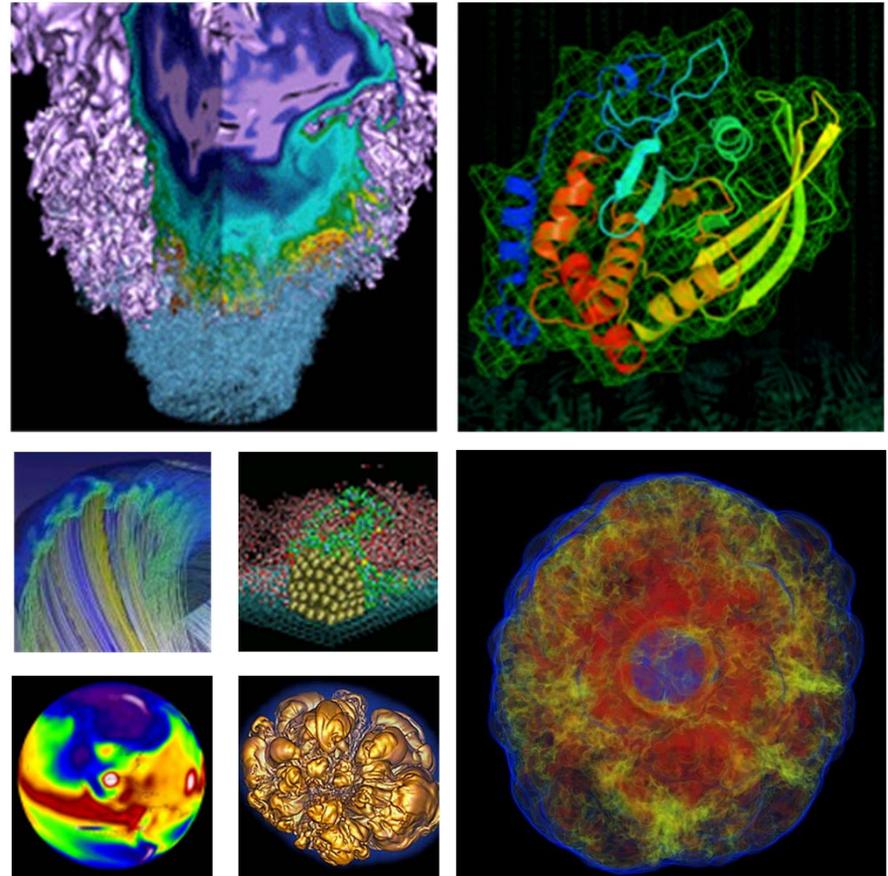


NERSC Spin Project Overview



Cory Snavelly

Lead, Infrastructure Services Group (ISG)
2018-02-07

What is Spin?



