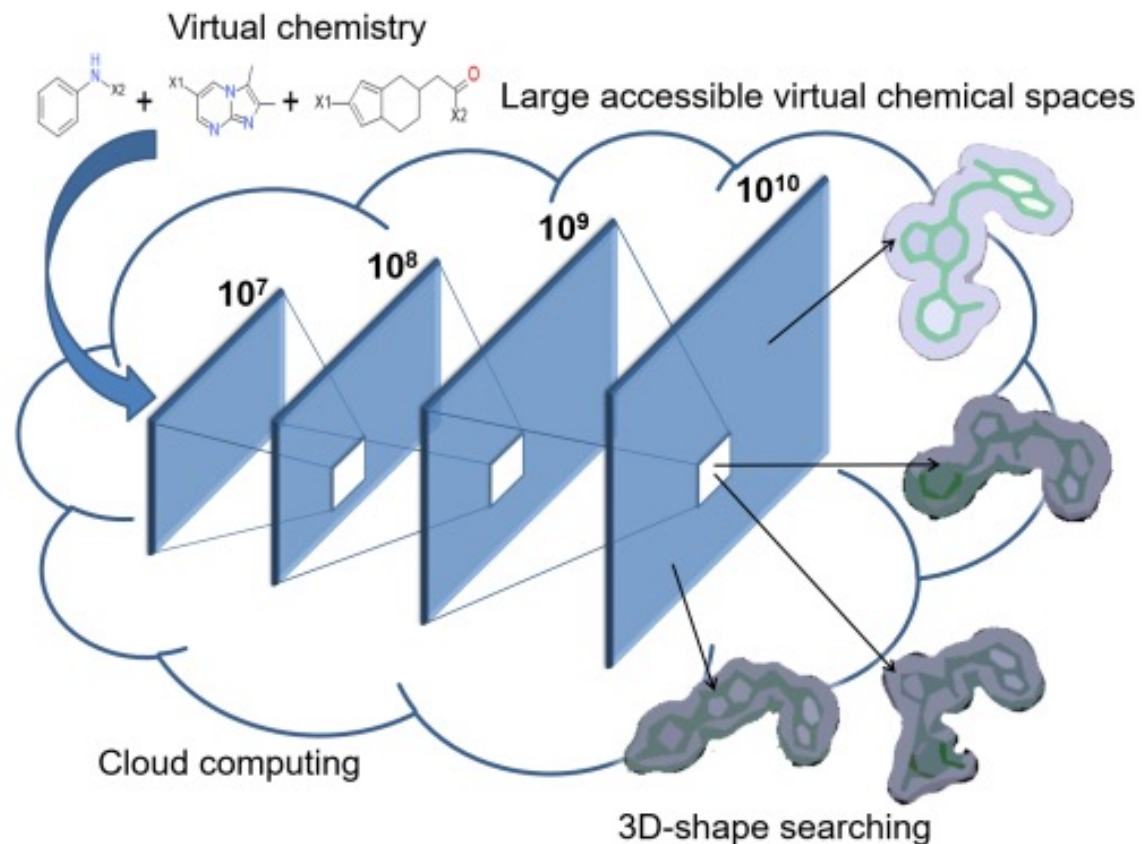


# Advances in Orion HPC

Andrew Shewmaker



# Tomorrow's Big Experiments Become Common

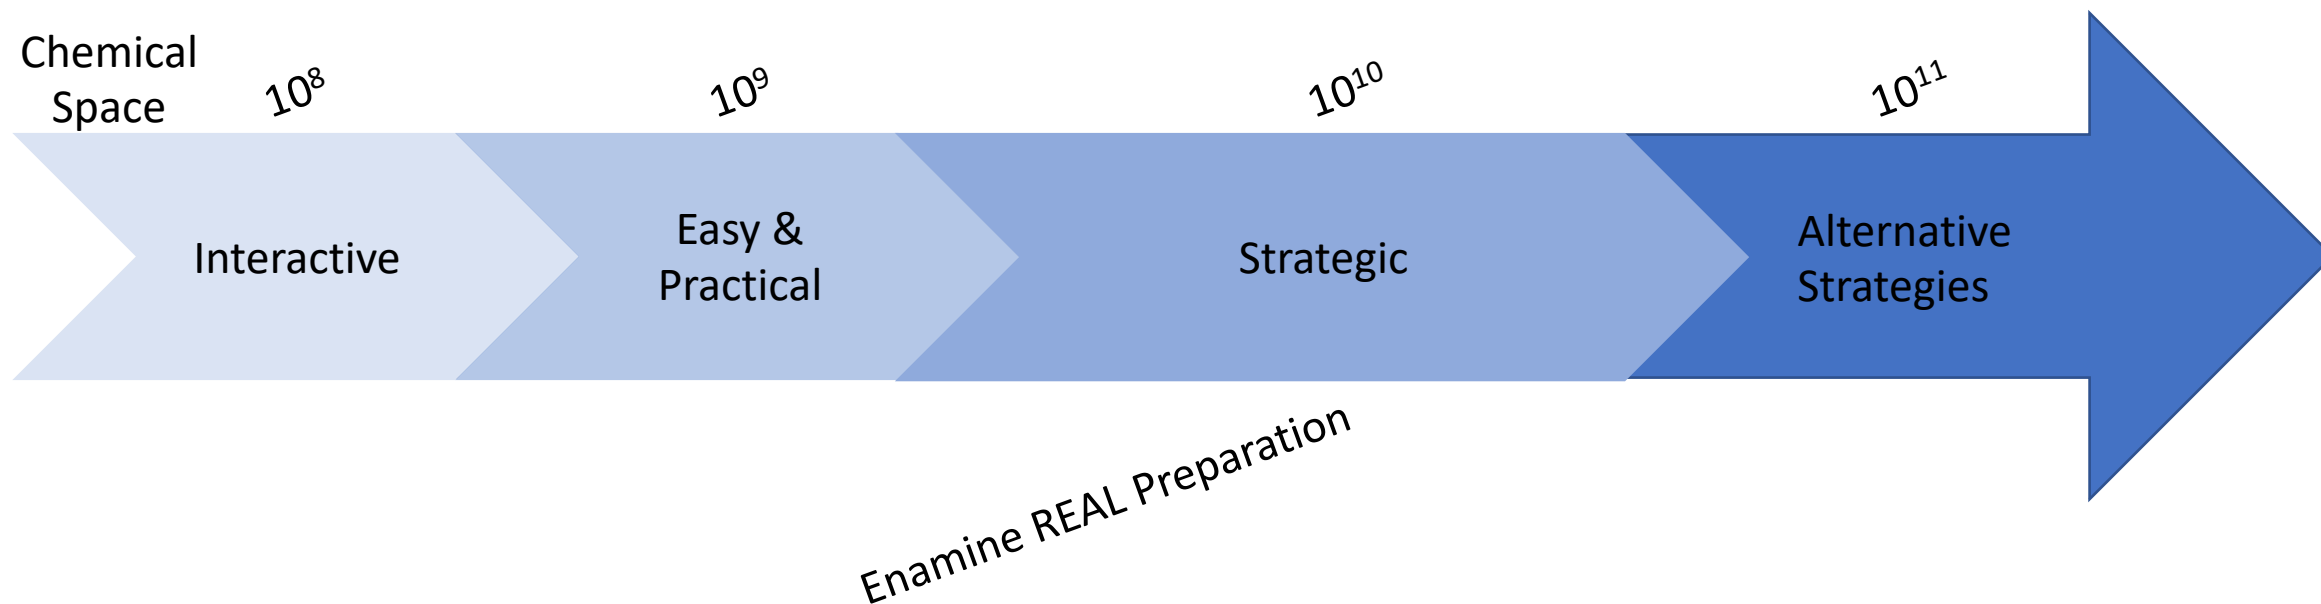


- $10^9$  molecule jobs are now easy and practical
- Next, make  $10^{12}$  molecule jobs possible so that  $10^{10}$  jobs can become common
- Brute force complements smarter algorithms

