Technical Scope of Work

Keith Chadwick
Scientific Computing Division
Scientific Technical Architecture Group

Work supported by the U.S. Department of Energy under contract No. DE-AC02-07CH11359

Fermilab

Acknowledgements

TSW process and documents are owned by the Service Manager (Tammy Whited) and the Service Level Manager (Brian McKittrick).

- The CS Experment Liaisons,
- Ruth Pordes,
- Margaret Votava,
- and I are working with them on this.

TSW

Technical Scope of Work aka TSW.

- Document formerly known as MOU,
- Redesigned in terms of the SERVICES that you are receiving from the Computing Sector.

Default is that services are provided under the “foundation” Service Level Agreement (SLA):

- CS DocDB #4042:
- http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=4042

Standard Support Hours

The following are default definitions for service support hours to encourage uniformity:

- 8x5 - Monday through Friday, 8am – 5pm U.S. Central Time, not including Fermilab work holidays.
- 12x7 - Every day, 8am – 8pm U.S. Central Time, including Fermilab work holidays. The response time however may be slower on weekends and Fermilab work holidays, which should be clarified by those services offering this support.
- 24x7 - Every day, all of the time. The response time however may be slower on weekends and Fermilab work holidays, which should be clarified by those services offering this support.

Service Response and Resolution Times are impacted by the stated service support hours. An 8 x 5 service with 8 hour response time is in effect promising to respond within 8 business hours (weekdays, 8am – 5pm, non-holidays), not 8 wall-clock hours. A ticket entered for this service on Friday at 2pm may not be responded to until Monday 1:30pm.

Incident Priority Table (taken from the foundation SLA)

Impact / Urgency	**Extensive** Service is out for Enterprise	**Significant** Service is out for many users or degraded for Enterprise	**Moderate** Service is out for 1 user or degraded for many	**Localized** Service is degraded for 1 user
Critical *Based on event*	**Priority - Critical** Respond – 1 H Resolve – 5 H	**Priority - High** Respond – 4 H Resolve – 35 H (1.5 D)	**Priority - Medium** Respond – 8 H Resolve – 97 H (4 D)	**Priority - Medium** Respond – 8 H Resolve – 97 H (4 D)
High *Required*	**Priority - High** Respond – 4 H Resolve – 35 H (1.5 D)	**Priority - High** Respond – 4 H Resolve – 35 H (1.5 D)	**Priority - Medium** Respond – 8 H Resolve – 97 H (4 D)	**Priority - Low** Respond – 8 H Resolve – 172 H (7 D)
Medium *Important*	**Priority - Medium** Respond – 8 H Resolve – 97 H (4 D)	**Priority - Medium** Respond – 8 H Resolve – 97 H (4 D)	**Priority - Medium** Respond – 8 H Resolve – 97 H (4 D)	**Priority - Low** Respond – 8 H Resolve – 172 H (7 D)
Low *Desirable*	**Priority - Medium** Respond – 8 H Resolve – 97 H (4 D)	**Priority - Low** Respond – 8 H Resolve – 172 H (7 D)	**Priority - Low** Respond – 8 H Resolve – 172 H (7 D)	**Priority - Low** Respond – 8 H Resolve – 172 H (7 D)

Service Availability & Support

It is important to note the difference between service availability and service support.

Consider the case of a typical service:

- The service availability is usually 24x7 – the service is not turned off outside of business hours.

Service Support Level:

- 8x5 – any issues with the service may wait until the next business day.
- 12x7 – any issues with the service may wait until the next calendar day.
- 24x7 – any issues with the service are supported “around the clock” with a typical 1 hour response time during business hours, note that the response time may be slower on nights, weekends and laboratory holidays.

A Comment on 24x7 Support

Very few services are offered with 24x7 service support.

In the case of the Fermilab Campus Grid services:

- The central “shared” services (VOMS, GUMS, SAZ, Squid, MyProxy) are supported 24x7,
- The remainder of the services (gatekeepers, condor, jobsub, ifdh, etc.) are supported 8x5.

TSW - Documents

Sign into sharepoint.fnal.gov with your services account and then navigate:

-> Organizations

-> Computing Sector

-> Service Management

-> ITIL Processes

-> Service Level Management

-> In Progress TSWs

Work is currently proceeding on NOvA, Minerva and microBoone TSWs.

Eventually, the Computing Sector will work together with all experiments and to either update or write their TSW (over the course of the next year).

