IBM Storage Defender

IBM Storage Defender - Clean room environments

Document version 1.0

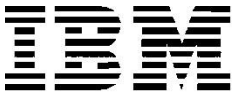

Contents

1 Introduction

IBM® Storage Defender® is an enterprise solution that provides end-to-end data resiliency across primary and secondary workloads. IBM Storage Defender brings together the capabilities that enterprises need to go beyond data protection to real cyber resilience. IBM Storage Defender can perform test recoveries to identify the backup copies that you can trust. To keep your test recoveries in a secure environment, it is preferred that you establish an isolated recovery environment called IBM Storage Defender clean room. This paper provides comprehensive description of the concepts and blueprints for the implementation, and a step-by-step documentation on how to setup an IBM Storage Defender clean room environment.

1.1 Who should read this paper?

This documentation addresses the following use cases:

1. **Understanding the IBM Storage Defender clean room concept:** The terminology and concepts used in the IBM Storage Defender clean room implementation depend on industry standards. To introduce yourself to the terminology and concepts used, refer to the following sections:
   - IBM Storage Defender and the role of the clean room
   - IBM Storage Defender ecosystem
   - Solution architecture
   - Relation between recovery group and isolated environment
   - Introducing IBM Expert Labs and IBM Client Engineering

2. **Implementing a clean room:** The implementation of an IBM Storage Defender clean room environment is based on blueprints that are presented by IBM. To prepare yourself for the implementation of an IBM Storage Defender clean room, you must read all the sections. However, the following sections are important in implementing a clean room:
   - Blueprints
   - Cluster configuration
   - Network configuration
   - Storage configuration

2 IBM Storage Defender and the role of the clean room

The clean room concept and setup that is described in this documentation plays an important role in the IBM Storage Defender concept. The clean room concept helps you to recover workloads into an isolated environment. This setup introduces the ability to safely operate on resources that might be contaminated with viruses, or other malware without the risk of infecting your production environment.

The following figure describes the relation between your production environment and the clean room environment. Virtual machines are protected and get recovered into the clean room for the verification before recovery into the production environment. IBM Storage Defender is connected to all the instances and provides observation and assistance.

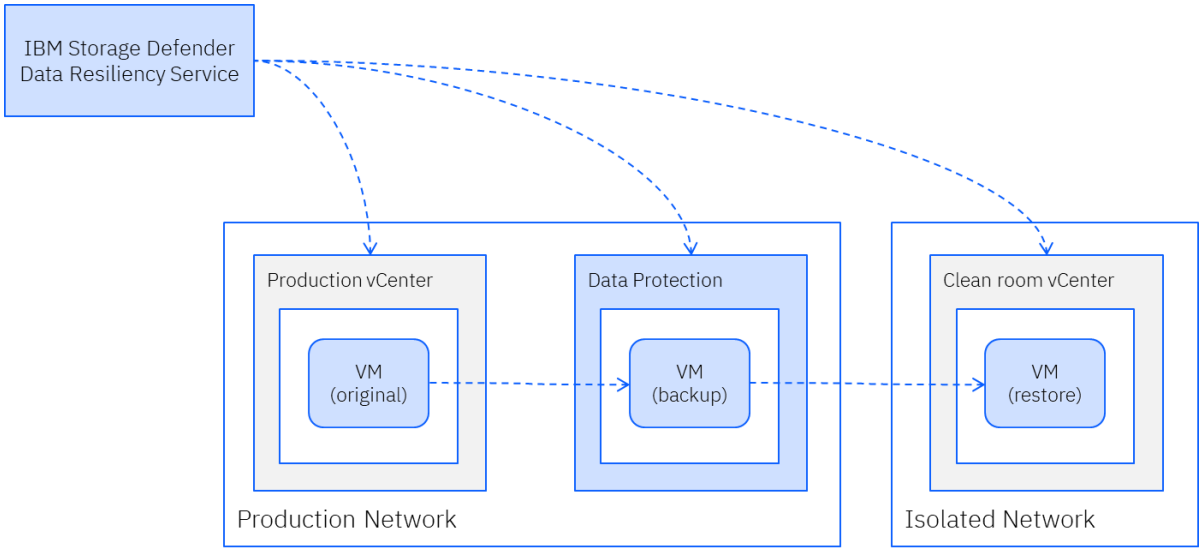

Figure 1: Role of the clean room in the IBM Storage Defender concepts

3 IBM Storage Defender ecosystem

IBM® Storage Defender® is enterprise solution that provides end-to-end data resiliency across primary and secondary workloads. IBM Storage Defender brings together the capabilities that enterprises need to go beyond data protection to real cyber resilience.

3.1 IBM Storage Defender Data Management Service

IBM Storage Defender Data Management Service is a component of IBM Storage Defender. Data Management Service is a cloud-based Software as a Service (SaaS) solution that allows the central management of your secondary data protection environments. This cloud-based portal accesses telemetry about your environment while the data stays on the premises, in the cloud, and in your hybrid data centre environment. IBM Storage Defender ensures that you have control of your data by bringing simple, efficient, and compelling tools that helps you manage both the storage and cyber-resiliency needs of the data.

3.2 IBM Storage Defender Data Resiliency Service

IBM Storage Defender Data Resiliency Service is an optional component of IBM Storage Defender. Data Resiliency Service is a combination of cloud-based Software as a Service (SaaS) managed by IBM, and an agent that manages communications from your data centre. The data centre agent is called Connection Manager which collects telemetry about your primary and secondary data while the data stays on premises. The telemetry communicated to the Data Resiliency Service helps you to plan, secure, and recover your important data.

