

# EMC Centera Virtual Archive 1.0

## *A Detailed Review*

---

### **Abstract**

This white paper introduces EMC Centera<sup>®</sup> Virtual Archive and discusses recommended configurations and planning guidelines for Virtual Archive version 1.0. It also summarizes the various components of Virtual Archive and the factors to consider before installing it in an existing EMC<sup>®</sup> Centera environment.

April 2010

---

---

Copyright © 2010 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on [EMC.com](http://EMC.com)

All other trademarks used herein are the property of their respective owners.

Part Number h7062

---

## Table of Contents

<b>Executive summary</b> .....	<b>5</b>
<b>Introduction</b> .....	<b>5</b>
Audience .....	5
Terminology.....	5
<b>Virtual Archive overview</b> .....	<b>7</b>
Federation.....	7
Software .....	7
Host System .....	8
Federated System .....	8
Campus Federation.....	8
<b>Virtual Archive recommended configurations</b> .....	<b>8</b>
<b>Sizing</b> .....	<b>9</b>
Number of access nodes.....	9
Number of management nodes.....	9
Cluster size.....	9
<b>Network and latency</b> .....	<b>10</b>
Bandwidth.....	10
Network Address Translation .....	10
Location of application servers and the Host System .....	10
Campus Federation.....	10
Virtual Archive and LAN .....	11
<b>Replication and failover</b> .....	<b>12</b>
Replication in a Federated environment .....	12
Failover in a Federated environment .....	13
<b>Use case considerations</b> .....	<b>14</b>
<b>Supported use cases</b> .....	<b>14</b>
Capacity expansion .....	14
<b>Use cases not supported for a Virtual Archive environment</b> .....	<b>16</b>
Merging two existing clusters .....	16
Virtual Archive in a WAN setup .....	16
Direct connectivity to a Federated System in a Federated environment .....	17
<b>Conclusion</b> .....	<b>18</b>

---

## List of Tables

Table 1. Terminology.....	5
Table 2. Host System – Access node sizing .....	9

## List of Figures

Figure 1. Virtual Archive data flow .....	8
Figure 2. Campus Federation .....	11
Figure 3. Federation in a multi-LAN environment .....	11
Figure 4. Virtual Archive setup in a unidirectional replicated environment .....	12
Figure 5. Replication setup on a Host System using CLI .....	13
Figure 6. Virtual Archive failover setup using CLI .....	13
Figure 7. Federation failover scenario .....	13
Figure 8. Customer environment before adding Virtual Archive .....	14
Figure 9. Adding a Host System and creating a Federation .....	15
Figure 10. Existing EMC Centera cluster added to a Federation as a Federated System .....	15
Figure 11. Virtual Archive connectivity via WAN is not supported.....	17
Figure 12. Direct connectivity to a Federated System is not supported .....	18

---

## Executive summary

EMC Centera® Virtual Archive is a new EMC technology that enables the aggregation of a set of EMC® Centera clusters. This set of clusters forms a virtual, tamper-proof archive with the capacity of all clusters aggregated and available to applications. Once connected to the Virtual Archive, EMC Centera functionality will be accessible to applications in a seamless manner. Now, capacity and processing power can be added beyond the existing boundaries of a single EMC Centera system or a single data center.

The new Centera Virtual Archive allows customers to federate multiple EMC Centera systems to create a single, highly scalable, distributed digital archive. Most significantly, Virtual Archive simplifies management of archives at scale and distance, delivering significant operating efficiencies and flexible allocation and reallocation of archive capacity.

## Introduction

This white paper provides an introduction to Virtual Archive. It starts with an overview of Virtual Archive and discusses typical use case scenarios, performance considerations, scalability/sizing, network latency, and supported use cases.

## Audience

This white paper is intended for customers, including storage architects and administrators who have a good technical background and understanding of storage technology. Familiarity with the existing EMC Centera components, concepts, and basic understanding of Centera features is recommended.

## Terminology

The following table summarizes the commonly used terms used in this white paper.

