



Vlaanderen  
is supercomputing

# VSC Cloud and Data platform

VSC User's Day, June 4 2019

VLAAMS  
SUPERCOMPUTER  
CENTRUM

*Innovative Computing  
for A Smarter Flanders*

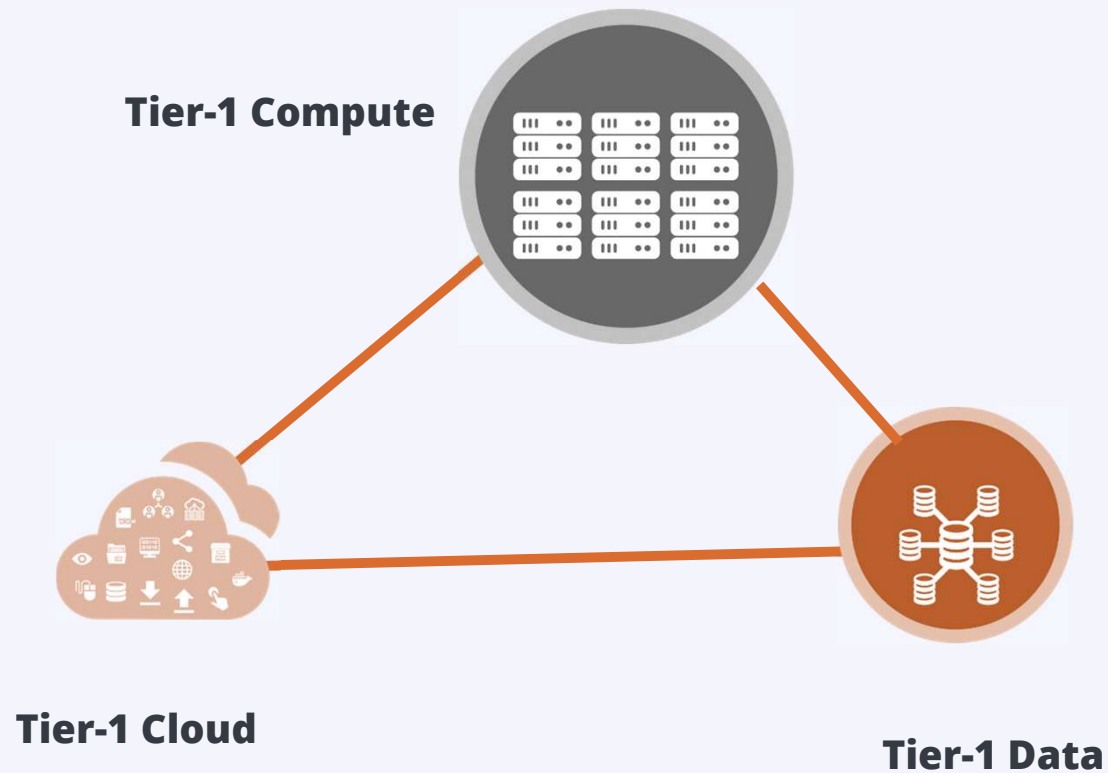
[vscentrum.be](http://vscentrum.be)



Tier-1 SaaS status update

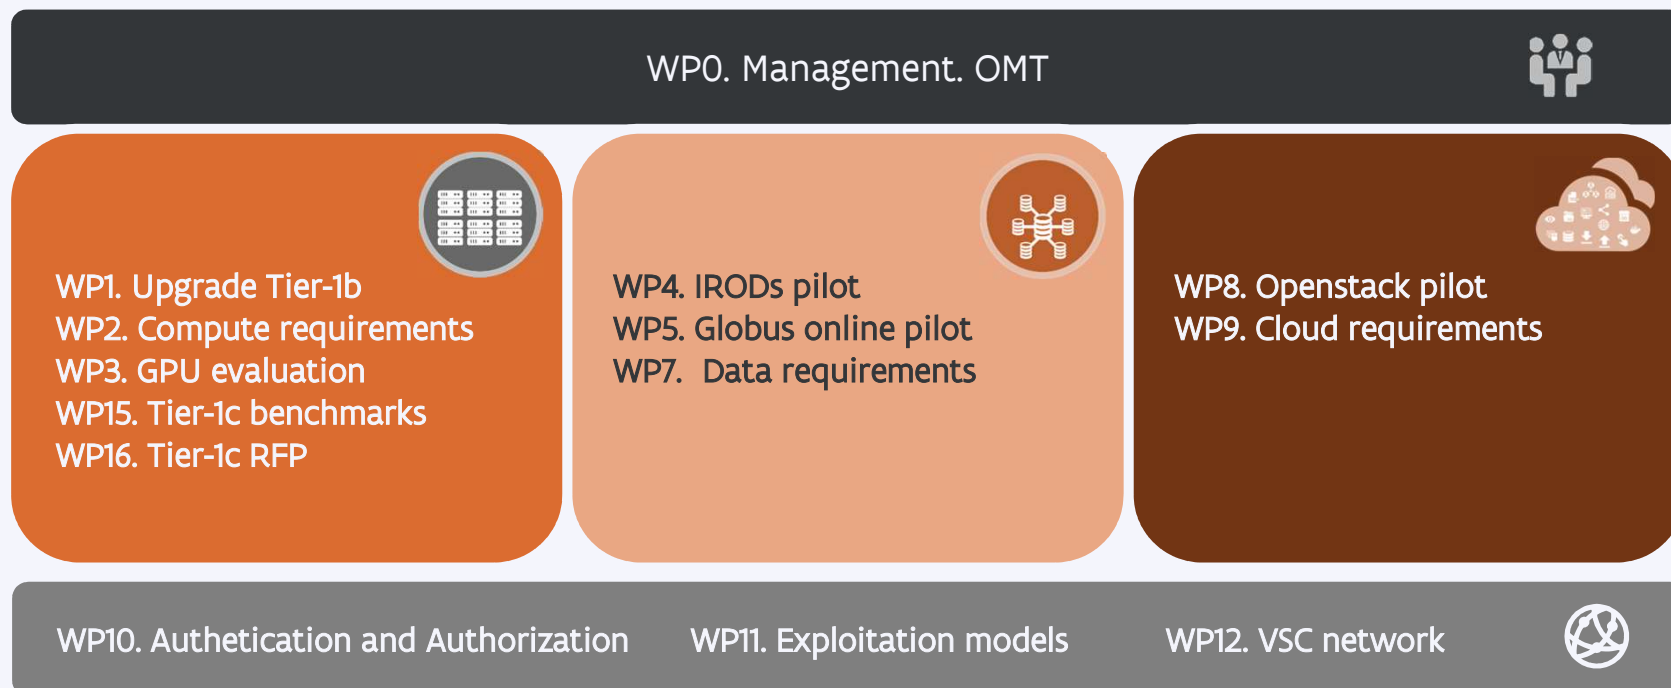
A year ago ...

# TIER-1 Supercomputing as a Service



**Highly integrated**  
**Flexible**  
**Easy and intuitive**  
**Secure**  
**Performant**  
**Well-documented**

# 2018: detailed action plan



16 WorkPackages: Concrete objectives

Mixed VSC Workgroups: Promote collaboration and Share of knowledge

# 2018: realizations

		jan	feb	mar	apr	mei	jun	jul	aug	sep	okt	nov	dec
COMPUTE	WP1				Datacenter aanpassingen							Configuratie en bestelling	
	WP2					GPU evaluatie							
	WP3				Enquete voorbereiden		Light Survey				Detailed survey		
DATA	WP4				IRODS piloot setup								
	WP5					Globus piloot setup							
	WP7				Enquete voorbereiden		Light Survey				Detailed survey		
CLOUD	WP8				Openstack piloot setup					Gesloten piloot			
	WP9				Enquete voorbereiden		Light Survey				Detailed survey		
INTEGRATIE	WP12									Netwerk testen en performantie			
	WP13									VSC Datacenter inventaris			
	WP14											Aankoop overeenkomst	

# 2018: milestones

- Extension BrENIAC
- Preparation Tier-1c procurement
- Users surveys:
  - light and detailed
- Openstack pilot
- IRODs and Globus pilots

# Action plan 2019: summary

		jan	feb	mar	apr	mei	jun	jul	aug	sep	okt	nov	dec
COMPUTE	WP1	DC aanpassingen	HW inst	SW inst	Bench	Pilot			Productie				
	WP2	workshops voorbereiden				workshops uitvoeren			workshops rapport	interviews voorbereiden		interviews	
	WP3	GPU software pakketten test				GPU rapport							
	WP16	RFP voorbereiden			Aaankoop procedure								
	WP15	Benchmarks studie en selectie											
DATA	WP4	IRODS test setup					Gesloten pilot			Open Pilot			
	WP5	Globustest setup				Gesloten pilot			Open Pilot				
	WP7	workshops voorbereiden				workshops uitvoeren			workshops rapport	interviews voorbereiden		interviews	
CLOUD	WP8	Gesloten pilot											
	WP9	workshops voorbereiden					workshops uitvoeren			workshops rapport	interviews voorbereiden		interviews
INTEGRATIE	WP10	Authorisatie en authenticatie											
	WP11	Exploitatie modellen											
	WP12	VSC netwerk testen en configuratie											

## Milestones:

- Procurement Tier-1c
- Close pilot Data platform
- Cloud in pre-production
- Exploitation models for Data and Cloud

(1) Cloud pre-production with some limitations: no SLAs, limited services available. System will be running in old hardware and with limited storage

# Tier-1 requirements: User's workshops

Light and Detailed surveys did not provide conclusive information  
A new approach will be used to gather the User's requirements:

Workshops for the following scientific domains:

- Computational Chemistry/Physics/Material Sciences
  - Life sciences/Medical Sciences/Hospitals/bioinformatics/Biology
  - Engineering/Technology
  - Social Sciences/Humanities
  - Earth sciences: geology, climate, weather
- 
- 10-15 relevant user's selected from the VSC sites
  - Compute, Data and Cloud discussed in a full day

# Workshop format

Participants will prepare upfront a short presentation of their workflow:

- Which input and output data do they have?
- Which are the processing steps and where and how they take place?
- Are there intermediate steps?
- Who need access to which data and when?