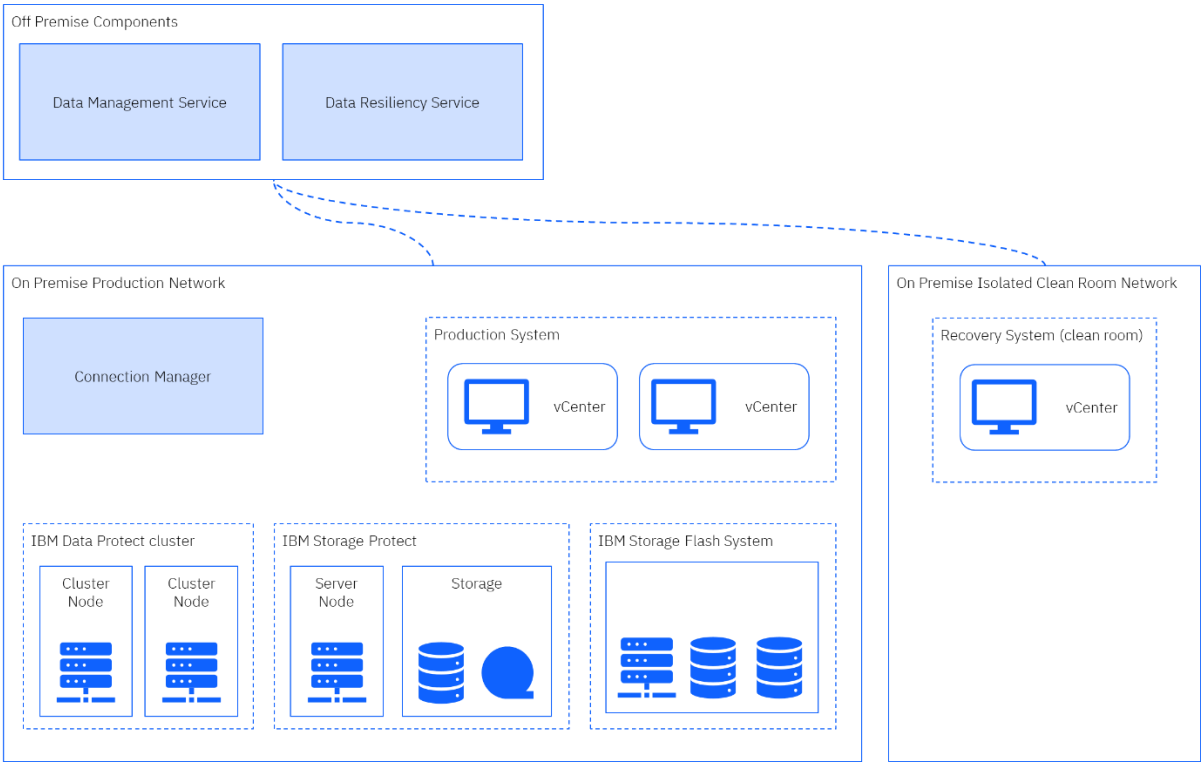

Figure 2: IBM Storage Defender ecosystem

3.3 Recovery Group concept

With the Recovery Group concept, IBM Storage Defender introduces a logical construct that you can use to group resources that belong to the same application. By defining multiple Recovery Groups, you can reflect complex applications and dependencies between applications when it comes to recoveries.

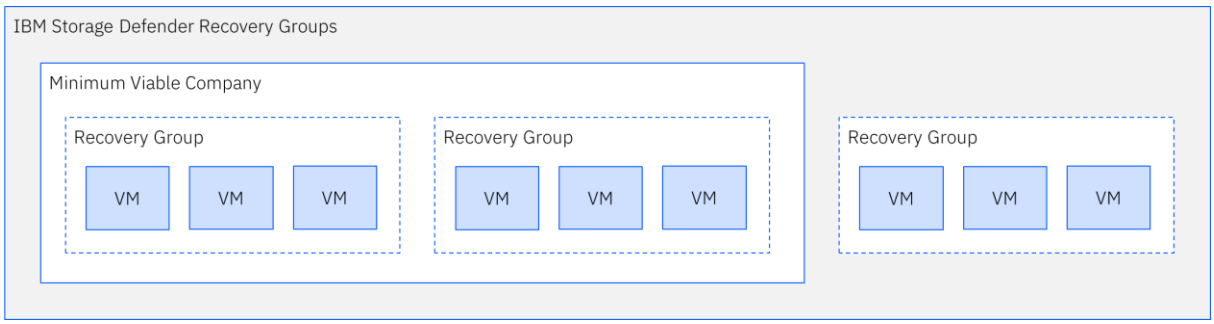

Figure 3: Recovery Groups & Minimum Viable Company

3.4 Recovery system versus lean room environment

IBM Storage Defender requires the setup of a system that can be used as a target for recoveries. In general, two different types of recoveries can be described: Recoveries that are used for the verification and test of backup copies, and recoveries after a cyber-attack that are used for temporary production.

The test recovery has a short lifecycle with small footprint, and it is typically limited to a single application. The temporary production recovery contains all the applications that you define as required for a temporary production of your company's IT infrastructure.

This environment can be described as recovery system or clean room environment. While a recovery system can be a physical part of the production environment, a clean room is defined as an isolated environment. This isolated environment prevents infection from the compromised production environment and cannot infect the production environment when compromised backup copies are recovered into it. In addition to physical separation, the separation of duties are typically implemented in a clean room environment to limit the influence of a malicious user.

3.5 Connecting IBM Storage Defender to a clean room environment

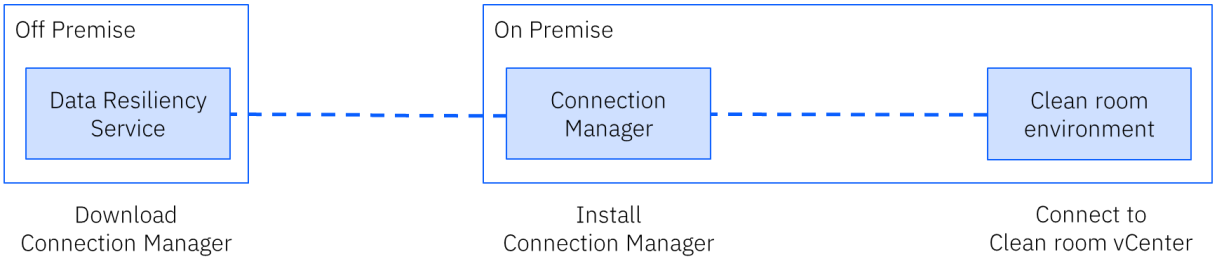

Figure 4: Connecting IBM Storage Defender to clean room

To connect the clean room environment to your IBM Storage Defender tenant, complete the following steps:

1. Download and install Connection Manager
   a) Log in to IBM Storage Defender and download the Connection Manager OVA software and generate the connection token.
   b) Download the certificate chain and install it in your vCenter Certificate Management.
   c) Install the Connection Manager.
   d) Update the SSH password of the Connection Manager.
   e) Connect the Connection Manager to IBM Storage Defender Data Resiliency Service.
2. Connect to clean room vCenter
   a) Log in to the Connection Manager.
   b) Add new clean room environment by using the IP or FQDN of the clean room vCenter and enter the username and password of the identity you want to use.

4. Solution architecture

A clean room environment setup has several similarities with a standard vCenter setup. The following figure displays the high-level structure of a clean room environment. Apart from the recovery groups that are restored by using datastores that are mapped from data protection solutions, a DMZ is implemented to allow access to the isolated portions of a clean room.

