Digital Research Alliance of Canada | Alliance de recherche numérique du Canada

Distributed Storage and Compute Grid

Call for Participation

2025-09-09

Contents

1. Introduction

The Alliance is piloting a nationally operated distributed storage and compute service to provide reliable, scalable, and secure data storage and AI compute resources (and supporting service nodes) through a distributed, redundant environment. This initiative, funded through the Government of Canada's Canadian Sovereign AI Compute Strategy, will be a uniquely Canadian service that leverages our large geography for true site diversity and high-speed network connectivity across the country.

The Pilot is the Alliance's response to needs repeatedly raised by the research community around limited options for active data storage, collaboration, compliance across jurisdictions, and rising local operating costs. The proposed operational model will aim to reduce local burden, strengthen interoperability, and generate the evidence on reliability, performance, usability, security, and total cost that will inform broader deployment.

2. Overview of Call for Participation

The Alliance aims to pilot colocating Alliance-owned infrastructure at partner institutions and linking these sites over CANARIE's National Research and Education Network (NREN) and associated regional area networks (RANs) to form a resilient grid. The service will be centrally managed by the Alliance through the creation of a Grid Operations team. Local data centre support will be intentionally minimal and limited to "remote-hands" activities (e.g., receiving and racking equipment at install, swapping replaceable components, etc.).

This Distributed Storage and Compute Grid (DSCG) will be comprised of a small number of storage/compute clusters distributed across the country using two reference deployment patterns:

- Multi-site stretched cluster: at least three geographically proximal interconnected data centres, each hosting a small footprint (e.g., ~two racks of compute/storage equipment), operating as one logical cluster with automatic data placement and tolerance of a site-level failure.

- Single-site cluster: a self-contained cluster at one data centre, used where a stretched deployment is not feasible, with replication and recovery contained within that single site.

This Pilot will comprise up to three clusters distributed across Canada and based on a mixture of single- and multi-site configurations. Policy driven synchronization between clusters will be enabled to provide cross-regional resilience and data locality.

Note: In the context of this Pilot, "cluster" refers to a Ceph and Kubernetes based cluster and not an HPC compute cluster.

2.1 Use Cases

The DSCG will form the foundation of the Canadian Research Data Platform (CRDP): a secure, scalable, and interoperable platform to streamline access, management, and analysis of research data across the Canadian Digital Research Infrastructure (DRI). Specifically, the DSCG will provide durable, scalable storage solutions to underpin services across the research lifecycle. Some of the use cases for the DSCG will include:

- Active Data Management Service (ADMS) will be the Canadian Research Data Platform's (CRDP) core service for storing, managing, sharing, and tracking research data during the active phase of the research lifecycle when data is being collected, processed, and analyzed. It will provide researchers with secure, scalable, and collaborative storage integrated within the CRDP ecosystem, ensuring that data is well-

managed from creation through to deposit, publication, and preservation.

- The DSCG will serve as a robust backend for participating data repository platforms in the CRDP, such as the Federated Research Data Repository (FRDR). These repositories will use the DSCG to store published datasets securely and reliably. Additionally, national preservation services will leverage the same infrastructure to manage long-term stewardship of Canadian research data. This preservation layer would replicate data across multiple nodes, ensuring that valuable research outputs remain intact and accessible beyond their initial publication period.

- Research-focused national data spaces will be decentralised infrastructures, where diverse researchers could share and use data in a secure, reliable and trustworthy manner following common governance.

- Research Platforms and Portals (RPPs) hosted on Alliance infrastructure will draw on the DSCG for reliable, high-performance storage backing services, datasets, and user workspaces.

2.2 Objectives

- Demonstrate real-world operation across a multi-site cluster with very high levels of redundancy and reliability.

- Demonstrate real-world operation across the distributed grid of clusters with support for policy-driven data movement.

- Verify the central operating model of using Alliance Funding Agreements (AFAs), continuous monitoring, and incident/change management to meet availability, performance, security, and low local-burden targets.

- Establish a repeatable onboarding pathway by codifying site readiness, installation, acceptance testing, and training so additional host sites can be added efficiently.

