Accelerated Network Transfer for Migration (On-premises to Power Virtual Server)

Do-it-Yourself solution guide

Version: 2

Prepared by:

Boby Kuruvila George
Senior Cloud architect,
IBM Cloud Infrastructure Solutions
boby.george@ibm.com

Bryan Buckland
STSM, IBM Cloud Infrastructure Solutions
IBM Cloud Infrastructure Solutions
bryan.buckland@ibm.com

Table of Contents

Introduction: 3

Flow chart for Migration: 5

Reference architecture: 6

- Architecture diagram at high level: 6
- Labelled overview of architecture 6

Actualising the architecture: 9

- Prerequisites: 9
- Three step procedure: 10
- Step A: Preparation at source (On-premises). 10
- Step B: Preparation at target (Power Virtual Server) 11
- Step C: Transfer & Access 13
- Transfer: 13
- Access: 17
- NFS server configuration (On the Aspera Server): 17
- NFS access from AIX (target VSI): 18
- NFS access from IBM i (target VSI): 20

Additional considerations & best practices: 22

Conclusion: 24

Introduction:

Transferring data for migration is a vital part of the typical journey to cloud. Data transfer into IBM cloud can be achieved through these methods:

1. Appliance-based transfer.
2. Network-based transfer.
    a. Aspera based transfer – High speed transfer solution in tandem with NFS content sharing.
    b. Traditional TCP-based transfers. For example, SCP, FTP, or SFTP.
    c. Host based block replication or transfer after initial OS build is performed in Power Virtual Server.
        i. IBM i Geo Mirror.
        ii. AIX GLVM.
    d. Backup or VTL replication.

Though there are several ways to transfer data over a network; higher efficiency helps you realise a faster time to value from the cloud and enhances the migration experience overall. It also speeds up the cutover.

Typical pain points: The below are few pain points that is commonly encountered during migration.

- ❖ **Transfer delays:** Transfer delays can affect the build of the virtual server instances. If the transfer takes longer than expected, it can cause issues such as stale data (requiring the backup & transfer to be repeated) and project overruns. It is important to ensure the transfer process is as efficient as possible to avoid these problems.
- ❖ **Continuous replication scenarios:** A delayed data transfer & subsequent build puts more strain on the source server's (on-premises server) storage pool to retain logs from the point of taking the original backup for establishing subsequent logical or native DB replication between source (on-premises) and target (Power Virtual server instance) in preparation for a future cutover. This could potentially cause performance issues if the storage capacity threshold is breached due to log accumulation. Similarly, if client has limited storage capacity, it may require the entire backup & transfer process to be repeated using more effective means.
- ❖ **Network conditions are not ideal:** When you face latency, competing workload, or network congestion, you would want to use the best tools available to migrate in an efficient manner rather than endure a painfully slow process. This is when Aspera with the NFS methodology is beneficial.

This solution guide will equip you with an architectural understanding along with a stepwise procedure for setting up the high-speed transfer solution.

Aspera leverages a proprietary FASP protocol for the transfer and has no theoretical throughput limit. In addition to network capacity, transfer rates can be controlled by user-configured rate settings and the resources of the local and remote machines. You can configure bandwidth usage limits and the number of concurrent FASP transfers that are allowed by the Desktop Client

Use the Aspera online calculator to estimate the transfer speed and time. It can also be used to compare with traditional TCP-based transfers.

Flow chart for Migration:

In preparation for the accelerated transfer, the necessary OS & application specific backup procedure needs to be performed on-premises. This would require generation of either OS or application specific backups which will then be transferred into the Power Virtual Server environment using the accelerated network transfer solution described in this guide. OS build and data restoration would be subsequent activities. The flowchart of the procedure is depicted below:

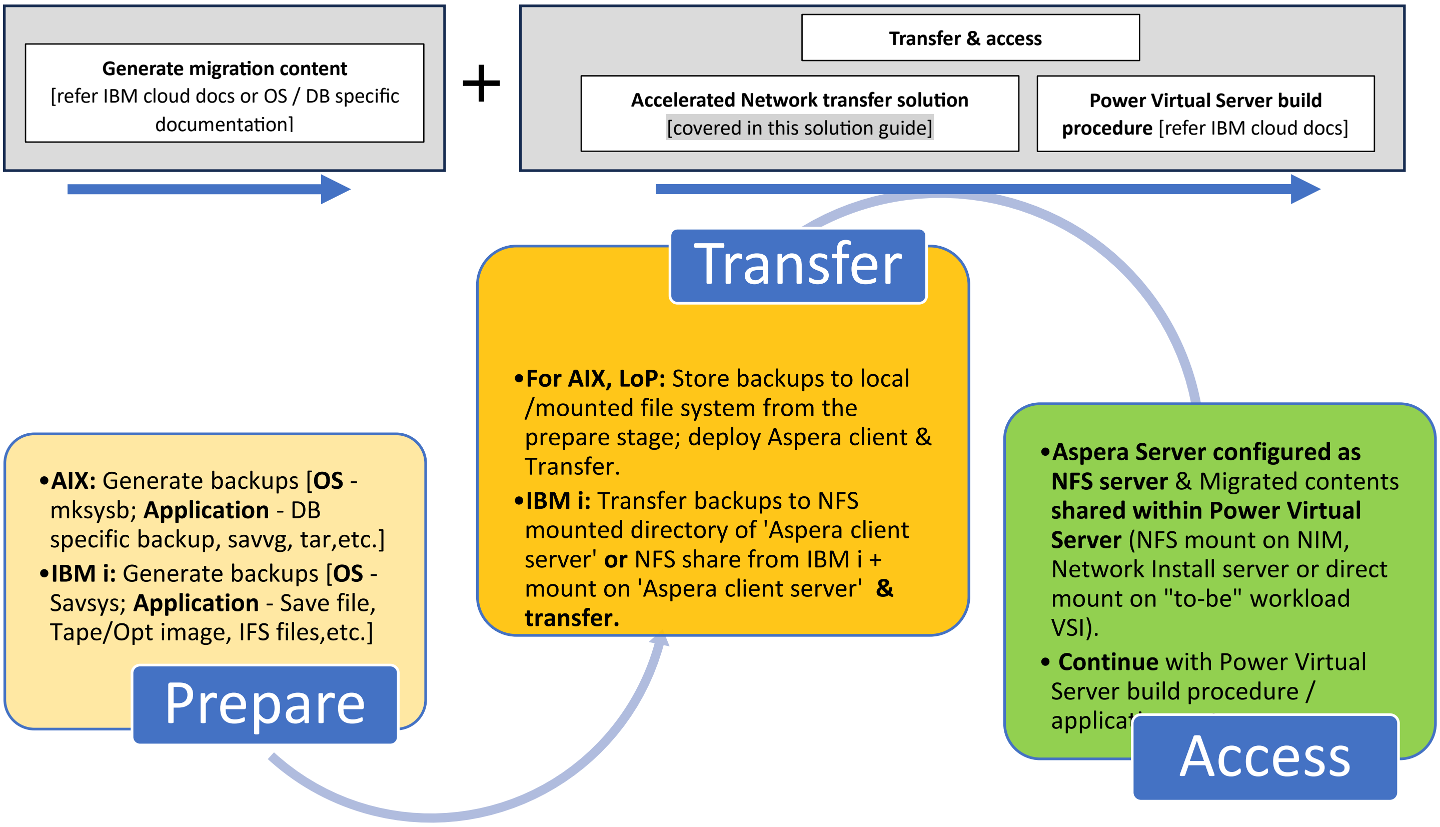

Generate migration content [refer IBM cloud docs or OS / DB specific documentation]

+

Transfer & access

Accelerated Network transfer solution [covered in this solution guide]

Power Virtual Server build procedure [refer IBM cloud docs]

Prepare

- **AIX:** Generate backups [**OS** - mksysb; **Application** - DB specific backup, savvg, tar,etc.]
- **IBM i:** Generate backups [**OS** - Savsys; **Application** - Save file, Tape/Opt image, IFS files,etc.]

Transfer

- **For AIX, LoP:** Store backups to local /mounted file system from the prepare stage; deploy Aspera client & Transfer.
- **IBM i:** Transfer backups to NFS mounted directory of 'Aspera client server' **or** NFS share from IBM i + mount on 'Aspera client server' **& transfer.**

Access

- **Aspera Server configured as NFS server** & Migrated contents **shared within Power Virtual Server** (NFS mount on NIM, Network Install server or direct mount on "to-be" workload VSI).
- **Continue** with Power Virtual Server build procedure / applicat

The following procedure is typically used to generate migration content – For example OS specific backup files for AIX and IBM i. For application or DB related backup procedures, It is recommended to refer OEM documentation specific to the version.