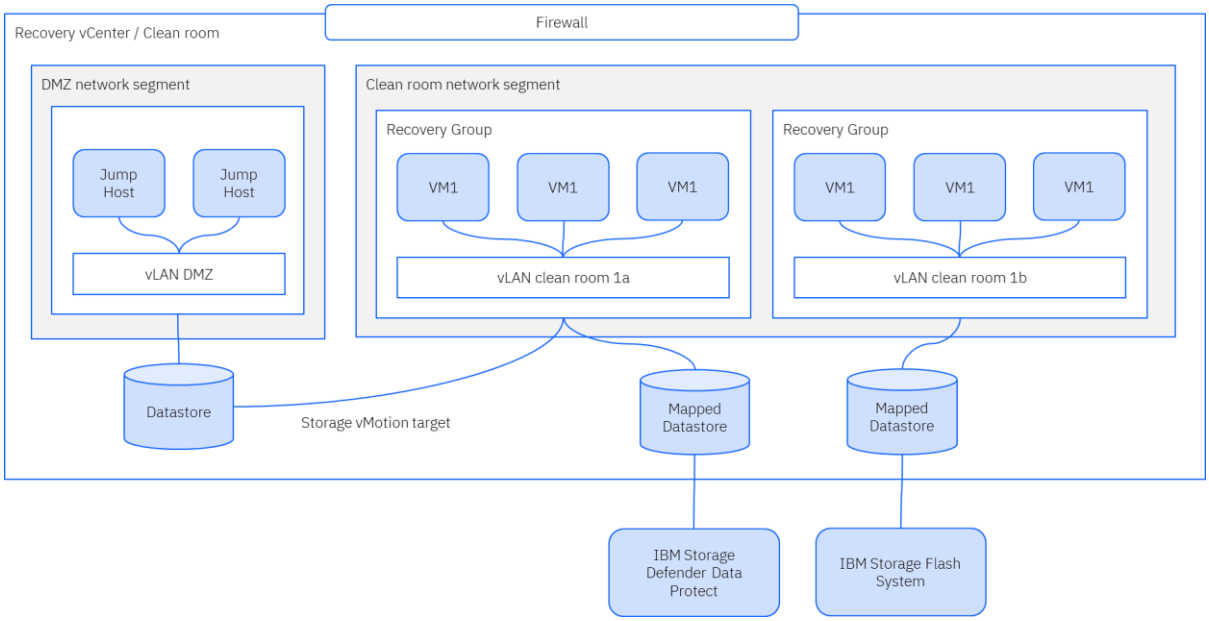

Figure 5: Clean room environment schema

Isolation is an important aspect when you implement a clean room. You need to consider isolation in multiple dimensions, such as isolation of infrastructure, network, and access management. In addition to the isolation, you must implement monitoring and logging of a clean room environment. The following sections describe the different aspects of isolation for a clean room environment.

4.1 Infrastructure isolation

The isolation of the infrastructure is an important aspect of a clean room environment. Isolation can refer to physical separation in form of a set of computer hardware that is used for a hypervisor that is independent from any production environment. When a cloud service provider is used, isolation refers to a logical separation in form of different cloud accounts.

4.2 Network segmentation and monitoring

Network segmentation has multiple aspects that are as follows:

- **Logical separation and subnetting:** In addition to the recovered virtual machines, the clean room environment contains systems that are used for tools and management. You need to separate groups of systems into network segments to prevent the breakout of malware from infected systems. If multiple recovery groups are recovered into the same clean room to establish a temporary production environment, you must use a dedicated VLAN for each recovery group. Apart from the breakout prevention, the positive impact on the administrative separation of duty is another important factor in this context.
- **Access control and firewalls:** Use firewalls and access control lists (ACLs) to control and monitor traffic between network segments. In addition, enhance security by enforcing rules that are based on source, destination, and port.
- **Security zones and critical infrastructure protection:** Establish security zones, including a De-Militarized Zone (DMZ), to separate public-facing servers and protect critical infrastructure components by limiting potential attack vectors.
- **Monitoring, encryption, and regular auditing:** Implement network monitoring tools and centralized logging to make sure visibility and timely detection of security incidents. In addition, implement secure communication between recovery groups in the same

clean room. If applications require interaction, you can use VPNs and encryption. If the clean room is used for temporary production, you need to conduct regular security audits.

4.3 Identity management and logging

The administrative separation of a clean room environment from a production environment helps to implement extra security protocols. It can be a different set of administrative identities or a total separation of the identity management in a separate directory service.

The logical separation of administrative roles for the production system and the clean room environment and strict limits on the user's permissions prevent a user to influence both environments.

The implementation of auditable logging for all operations in the clean room makes sure that any operation on the recovered data is traceable. It includes the creation and configuration of the clean room, clean room operations, such as recovery, data masking, anonymization, or temporary production usage of the data.

4.4 Compliance and legal compliance

The bounded usage scope of a clean room environment allows comprehensive documentation of all operations in the clean room. In addition, you need the auditable logging that is mentioned earlier in the configuration of the clean room environment. The usage scope such as temporary production or test recovery and the operations on the data in the cleanroom, such as analysis or development can be documented. A comprehensive documentation reveals the regulatory compliance requirement of a company.

5. Relation between recovery group and isolated environment

IBM Storage Defender Data Resiliency Service implements the logical construct of a recovery profile. The recovery profile defines the logical connector between the recovery group and the clean room network segment to be used for recoveries.

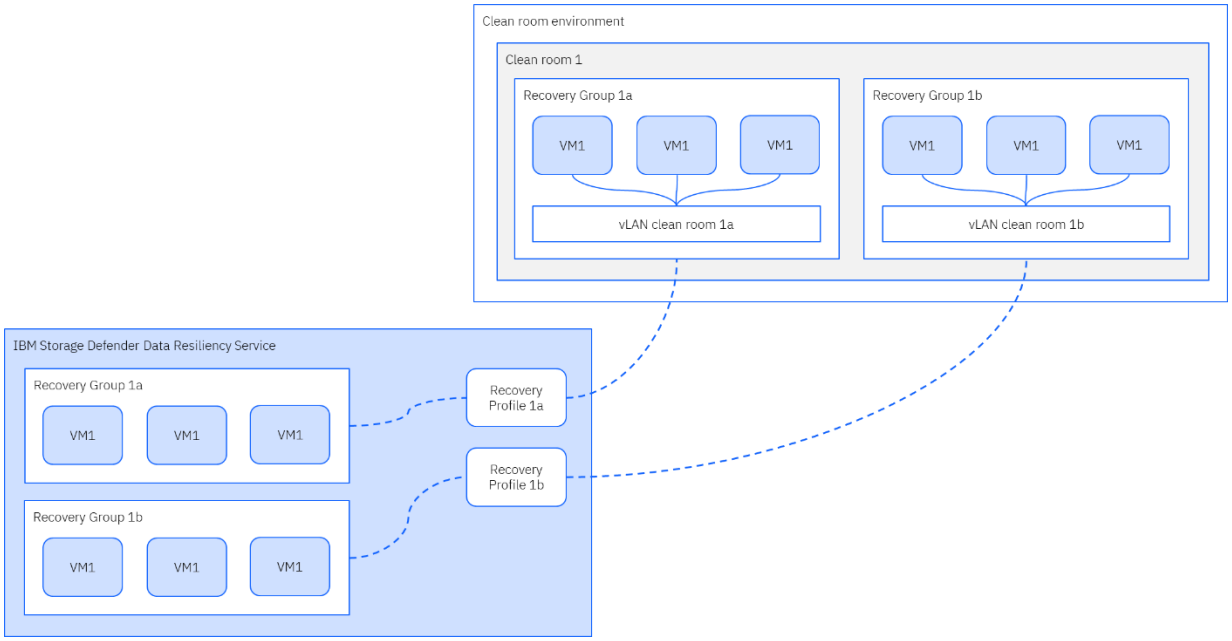

Figure 6: Relation between recovery groups and isolated environment

You can configure multiple recovery profiles in IBM Storage Defender Data Resiliency Service. You can use each recovery profile multiple times. You can use the same clean room VLAN network for multiple recovery groups. To configure the environment more granular, you can use one dedicated VLAN per recovery group.

For temporary production clean rooms, you need to rebuild the configuration that you use in your production environment. When you isolate virtual machines of a recovery group in your production environment with a dedicated VLAN setup, implement the same configuration in the clean room.

In some cases, you need to bring multiple recovery groups into the same VLAN. For example, the recovery groups that have dependencies and must run in the same network segment when you start the virtual machines.

For test recovery clean rooms that are used for recovery validation for a limited time frame, you must prevent the multiple recovery groups to use the same network segment at the same time frame. It prevents cross recovery group infections.

You can configure the recovery groups and recovery profiles for the test recovery use case. At the same time, you can prepare multiple recovery profiles for the temporary production clean room. When you need the recovery of multiple recovery groups into the temporary production clean room, you need a manual recovery operation in the Data Resiliency Service. You can switch recovery groups that are configured for test recovery profiles to temporary production profiles at the time of the operation. To complete the manual recovery operation, see IBM Documentation[1].