**Table 1. Terminology**

Terms	Definitions
ARM	Advanced Retention Management feature of EMC Centera
BLOB	Denotes the finite bit string that is the customer data. A single C-Clip™ can contain pointers to multiple BLOBs. The content of a BLOB is opaque to CentraStar®. BLOBs also have associated Content Addresses, but they are not exposed to the applications accessing EMC Centera. This is to ensure the path to the data is always through the appropriate pointer (C-Clip)
Centera SDK	Centera Software Development Kit
Centera Viewer (CV)	Management tool used for EMC Centera configuration
Centera Virtual Archive	A collection of Federated clusters with Virtual Archive software
Centera Virtual Archive software	The Virtual Archive software that is hosted on one (and later more) clusters
Cluster	A cluster consists of one or more EMC Centera cubes interconnected to present a single storage area
Cube	A single EMC Centera unit containing a minimum of four and a maximum of 32 <sup>1</sup> nodes including a set of cube switches
EMC CentraStar	EMC Centera firmware running on the Centera hardware
EOSL	End of Service Life
Federation	A collection of clusters that together form a virtual cluster

---

<sup>1</sup> Gen3 or older clusters have 32 nodes in a cube. Gen4 and Gen4LP clusters have 16 nodes in a cube.

---

Federation Source	The primary Federation in a replication configuration that is used by the application to store and retrieve content
Federation Target	The Federation in a replication configuration that receives a copy of the content written to the primary Federation by the application
Federated System or System	A cluster that is part of a Federation
Gen4LP	Generation 4 Low Power Centera hardware
Host System or Host	A cluster that is Federated and hosts the Centera Virtual Archive software
LAN	Local Area Network
Replication source or Source	The primary cluster in a replication configuration that is used by the application to store and retrieve content
Replication target or Target	The cluster in a replication configuration that receives a copy of the content written to the primary cluster by the application
Restore source	The cluster where restore is started, from which content is copied to a specified restore target
Restore target	The cluster that receives a copy of the content from a cluster running restore
VIM	Vendor Interface Module
WAN	Wide Area Network

---

## Virtual Archive overview

Virtual Archive is designed to provide the ability to address increasing storage demands by adding more EMC Centera storage in an effective manner. Virtual Archive addresses the need to expand archives into other data centers and increases the abstraction of the technology to also include software versions and other architectures. The latter is needed to ensure seamless access to a long-term archive and to deal with the unknown (not-yet-known new storage technology).

Virtual Archive is software-based and can be installed on any Gen4LP or latest EMC Centera hardware running EMC CentraStar version 4.0 SP2 or later. Virtual Archive version 1.0 allows for scalability of EMC Centera clusters up to 512 nodes.

This section covers the main components and concepts of Virtual Archive used throughout the white paper.

### ***Federation***

A Federation is a collection of clusters that together form a virtual cluster. It is comprised of a single Host System and one or more Federated Systems. After the Virtual Archive software is installed on a standard EMC Centera cluster (Gen4LP or later), your EMC Centera administrator will need to perform a trivial process of creating the Federation on the Host System and then adding additional EMC Centera hardware as Federated Systems. The process of adding an existing EMC Centera cluster to a Federation is explained in later sections.

### ***Software***

Virtual Archive software resides alongside CentraStar and takes the responsibility of routing data traffic between the Host and Federated Systems. When a request is received from an application server, the Host System decides which Federated System will service the request. The request is then routed back via the Host System to the application servers.

Figure 1 shows the Virtual Archive data flow in a Federated environment comprised of a Federation of a Host System and a Federated System.

