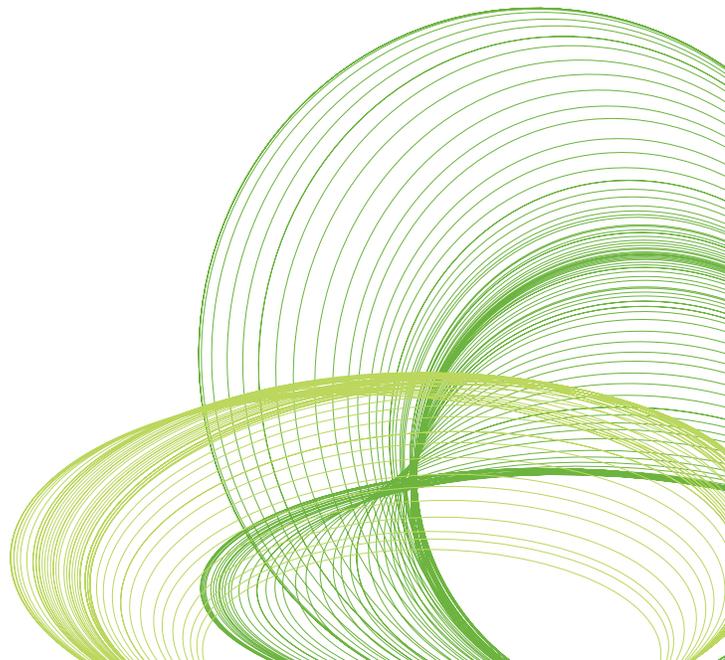




QlikView Technical Brief

February 2012

qlikview.com



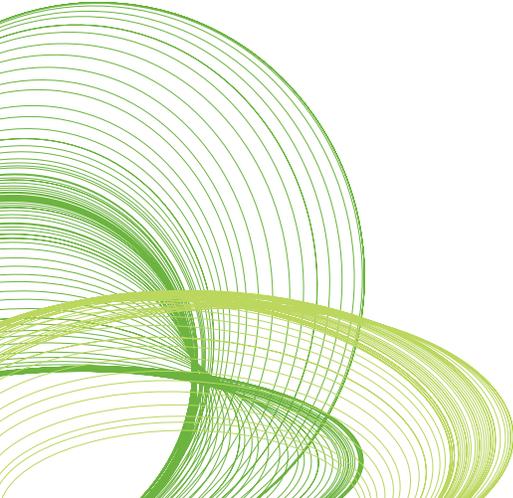
This document will highlight some of the conceptual differences between the terms SAN and NAS and describe how they fit into a QlikView Server environment by suggesting some recommended implementations. It will also describe the definition of DAS (Direct-Attached Storage) and how it applies in SAN configurations.

Table of Contents

SAN - Storage Area Network	4
Host adapters	4
Typical architecture	5
NAS - Network Attached Storage	7
Typical architecture	7
DAS - Direct-Attached Storage	8
• • • • • -€ -	,

QlikView Server and Distribution Service cluster inter-communication	8
QlikView Server file storage conformity	9
QlikView file types explained	9
A note on the usage of .shared files	10
Considerations	10
Conclusions	11
Some suggestions	12
•• -€ - •f•	„••

QlikView Server cluster utilizing a virtual Windows based NAS using a SAN	12
QlikView Server cluster utilizing a physical NAS using a SAN	13
Single QlikView Server using a SAN as local storage	14
QlikView Server using a pure Windows based NAS as local storage	14



Definitions

It is important to differentiate between the two terms SAN and NAS. Also, the term DAS will be involved to nuance the usage of a SAN.

† ‡

A typical definition of a SAN could be;

“SANs primarily are used to make storage devices (such as disk arrays, tape libraries, and optical jukeboxes) accessible to servers so that the devices appear as locally attached to the operating system.” (Wikipedia.org)

SANs are often used in larger virtualization scenarios or where larger areas of high-speed access storage space is required. A SAN typically does not interface with other devices (such as servers) via regular network communication, but shares its resources via a dedicated high-speed data connection that is explicit to its connected clients - hence the use of the word “Network” in this term. This connection can be a variety of types and use different communication protocols, but some of the common ones are (simplified);

- Fibre Channel
- iSCSI
- ATA over Ethernet (AoE)
- HyperSCSI

Depending on what type is used, different types of cabling, protocols and routing equipment are used. For example, Fibre Channel is most commonly used via the FCP protocol and transferred over Fibre Channel cables and switches. iSCSI on the other hand, carries SCSI commands over TCP/IP network, making it possible to create a SAN connection over regular (but dedicated) gigabit Ethernet connections.