Based on 2-3 selected workflows per session will analyze the requirements for the 3 platforms

On each block questions as the ones in the surveys will be asked. Participants will write their answers in post-its and then cluster the requirements

Eventually a final moment to define features priorities

## Proposed session agenda


9h30	Introduction
9h45	Presentations participants
4 blocks of 1h (compute/data/cloud/integratie)	
11h15	compute
12h15	lunch/networking
12h45	data
13h45	cloud



Tier-1 SaaS: Compute component




580 nodes  
Intel E5-2680v4 (Broadwell)  
28 cores/node  
128/256 GiB RAM  
EDR Infiniband



408 nodes  
Intel Gold 6132 (Skylake)  
28 cores/node  
192 GiB RAM –  
EDR Infiniband

988 nodes  
27,664 cores  
Peak: 1,46 PFlop/s (DP)  
Linpack: 1 PFlop/s (DP)

IB EDR network



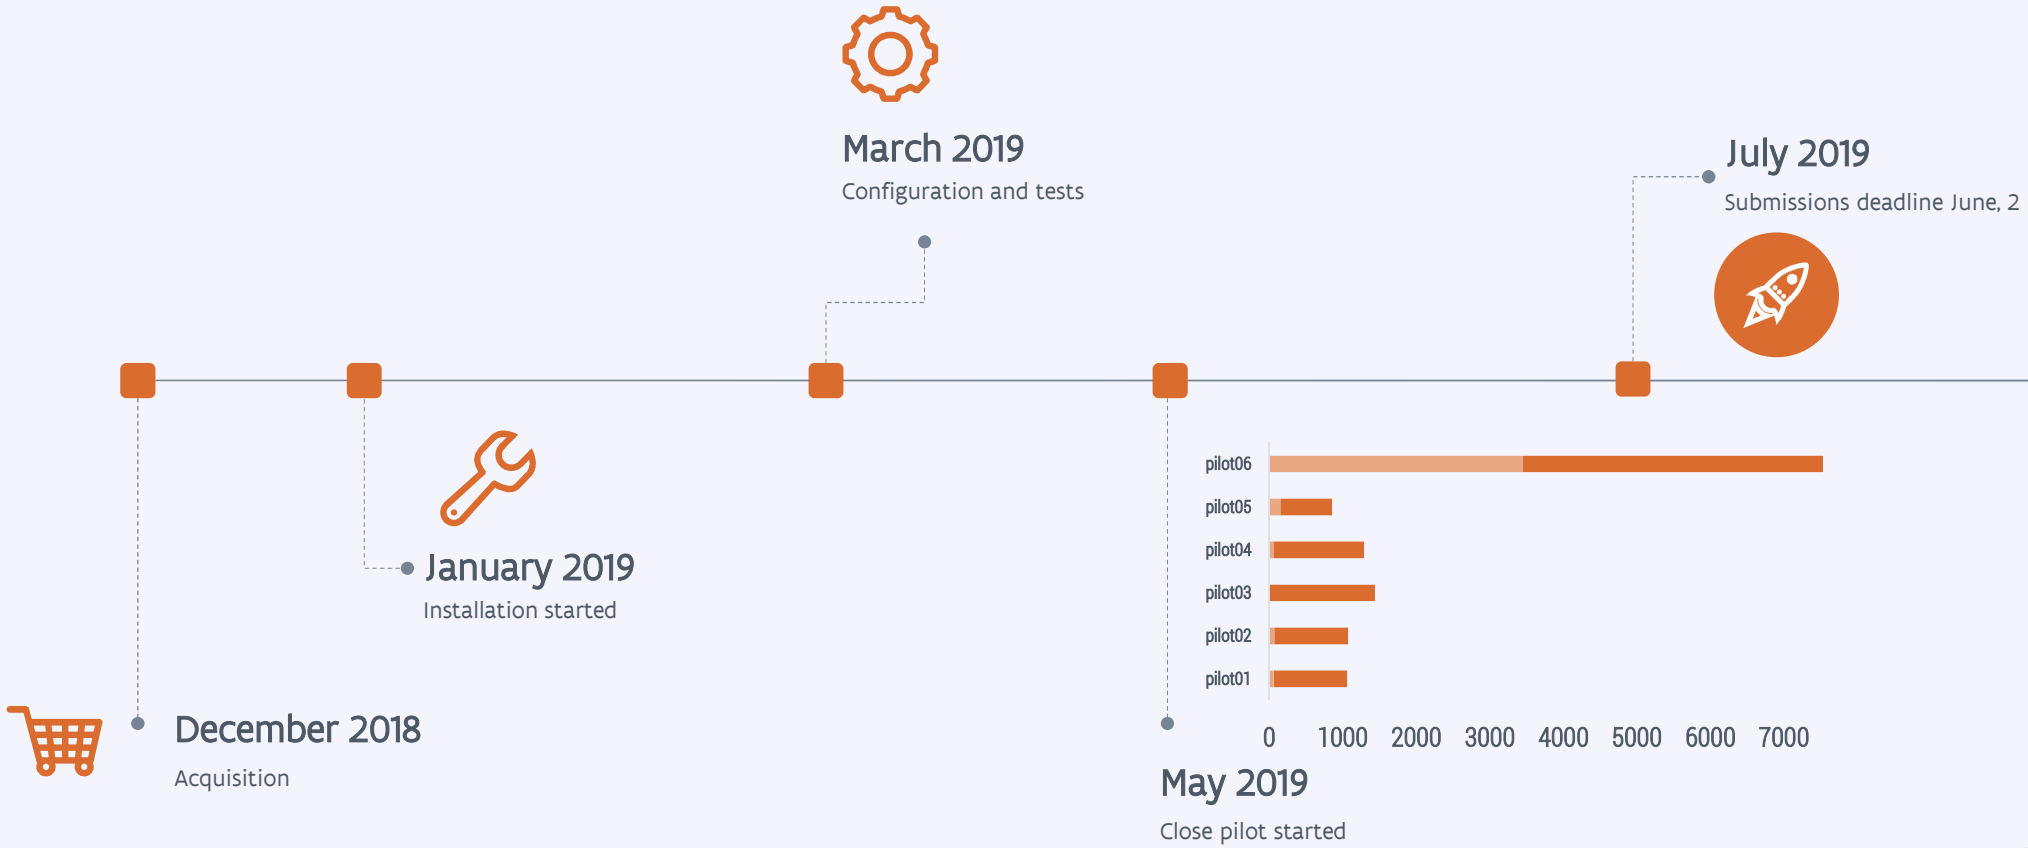
GPFS – 1,3 PiB  
BW – 20 GB/s

Eth network



NFS

# BrENIAC extension



# Closed pilot: first experiences

M87\*

April 11, 20

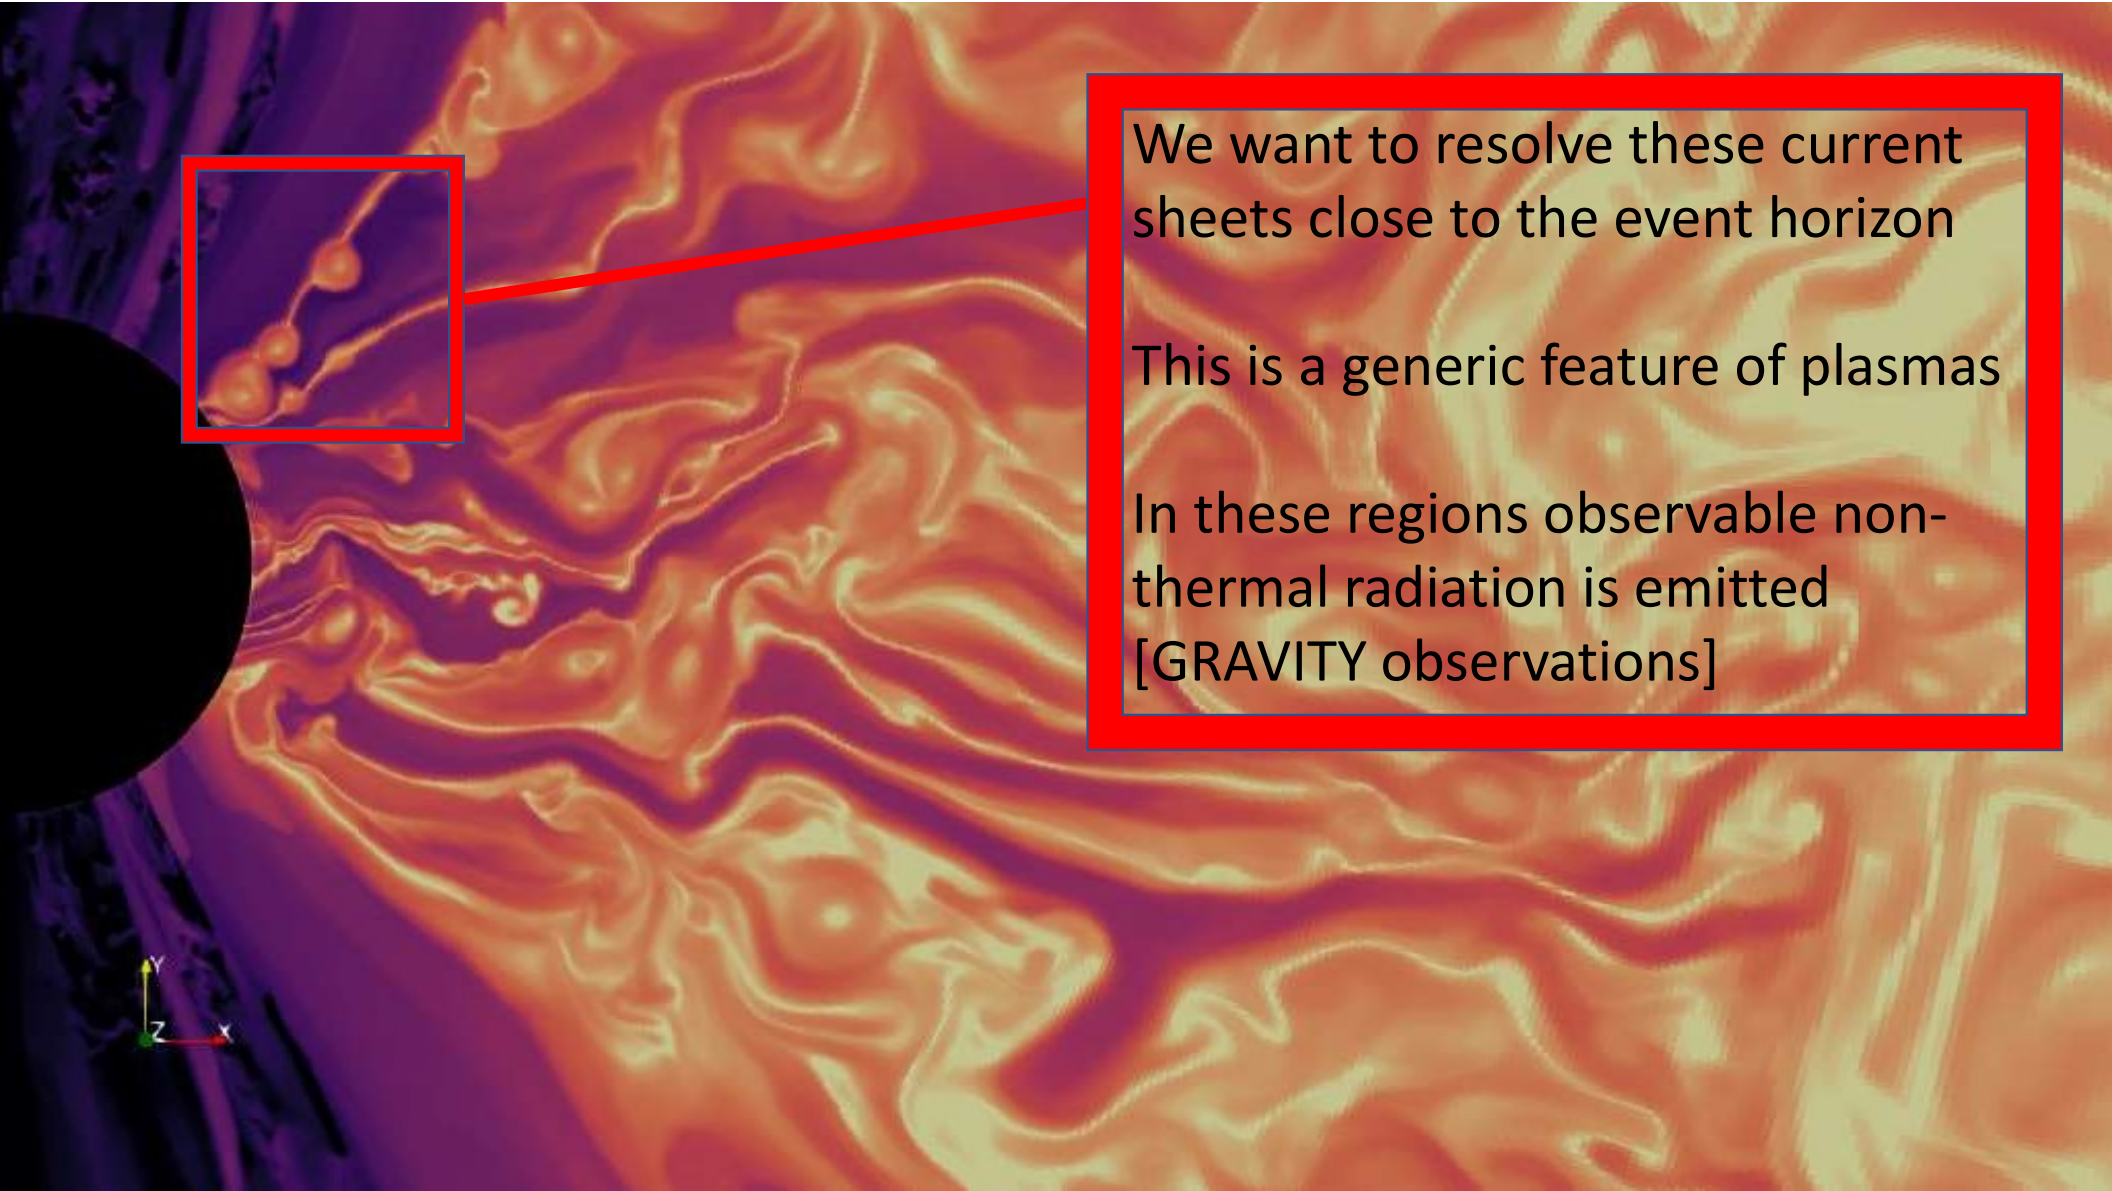
# Black Hole Astrophysics on BrENIAC

Bart Ripperda

in collaboration with Rony Keppens, Thomas Hertog,  
Bert Vercnocke, Fabio Bacchini, Jannis Teunissen,  
Sasha Philippov, Jordy Davelaar, Oliver Porth, Hector  
Olivares, The EHT Collaboration

50  $\mu\text{as}$


The computational resources and services used in  
this work were provided by the VSC



We want to resolve these current sheets close to the event horizon


This is a generic feature of plasmas

In these regions observable non-thermal radiation is emitted  
[GRAVITY observations]



6160 cores (220 nodes) per job  
2,6 days walltime per 1 job

- In accretion disks and jets as observed by the Event Horizon Telescope magnetic flux tubes can kink and produce current sheets