---

**Note:** Applications will always need to connect to the Host System regardless of its location.

---

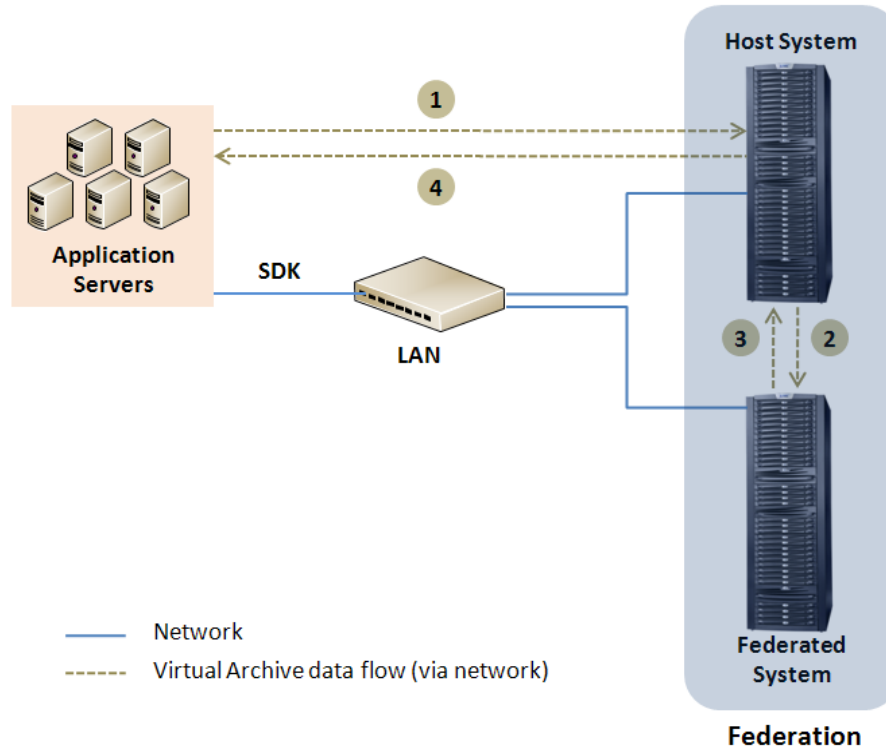


Figure 1. Virtual Archive data flow

## Host System

A Host System is the cluster in a Federation on which the Virtual Archive software is installed. The software is active on all nodes that have an access and/or management role. Applications connect to the Federation through the Host System. Virtual Archive software smartly routes and balances the application data streams across all clusters in the Federation, including itself. With version 1.0 of the Virtual Archive software, only one Host System is supported in a Federation.

With Virtual Archive 1.0, the Host System must be a new EMC Centra cluster that has the Virtual Archive software installed. Virtual Archive software can be installed on Gen4LP hardware (or the latest) running CentraStar 4.0 SP2 or later.

## Federated System

A cluster that is added to a Federation is called a Federated System. This cluster can be elected by the Virtual Archive software to store new content received from the application.

## Campus Federation

A campus is defined as being comprised of data centers that are interconnected and accessible via a LAN. With Virtual Archive 1.0 the Host and Federated Systems need to be located within a certain proximity, and the term used to define the proximity is known as *Campus Federation*.

## Virtual Archive recommended configurations

This section provides a summary of the recommended best practices that should be considered when deploying Virtual Archive.

---

## Sizing

Many factors influence sizing and the scenario could vary for different environments. Factors influencing the sizing decision may include, but are not limited to, the following: application load; total number of access, storage, and management nodes required on a Host System; number of open pool connections by the applications; and number of applications accessing the EMC Centera system. These factors are discussed in this section.

### Number of access nodes

The number of access nodes on a Host System should be planned carefully. The cluster that hosts the Virtual Archive software requires at least the same number of nodes with access roles as found on other clusters in the Federation. If an application is designed to open up many pool connections to the EMC Centera (going beyond the recommended best practice of 25 connections per access node) and has a high volume ingest/retrieval rate, then with the introduction of Virtual Archive, the number of access nodes should be increased on the Host System.

Table 2 summarizes the access node sizing of a Gen4LP Host System.

**Table 2. Host System – Access node sizing**

Gen4LP Centera	# Nodes	# Access Nodes			
Cube A	16	2	4	6	8
Cube B	16	2	4	6	8
Host System	16	4	8	8 <sup>2</sup>	8 <sup>2</sup>

### Number of management nodes

Virtual Archive software uses management nodes for the purpose of keeping the Federation in sync. Multiple management nodes should be configured on each Host and Federated System for redundancy. A minimum of two management nodes are recommended on all systems in the Federation.

### Cluster size

In order to choose the Virtual Archive cluster size users should look at the following areas:

- Application load and number of applications
- Content protection scheme

### Application load

An important factor to consider when sizing an environment for Virtual Archive use is the application load and number of applications. Applications behave differently and ingest/retrieval rates could vary from one application to another. Therefore, understanding application load helps assess whether the EMC Centera system is being used most effectively.

If Virtual Archive is used as a means of adding more capacity, then the number of access nodes may need to be increased to get the same or better performance for the existing applications and application load. The same sizing rules as for a single cluster apply in this case.

When installing Virtual Archive in your EMC Centera storage environment that has multiple applications and/or has high ingest rates (for example, consolidating multiple applications to use a single Federation), it

---

<sup>2</sup> A maximum of eight access nodes are allowed on a 16-node EMC Centera.

---

is recommended that a Host System with a large number of access nodes be used. This will allow having an equal or better performance experience after installing Virtual Archive.

Table 2 has more on access node sizing information.

