



InfiniBox 3.0

Best practices for setting up services

About INFINIDAT

INFINIDAT features high-end, hyperscale storage with highly-dense capacity, high performance, and infinite use cases. INFINIDAT offers you an environment-friendly system with low power consumption, low CapEx & OpEx and unmatched TCO.

With more than 90 patent applications (many of which are approved), INFINIDAT is a global leader in the areas of data protection, performance, automation, energy savings, ease of management, and scaling.

About this manual

This publication provides best practices for setting up Infinibox services (iSCSI, NAS, replication).

Date	Product version	Content
Dec-2016	3.0	This guide consolidates exiting guides for Replication and NAS with newly introduced content for setting up an iSCSI service.

Contents

Introduction.....	1
<i>INFINIBOX ETHERNET PORTS.....</i>	<i>1</i>
<i>HOW THIS DOCUMENT IS ORGANIZED</i>	<i>1</i>
Creating a network space.....	2
<i>CREATING A NETWORK SPACE FOR ISCSI.....</i>	<i>2</i>
<i>CREATING A NETWORK SPACE FOR NAS.....</i>	<i>5</i>
<i>CREATING A NETWORK SPACE FOR REPLICATION.....</i>	<i>8</i>
Creating a service	11
<i>CREATING AN ISCSI SERVICE.....</i>	<i>11</i>
<i>CREATING A NAS SERVICE</i>	<i>15</i>
<i>CREATING A REPLICATION SERVICE</i>	<i>17</i>
Deploying multiple services together	20

Introduction

InfiniBox deploys IP-based data services (iSCSI, NAS, replication) on top of Network Spaces. A network space is a logical entity that makes use of the InfiniBox networking hardware.

InfiniBox Ethernet ports

InfiniBox has either 2 or 4 Ethernet ports per node (6 to 12 ports per InfiniBox). The network layout – the way InfiniBox connects to external resources – depends on the deployed service and on the number of external switches. A detailed discussion on the available network layouts can be found on the Support site (Deployment considerations for services).

How this document is organized

This document describes the best practices for setting up services in the following order:

- [Creating a network space](#) – what is required for creating a network space for each of the services
- [Creating a service](#) – how to create which of the services
- [Considerations for coexistence of services](#) – how to deploy several services on the same InfiniBox

Creating a network space

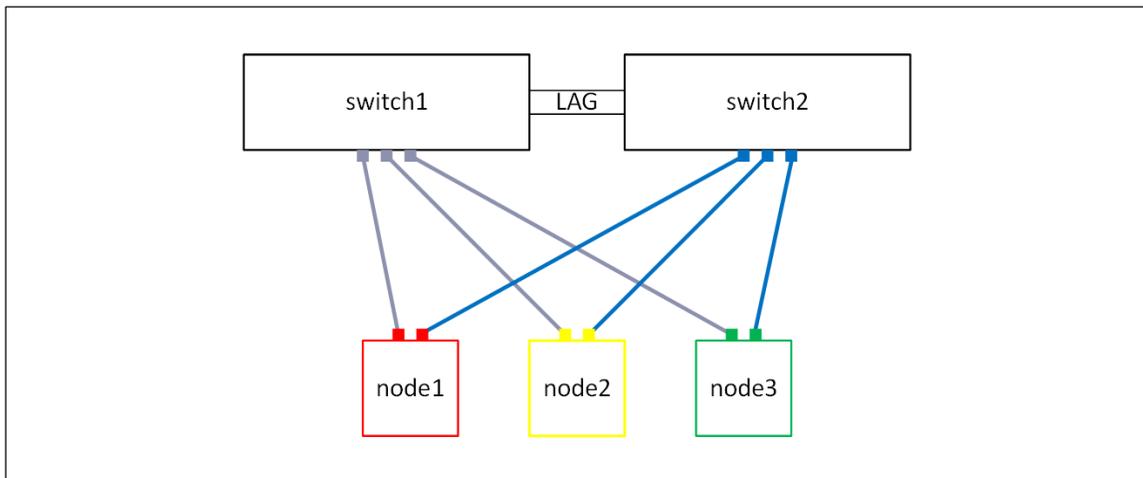
Creating a network space for iSCSI

It is required to connect the iSCSI service to switches. The switches can be linked through a Link Aggregation Group, but this is not a recommendation.

The iSCSI service runs independently from replication and NAS.

- Each network space includes one port on each node.
- Each network space is connected to one switch.
- Each InfiniBox port is part of a network space, and each network space is connected to each switch.

Note: it is not recommended to use server-side Link Aggregation for iSCSI. The high-availability is provided by the multi-path.



PREREQUISITES FOR USING VLANS

Make sure that your network is configured to allow access to the specific VLAN (the network administrator allows the ports that are connected to InfiniBox to send VLAN-tagged data to the relevant VLANs).

