

Digital Preservation in HPC Environments

Amanda Tomé, Digital Preservation Coordinator
Community Owned Workflows - March 4, 2026



Digital Research
Alliance of Canada

Alliance de recherche
numérique du Canada



Agenda.


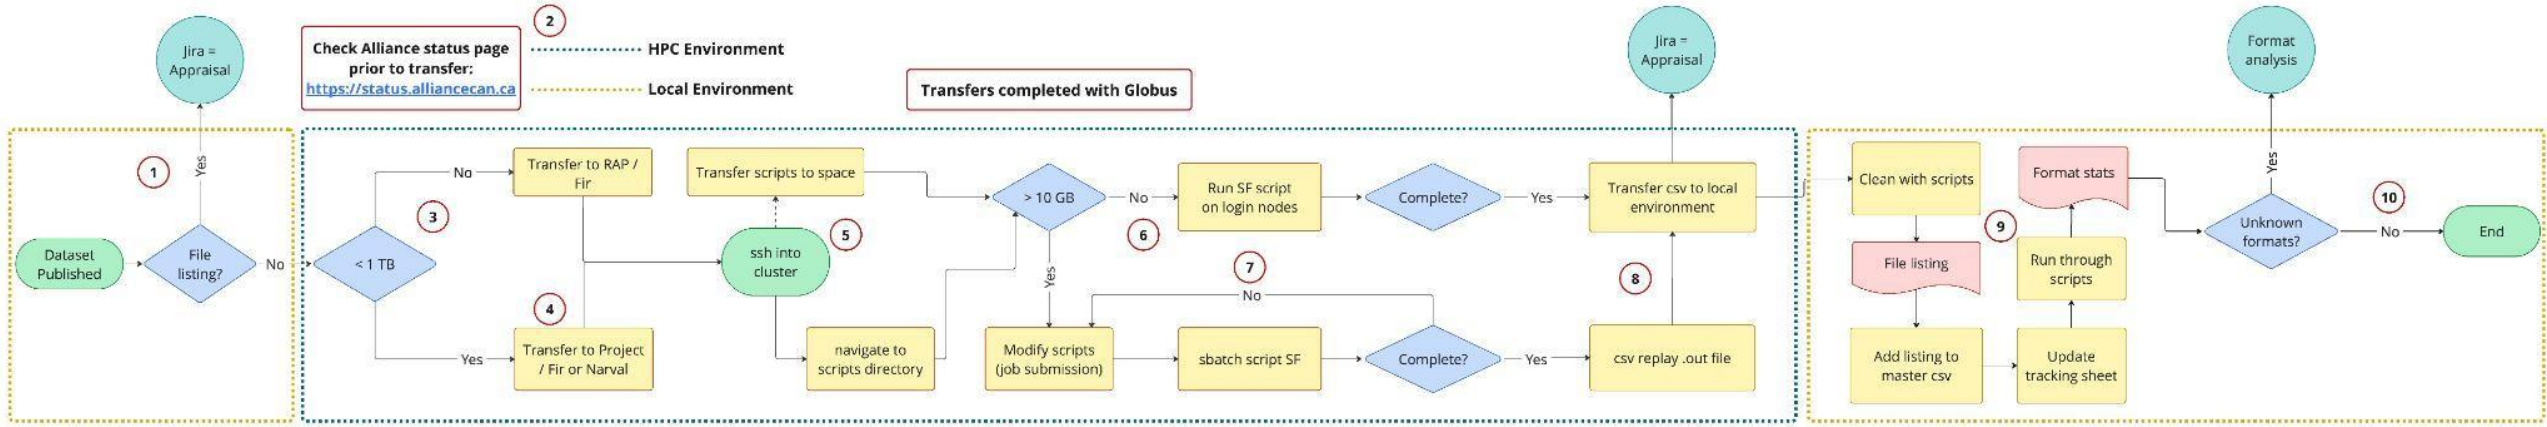
1. Broad workflow overview and explanation
2. Glossary of terms
3. Setup
4. Workflow steps





Draft Workflow.

File Listing Workflow - DRAFT





File Listing.

Publication ID	The internal ID associated with the dataset in FRDR
File Path	The full path to the individual file
Format	The format identification result for the file
Extension	The extension for the file

“...[gathers] information about digital content to support an understanding of the digital content held by an organization, and to facilitate efficient management of its preservation.”




- Digital Preservation Coalition (2024). *Digital Asset Register Toolkit, 1st Edition*, [<http://www.doi.org/10.7207/dartool24-01>]

Glossary.



High-Performance Computing (HPC) / Clusters	High Performance Computing - powerful computing systems composed of many interconnected computers—called nodes—that work together in parallel to solve complex, data-intensive problems much faster than a single computer could.
Login nodes	The primary point of access to the clusters. Login nodes are a shared resource and are only meant for lightweight tasks.
Compute nodes	Responsible for executing the bulk of the compute workloads. The compute nodes are optimized for performance with large CPUs and RAM allocations.
Secure Shell (SSH)	Secure Shell is used to create a secure connection between local and remote servers.
Slurm	Slurm workload manager is a job scheduler used by supercomputers and compute clusters