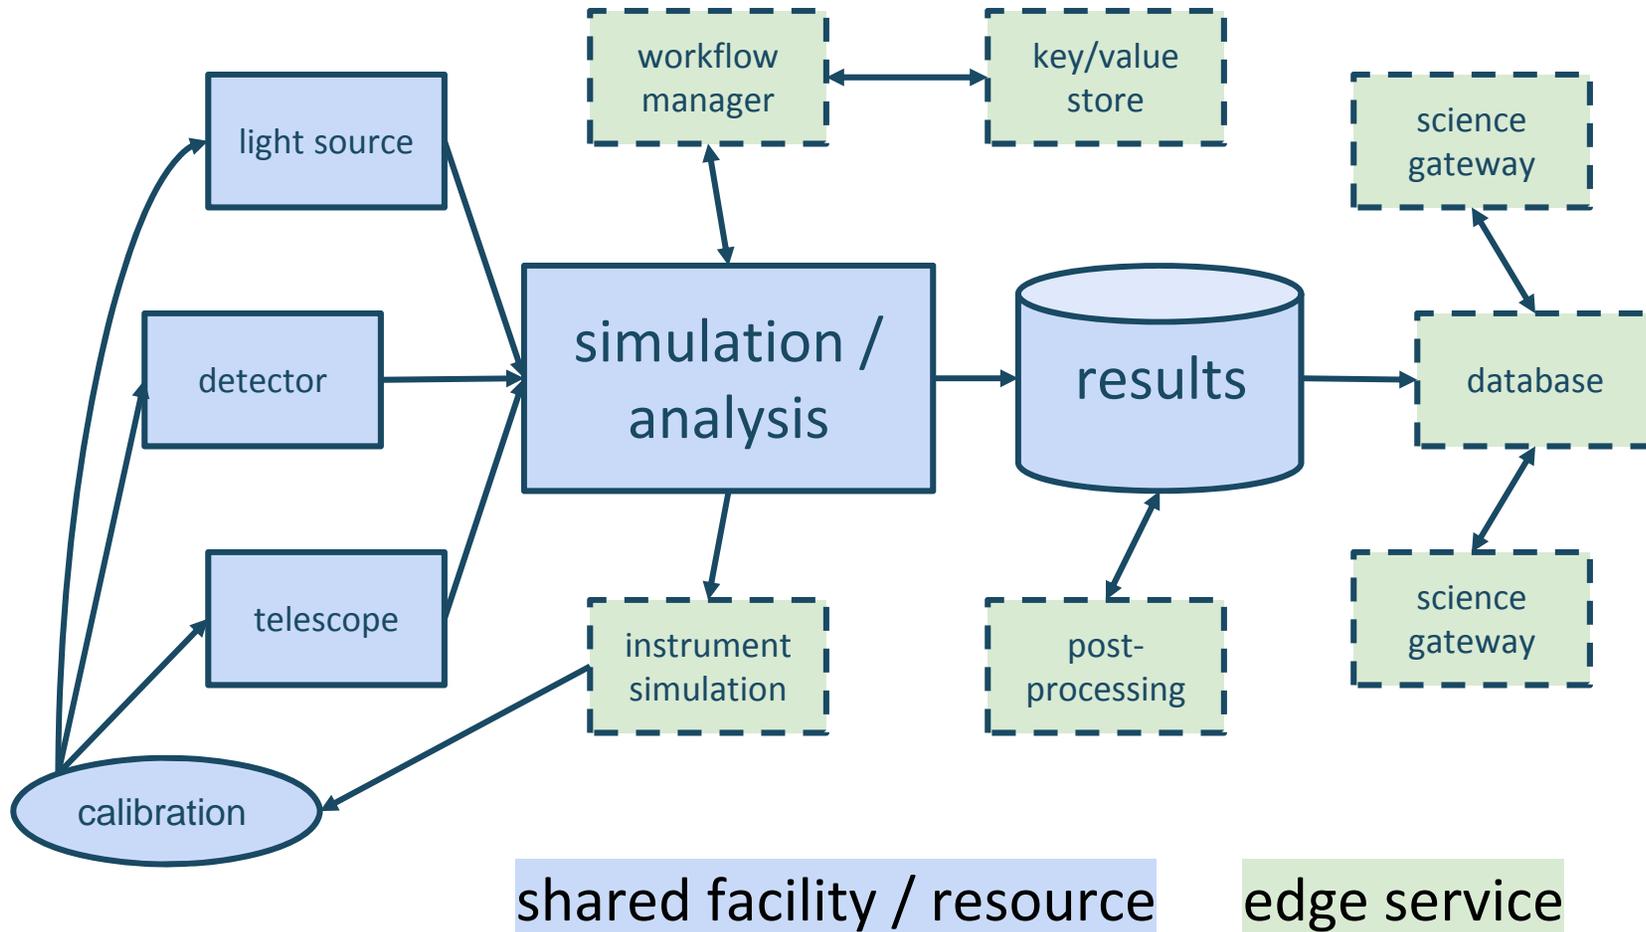
Spin* is a system to build and deploy science gateways, workflow managers, databases, and other “edge services” using Docker containers.

It is designed to be **flexible, scalable, and integrated tightly** with NERSC resources:

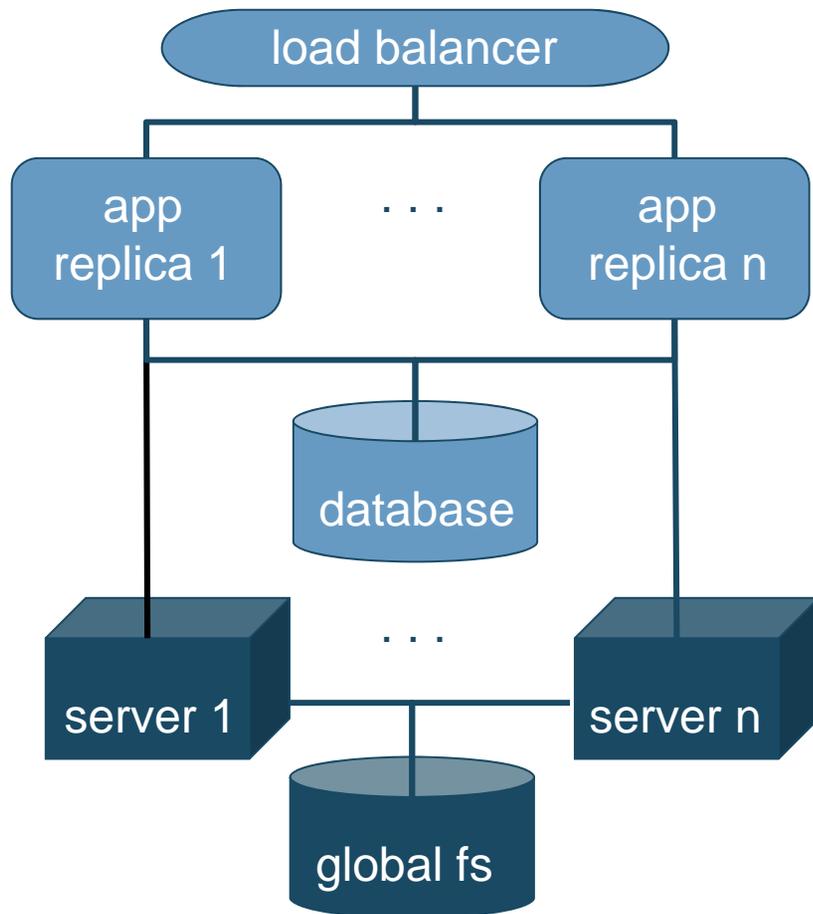
- Develop on your laptop; **deploy in minutes.**
- **Scale out** for performance.
- Access **HPC networks and file systems.**
- NERSC manages everything under the hood.

* Scalable Platform Infrastructure at NERSC

Edge Services need a real home.



What do services look like in Spin?



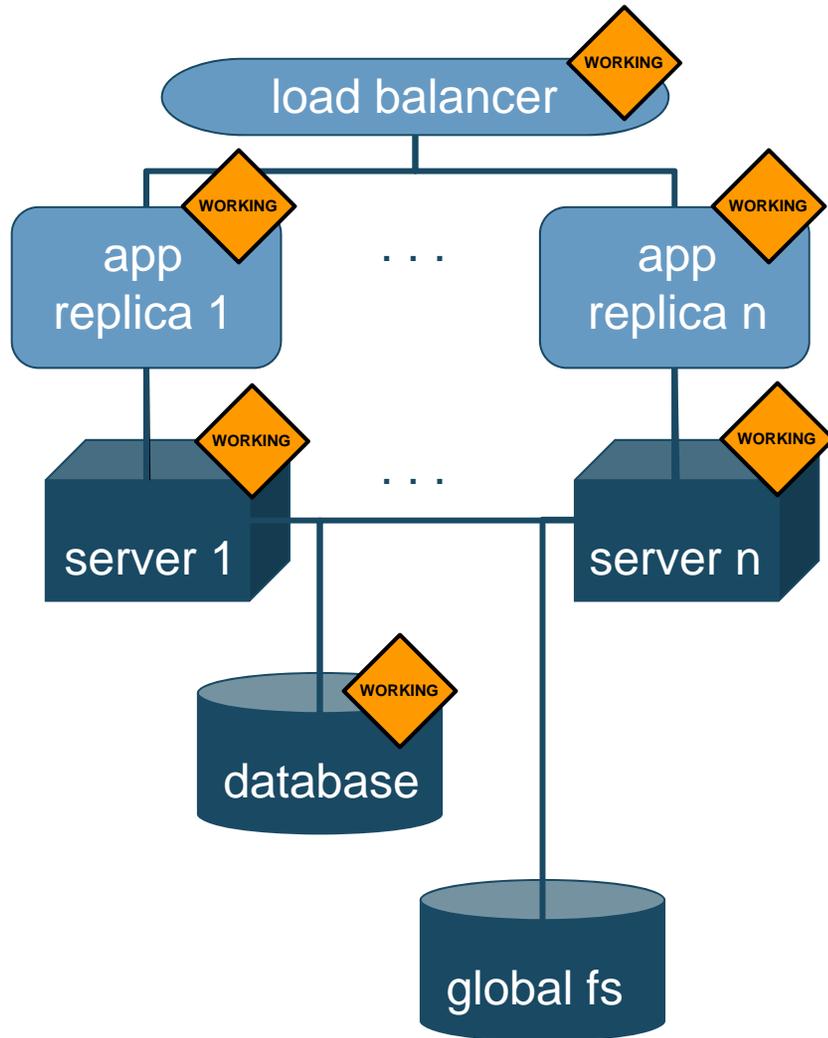
Deploying is easy:

1. A software load balancer is configured automatically.
2. Containerized app is supplied by a user (or NERSC).
3. Start multiple instances to scale performance.
4. Add a dedicated database.
5. A private network for the service is plumbed on demand.

Everything is a container.

Service up in $O(\text{seconds})$.

Whereas, conventionally...



Even with automation, we still:

1. Configure and test the load balancer appliance.
2. Install apps and dependencies, resolve conflicts, configure a startup script, etc.
3. Provision access to a shared database, etc.
4. Change firewall or network configuration as needed.

Service up in $O(\text{days})$ or longer.

Don't despair ... declare!



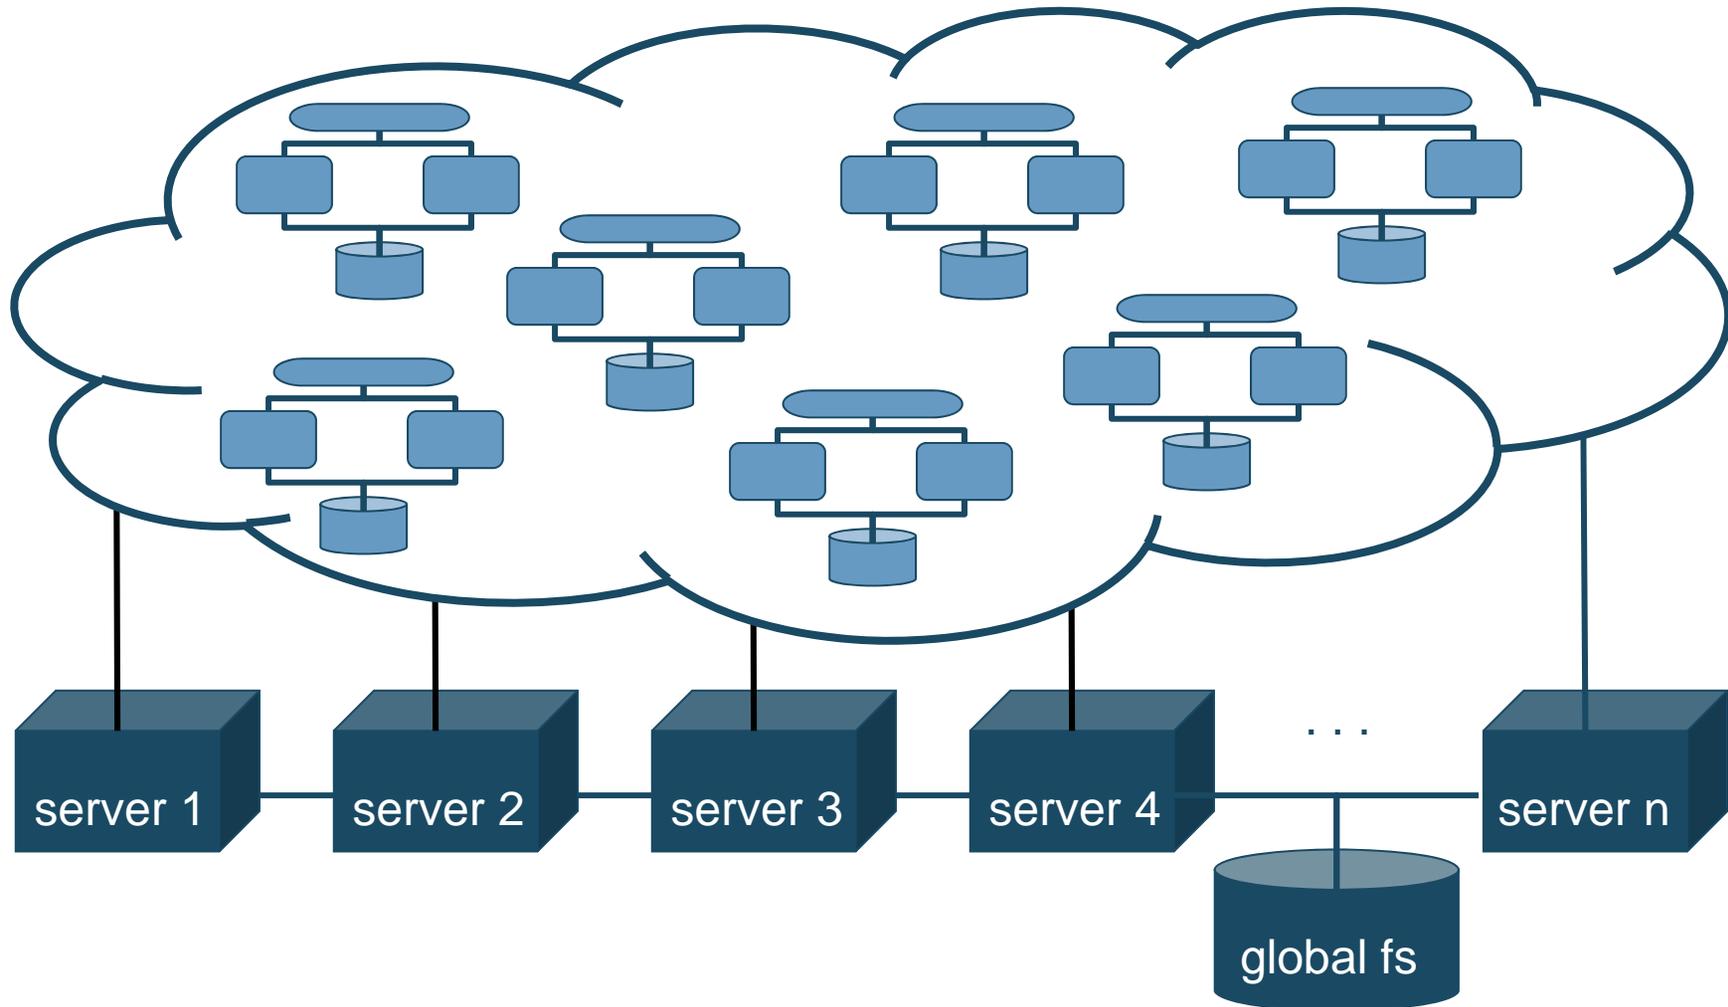
docker-compose.yml

```
version: '2'
services:
  web:
    image: httpd
    tty: true
    links:
      - db:db
  db:
    image: mysql
    environment:
      MYSQL_DATABASE: magic
      MYSQL_USER: cory
      MYSQL_PASSWORD: MyOwnPass
    tty: true
```