- Build the evidence base for scale-up deployment by collecting metrics on reliability, performance/latency, usability, security posture, adoption, and total cost of ownership to inform national rollout and investment decisions.

2.3 Benefits of Participation

This is an exciting new initiative where participating institutions will be at the forefront of defining a new national service and address longstanding researchers' needs. This Pilot will move DRI away from the current focus on a small number of very large research computing systems to a uniquely Canadian system distributed across the country.

Beyond the opportunity to shape the future of DRI, researchers from participating sites will be given preferential access to the pilot. The Alliance will invite projects from researchers at the participating sites to test the infrastructure. Generally, the focus will be on projects which either

use AI or undertake basic AI research. Up to 30% of the available resources will be made available to such projects.

3. Host Sites Eligibility Criteria and Requirements

3.1 Eligibility

Participation in this Pilot is open to:

- Universities, colleges, and research hospitals in Canada with active research programs or demonstrated capacity to conduct meaningful research.
- Not-for-profit organizations in Canada that conduct, or can demonstrate the capacity to conduct, meaningful research.
- Institutions or not-for-profit organizations in Canada whose activities enable or support Canada's digital research infrastructure (e.g., data, computing, software, networking, cybersecurity).

Participating hosts must, at a minimum, be able to meet the Technical Requirements (3.2) and agree to work with the Alliance to confirm operational procedures and service levels during the Pilot.

3.2 Technical Requirements

In addition to the eligibility described above, participating hosts must provide the following services as a colocation facility hosting the rack(s).

Space	Appropriate data centre space for the rack(s) including power and cooling. • Stretched Cluster: two racks. • Traditional Cluster: at least four racks. • Up to 800 kg per rack. • Estimate of 24 kW per rack.
Racks	The Alliance expects to work with sites to procure suitable racks either independently or through preferred suppliers.
Power Distribution Units (PDUs)	The site should provide two metered PDUs per rack with sufficient capacity for each rack of equipment hosted by the site. Servers will be connected in a redundant fashion.

Uninterruptible Power Supply (UPS)	There should be sufficient UPS capacity to gracefully power down a rack when mains power is interrupted. Longer-term generator power is not necessary as the Grid-Cluster architecture provides the necessary redundancy. However, diesel generator power would be an advantage if it is available.
CANARIE Network	A connection to the provincial/territorial NREN partner (e.g., BCNET, SRNet, ORION, ACENET, Cybera, etc.), which provides the on-ramp to the CANARIE backbone for research and education traffic. To support the sustained transfers typical of distributed storage and compute workflows the Alliance expects to work with CANARIE and the NREN partners to provision a minimum of 10 Gbps with upgrades to 100 Gbps where necessary for the stretched clusters. The stretched clusters are expected to be based around a layer 2 VPN service (eg. VPLS) operating at 100Gbps and <10ms RTT.
Commercial Network	The site should be able to provide a commercial internet connection of at least 1 Gbps. Our assumption is that institutional data centres will have commercial internet connections and will be able to provide access for the DSCG at an advantageous price.
IP Addressing	The site should be able to provide publicly routable IPv4 addresses such as a /27. IPv6 would be an asset.
Network Security	The site should be able to provide basic network security facilities through the site's normal firewall and monitoring capabilities.
Network Usage Monitoring	The site should be able to provide network (CANARIE and commercial) usage monitoring to the Alliance's proposed Grid Operations Centre.
Data Centre Operation Support	The site should be able to provide operators who can undertake day-to-day on-site tasks. This includes "rack-and-stack" initial installation, physical monitoring for visual and audible alarms and responding to ticket requests from the Grid Operations team for tasks such as replacing faulty equipment, power cycling, basic hardware troubleshooting, etc. 5x8 support is sufficient but 7x24 would be preferred.