SAN solutions are in almost every scenario considered DAS (Direct-Attached Storage) - more on DAS later.

^

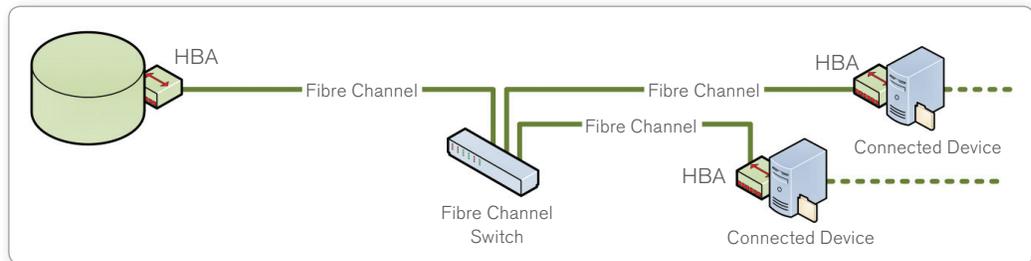
If the SAN is using a kind of connection that requires specific hardware to communicate, the client adapters are usually called Host Adapters, or Host Bus Adapters (HBA). Specific HBAs often (but not necessarily) requires the use of compatible switching and/or routing equipment for communication infrastructure within the SAN, like for example FibreChannel switches or Gigabit Ethernet switches in an iSCSI architecture. In iSCSI, the HBA is the network interface that communicates with the SAN.

There are a multitude of different configuration possibilities in a SAN. The following images could describe a number of typical SAN architectures using Fibre Channel or iSCSI.

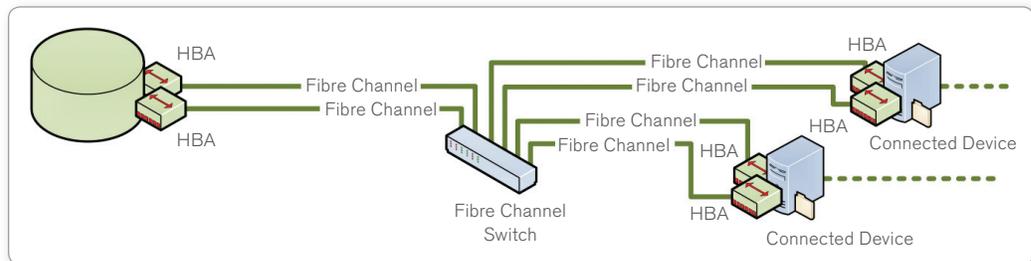
Fibre Channel SAN - Point-to-point topology



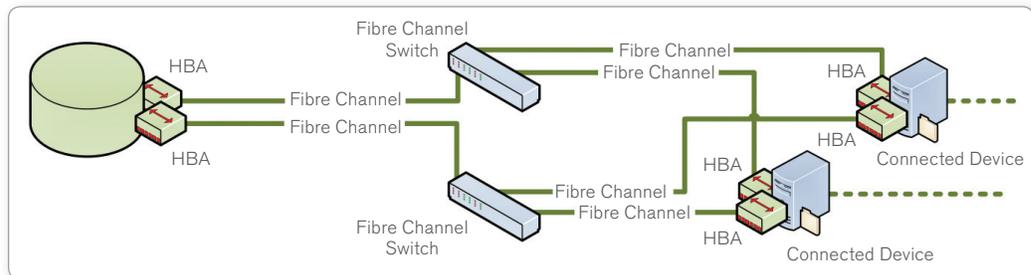
Fibre Channel SAN - Switched fabric topology



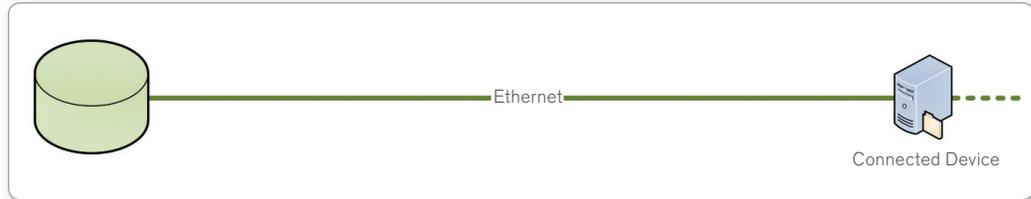
Fibre Channel SAN - Switched fabric topology, arbitrated loop, scenario 1