- Current sheets are characterised by electric field parallel to the magnetic field  $\vec{E} \cdot \vec{B} < 0$  (blue)
- In ideal parts  $\vec{E} \cdot \vec{B} \geq 0$  (white/red)  $\rightarrow$  no extreme resolution needed
- Base level 128x128x64 with 3 AMR levels (up to 512x512x256) to resolve linear and ideal physics
- 4 additional AMR levels (up to 8192x8192x4096) to resolve current sheets  $\rightarrow$  activated by  $\vec{E} \cdot \vec{B} \neq 0$






Tier-1 SaaS: Cloud component

TEAM WORK

# Cloud workgroup.



System admim

Álvaro Simón Garcia



User Support

Thomas Danckaert



System admin

Stephane Gerard



User Support

Alexander Vapirev

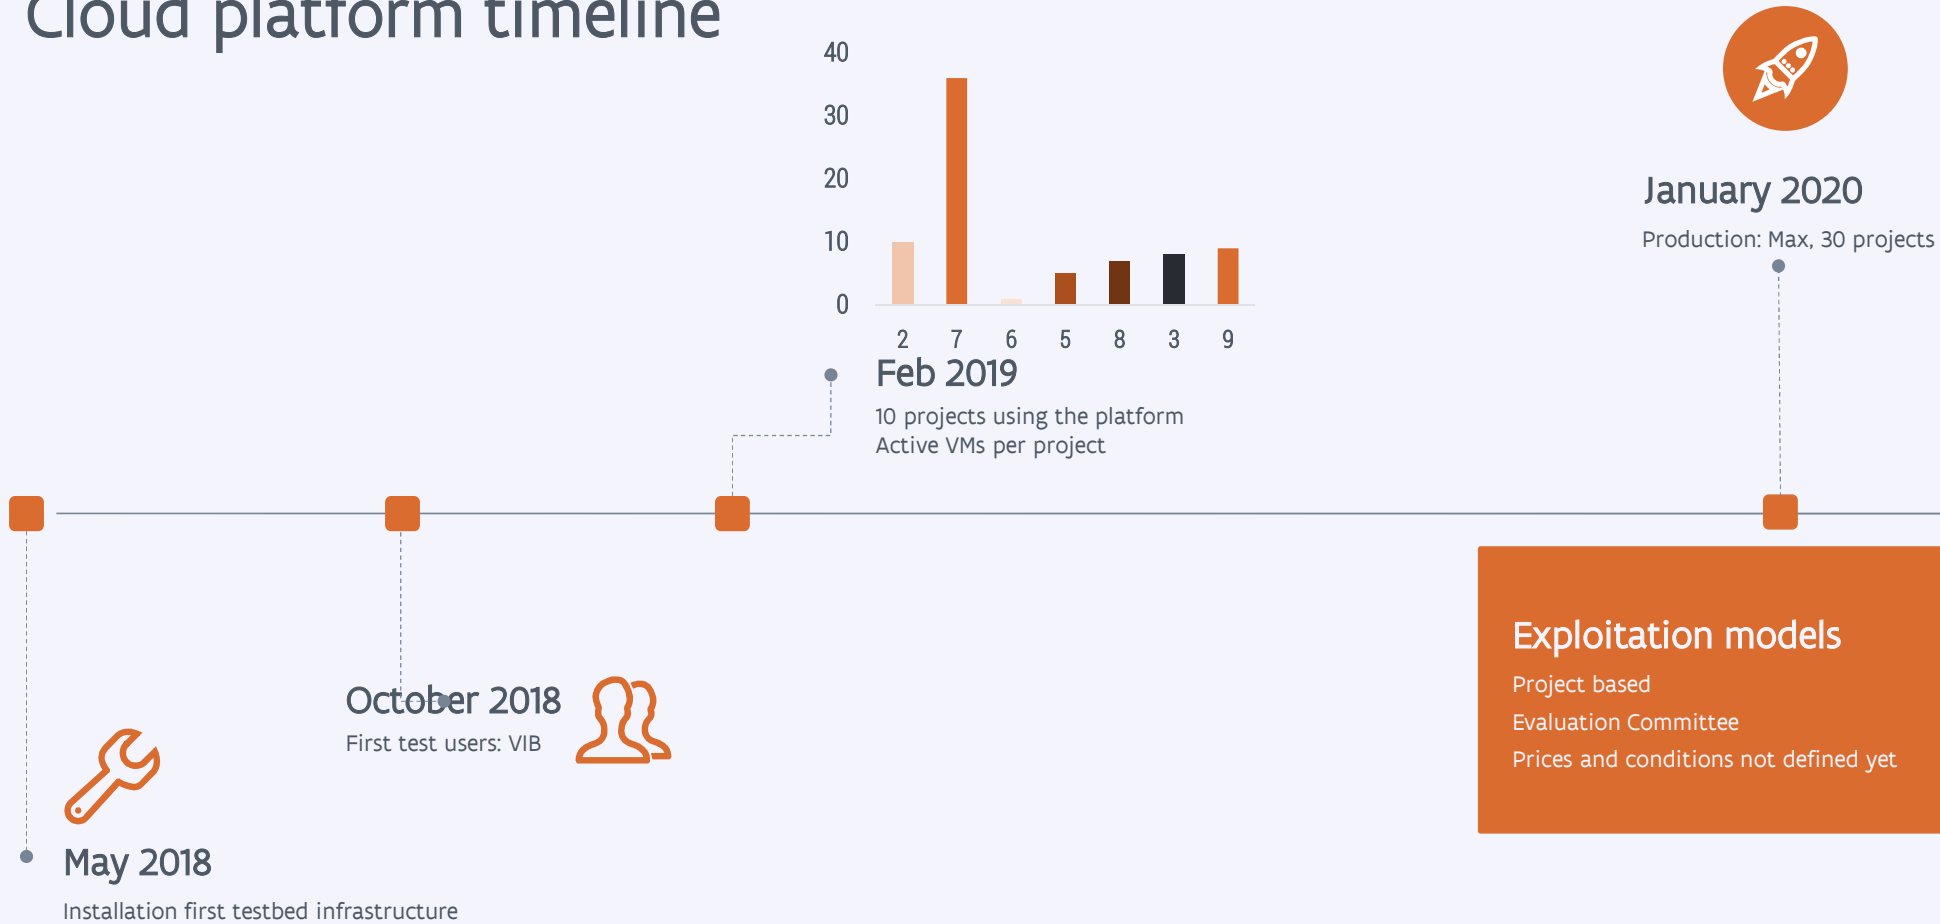


WP Lead

Stijn De Weirdt



# Cloud platform timeline




## Authentication

### Tier1-Cloud Dashboard

### VSC SSO Account Portal

## Orchestration



- ✓ An easy way to generate complex cloud infrastructures as user with YAML templates.
- ✓ Compatible with AWS CloudFormation templates.
- ✓ It is possible to upload your own templates use the wizard or download templates directly from an external URL.

## Compute / Network



- ✓ 26 Hypervisors (old Tier-1 Muk compute nodes).
- ✓ Intel(R) Xeon(R) CPU E5-2670 0 @ 2.60GHZ.
- ✓ 16 Cores and 128Gb ram per hypervisor.
- ✓ At least one public IP per project with port forwarding support.
- ✓ Several private class C networks per project.

## Storage



- ✓ 100 Tb storage using Ceph as backend.
- ✓ Public VM images maintained by the VSC Cloud team (Ubuntu, CentOS..).
- ✓ Users can manage their own shared filesystem between VMs.



# Tier-1 SaaS Cloud component: demo




## Log in

**Authenticate using**

VSC Accountpage ▾

If you are not sure which authentication method to use, contact your administrator.