Tape Facilities	While redundancy within a cluster and over the DSCG should be sufficient for almost all purposes, there may be need for storage for cold or inactive data. While not necessary, the Alliance is therefore interested in tape facilities at one or more sites within the DSCG.
Environmental Monitoring	The site should be able to provide sufficient monitoring to identify and report environmental issues to the Grid Operations Centre including: • Power outages • Loss of cooling • Cooling leaks • Electrical Issues • Physical Access • Fire and smoke • Other
ITSM	The Alliance is installing an IT Service Management system (Xurrent) which will be used by the Grid Operations Centre. The Alliance therefore expects that the site incident management, change management, asset management and risk management information will be shared with the Alliance. The Alliance would also be open to the site using the Alliance's ITSM system directly but understands that the site may have their own ITSM system.
Receiving, and Staging of Equipment	The provider must be able to accept inbound deliveries of servers and related IT equipment via a suitable loading dock during business hours. Upon receipt, shipments must be inspected, documented, and secured in an access-controlled, monitored storage area until they are ready for unpacking and installation.

4. Roles and Commitments

4.1 The Alliance

- Pay for colocation costs, including rack/room charges, power, cooling, and required network cross-connects. Manage delivery and installation logistics.
- Procure, insure, and manage Alliance-owned hardware for the Pilot term. Alliance will manage costs of shipping and insurance for Alliance-owned equipment, and vendor installation support as needed.
- Operate the service centrally: monitoring, provisioning, capacity planning, security baselines and patching, incident response, telemetry, and change control.

- Pay for and coordinate networking between sites with CANARIE and provincial partners and complete acceptance testing for throughput and resiliency.
- Provide onboarding materials, AFAs, training for local contacts, periodic reporting on reliability/performance/security/cost, and clear end-of-pilot options (extend, scale, relocate, transfer, or decommission).

4.2 Hosting Institution

- Meet the Technical Requirements (3.2)
- Work with the Alliance to coordinate the delivery and installation of equipment.
- Participate in Pilot reviews to confirm operational procedures and service levels.
- Collaborate with the Alliance and other hosting sites to ensure the success of the Pilot.

Note: This is a call for participation, not a grant program. The Alliance will retain ownership and lifecycle responsibility for Pilot hardware and cover agreed colocation costs.

5. Host Sites Selection

The Alliance will identify a small cohort of host sites to establish two to three clusters for the Pilot. Because this is a call for participation, the goal is to assemble a viable configuration rather than to rank applicants. Nonetheless, selection will include technical suitability and readiness alongside portfolio factors (e.g., geographic/network diversity and the ability to form performant stretched clusters). The process will include:

1. Qualify sites for technical and site readiness according to section 3.2.
2. Develop 2 or 3 grid architecture scenarios depending on site capabilities and geographic location, and availability of suitable network connections. Each scenario will include at least one multi-site cluster and one single site cluster.
3. Review each scenario for total cost and choose the scenario which maximises resources within the available budget.

Depending on the mix of Respondents and available capacity, some eligible sites may be invited to proceed immediately; others deferred to a later phase or assigned roles that differ from those proposed.

6. Procurement, Delivery and Installation

Procurement

The Alliance will procure all Pilot hardware, retain title and insurance, and coordinate all vendor logistics with each host site. Deliveries to participating sites will be completed no later than March 31, 2026.

Delivery and Installation

Prior to delivery, host sites will be expected to confirm readiness in accordance with the Technical Requirements (3.2). The site should be prepared with a network drop, racks, PDUs and appropriate cooling so equipment can be installed after delivery.

Upon delivery, host sites will be responsible for receiving shipments, inventorying all equipment, unpacking and installing equipment.

The Alliance will coordinate shipping between vendors and host sites, validate site prerequisites, bring grid online, configure switches and firmware, install and configure appropriate software, build and integrate the cluster, and run acceptance tests.

7. Timeline

This is a Pilot with guaranteed funding until March 31, 2028. However, the Alliance expects that the Pilot will evolve to a basic production service and will be making a strong case for continued operational funding.