Provided below are reference URLs to procedure for migrating from your on-premises environment into Power Systems Virtual Servers. Once the migration content is generated, it can then be transferred into the Power Virtual Server environment using the “accelerated transfer procedure” described in the next section.

Sample procedure for AIX:

1. https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-move-data-to-cloud
2. https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-restoring-aix-mksysb-image#creating-the-mksysb-image-on-the-source-aix-instance

Sample procedure for IBM i:

1. https://cloud.ibm.com/media/docs/downloads/power-iaas-tutorials/PowerVS_IBMi_Migration_Tutorial_v1.pdf
   → Pg 53

Reference architecture:

Architecture diagram at high level:

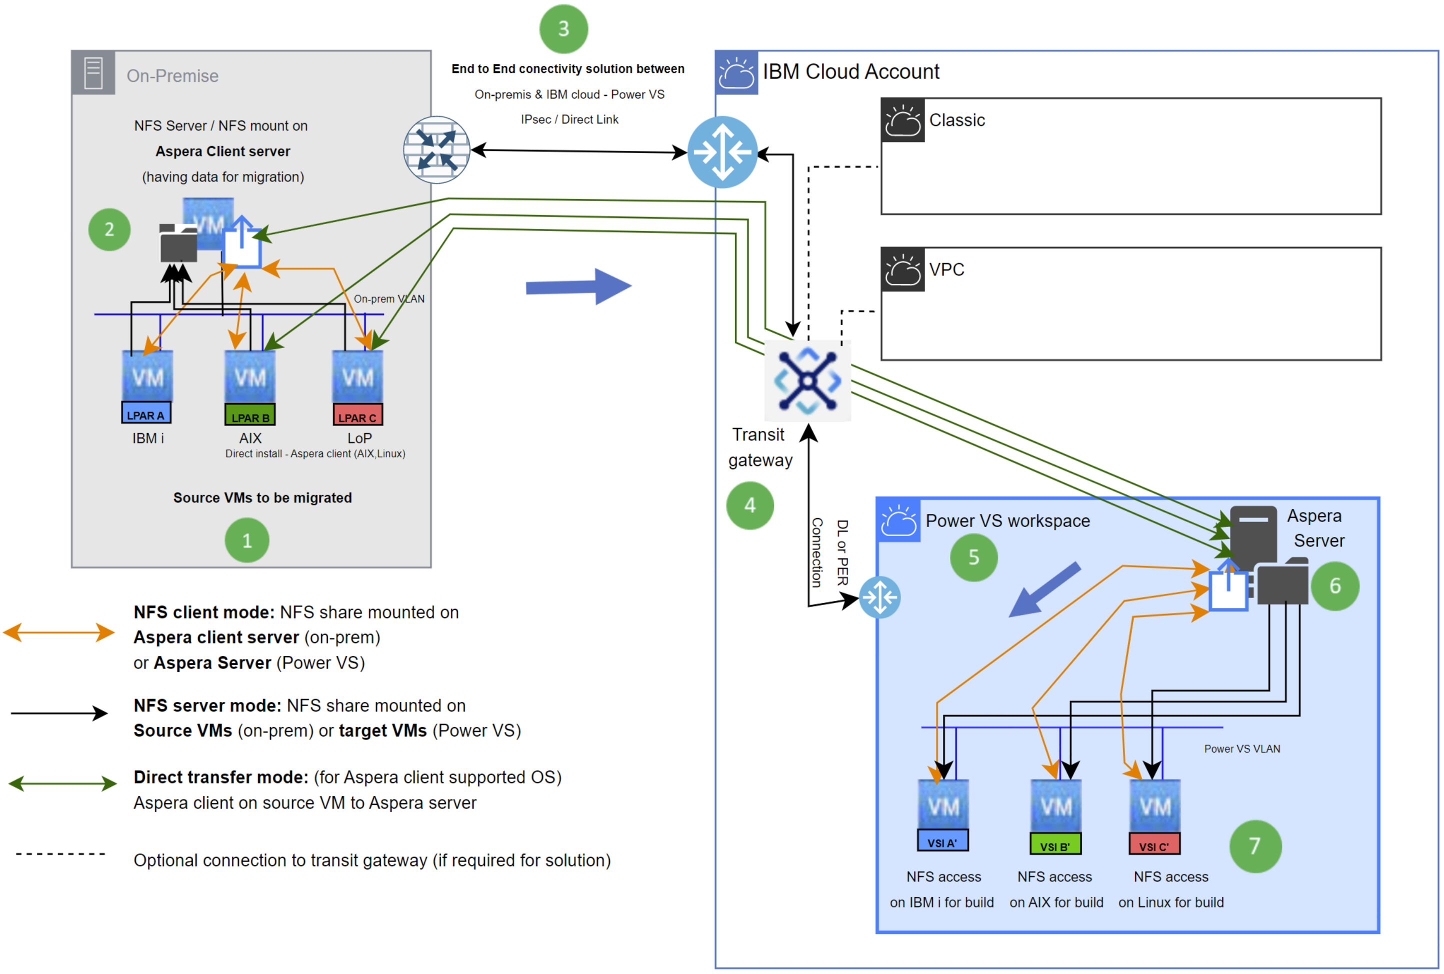

NFS client mode: NFS share mounted on **Aspera client server** (on-prem) or **Aspera Server** (Power VS)

NFS server mode: NFS share mounted on **Source VMs** (on-prem) or **target VMs** (Power VS)

Direct transfer mode: (for Aspera client supported OS) Aspera client on source VM to Aspera server

Optional connection to transit gateway (if required for solution)

Labelled overview of architecture

1 **Depiction of on-prem DC:** The diagram represents all three flavours of operating systems that is supported on POWER systems platform. In this reference architecture, these three source VMs named as LPAR A (IBM i), LPAR B (AIX) and LPAR C (Linux on POWER) are to be migrated to Power Virtual Server. These Operating systems can host any type of workload and could either be virtualised logical partitions on a POWER system host or native single OS deployments on one host.

2 **On-premises Aspera client server (optional for AIX / Linux)** dedicated for hosting the Aspera client. This source environment can be deployed in the following three modes for facilitating the transfer:

A. **Direct transfer mode:** (for Aspera client supported operating systems – AIX & Linux). With this mode, the Aspera desktop client software is installed directly on source VM (AIX or Linux on Power) for transferring OS/Application backups or other migration data to the Aspera server hosted in Power Virtual Server.

B. **NFS server mode:** NFS share mounted on Source VMs (on-premises). With this mode, the Aspera client server is allocated with sufficient storage space for migration. The block storage capacity of this Aspera Client server can be made available to the source VMs via NFS (by mounting on IBM i, AIX or Linux). Once the OS/application backups from each of the client VMs are either directly backed up to this mounted location or transferred to it, the Aspera client utility can then transfer it to Power Virtual Server.

C. **NFS client mode**: NFS share served by an IBM i, AIX or Linux VM running the NFS service is mounted onto the Aspera client server (on-premises).
OS/application backup or files for the purpose of build/migration made available on this NFS share directory which when mounted to the Aspera client server can then be transferred into Power Virtual Server.

3 High level depiction of private interconnectivity between on-premises DC and IBM cloud. This connectivity can be achieved using either IPsec VPN (Site to Site) or via IBM cloud Direct link. Bandwidth considerations should be considered based on the quantum of capacity and time available for data transfer/migration and considering post cutover access and usage scenarios.

4 **Transit gateway:** The transit gateway is often used to interconnect Intracloud environments such as classic resources, VPC resources and Power Virtual Server workspaces. If using VPN for VPC, a transit gateway is required to interconnect the VPC environment that hosts the VPN service with the target Power Virtual Server workspace.

This Power Virtual Server workspace could either be a PER (Power Edge Router) enabled workspace or a standard Power Virtual Server workspace. Standard Power Virtual Server workspaces are interconnected to the Transit gateway via Cloud connection/Direct links.

Various connectivity models are available for establishing private end to end connectivity between on-premises and Power Virtual Server. For more information, refer to IBM cloud documentation.

5 The Power Virtual Server workspace in the IBM cloud that will host the Aspera server and the “to-be” migrated VMs.

6 Aspera server hosted in the Power Virtual Server workspace. Like the modes explained in label 2, this Aspera server in Power Virtual Server can be configured with the below two modes:

A. **NFS server mode:** With this mode, the Aspera server is allocated with sufficient storage space for receiving the transferred files and is also configured to be an NFS server. Post successful transfer of migration content, the share is mounted on the target “to-be” VMs or either Network install server/NIM server for facilitating the build.

