

# Supercomputing, Simplified.

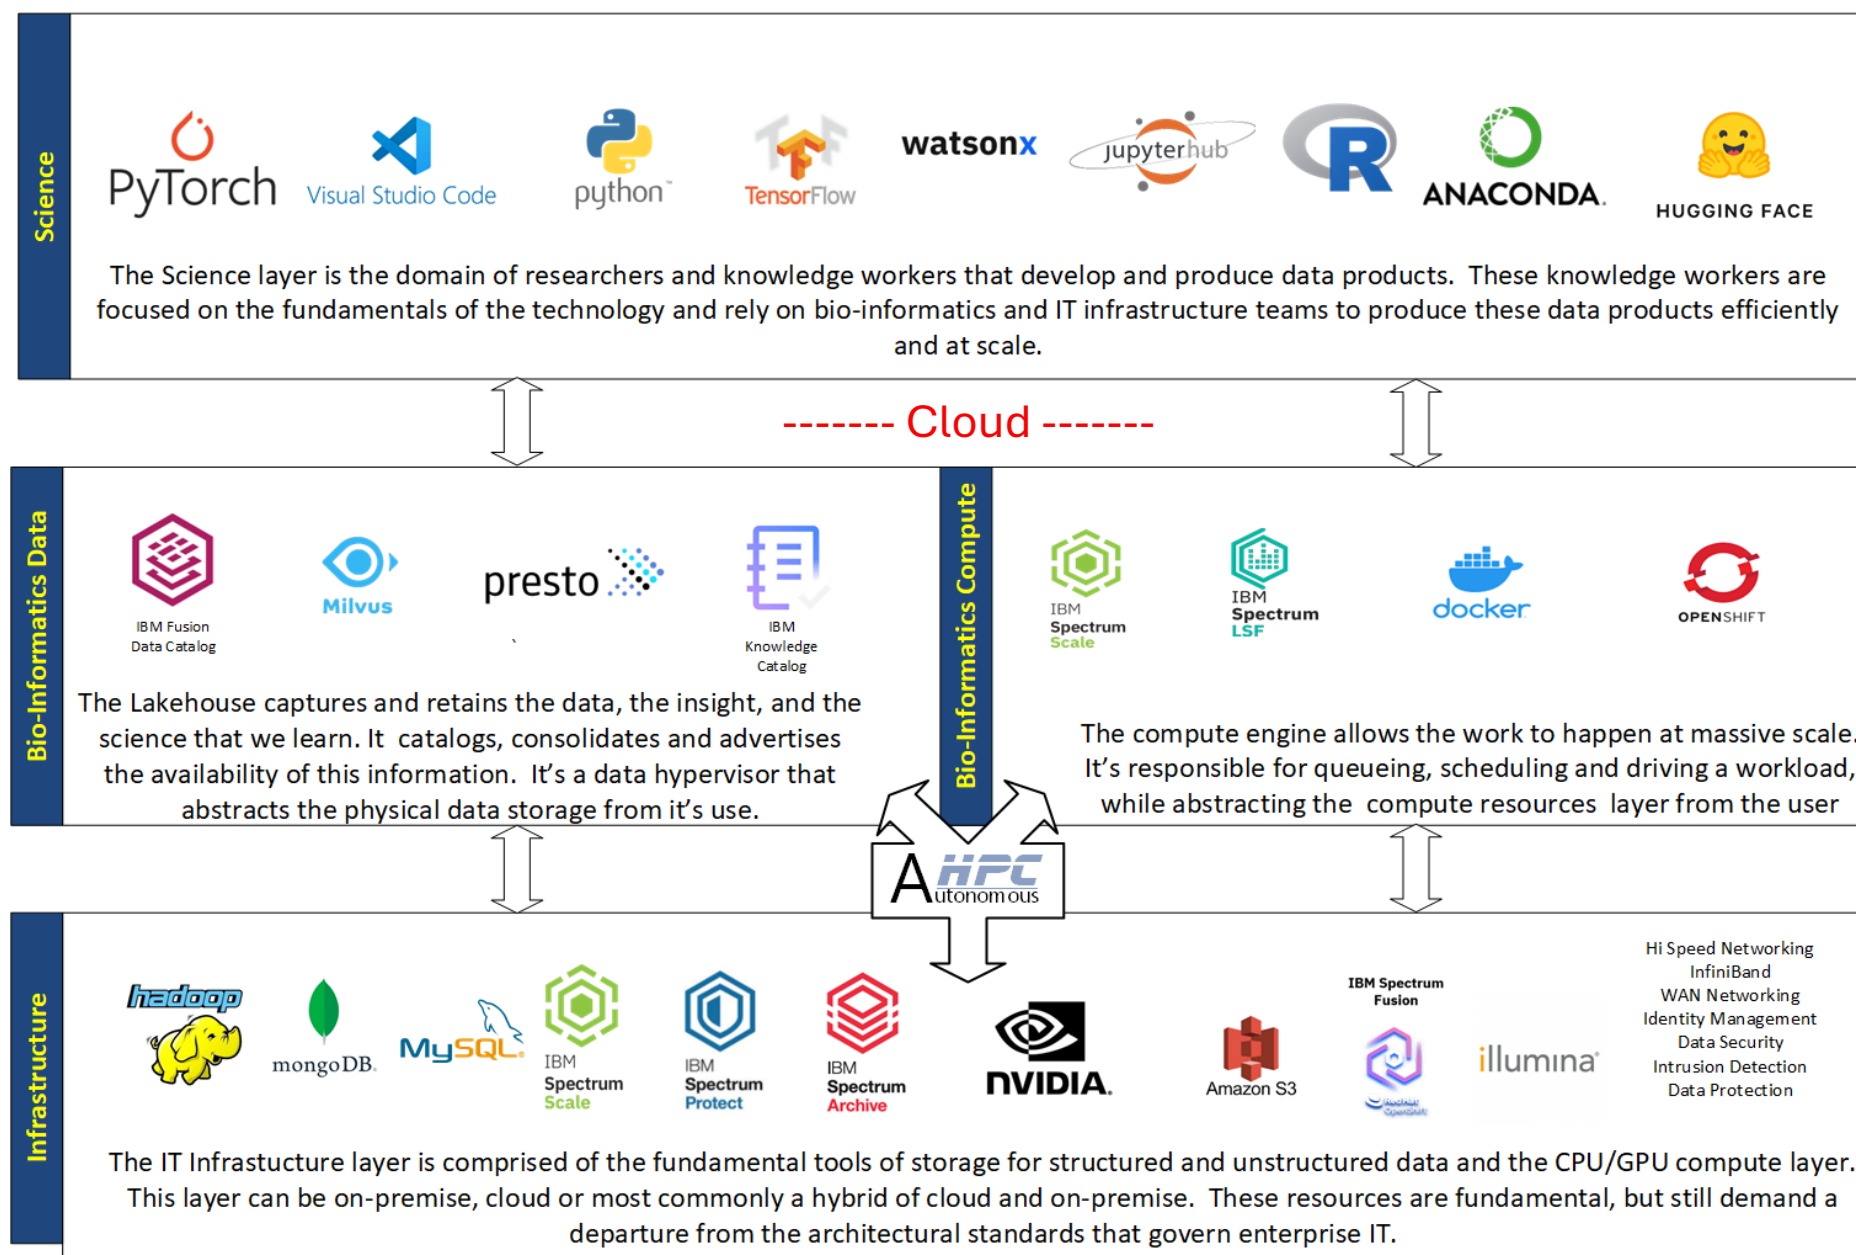
---

# Where does the simplicity leak out?

Simplicity isn't avoiding or masking complexity...