### **Content protection scheme**

EMC Centera offers two protection schemes, Content Protection Parity (CPP) and Content Protection Mirroring (CPM). CPP stores six data fragments and one parity fragment, and CPM stores data on two unique nodes. A minimum of an eight-node configuration is required for CPP and a four-node configuration for CPM.

As a standard best practice, applications having small objects use CPM and applications having large object sizes use CPP. However, Virtual Archive will not elect a Federated System based on file size. It does not enforce any restrictions on the protection scheme selection and it is not required that the Host and Federated Systems have the same protection scheme. Best practice recommendations for having a CPP or CPM cluster should be followed according to application recommendation.

---

**Note:** More information about capacity and protection schemes is available in the white paper *EMC Centera Capacity Reporting for Gen4LP on CentraStar 4.0 - A Detailed Review* available on <http://www.emc.com> and <http://powerlink.emc.com>.

---

## ***Network and latency***

As a general networking concept, poor network connectivity adds to the latency experienced by end users. This section provides best practices recommendations on how to avoid network issues and latency.

### **Bandwidth**

The network between the clusters in the Federation must accommodate the application writes and reads to the Federation. This means that the network between the clusters in the Federation must support the aggregate bandwidth required by all applications. The best practice for Federations with high application loads is to support twice the aggregate bandwidth required by all applications.

### **Network Address Translation**


EMC Centera Virtual Archive 1.0 does not support the use of Network Address Translation (NAT) between Federated Systems. If NAT is used, all Federated Systems must be on the same logical side of the NAT boundary.

### **Location of application servers and the Host System**

EMC Centera is an IP-based device and hence uses IP network connectivity for all its data operations. Virtual Archive is additional software installed on the EMC Centera and uses the underlying EMC Centera functionality. Applications will only connect to the Host System as it is responsible for making sure that data is routed between the Federated Systems and application servers. Therefore, use general network design practices to keep latency low, and avoid packet loss and congestion.

### **Campus Federation**

With Virtual Archive 1.0 it is recommended that the application servers, Host System, and Federated Systems be located within a company campus. If not, systems performance could be affected due to latency introduced by additional network hops. Figure 2 shows an example of a Campus Federation whereby all the application servers and the Host System and Federated System are located within the same data center and interconnected via a single LAN switch.

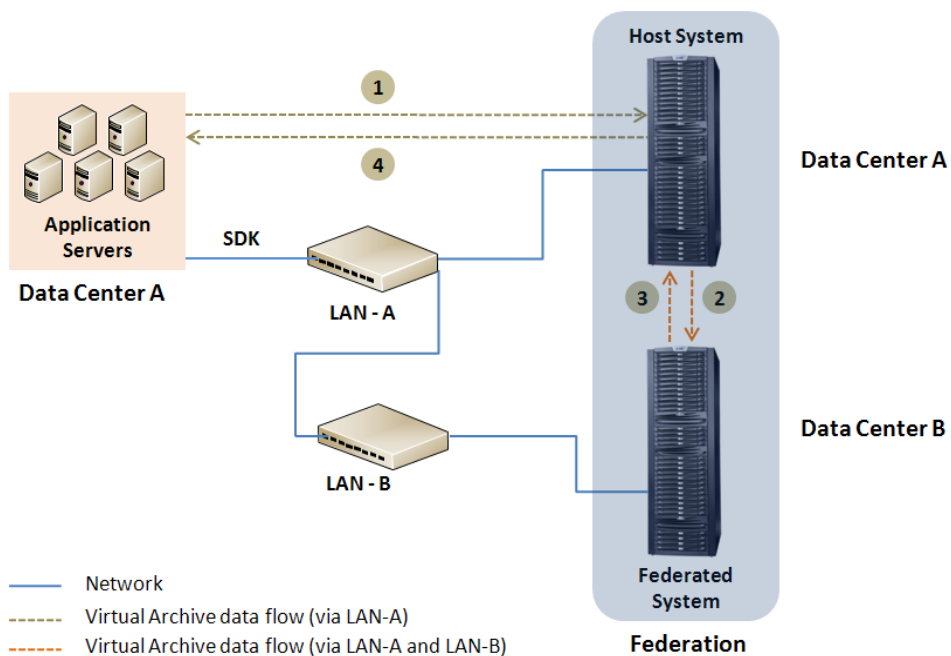


**Figure 2. Campus Federation**

### Virtual Archive and LAN

Within a Campus Federation it is possible to have the Host and Federated Systems located in different data centers provided there is acceptable LAN connectivity.