B. **NFS client mode**: NFS share from target VMs (within Power Virtual Server) is mounted onto the Aspera Server. **In this mode the “being” built/workload server (**IBM i, AIX or Linux VM) **provisioned with necessary storage capacity is configured to be the NFS server** and its share is mounted on the Aspera server. The mounted location is configured to be the target location for the migration content being transferred by the Aspera client from on-premises.

When to use NFS server mode/client mode?

1. NFS server mode is beneficial when the “target/workload VSI’s” **haven’t been built yet** and the **first set of migration content is required for primary OS build** (for example, mksysb or savsys backups) or when the build is scheduled for later.
2. NFS server mode is **likely to render better transfer performance** since the write is local to its block device rather than a remote write via NFS client.
3. NFS server mode is beneficial when migration content is being planned to be **centrally located within Power Virtual Server** for build purposes for more than one VSI.
4. With NFS server mode, there is **only one NFS server configuration** which **simplifies the access solution** post transfer, while with NFS client mode each target/workload VSI will need to be configured as NFS server.
5. NFS **client mode is beneficial to avoid a second copy** from the Aspera/NFS server into the “to-be” built VSI. It is **potentially beneficial for Database restores, post OS build.**
6. NFS client mode would require **less block storage on the Aspera server.**

7 **“To-be” or “being built”** target Power Virtual Server hosted **workload** VMs accessing / mounting the NFS share from the Aspera Server. These VMs could also be an AIX NIM server or IBM i Network install server.

Actualising the architecture:

Prerequisites:

1. **Aspera HSTS license.**
2. **Establish end-to-end connectivity** between on-premises and IBM Cloud into Power Virtual Server. Any of the below connectivity models are typically implemented:
    a. **IBM cloud direct link:** Private connectivity into IBM cloud using one of the Direct link offerings (Direct link dedicated or connect) + intracloud connectivity into Power Virtual Server (with or without Transit Gateway) leveraging PER or DLs as applicable.
    b. **Megaport:** Megaport has established Network to Network integrations into certain specific Power Virtual Server Colo's (DAL12, DAL13, FRA05, LON06, MON01, SYD05, OSA21, WDC04, and WDC06 data centres). Customers could choose to interconnect with Megaport as an NSP to then establish a virtual cross connect (VXC) into the Power Virtual Server workspace.
    c. IPsec from on-premises leveraging **IBM cloud VPC VPN gateway** (Policy based IPsec site-to-site) or host to Site.
    d. IPsec from on-premises to **Power Virtual Server hosted - IPsec VPN**
    e. IPsec from on-premises to via IBM cloud hosted BYO – **VRA in Classic / VPC**
3. **Infrastructure components:**
    a. Power Virtual Server workspace with VLAN.
    b. **Infra to host Aspera & NFS: At source (On-premises):** LPAR / VSI on prem that supports Aspera desktop client or direct install of Aspera desktop client for Linux on Power (LoP) or AIX; **At target (Power Virtual Server)**: Linux VSI or AIX VSI to host the Aspera HSTS server.
    c. **OS requirement:** Supported operating system – See release notes specific to the Aspera HSTS / desktop client version being installed.
    d. **SSH Server:** Version 7.0 or higher is recommended. (To check, run ssh -V)
    e. **Connectivity components (as per connectivity solution):** Transit gateway; IBM cloud VRA/Customer side Gateway / Firewall.
4. **Firewall Considerations:** Typically, port TCP 22 (default port for SSH) and UDP port 33001 are to be allowed on the firewalls enroute between the source (on-premises) and Target (Aspera HSTS server hosted in Power Virtual Server). SSH is used for the Aspera client to open a session, thereafter data transfer will occur through the designated UDP port. For more information, kindly refer the below link:
https://www.ibm.com/support/pages/firewall-considerations

Three step procedure:

The below three step procedure can be referred to for setting up the network based accelerated migration/transfer solution **‘after the prepare phase’** is completed.

Step A: Preparation at source (On-premises)
Step B: Preparation at target (Power Virtual Server)
Step C: Transfer & access

Step A: Preparation at source (On-premises).

Installation of Aspera desktop client (On-premises):
Aspera desktop client can be installed on POWER Linux, AIX platforms. Its currently not supported on IBM i.
NFS based approaches can be leveraged for IBM i transfers. IBM i NFS server shares the directory containing the migration content, which is then mounted onto a Linux server that is deployed with the Aspera client. The aspera client can then be used to transfer the mounted contents to the Power Aspera server which is hosted in the Power Virtual Server workspace.

a) Download the appropriate Aspera desktop client software:

All available client software packages:
https://www.ibm.com/products/aspera/downloads

Navigate to “**IBM Aspera Desktop Client**” and click on “**Download now**” and on the next page that lists the softwares available by platform, filter by platform “AIX” or “LinuxPPC” as applicable.

Either download directly to the source on-premises - Aspera client server or download to a PC/workstation and then SCP the file over to the Aspera client server.

b) To install,

```
rpm -ivh < Aspera client installer RPM file name>
```

For example,

```
rpm -ivh ibm-aspera-desktopclient-4.4.2.572-linux-ppc64le-release.rpm
```

AIX: If installing directly to the source AIX LPAR, then download/transfer the appropriate binary and perform the installation:

```
bash <Aspera client installer file>
```

For example,

```
bash ibm-aspera-desktopclient-4.4.1.23-aix-7.1-ppc64-release.sh
```

c) Post installation, to verify that the installation has been successful:

```
ascp -A
```

You should get a similar output.

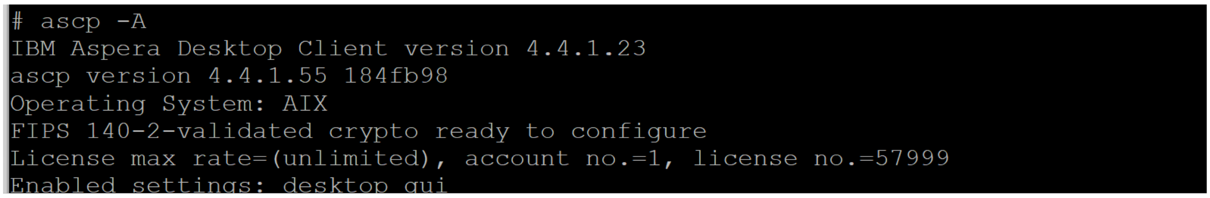

```
# ascp -A
IBM Aspera Desktop Client version 4.4.1.23
ascp version 4.4.1.55 184fb98
Operating System: AIX
FIPS 140-2-validated crypto ready to configure
License max rate=(unlimited), account no.=1, license no.=57999
Enabled settings: desktop gui
```

Step B: Preparation at target (Power Virtual Server)

Given that this guide is for direct transfer – i, e. point to point from on-premises into Power Virtual Server, a network connectivity between the source and target environment is a prerequisite. There are various connectivity models to achieve this (IPsec, Direct links, etc.).

For Step B, in your Power Virtual Server workspace, we will now need to deploy a Linux VSI along with installing / configuring the Aspera HSTS server for which you have two deployment choices:

Option 1 – Automated deployment: Very easy to use click & deploy terraform based automated deployment via IBM cloud Schematics.

With this simplified automation, all you need to do is to create a schematics workspace, point to the below Github repo that has the terraform code and provide configuration inputs. For further information, refer to the readme section of the Github repo.

https://github.com/IBM/power-aspera-server

or

Option 2 – Manual deployment: Perform a manual deployment on the IBM cloud portal by following the procedure provided in the "Manual configuration" section below:

Manual configuration

These are the steps to be done in the target system (Power Virtual Server) in IBM cloud.

1. Provision a Linux VSI in your Power Virtual Server workspace with the required storage capacity necessary for migration assigned to it. As a starting point you may choose to allocate 2 cores (POWER 9) and 8 GB of Memory which can then be dynamically changed after monitoring the utilisation.

2. Download the Aspera HSTS server rpm file from https://www.ibm.com/products/aspera/downloads and transfer it to the provisioned Linux VSI.
   Navigate to the Aspera Server software section → click on "**download now**" → in the next page from the "**change your selection**" section, set the platform as "**Linux PPC**" which will list all the HSTS Aspera server versions available for download → Select and download the desired version required, preferably the latest.

   Example: Download link of HSTS Version 4.4 for POWER Linux.
   https://www.ibm.com/support/fixcentral/swg/downloadFixes?parent=ibm%7EOther%20software&product=ibm/Other+software/IBM+Aspera+High-Speed+Transfer+Server&release=All&platform=All&function=fixId&fixids=pvt_ibm-aspera-hsts-4.4.1.96-linux-ppc64le-release%3A977657635136772096&includeRequisites=1&includeSupersedes=0&downloadMethod=http

3. Installation of Aspera:
   Regular rpm install (`rpm -ivh <file name>`)
   For example, `rpm -ivh ibm-aspera-hsts-4.4.1.96-linux-ppc64le-release.rpm`

4. Apply License:
   a. Obtain the Aspera HSTS license key ().
   b. Create a new file named "aspera-license" in "/opt/aspera/etc/"
   c. Paste the license into the newly created "aspera-license" file.