I'll be contacting each experiment through their corresponding CS Liaison as their TSW comes up.

FIFE Architecture

Keith Chadwick
Scientific Computing Division
Scientific Technical Architecture Group

Work supported by the U.S. Department of Energy under contract No. DE-AC02-07CH11359

Fermilab

Acknowledgements

I'd like to acknowledge the contributions of the FIFE Architecture Committee lead by Stephan Lammel:

- Elizabeth Buckley-Geer, Glenn Cooper, Dave Dykstra, Stuart C. Fuess, Oliver Gutsche, Burt Holzman, Stephan Lammel, Dmitry Litvintsev, Adam Lyon, Igor Mandrichenko, Ray Pasetes, Andy Romero, Stephen Wolbers

The FIFE design is now being stewarded by the FIFE Working Group lead by Mike Kirby.

- I'm serving as the document secretary/editor.

FIFE Architecture Design Report

FIFE = FabrIc for Frontier Experiments

CS DocDB #5180:

- http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=5180

Sign into sharepoint.fnal.gov with your services account and then navigate:

-> Organizations
  -> Computing Sector
  -> Divisions
  -> Scientific Computing Division
  -> Projects

- https://sharepoint.fnal.gov/org/scd/fife/SitePages/Home.aspx

Overall FIFE Architecture Design Considerations

Network Centric

- an Internet Protocol based network binds the components into a system and thus equalizes naturally on-site and off-site activities;

Modular

- switching the implementation of a component will keep the overall architecture intact and thus bring longevity across technology updates;

Scalable

- allows data/CPU intensive components to be implemented as distributed services;

Fault Stable

- components have minimal inter-dependency and maximum isolation so failures will stay contained and not escalate to unrelated components/services.

High-level user view of the components grouped into interactive computing, batch computing, storage, software and database groups