Sign In



# Questions?



Tier-1 SaaS: Data component

TEAM WORK

# Data workgroup



System admim

Kenneth Waegeman




System admin

Yorick Poels




System admin

Peter Verraedt



Storage admin

Tom Van Mirlo



System admim

Kristoff Van Buggenhout



User Support

Kurt Lust



User Support

Ward Poelmans



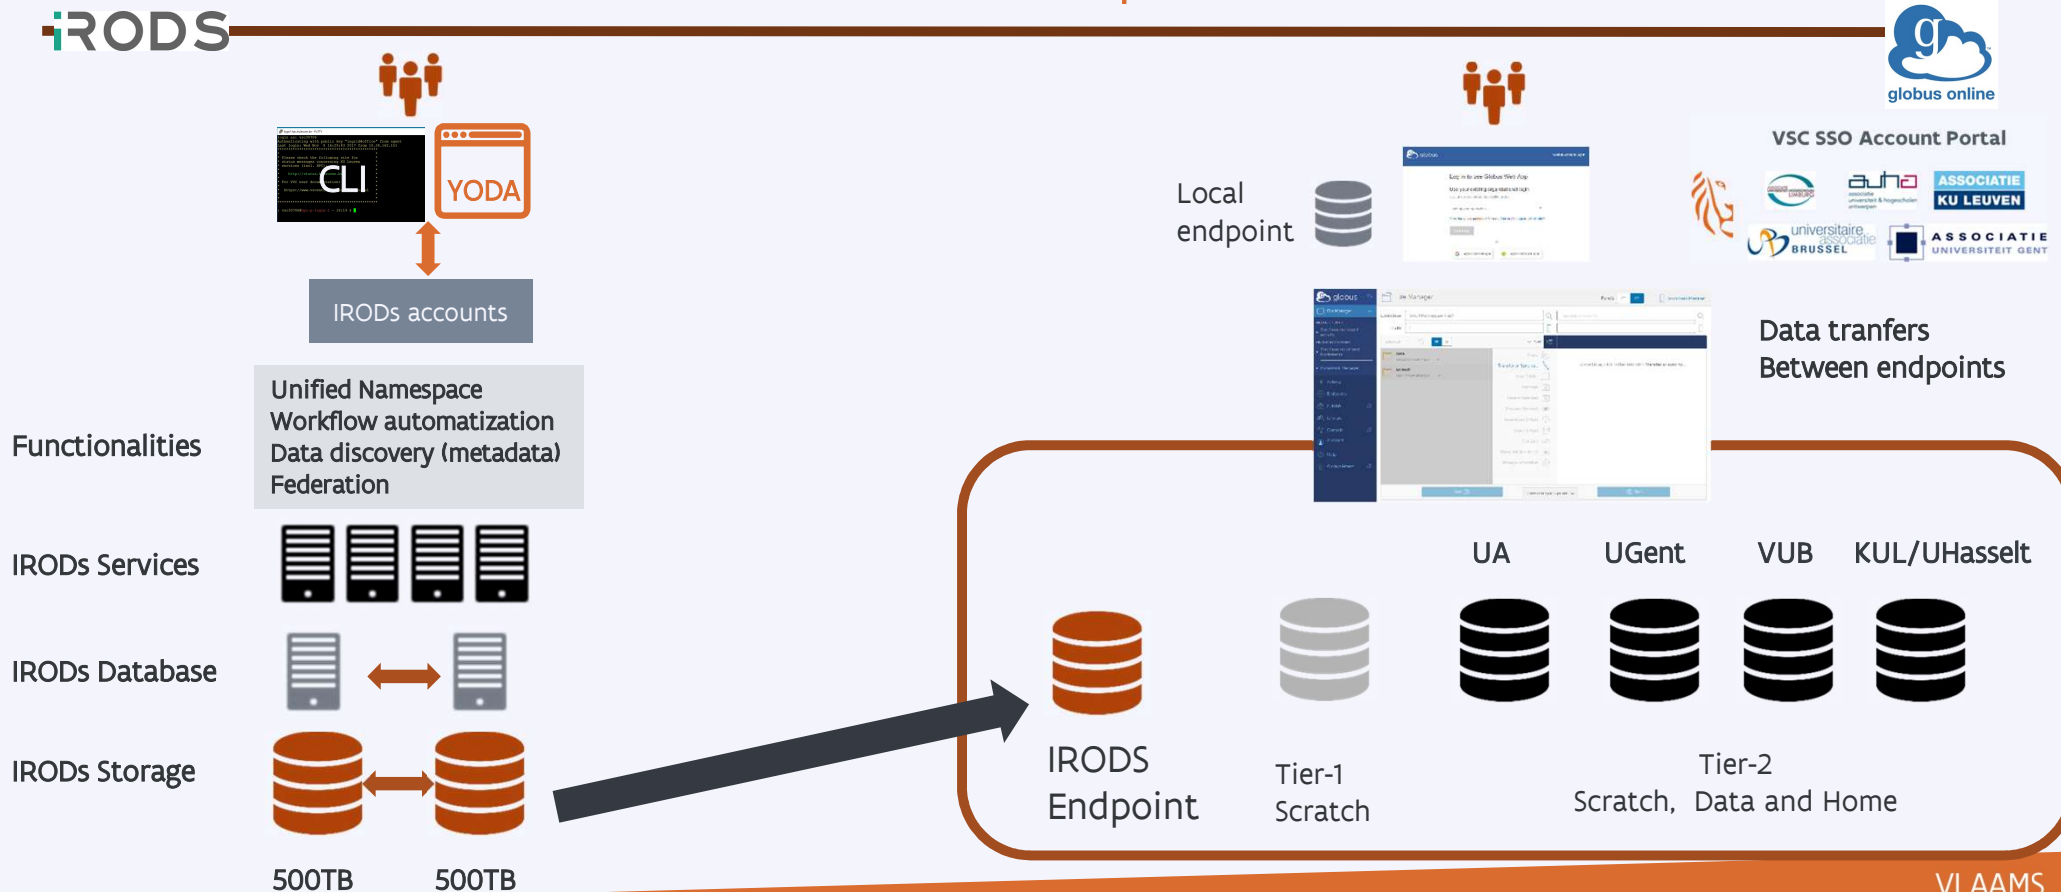
WP Lead

Stijn De Weirdt




# Tier-1 SaaS Data component: architecture

## Tier-1 Data component



# Data platform timeline



# What is Globus?

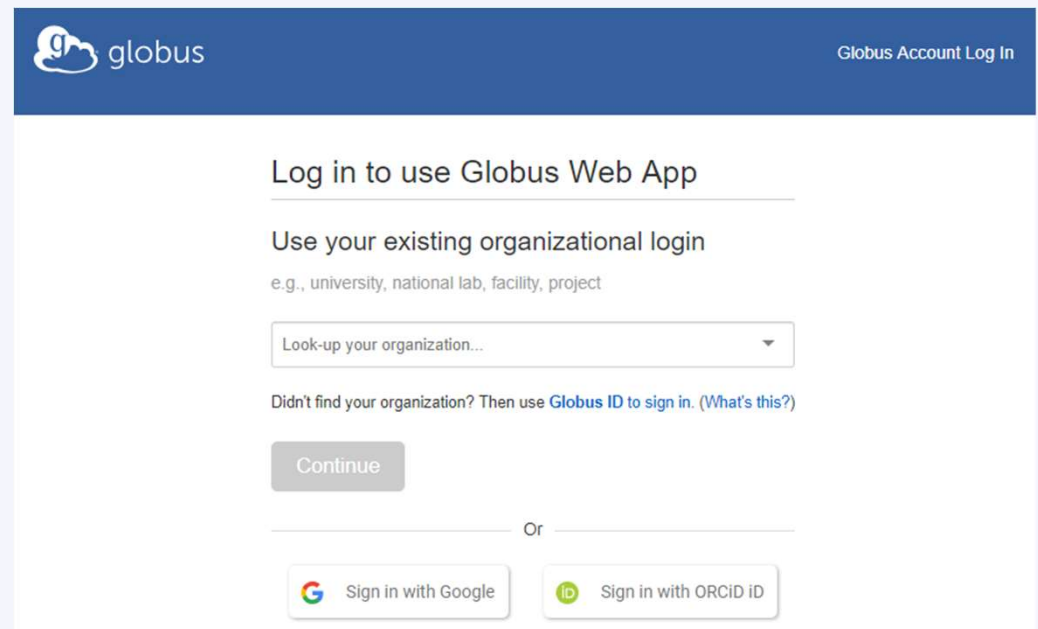
- A data management cloud
- Enables researchers to transfer and share data
- Search capabilities within their own research data applications and services
- Utilizes advanced identity management, single sign-on, and authorization capabilities



globus.org

# Access

- Go to *globus.org*



The screenshot shows the Globus Web App login interface. At the top left is the Globus logo, and at the top right is a link for "Globus Account Log In". The main heading is "Log in to use Globus Web App". Below this, it prompts the user to "Use your existing organizational login" with examples like "university, national lab, facility, project". There is a dropdown menu labeled "Look-up your organization...". A link says "Didn't find your organization? Then use Globus ID to sign in. (What's this?)". A "Continue" button is present. Below this, separated by "Or", are two buttons: "Sign in with Google" and "Sign in with ORCID iD".

# Access

- Go to *globus.org*
- Log in with your institution credentials

globus

Globus Account Log In

Log in to use Globus Web App

Use your existing organizational login  
e.g., university, national lab, facility, project

Look-up your organization...

Didn't find your organization? Then use [Globus ID to sign in.](#) (What's this?)