5. Run the below command to verify a successful installation.
   ascp -A

The output should list the license details of the Aspera server. Example of perpetual license below:

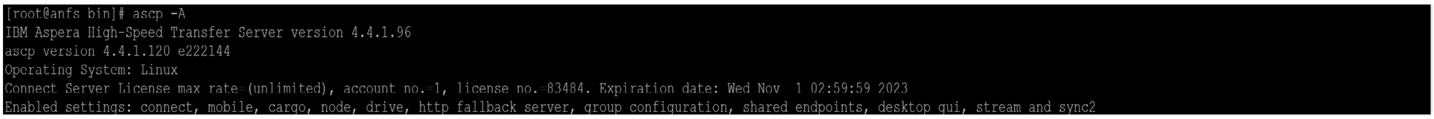

```
[root@anfs bin]# ascp -A
IBM Aspera High-Speed Transfer Server version 4.4.1.96
ascp version 4.4.1.120 e222144
Operating System: Linux
Connect Server License max rate=(unlimited), account no.=1, license no.=83484. Expiration date: Wed Nov  1 02:59:59 2023
Enabled settings: connect, mobile, cargo, node, drive, http fallback server, group configuration, shared endpoints, desktop gui, stream and sync2
```

6. After verifying that the installation is successful, create a directory for housing the “to-be” migrated content, based on the type of transfer mode desirable:
    a. **NFS server mode:** This is when the Aspera server’s block storage will be the target landing storage area for the “to-be” transferred migration data/files.
    Aspera target directory: In the mount point of the file system which has sufficient storage space, create a directory in it. For example, I am creating a directory in / file system.

    ```
    mkdir /asdir
    ```

    b. **NFS client mode:** This mode is an NFS mount of a directory shared from the Target “to-be” VSI.
    Once mounted, the procedure is as per previous step, or you can have Aspera transfer to a network share location.

Configuration post installation:

Configure the aspera server with the target destination directory. Also configure a standard OS user to use aspera. In the below example, I am setting the “root” user for Aspera.

```
asconfigurator -F "set_user_data;user_name,<user id>;absolute,<path to target directory>"
```

```
asconfigurator -F "set_user_data;user_name,root;absolute,/asdir"
```

For more information, refer to the examples in the transfer section provided in Step C.

Step C: Transfer & Access

NFS server mode: Configure the Aspera Server VSI to also function as an NFS server to then share the directory containing the migration content. This directory can then mounted/accessed from the VMs (target workload VMs) being built or from IBM i Network install server/NIM server.

Post completion of Step A, Step B, and after ensuring that the end-to-end connectivity has been established, we are now ready to transfer the migration content from on-premises into the Power Virtual Server environment. The Aspera command (ascp – aspera secure copy) can be configured with a variety of options that can be found in the Aspera HSTS guide.

Provided below are few examples with the “ascp” executable which is command line program. Another command line FASP transfer program is “ascp4”, that is optimized for the transfers of many small files. It has similar capabilities as ascp.

Transfer:

From the on-premises source system, which could either be the source AIX / Linux server or any OS platform that is supported by the Aspera desktop client utility, file transfer of migration content is initiated into the Aspera server which is hosted in the Power Virtual Server workspace.
The ascp command provides a wide range of controls / settings that are customisable. Provided below are few examples.
In some of the examples, I am authenticating using a root user of the Aspera server given that the Aspera server in Power Virtual Server is solely used for the transfer. However, you can configure non-root system users as Aspera transfer users by following the instructions in this link (see example 2).
https://www.ibm.com/docs/en/ahts/4.4.x?topic=suug-setting-up-transfer-users-1
Additionally, key based authentication is also supported and is more secure. If password-based authentication is preferred, then ensure a strong password is used.
For more information on setting up key based authentication please see the documentation:
https://www.ibm.com/docs/en/ahts/4.4.x?topic=suug-setting-up-users-public-key-server-3
https://www.ibm.com/docs/en/ahts/4.4.x?topic=line-creating-ssh-keys”

Example 1: To transfer a single file to the Aspera server
ascp <path to source file> <user>@<Aspera Server IP>:/
For example, transfer a file named 1GB.dat
Aspera source directory: /aspsrc
ascp -l 1g /aspsrc/1GB.dat root@192.168.50.243:/
*In the above example the file could also be from an NFS mount point

Example 2: To transfer a single file to Aspera server using a non-root user and to a directory in the user’s home location.
For additional considerations & options, do refer:
https://www.ibm.com/docs/en/ahts/4.4.x?topic=suug-setting-up-transfer-users-1

On the Aspera server, as a root user run

asconfigurator -F "set_user_data;user_name,<user id>;absolute,<docroot>"

For example, transfer a file named 1GB.dat using a non-root user.

Aspera source directory: /asdir

On the target Aspera server containing the user:

$mkdir /home/boby/asusertest

Now as a root user run,

asconfigurator -F "set_user_data;user_name,boby;absolute,/home/boby/asusertest"

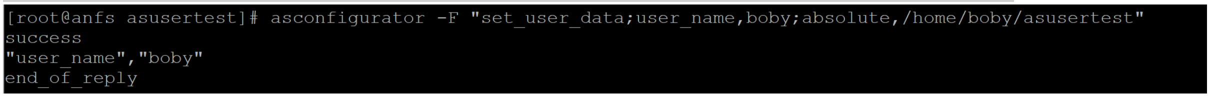

```
[root@anfs asusertest]# asconfigurator -F "set_user_data;user_name,boby;absolute,/home/boby/asusertest"
success
"user_name","boby"
end_of_reply
```

From the Aspera client machine, initiate a transfer with the non-root user.

ascp <path to source file> <user>@<Aspera Server IP>:/

For example, ascp 1GB.dat root@10.50.50.74:/test1GB.dat

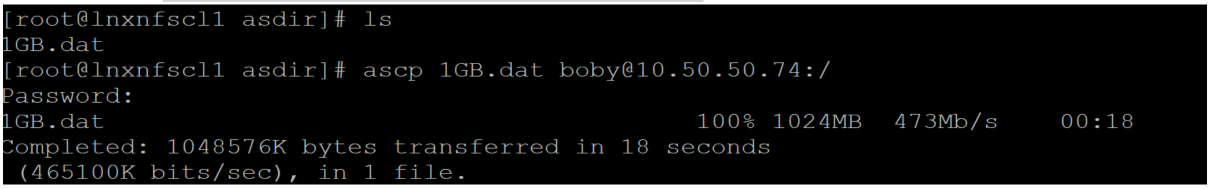

```
[root@lnxnfscl1 asdir]# ls
1GB.dat
[root@lnxnfscl1 asdir]# ascp 1GB.dat boby@10.50.50.74:/
Password:
1GB.dat                                         100% 1024MB   473Mb/s    00:18
Completed: 1048576K bytes transferred in 18 seconds
 (465100K bits/sec), in 1 file.
```

Verify that the file has been transferred to the user's target directory of the Aspera server.

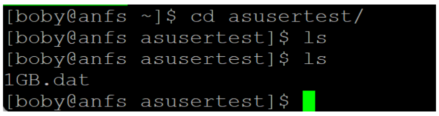

```
[boby@anfs ~]$ cd asusertest/
[boby@anfs asusertest]$ ls
[boby@anfs asusertest]$ ls
1GB.dat
[boby@anfs asusertest]$
```

Example 3: To transfer at rates "up to" the specified target rate.

