

ATLAS Network Requirements

Alastair Dewhurst, Roger Jones, Shawn McKee on
behalf of ATLAS

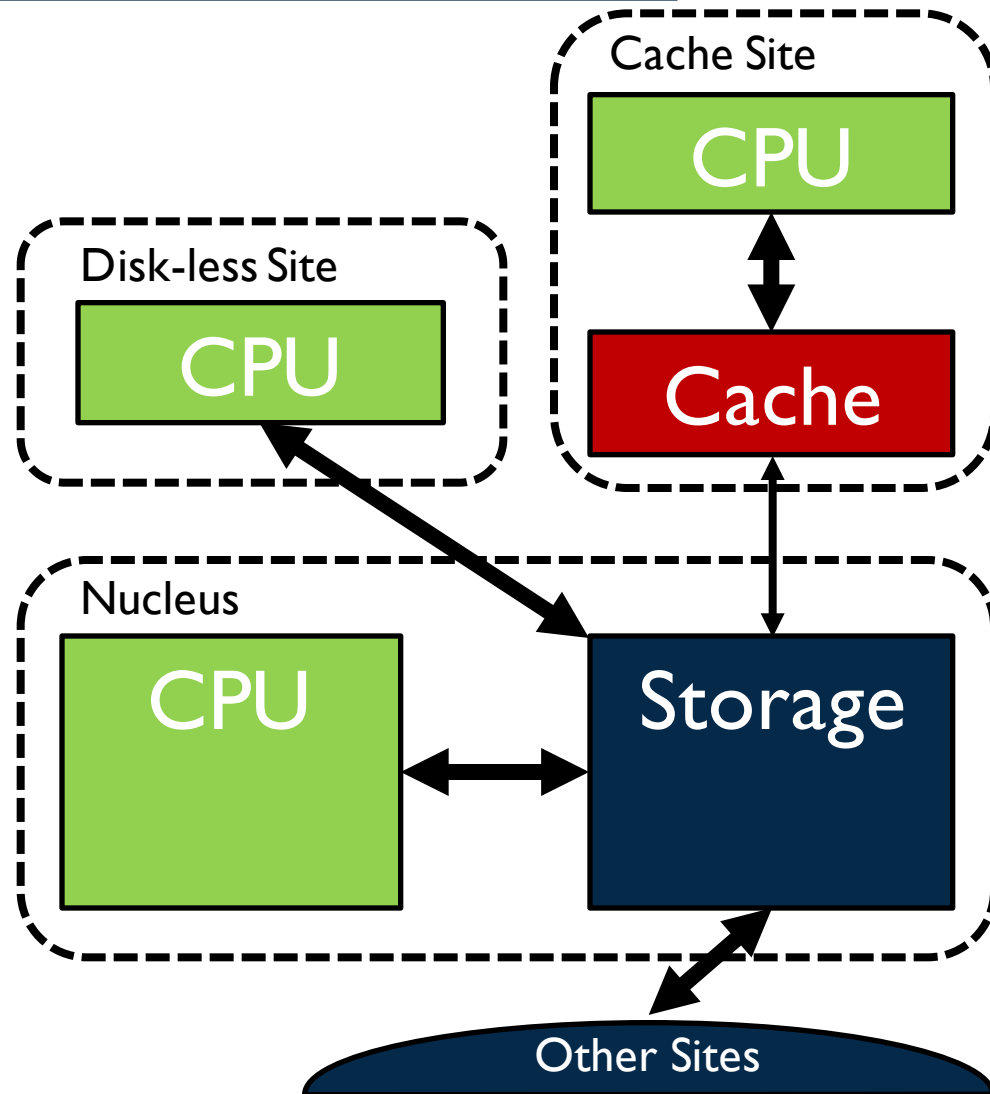


ATLAS Model

- ATLAS are moving to a Nucleus and Satellite Model.
- Sites meeting certain criteria can be considered Nucleus:
 - Storage capacity > 1PB
 - Good Network throughput (Closeness metric)
 - Site availability: > 95% (using ASAP metric)
- We will continue to make use of site whose storage doesn't meet these requirements, but the type of work run may be limited.
- Each Task is assigned a Nucleus site and 0 or more Satellites.
 - Note: a Nucleus site will often act as a satellite for some tasks it runs.
 - PanDA dynamically assigns satellites to nuclei on a per task basis.
 - **ATLAS is already automatically balancing work between Disk and CPU rich sites.**