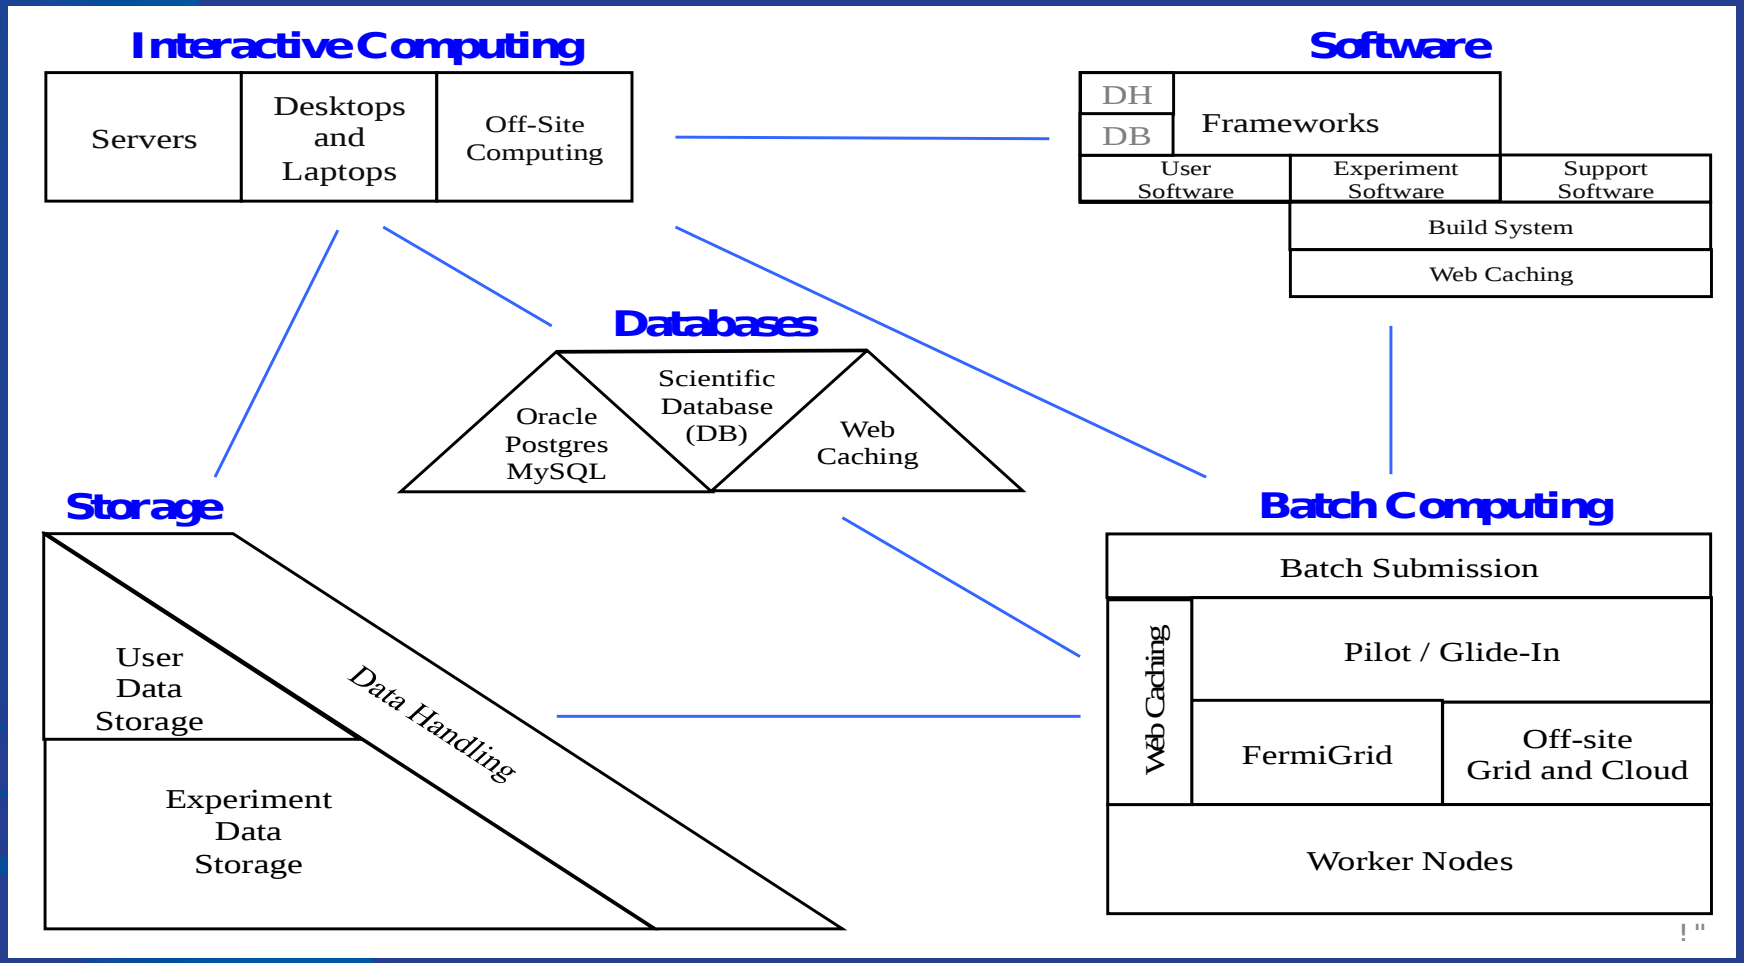

Components of the FIFE architecture organized in layers and with relations of the high-level components shown. Components depend only on components in lower layers