Default is 10000 Kbps. This option accepts suffixes G or g for giga, M or m for mega, K or k for kilo,

For example, transfer a single file named 1GB.dat

Aspera source directory: /local-dir

To set as 100 Mbps

ascp -l 100m /local-dir/file1.dat root@10.0.0.2:/

To set rate as 1 Gbps

ascp -l 1G /local-dir/file1.dat root@10.0.0.2:/

To set rate as 1.3 Gbps

ascp -l 1.3G /local-dir/file1.dat root@10.0.0.2:/

Example 4: To transfer (upload & download examples) "all files" in a source directory.

ascp <path to source directory> <user>@<aspera server IP address>:/

Upload

For example, Aspera source directory: /asdir

ascp /asdir boby@10.50.50.74:/

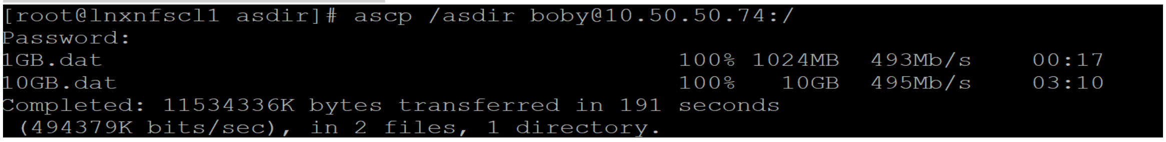

```
[root@lnxnfscl1 asdir]# ascp /asdir boby@10.50.50.74:/
Password:
1GB.dat                                        100% 1024MB   493Mb/s    00:17
10GB.dat                                       100%   10GB   495Mb/s    03:10
Completed: 11534336K bytes transferred in 191 seconds
 (494379K bits/sec), in 2 files, 1 directory.
```

Download:
Similarly, to reverse the above – for any scenario requiring a directory to be transferred back to on-premises, refer to the below example:
ascp -d -l <transfer speed> <user>@<IPaddress>:/<directory to transfer>/ /<download destination target directory>

For example, ascp -d -l 500M root@10.50.50.74:/aspulldir/ /asdir/newdir
Aspera server containing directory to download:

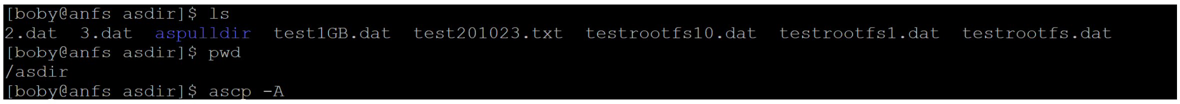

```
[boby@anfs asdir]$ ls
2.dat  3.dat  aspulldir  test1GB.dat  test201023.txt  testrootfs10.dat  testrootfs1.dat  testrootfs.dat
[boby@anfs asdir]$ pwd
/asdir
[boby@anfs asdir]$ ascp -A
```

Aspera client server invoking the pull / download.

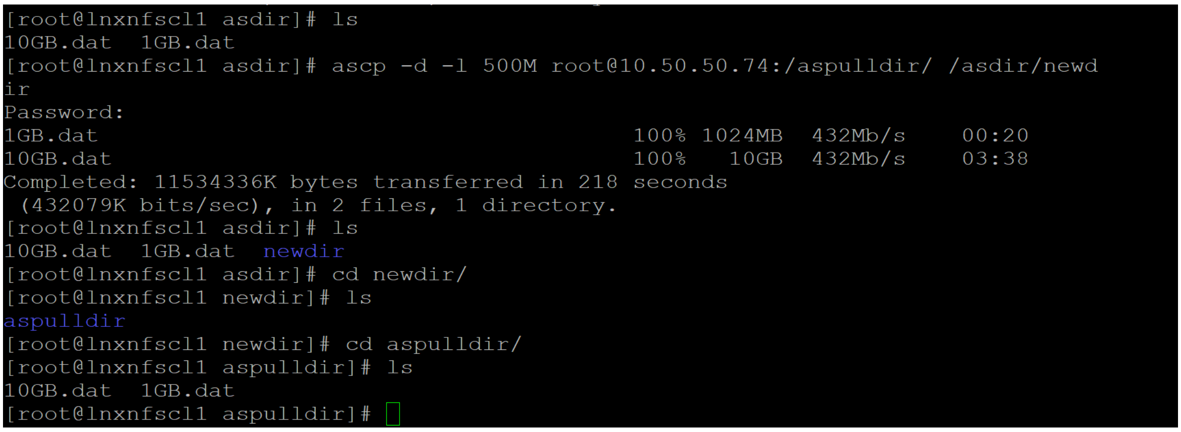

```
[root@lnxnfscl1 asdir]# ls
10GB.dat  1GB.dat
[root@lnxnfscl1 asdir]# ascp -d -l 500M root@10.50.50.74:/aspulldir/ /asdir/newd
ir
Password:
1GB.dat                                      100% 1024MB  432Mb/s    00:20
10GB.dat                                     100%   10GB  432Mb/s    03:38
Completed: 11534336K bytes transferred in 218 seconds
 (432079K bits/sec), in 2 files, 1 directory.
[root@lnxnfscl1 asdir]# ls
10GB.dat  1GB.dat  newdir
[root@lnxnfscl1 asdir]# cd newdir/
[root@lnxnfscl1 newdir]# ls
aspulldir
[root@lnxnfscl1 newdir]# cd aspulldir/
[root@lnxnfscl1 aspulldir]# ls
10GB.dat  1GB.dat
[root@lnxnfscl1 aspulldir]#
```

Example 5: Public key authentication (SSH Key)
ascp -i <location to private key> <path to source file> <user>@<Aspera server's IP address>:/<target file name>
In the below example, I am authenticating as a root user with the private key. You can do the same for a standard user as well. Additionally, in the example, I am setting the desired transfer rate as 800 Mbps and renaming the transferred file once it's been sent to the target.

ascp -l <target transfer rate> -i <location to private key> <path to source file> <user>@<IP address>:/<target file name>
For example, ascp -l 800M -i ~/.ssh/Prvt_key /asdir/1GB.dat root@10.50.50.74:/1GBasp.dat

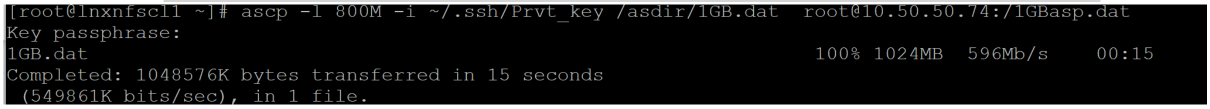

```
[root@lnxnfscl1 ~]# ascp -l 800M -i ~/.ssh/Prvt_key /asdir/1GB.dat  root@10.50.50.74:/1GBasp.dat
Key passphrase:
1GB.dat                                                   100% 1024MB  596Mb/s    00:15
Completed: 1048576K bytes transferred in 15 seconds
 (549861K bits/sec), in 1 file.
```

For additional info, kindly refer:
https://www.ibm.com/docs/en/ahts/4.4?topic=suug-setting-up-users-public-key-server-3
https://www.ibm.com/docs/en/ahts/4.4?topic=line-creating-ssh-keys
Other useful "ascp" command options:
a. Restricting Aspera user permissions with aspshell.
By default, all system users can establish a FASP connection and are only restricted by file permissions. You can restrict the system user's file operations by assigning them to use the aspshell, which allows limited operations such as running Aspera uploads/downloads, browsing, listing, creating, renaming, or deleting contents.
To change the user login shell to aspshell:

```
# sudo usermod -s /bin/aspshell <username>
```

Confirm that the user's shell is updated by running the following command and looking for “/bin/aspshell” at the end of the output:

```
# grep <username> /etc/passwd
```

Output:

```
<username>:x:501:501:...:/home/username:/bin/aspshell
```

b. Resume transfer

-k {0|1|2|3}

Enable the resuming of partially transferred files at the specified resume level. (Default: 0).
Specify this option for the first transfer or it doesn’t work for subsequent transfers. Resume levels:

-k 0 – Always retransfer the entire file.

-k 1 – Compare file attributes and resume if they match, and retransfer if they do not.

-k 2 – Compare file attributes and the sparse file checksum; resume if they match, and retransfer if they do not.

-k 3 – Compare file attributes and the full file checksum; resume if they match, and retransfer if they do not.

If a complete file exists at the destination (no .aspx), the source and destination file sizes are compared. If a partial file and a valid .aspx file exist at the destination, the source file size and the file size that is recorded in the .aspx file are compared.

Reference for additional options & commands:

The below link to Aspera documentation provides wealth of information and advanced transfers options. The below link is for Aspera version 4.4 which can be changed to your appropriate version on the page.

https://www.ibm.com/docs/en/ahts/4.4?topic=atfcl-ascp-command-reference-1
https://www.ibm.com/docs/en/ahts/4.4?topic=atfcl-ascp-general-examples-1

Access:

These are sample steps for configuring access via NFS. You may want to tailor your NFS configuration to suit your organisation's security or policy requirements. Kindly refer to the NFS documentation for the specific OS.

NFS server configuration (On the Aspera Server):

The below example is for RHEL / Cent OS and for NFS V3. NFS v4 is similar and has a domain-based concept with several enhancements. By default, NFS utility should already be installed in the OS, if it is not available for some reason have it installed by referring to the OS and version specific documentation.

1. Start the NFS server in the Aspera server.

   ```
   systemctl start nfs-server
   ```

   Check the status after starting.