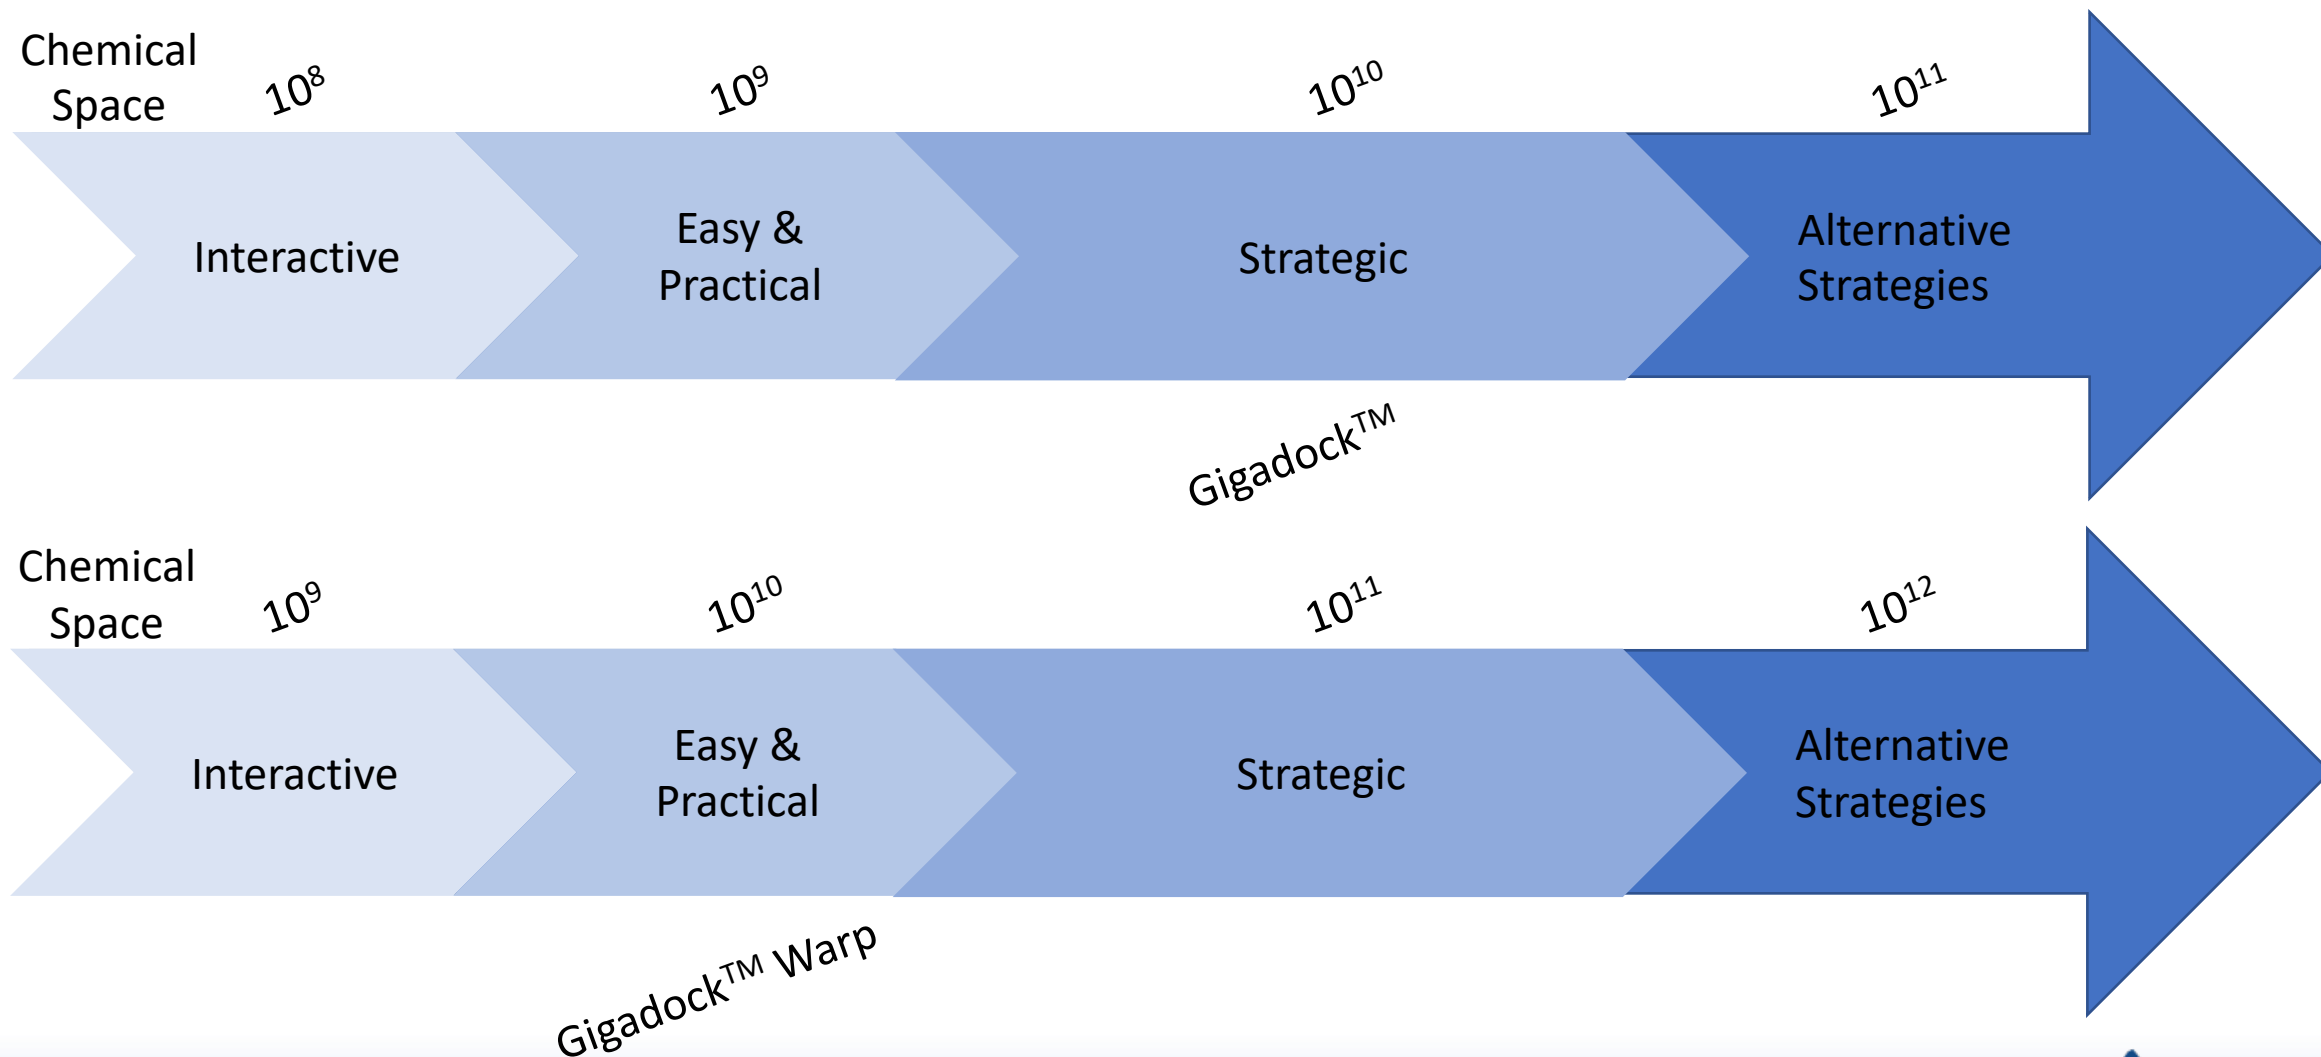
Grebner, C., Malmerberg, E., Shewmaker, A., Batista, J., Nicholls, A., & Sadowski, J. (2019). Virtual screening in the cloud: how big is big enough?. *Journal of Chemical Information and Modeling*, 60(9), 4274-4282.