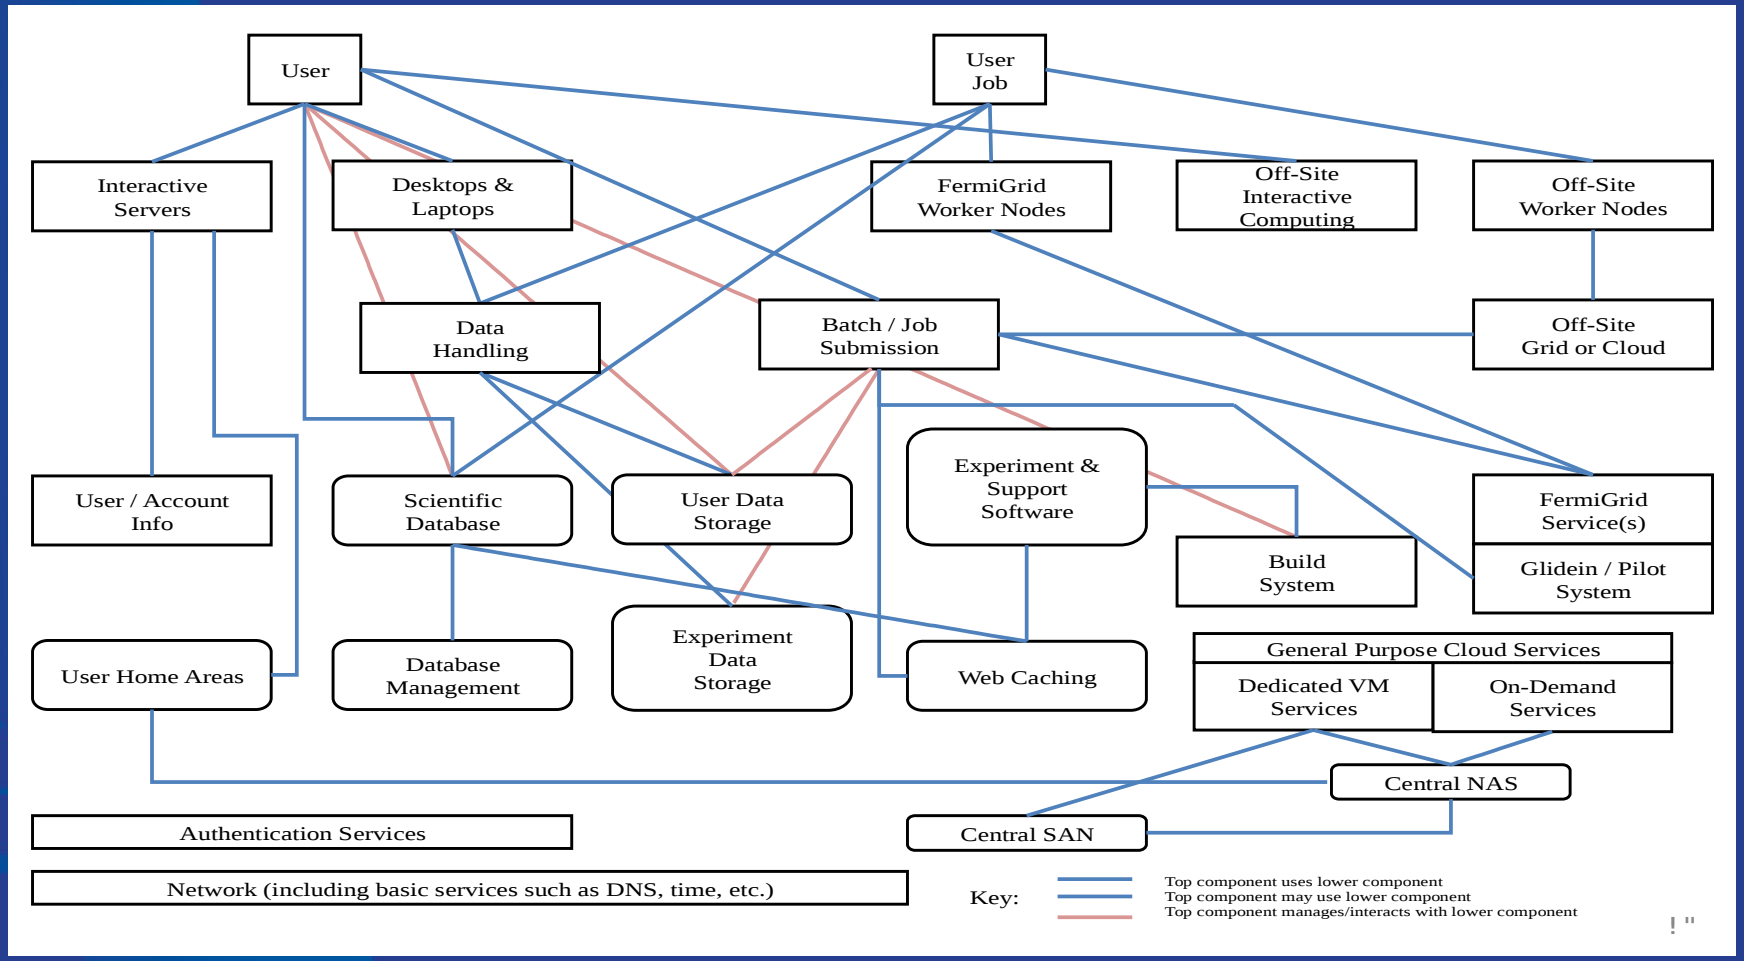

Fermilab

Storage Access

Central network attached storage, NAS (BlueArc), should not be used to store scientific data, and direct access from the large number of worker nodes should be eliminated.