...it's ensuring the availability of resources to manage it

---



Re-Store doesn't tell scientists HOW to do science, we tell them WHERE to do science.

The Bio-Informatics layer aggregates, consolidates and abstracts the infrastructure layer from the knowledge workers

AutonomousHPC and AfterCare service helps the Infrastructure be responsive to Bio-Informatics.

# The inherent complexity of supercomputing

- Supercomputers are not static machines — they are living, evolving ecosystems.
- Each system is a unique blend of
  - architecture
  - workload
  - environment
  - user community.
- Their complexity demands more than technical knowledge; it requires a
  - deep, adaptive understanding of how systems behave under pressure
  - how they grow
  - how to anticipate their needs.
- Supporting them calls for a rare fusion of
  - Engineering
  - systems integration
  - performance optimization
  - system-specific information
  - human intuition — a craft honed through experience and constant learning.

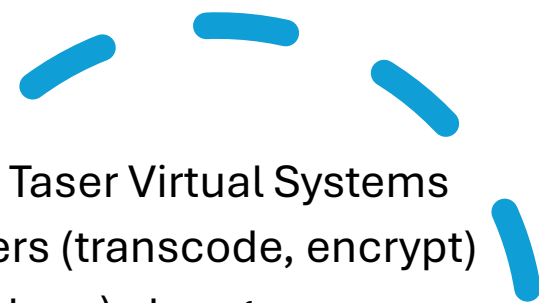
# Re-Store Mission Statement

Re-Store is committed to delivering the deepest skills in the industry to design, sell, install, operate, manage, troubleshoot, and orchestrate Storage Scale-based supercomputers for maximum business value.






# Re-Store history with GPFS

- 
- We did our first project in 2008 with Taser Virtual Systems
    - multiple supercomputer clusters (transcode, encrypt)
    - Massive, 15 PiB, tiered (HDD & tape) storage
  - In 2012, we became IBM's only partner exclusively focused on GPFS
    - The software was surprisingly mature and robust
    - There were very limited IBM resources
    - GPFS was the domain of national labs. No best-practices in enterprise computing.
  - In effect, we were James T Kirk "To boldly go where no man has gone before"



# Lessons Learned from a GPFS Pioneer (2008)

How Re-Store became  
Re-Store



# Lesson #1 – Installation is not integration

If you don't assign an Integrator  
you become the Integrator





## Lesson #2

# Support is Critical, but Not Complete

Vendor support is designed to provide post-mortem break/fix on their component.

The issue is often between all of islands of vendor support



# AfterCare

---


AfterCare is comprehensive set of practices, strategies, training, mentoring and services implemented after the initial installation/integration. It is the critical phase that bridges the gap between the system design and the long-term alignment with business goals

- ecosystem & technology focused – not vendor specific or product focused
- strives to provide optimal, reliable operation that's aligned to the business
- born to close the gaps between all of the vendors, hardware, software network, and to supplement customer skills.
- proactive, hands-on, contextual & personal approach to prevent or predict those issues
- does not overlap vendor support. It's a liaison into the vendor support systems and to coordinate multiple vendors.
- custom and ranges from an advisory role to a complete outsource of supercomputer operation


## Lesson #3

# System data is critical

We cannot manage a system unless we first understand it.