CREATING AN ISCSI NETWORK SPACE

- Step 1
- Click **Settings** on the InfiniBox menu
 - Click the **Network Space** tab
 - Click **Create**
The **Create Network Space** screen opens

The screenshot shows the 'Create Network Space' configuration window. The window has a title bar with a gear icon, the text 'Create Network Space', and a close button. On the left side, there is a sidebar with 'Ethernet Interfaces' and 'IP Configuration' sections. The main area contains the following fields and controls:

- Network Name:** A text input field.
- Service:** A dropdown menu with 'iSCSI' selected.
- Rate Limit Per Node:** A dropdown menu with 'Optional (Mbps)' selected.
- MTU:** A text input field with '1500' entered.
- Ethernet Interfaces:** A section containing three dropdown menus for 'Node 1', 'Node 2', and 'Node 3', each with 'PG1' selected.
- Buttons:** 'CREATE INTERFACES' and 'CREATE VLAN' buttons are located below the interface selection fields.
- Footer:** 'CANCEL' and 'NEXT' buttons are located at the bottom right of the window.

- Step 2
- Fill in the network space name and MTU as required.
 - Select **iSCSI** from the **Service** drop-down list

Step 3 Either select an Ethernet interface or click **Create Interfaces**.

1. You can rename the default Interface name
2. Select from the available Ethernet ports (port that are already taken by other interfaces are greyed-out)
3. Click **Create**
4. You are returned to the **Create Network Space** screen.
Continue to create interfaces as needed (for example, when creating an interface per InfiniBox node)

The screenshot shows the 'Create Network Space' dialog box. The title bar includes a gear icon, the text 'Create Network Space', and a close button. The left sidebar has 'Ethernet Interfaces' and 'IP Configuration'. The main content area is titled 'Create Ethernet Interfaces' and contains the following fields and options:

- Interface Name**: A text input field containing 'PG2'.
- Node 1**: A text input field containing 'N1PG2'.
- Node 2**: A text input field containing 'N2PG2'.
- Node 3**: A text input field containing 'N3PG2'.
- Ports**: A list of checkboxes for 'ETH1', 'ETH2', 'ETH3', and 'ETH4'. 'ETH1' is checked with a green checkmark.
- Buttons**: 'CANCEL' and 'CREATE' buttons are located at the bottom left. 'CANCEL' and 'NEXT' buttons are located at the bottom right.

- **Note:** the Ethernet interfaces can be created beforehand, from the **Network Interfaces** tab.

Step 4 (Optional)

Click **Create VLAN** in order to group interfaces into a Virtual LAN.

Step 5 Click Next to move to the **IP Configuration** screen.

Step 6 Fill in the networking data:

- Network
- Netmask
- IP Addresses – click **Add** verify the validity of the IP Addresses

Click **Finish**.

The Network Space is created and is visible on screen.

Creating a network space for NAS

NETWORK CONFIGURATION

Every InfiniBox system includes 3 controller nodes. Every node has access to all drives. Depending on the available InfiniBox model, one may have either 2 (F61xx/F21xx) or 4 (F62xx/F22xx) 10Gbps Ethernet ports per node.



Fig.1: 4 port Copper node

Fig.2: 4 port Optical node

All new InfiniBox models (F62xx/F22xx) are shipped with 4 ports per node (total 12 ports for an InfiniBox system). This article covers 4-port configurations only.

RECOMMENDED PHYSICAL NETWORK LAYOUT OPTIONS

Different InfiniBox services require a network space of their own. For example, a separate network space is required for the replication service (available as of InfiniBox v2.0). A separate network space will be required for the iSCSI service (available in an upcoming release). IPs assigned to different network spaces can't overlap. Plan your network layout accordingly.

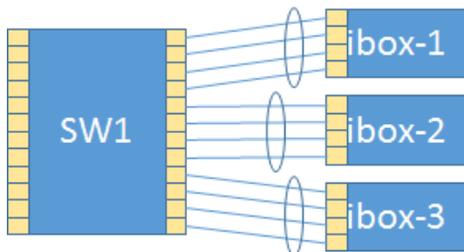
Best planning for 2 services (Replication & NAS) would consider isolation of services with full redundancy which includes 6 LACP (Link Aggregation Control Protocol) Bonds (two per node), and 6 IP addresses per service.

Alternatively, sharing interfaces for 2 services will provide bandwidth optimization.

InfiniBox allows using either single interface or LACP to increase available bandwidth and redundancy. Multiple network layouts are supported; some of the most common are shown below.