Scientific data should be stored in the mass storage system and accessed through the distributed disk cache system via a data-handling layer.

For the conceptual design EnStore, dCache and SAM-Lite were chosen as implementation.

Central NAS (BlueArc)

Services should be made independent of the central NAS as much as possible.

- The “core” FermiGrid services (VOMS, GUMS, SAZ, Squid, MyProxy) used by all of the Grid clusters at Fermilab have always been independent of the BlueArc.
- Batch job services were isolated from the BlueArc in July 2013.
- Batch jobs and GridFTP services can still access the BlueArc (for the moment).

Production and Analysis Computing

Grid computing facilities on-site and off-site provide the computing power for all data processing, selection and analysis:

Fermilab Grid Clusters	# of Batch Slots
CDF Grid Cluster	5,098
CMS Grid Cluster	12,002
D0 Grid Clusters	5,776
GP Grid Cluster	7,690

These facilities are accessed via a batch/job submission service that is based on a pilot/glide-in system (jobsub).

- The jobsub framework can also be used to make (mostly) transparent use of Cloud Computing Resources (FermiCloud, Amazon EC2, etc.).

Databases

Database read access should be tiered.

- Ideally, database accesses should be via a scalable web based REST method such as the Frontier system used by both CDF and CMS.

Data Handling

There are opportunities at both handling of user generated scientific data and experimental software to improve the analysis experience of users.

A dCache based user data storage/project space is part of the conceptual design. Experiment software builds are moved from the interactive GPCF machines to a dedicated build system where a large number of processors can be used to build software releases in parallel.

Experiment and support software is then distributed and accessed on/off-site, interactively and in batch via CVMFS.