**AHFC**  
autonomous



Monitor

Optimize

Automate

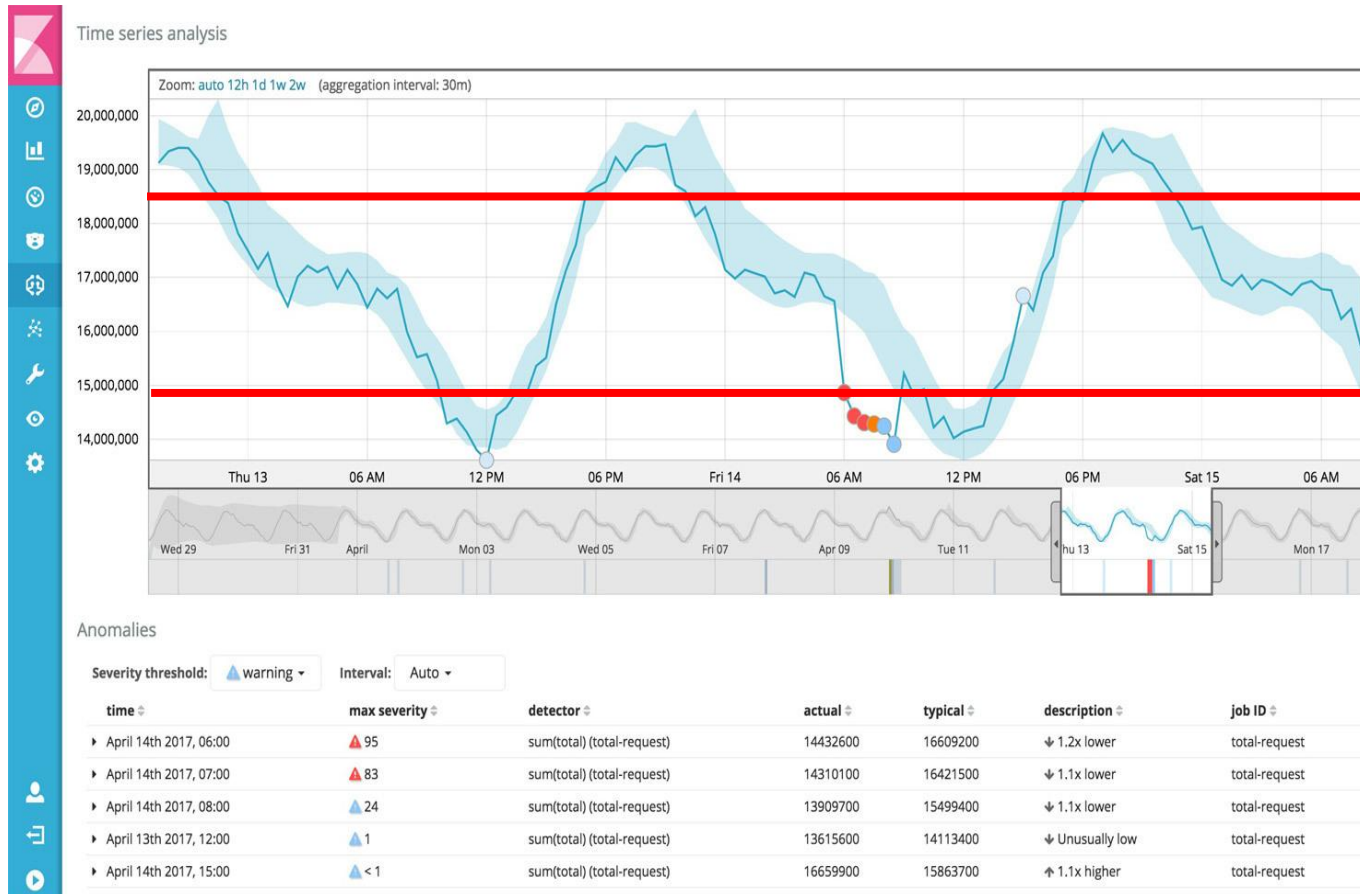
# Every installation has monitoring

## That doesn't solve the problem.



- Monitoring collects telemetry
  - Speeds
  - Feeds
  - Capacity
- 
- Focuses on data typical to
  - Enterprise computing
- 
- Raises alerts on static thresholds


# Add: Machine learning




- Standard 80/20 monitoring rules don't provide the visibility required to manage dynamic systems. In this case, the system would trigger min and max alerts each day, even though the system is performing normally.

