

The March with the Clouds

By Brian Dobbins

The March Toward the Clouds: MF Perspectives

March 2, 2022

CI Compass Cyberinfrastructure for NSF Major Facilities Workshop

With
The March To ~~the~~ the Clouds: NCAR

Brian Dobbins
CI4MF 2022
March 1st, 2022



The National Center for Atmospheric Research


Supercomputers



“Cheyenne”

- 5.34 PF
- 145K cores
- ~40PB

Open-Source Models



CESM
WRF
MPAS-A
+ *others..*

Two of our key cyberinfrastructure focuses – a supercomputing facility, and *leading-edge, open-source, community models.*

Cloud Cost Comparison

On-Demand Cloud Pricing (2022)

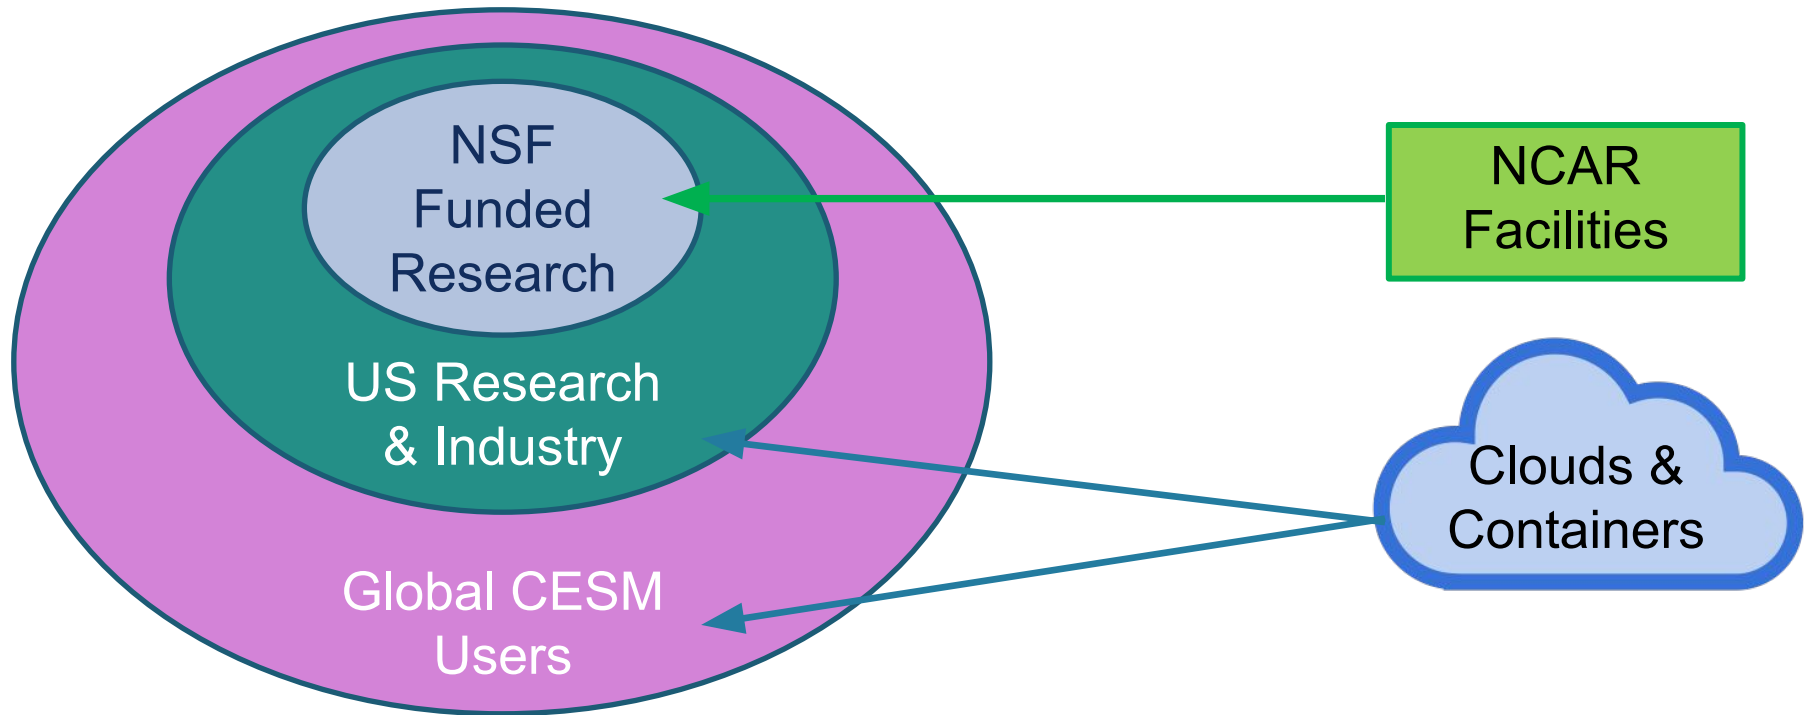
Node	AWS C5N	AWS HPC6a	Azure HBv3
Price (per core-hour)	0.11	0.03	0.03