   ```
   systemctl status nfs-server.service
   ```

2. Edit the exports file to set the file system/directory that we created in Step B. For all options and for a complete description, check the man pages **(man exports)**

   ```
   vi /etc/exports
   ```

   In my example: I intend to share the directory containing the migrated content ("/asdir") to any machine on these two Power Virtual Server networks (Network with CIDR - 192.168.50.0/24 and 10.50.50.0/24)

   ```
   /asdir 192.168.50.0/24(rw,sync)
   /asdir 10.50.50.0/24(rw,sync)
   ```

3. Run the command exportfs -arv to have the **nfsd deamon** read the `/etc/exports` file for it to take effect. For every change done to the /etc/exports file, we need to run the command again.

4. Verify that the NFS server is configured as intended.

   ```
   exportfs -s
   ```

   You should get an output like below:

   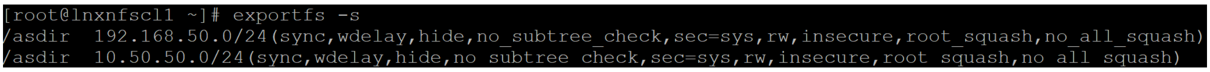

   ```
   [root@lnxnfscl1 ~]# exportfs -s
   /asdir  192.168.50.0/24(sync,wdelay,hide,no_subtree_check,sec=sys,rw,insecure,root_squash,no_all_squash)
   /asdir  10.50.50.0/24(sync,wdelay,hide,no_subtree_check,sec=sys,rw,insecure,root_squash,no_all_squash)
   ```

   Depiction of contents within /asdir that will be shared via NFS.

   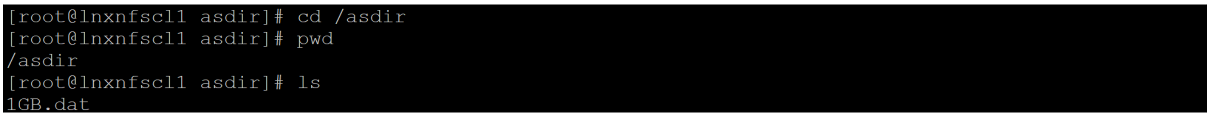

   ```
   [root@lnxnfscl1 asdir]# cd /asdir
   [root@lnxnfscl1 asdir]# pwd
   /asdir
   [root@lnxnfscl1 asdir]# ls
   1GB.dat
   ```

NFS access from AIX (target VSI):

This section assumes that an AIX VSI is provisioned within the Power Virtual Server workspace. To access the share on the AIX machine, be it from a NIM server or the target workload VSI in Power Virtual Server.

1. To list all the mounts that are available on the server by running the following command.
   showmount -e Public_IP_address_of_the_NFS_server
   For example,
   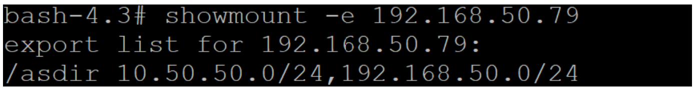
   ```
   bash-4.3# showmount -e 192.168.50.79
   export list for 192.168.50.79:
   /asdir 10.50.50.0/24,192.168.50.0/24
   ```

2. Create a local directory and map the NFS server directory to it.
   For example,
   mkdir <directory>
   mount <NFS server IP>:<path of shared directory> <local directory to
   mkdir /asmount
   mount with defaults:
   mount 192.168.50.79:/asdir /asmount
   mount with specify wsize and rsize
   mount -o wsize=1024,rsize=1024 192.168.50.79:/asdir /asmount
   wsize and rsize specify the buffer size to use for read and write request.
   Reference:
   NFS mount command
   https://www.ibm.com/docs/en/aix/7.2?topic=m-mount-command

   NFS tuning on client for AIX
   https://www.ibm.com/docs/en/aix/7.2?topic=performance-nfs-tuning-client
   https://www.ibm.com/docs/en/aix/7.2?topic=client-read-write-size-adjustments

3. To list the mounted file systems, enter either of the below commands:
   mount or the df -g
   For example,

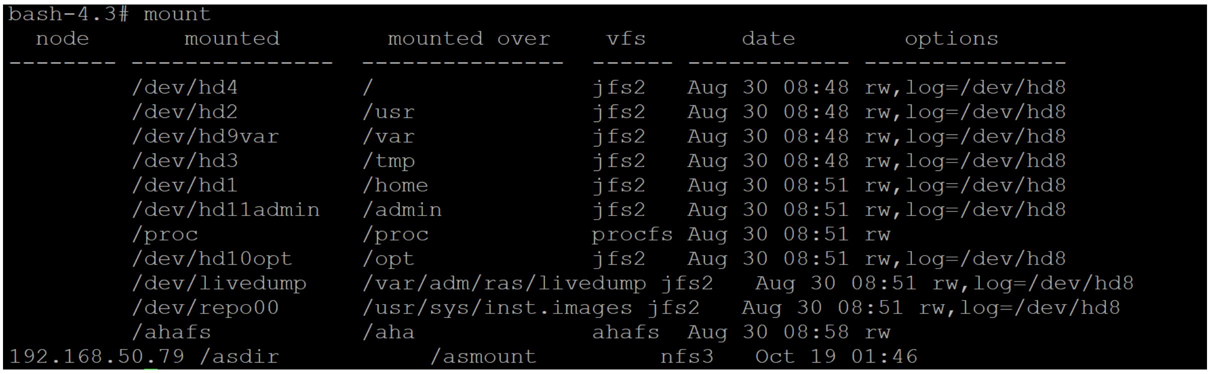

```
bash-4.3# mount
  node       mounted        mounted over    vfs       date        options
-------- --------------- --------------- ------ ------------ ---------------
         /dev/hd4        /               jfs2   Aug 30 08:48 rw,log=/dev/hd8
         /dev/hd2        /usr            jfs2   Aug 30 08:48 rw,log=/dev/hd8
         /dev/hd9var     /var            jfs2   Aug 30 08:48 rw,log=/dev/hd8
         /dev/hd3        /tmp            jfs2   Aug 30 08:48 rw,log=/dev/hd8
         /dev/hd1        /home           jfs2   Aug 30 08:51 rw,log=/dev/hd8
         /dev/hd11admin  /admin          jfs2   Aug 30 08:51 rw,log=/dev/hd8
         /proc           /proc           procfs Aug 30 08:51 rw
         /dev/hd10opt    /opt            jfs2   Aug 30 08:51 rw,log=/dev/hd8
         /dev/livedump   /var/adm/ras/livedump jfs2   Aug 30 08:51 rw,log=/dev/hd8
         /dev/repo00     /usr/sys/inst.images jfs2   Aug 30 08:51 rw,log=/dev/hd8
         /ahafs          /aha            ahafs  Aug 30 08:58 rw
192.168.50.79 /asdir          /asmount        nfs3   Oct 19 01:46
```

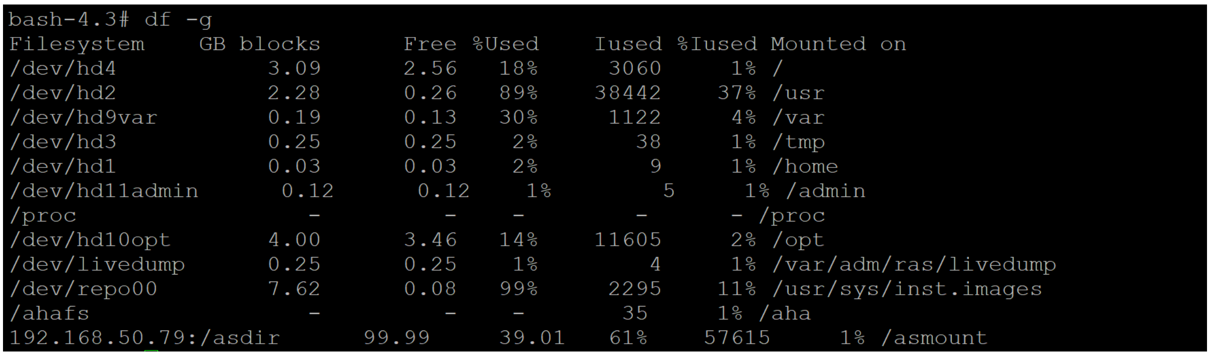

```
bash-4.3# df -g
Filesystem    GB blocks      Free %Used    Iused %Iused Mounted on
/dev/hd4           3.09      2.56   18%     3060     1% /
/dev/hd2           2.28      0.26   89%    38442    37% /usr
/dev/hd9var        0.19      0.13   30%     1122     4% /var
/dev/hd3           0.25      0.25    2%       38     1% /tmp
/dev/hd1           0.03      0.03    2%        9     1% /home
/dev/hd11admin      0.12      0.12    1%        5     1% /admin
/proc                 -         -    -        -      - /proc
/dev/hd10opt       4.00      3.46   14%    11605     2% /opt
/dev/livedump      0.25      0.25     1%       4     1% /var/adm/ras/livedump
/dev/repo00        7.62      0.08   99%     2295    11% /usr/sys/inst.images
/ahafs                -         -    -       35     1% /aha
192.168.50.79:/asdir      99.99     39.01   61%    57615     1% /asmount
```

Accessing the contents of the share:

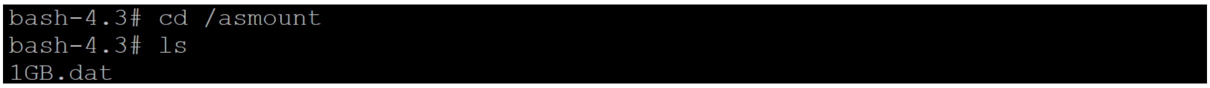

```
bash-4.3# cd /asmount
bash-4.3# ls
1GB.dat
```

NFS access from IBM i (target VSI):

This section assumes that an IBM i VSI is provisioned within the Power Virtual Server workspace. This VSI could be an IBM i Network install server or a target IBM i server that is already built and now needs the migrated data for application / data restoration purposes.