Figure 6 shows a sample configuration with two recovery groups that belong to dedicated recovery profiles and are restored into dedicated clean room network segments at the time of the recovery.

[1] IBM Documentation: Initiate the activation of a recovery point

6. Blueprints

You can implement clean room environments with different capacity according to as your capacity requirements. To help you with decision making, the four blueprint sizes are precalculated.

6.1 Considerations

6.1.1 Shared versus dedicated vCenter

Do you plan to implement a dedicated vCenter for the clean room to make sure resource isolation on VMware level?

In this case, you need to size the clean room and add capacity that you need for the vCenter appliance. The blueprints that are listed in the following table assume that the first host in the cluster contains the vCenter appliance.

Note: The goal of the blueprint is to provide a step-by-step documentation to implement an isolated environment. You can consider installing the vCenter appliance on a host that is independent from the clean room ESXi hosts. It is helpful when the clean room environment itself must be recovered. When you use the clean room for temporary production, such separation can be useful as it introduces a higher level of recoverability of the environment.

6.1.2 Recovery group size

How many VMs do you have in one recovery group?

The number of VMs per recovery group have a direct impact on the clean room size. If multiple VMs are recovered at the same time into a clean room, the CPU, RAM, and storage resources must be considered. The blueprints that are listed in the following table defines the approximate number of VMs that are used for testing the setup.

6.1.3 Shared versus dedicated clean room

Do you plan to use the clean room for one or multiple recovery groups?

Similar to the recovery group size, clean room for single or multiple recovery groups has an impact on the clean room size. You need to calculate the total number of VMs you want to recover into a clean room by multiplying the number of VMs per recovery groups with the number of recovery groups. You can consider implementing the multiple clean rooms with different sizes to achieve the optimal setup model.

6.1.4 Lifecycle of VMs in the clean room

The clean room can be used for test recoveries and for temporary production. Depending on how you want to use the clean room, VMs are recovered to the clean room concurrently. It impacts the resource requirements for the clean room.

6.2 Configuration

The calculation examples that are listed in Table 1 describe the different options into consideration. The calculations based on the number of VMs that are recovered concurrently into the same clean room environment.

Table 1: Clean room configuration blueprints

Clean room size	**Small**	**Medium**	**Large**	X-Large
Number of VMs concurrently in clean room	100	200	500	1000

Number of ESXi hosts	1	2	5	10
Number of jump hosts per clean room	2	2	2	2
VMware vCenter appliance size	Small	Small	Small	Small
Physical RAM per host if shared with vCenter appliance (GB)	512 GB	512 GB	512 GB	512 GB
Physical RAM per host if guest systems only (GB)	n/a	512 GB	512 GB	512 GB
Physical CPU cores per host if shared with vCenter appliance	48	48	48	48
Physical CPU cores per host if guest systems only	n/a	48	48	48
Number of physical network adapter per host	2x10 GbitE	2x10 GbitE	2x10 GbitE	2x10 GbitE
Storage per host if shared with vCenter appliance (GB)	6243 GB	11222 GB	26428 GB	51977 GB
Storage per host if guest systems only (GB)	n/a	10528 GB	25734 GB	51283 GB
Overall storage per clean room	6358 GB	21750 GB	129366 GB	513526 GB

6.3 VMware vCenter appliance capacity considerations

The vCenter appliance deployment size relates to the number of VM restored to a clean room concurrently. For clean room environments that are beyond the blueprint size X-Large or for scalability in your production environment, you can consider larger vCenter appliance deployments. Table 2 describes the VMware calculations.

Table 2: VMware vCenter appliance capacity

VMware vCenter appliance size[2]	**Small**	**Medium**	**Large**	**X-Large**
Number of host / virtual machines	100/1000	400/4000	1000/10000	2000/35000
CPU	4	8	16	24
RAM (GB)	19 GB	28 GB	37 GB	56 GB
Default Storage (GB)[3]	694 GB	908 GB	1358 GB	2283 GB

7. Cluster configuration

The section describes the sample implementations of the clean room architecture. The following aspects are considered in clean room implementation architecture:

- The data that is recovered into the clean room has two sources, IBM Storage Flash System and IBM Storage Defender Data Protect. To allow recovery from both sources, Fiber Channel Adapters, iSCSI, and Ethernet adapters are included.
- The clean room is connected to the client private ethernet network and storage area network to allow management and temporary production.

[2] VMware Documentation: Hardware Requirements for the vCenter Server Appliance
[3] VMware Documentation: Storage Reuirements for the vCenter Server Appliance

- Both sources export the volume that contains data, and these exported volumes are integrated as datastores into the VMware infrastructure.
- The Jump (Bastion) Host systems are part of the clean room environment and require their own datastore space.
- The segmentation of the networks is realized implementing VLAN and vSwitch instances.

Note: Your implementation must resemble your production environment. If application runs in a dedicated VLAN in your production environment, you need to configure a dedicated VLAN for that application in the clean room.