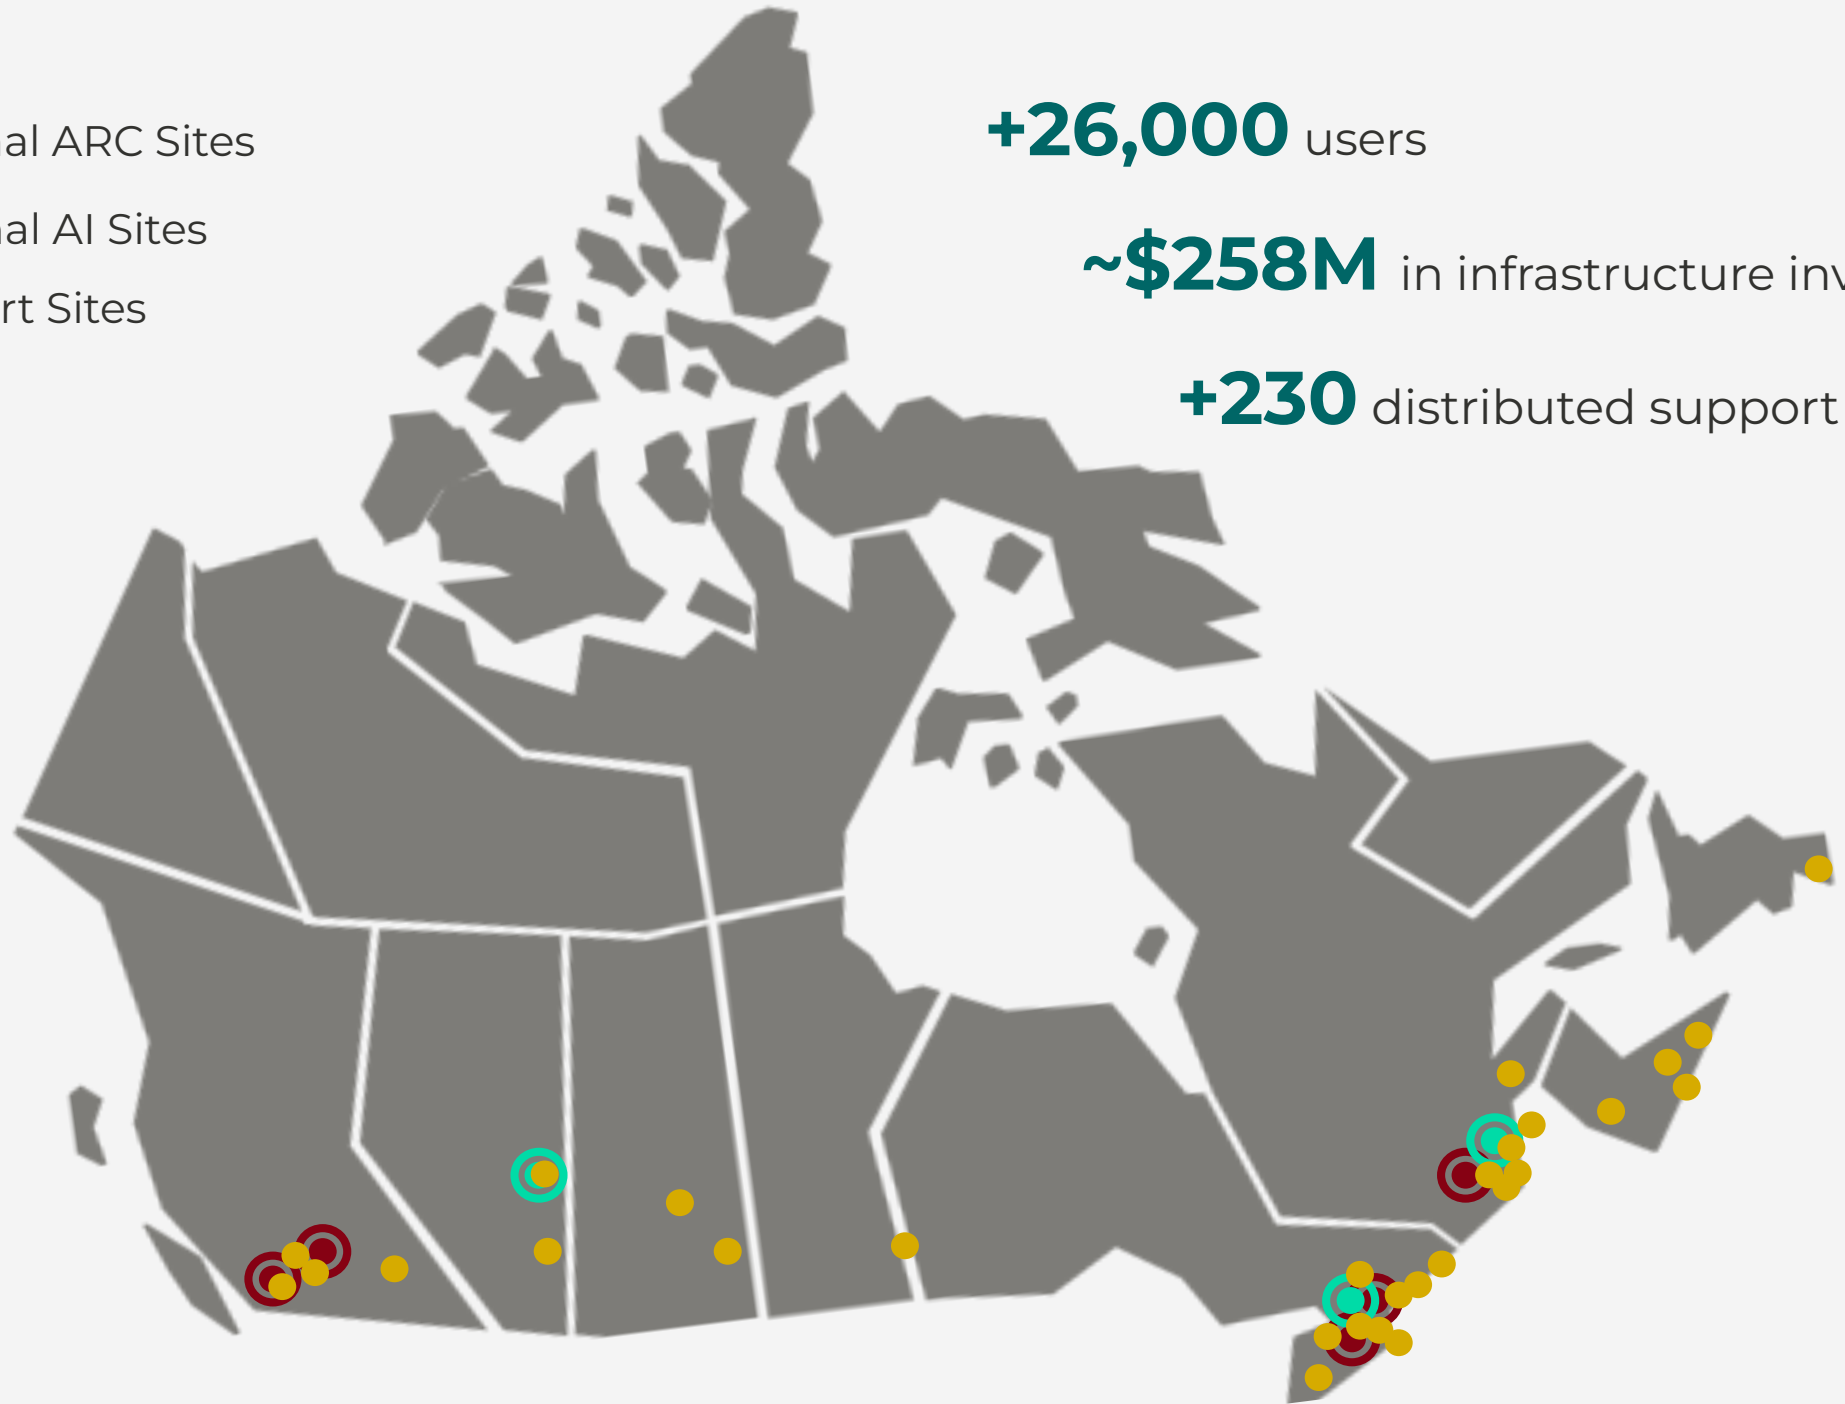
Why use HPC clusters?