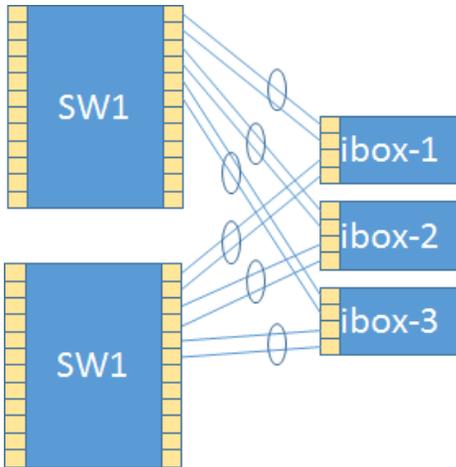
SINGLE NETWORK, SINGLE SWITCH

All InfiniBox ports are connected to a single switch. Single network space and single set of NAS IPs is created with 3 LACPs (switch vulnerable):



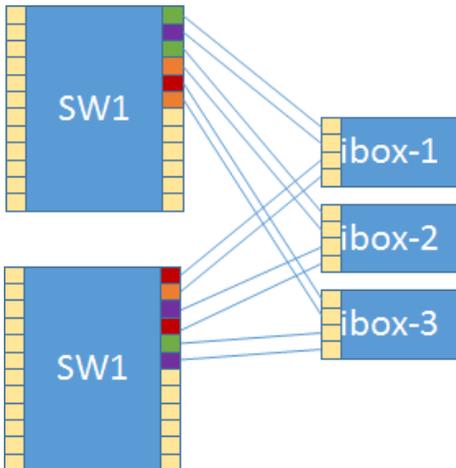
(BEST) - SINGLE NETWORK, TWO SWITCHES

2 network switches are used for increased redundancy. Single network space and single set of NAS IPs is created with 6 Active LACP Bonds:



MULTIPLE NETWORKS, TWO SWITCHES (VLAN WORKAROUND)

2 network switches, up to 4 separate networks. Note - in this scenario, bandwidth for each network will be limited to a single 10GbE connection:



In case of a switch failure InfiniBox 2.2.x will detect a port failure (link down) and will initiate failover. Future code revisions may handle mis-configurations such as:

- Incorrect MTU set on switch
- Incorrect LACP set on switch
- One port was accidentally connected an incorrect VLAN
- Ports from nodes 1&2 connected to switch 1, ports from node 3 connected 3 and the ISL between the switches fails, requiring the IPs to failover.

INSTRUCTIONS FOR CREATING A NETWORK SPACE

A proper network configuration for NAS requires a Network Space (see Fig.3). The network space definition includes:

- Service type (NAS)
- Network definition (Network addresses, netmask, default gateway)
- 3 interfaces (1 per node)
- MTU size

It is recommended to assign 6 separate IPs to the network space service definition (2 per node).

If a node were to fail, its two IPs would be automatically redistributed between the two remaining nodes, resulting in a balanced configuration with three IPs per node.

The following image displays the network space definition on the InfiniBox GUI.

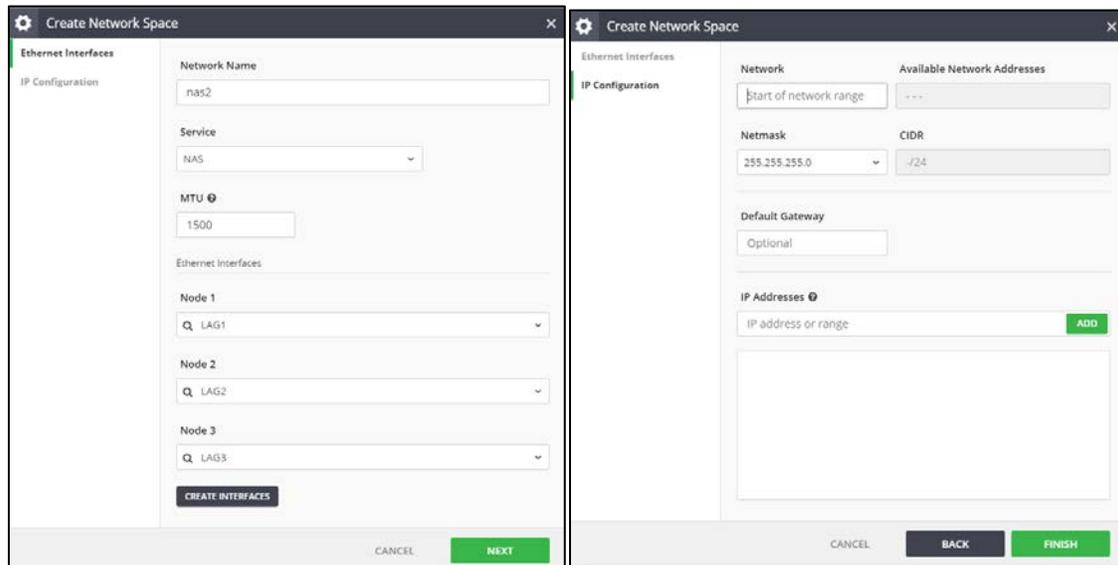


Fig.3 - InfiniBox GUI for Network Space definition

Creating a network space for replication

- Interface - One or more Ethernet Ports that are grouped together for the purpose of redundancy and/or performance using LACP.
- Network Space - A set of definitions for a network address space, annotated using CIDR.
- Link - A bidirectional means of communication between local and remote systems.
- Replica - An entity that matches a pair of source and target volumes and other settings essential for the replication.
- Replica attributes
 - Sync interval - the interval between two consecutive sync jobs
 - RPO (Recovery Time Objective) - the lag between the source and target sites

BEFORE YOU BEGIN

Make sure you have accessibility (admin privileges) to both local and remote systems. On each of the systems:

Functionality	Connected ports	IP addresses	Firewall ports
Management	1	1	80,443
Data	3 to 6 (3 * 1 to 2)	3 to 6	8067

INSTRUCTIONS

- Repeat the following task sequence for each of the sites (local and remote)
- The ports that are available for the replication are listed on the **Network Interfaces** tab

Step 1 On the GUI, click the **Settings** icon on the toolbar on the left, and then click the **Network Spaces** tab.