NCAR “Derecho” (2022)

Cost	~\$35M-\$40M
Power / Year	~\$1.25M
Lifetime	~5 years
Cores	~320K

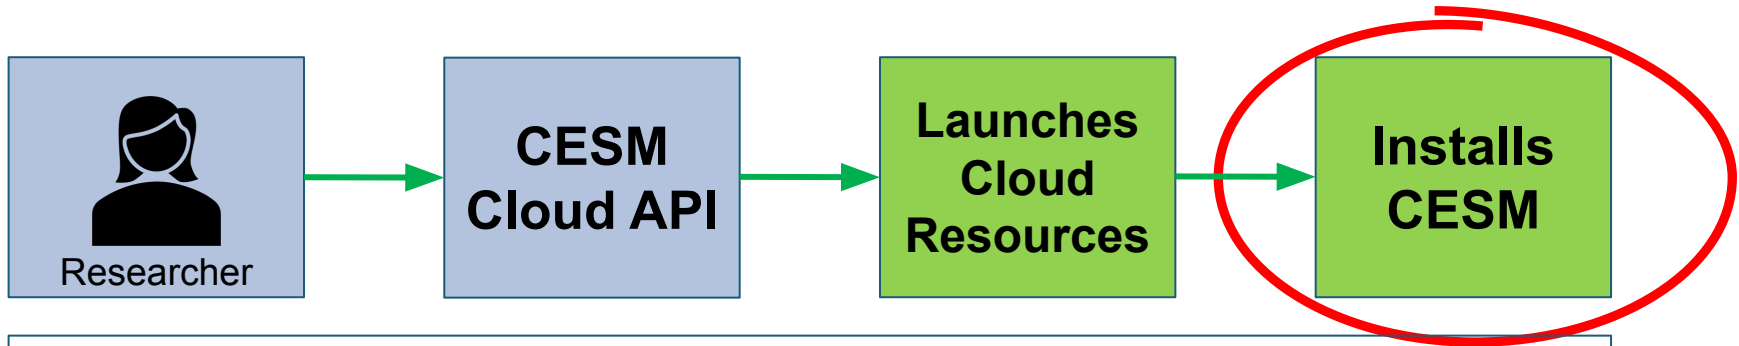
Derecho Cost (per core-hour): **~\$0.003**

Supporting Users Communities via the Cloud



Our facilities serve a *subset* of our users; cloud technologies enable the *rest* of that community *easy access to ready-to-use* modeling capabilities.

Cloud & Container 'Platforms'



The CESM Cloud API actually does two things:

- It handles the configuration and launching of cloud resources
- And *installs a preconfigured, ready-to-run CESM environment*

This full environment, or 'platform', is *standardized*:

- Enables easier training & educational materials
- Enables *shared scientific workflows*
- Improves support ('snowflake HPC systems')
- Containers bring these benefits to owned (non-cloud) systems

Summary on Model Enablement (NCAR)

- Cloud is *complementary* to our on-prem resources, not a replacement for them.
- Cloud *enables* widespread use of scientific computing resources... but complexity is a barrier without APIs / gateways.
- The opportunity to *standardize* science platforms on clouds *and* containers is transformative – let scientists do science, not IT!