-  National ARC Sites
-  National AI Sites
-  Support Sites

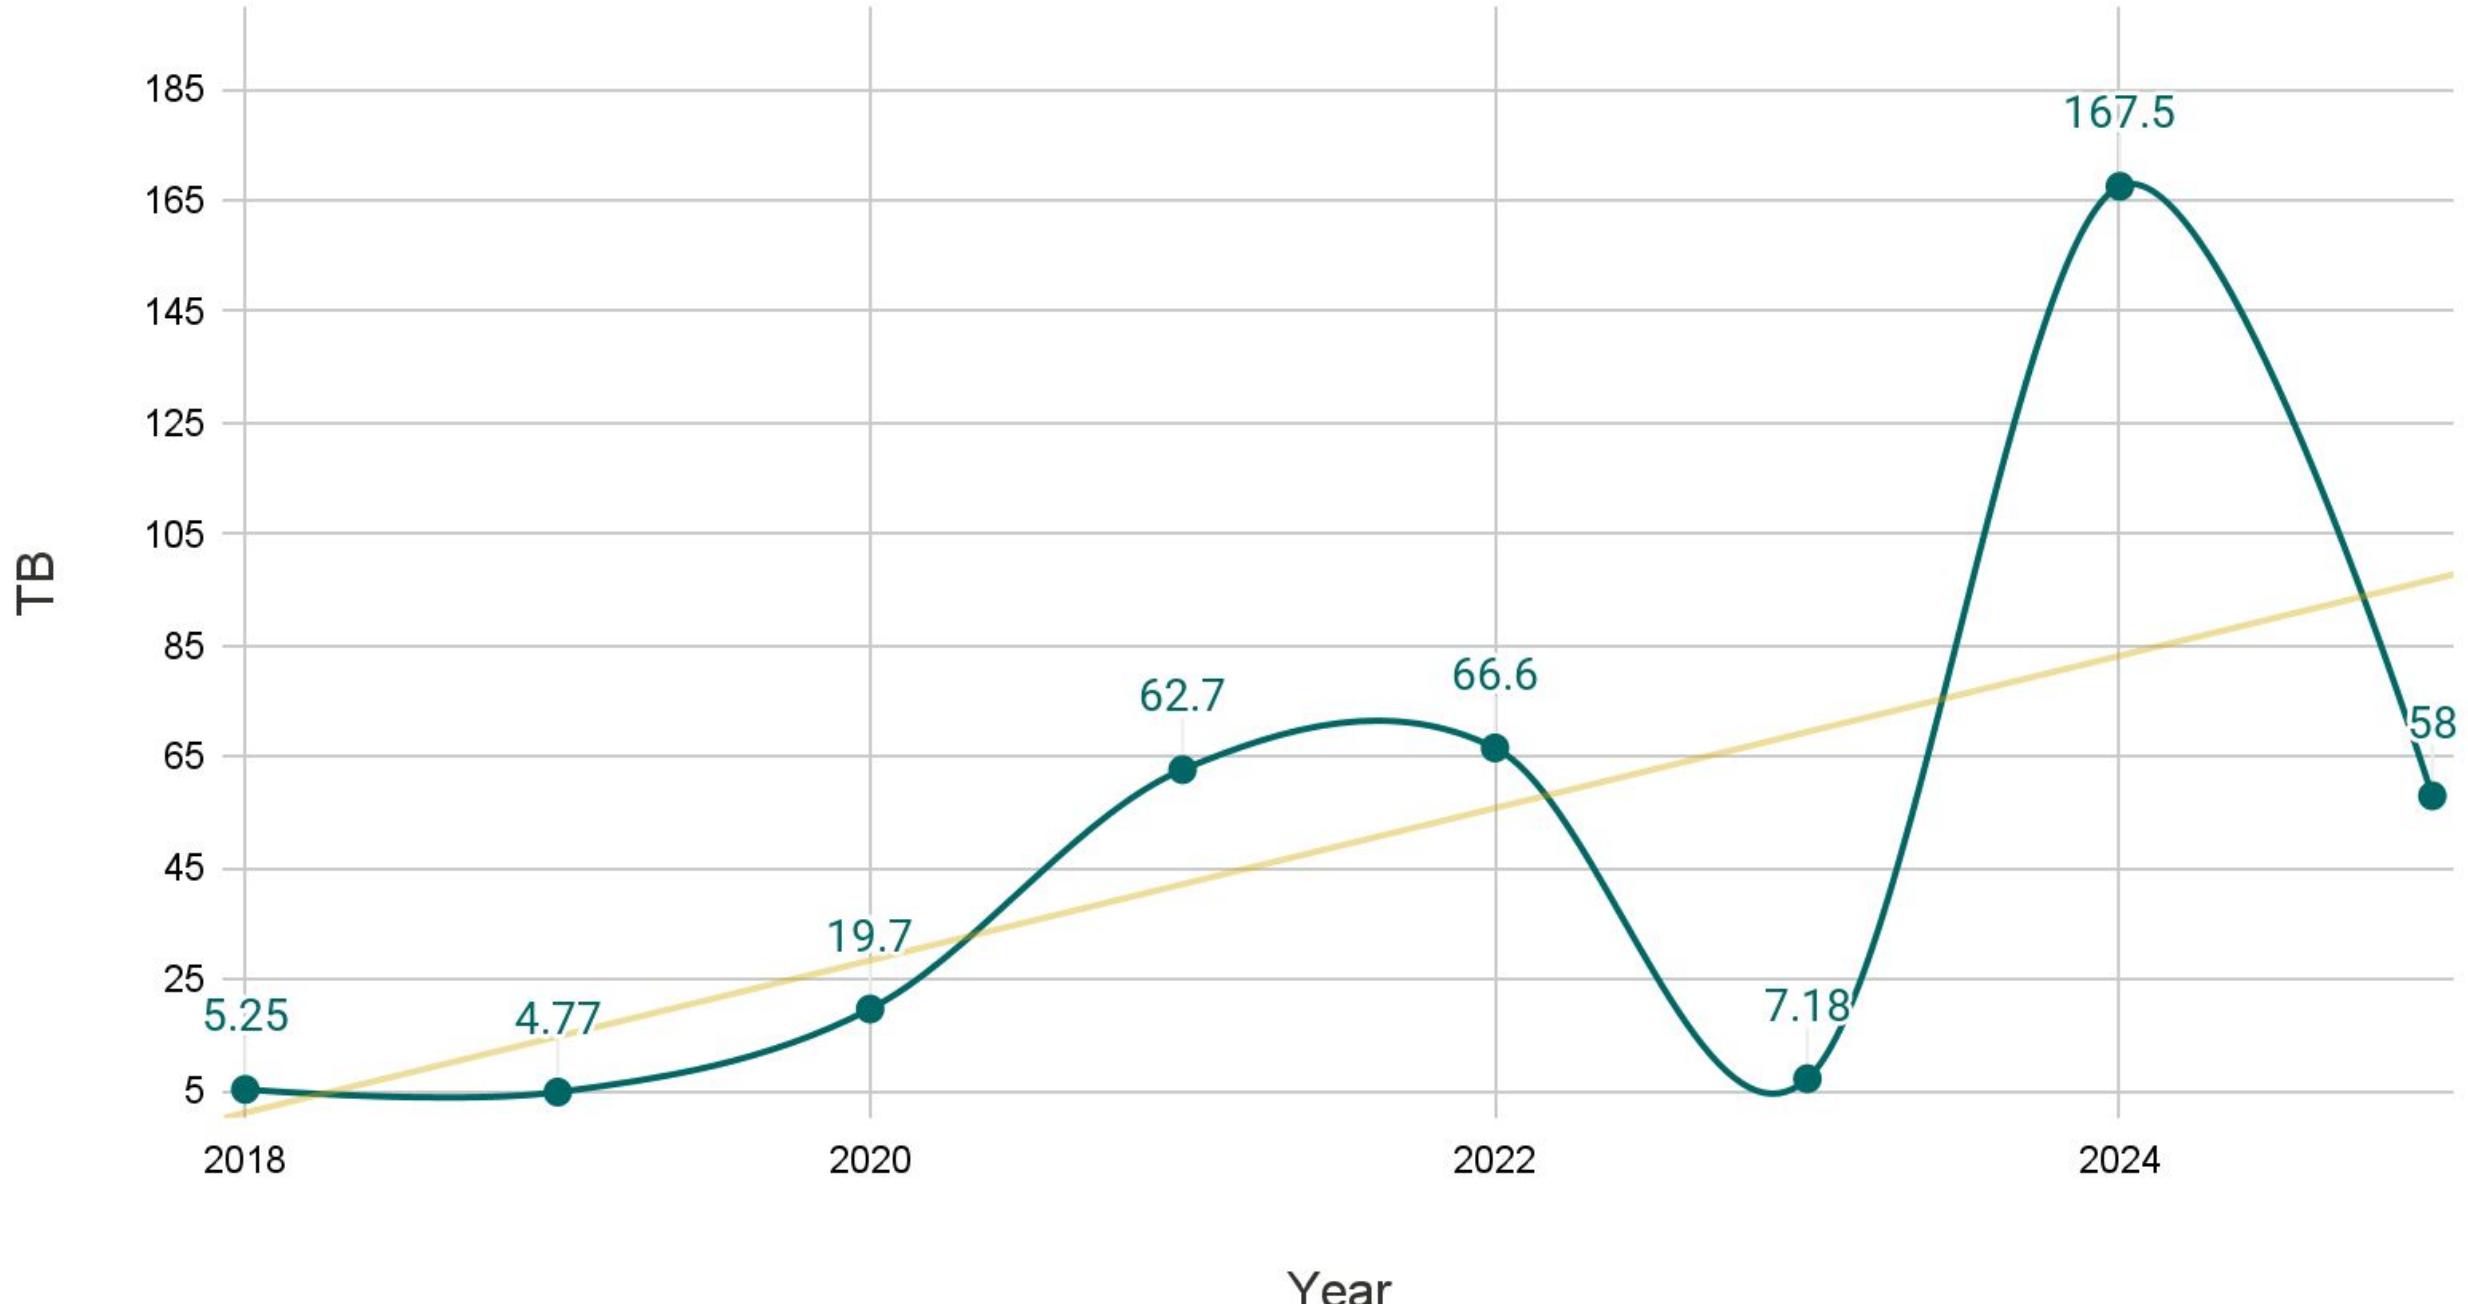
+26,000 users

~\$258M in infrastructure investments

+230 distributed support staff



Growth Over Time (Size)



Setup

Tools / Space / New Skills



Getting Ready.

Setup - Siegfried

- Installation of file identification tool Siegfried
 - Cluster wide
- Discussed with colleagues the possibility of installation
- Other digital preservation professionals can also use the tool



Getting Ready.

Setup - Resource Allocation Project (RAP)

- Designated set of compute and storage resources assigned to a specific research project or group
- 50 TB, 50 cores, 5,000,000 files



Storage Allocations.

Filesystem	Default Quota	Backup	Purge
Nearline*	50 GB / 500,000	Yes	No
Project*	1 TB / 500,000	Yes	No
Scratch*	20 TB / 1,000,000	No	Yes
RAP**	50 TB / 5,000,000	No	No

*Per user

**All users / Resource Allocation Project

[Storage and file management - Alliance Doc](#)



Getting Ready.

Setup – New Skills

- Time was taken to understand how to submit a job on the cluster
- Crafting the script for Siegfried required a deep dive into the documentation and help from the community
- Training coordinator from partner site took time to answer our questions