Click **Create**.

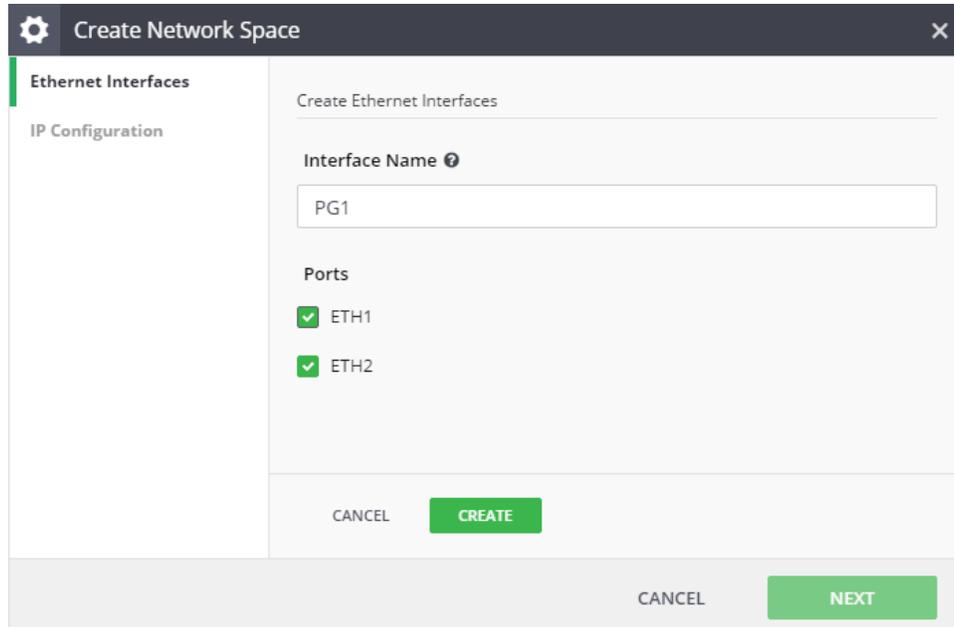
The **Create Network Space** wizard opens. Fill in the replication network space details.

The screenshot shows the 'Create Network Space' wizard. The 'Ethernet Interfaces' tab is active. The 'Network Name' is 'Replication'. The 'Service' is 'Replication'. The 'Rate Limit Per Node' is 'Optional (Mbps)'. The 'MTU' is '1500'. The 'Ethernet Interfaces' section shows 'Node 1' set to 'n1ethdata1'. There are 'CREATE INTERFACES' and 'CREATE VLAN' buttons. At the bottom, there are 'CANCEL' and 'NEXT' buttons.

Step 2 Click **Create Interfaces**.

The **Create Ethernet Interfaces** screen opens. Select ports for each of the interfaces and click **Create**.

You are returned to the **Ethernet Interfaces** screen.



The screenshot shows a configuration window titled "Create Network Space" with a sub-tab "Ethernet Interfaces". The main area is titled "Create Ethernet Interfaces". It features a text input field for "Interface Name" containing "PG1". Below this is a "Ports" section with two checked checkboxes: "ETH1" and "ETH2". At the bottom of the main area are "CANCEL" and "CREATE" buttons. At the bottom of the entire window are "CANCEL" and "NEXT" buttons.

Click **Next** to move to the **IP Configuration** screen.

Step 3 Fill in the following properties:

- Network
- Netmask
- Default Gateway
- IP Addresses

Click **Finish**.

The wizard screen closes and the network space is displayed on the **Network Spaces** tab.

Creating a service

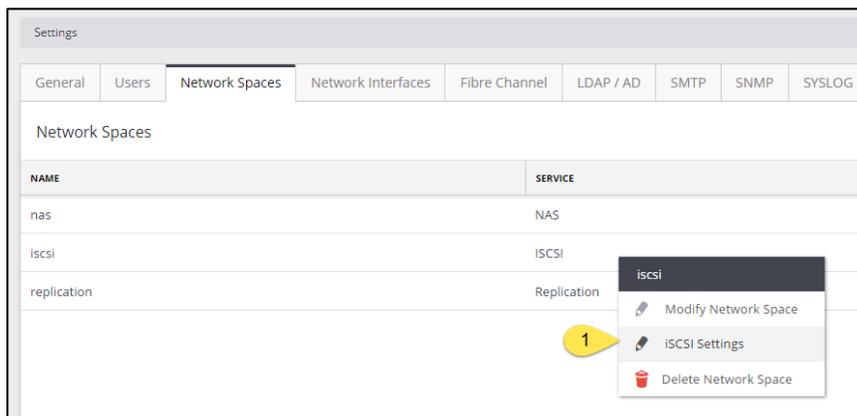
Creating an iSCSI service

ADVANCED ISCSI CONFIGURATION

Specify your port (if not using defaults) as well as an iSNS Server IP (if you have one). Or simply accept defaults.