Connect to the IBM i server either through the Power Virtual Server console or IBM i Access Client Solutions (ACS). Login to the OS as an admin user to

1. Add a host table entry for the Aspera server.
   CFGTCP → Option 10 and Option 1 to Add.
2. Start the below NFS daemons, one being a Block I/O daemon to handle bulk I/O traffic while the second being an NFS RPC bind daemon.
   STRNFSSVR SERVER(*BIO)
   STRNFSSVR SERVER(*RGY)
   Alternatively, run STRNFSSVR SERVER(*ALL) to start all daemons.
3. Create a directory in IFS to mount over:
   CRTDIR DIR (‘<directory name>’)
   For example,
   CRTDIR DIR('/asmount')
4. Mount the file system onto the directory created in the previous step:
   ADDMFS TYPE(*NFS) MFS(‘<NFS server:/directory to mount>’) MNTOVRDIR('</local directory created to mount over>')
   OPTIONS('rw,suid,retry=5,rsize=32768,wsize=32768,timeo=20,retrans=5,acregmin=30,acregmax=60,acdirmin=30,acdirmax=60,hard,async,sec=sys,vers=4,nocache')

   For example,
   ADDMFS TYPE(*NFS) MFS('nfsserver:/asdir') MNTOVRDIR('/asmount')
   OPTIONS('rw,suid,retry=5,rsize=32768,wsize=32768,timeo=20,retrans=5,acregmin=30,acregmax=60,acdirmin=30,acdirmax=60,hard,async,sec=sys,vers=4,nocache')

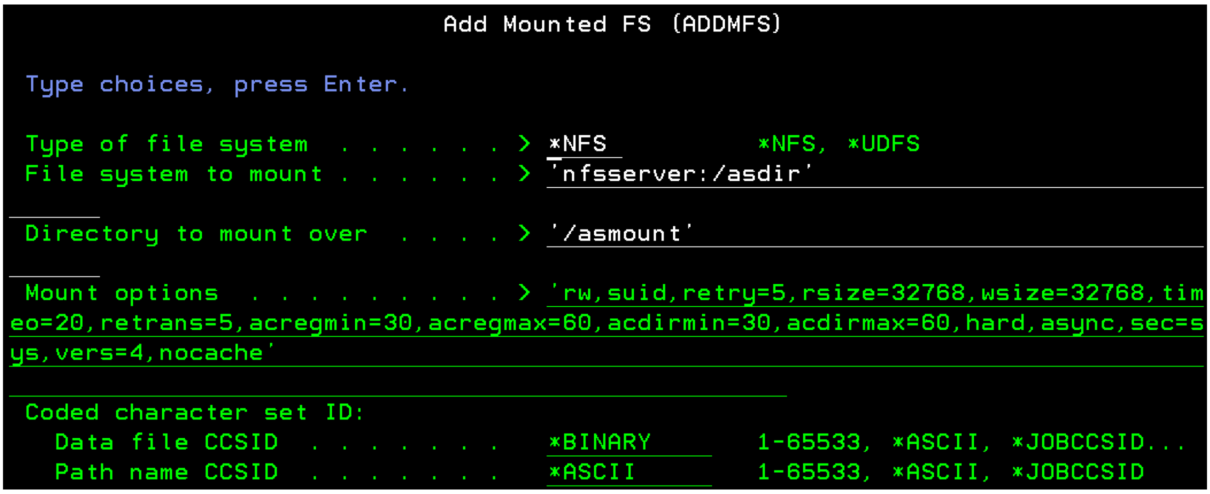

```
                            Add Mounted FS (ADDMFS)

 Type choices, press Enter.

 Type of file system  . . . . . . > *NFS          *NFS, *UDFS
 File system to mount . . . . . . > 'nfsserver:/asdir'

 Directory to mount over  . . . . > '/asmount'

 Mount options  . . . . . . . . . > 'rw,suid,retry=5,rsize=32768,wsize=32768,tim
eo=20,retrans=5,acregmin=30,acregmax=60,acdirmin=30,acdirmax=60,hard,async,sec=s
ys,vers=4,nocache'

 Coded character set ID:
   Data file CCSID  . . . . . . .   *BINARY       1-65533, *ASCII, *JOBCCSID...
   Path name CCSID  . . . . . . .   *ASCII        1-65533, *ASCII, *JOBCCSID
```

5. Once the mount is successful, you can access the contents via the IFS.
   WRKLNK OBJ('</local directory mounted>')
   For example,
   WRKLNK OBJ('/asmount')

Access the contents of the share or copy it into the respective area of the file system for further restore / application use.

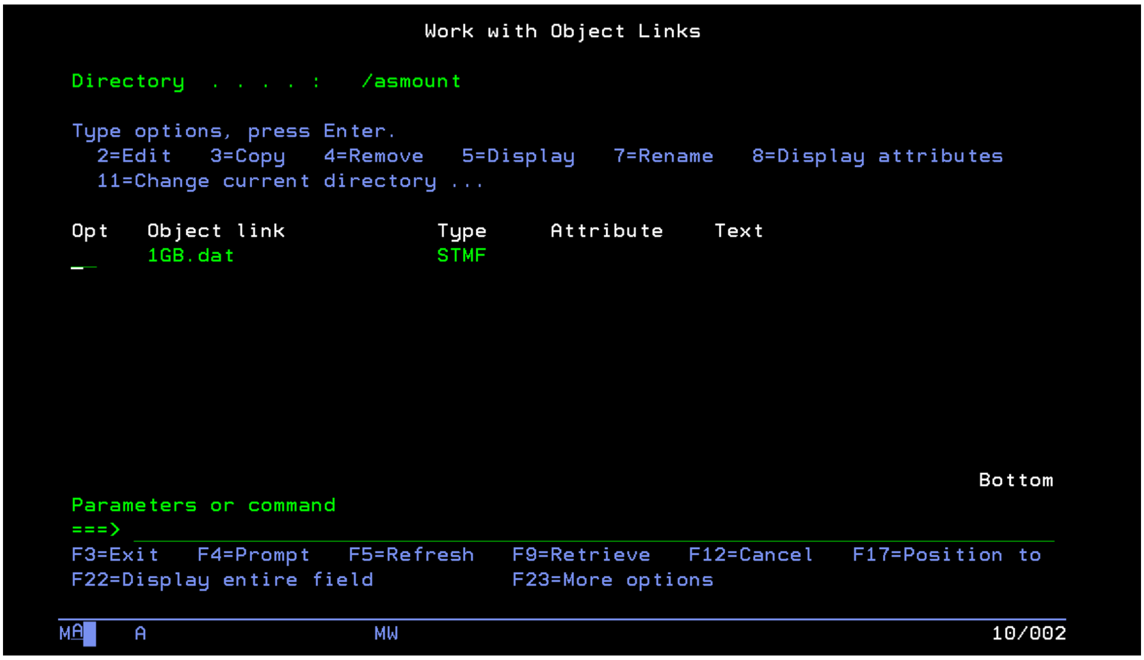

```
                            Work with Object Links

   Directory  . . . . :   /asmount

   Type options, press Enter.
     2=Edit   3=Copy   4=Remove   5=Display   7=Rename   8=Display attributes
     11=Change current directory ...

   Opt   Object link            Type     Attribute    Text
   _     1GB.dat                STMF

                                                                        Bottom
   Parameters or command
   ===> ____________________________________________________________________
   F3=Exit   F4=Prompt   F5=Refresh   F9=Retrieve   F12=Cancel   F17=Position to
   F22=Display entire field           F23=More options

 MA     A                  MW                                          10/002
```

Note: The Mount options can be configured and tuned for specific performance or security needs though I have set the default in the above example.