Network transfers



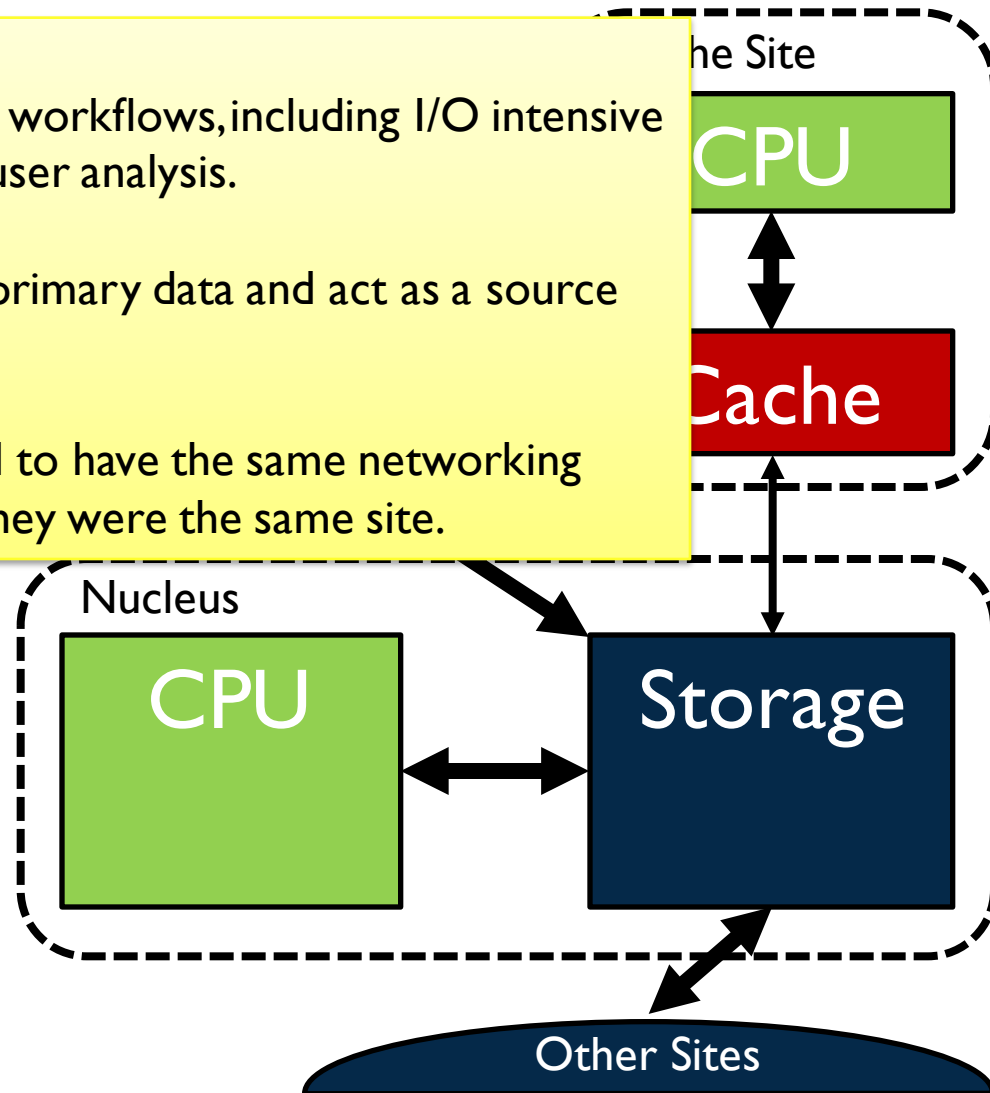
Network transfers

Nucleus sites

Nucleus sites will run all workflows, including I/O intensive such as merge jobs and user analysis.

Nucleus sites will store primary data and act as a source for data distribution.

Federated sites will need to have the same networking between sub-sites as if they were the same site.




Network transfers

Disk-less sites

Will run most workflows but not the most I/O intensive such as merge or user analysis.

WN talk directly to Nucleus Storage.