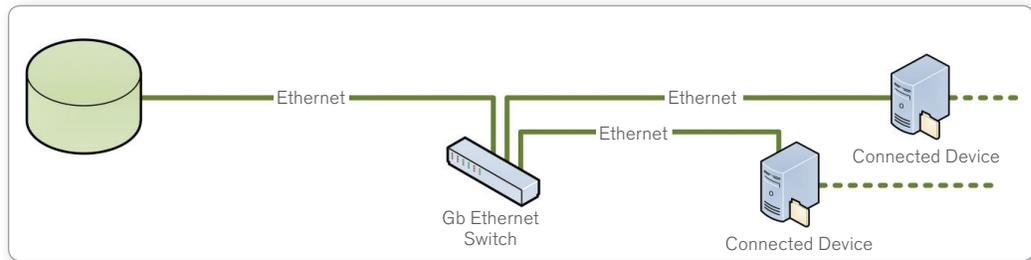
Fibre Channel SAN - Switched fabric topology, arbitrated loop, scenario 2



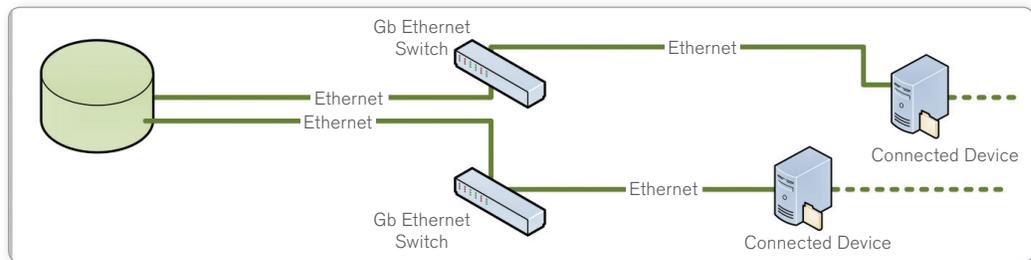
iSCSI SAN - Point-to-point topology



iSCSI SAN - single or multiple client devices



iSCSI SAN - Multipathing topology



† ‡

A NAS per definition is;

“Network-attached storage (NAS) is file-level computer data storage connected to a computer network providing data access to heterogeneous clients. [..]” (Wikipedia.org)

or;

“NAS devices [..] are typically storage arrays with direct attached storage that communicate with application servers using file-level protocols, such as CIFS or NFS.” (Wikipedia.org)

NAS units differ from a SAN in the way they distribute storage to other devices. Unlike a SAN, the NAS shares its storage with standard network protocols, and thus are in no need of specific host adapters other than the standard ones to reach the network. Sharing of the data on the storage is then performed under the governing rules of the NAS file system. Authentication is usually performed on the same level.

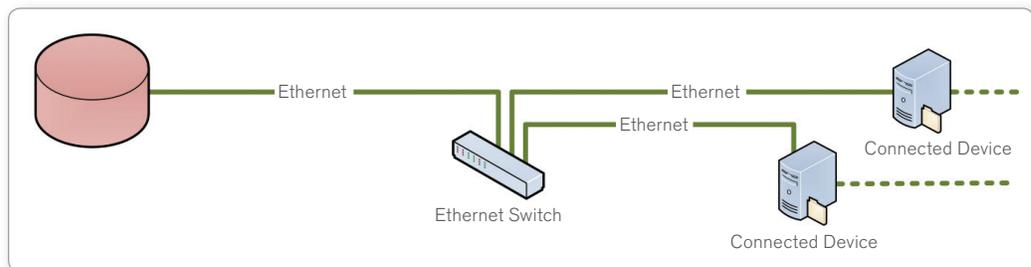
For example, a NAS might be a regular Windows server that is using storage space in a SAN, and then sharing it on the network for other devices to use. A NAS might as well be using internal storage like local disk arrays or similar.

NAS solutions might be considered DAS (Direct-Attached Storage) in certain configuration scenarios.

^

The following image could describe a typical NAS configuration.

