





Consolidate the HiPS network

IVOA Santiago Interop – 27 & 30 October 2017

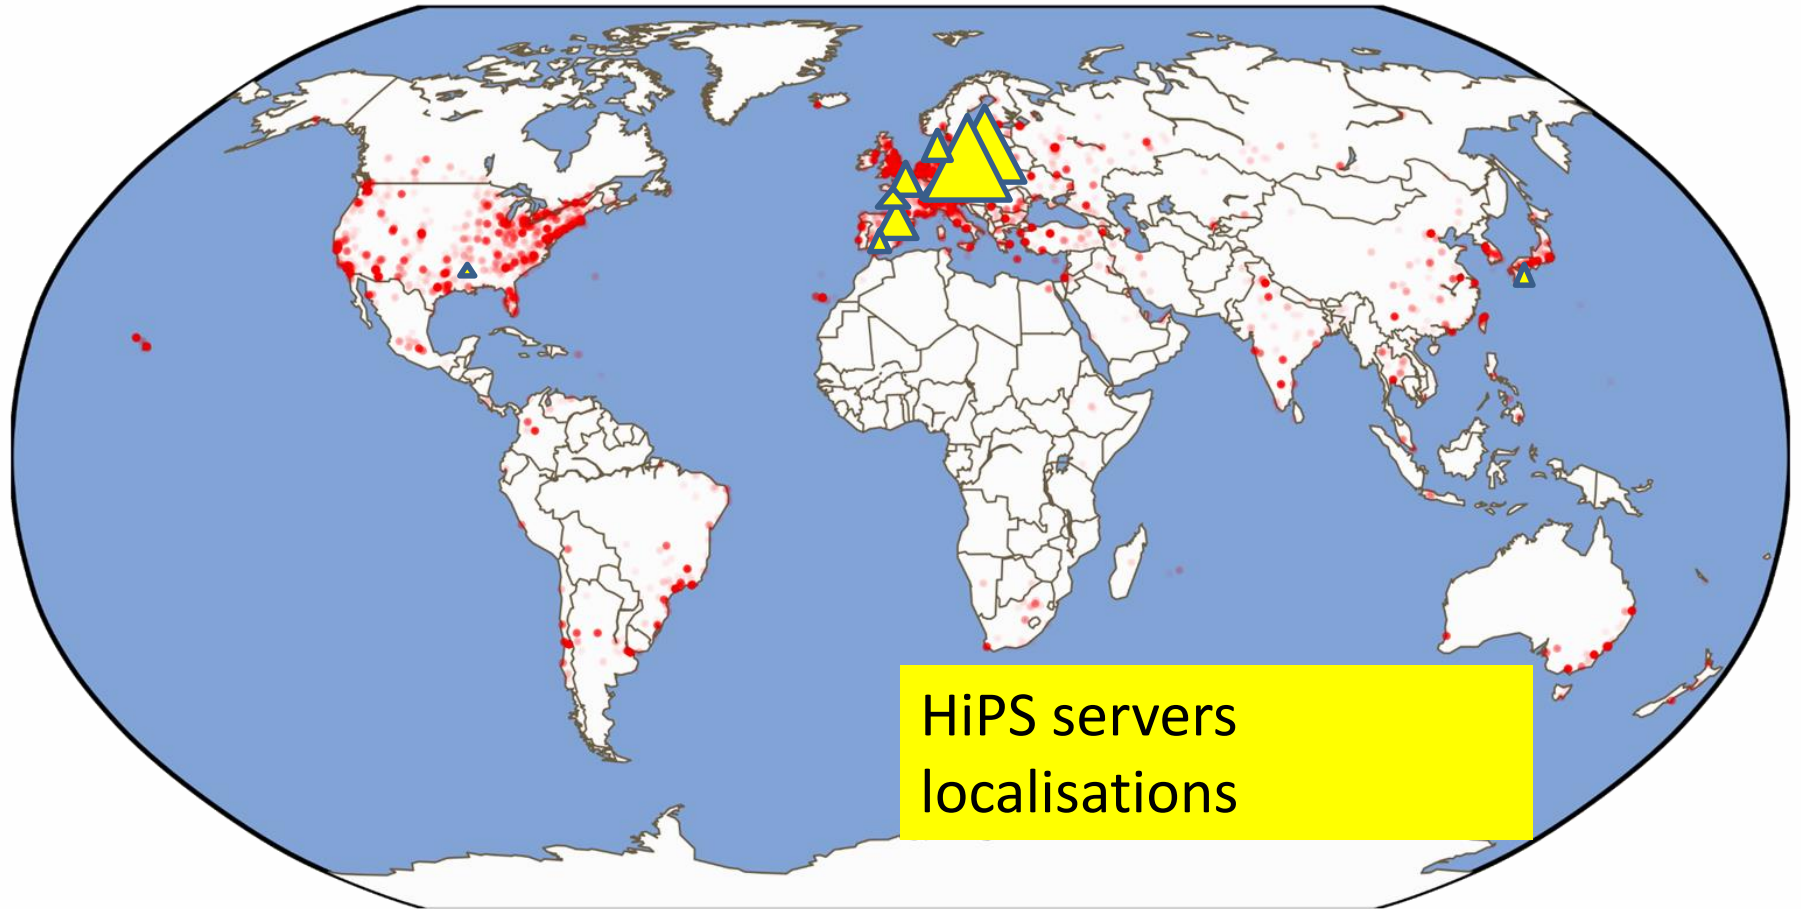
Pierre Fernique




□ They are HiPS clients...



□ ... and they are HiPS servers



□ Where are you?

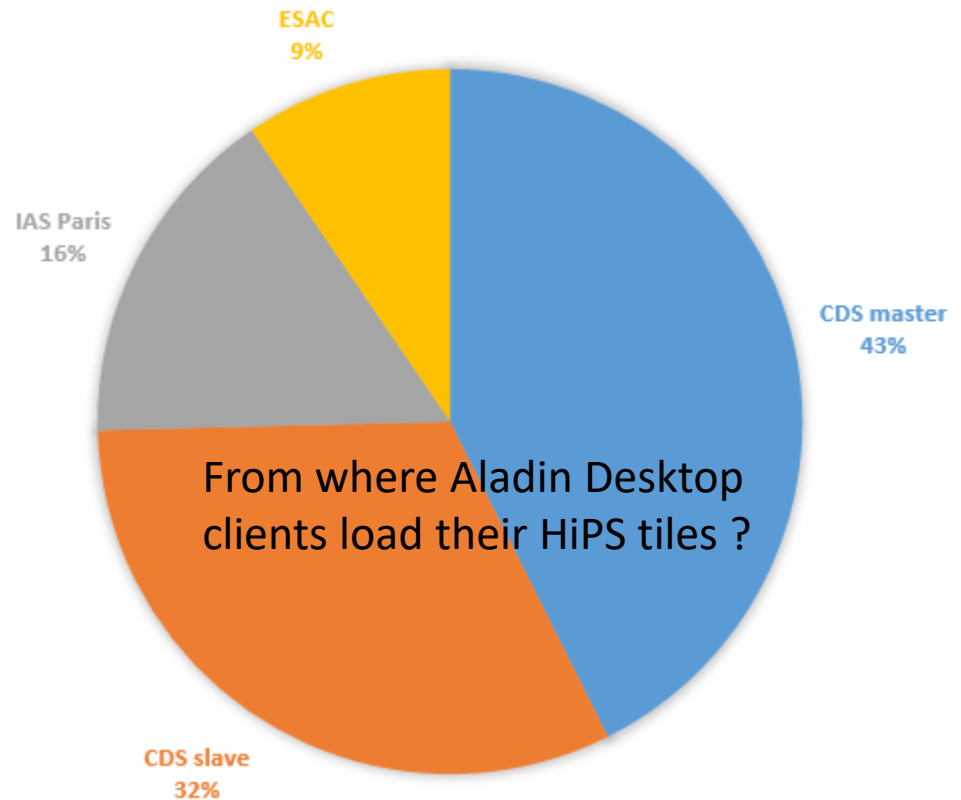


□ HiPS key figures...

- **400** HiPS for **150TB** of data
- But among them, **10** HiPS (colored=compressed) are heavily used (representing less than **3TB**)
=> **DSS**, SDSS, 2MASS, Spitzer, Herschel, PLANCK, ...

From where are you loading your HiPS ?

- Clients may manually or automatically choose a HiPS server according to the HiPS they want and the number of mirror sites for this HiPS
- **Presently ~75%** of the HiPS requests arrived on CDS HiPS servers at Strasbourg/France (rough approximation based on Aladin Desktop logs)




□ What's the challenge ?

- **150 000** tiles/day displayed by Aladin Lite (& derived tools)
+ **150 000** tiles/day by Aladin Desktop
- Thanks to user cache, only half of them are really requested, and only **~30GB/day** are effectively downloaded (upper estimation)
- Not a problem of server load, but more a need to reduce the **response time between the client and the closest HiPS server**


□ Consolidate the HiPS network

- In a perfect HiPS world, we would have:
 - one (or two) HiPS nodes per continent
 - with a copy of the 10 more used HiPS



□ What is a HiPS node ?


Nothing more:




- than an **http server** (apache for instance)
- Distributing collections of **regular files** (HiPS tiles)

*(no data base, no CGI script,
and even... no IVOA standards to read & understand, ...)*

□ Become a HiPS mirror nodes...





you can...