Continue

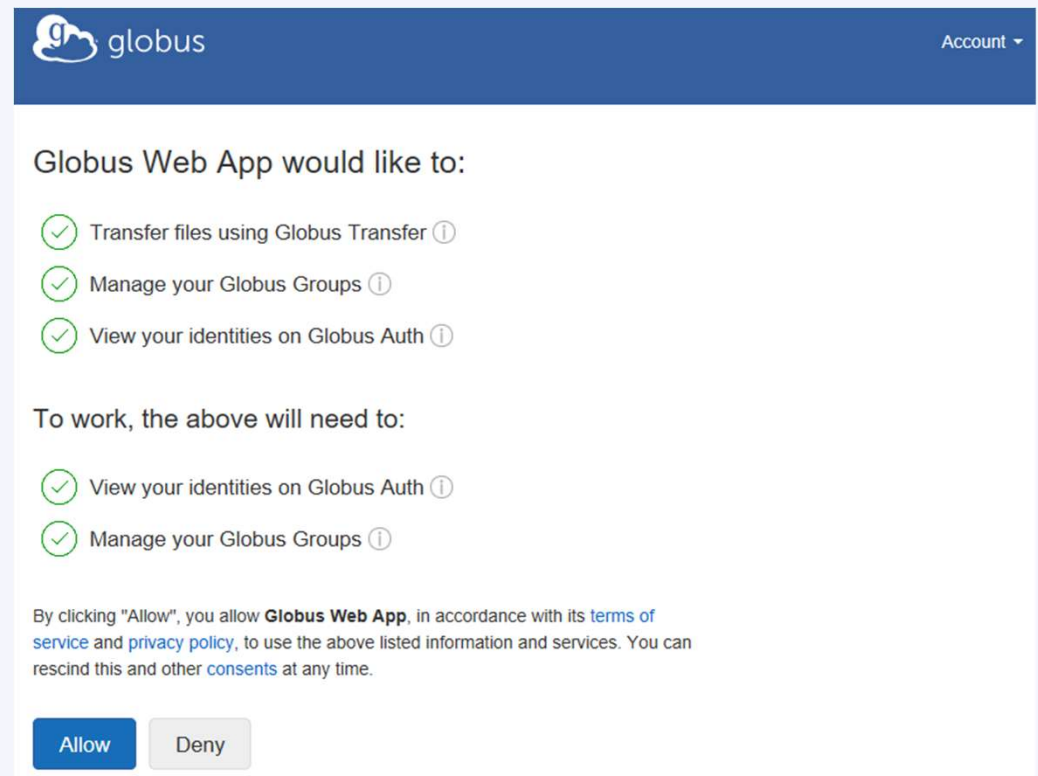
Or

Sign in with Google

Sign in with ORCID iD

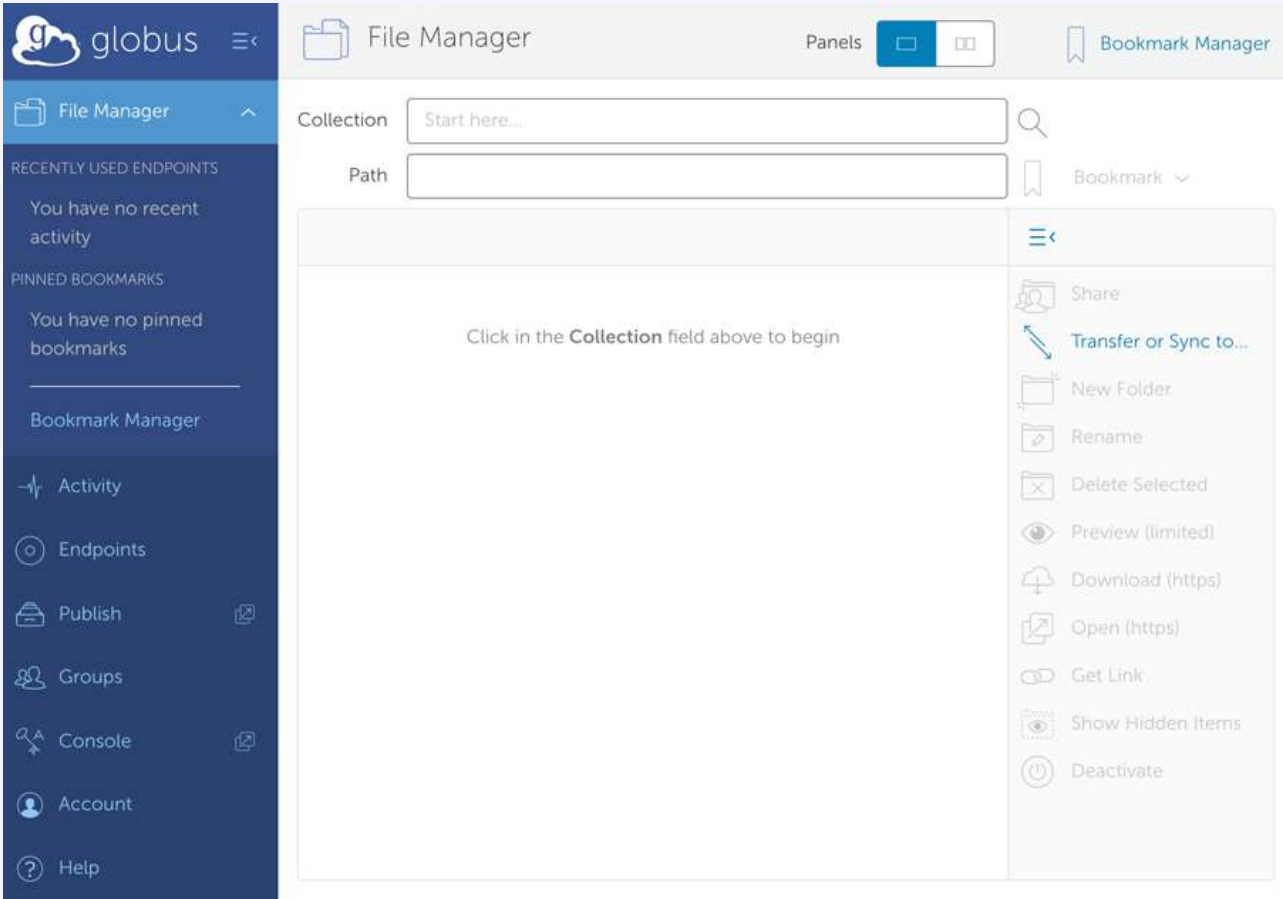
# Access

- Go to *globus.org*
- Log in with your institution credentials
- Give Globus permission to use your identity to access information and perform actions

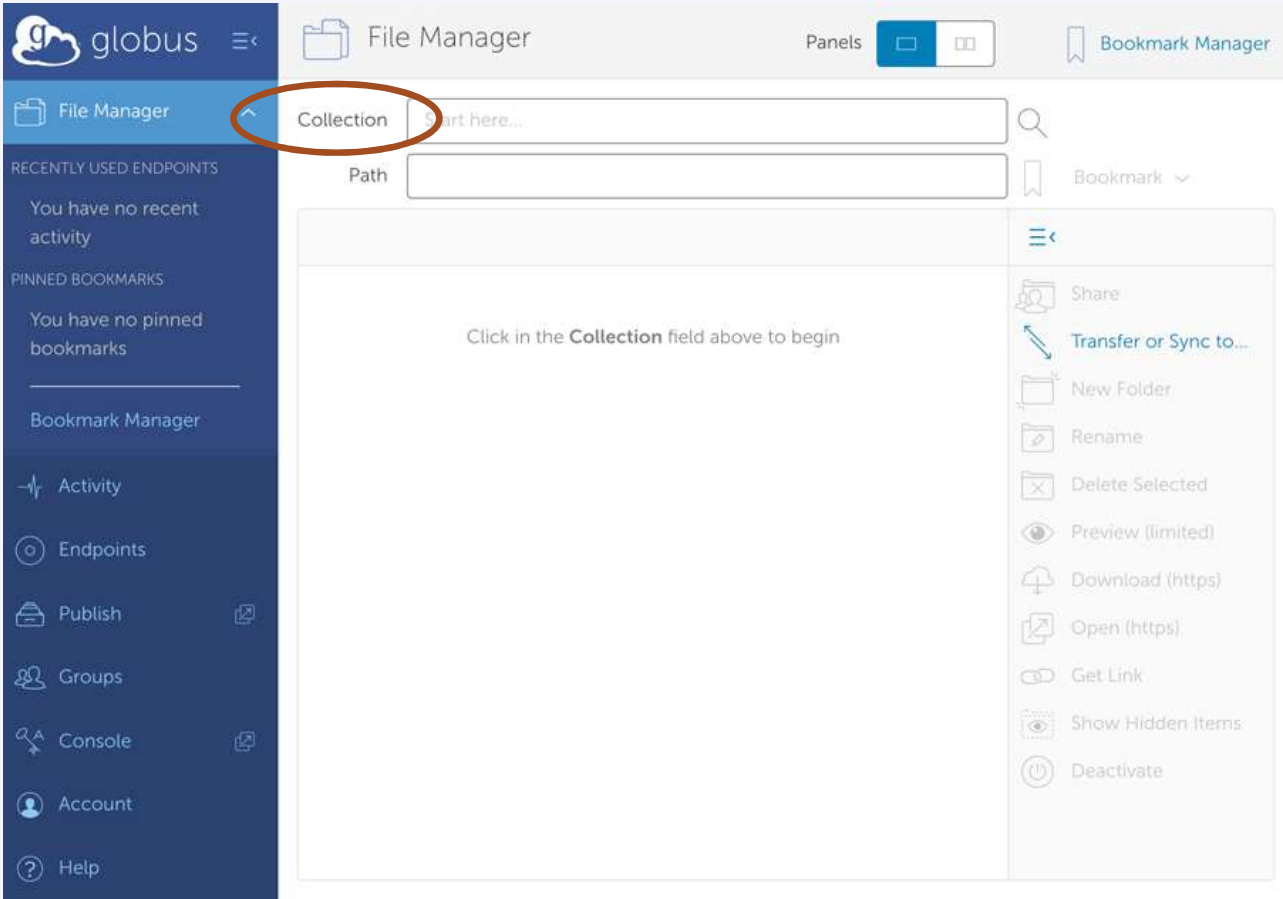


The screenshot shows the Globus Web App permission dialog box. At the top, there is a blue header with the Globus logo and the text "globus" on the left, and "Account" with a dropdown arrow on the right. Below the header, the text "Globus Web App would like to:" is followed by three items, each with a green checkmark icon and a help icon (i): "Transfer files using Globus Transfer", "Manage your Globus Groups", and "View your identities on Globus Auth". Below this, the text "To work, the above will need to:" is followed by two items, each with a green checkmark icon and a help icon (i): "View your identities on Globus Auth" and "Manage your Globus Groups". At the bottom, there is a paragraph of text: "By clicking 'Allow', you allow **Globus Web App**, in accordance with its [terms of service](#) and [privacy policy](#), to use the above listed information and services. You can rescind this and other [consents](#) at any time." Below the text are two buttons: "Allow" (blue) and "Deny" (grey).

