Clustering solutions (see below). NQ Storage can be integrated in either a Standalone Server or a Clustering solution.

NQ Storage as a standalone server

The model for using NQ Storage as a Standalone Server is illustrated in Figure 1.

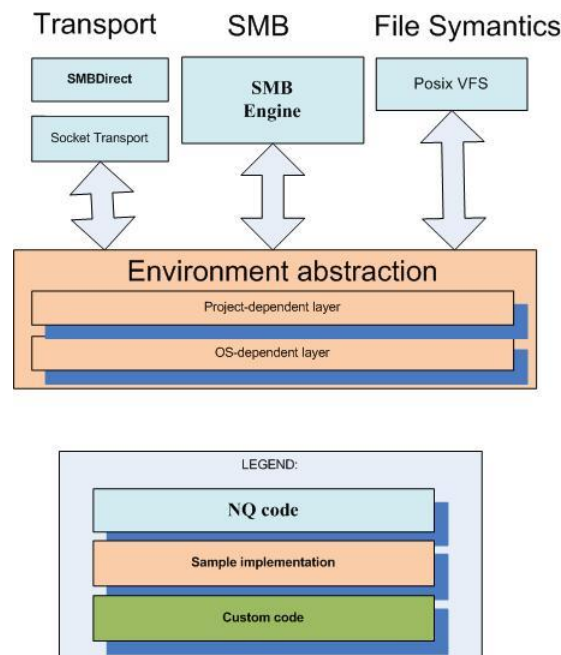


Figure 1: NQ Storage architecture for a standalone server

The SMB Engine module in Figure 1 is the central component of the entire architecture. It is responsible for SMB syntax as well as for some non-file semantics. File semantics are maintained in the VFS module called POSIX VFS. Two transport modules – Socket Transport and SMB Direct – are responsible for networking. The Environment Abstraction component maps an abstract system API on the exact operating system calls. Visuality Systems distinguishes between Project-Dependent (UD) and System-Dependent (SY) layers. Recalling the difference between porting and integration, SY corresponds to Porting while UD corresponds to Integration.

NQ Storage as a Clustering Storage system

Another approach is integrating NQ Storage into a Clustering Storage system as illustrated in Figure 2.

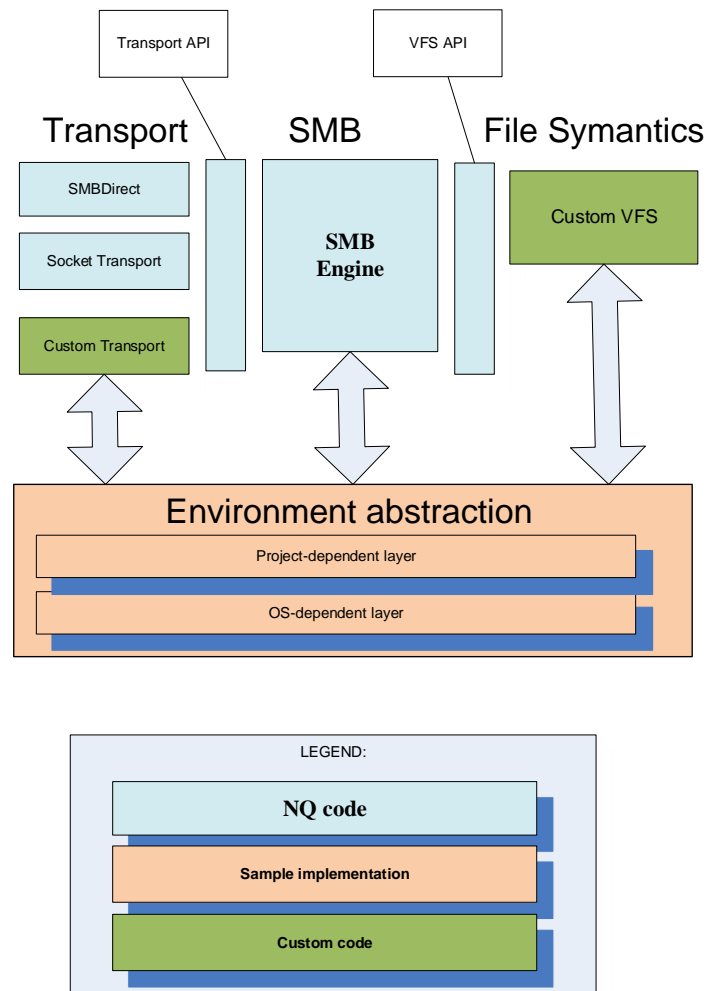


Figure 2: NQ Storage architecture for a clustering storage system

Contrary to the Standalone case, the file system semantics illustrated in Figure 2 crosses not only the protocol borders, but also node borders. Entities such as Persistent Handles, Continuously Available shares, etc., cannot be handled inside one node, and therefore, fall out of NQ Storage's scope. The Porting process is the same as for Standalone, while VFS Integration requires significant efforts. This happens because for a Clustering solution, POSIX VFS must be replaced with a custom VFS implementation. As NQ Storage exposes a well-defined VFS API, implementing a custom VFS involves connecting the SMB Engine with the Clustering File System (FS), which most storage providers already possess. To gain the best performance, most VFS API calls are made asynchronously with completion routines (callbacks). This approach allows the user to fully benefit from multi-threaded architecture (see below).

