

Edge Services Overview

Topic Leads:

Ofer Rind <rind@bnl.gov>

Benjamin Weaver <benjamin.weaver@noirlab.edu>



Introduction

This presentation follows from a pre-workshop meeting held on April 1st. Thank you to Doug Benjamin, Rob Gardner, and Dustin Lang for that fruitful discussion!

[Slides](#) and [notes](#) from that meeting are available.

Edge services is a young, broad, and evolving topic with numerous applications and issues. It is unlikely that we have exhausted its breadth.

We welcome and encourage participation from all interested parties!

Outline

- Definition and Relevance of Edge Services
- Existing Examples
- Common Themes
- Concerns and Future Directions
- Specific Topics
- Summary

Defining Edge Services

“Services that operate at the interface between a data center and the wide area network, separated from the data center’s core services. This includes middleware that facilitates user access between the data center and external systems (e.g. storage, databases, workflow managers). These services may be managed externally in partnership with the data center and federated across multiple data centers.”

Definition Discussion

- This is our latest attempt at a working definition for the purpose of this meeting.
 - It may differ from the definition of Edge Services outside the scientific community
 - It may not cover all the use cases we want to address
- These services will have overlap with other Computational Frontier topics.
- Discussion topic: Is this a reasonable definition for the formal report?
 - How does this definition compare to definitions from outside the scientific community?
 - Do we need to distinguish between “edge services” and “edge computing”?
 - Does the definition adequately capture user-supplied services?
 - Does it need to address roles and responsibilities? Security?

Impact of Edge Services on Scientific Computing



From the Snowmass perspective, where do Edge Service fit in?

- Trajectory of Edge Services has been driven by the explosive growth of scientific data sets, by the increasing complexity of analyses and workflow systems, and by the development of enabling technologies.
- Their continued evolution over the next decade will have a significant impact on how and where science is done
- In many cases, they represent a fundamental shift in the User-Data Center relationship, raising questions about roles/responsibilities and security
- As a community, we need to determine how we should be driving this evolution, where there are remaining functionality gaps, and what resources are needed to keep these efforts on track.

Potentially Related White Papers

This list is not intended to be comprehensive.

- Alex Sim, *et al.*, Deploying in-network caches in support of distributed scientific data sharing, [arXiv:2203.06843](https://arxiv.org/abs/2203.06843).
- Dave Casper, *et al.*, Software and Computing for Small HEP experiments, [arXiv:2203.07645](https://arxiv.org/abs/2203.07645).
- Doug Benjamin, *et al.*, Analysis Facilities for HL-LHC, [arXiv:2203.08010](https://arxiv.org/abs/2203.08010).
- Tom Lehman, *et al.*, Data Transfer and Network Services Management for Domain Science Workflows, [arXiv:2203.08280](https://arxiv.org/abs/2203.08280).
- K. Herner, *et al.*, Towards an HPC Complementary Computing Facility, [arXiv:2203.08861](https://arxiv.org/abs/2203.08861).
- Maria Acosta Flechas, *et al.*, Collaborative Computing Support for Analysis Facilities Exploiting Software as Infrastructure Techniques, [arXiv:2203.10161](https://arxiv.org/abs/2203.10161).
- SLATE white paper draft in preparation

Existing Examples


- [Spin](#) is a container-based platform at NERSC that allows users to deploy individualized application stacks, including databases, data access tools and analytics.
- Services such as [My NERSC](#) provide individual and collaboration-scale access to metadata, allowing for long-term storage planning, or tools to manage access permissions.
- Data transfer services such as [Globus](#) are routinely implemented in edge environments such as a dedicated DMZ or data transfer node.
- [NSF's SLATE](#) is intended to provide “DevOps”-like service deployment federated across many possible data centers (see lightning talk to follow).
- [ServiceX](#) combines both data extraction and delivery for high-energy physics data sets.

Existing Examples

- [XCache](#) is a data caching service based on [XRootD](#). It is used by the LHC ATLAS and CMS collaborations for data distribution.
- [Open Science Grid \(OSG\) Data Federation](#) enables users and institutions to make datasets available to distributed HTC environments.
- [Frontier Squid](#) is an OSG implementation of the [Squid](#) cache proxy software.
- [CERN OpenStack Private Cloud](#) is an Infrastructure-as-a-Service integrated with CERN computing. Similar to Spin, but VM-based.
- [Chameleon](#) is a cloud testbed service that Internet2 [OESS](#) orchestration.
- [FABRIC](#) combines testbeds (e.g. Chameleon) with CI services.