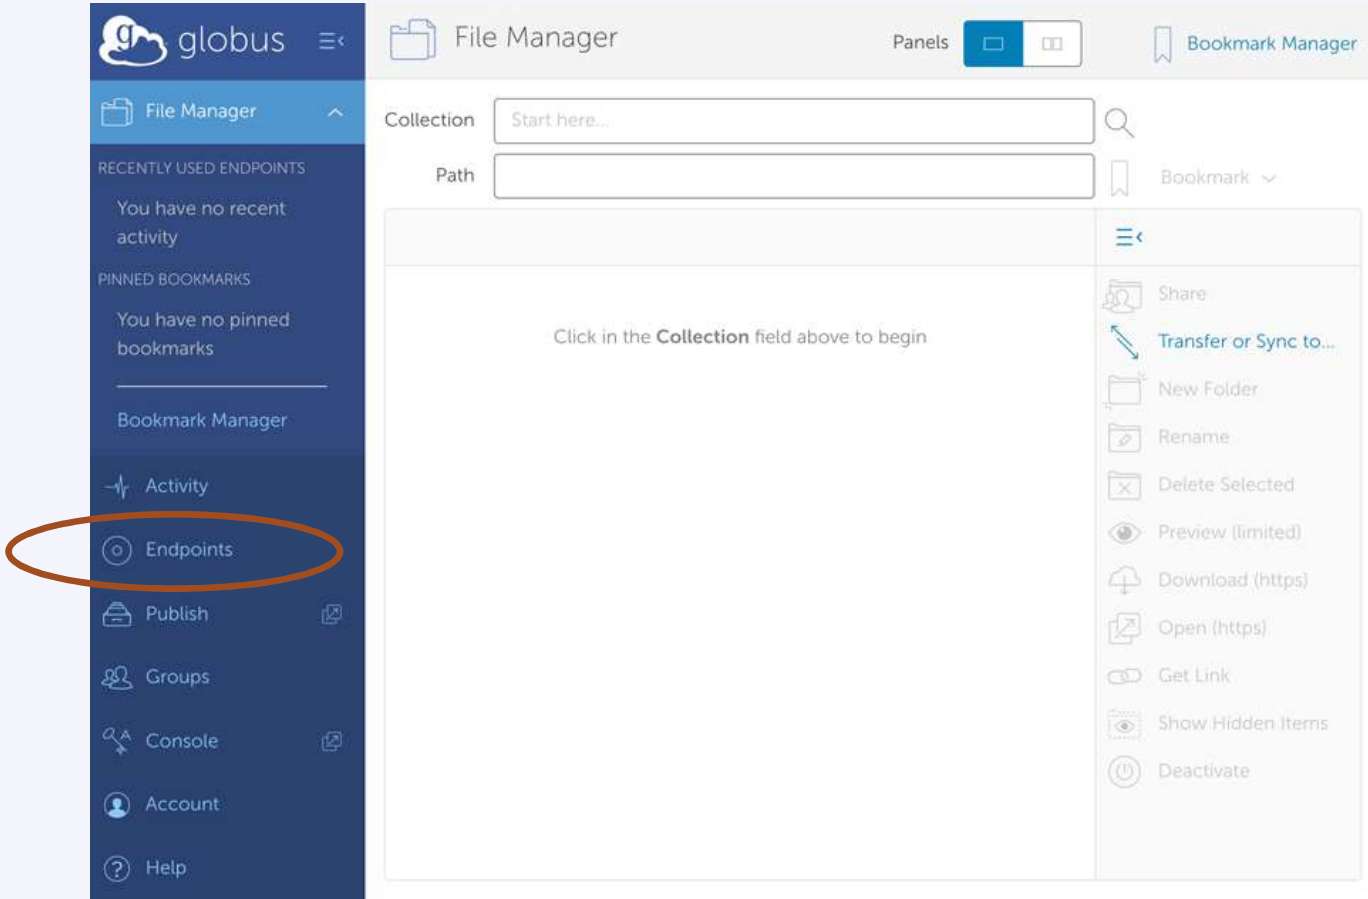
# Globus File Manager



# Globus File Manager



# Globus File Manager



# Collections and Endpoints

- Within Globus user data resides in a location called Globus Collection

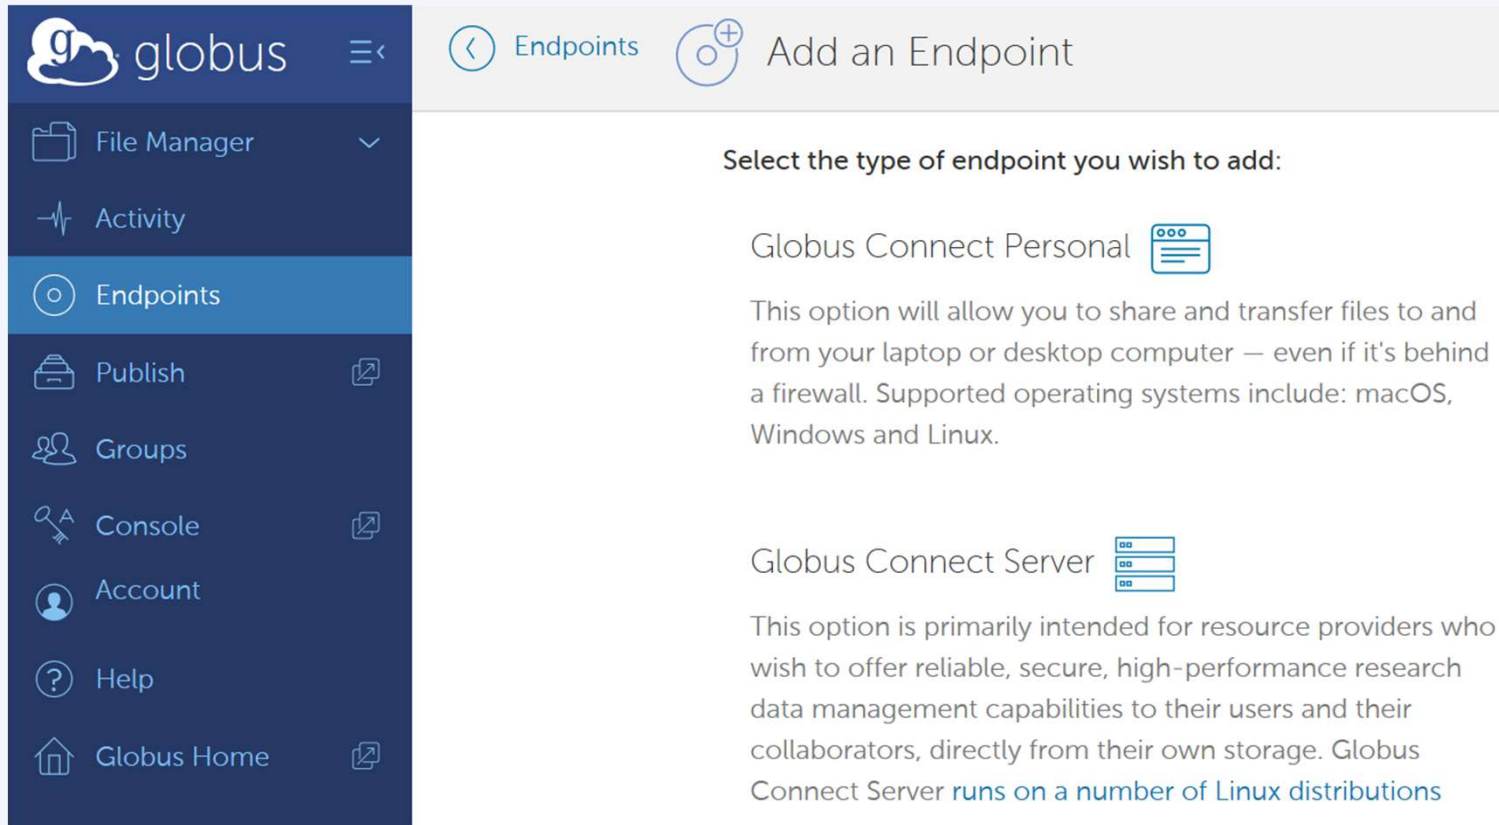


# Collections and Endpoints

- Within Globus user data resides in a location called Globus Collection
- An Endpoint is a resource (server, cluster, storage system, laptop, etc...) that hosts collections



# Create an Endpoint



The screenshot shows the Globus Connect web interface. On the left is a dark blue sidebar with the Globus logo and a menu containing: File Manager, Activity, Endpoints (highlighted), Publish, Groups, Console, Account, Help, and Globus Home. The main content area has a light grey header with a back arrow, the text 'Endpoints', and a plus icon followed by 'Add an Endpoint'. Below the header, the text 'Select the type of endpoint you wish to add:' is displayed. Two options are listed: 'Globus Connect Personal' with a laptop icon and a descriptive paragraph, and 'Globus Connect Server' with a server rack icon and another descriptive paragraph.

globus

Endpoints Add an Endpoint

Select the type of endpoint you wish to add:

Globus Connect Personal

This option will allow you to share and transfer files to and from your laptop or desktop computer – even if it's behind a firewall. Supported operating systems include: macOS, Windows and Linux.

Globus Connect Server

This option is primarily intended for resource providers who wish to offer reliable, secure, high-performance research data management capabilities to their users and their collaborators, directly from their own storage. Globus Connect Server runs on a number of Linux distributions

# Create an Endpoint

The screenshot shows the Globus Connect Personal web interface. On the left is a dark blue sidebar with the Globus logo and a menu containing: File Manager, RECENTLY USED (with a note 'You have no recent activity'), PINNED BOOKMARKS (with a note 'You have no pinned bookmarks'), Bookmark Manager, Activity, Endpoints, Publish, Groups, Console, Account, Help, and Globus Home. The main content area is titled 'Create New GCP Endpoint' and is divided into three steps:

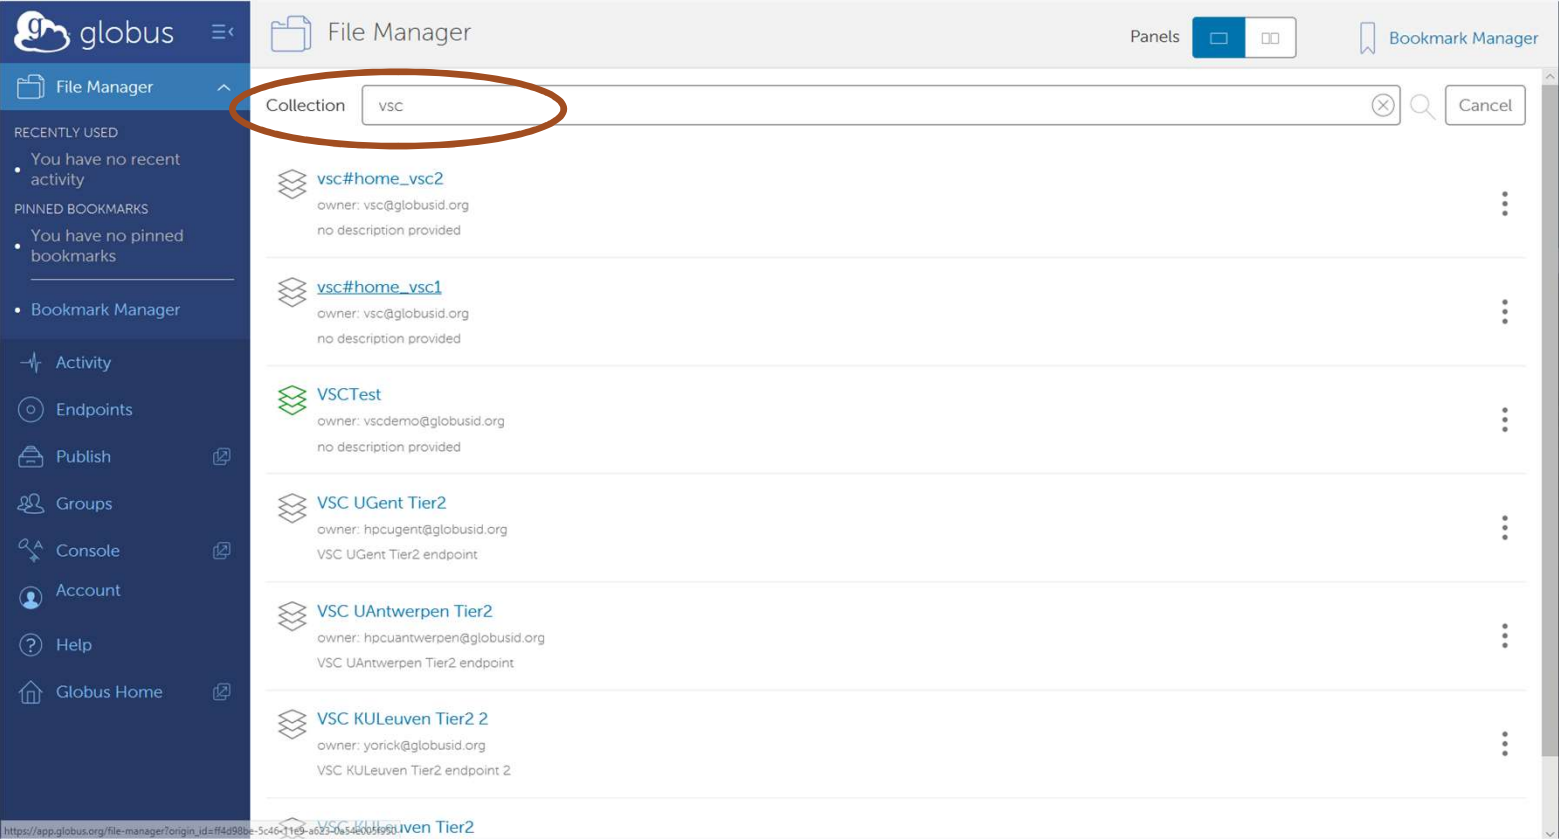
- STEP 1 Name Your Endpoint**
  - Endpoint Display Name:  (with a help icon)
  - Text: Please enter a human-friendly name for your Globus Connect Personal endpoint to help you identify it.
  - Checkbox:  this will be a high assurance endpoint
  - Text: Once this installation process is completed you will have to get this endpoint certified by your IT department before high assurance data such as HIPAA (with BAA), PHI, PII or controlled unclassified is allowed to be accessed.
  - Associated Identity: u00000000@kuleuven.be
- STEP 2 Generate a Setup Key for Installation**
  - Text: This setup key will be required during the installation of Globus Connect Personal on your computer.
  - Button: Generate Setup Key
- STEP 3 Download & Install Globus Connect Personal**
  - Text: Click one of the buttons below to download and install Globus Connect Personal for your operating system.
  - Buttons: [for Mac OS X](#), [for Linux](#), [for Windows](#)
  - Text: Once downloaded, run the installer. When prompted, paste in the Setup Key to complete the installation.