An iSNS server provides a central management point for the group by dynamically maintaining up-to-date information about the iSCSI target names for group volumes. Once the IP address of an iSNS server is entered in an iSCSI initiator's configuration utility, the setting is persistent across initiator sessions.

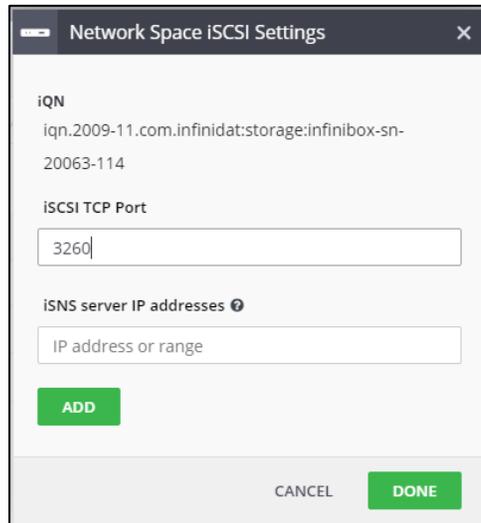
Step 1 Right-click the iSCSI network space and select **iSCSI Settings** from the pop-up menu.



The **iSCSI Settings** screen opens.

Step 2 Fill in the:

- iSCSI TCP Port
- iSNS server IP addresses



Network Space iSCSI Settings

iQN
iqn.2009-11.com.infinidat:storage:infinibox-sn-20063-114

iSCSI TCP Port
3260

iSNS server IP addresses ⓘ
IP address or range

ADD

CANCEL DONE

Click **Done**.

ISCSI CLIENT CONFIGURATION

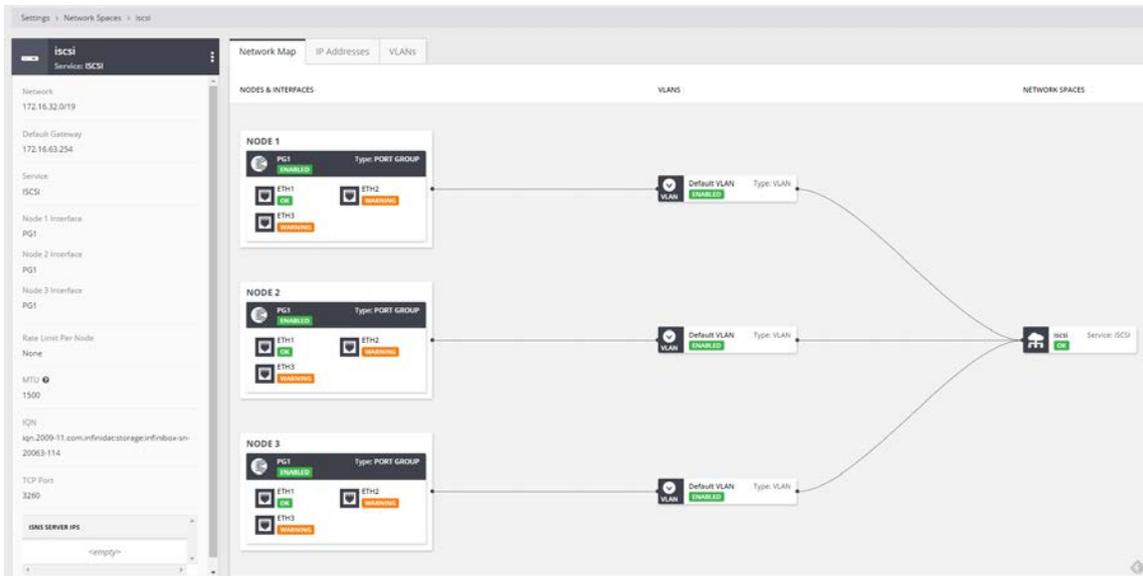
Our Host PowerTools with client registration automatically exposes the IQNs to the Server and provides an extremely simple vehicle for proper client configuration and Server/Client configuration validation. Mapping clients to volumes allows you to simply pick or add the Client iQNs for port connectivity. Advanced performance metrics and Analytics are automatically available for FC, iSCSI, and NAS. Best Practices suggest the client use MPIO for balancing IO across the Server IPs (all three controllers).

iQN = iSCSI Qualified Name - similar to the WorldWide Name (WWN) associated with Fibre Channel devices and is used as a way to universally identify the node.