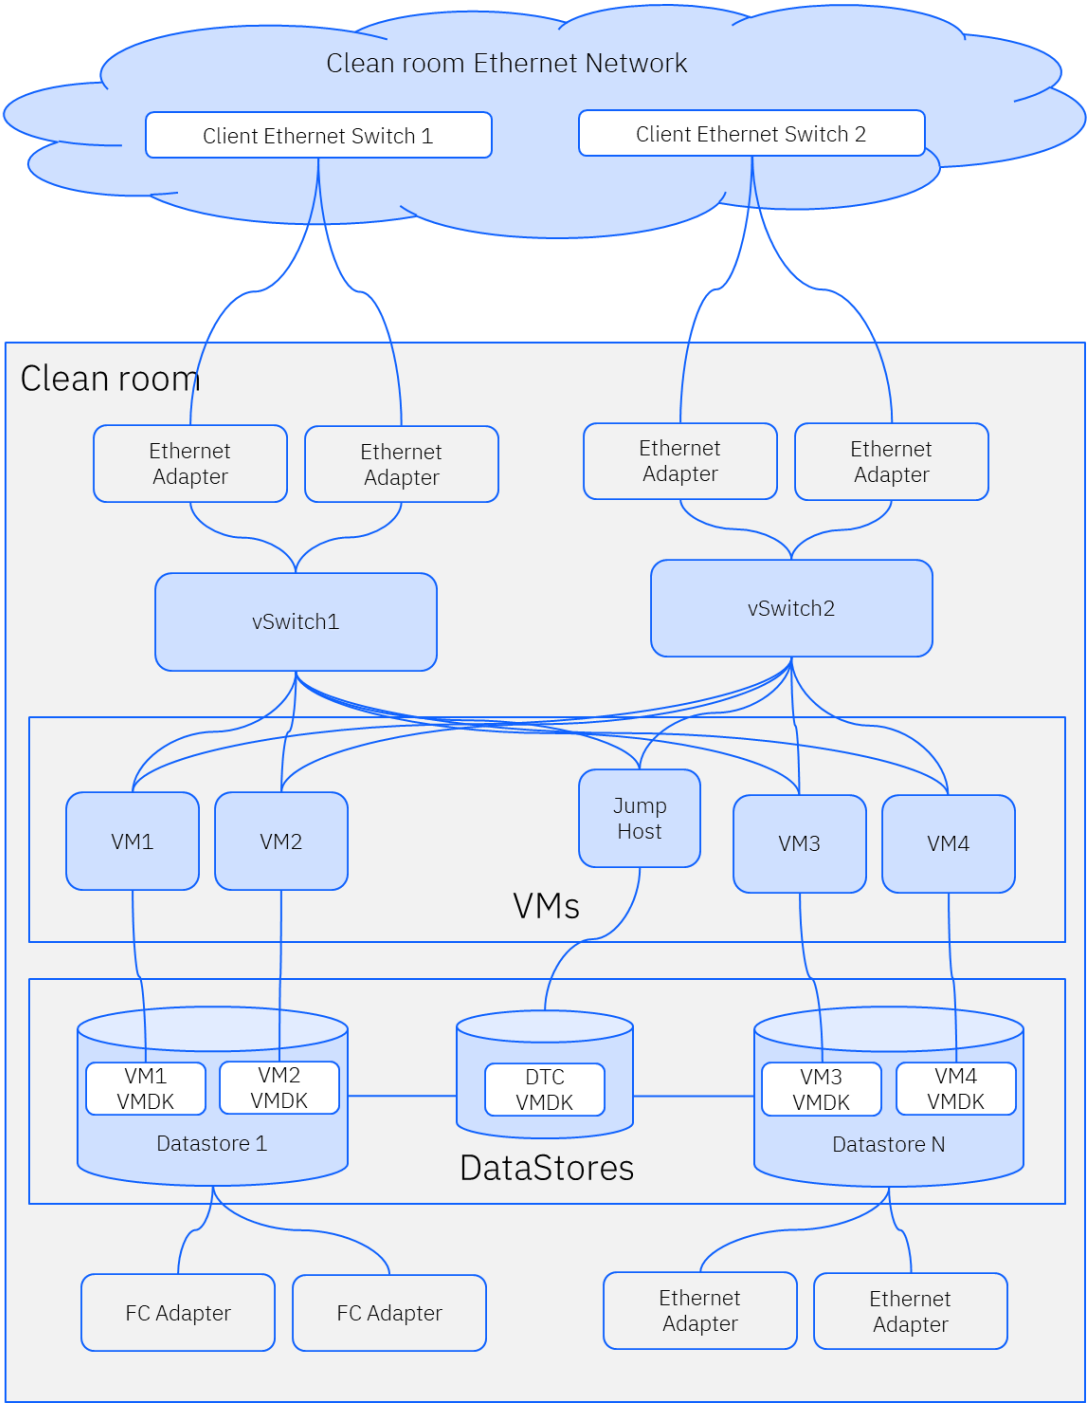

Figure 7: Clean room sample implementation

The following figure shows implementation of a dedicated ESXi cluster for clean room operations.

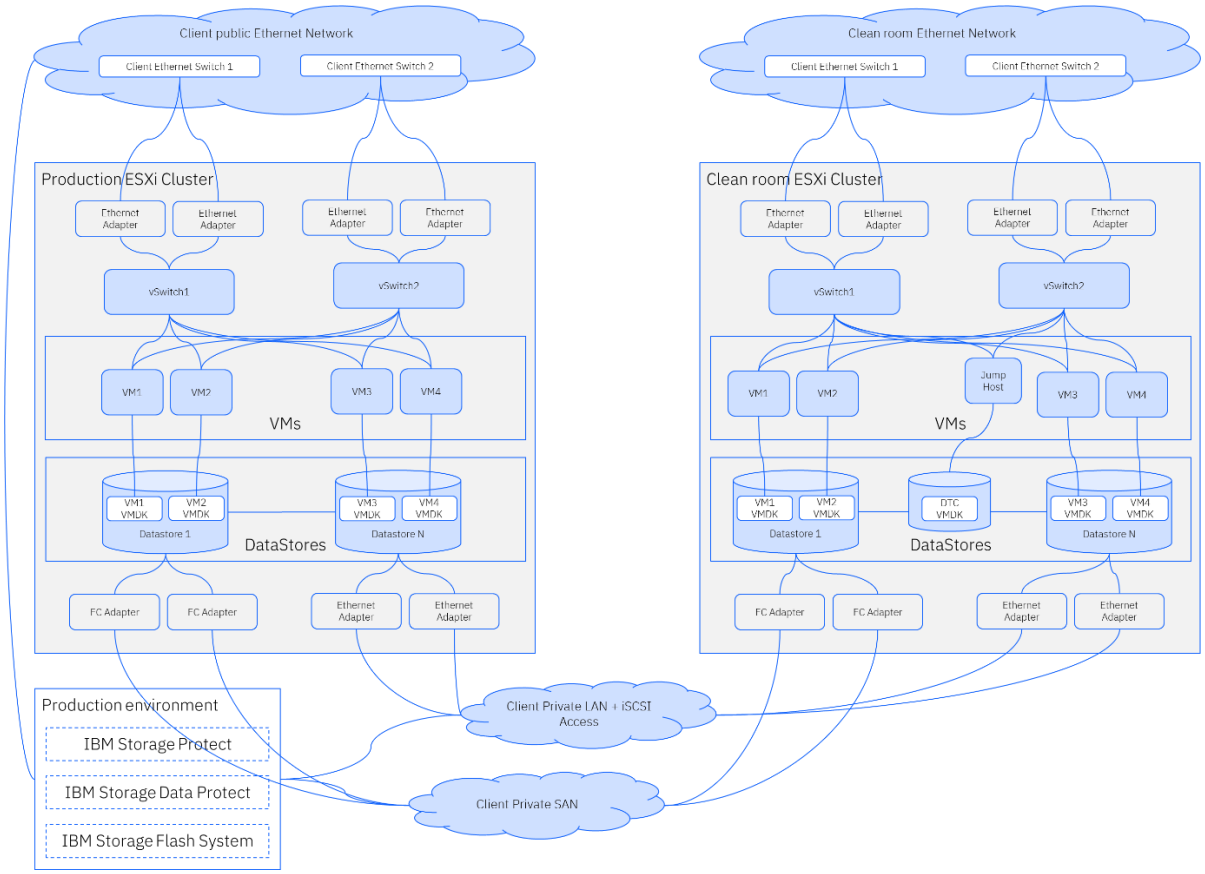

Figure 8: Clean room on dedicated ESXi cluster

7.1 Setup a clean room on a dedicated ESXi cluster

The section describes the setup of clean room components. Table 3 lists the steps to setup a clean room on a dedicated ESXi cluster. These steps include the setup of the physical ESXi host system and the deployment of the vCenter server. Table 3 provides a checklist for the setup operation and VMware documentation links for reference.

Table 3: Setup VMware infrastructure for clean room

Operation	Link to available documentation	Comments
Physical Setup of SAN and Network	Introduction to Storage Introduction to vSphere Networking	For various storage and networking options, see the VMware documentation.
Install and set up vSphere	How to Install and Set Up vSphere	
Install and setup ESXi Hosts	Installing and Setting Up ESXi	
Install vCenter server appliance	Deploying the vCenter Server Appliance	

Creating data centers	How Do You Create a vSphere Data Center	
Adding hosts to the data centers	How to Add a Host to Your vSphere Data Center or Folder	

The operations in the checklist are common tasks. The vSphere administrator in accordance with storage admin and network admin complete these tasks for providing the physical setup of the SAN and network.

8. Network configuration

You can set up and configure networking in a vSphere environment with VMware vSphere virtual switches. This section introduces the terminology and describes the detailed steps to setup the networking for a clean room.