Figure 3 shows the scenario where the application servers and Host System are located in data center A, and the Federated System is located in data center B. However, since all components are located in the same location and interconnected via LAN, this configuration meets product specifications.



**Figure 3. Federation in a multi-LAN environment**

---


## Replication and failover

EMC Centra allows content to be replicated in unidirectional, bidirectional, chain, and star topologies. However, in the case of Virtual Archive 1.0, the only supported replication configurations are unidirectional and bidirectional. Virtual Archive requires a one-to-one mapping of primary to replica clusters in a Virtual Archive Primary Federation and Virtual Archive Replica Federation.

The following scenario explains the requirements for setting up replication in a Federated environment with Virtual Archive 1.0.

### Replication in a Federated environment

In order to set up replication in a Federated environment your EMC Centra administrator will need to make sure that the Federation Target site has an equal number of clusters available. In other words, there needs to be a one-to-one mapping between the Federation Source and Federation Target site clusters. Figure 4 gives an example of how replication would look in a Federated environment.



**Figure 4. Virtual Archive setup in a unidirectional replicated environment**

In Figure 4, replication will need to be set up between Alpha and Charlie (Host Systems at the Federation Source and Federation Target sites, respectively) and also between Baker and Delta (Federated Systems on the Federation Source and Federation Target sites, respectively).

However, when setting up replication between Alpha and Charlie your EMC Centra administrator will need to make sure that the correct data port is used. After Virtual Archive software is installed, it takes over port 3218 (the data port) and therefore port 13218 should be used when supplying the Federation Target address as shown in Figure 5:

```

Config# set cluster replication
Replication Enabled? (yes, no) [no]: yes
Replication Address: 10.241.44.19:13218,10.241.44.20:13218
Failover Address [10.241.44.19:13218,10.241.44.20:13218]:
10.241.44.19:13218,10.241.44.20:13218
Replicate Delete? (yes, no) [no]:
Profile Name: replication
Location of .pea file [prompt]: C:\replication.pea
Config#

```

**Figure 5. Replication setup on a Host System using CLI**

## Failover in a Federated environment

Failover in a Federated environment is different from the traditional EMC Centera SDK failover in that the failover EMC Centera address needs to be manually entered on the Host System. Figure 6 shows configuration of a failover address on a Host System.

```

Config# federation set failover
Failover address [not set]: 10.241.44.19:13218,10.241.44.20:13218
Issue the command?
(yes, no) [no]:yes
Config#

```

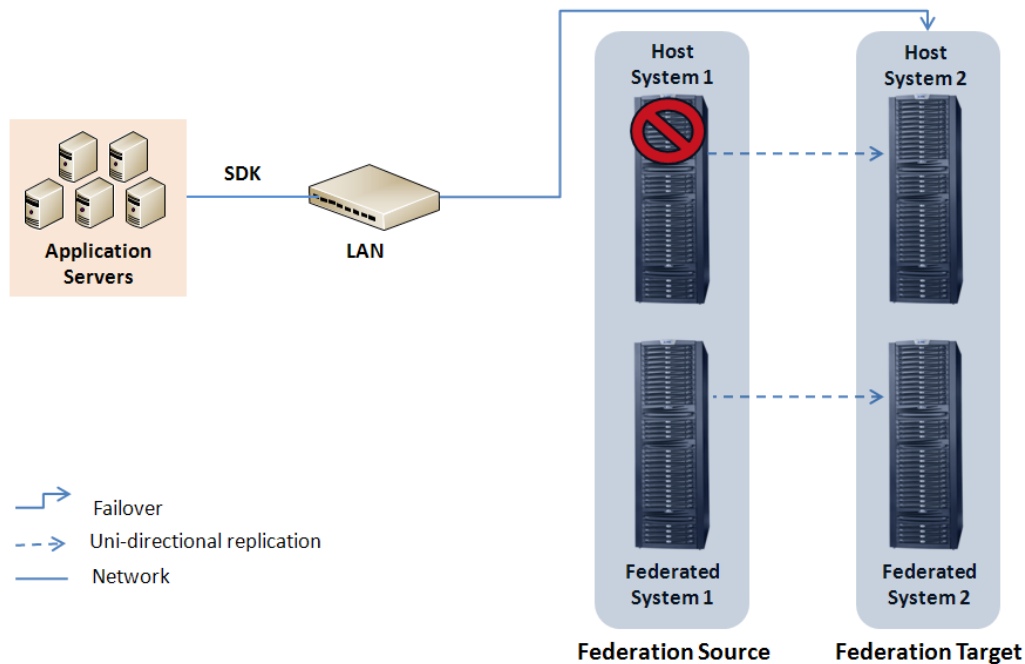
**Figure 6. Virtual Archive failover setup using CLI**

---

**Note:** The failover address is the address of the Host System at the Federation Target.

---

If the Federation Source is inaccessible for any reason, the EMC Centera SDK will fail over to the Federation Target. This is shown in Figure 7.