iqn.yyyy-mm.naming-authority:unique

VIEWING THE ISCSI NETWORK SPACE MAP

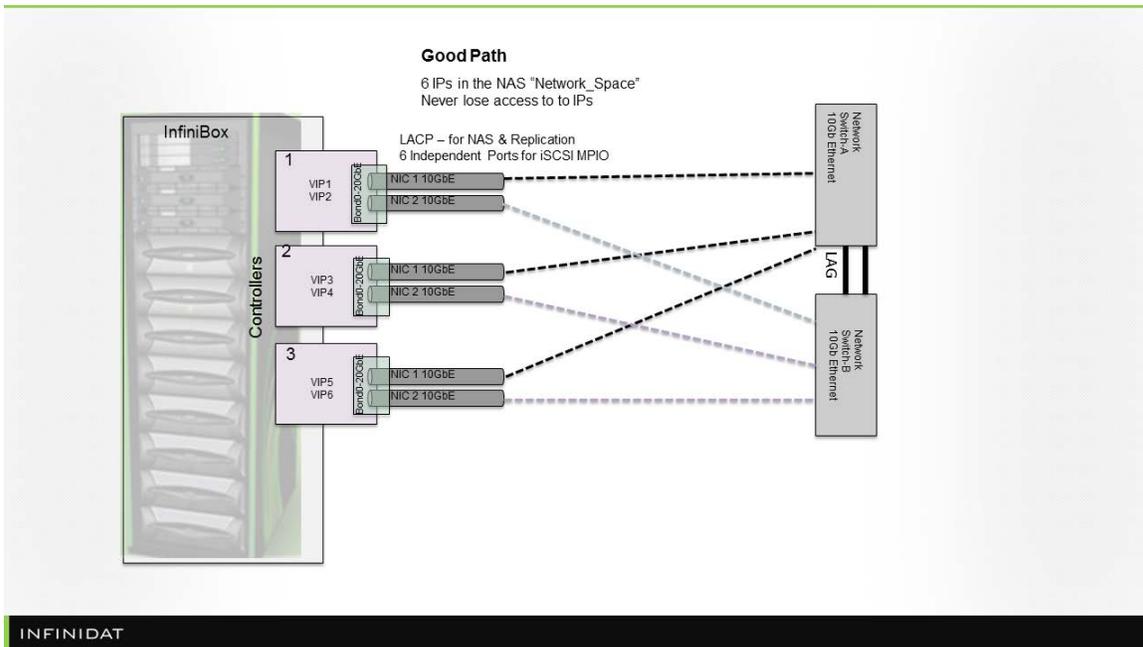
Click the iSCSI service to open its screen. A visual map of the network space and the way it connects the INfiniBox nodes to the switch is displayed on screen.



ASSURING REDUNDANCY

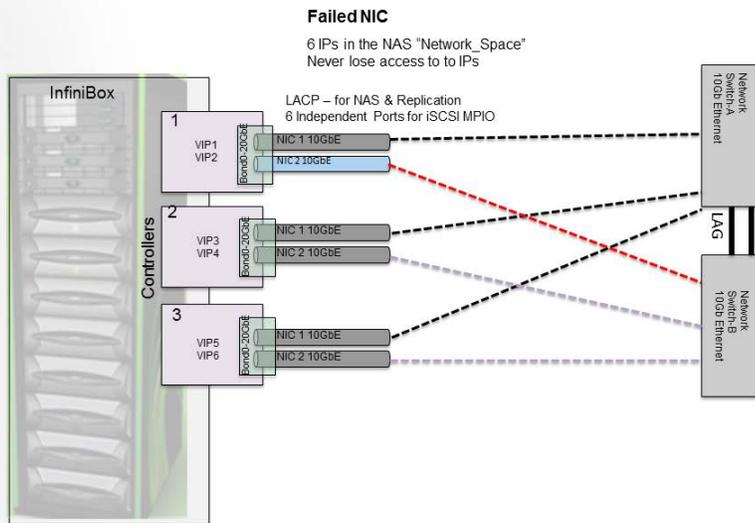
Good Path ensures full 20GbE aggregate performance/node and redundancy to all Virtual IPS (NAS/Replication)– always.

For iSCSI 6 IPs are also virtually Applied - however each connection is independently managed with MPIO on the client.



Failed NIC

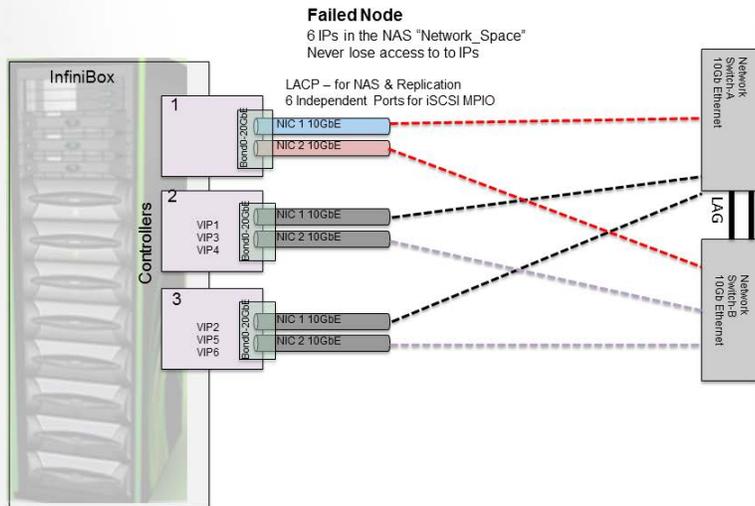
1 Failed NIC maintains access to all Virtual IPs albeit 1 Controller has a Limited Max BW equal to the aggregation of remaining ports in the surviving port.