# Today's Strategic, Tomorrow's Easy & Practical

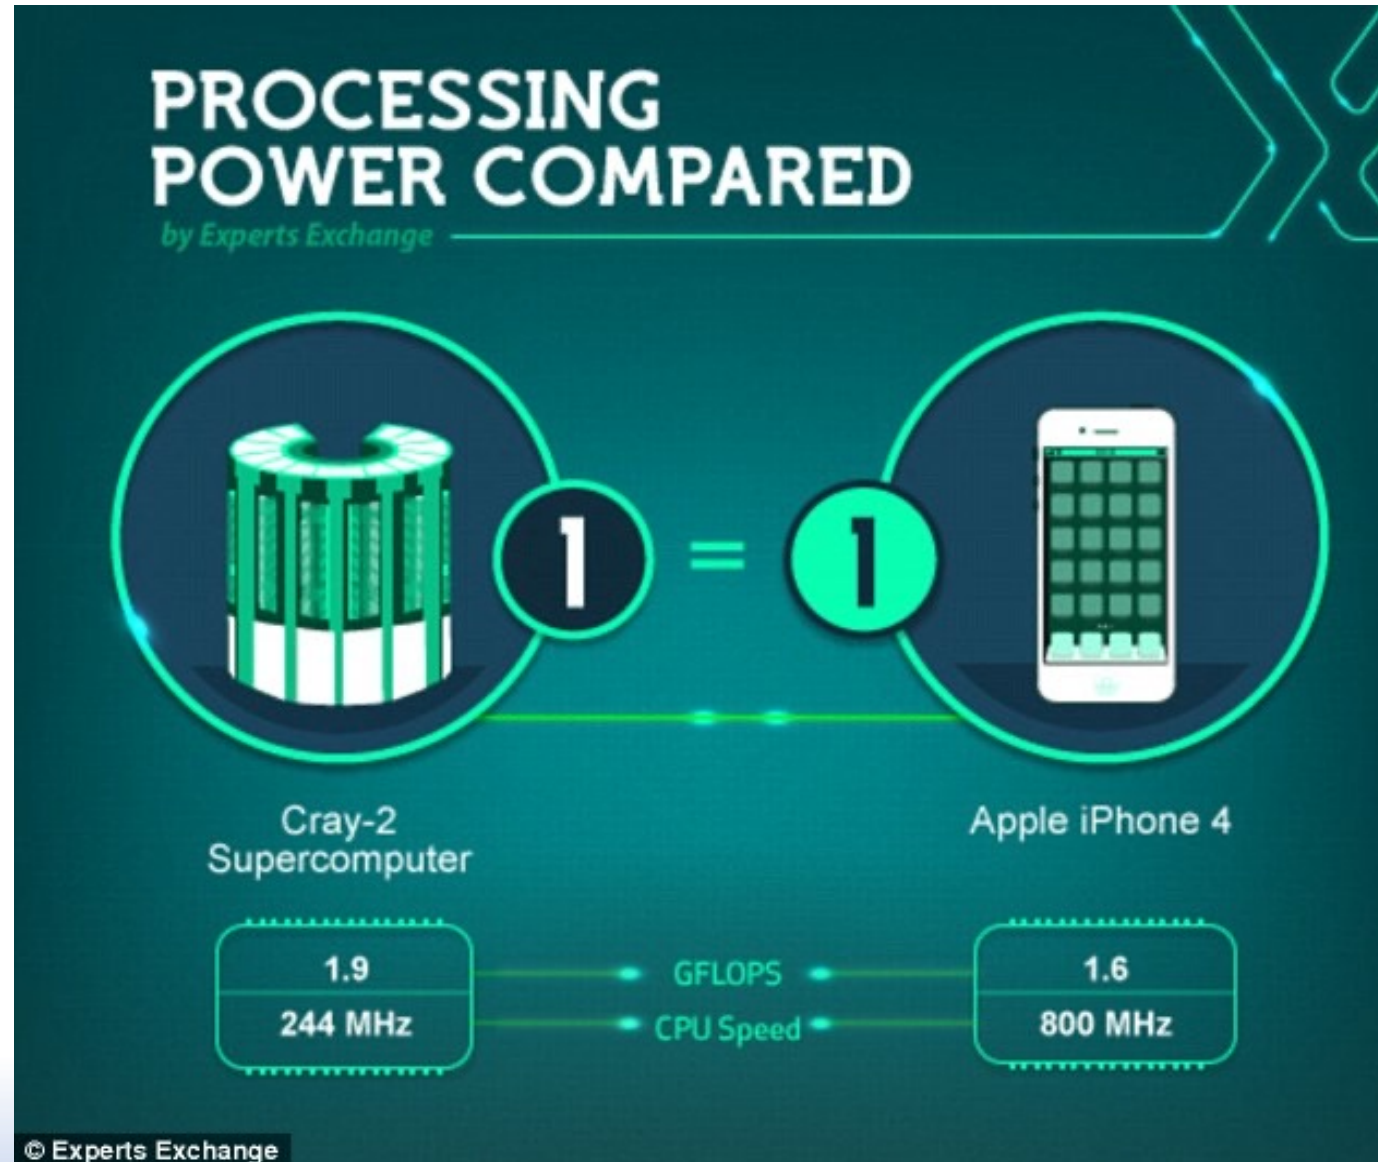
Orion HPC provides quantity to complement scientific quality