**Figure 7. Federation failover scenario**

---

## Use case considerations

### ***Supported use cases***

This section provides information on the typical use cases that are supported by Virtual Archive.

---

**Note:** To discuss use cases not listed in this section please contact your EMC Sales Representative.

---

### Capacity expansion

This section will discuss the use cases where Virtual Archive can be used as a means of expanding capacity on the existing EMC Centera cluster.

#### **Adding capacity to an existing cluster**


There is often a need to increase the storage capacity of existing EMC Centera clusters, which is achieved by adding more EMC Centera storage to the existing environment. In most cases you can expand capacity by adding additional nodes to your existing clusters, but if there is a need to expand to multiple cabinets you now have the option to use Virtual Archive. Capacity expansion is the most common use case for Virtual Archive.

#### **Steps for adding Virtual Archive in an existing environment**

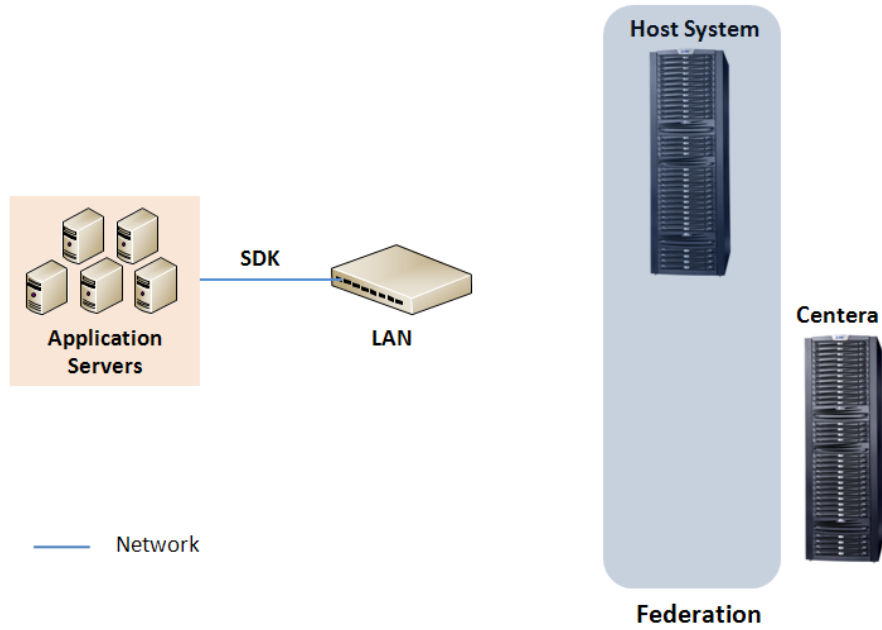
In order to be able to use Virtual Archive in an existing environment, the following steps need to be taken into consideration:

1. Capacity can be added to any cluster in the Federation. However, a Federation can only be created by adding a new cluster to an existing cluster. The new cluster must run the Virtual Archive software. Once the Federation is established, additional new clusters can be added. Note that the new cluster will be the latest hardware and will meet the requirements for Virtual Archive.
2. Configure the new cluster to match the settings of the existing cluster.
3. Install Virtual Archive software on the new cluster.
4. Create a Federation on the Host System and add an existing cluster as a Federated System.
5. Update the application connection string so that it now uses the Host System IP.