8.1 Terminology

8.1.1 Standard switch versus distributed switch

A network switch in vSphere contains a management plane and a data plane. The data plane implements packet tagging, filtering, and switching. The management plane implements the control structure that is used to configure the data plane.

A standard switch includes the data plane and the management plane. You can configure each standard switch individually. In a distributed switch, the data plane is separated from the management plane. The data plane is implemented in the switch that is local to the ESXi host system, and the management plane resides in the central vCenter server system. The data plane section of the distributed switch is also called a host proxy switch. All the host proxy switches automatically inherit the configuration from the central management plan.

8.1.2 Distributed port groups and configuration inheriting

An important factor in the clean room isolation is the separation of networks for each recovery group or on clean room level. To apply network isolation, you can use the distributed port groups. You can configure NIC teaming, failover, load balancing, vLAN, security, traffic shaping, and other policies on distributed port groups. Each distributed port group has a unique network label which is used to identify the port group. The virtual ports that are connected to a distributed port group share the same properties that are configured in the distributed port group. The configuration that is set up for a distributed port group in the central management plane is automatically inherited by all the host proxy switches, and all the virtual machines that are associated with the same port group.

8.1.3 Uplink ports

You can use uplink ports to map physical NICs of hosts to uplink to the distributed switch. Each physical NIC on an ESXi host system is connected to an uplink port. In addition to the mapping of the physical host NICs, the uplink ports are templates for failover policies and load balancing policies. Similar to the distributed ports, the configuration of the uplink ports is set up at the central management plane of the distributed switch, and it is propagated automatically to the host proxy systems.

8.2 Network bandwidth considerations

The clean room environment has two major use cases. The primary use case is the test recovery of virtual machines. The secondary use case is the temporary production running in the clean room environment.

For a test recovery, the volume that contains the virtual machines to be recovered is mounted as a datastore to the clean room environment. IBM Storage Defender data protect performs a storage vMotion when it restores the VM. While recovering the virtual machines, the VM data that is updated for the startup or any operations on the virtual machine is stored in a file called the delta file in the target datastore used for the recovery. This implies that all the operations like scanning a file system is compute bound on the ESXi host system, and IO bound on the target datastore. On the contrary to this approach, the data resides on the mounted datastore when you recover from IBM Storage Protect Flash System. Depending on the VM size, different amounts of data must be transported over the network. Another important factor is the time you plan to finish the operations and the number of recoveries you plan to process in parallel.

You must prepare the network bandwidth for the largest recovery group in your environment or the largest expected concurrent operations process, wherein large refers to both factors size of a single virtual machine and number of virtual machines to be processed in parallel.

For example:

You are processing two recovery groups in parallel on the same clean room ESXi cluster where recovery group 1a contains 10 VMs of 256 GB each, and recovery group 1b contains two VMs of 1 TB each. The goal is to finish the recovery operation in less than 1 hour.

To accomplish this, you must transport 10 * 256 GB + 2 * 1 TB = 4608 GB per hour. Breaking down the GB per hour to GB per seconds, the network must be prepared to transport 1,28GB/s which will require a 10 GbitE (or comparable performance using FCoE, FC connection) connection to this clean room ESXi cluster referred in the Blueprints section.

8.3 Setup a distributed switch

The first image shows the logical configuration of a standard switch in a vCenter environment. The second image shows a comparable setup using a distributed switch.

Note: You must use separated networks for production data and VM management. If this recommendation is reflected in your production setup, you can create the equivalent setup in the clean room. The second switch for the management network segments is shown in the image.

The number of available physical NICs and the related network bandwidth limit the overall throughput to the isolated clean room environment. You must use the standard switch or distributed switch setup depending upon the expected recovery workload. The approach may vary when multiple clean room environments are setup.