The Workflow

This is a work in progress

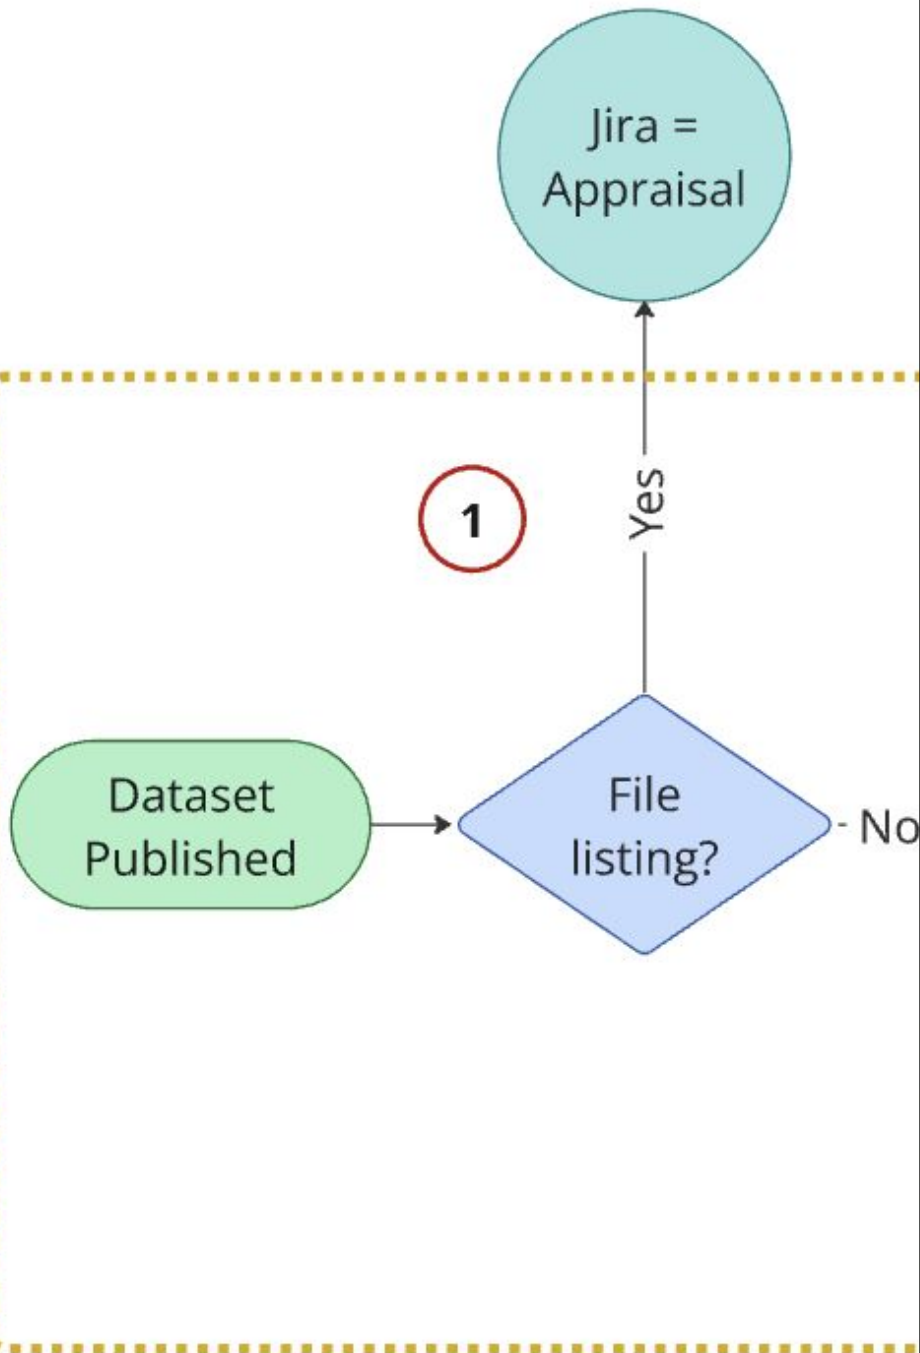


Step 1 - Start.

Dataset is published

Dataset is now part of the preservation queue

1. Assess whether a file listing exists
 - If **yes**, dataset moves to Appraisal Status in the preservation queue
 - If **no**, dataset moves into the File Listing Status in the preservation queue



FPRES.

Jira FPRES Board

- Used to track the status of datasets
- Used to assign tasks to various individuals

The screenshot displays a Jira workspace titled "FRDR Digital Preservation". The navigation bar includes "Summary", "Backlog", "Board" (selected), "All work", and "Shortcuts". A search bar and a filter button are visible. The board is divided into two columns: "FILE LISTING 8 OF 8" and "APPRAISAL 8 OF 8".

FILE LISTING 8 OF 8


- CLONE - Pub 1486 - RADseq data for Atlantic and Indian Yellow-nosed albatross 10.20383/103.01491 [edit]
Jan 13, 2027 [checkbox] 1 [AT]
- CLONE - PUB 1388 - Polyester Microfibers in the Salish Sea 10.20383/103.01393
Jan 19, 2027 [checkbox] 1 [AT]

APPRAISAL 8 OF 8

- CLONE - Pub 1088 - Newfoundland and Labrador Breast Screening (NLBS) Dataset (NL-Breast-Screen) 10.20383/103.01093
Mar 13, 2026 [checkbox] 2 [AT]
- CLONE - Pub 1312 - SOCC - SFU Opinion and Comments Corpus 10.20383/103.01317
Jun 19, 2026 [checkbox] 1 [AT]



Step 2 - Status.

 A service is experiencing an outage. Please see below for additional information.		
✔ Operational ⚠ Available with conditions 🛑 Outage 📅 Scheduled events		
Service	Status	Current incidents
Arbutus ?	🛑	Arbutus CephFS and Object Storage maintenance
Béluga ?	⚠	Decommissioning: 2026 End of Service Timeline - Fermeture : calendrier de fin de service pour 2026
CCDB ?	✔	
Cedar ?	🛑	End of Service, September 12, 2025 - Fin de service, 12 septembre 2025
Cedar Cloud	✔	
DMP Assistant - L'Assistant PGD ?	✔	
FRDR-DFDR ?	✔	
Fir ?	✔	
Globus Service	✔	

Check Alliance status page prior to transfer:
<https://status.alliancecan.ca>

Status Check

1. Review status of the Alliance systems, primarily the HPC Clusters

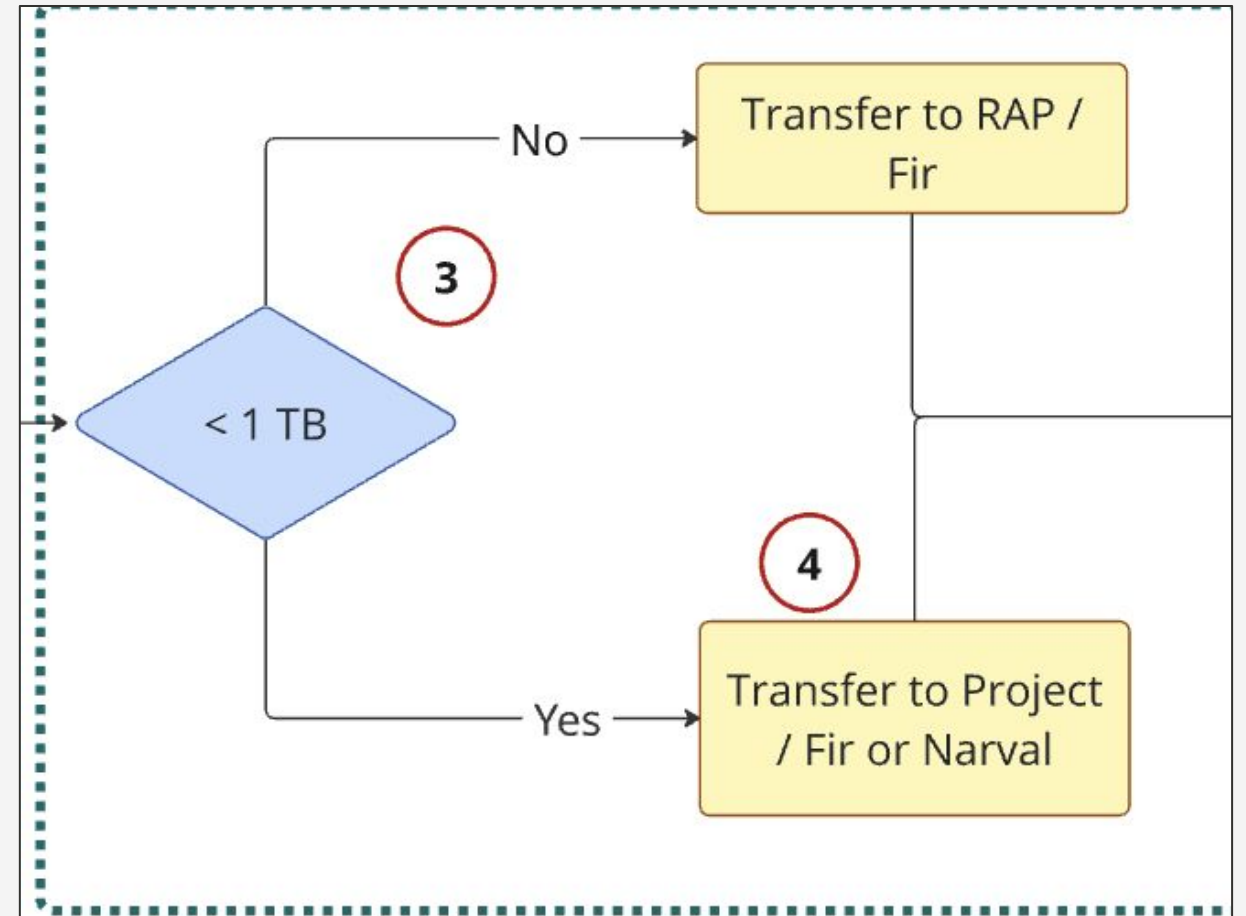


Step 3 - Choose a location.

Considerations

1. Is the dataset smaller than 1 TB?
 - If **yes**, transfer dataset to Project space on Fir or Narval
 - If **no**, transfer dataset to Fir Resource Allocated Project (RAP)

*Record space in Jira task





HPC Clusters.