# Today's Strategic, Tomorrow's Easy & Practical



# Consider the Supercomputer in your pocket






Image taken from Minnesota StarTribune

If you were plowing a field,  
which would you rather use?


... 2 strong oxen or 1024  
chickens?

Seymour Cray, 1996-ish



# Orion Harnesses Oxen and Chickens\*

Why not use whichever proves effective?



\* Ask David LeBard about his thoughts on the subject





# How do we make HPC more effective?

- Reliability
- Scalability
- Cost Control
- Developer Experience
- Data Organization



# Orion Is Now Even More Reliable

Orion has always retried parallel work

Now Orion actively monitors for runtime/hardware errors, draining work from problematic instances

- AWOL drivers
- GPU double bit errors
- Internal Docker bugs




# Orion Is Now Even More Scalable

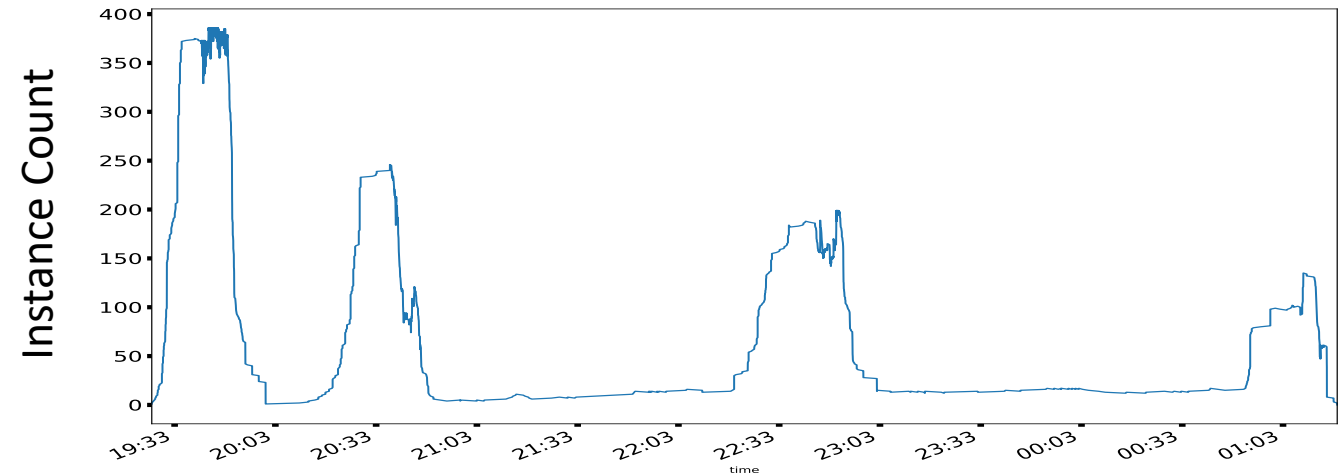
Orion's original scaling algorithm worked well for a long time, but needed updated as certain instance types became less available