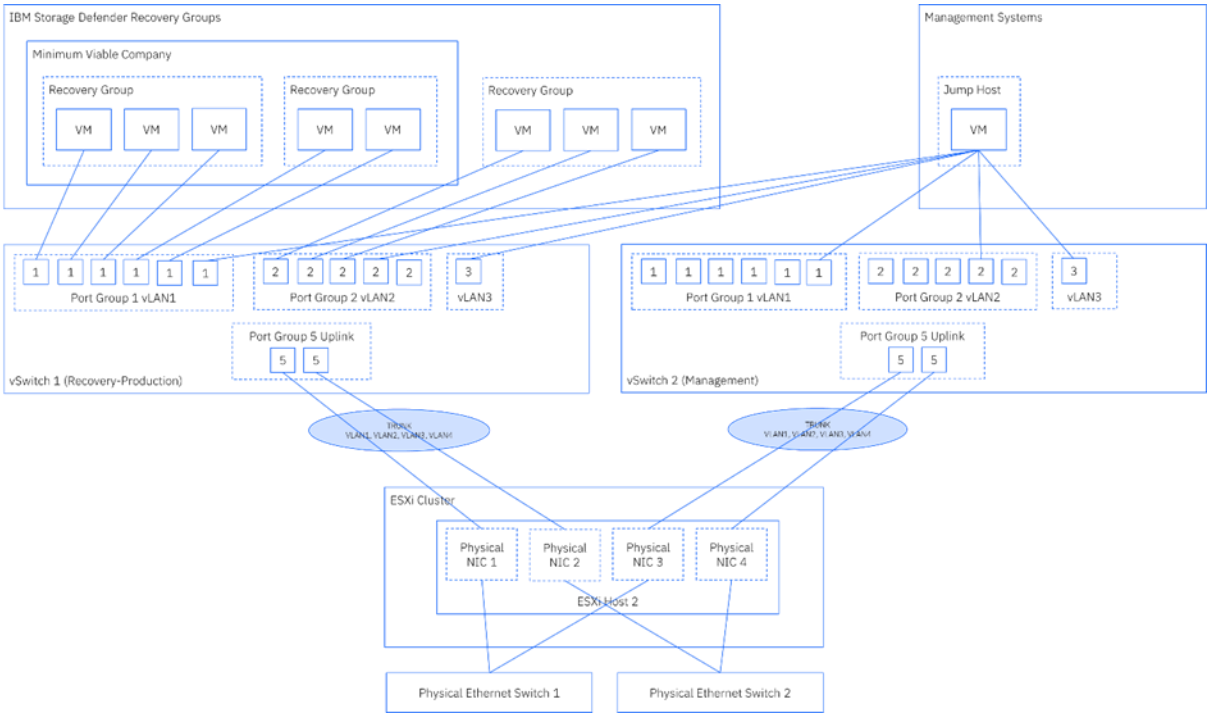

Figure 9: Cleanroom network using standard vSwitch

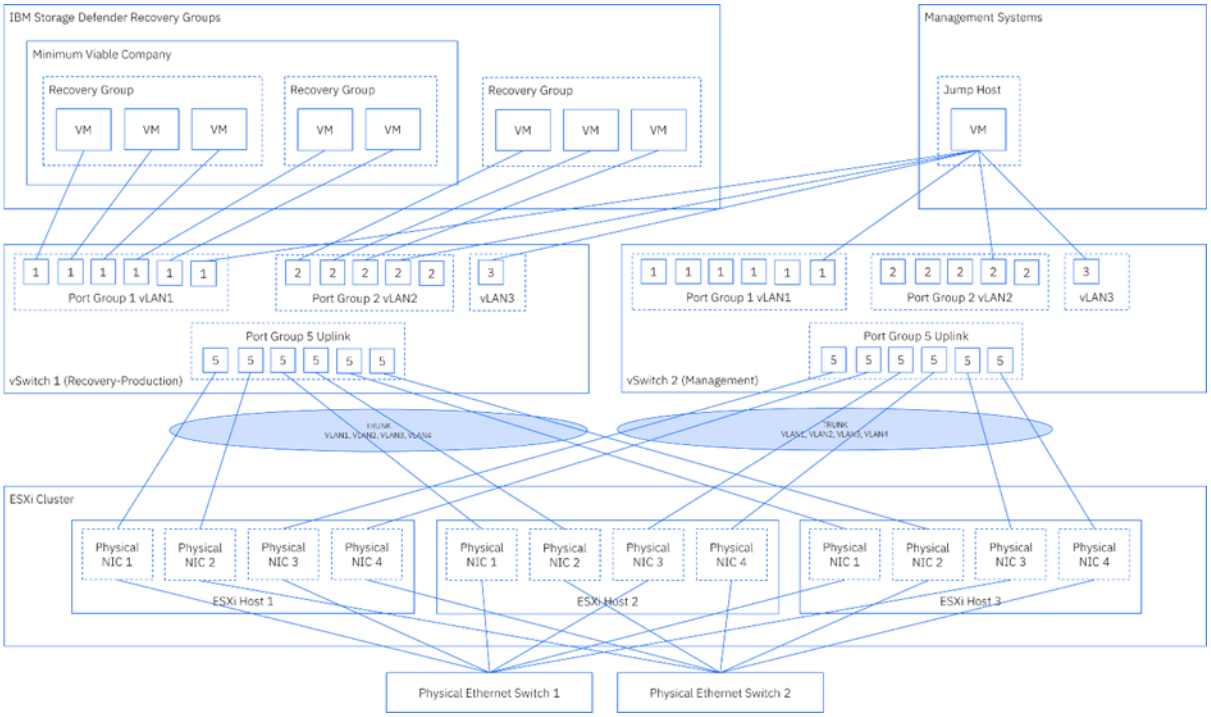

Figure 10: Clean room network using distributed vSwitch

When you setup up a vSphere distributed switch, you must also set up the physical ethernet switches accordingly. It is important that the separation of single networks in the clean room are also considered when you configure the physical ethernet switches together with the cabling of the physical network interfaces of the ESXi hosts.

8.4 Network configuration checklist

The following table represents a checklist for the setup operation. Refer to the following links for VMware documentation.

Table 4: Network configuration

Operation	Link to available documentation	Comments
Create a vSphere distributed Switch for recovery production networks and the management network	Create a vSphere Distributed Switch	
Add Hosts to vSphere distributed Switches and define uplink adapters.	Add Hosts to a vSphere Distributed Switch	Make sure to assign the VMKernel adapters to the port group in the management vSwitch.
Create distributed port groups for every dedicated network by making use of vLANs.	Add a Distributed Port Group	
Create distributed port group in the distributed vSwitch used for management with its own vLAN for the Jump Host (DMZ)	Add a Distributed Port Group	Note: The jump host must implement dedicated network adapter to get connected and access the vLANs of the recovery groups.

The figures below show a screenshot from a vCenter that was used for the test of this documentation. The figures show the topology after the setup for isolated recovery groups, uplink ports and VMKernel ports in the test clean room.

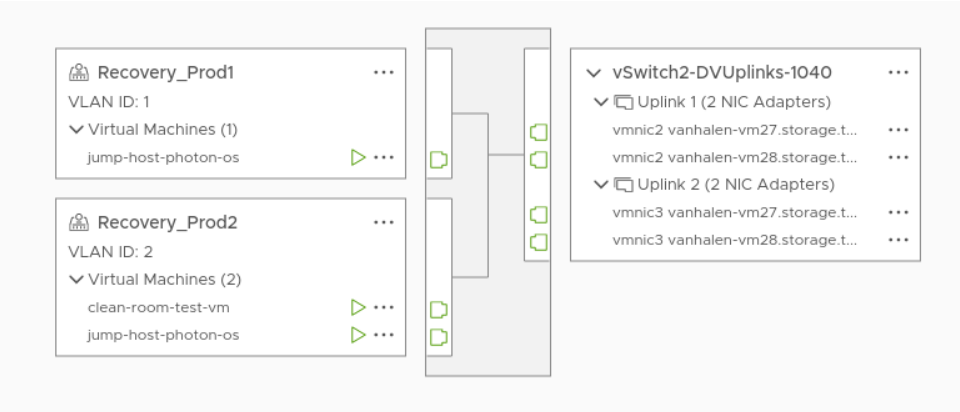

Figure 11: Example Topology for a distributed vSwitch with 2 vLANs for Recovery and the Uplink portgroup

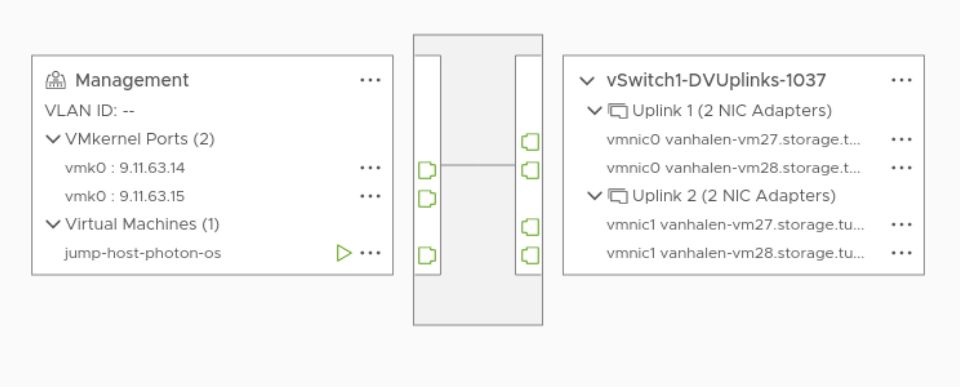

Figure 12: Example Topology for a distributed vSwitch with a Management Network with VMKernel Ports configured

9. Storage configuration

This section describes the configuration of the datastore required for jump hosts (bastion hosts) that must be inside the clean room network. This section also describes the configuration of the SAN connection required for the clean room to allow recovery from the hardware snapshots. The creation of the vSphere cluster is also described that consolidates all the created resources in a common logical construct.

9.1 Storage IO and ESXi compute considerations

As described in section Network bandwidth considerations, the operations performed on the recovered virtual machines are IO bound on the target datastore for IBM Storage Defender data protect recoveries, and IO bound on the mounted datastore when recovering from IBM Storage Flash System. In addition, the compute is bound on the ESXi host system where the operation is processed. This implies that not only the network but the storage subsystem and the compute on the ESXi host must also be able to cover the workload. Referring to the following example used in the network bandwidth section:

You plan to process two recovery groups in parallel to the same clean room ESXi cluster where recovery group 1a contains 10 VMs of 256 GB each, and recovery group 1b contains 2 VMs of 1 TB each. You want to scan the virtual machines using the virus scanner that you have installed on the endpoint. The goal is to finish the scan operation in less than 1 hour.

Comparable to the network bandwidth, the storage sub system must be able to provide the IO performance required. In this example it would be 1.28GB/s read. Since the operation is a scan operation, the write performance is secondary. What must be considered is that the same storage sub system might be used for other recovery operations at the same time and potentially production and backup operations as well. The additional workload must be added to the consideration.