INFINIDAT

Failed Node

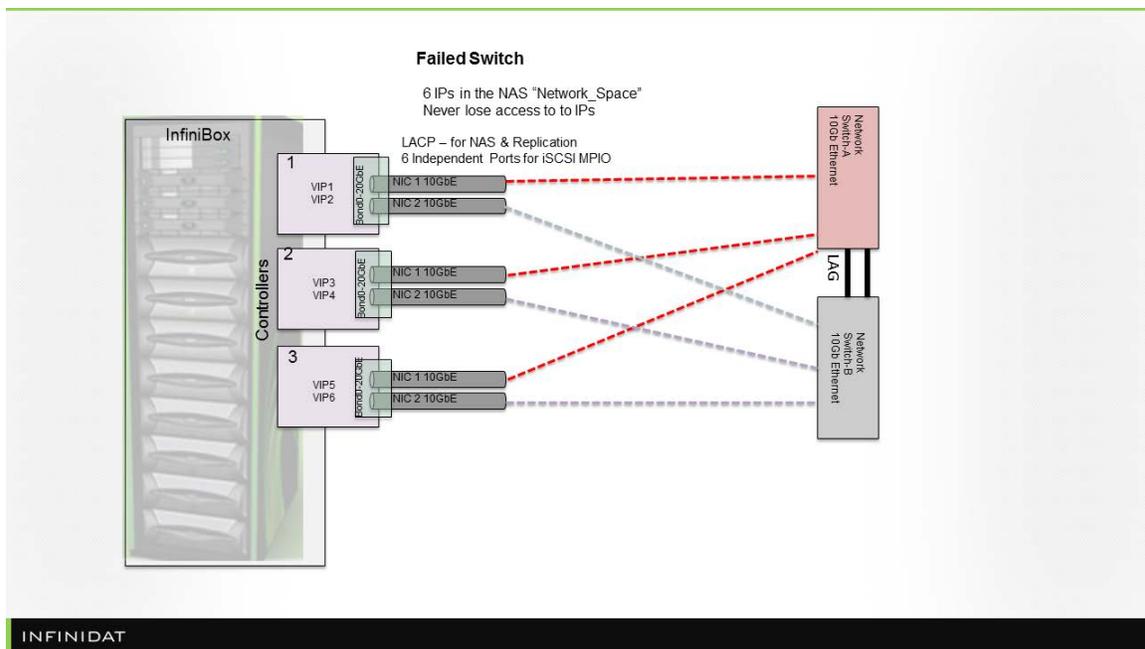
Node Failure maintains access to all IPs although it causes Virtual IP failover with redistribution of IPs across the remaining controllers. Bandwidth is limited to the aggregate Max Bandwidth of the remaining Controller ports.



INFINIDAT

Failed Switch

A Failed Switch maintains access to all Virtual IPs albeit max Aggregate performance is reduced to 50% of overall Port Max Bandwidth on each controller. Note, most clients do not operate at the ceiling level therefor impact is commonly a low impact to client performance.



Creating a NAS service

LOAD BALANCING

DNS Round Robin is the best and simplest way to enable a balanced distribution of mounts between 6 IPs within the network space. Simply define multiple A records with the same name and different IPs in your DNS zone file:

```
; zone file fragment
ibox IN A 192.168.0.4
ibox IN A 192.168.0.5
ibox IN A 192.168.0.6
ibox IN A 192.168.0.7
ibox IN A 192.168.0.8
ibox IN A 192.168.0.9
```

With this setup, the name "ibox" has been assigned as the hostname. Each sequential name lookup will result in the next IP being used, so client connections should be distributed across the available ports and IPs.

Automounter can be used to balance mounts between multiple IPs:

```
fs01 -rw 192.168.0.4,192.168.0.5,192.168.0.6,192.168.0.7:/fs01
```

Manual IP balancing could also be used with small amount of NFS clients

NFS SERVER CONFIGURATION OPTIONS

NFS performance highly depends on the actual application workload. As a rule of thumb, if you are dealing with a large file IO it is recommended to decrease the fragmentation of packets. When jumbo frames are used, it is important to preserve consistent settings on the client, network and file server to get the best performance.