Network transfers

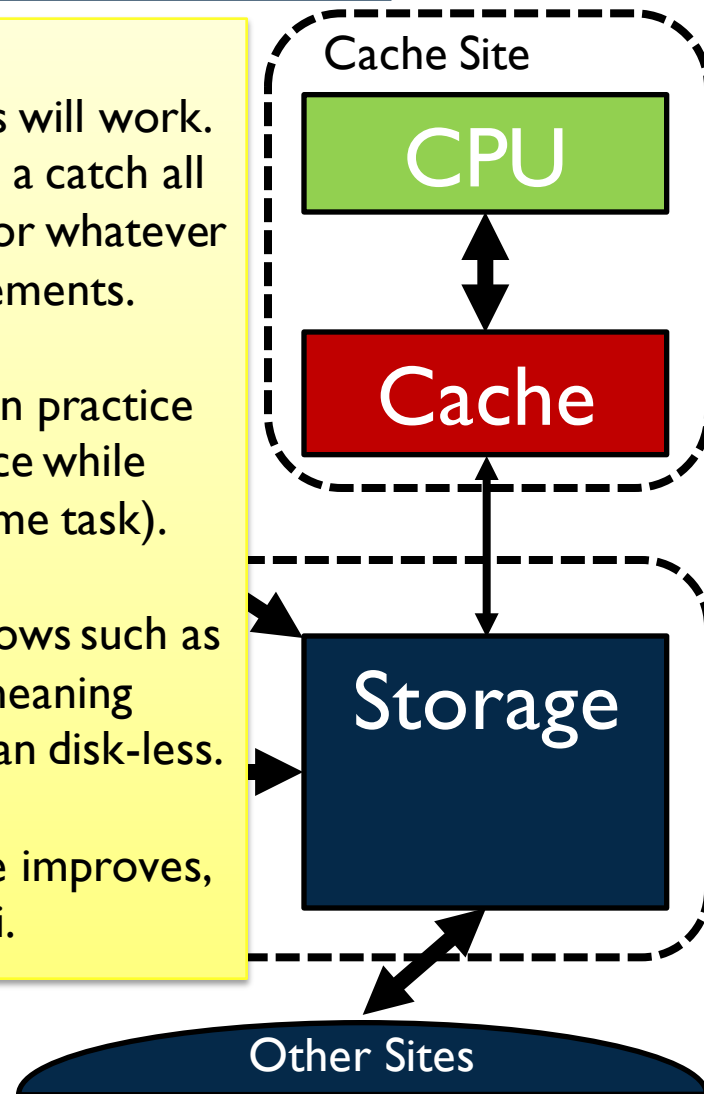
Cache Site

We don't know exactly how cache sites will work. The word "cache" is being used here as a catch all term for sites with some storage that for whatever reason doesn't meet the Nuclei requirements.

However cache sites are implemented in practice the storage should increase performance while reducing overall throughput (for the same task).

Some of the more I/O intensive workflows such as user analysis are likely to be run here meaning network requirements won't be less than disk-less.

As the size and reliability of a cache site improves, the more it will be treated like a Nuclei.



Network Requirements

Nucleus	Now	5 year (2022)	10 year (2027)
Storage Capacity (PB)	2	5	12.5
Total CPU (kHS06)	40	100	250
LAN (Gb/s)	40	200	1000
WAN (Gb/s)	20	60	200

Disk-less	Now	5 year (2022)	10 year (2027)
Total CPU (kHS06)	20	50	125
WAN (Gb/s)	4	20	100

- Storage and CPU numbers are example values, scale appropriately for different sites.
- Storage and CPU increasing by 20% year (flat cash assumption).
- We expect network usage to grow slightly faster.
- We expect Nucleus storage to be dual stack by 2019 (End of Run 2).
- In 5 years we recommend 100Gb/s WAN connectivity for Tier 2s.



Backup



Network Monitoring

- ATLAS would like better visibility into our networks to improve our ability to resolve network issues.
- Alerting/alarming on network problems:
 - Large packet loss, low throughput, missing data, infrastructure problems.
- perfSONAR 4.0 campaign
 - Nodes not updating, services not configured properly, services and/or nodes down, problematic test results
- ATLAS would like to see network traffic data from RENs
 - Dashboards to display relevant data together.
- Tools to better identify network monitoring infrastructure problems.



Network Monitoring

- In-ompl-te m-nitor-ng m-kes it si-nif-cantly h-rder -o intr-pr-t th- r-sul--.



Orange = no data!