The compute power (RAM and CPU) required on the ESXi host can be calculated as well when the compute characteristic of the scanner used is known. Typically, the vendor of an anti-virus scanner publishes this information in the hardware and software specification of the product. For the calculation, it must be considered that the same ESXi host system might be used for multiple scan operations at the same time. In the above example, this means that the time it takes to scan the 256GB virtual machines 12 machines might be scanned in parallel.

This implies that the hardware specification documented for the scanner must be multiplied with 12. It must also be considered that other vCenter and VM operations might run on the same host concurrently. In addition to the overall number of CPU cores available for scanning, the CPU type plays a role as well. Specifically, for CPU intensive workloads, a faster low-core CPU must be preferred compared to a slower high-core CPU.

9.2 Jump host considerations

The jump host must be located on a separate datastore which is persistently available in the clean room environment and not local to one of the ESXi hosts. This datastore could be backed by any type of physical storage available in vSphere environments like local storage, fiber channel, iSCSI, or NFS. You can use the same shared storage system for the datastore that backs the vCenter appliance.

You can use the operating system of your choice for the setup of a jump host. In many cases, Linux distributions are used. Since the system requirements for a jump host are insignificant, the minimal installation option that several Linux distributions contain can be used. As a jump host acts as an access point to the isolated clean room network, it is essential that the firewall configuration is very restrictive and allows access only to the necessary services.

Note:

- You must keep your operating system up to date.
- You must create and use low privileged user accounts for login that have no sudo permissions. You must administer the system login with a low privileged user account and escalate to a sudo user locally.
- For frequent SSH login, you must use SSH keys instead of user and password combinations. You must use different SSH keys for different system that access your jump host. You must use a keystore to manage the keys that require MFA.
- You must block SSH access to any account except the low privileged account that you created.
- You must configure a firewall so that all unnecessary ports are closed.
- You must consider implementing a tool that analyzes application logs and can create firewall rules dynamically to prevent brute-force attack by blocking misbehaving IP’s.
- You must consider agent forwarding or port forwarding when you login to the jump host and then login to the target host is too time consuming.
- On your target host, you must configure the firewall to allow access from the jump host IP only.
- You must ensure that connections from your clean room hosted applications cannot connect back to the jump host. Connections must only initiate from the jump host into the isolated clean room.

9.3 Connection to SAN

The clean room ESXi hosts must be connected to the Production SAN by Fibre-Channel Adapters or iSCSI. SAN Zoning and iSCSI configuration must allow to map SAN volumes, from the Storage Subsystems used in the Production Environment to the clean room ESXi hosts. The following diagram shows a Production ESXi Cluster with SAN connections to a clean room ESXi cluster. The clean room has only one persistent datastore for hosting the jump host and the VMs recovered using storage vMotion and several datastores that could be mounted from the Storage subsystems.

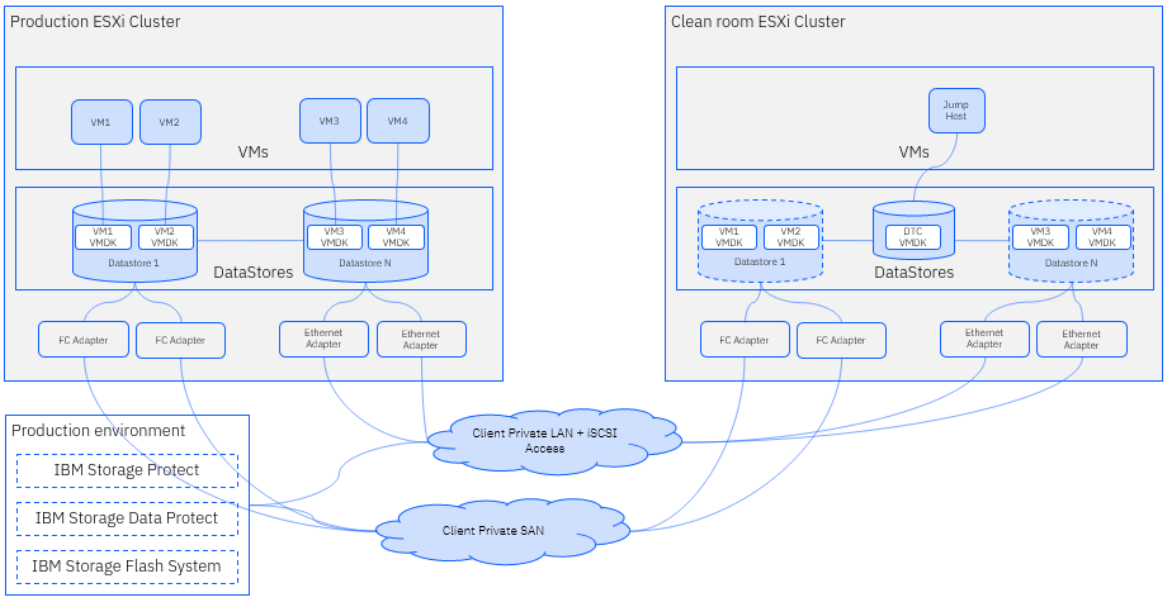

Figure 13: Storage Area Network connections to clean room environment

9.4 Storage Configuration and jump host deployment

You must refer to the following checklist for preparing the storage connections, creating a datastore, and deploying the jump host for a clean room environment.

Table 5: Storage and Cluster configuration

Operation	Link to available documentation	Comments
Configuring storage area network connection to the ESXi host.	Using ESXi with Fibre Channel SAN Using ESXi with iSCSI SAN	
Creating clusters to consolidate the resources of multiple hosts and virtual machines.	How Do You Create and Configure Clusters in the vSphere Client	
Create datastore for jump host (bastion host) system.	Creating vSphere Datastores	
Deploy jump host (bastion host) on the datastore.	Deploying Virtual Machines Deploy and Export OVF and OVA Templates	

10. Introducing IBM Expert Labs and IBM Client Engineering

Implementing one or more clean rooms can be challenging when it comes to sizing, provisioning, or deployment of the required components. IBM provides experienced teams

that help you to get support for the clean room setup and maintenance. These teams are part of the IBM Expert Labs and IBM Client Engineering.

IBM Expert Labs

IBM Expert Labs is a professional services organization powered by an experienced team of product experts. This knowledgeable team brings deep technical expertise across software and infrastructure, including IBM Data and AI, Automation, Sustainability, Security, Software Defined Networking, IBM Power®, IBM Storage, IBM Z® and LinuxONE, IBM GDPS® and IBM Cloud®.

IBM Client Engineering

IBM Client Engineering delivers meaningful and scalable business outcomes across all industries. With our deeply skilled multi-disciplinary squad and human-centered approach, we provide value-based experiences and solutions catered to your organization’s needs. Whether a custom demo in your environment or an MVP to prove value, we meet you where you are and work with your organization at any stage of its digital transformation journey. Client Engineering is an investment in you to co-create and innovate leveraging IBM technology and methodologies.

Table of figures

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing

IBM Corporation

North Castle Drive, MD-NC119

Armonk, NY 10504-1785

US

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing

Legal and Intellectual Property Law

IBM Japan Ltd.

19-21, Nihonbashi-Hakozakicho, Chuo-ku

Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing

IBM Corporation

North Castle Drive, MD-NC119

Armonk, NY 10504-1785

US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBMproducts. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.