Fir

- Simon Fraser University in Burnaby, BC
- Resource Allocation Project location
- Deployed 2025

Narval (Narwhal)

- École de technologie supérieur in Montreal, QC
- Maintained by Calcul Québec
- Deployed 2021

Transfer to RAP /
Fir

Transfer to Project
/ Fir or Narval

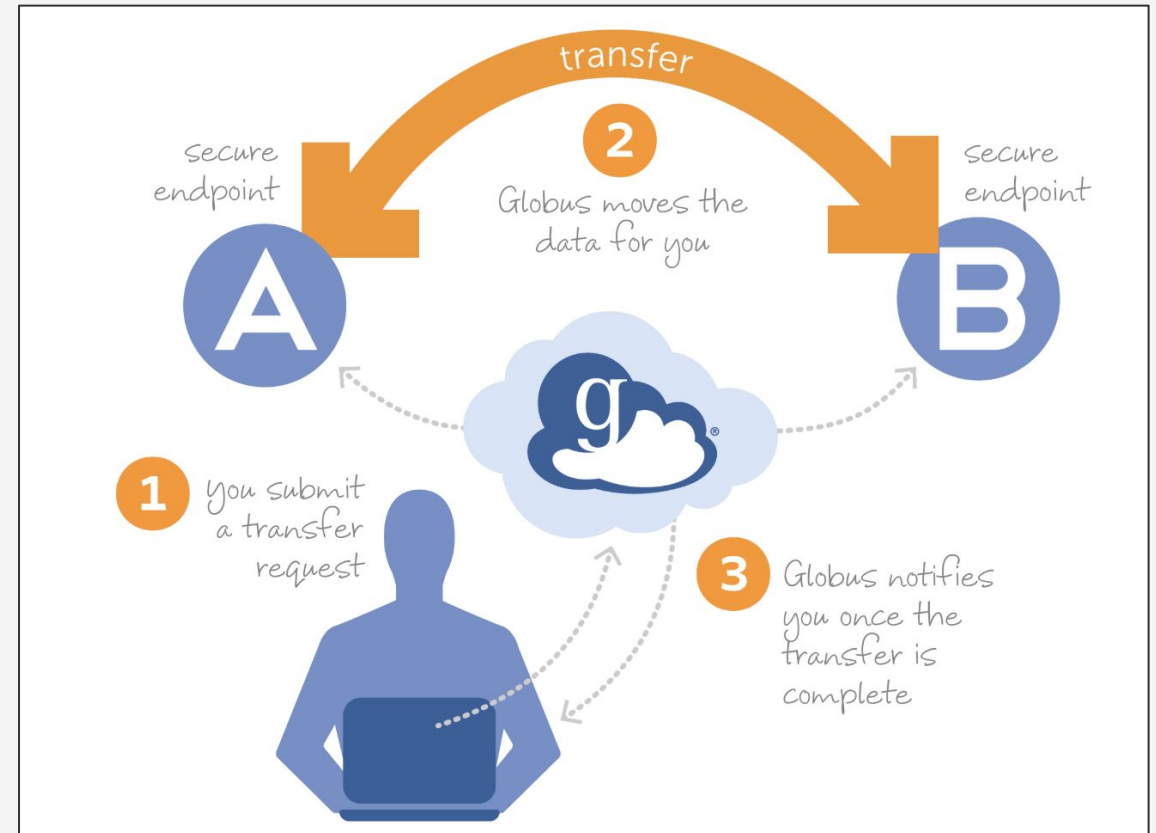
Step 4 - Transfer.

Globus

- Developed by University of Chicago
- Reliably and securely move, share and discover data
- Verifies data integrity at end of transfer

■ ■ ■ ■ ■ HPC

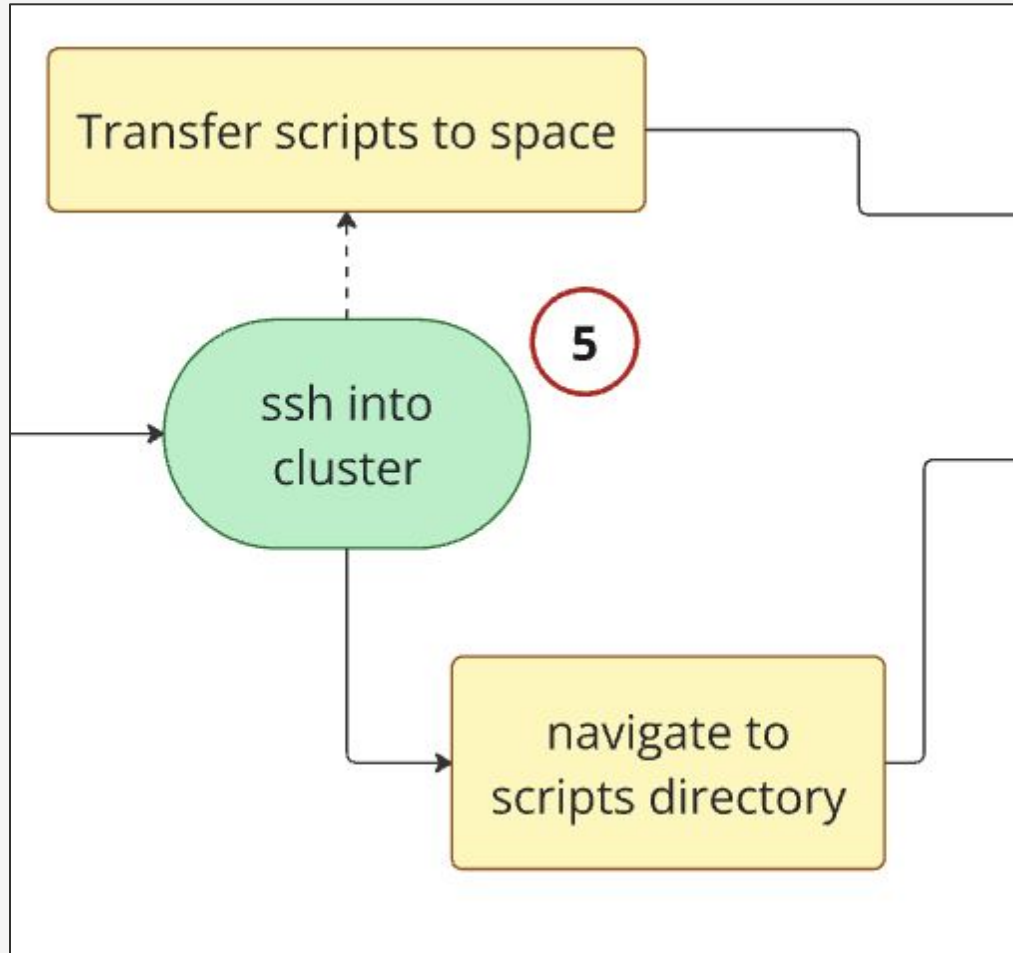
Transfers completed with Globus



<https://www.globus.org/data-transfer>



Step 5 - SSH.



Access Clusters

1. Open terminal
2. ssh into cluster
3. `username@fir.alliancecan.ca`
4. Navigate to directory

Transfer

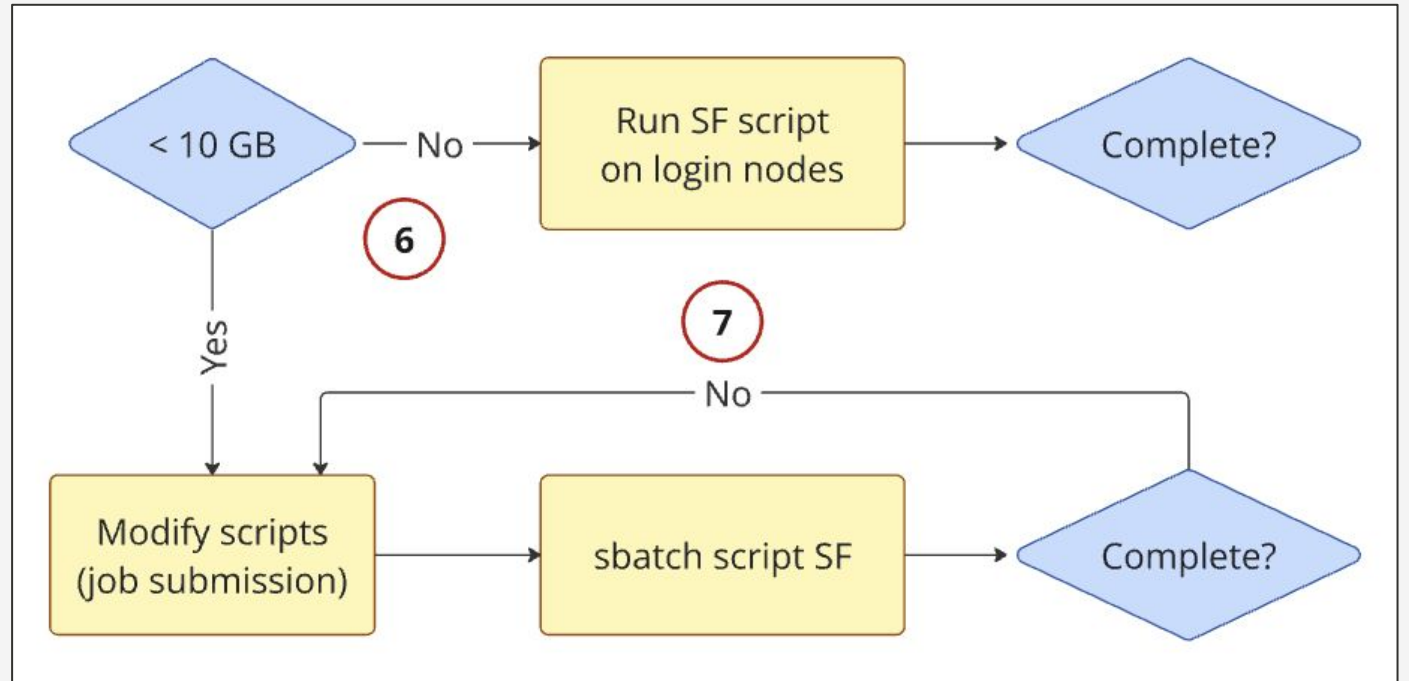
- Scripts needed are transferred into space



Step 6 - Running Scripts.