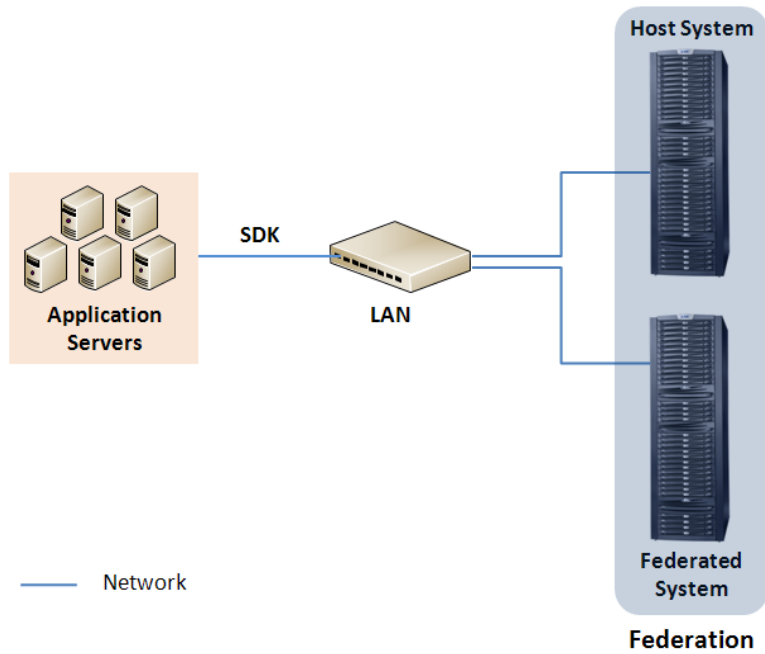
Figures 8 through 10 describe the process of adding capacity by having a new Host System and then adding the existing EMC Centera cluster as a Federated System.



**Figure 8. Customer environment before adding Virtual Archive**



**Figure 9. Adding a Host System and creating a Federation**



**Figure 10. Existing EMC Centera cluster added to a Federation as a Federated System**

**Adding capacity to an existing cluster (replicated environment)**

In the scenario where you have an existing replicated setup (either in uni- or bidirectional replication), Virtual Archive needs to be added at both the Primary and Replica sites.

---

### **Steps for adding Virtual Archive in a replicated environment**

The following steps need to be followed in order to add capacity to the primary and replica sites by using Virtual Archive:

1. Capacity can be added to any cluster in the Federation. However, a Federation can only be created by adding a new cluster to an existing cluster. The new cluster must run the Virtual Archive software. Once the Federation is established additional new clusters can be added. Note that the new cluster will be the latest hardware and will meet the requirements for Virtual Archive.
2. Configure the new cluster (at the primary site) to match the settings of the existing cluster.
3. Install Virtual Archive software on the new cluster (to be used as a Host System at the primary site – the Federation Source).
4. Create a Federation on the Host System and add the existing cluster as a Federated System.
5. Perform steps 1 through 4 for the replica site (Federation Target).
6. Establish replication between Host Systems at the primary (Federation Source) and replica (Federation Target) sites. Use port 13218 when setting up replication between the Host System at the Federation Source and Federation Target sites and provide a failover address
7. Update the application connect string so that it now uses the Host System IP at the Federation Source.

### ***Use cases not supported for a Virtual Archive environment***

The following use cases are not supported.