# Correlate Events with Telemetry with Artificial Intelligence

- Log entries describe discrete events
- A small customer cluster with consolidated syslogging sends over 1,000,000 event log entries per hour!
- A human being cannot derive meaningful insight from that volume of data.
- AutoHPC is leveraging IBM's Granite model to distill actionable insight from the logs and fuse it with the observed telemetry.



Granite

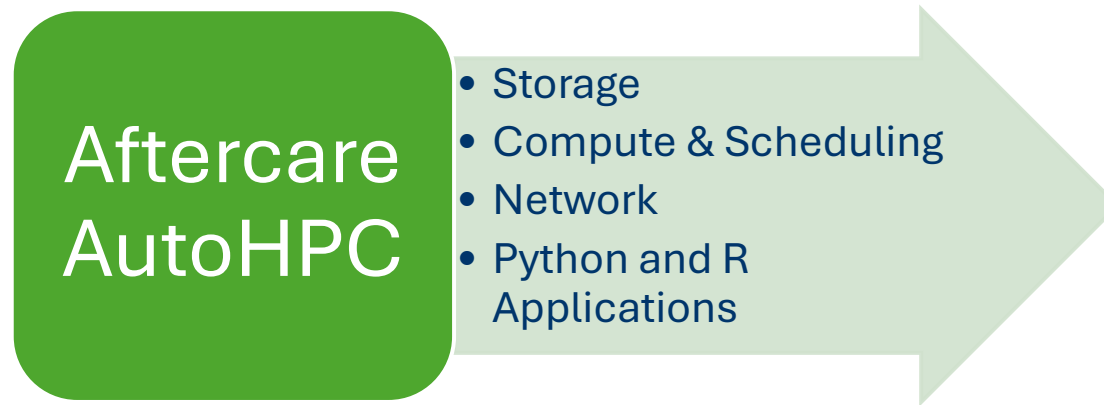




# Demonstrable results with AutonomousHPC and AfterCare at Guardant Healthcare

## All systems tuned and operational

- ✓ Spectrum Scale
- ✓ InfiniBand
- ✓ Sun Grid Engine
- ✓ Applications
- ✓ Compute



Pipeline runtime  
From 32 hours  
To 8 hours

# Benefits of Optimization



---

Better Patient Outcomes

---

Customer Satisfaction

---

Competitive Edge in Marketplace

---

75% reduction capital expenses

---


75% reduction operating expenses

---

Lower administrative costs

---

Improved documentation and reliability



# How to work with Re-Store

Re-Store as your reseller

Re-Store + your reseller

Re-Store + IBM Direct

Michael Sedlmayer  
President, Architect  
[sedl@re-store.net](mailto:sedl@re-store.net)  
858.375.4321

John Aiken  
Dir of Sales  
[john@re-store.net](mailto:john@re-store.net)  
714.402.4911

Bernard Shen  
Dir of Engineering  
[bernard@re-store.net](mailto:bernard@re-store.net)  
310.993.6500

Will AI takeover  
my job?

The new HPC support  
frontier as we see it.

# Yes, but how?

- Real-time determination of root cause, then lead up to our ability to do predictive failure
  - when we can predict failure, then we can prevent it
  - it is hard because every HPC is different and operational pattern and workloads are different
  - real-time tracking of all error and warning events will tell us where problems are brewing
- Fine-grain performance and capacity analysis, and real-time tracking of bottlenecks
  - are all sub-systems running optimally and per vendor spec? if not, why not?
  - what resources are consumed by which type of jobs?
  - is the storage serving data optimally? if not, why not?
  - is my network running optimally? if not, why not?
- HPC problems are not contained within sub-system/vendor boundaries
  - some problems we can anticipate, some problems we don't even know exist
  - need to monitor/track/analyze all discrete events and system telemetry in real-time
  - human experts do not have the ability to track and analyze this 24/7 and in real-time
  - this is the best job for a computer to do

# work-in-progress (example)