Existing Examples: a Spin implementation

- [Legacy Survey Sky Browser](#)
- Browse through imaging from the [DESI Legacy Imaging Surveys](#).
- User provides containers; no need to ask for special resources.



Common Themes

These Edge Services characteristics and topics come up a lot, and are likely to stick around, even if the specific implementations change over time.

- Containers and container orchestration
- Configuration with declarative language
- DevOps
 - Specifically with version-controlled repositories
 - “System as software”
 - “Infrastructure-as-a-Service”
- Federation of Services

Concerns

- Existing systems with high maintenance cost.
 - Future maintainers should not be subject to highly specialized software & systems.
- Grid services (from c. 2010) shrink as centers lose personnel and services become unmaintainable.
- Definitions are evolving rapidly.
 - In the c. 2018 era, “Edge Services” primarily meant data caching services related to data delivery.
- Is there duplication of effort?
 - Many papers mention Rucio+FTS, how is this significantly different from Globus?
 - Globus & OSG Data Federation also have similar capabilities.

Future Directions

- Platforms that scale beyond individual data centers
 - Existing needs are already larger than individual center capacities.
- Changing incentives of HPC center managers.
- Sharing edge services outside of “pure” HEP.
 - Astrophysics
 - Genomics
 - Microbiology
- Glue between individual institutions that may have very different internal policies
- Security is much, much more important than even two years ago.
- Need easier-to-use platforms.

Specific Topics

- Kubernetes
- Federated Systems & Security
- Looking Outside High Energy Physics

Will Kubernetes solve everything?

- What is the true cost of deploying Kubernetes?
 - Anecdotally, at least, a lot! Not widely known for ease of use.
 - Originally designed for a single-owner, single-cluster.
 - Additional functionality needed on top (*i.e.* superset) to support federated or cloud environments: multi-owner, multi-cluster.
 - Still new, with a steep learning curve for facilities.
 - Kubernetes developers support use cases very different from HPC, scientific community.
 - How will it perform on an over-subscribed resource?
 - Security in a multi-user, multi-cluster environment.
- But even given all of the above, what is the alternative?
 - Facilities need more open source, not less.

Federated Systems

- HPC managers focus on improving their own systems.
 - Federation is not a priority.
- Also not a priority for end users
 - “It works or it doesn’t.”
- NSF is trying to change this.
- Less pressure from DOE.
 - Yet effectively all HEP is multi-national and multi-institution.
- Lack of federation results in large developer staff load at each institution.
- Federated Security
 - Probably the most important!
 - Globus is fairly far along.
 - NERSC starting to experiment.

Looking Outside High Energy Physics

- Examples from Astrophysics: Edge services can “point” inward.
- Example: workflow manager for $\sim 10^7 - 10^9$ jobs.
 - You can't submit all of those jobs all at once.
 - Workflow manager tracks completed jobs; submits new ones.
 - Implementation in NERSC Spin system; not intended for use outside the HPC center (but could be!).
- Example: bespoke databases for instrument operations.
 - No need to ask center management for database access, just start up a PostgreSQL container.
 - Access to instrument itself is difficult in both physical and network senses: copy data to HPC center.
 - Again, not intended for use outside the HPC center (but could be!).

Summary

- There are many types of Edge Services.
 - Far more than we can describe in just these slides.
 - Definitions changing year-by-year.
- A few common themes form the foundation of future ideas.
- Don't go down a path that nobody will use even a few years from now.
- Do make life easier on support staff by using open source, open protocols & federated resources.
- Many ideas outside of DOE/HEP.
- Security must be an important part of all decisions.
 - Federated identity & security can help bridge different institutional requirements.
- Discussions welcome!

Supplemental Material

- [Raw notes from Informal Edge Services Meeting, 2022-04-01](#)
- [Introductory slides provided by Joe Breen <joe.breen@utah.edu> for the August 2020 Snowmass Computational Frontier Workshop](#)