NAS - single or multiple client devices



† † ‡

When talking about DAS, it is important to understand that it can be a multitude of different things, but mainly comes down the following definition;

“A typical DAS system is made of a data storage device [...] connected directly to a computer through a host bus adapter (HBA).” (Wikipedia.org)

So, a DAS can describe both a SAN and a NAS configuration, dependant of the architecture used.

For example, when configuring a SAN to communicate with one or more clients devices that has HBAs, and that is directly attached to the SAN, that storage is considered DAS. A NAS can also be considered DAS if for example attached to only one device via a dedicated network connection (usually a dedicated network interface) - which means this device is the only one attached to the storage.

So, in conclusion; DAS is more a concept than it is a type of storage and can be applied in scenarios with both NAS and SAN configurations, and in scenarios where only one device or many access the storage.

Using SAN/NAS and QlikView Server

†

The two main cluster technologies in QlikView Server that we will discuss in this document, are;

1. QlikView Server clustering
2. QlikView Distribution Service clustering

•• -€ - %o Š ‹•f - •

A QVS cluster routes inter-communication between cluster nodes via the server PGO (Persistent Global Object) files. The master PGO files reside in the QVS document root folder. A synchronized backup of the PGO files are always kept in the local QVS app data on each server.

In a QVS cluster, the document root folder must be the same on all nodes. The master PGO files for the cluster will reside in this folder. Since the clustered services rely on instantaneous read and write access to the PGO files, this shared location must be accessible all the time and with high resilience.

•• -œf -€ ‹ %o Š ‹•f - •

The inter-communication between the nodes in a QDS cluster is via files on a shared location. The QDS cluster does not utilize any PGO files and does not share any like resources with QVS, but keeps its own set of files that are in XML format and only has one location where the files reside - the cluster root folder configured for the QDS cluster.

The root folder for the QDS cluster must be configured the same on all nodes. Since the QDS cluster services rely on instantaneous read and write access to the XML files, this shared location must be accessible all the time and with high resilience.

‡

QlikView Server currently only conforms with Windows File Shares. This means that storage must be owned, governed and shared by a Windows instance. The underlying storage type should not matter.

^ Ž^

“€

The Qlikview Document. An on-disk resident file for holding data model, sheets, sheet objects, document settings and document objects such as document bookmarks and reports.

‘ “_ •

Server object container for the .qvw file. Contains server objects such as collaboration objects and server bookmarks. Is always located in the same folder as the .qvw file it belongs to. This file is only used by Qlikview Server.

” •

Meta data container for the .qvw file. Contains .qvw category assignment and other meta data. This file is only used by Qlikview Server and is regenerated at .qvw document distribution.

‘••

Persistent Global Object files for Qlikview Server(s). Contains license CAL data and server/cluster global settings. Used when a cluster contains data and settings that need to be shared amongst all QVS nodes in the cluster.

‡ ‘

Please be aware of the fact that the .shared files are used to a great extent: The .shared files works as supporting storage environment to satisfy a range of functionality in QlikView Server. Amongst those functional areas are the following:

- All bookmarks for AjaxZfc clients (shared or not)
- All shared and/or server bookmarks for all other clients (Desktop, IE Plug-in, and so on)
- “Session Recovery” functionality
- Input fields (when changing values as client against server)
- Collaboration (creating any sheet object as client against server)
- Annotations
- “Email bookmark as link” functionality
- Document chaining (also known as Transfer state)

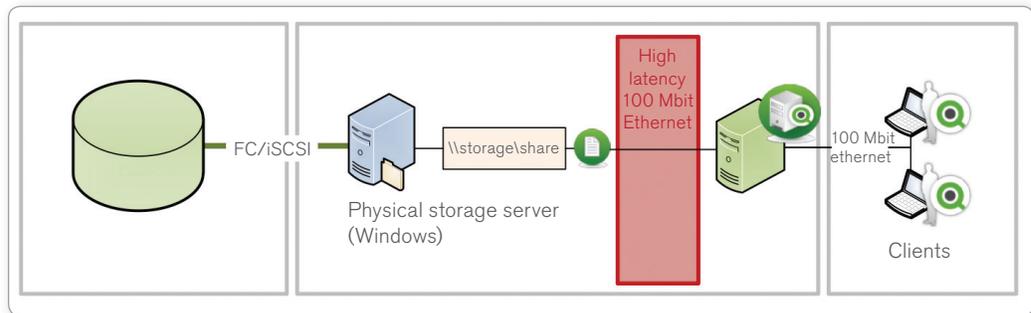
Also, when using an external storage like a SAN attached to a Windows File Share or a Windows based NAS to host QlikView Server resources, it is important to make sure that access times to the storage is on acceptable levels. QlikView Server and Distribution Service clustering requires fast access to files and documents on storage to operate properly. If, for some reason, connection to the storage is rendered unavailable or are experiencing delays or latency, functionality in the cluster will be severely impacted. For example, if a SAN is directly connected to a physical share server (a Windows based NAS, in this case), which in turn is connected to the QlikView Server via regular network, it is not only the DAS connection that needs to be considered. Speed and stability of the network connection between the QVS and the share (Windows based NAS) must also be secured to maintain cluster inter-communication and QVS I/O operations.