Conceptual Design

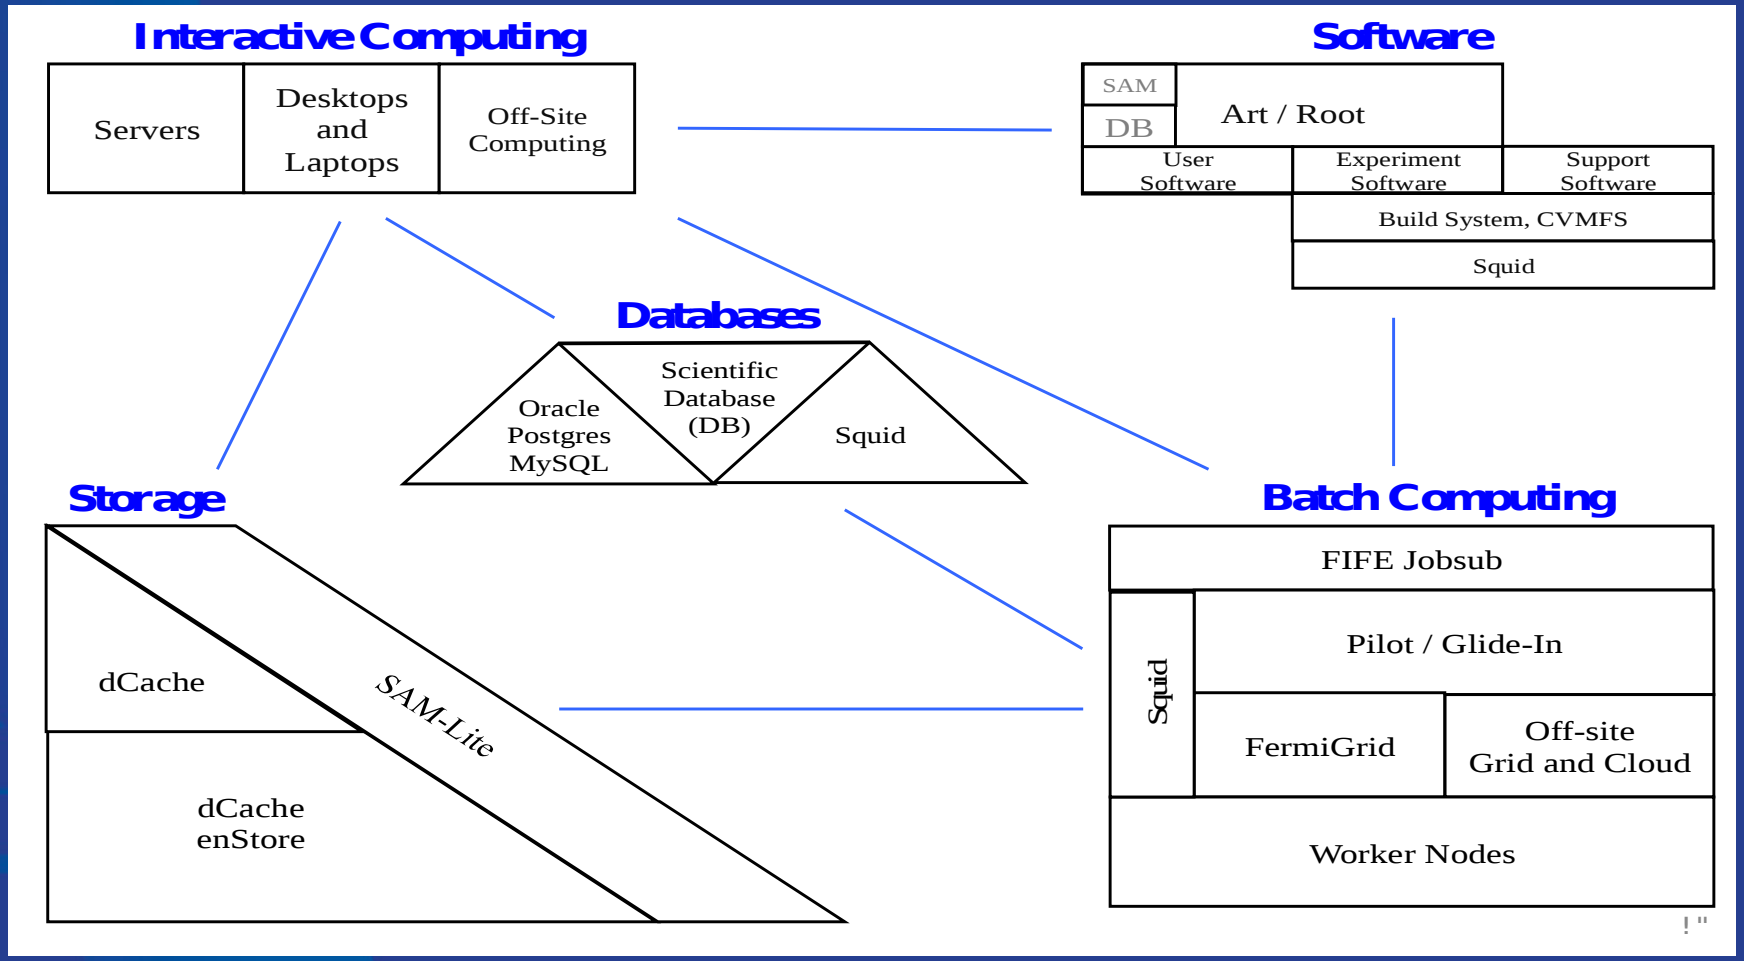

Web Caching

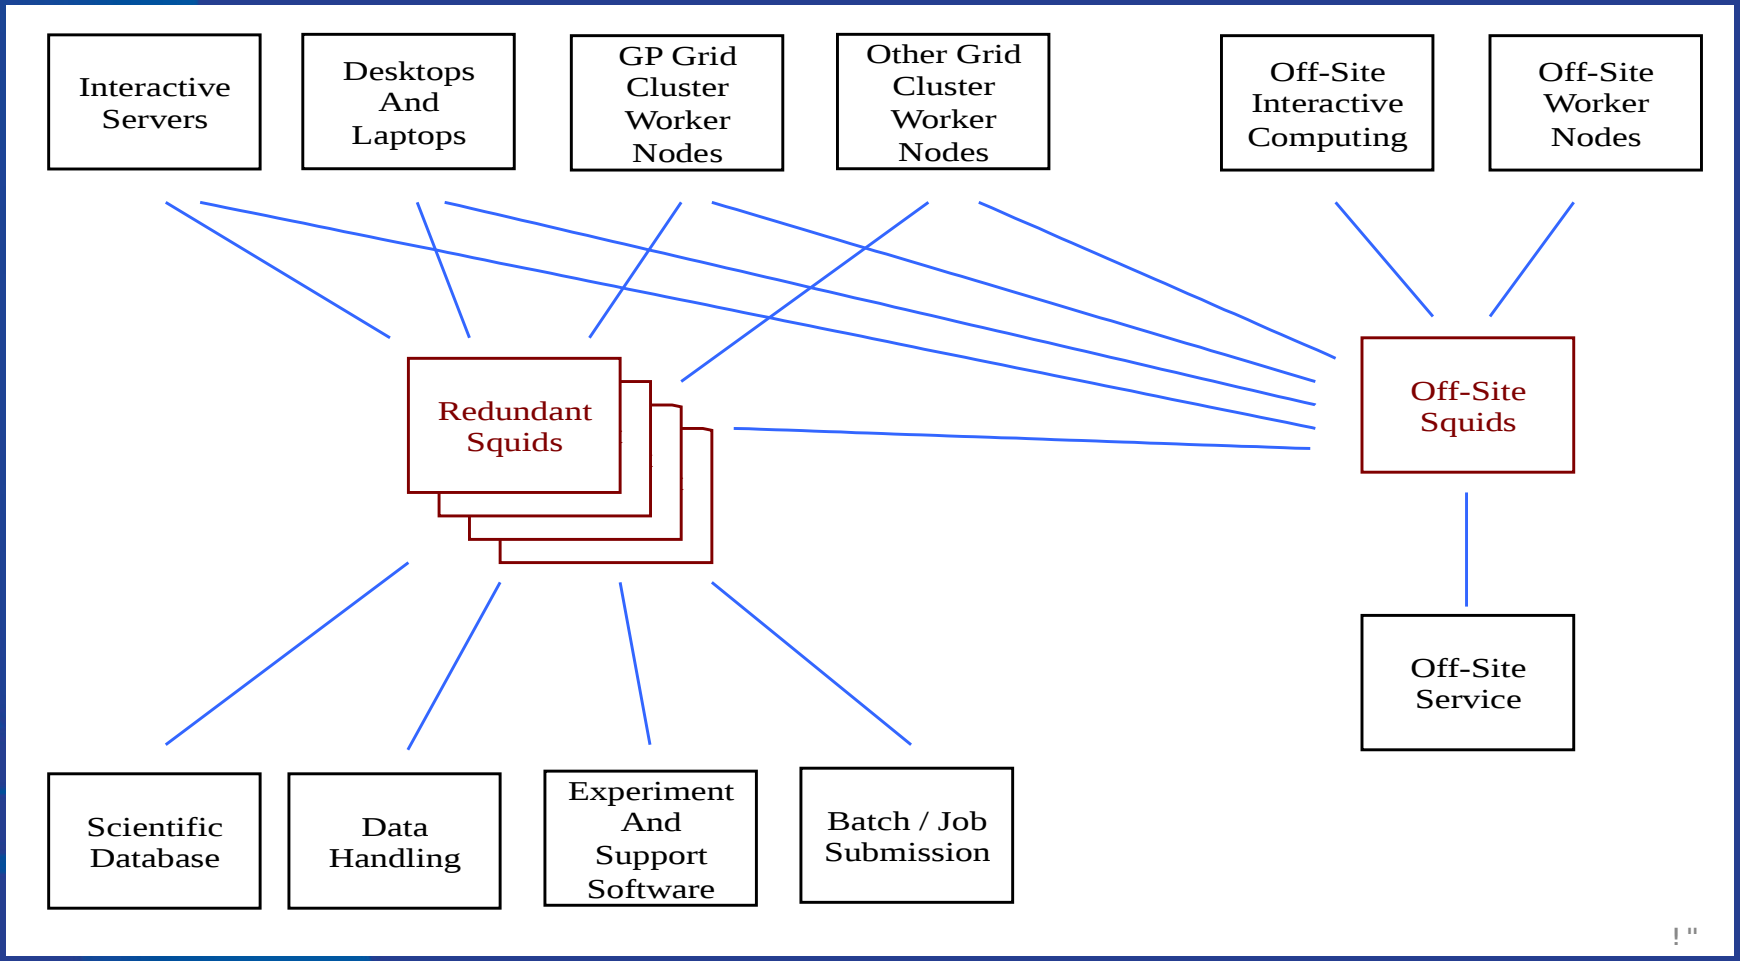
Interactive Servers
Desktops And Laptops
GP Grid Cluster Worker Nodes
Other Grid Cluster Worker Nodes
Off-Site Interactive Computing
Off-Site Worker Nodes
Redundant Squids
Off-Site Squids
Off-Site Service
Scientific Database
Data Handling
Experiment And Support Software
Batch / Job Submission

Thank You

Any Questions?