The entire architecture of NQ™ Storage was designed with performance in mind. Visuality Systems achieved the highest performance in its product by using the following techniques:

- **Thread pools.** Using aspect of threads to the fullest allows complete benefit from SMB concurrency.
- **Internal memory management.** Most buffers are pre-allocated.
- **Zero-copy.** NQ™ Storage avoids copying payload of read and write operations.
- **Server-side operation.** NQ™ Storage uses different techniques to minimize network traffic by implementing server-side file copying, zeroing, remapping, etc. When a client machine is running advanced applications such as HyperV, such procedures boost performance.

NQ™ Storage can enable excellent performance only in conjunction with fast networking. Since networking may introduce latencies, the Socket Transport module should not necessarily use the default BSD sockets. It may be mapped on a user-space solution like DPDK, which grants very low latencies. This is achieved by proper Porting as explained above. By freeing the bottleneck networking may create, Visuality Systems grants an additional performance boost.

The list of transports may be optionally extended with a custom transport module. For this purpose, NQ Storage exposes Transport API. A custom Transport may utilize low-latency networking (like RDMA or DPDK), thus providing the best performance.

5. Quality, performance and scalability

Experienced software developers agree that the strength of a software company as well as the quality of its products depends primarily on the strength of its QA team and the level of technologies used. To achieve the best quality, Visuality Systems' QA engineers execute the following techniques:

- **SMB protocol validation.** Microsoft Protocol Test Suite is used to validate the SMB implementation against Microsoft Protocol Specifications.
- **User experience.** Ranorex™ is used to simulate typical scenarios.
- **Negative testing.** Defensics™ is used to validate NQ™ Storage's ability to handle malformed packets.
- **Load testing.** The Load Dynamix appliance (by Virtual instruments) in Visuality Systems' lab allows simulating up to 10,000 user connections with read/write activities on each of them. With other tools, NQ™ Storage is tested under different types of load.

With performance, Visuality Systems assumes that both the networking and the File System are adequate, and thus, does not create congestions. Of course, NQ™ Storage will perform to its fullest when it is integrated into a platform with fast, low-latency networking. A fast FS is equally important. All the advantages of NQ™ Storage may be lost when using slow HDD instead of SSD. On an adequate platform, NQ™ Storage knows how to fully utilize the SMB traffic with low CPU usage and high throughput. This becomes possible by meeting the client's capabilities, such as SMB concurrency, RDMA, etc. An SMB server implementation can be considered scalable when its performance behaves well on increasing loads. As shown in Figure 3, NQ™ Storage responds with an increasing performance on a growing load, while on higher loads its performance remains stable. As NQ™ Storage is thread-based (contrary to, for instance, Samba, which is process-based), it is capable of handling heavy loads in terms of thousands of client connections, numerous open files, etc. For an example of NQ™ Storage handling read/write load on 10,000 concurrent connections, see [figure 3](#).




Figure 3: NQ v/s Samba on multiple connections

6. Compliance

NQ™ Storage fully complies with Microsoft SMB2, SMB3 and SMB3.1.1 Specifications. We certify this compliance by testing NQ™ Storage with Microsoft Protocol Test Suite. This grants connectivity from all latest client versions of Microsoft, Apple Macintosh and Samba.

Client	Latest SMB dialect	Negotiated SMB dialect
Windows Vista Windows 2008	2.0	2.0
Windows 7 Windows 2008R2	2.1	2.1
Windows 8 Windows 2012	3.0	3.0
Windows 8.1 Windows 2012R2	3.02	3.02
Windows 10 Windows 2016	3.1.1	3.02
High Sierra	3.02	3.02
Samba 4.7.3	3.1.1	3.1.1

Figure 4: NQ Storage capacity

The table in figure 4 demonstrates connectivity between the most common SMB clients and NQ™ Storage. In all cases, NQ™ Storage negotiates the latest SMB dialect supported by the respective client.

Some applications, commonly running in Data Centers, utilize the SMB protocol to its fullest, causing both intensive and comprehensive traffic, including server side operations, notifications, etc. To support this trend, NQ™ Storage complies with such tools as MS SQL Server, HyperV and others to ensure that not only the SMB syntax, but also the least SMB semantic are supported.

7. Conclusion

Nq Storage features high performance, scalability and configurability. It also supports a rich list of VFS transactions, which makes it suitable for any storage system.

NQ™ Storage is available for integration as source code.