1. At the Ethernet frames level - configure MTU=9000. For InfiniBox NAS, it can be done as part of the Network Space creation
2. At the NFS level - configure rsize and wsize to at least 65536 as part of NFS export creation. The same values should be used on the NFS client during mount:

INFINIDAT NAS SYSTEM DEFAULTS

There are several system-level defaults which can be modified by INFINIDAT support:

Setting	Parameter name	Default value
Export permissions for newly created exports	mgmt/file.export_permissions_default	[{"access": "RW", "client": "*", "no_root_squash": true}]
Access permissions for newly created filesystem	mgmt/file.root_mode_default	755

By default, new file systems are created with root:root ownership and permissions 755:

```
$ ls -ld /mnt/fs01
drwxr-xr-x 2 root root 0 Sep 12 10:10 /mnt/fs01
```

The file system ownership and/or permissions should usually be modified on the client with root access to the file system as needed:

```
$ sudo chown john:it /mnt/fs01
$ sudo chmod 2750 /mnt/fs01
$ ls -ld /mnt/fs01
drwxr-s--- 2 john it 0 Sep 12 10:10 /mnt/fs01
```

Creating a Replication service

Remote replication is the process of copying a dataset from a local InfiniBox to a remote InfiniBox for the purpose of data protection or disaster recovery. The replication ensures uninterrupted business operation in case of a site failure. The INFINIDAT Remote Replication technology asynchronously replicates data from a source InfiniBox to a target InfiniBox. The replication is defined per volume.

This task replicates a volume, or a consistency group, from a source InfiniBox to a target InfiniBox.

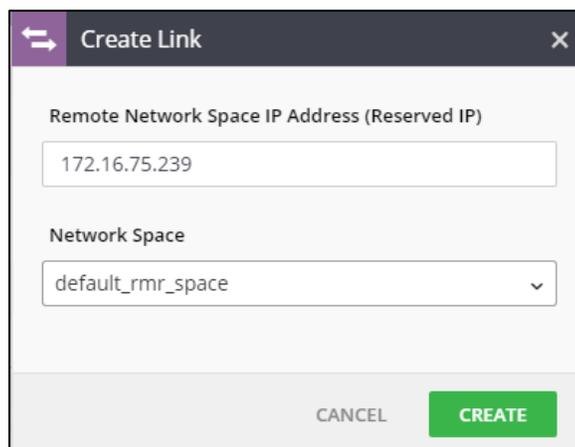
CREATING A LINK

Step 1 Click the Replication icon at the toolbar on the left. The **Replication** screen opens. Click the **Links** tab.

Step 2 Click **Create Link**. Insert an IP address of the remote system and click **Create**. Select the local network space that runs the Replication service.

In case that the remote system credentials are not identical to the credential of the local system, you will be asked to enter them.

The two systems are now connected.



The Links tab displays the link status.



CREATING A REPLICATION

Based on fulfilling the steps above, you can now create a replica on the source system.

Step 1 Click the **Replica** tab. Click **Create Replica**. Fill in the following details:

- Replica type – either a volume or a consistency group
- Source – the name of the volume or consistency group
- Remote system – the system that stores the target of the replica
- Remote target – select either of the following:
 - Create new – let InfiniBox to automatically create the replica entity on the target
 - Select existing – select an entity on the target InfiniBox
- Remote name – filled automatically if Create new was selected
- Remote pool – select the pool that will store the remote entity
- Interval
- RPO

Click **Create**.

The replica is created and is visible on the **Replication** screen.

The replica automatically initializes and its replication progress is visible on screen.

← Create Replica
×

Replica type

Consistency Group ▼

Source

Q cg2 ▼

Remote System

Q box-avi ▼

Remote target

Create new
 Select existing

Remote CG Name

cg2-target

Remote Pool

Q p-box-avi ▼

Interval

60

Sec ▼

RPO

5

Min ▼

📍 LOCAL
🌐 REMOTE
⋮

📊 cg2
Pool: p-box-avia

🗄️ v-7
Size: 15 GB

🗄️ v-6
Size: 15 GB

🗄️ v-9
Size: 15 GB

🗄️ v-8
Size: 15 GB

🗄️ v-10
Size: 15 GB

📊 cg2-target
Pool: p-box-avi

🗄️ v-7-target
Size: 15 GB ✎

🗄️ v-6-target
Size: 15 GB ✎

🗄️ v-9-target
Size: 15 GB ✎

🗄️ v-8-target
Size: 15 GB ✎

🗄️ v-10-target
Size: 15 GB ✎

CANCEL
CREATE

Deploying multiple services together

Deploying multiple services on the same InfiniBox system requires the services to share both InfiniBox and networking resources. On InfiniBox, each service has evenly divided access to each of the three INfiniBox nodes. On the networking end, we separate services according to the number of switches that they require.

- Single switch – the NAS and replication services can connect to a single switch
- Multiple switch – the iSCSI service greatly benefits from connecting to two switches, preferably switches that maintain a Link Aggregation Group between them

For a broad discussion on that various ways services can be deployed, see the [Deployment Consideration for Services](#) article on the Support site.