NFS client mode: This mode entails each "target/workload" VSI (IBM i, AIX or Linux) or IBM i Network install / NIM server in the Power Virtual Server workspace to be configured as the "NFS server", then make available a "share with sufficient storage capacity" to hold the migration data to be transferred. This network share is mounted on the Aspera server and then used as a destination for data transfer from on-premises, effectively ensuring the write happens on the "target/workload" VSI. Kindly refer NFS documentation for AIX or IBM i for the NFS configuration procedure.
A similar approach can be done on-premises as well as depicted in the below architecture for IBM i:

NFS client mode architecture for IBM i:

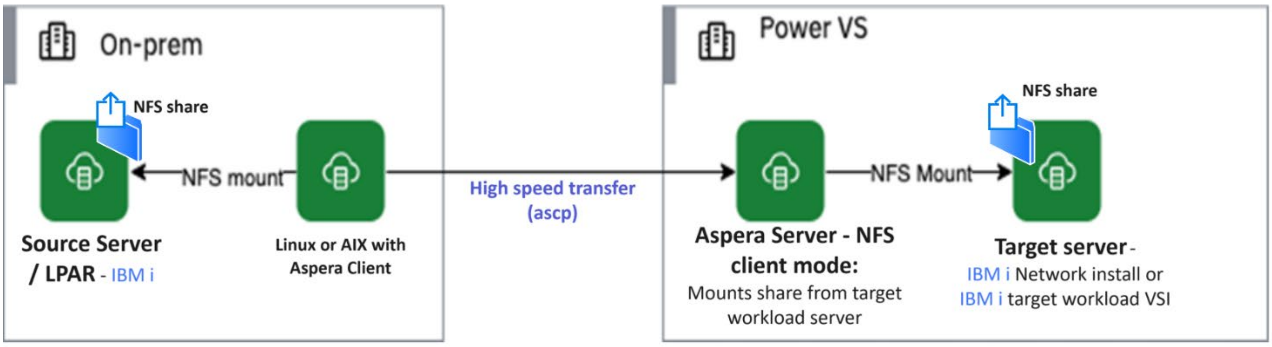
On-prem
NFS share
NFS mount
Source Server / LPAR - IBM i
Linux or AIX with Aspera Client
High speed transfer (ascp)
Power VS
NFS share
NFS Mount
Aspera Server - NFS client mode:
Mounts share from target workload server
Target server - IBM i Network install or IBM i target workload VSI

NFS client mode architecture for AIX or Linux:

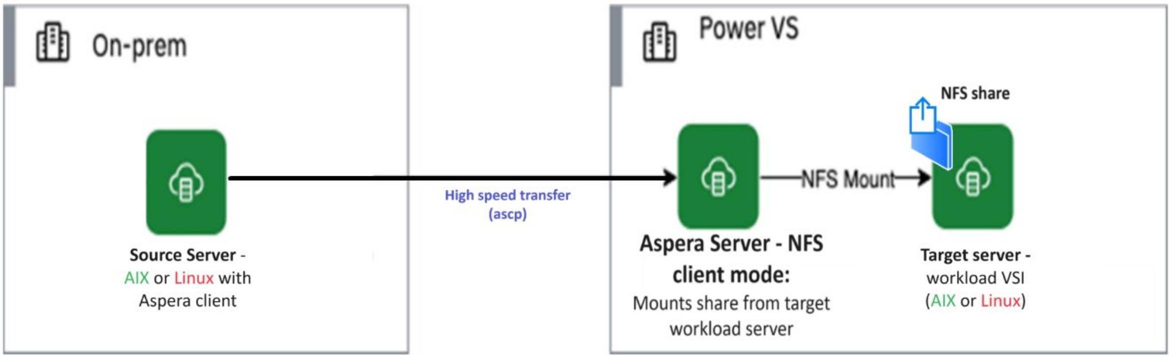
On-prem
Source Server - AIX or Linux with Aspera client
High speed transfer (ascp)
Power VS
NFS share
NFS Mount
Aspera Server - NFS client mode:
Mounts share from target workload server
Target server - workload VSI (AIX or Linux)

Build procedure post accelerated transfer:

Build procedure for AIX: The following procedure is typically followed to restore a mksysb image once the image is transferred into Power Virtual Server. This procedure leverages a helper VM / alternate disk install procedure that is described in cloud docs.

1. https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-restoring-aix-mksysb-image#create-aix-conversion-vm

Build procedure for IBM i: Similarly, for IBM i build a helper VM is also leveraged with the installation of the target workload VM done through an IBM i Network install server. Reference URLs below:

1. https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-preparing-install-server
2. https://www.ibm.com/docs/en/i/7.4?topic=storage-virtual-optical-using-network-file-system
3. https://www.redbooks.ibm.com/redpapers/pdfs/redp4937.pdf
4. https://cloud.ibm.com/media/docs/downloads/power-iaas/Cloud_Optical_Repository.pdf

Additional considerations & best practices:

1. **Aspera Performance:** There are many factors that influence transfer performance. Though the available end to end network bandwidth between source and destination is one key factor, other factors such as I/O performance for read and write and compute/memory resources on the hosts (both source and target). The below link provides useful information on how to determine bottleneck for troubleshooting purposes.

   https://www.ibm.com/support/pages/performance-testing

   a) **Host resources:** Ensure that the server that will be used as the source is scalable (if virtual) or well-resourced (if physical server). The compute and memory capacity required will depend on additional factors (For example, amount of data to backup, write / read throughput, Network speed, workload levels (if aspera client is being installed to source LPAR (AIX. LoP) etc.).

   b) **I/O performance of disk subsystem**: Recommend that the NFS storage at source is backed by high performance disk (For example, NVMe SSD) either directly attached, with high performance raid controller or via high performance SAN to maximise transfer rates.

c) **Network**: High performance network from source LPAR/server to the switch and beyond up to the destination in Power Virtual Server as per the network connectivity model being implemented.

The below link provides comprehensive info for identifying performance bottlenecks and for taking corrective actions.
https://www.ibm.com/support/pages/identifying-and-understanding-performance-bottlenecks-aspera-transfers

2. **Compression:** To optimise the transfer, it's recommended to compress the source files (backup files – from prepare phase) prior to initiating the transfer. Compression reduces the file size, ensuring only compressed content is transferred via the network.

   For non-Aspera based traditional TCP based transfer, for example, SCP - use the "-C" option to enable compression.

3. **Security:** Security is a broad topic, while it's beyond the scope of this document. Adhering to security best practices helps protect against threats. This transfer solution leverages two key technologies – the first being Aspera for accelerated transfer and the second being NFS that is used either for accessing content for the transfer or after the transfer.
   Provided below are few points to take into consideration, for additional information kindly refer to OS specific documentation.

   a) **Disable public access:**
   As a best practice, do not "toggle on" the Public IP of the Aspera Server to prevent external threats, given that this point-to-point transfer solution is private network based without the need for a public IP to be configured to the VSI.

   b) Aspera & SSH security:

   Review how to improve ssh security for the Aspera server that is being hosted in Power Virtual Server. Its importance increases should you decide to toggle on the Public IP of the VSI.

   https://www.ibm.com/docs/en/ahts/4.4?topic=upgrades-installing-configuring-ssh-server#task_sdv_dnb_rpb__title__1

   Changing TCP port number for SSH:
   https://www.ibm.com/docs/en/ahts/4.4?topic=iu-configuring-ssh-server-1#task_jb2_3ms_gx__title__1

   c) NFS security:
   Even though the risk is low within a private network setup and for a Linux / AIX / IBM i OS environment, ransomware could encrypt all mounted remote drives. You can reduce your risk by mounting only the NFS drive for the duration of the copy/transfer. Then, remove the NFS-mounted drive for daily operations. Additionally, NFS exports can be done with least privileges such as read only, etc. access can be granted to specific hosts, etc. NFS security is dependent on the version of NFS being deployed. NFS version 4 offers enhanced security

capability. For configuration aspects, refer to the NFS documentation for the specify OS.

4. **Consolidated listing of Aspera related best practices: The below link summarizes settings for consideration.**

   https://www.ibm.com/docs/en/ahts/4.0?topic=a-aspera-ecosystem-security-best-practices-2

Conclusion:

Enterprises require the best migration and data transfer solution, one that can be tailored to deliver a dependable & predictable outcome. This need is further reinforced when the environment being migrated is complex – either higher number of LPARs, moderate to large capacity or a combination of both.

This next generation solution guide helps accelerate data transfers, helping get you across the line in a friction free manner rendering the best experience that you deserve.

****