GPFS and Single Name Space
Guardant Health

Nov 17, 2024

What is Guardant Health?

Our mission is to ***conquer cancer with data***. It's an ambitious pursuit. It requires a singular focus, the contributions of many, small steps and bold leaps, and a relentless commitment to patients.

Helping Patients at All Stages of Cancer

Preventive Screening
- CRC
- Other cancers in future

Precision Oncology
- Oncology treatment
- Monitoring
- Tissue & Pharma

GUARDANT

Guardant Health Patient and Data Journey

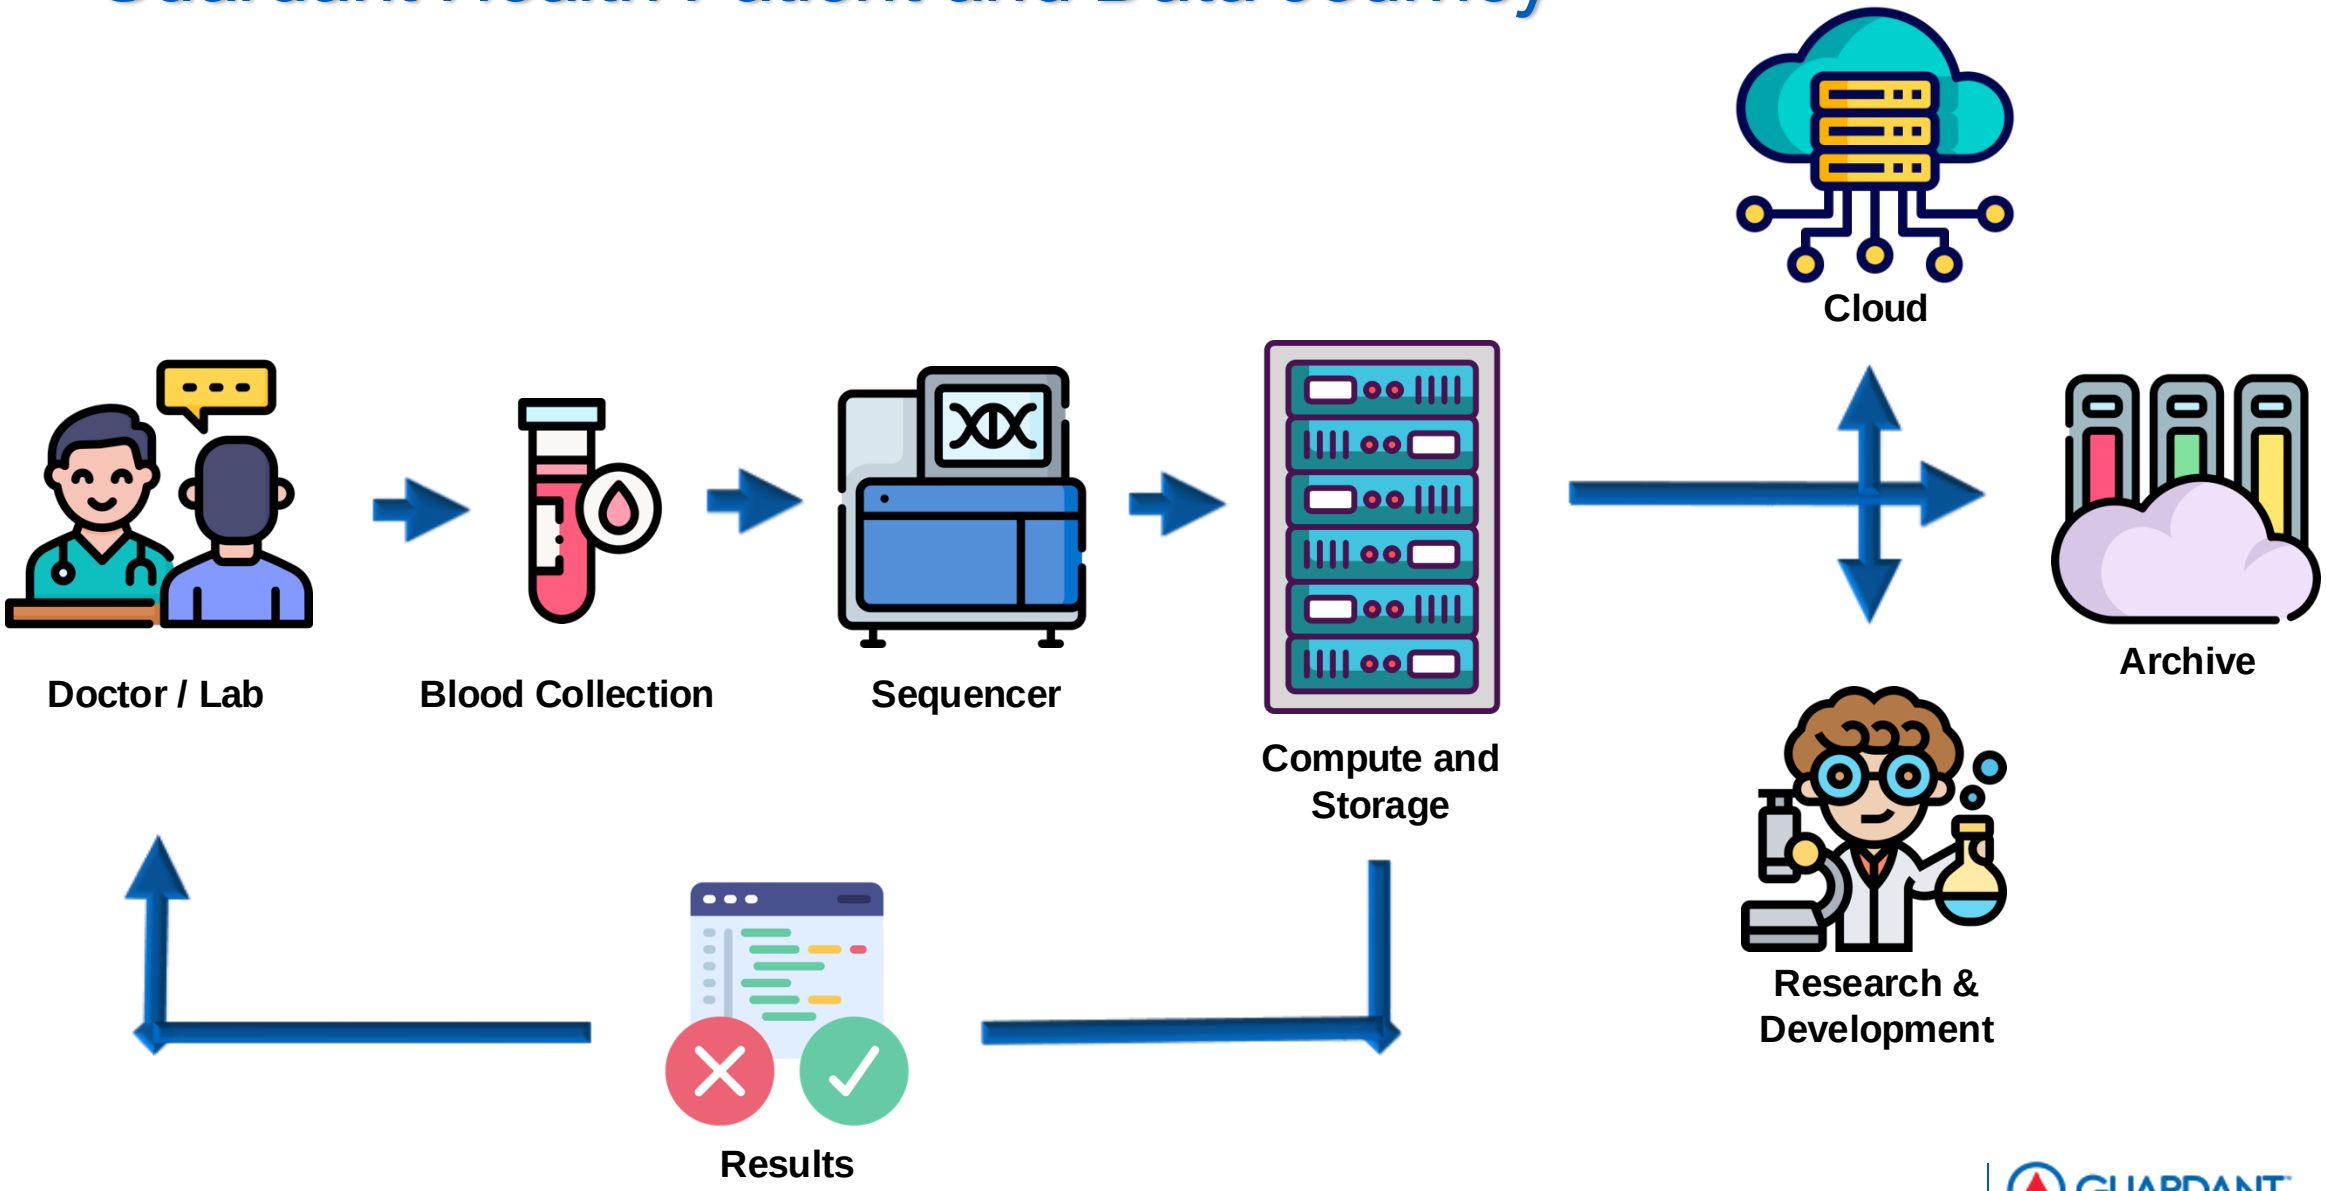

Guardant Health Scaling Issue

8 petabytes

80 petabytes
with 600K+ patients

5 million patients annually
700 petabytes annually?!

Retention at that rate: 30 years … 21 exabytes?!

Well, can't we archive the data? *Um ... about that ...*

An audit found that 80% of the data generated in a three-month period was accessed.
We have a data download problem...

GUARDANT

Restating the Problem

Data is Regularly Accessed: We can't use slow tape storage because our data needs to be available all the time.

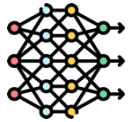