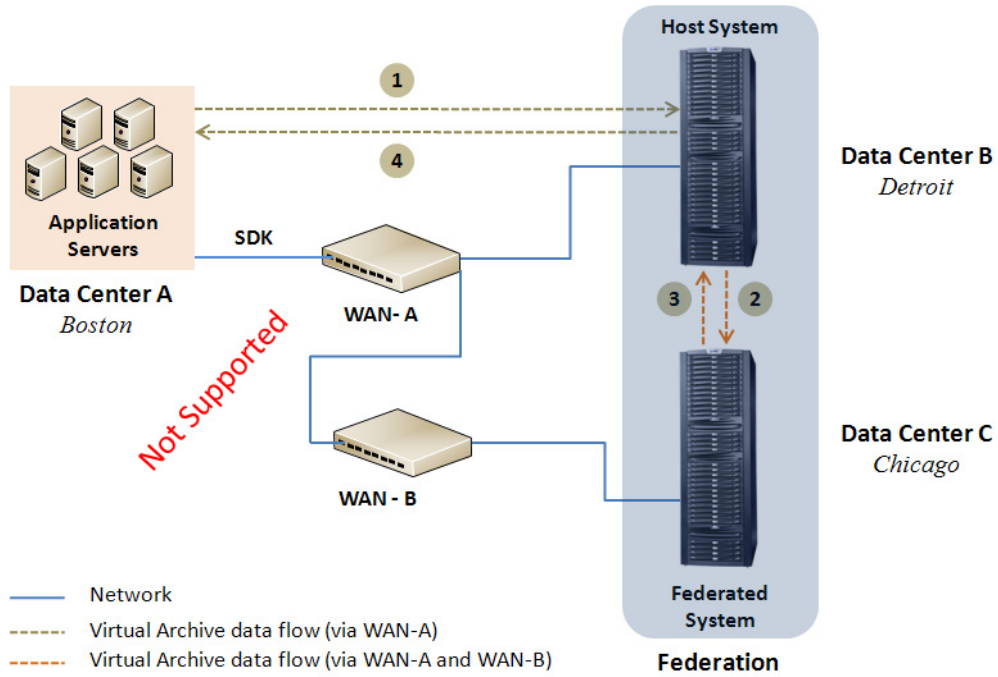
#### **Merging two existing clusters**

With Virtual Archive 1.0, the merging of two existing clusters is not a supported use case. If you fall in this category and would like to discuss the possibility of using Virtual Archive in your existing environment, please contact your EMC Sales Representative.

#### **Virtual Archive in a WAN setup**

Virtual Archive 1.0 is not suited for any clusters that are geographically dispersed (different cities, states, countries, and so on) and require going from a WAN to a WAN connection. Going over the WAN can add huge network latency and could potentially result in a degraded performance. It is therefore not recommended to have Virtual Archive in an environment where the Host System, Federated Systems, and application servers are dispersed across different cities, states, or countries.

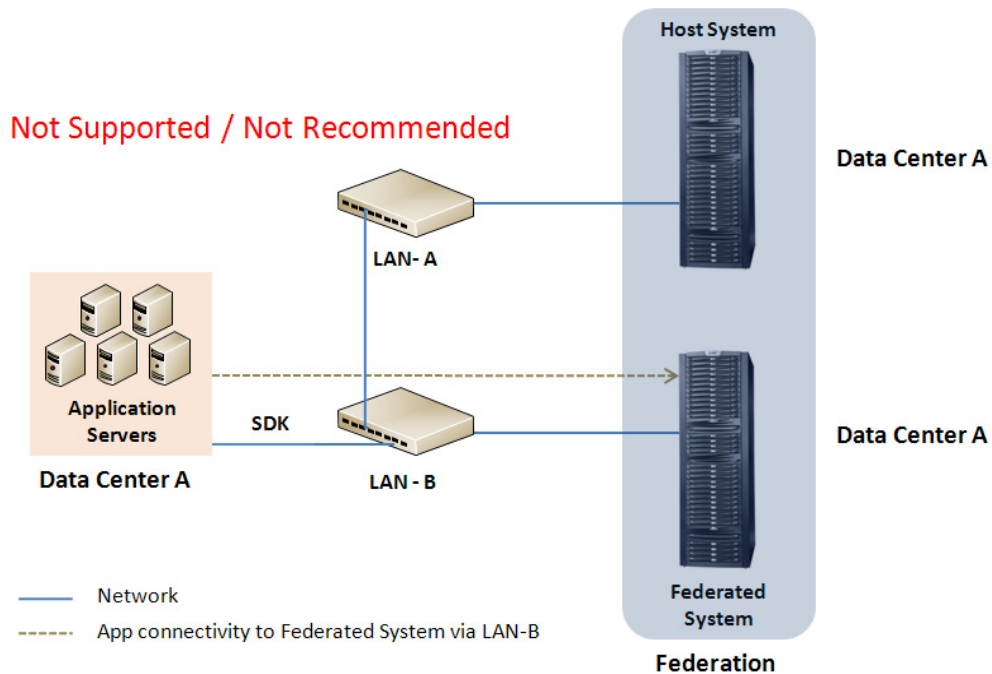
Figure 11 shows an example of a configuration that is not supported and should be avoided.



**Figure 11. Virtual Archive connectivity via WAN is not supported**

### Direct connectivity to a Federated System in a Federated environment

In a Federated environment, the application servers should always connect to the Host System. The Host System is the only cluster that has knowledge about the Federation and therefore is responsible for the routing of data transactions. Although not enforced, it is not recommended for application servers to connect directly to the Federated Systems for data operations as they will not be able to read data that was routed by the Host System to Federated Systems. Figure 12 presents this graphically.



**Figure 12. Direct connectivity to a Federated System is not supported**

## Conclusion

This white paper is intended to provide an introduction and recommended configurations and best practices for EMC Centera Virtual Archive 1.0. It is expected that after reading this paper the reader will have a better idea of this technology and how can it be deployed in their existing environment, and the various factors they need to consider.