# Create an Endpoint

The screenshot displays the Globus Endpoints management interface. On the left is a dark blue sidebar with the Globus logo and navigation items: File Manager, Activity, Endpoints (highlighted), Publish, Groups, Console, Account, Help, and Globus Home. The main area is titled 'Endpoints' and features a search bar 'Search all endpoints' and a '+ Create new endpoint' button. Below these are filter tabs: 'Recently Used', 'In Use', 'Shareable by You', 'Shared with You', and 'Administered by You' (which is selected and highlighted with a dashed border). A 'Filter administered by you' dropdown is also present. A table lists the endpoints with columns for ENDPOINT, STRICT, STATUS, ROLE, and SHARED. One endpoint is listed: 'MyGlobusFiles' (Globus Connect Personal) with a status of 'ready' and a role of 'Administered by You'. The table also includes icons for sharing and navigation.

ENDPOINT	STRICT	STATUS	ROLE	SHARED
MyGlobusFiles Globus Connect Personal		ready		

# Select a Globus Collection



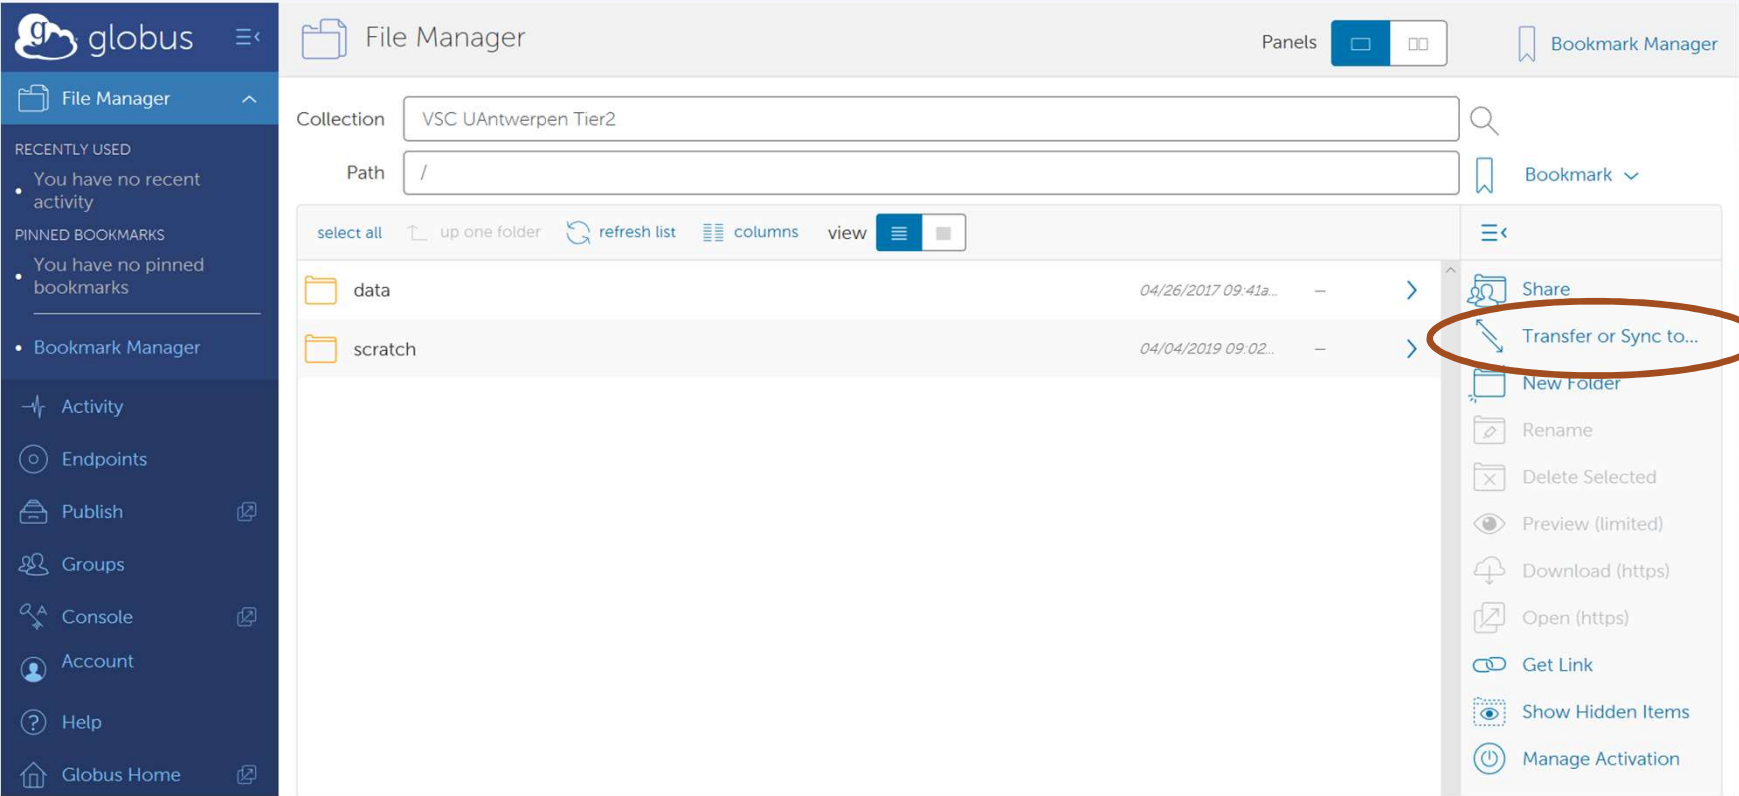
# Select a Globus Collection

The screenshot shows the Globus File Manager interface. On the left is a dark blue sidebar with the Globus logo and a menu containing: File Manager, RECENTLY USED (You have no recent activity), PINNED BOOKMARKS (You have no pinned bookmarks), Bookmark Manager, Activity, Endpoints, Publish, Groups, Console, Account, Help, and Globus Home. The main area is titled 'File Manager' and contains a search bar with 'Collection' and 'Path' labels. The 'Collection' field contains 'VSC UAntwerpen Tier2'. Below the search bar is a large white area with a power icon and the text: 'Please authenticate to access VSC UAntwerpen Tier2. When you press the CONTINUE button below you will be redirected to the collection's login page. After logging in, you will be returned here.' A blue 'Continue' button is centered below this text. On the right side of the main area, there is a 'Panels' toggle and a 'Bookmark Manager' link. A context menu is open on the right, listing actions: Share, Transfer or Sync to..., New Folder, Rename, Delete Selected, Preview (limited), Download (https), Open (https), Get Link, Show Hidden Items, and Manage Activation.