Increased Data Access with ML: Machine learning is making us access more data, even as we scale up.

Massive Storage Needs: Storing **700 petabytes** of data would need about **50 racks** of equipment … that's a lot!

Additional Costs:

Data Center Expenses: Housing all that equipment isn't cheap.
Support Requirements: We'd need ongoing maintenance and support.
Personnel Needs: Managing this can't be done by a few people (and unfortunately, hiring interns on roller skates isn't an option).

Long-Term Storage: We must keep all this data for a long time.

There must be a way...

GUARDANT

Enter Single Name Space

Utilizes AFM cache on object storage and gives bioinformatics and scientists **access to the same data on all sites/clusters (on-prem and/or cloud)**

- Gives the ability to write into a single object store location (using S3 where data is locally cached when accessed, and metadata gets propagated on all clusters automatically)

- Two interesting accessibility modes;

  - Read/Write (IW): *Allows all clusters to read and write to their local caches independently while only pushing changes to other cache sites*
  - Single Write and Read Only: *One cluster writes to cache while the rest access the data as read only*

- Automates data synchronization across sites, removing the need to manually copy data and preventing duplicates

- Presents data as a fileset within the filesystem, appearing as a standard directory for ease of use

- By caching data locally and only transferring changes, data movement over network is reduced (less congestion)

- Cost benefit of moving data to any S3 provider

- **Example of an AFM-CO under filesystem:** /FS/dev-s3 (where "dev-s3" fileset would be home cache mounted as local cache)

GUARDANT

Isolated FS vs Single Namespace

Isolated File Systems

Palo Alto

Compute ⇌ Storage

```
/FS/dev
├── diagnostics
│   ├── molecular23
│   └── molecular24
└── equipment
```

Redwood City

Compute ⇌ Storage

```
/FS/dev
├── applications
│   └── app1
├── db_list
└── app_list
```

Colo

Compute ⇌ Storage

```
/FS/dev
├── projects
│   └── project_alpha
└── reports
```

AWS

Compute ⇌ Storage

```
/FS/dev
├── services
└── environments
```

Single Namespace (AFM-COS IW Mode)