Login Node

- Login Node
 - Main gateway to clusters
 - Lightweight tasks
 - Under 10GB
 - Scripts may need some modification





Example: Running a script

1. **nano** <file_name>.sh
2. Modify the script as needed
3. Save the file
4. **chmod +x** <file_name>.sh
5. **./**<file_name>.sh

```
UW PICO 5.09 File: 02-sf_file_id.sh Modified
#!/bin/bash
#SBATCH --account=xxx-xxxxxx
#SBATCH --ntasks=1
#SBATCH --nodes=1
#SBATCH --mem=4000
#SBATCH --time=0-00:05

#SBATCH --mail-type=ALL
#SBATCH --mail-user=xxxxx.xxxxxxxx@alliancecan.ca
#SBATCH -o output-file-%j.out
#SBATCH --error=slurm-%j.err
#SBATCH --job-name=sf-file-identification

#load module
module load siegfried

#user enters file path
DIRPATH="$1"

^G Get Help ^O WriteOut ^R Read File ^Y Prev Pg ^K Cut Text ^C Cur Pos
^X Exit ^J Justify ^W Where is ^V Next Pg ^U UnCut Text ^T To Spell
```



Step 7 - Submitting Jobs.

Compute Nodes

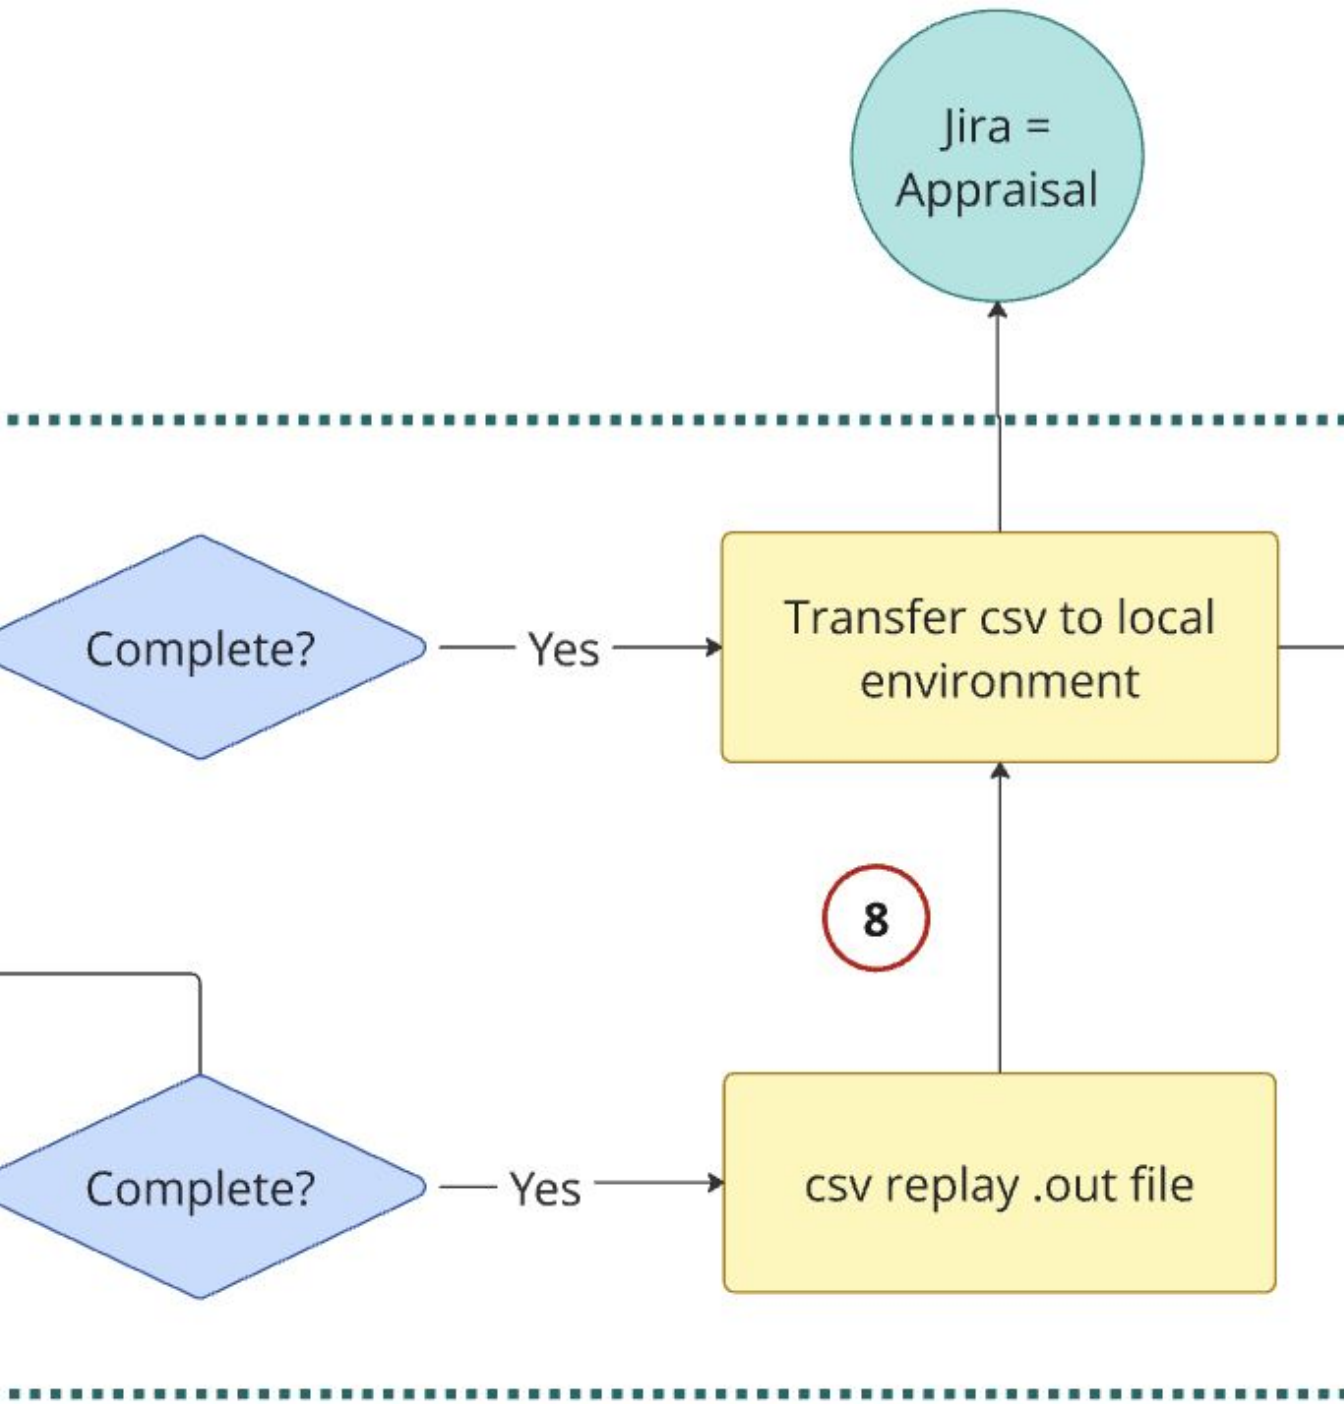
- Uses **Slurm** to schedule submissions to the clusters
- Use **sbatch** for job submission
- Need to specify processing memory and time needed
- **.out** files automatically generated
- If job fails - needs to be resubmitted

```
#!/bin/bash
#SBATCH --account=def-xxx - specify account
#SBATCH --ntasks=1
#SBATCH --nodes=1
#SBATCH --mem=4000 - specify memory
#SBATCH --time=0-00:05 - specify time

#SBATCH --mail-type=ALL
#SBATCH --mail-user=xxxxxxx@xxxxxx.ca
#SBATCH -o output-file-%j.out
#SBATCH --error=slurm-%j.err
#SBATCH --job-name=sf-file-identification
```



Step 8 - CSV.