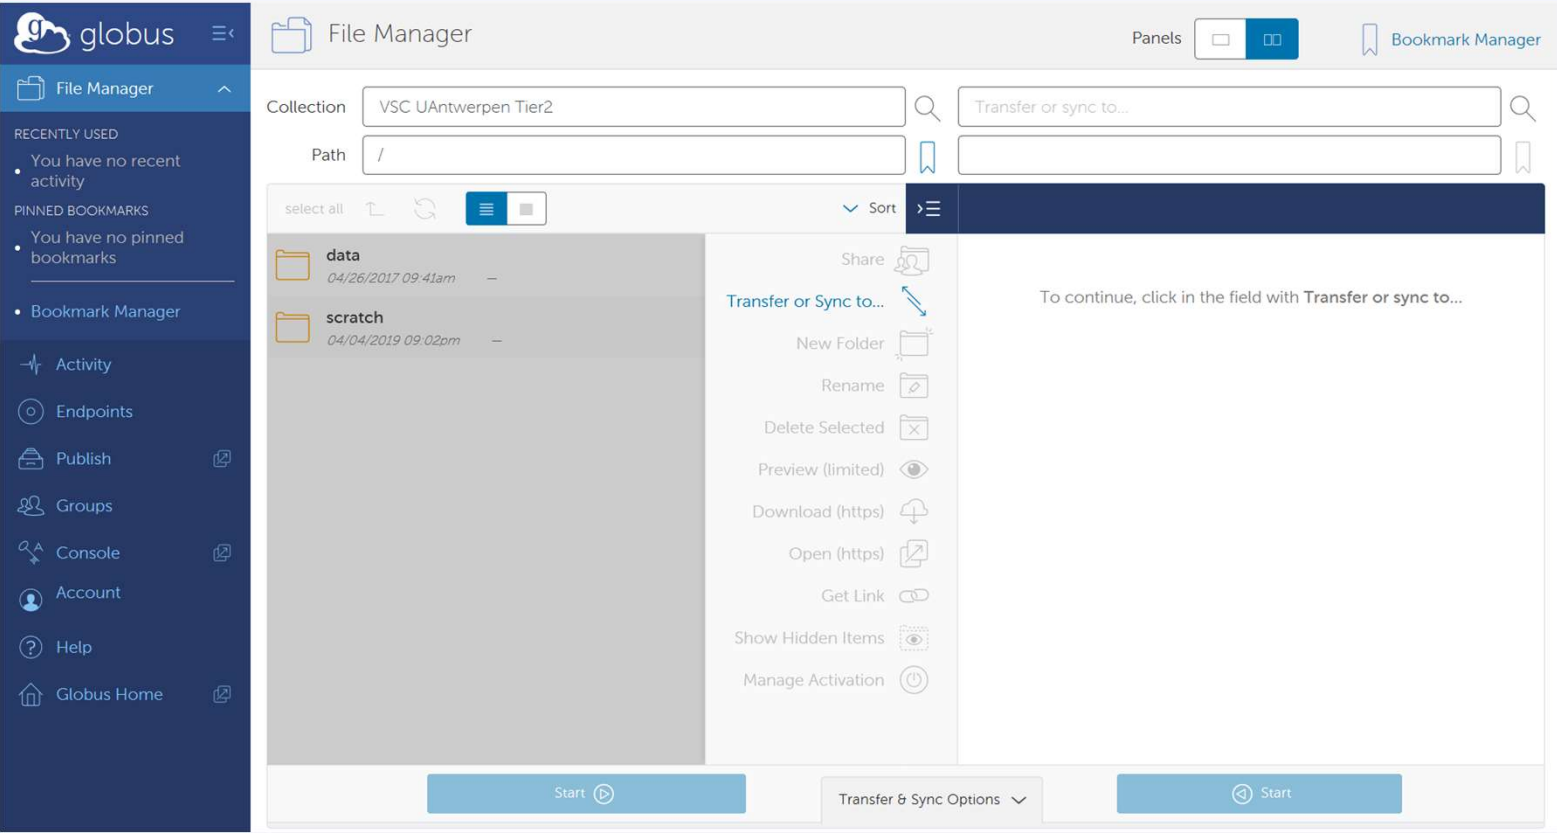
# Select a Globus Collection

The screenshot displays the Globus File Manager interface. On the left is a dark blue sidebar with navigation options: File Manager, RECENTLY USED (no recent activity), PINNED BOOKMARKS (no pinned bookmarks), and Bookmark Manager. Below these are icons for Activity, Endpoints, Publish, Groups, Console, Account, Help, and Globus Home. The main area is titled 'File Manager' and shows the 'Collection' as 'VSC UAntwerpen Tier2' and the 'Path' as '/'. Below the path are controls for 'select all', 'up one folder', 'refresh list', 'columns', and 'view'. A table lists two folders: 'data' (dated 04/26/2017 09:41a...) and 'scratch' (dated 04/04/2019 09:02...). On the right, a context menu is open, listing actions such as Share, Transfer or Sync to..., New Folder, Rename, Delete Selected, Preview (limited), Download (https), Open (https), Get Link, Show Hidden Items, and Manage Activation.

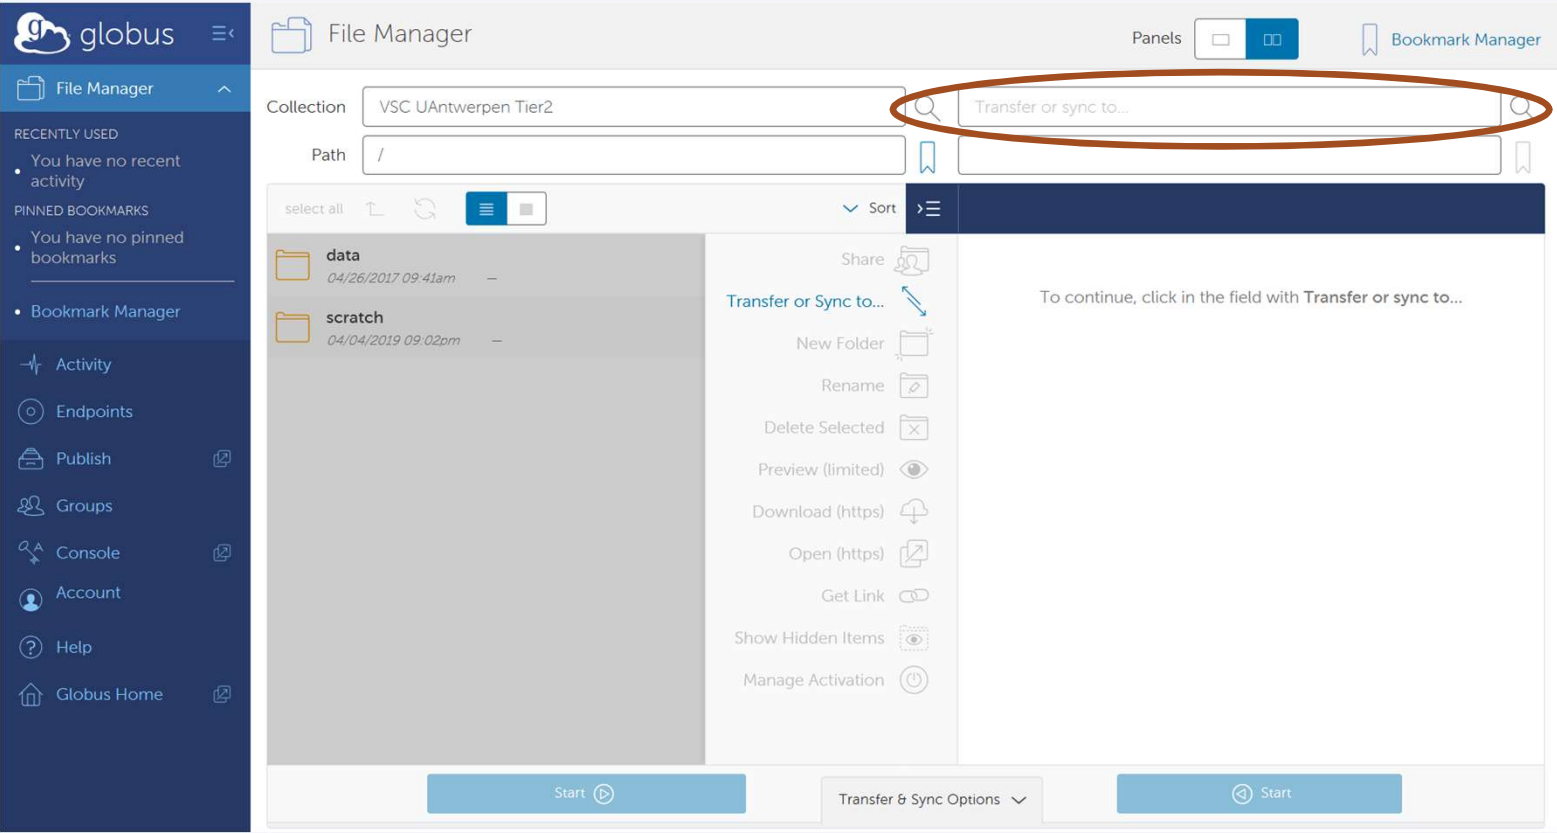
# Select a Globus Collection



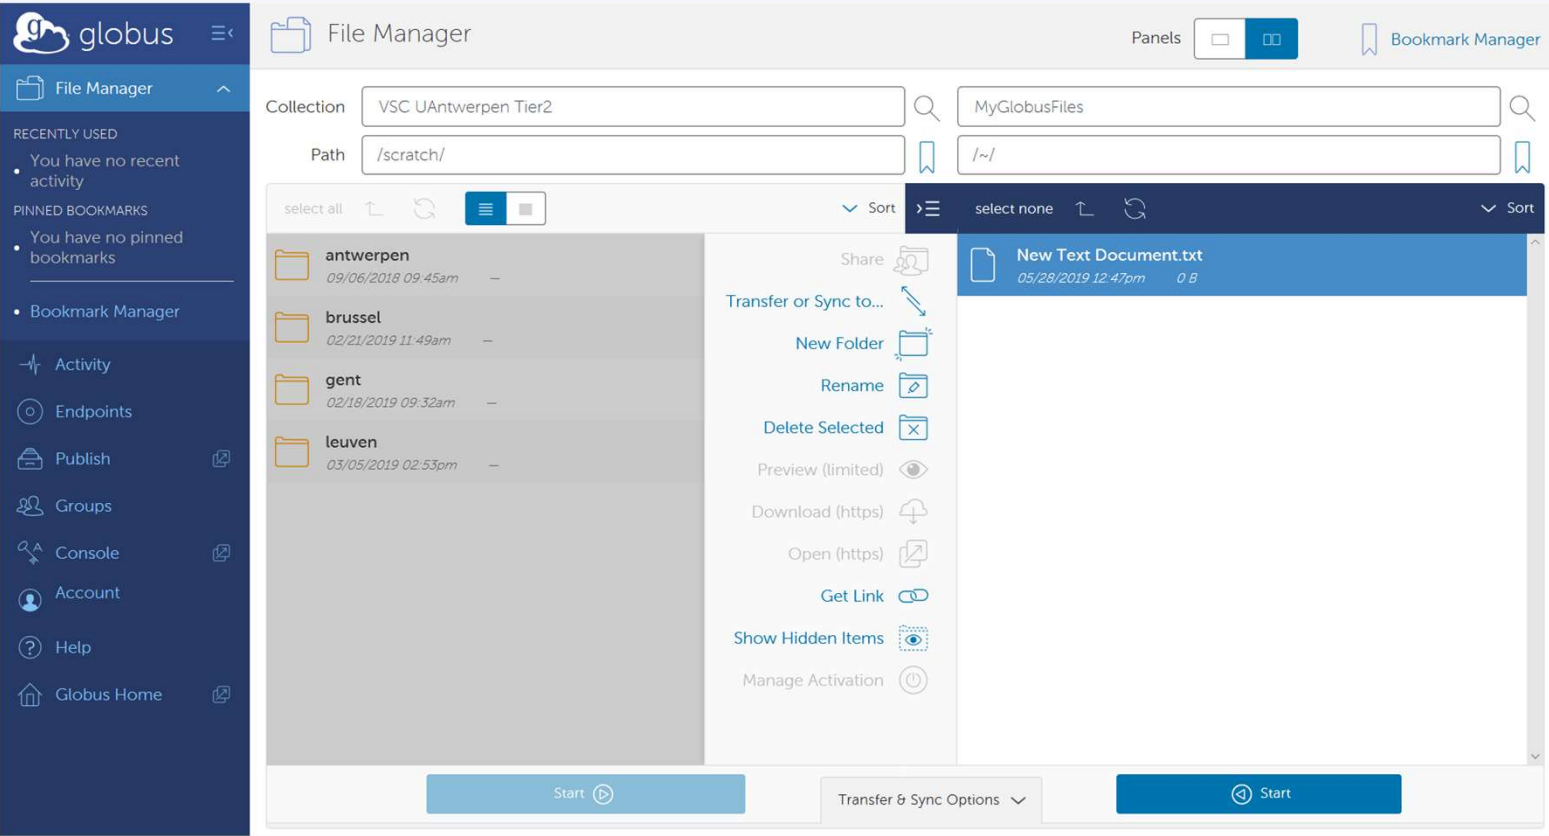
# Transferring Files



# Transferring Files



# Transferring Files




# Are you interested in using it?

Contact us:

[data@vscentrum.be](mailto:data@vscentrum.be)

Documentation: to be available soon at VSC documentation website



# Questions?