Furthermore, if using a SAN (and consequently one or more Windows based NAS units) for multiple QVS- or QDS clusters, make sure that the underlying storage itself does not become overloaded. Having two busy QVS/QDS clusters using separate Windows based NAS units as shares, but where both configurations use the same physical storage location (disk) in the SAN, might cause a performance impact in the form of I/O overload or system choke on the Windows based NAS.

It is also important to note that the QVS cluster root and the QDS cluster root should not be configured to reside in the same folder.

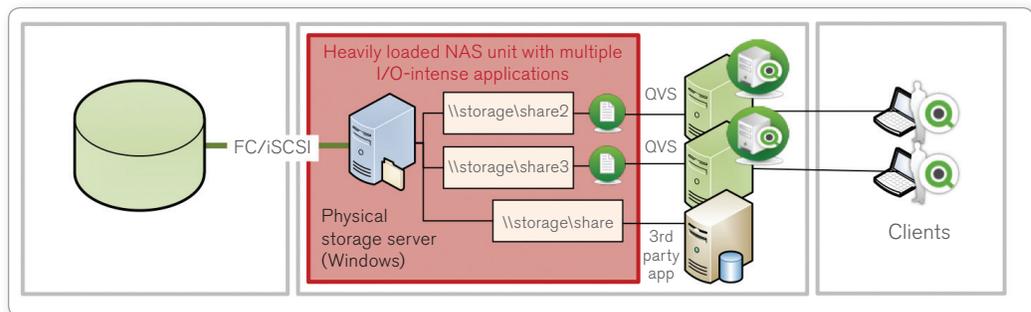
For example, if one theoretically has a 100 Mbit/sec network connection with high latency or packet loss between a QVS and the Windows based NAS, the bottleneck of that configuration would likely be the area marked in red in the design below.

Example of network bottleneck



Likewise, if a SAN is overloaded with I/O or requests, it might become a bottleneck as well.

Example of Windows based NAS overload



It's important to take the following into consideration when using shared or remote storage like SAN/NAS configurations;

- QlikView Server 9, 10 and 11 only supports storage on Windows File Shares.
- Do not point the QVS cluster root and QDS cluster root in the same folder.
- Inter-communication between QVS- and QDS clustered services/machines go via the shared storage - make sure that the connection does not have low throughput, packet loss, high latency or I/O overload.
- Most SAN types, manufacturers and hardware specifications will work in accordance to QlikView requirements - but it's the NAS layer's compatibility that's important. Again, only Windows OS based File Shares are supported.
- Heavy load on multiple Windows based NAS shares that inherit from the same storage space in a SAN might affect QVS performance if multiple QVS- or QDS clusters (or other services) use the same resources.

‡‡

- If there exists a high load in the communication between QVS and the Windows based NAS, consider using teamed network adapters or upgrade to gigabit speed interfaces.
- Verify duplex modes on the Windows based NAS and QVS network interfaces to verify that they are either in a compatible (and highest possible) negotiation duplex mode or set fixed over all connected devices.
- Consider the benefits of separating the cluster inter-communication and the client data communication on separate network adapters - routing the clients on one interface and the QVS/QDS communication on another might lighten the load on one or both interfaces. This configuration may also be combined with teaming of network adapters, under extremely high load.
- Monitor the throughput of all network interfaces and do scheduled reviews of traffic peaks.