- Aggressively suspend groups with lack of instances
- Distribute scale-up across multiple groups
  - 60% of increase is assigned to the best group
  - Remaining distributed among groups of similar cost

# Previous Scaling for NES



**Good AWS day**




**Bad AWS day**

- 100's of GPU spot instances lost and regained in a “bad day” run
- But they **all finished without failures!**



# January 21, 2022 - A Good Day On AWS

Peak of 1015 GPU instances, made up of g5, g4, and g3



# February 25, 2022 - A Less Good Day On AWS

Peak of 564 GPU instances, made up of g5, g4, and g3

367 g3.4xlarge

206 g5.4xlarge


28 g3.8xlarge

21 g5.16xlarge

5 g5.8xlarge

1 g4dn.4xlarge

66 instances taken away  
990 instances not given

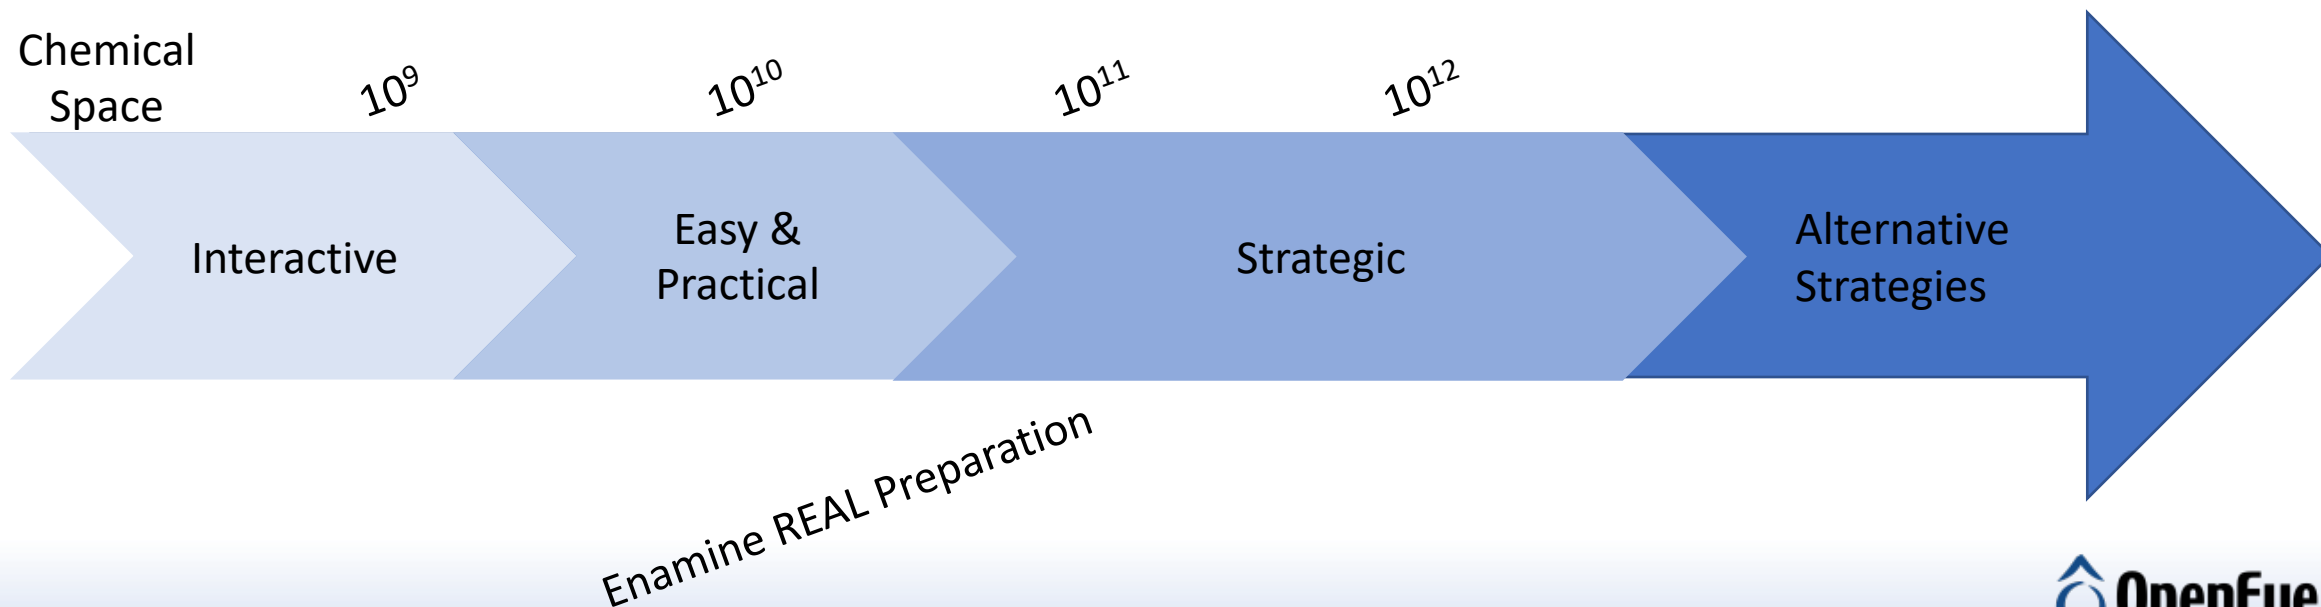


# Orion Scalability In The Future

Orion currently handles jobs with millions of objects in S3, but we intend to scale to billions


Future work includes:

- Collections metadata
- Scheduler metadata





# Orion Cost Control




# Orion Now Includes More Development Tools

- Build requirement-based floe packages with `ocli`
  - No additional cookiecutter package necessary
- Build your own image-based floe package with `ocli`
  - Embed binaries
  - Debug and build ahead of time