Login Nodes

1. Transfer csv to local environment

Job submission

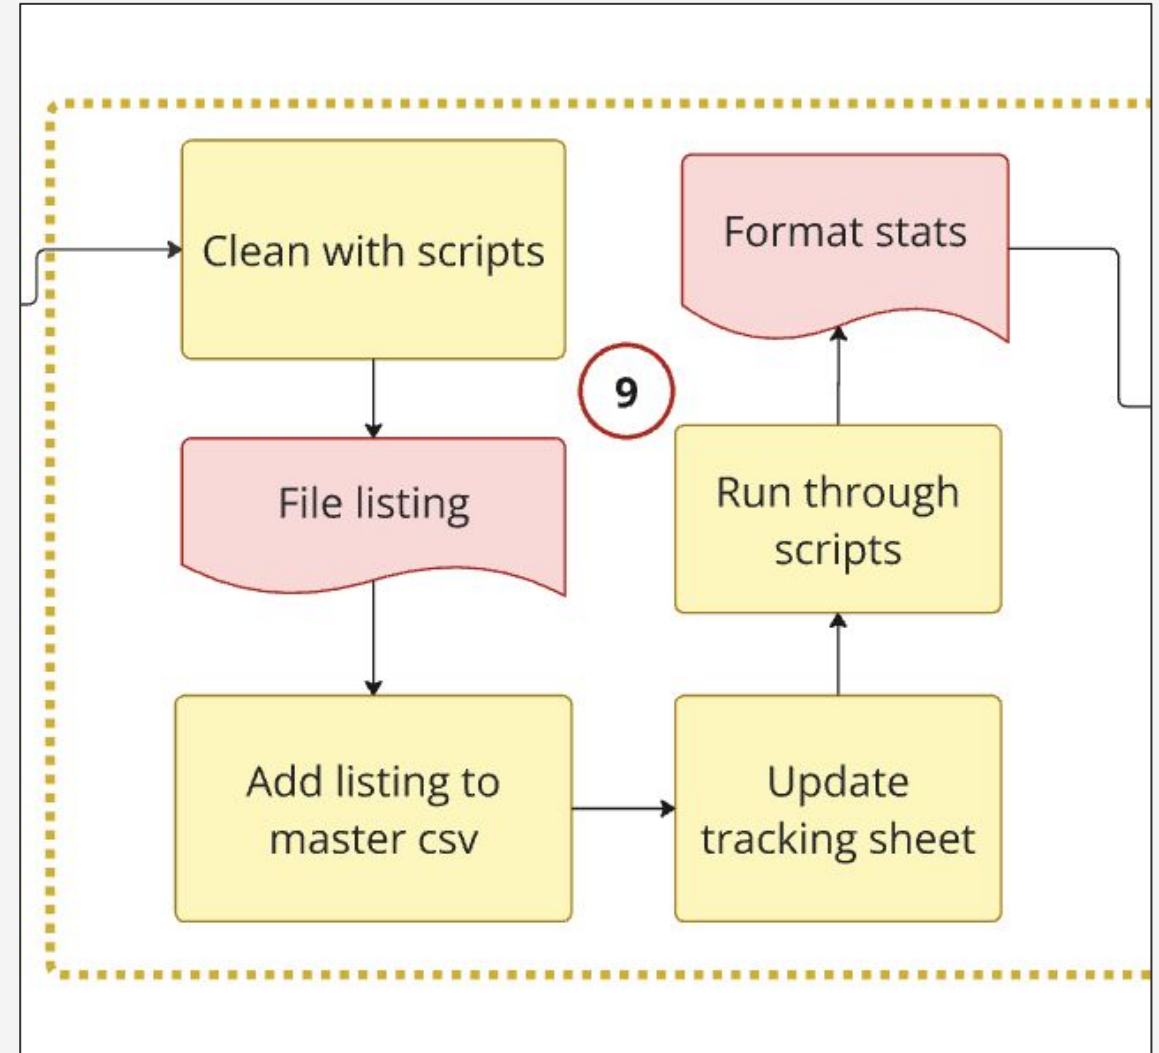
1. Use `sf -csv -replay <output-file.out> pub_xxx.csv`
2. Transfer to local environment
3. Dataset moves to appraisal



Step 9 - CSV Cleanup.

Local Environment

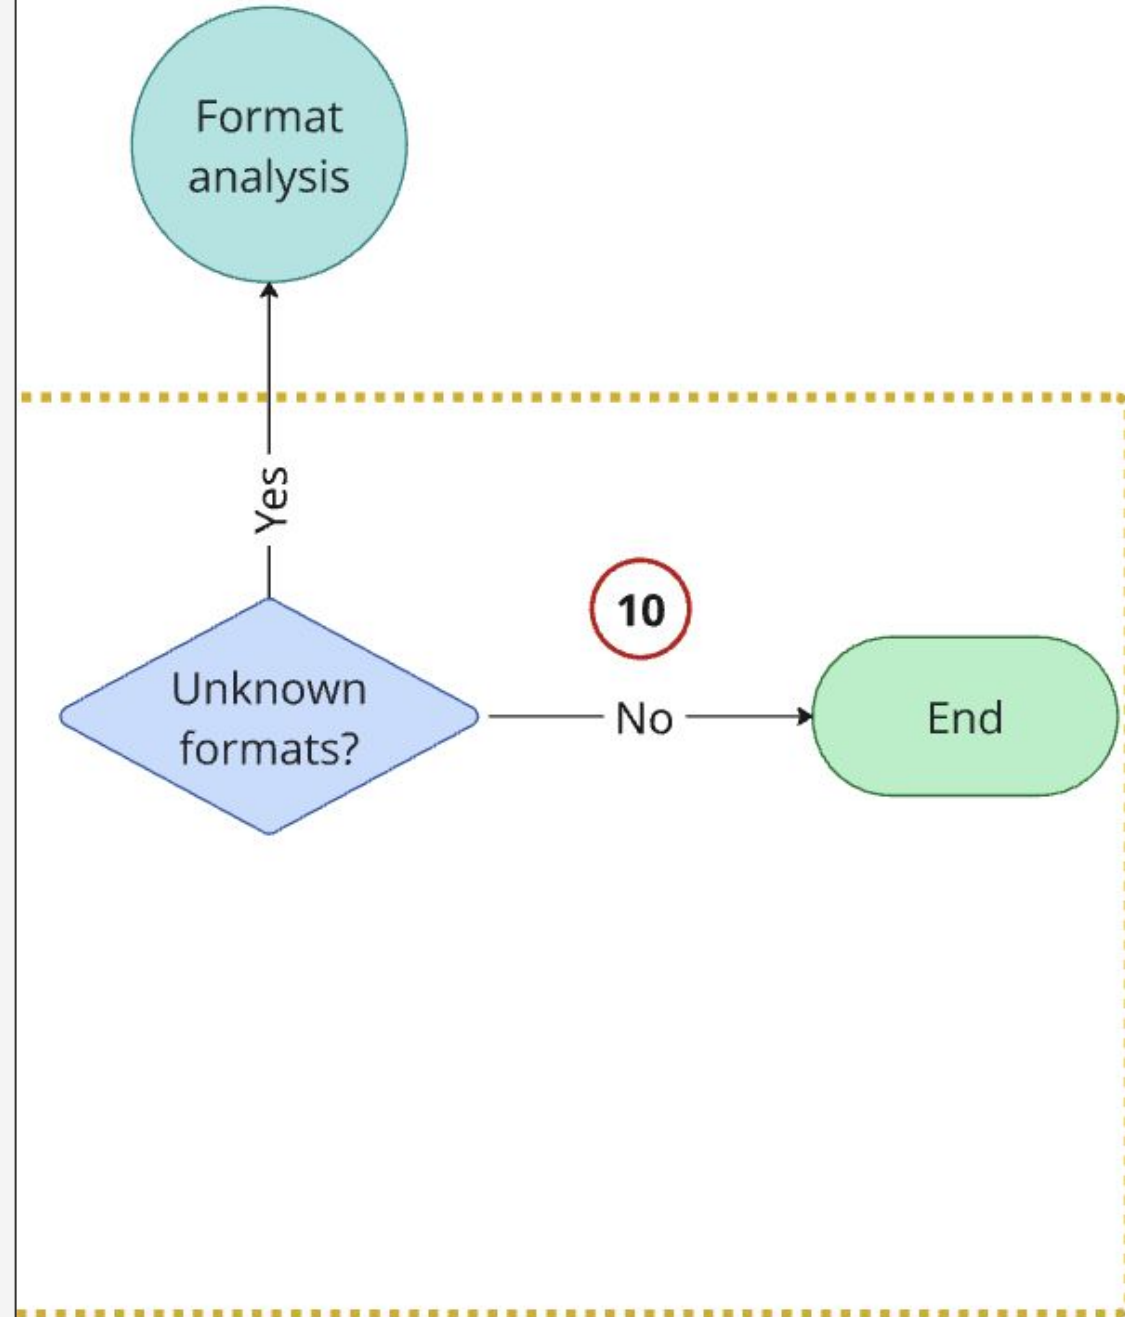
1. Run CSVs through scripts
2. Generate final file listing
3. Add the listing to the full CSV list
4. Update tracking sheet
5. Run scripts
6. Generate format stats