QlikView Server Configurations

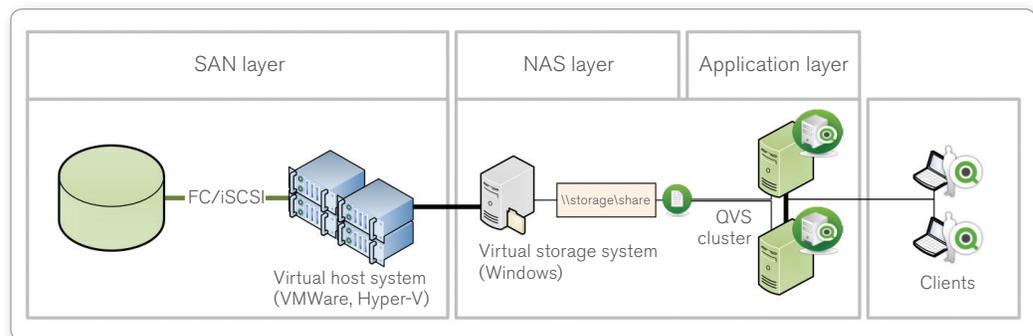
To demonstrate how SAN, NAS and DAS may be combined in different QlikView scenarios, we are going to take a look at some configurations where they are all applied in different ways. All configurations describe scenarios where Windows File Shares are used.

– ‡

‡

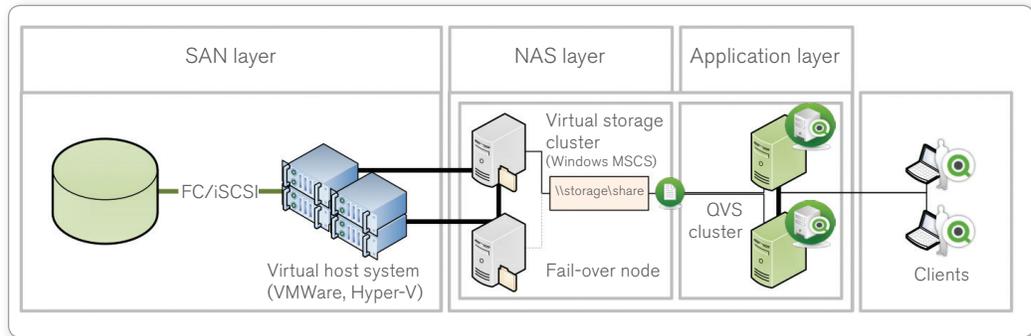
The following configuration shows how a SAN is utilized by multiple devices, in this case physical virtualization hosts. They in turn host a virtual Windows Server that via the virtualization hypervisor uses disk space in the SAN, that they share on the regular network. The QlikView Server cluster is then configured to use the shared disk space to host QlikView documents and cluster data, which the QlikView clients then access via QlikView Server.

Virtual host environment with SAN as DAS, with single virtual NAS



In this same configuration, you can also for example increase redundancy on the storage even further, by using clustering services to create a fail-over Windows based NAS unit in the form of an additional virtual Windows machine and Microsoft Clustering Services (MSCS).

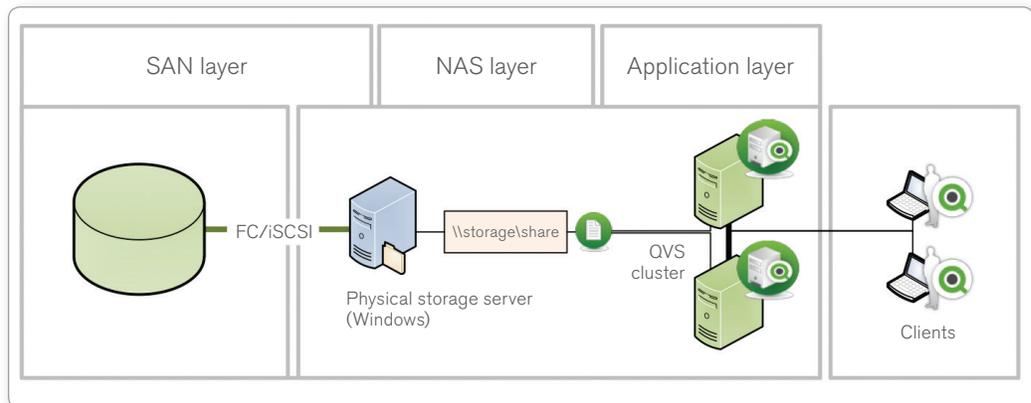
Virtual host environment with SAN as DAS, with virtual fail-over clustered NAS



– † ^ †

This configuration shows how a physical storage server uses the SAN. This will imply that the storage server has the appropriate HBA and is directly connected to the SAN - hence this configuration can be classified as a DAS configuration. The physical storage server uses disk space in the SAN, that it shares on the regular network. The QlikView Server cluster is then configured to use the shared disk space to host QlikView documents and cluster data, which the QlikView clients then access via QlikView Server.