  - Incremental builds in the future



# See Resource Usage With Metrics

GMX Bound NES  
is well-behaved



Use metrics to diagnose resource exhaustion.

An API request is made whenever a Cube asks for data managed by Orion. A high rate can slow down your job and even impact other jobs!

For more, check out job debug info. Includes tracebacks.

Close



# Orion Now Organizes Data Better

```
> ocli ls "/organization/OpenEye Data"  
/organization/OpenEye Data  
    FastROCS Collections/  
    Gigadocking Collections/  
    Tutorial Data/  
    SiteHopper Collections/  
    Spruce Data/  
    Generative Design Data/  
    MMDS Data/  
    Brood Data/
```

Don't worry, I'll let Joe demonstrate instead of showing dozens of command line screenshots

# Orion HPC Makes Big Experiments Practical

I look forward to hearing about your big experiments.

**Thank You**

**Special thanks for help with the first  $10^{10}$**

Christoph Grebner

Jharrod LaFon

Jose Batista

Forrest York

Erik Malmerberg

Jay Shankar

Florian Brozek

**Thanks for the rocket chicken**

David LeBard, et al.

**Thanks for the NES example**

Christopher Bayly

Gaetano Calabro

# Today's Strategic, Tomorrow's Easy & Practical

Orion HPC provides quantity to complement scientific quality