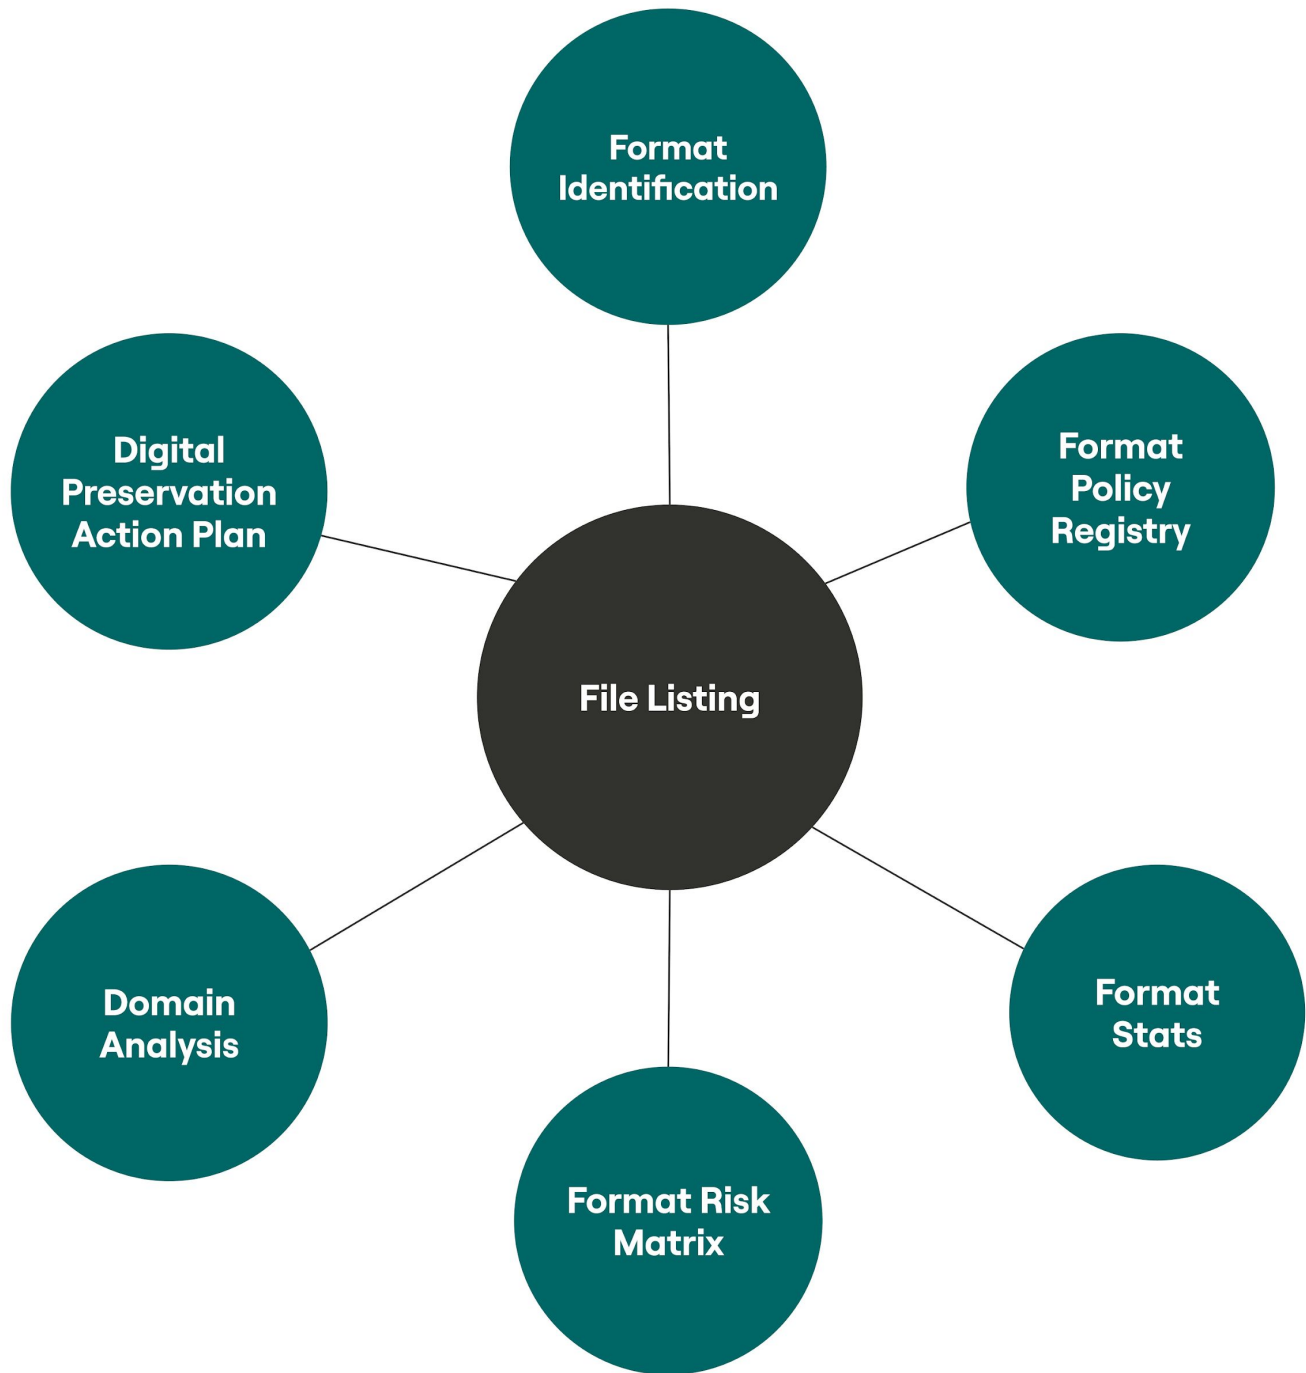


Step 10 - End.

Unknown Formats

- If **yes**, format analysis process
- If **no**, End





Final Listing.



- Format Risk Matrix
- Format Stats
- Format Identification
- Domain Analysis
- Format Policy Registry
- Digital Preservation Action Plan

Next Steps



Job Array.

- Allows users to submit similar jobs at once
- Potentially save time - jobs can run simultaneously
- Less tasks



Scripts on Clusters.

- Use Jupyter Lab to run cleanup scripts
- Keep everything in one place
- Modify Siegfried script to include replay option?

Thank You!



	A	B	C	D
1	publication_id	file_name	format_string	extension
2	864	/globusdata/5/published/publication_864/submitted_data/videos/trial_videos/6QP11_53_7-21-2021-214444.mp4	MPEG-4 Media File	mp4
3	864	/globusdata/5/published/publication_864/submitted_data/videos/trial_videos/6QP14_54_9-2-2021-183856.mp4	MPEG-4 Media File	mp4
4	864	/globusdata/5/published/publication_864/submitted_data/videos/trial_videos/6QP20_53_1-6-2022-164810.mp4	MPEG-4 Media File	mp4
5	864	/globusdata/5/published/publication_864/submitted_data/videos/trial_videos/6QP09_53_7-7-2021-54923.mp4	MPEG-4 Media File	mp4
6	864	/globusdata/5/published/publication_864/submitted_data/videos/trial_videos/6QP11_03_7-14-2021-42548.mp4	MPEG-4 Media File	mp4
7	864	/globusdata/5/published/publication_864/submitted_data/videos/trial_videos/6QP15_01_8-13-2021-32252.mp4	MPEG-4 Media File	mp4
8	864	/globusdata/5/published/publication_864/submitted_data/videos/trial_videos/6QP21_03_10-13-2021-195528.mp4	MPEG-4 Media File	mp4
9	864	/globusdata/5/published/publication_864/submitted_data/videos/trial_videos/6QP23_02_9-20-2021-224335.mp4	MPEG-4 Media File	mp4
10	864	/globusdata/5/published/publication_864/submitted_data/videos/trial_videos/6QP06_53_8-3-2021-53048.mp4	MPEG-4 Media File	mp4



**Digital Research
Alliance** of Canada

**Alliance de recherche
numérique** du Canada



alliancecan.ca



[/AllianceCan](https://www.linkedin.com/company/alliancecan)



preservation@frdr-dfdr.ca