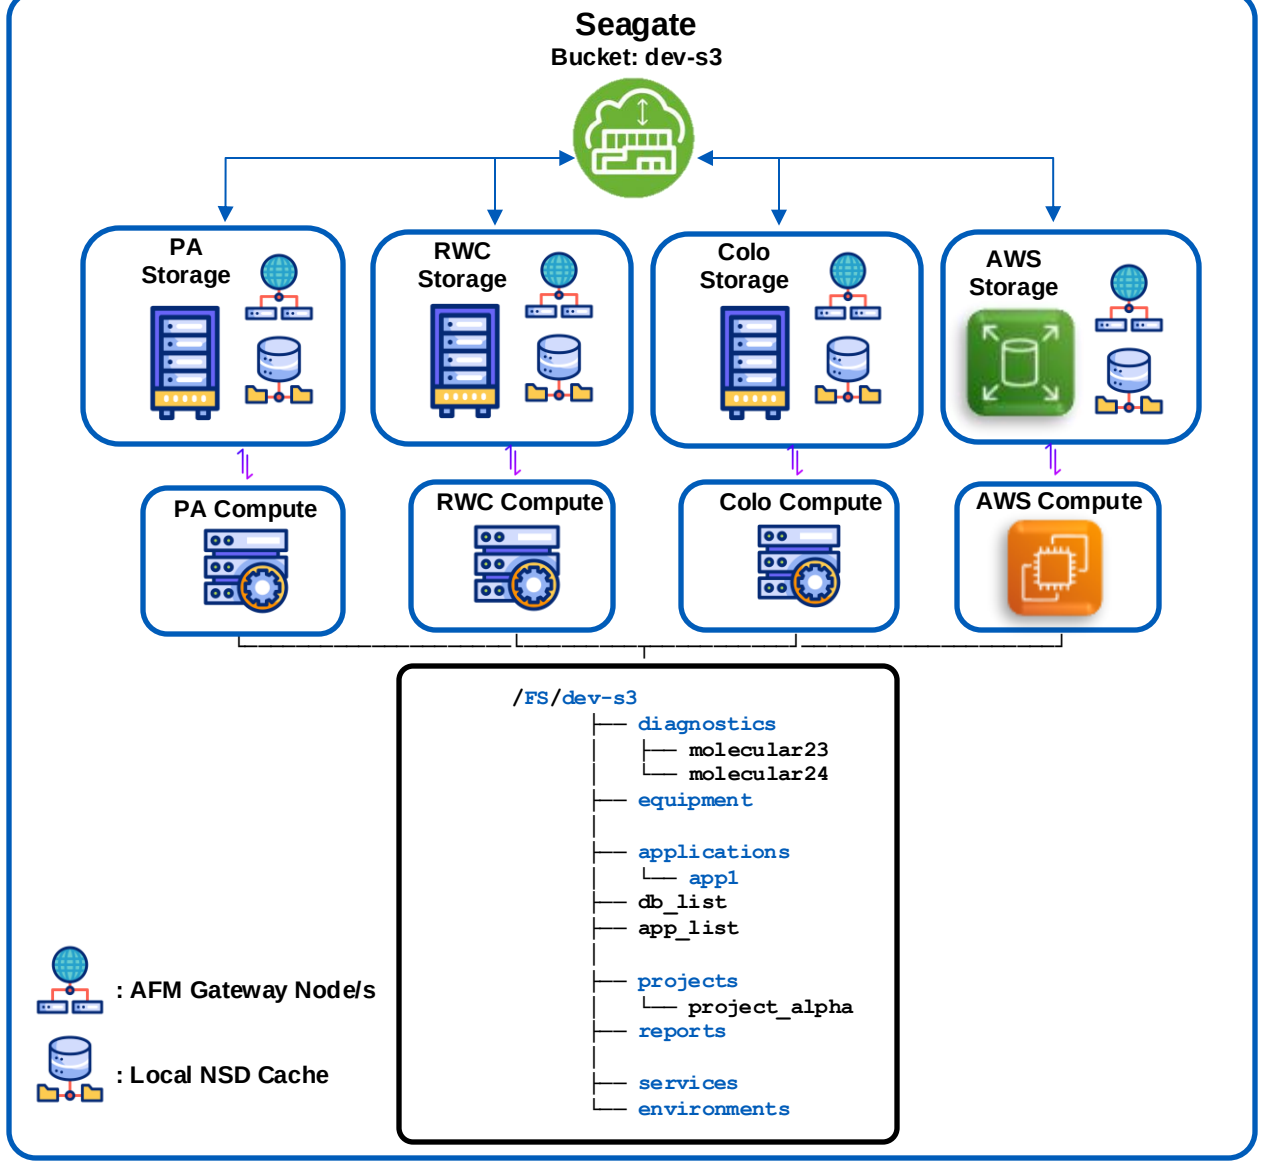

- Without SNS, data from other clusters must be accessed by manually copying files or using tools like rsync between different compute clusters.
- With SNS, the same data is accessible from all clusters.

GUARDANT

Guardant's Prototype

- 2 ½ petabytes in object storage (Uses GPFS with AFM-COS)

- SNS Clusters in;
  - Redwood City
  - Palo Alto
  - Santa Clara
  - San Diego

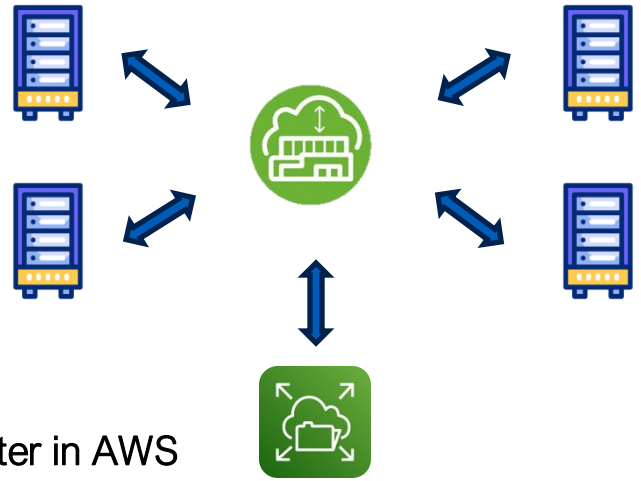

  - We also have a cluster in AWS

- Improved from SC23, but…

IBM Spectrum Storage

PACKETFABRIC

RStor

SEAGATE

aws

ddn

References:

https://www.spectrumscaleug.org/wp-content/uploads/2023/11/SSUG23SC-Guardant-Health-One-Namespace-to-Bind-Them.pdf

https://www.spectrumscaleug.org/wp-content/uploads/2023/11/SSUG23SC-Guardant-Health-On-Demand-Bandwidth-Implementation.pdf

GUARDANT

Limitations

- **Last writer "wins" in IW mode (Read + Write on fileset on all clusters)**
  - Low risk for prod: Any cluster where data originates is has single write permissions; all other SNS clusters would be read only

- **Bandwidth limitations today** (our RWC server rooms for example, have much lower bandwidth than Palo Alto)

- **Dependency on IT Network health**

Observations

- **Provides scalability,** add more buckets on the back end

- **Disaster resilience and recovery**
  - Even when connectivity goes down, you can keep writing data to local cache; all of it will resync automatically once connectivity to the S3 endpoints is re-established on the AFM gateway nodes

- **How fast is it?** In our current environment a 1GB file created on any of the SNS IW filesets becomes accessible on others within 10-15 seconds (dependent on how many files are being written on local cache for AFM to push to the bucket, in some cases files may take longer to appear)

- **Provides parallel data movement (**available in v5.1.9-x**)**
  - Uses round robin, reads and writes using multiple gateways, requires parallelIO mapping (fileset based)

- **Once your data is everywhere, anything can run anywhere …**

GUARDANT

The Single Namespace Working Group

SNS Working Group Goals

Goal is to produce interoperability standards for…

- **Standards for many things**
  - **object formats**
  - **metadata**
  - **APIs**
  - **File system presentation**
  - **Locking**
  - ***more***

- ***Working group represents vendors & end users***

- **Above all else:**
  - **An *Open* standards…**

GUARDANT

GUARDANT