rancher-compose.yml

```
version: '2'
services:
  web:
    scale: 2
    start_on_create: true
  db:
    scale: 1
    start_on_create: true
```

“Orchestrating” with Rancher



Technology Gaps



Multi-tenancy is immature. We're building our own

- access controls
...so I can't change (or even see) your services.
- policy enforcement
...so you can't run as root if you access that filesystem.

User namespaces break assumptions for

- filesystem access
...because UID 47284 inside might be somebody else.
- remote shell access
...because you can't trust a container like you trust .shosts.

Hence ... more access controls!

What do our users get?



- Focus on developing an app (or reusing a prepared one).
- Ignore the underlying infrastructure.
- Get scalability and fault tolerance for free.
- Get more cores and memory on larger servers.
- Increase utilization and spend less on equipment.
- Free beer and world peace.

The benefits of VMs without the overheads.
The benefits of cloud with tight integration to NERSC.

What does NERSC get?



- Easier to respond to a broader range of (unique) requests.
- Flexible resource allocation; no more islanded capacity here and resource exhaustion there.
- Monitor the entire infrastructure holistically.
- Non-disruptive upgrades and periodic maintenance.
- Better security with each service isolated in containers and private networks.

We're using Spin internally for our services, too!

"Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Networking and Information Technology Research and Development Program."

The Networking and Information Technology Research and Development
(NITRD) Program

Mailing Address: NCO/NITRD, 2415 Eisenhower Avenue, Alexandria, VA 22314

Physical Address: 490 L'Enfant Plaza SW, Suite 8001, Washington, DC 20024, USA Tel: 202-459-9674,
Fax: 202-459-9673, Email: nco@nitrd.gov, Website: <https://www.nitrd.gov>