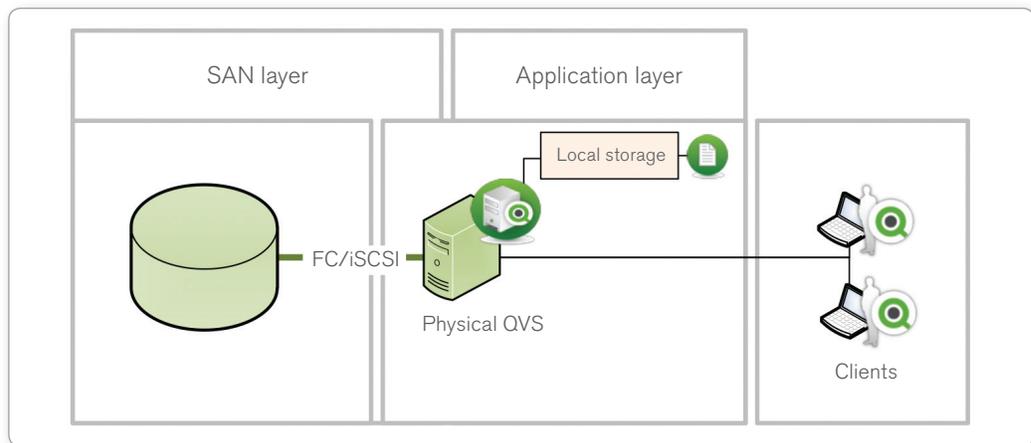
Physical NAS with SAN as DAS



‡ ‡ ‡

Not very common, but in this configuration the QVS is physical and directly connected to the SAN via a HBA, and again may classify as DAS. The QVS has no shared storage on the network (no NAS), but has document on local disk hosted on the SAN. Clients access data via the QlikView Server.

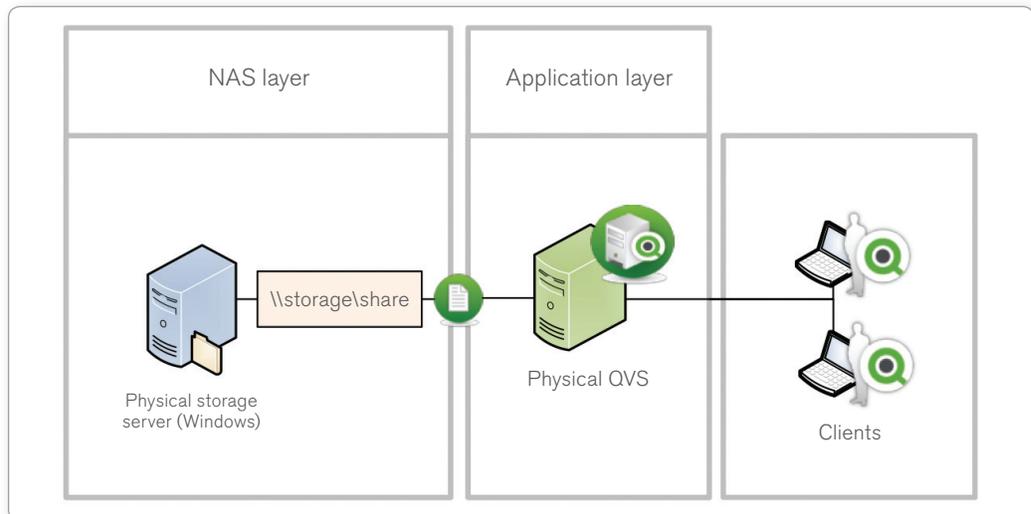
Physical QVS with SAN as DAS, no NAS



‡ ^
‡

In this configuration, a physical Windows storage machine is used as NAS, and no SAN is involved. The NAS is a Windows machine with local disks which shares a folder on the network, that QVS use for documents and cluster files (if clustered).

Physical Windows storage server as NAS



Physical Windows storage server as NAS with QVS cluster