Activity	Estimated Timeline	Description
Call for Participation Issued	September 9, 2025	
Information Session	September 22, 2025	
Submission Deadline	October 7, 2025 23:59 Pacific Time Zone (PT)	
Review and Feedback	October 8 – November 4, 2025	
Participation Confirmation Issued to Selected Respondent(s)	November 2025	
Finalize Hosting Agreement(s)	November 2025	

Delivery of Equipment	By March 31, 2026	Equipment delivery is required by this date.
Installation	April 1 - July 1, 2026	Each site will install the equipment in their data centre.
Configuration	July 1 - September 30, 2026	The Alliance Grid Operations Team will install and configure the Ceph, Kubernetes and Grid software.
Go-live	October 1, 2026	
Pilot Review	July 30, 2027	A formal review of Pilot results and funding availability will be required to either proceed to scale or initiate a wind-down.

8. Reporting

Host sites will be required to submit quarterly performance reports using a reporting template provided by the Alliance. Performance reports will include information on power and network outages, UPS/generator events, latency events, bandwidth utilization, staffing levels, maintenance activities, physical security events, compliance audits/findings, environmental monitoring, customer impact, upcoming planned activities.

9. Costs

The Alliance will cover reasonable and customary colocation charges and installation costs and will contract directly with the host sites. No cash awards will be made beyond these costs.

9.1 In Scope

- Recurring colocation/rental fees for the Pilot footprint such as rack/room charges, power and cooling, and cross-connects to the research and commercial networks.
- Commercial network costs if the site provides additional bandwidth through their existing connections.
- Local network infrastructure (e.g., cables, switches) that would be required for an upgraded connection to an NREN/RAN service at up to 100 Gbps.
- One-time installation charges related to the Pilot footprint (e.g., cross-connect turn-up, delivery handling where applicable).

9.2 Out of Scope

- Facility upgrades required to meet minimum site readiness (e.g., additional cooling capacity, room build-outs).
- Local staffing beyond remote-hands activities defined in the operating model.
- Optional tape services unless explicitly agreed (see below).

9.3 Cost Information

Respondents should provide a pricing schedule in the Cost Schedule (Appendix 2) covering the following items:

9.3.1 Recurring monthly charges

- Space (rack/room) rental costs.
- Power and cooling costs (state basis: per-rack, metered kWh, or facility tariff).
- Commercial network costs (if the site can provide commercial network bandwidth).
- Data centre operator support/remote-hands hourly costs.

9.3.2 One-time charges

- Installation costs to place and connect equipment in the data centre (including cross-connect turn-up).
- Infrastructure update costs, if required to meet technical requirements (e.g., network cross-connects for 100 Gbps); these will be out of scope unless pre-approved.

9.3.3 Optional services

If a tape library is available, Respondents should provide tape system costs including day-to-day operations, media purchases, and any additional capabilities (e.g., extra tape drives). Indicate whether these services would be made available to the Pilot and under what commercial terms.

9.4 Administrative Details

Pricing should be valid through the Pilot end date, with quarterly invoicing and applicable taxes and a finance/contract contact.

9.5 Billing and Invoicing

- Host sites will be required to provide an itemized rate sheet and submit invoices to the Alliance quarterly.
- Rates will remain fixed for the Pilot term unless otherwise agreed in writing.
- Any additional cost not listed on the rate sheet will require prior written approval.
-

10. Response Submission

Responses must be submitted to the Alliance by email to funding-subventions@alliancecan.ca.

Responses without a letter of certification will not be assessed.

> **Subject:** Response for Alliance Distributed Storage and Compute Grid
> **Email:** funding-subventions@alliancecan.ca

Responses should be submitted no later than October 7, 2025 at 23:59 Pacific Time Zone (PT).

11. Submission Checklist

- ❑ **Appendix 1.** Completed Response Form
  - Response must be formatted in Arial 11, single-spaced. (10 pages max)
  - Headings and subheadings are **bolded** adhering to the respective word/page requirements.
- ❑ **Appendix 2.** Completed Cost Schedule Form.
- ❑ Your institution's colocation or hosting agreement template, if any.
- ❑ An official letter, certifying the response, must be signed by one or more authorized signatories as defined by your institutional policies and the technical lead. The letter must include a statement confirming your organization's interest in participating the Pilot and that the information provided is accurate. Responses without a letter of certification will not be assessed.

12. Confidentiality and Lobbying Restrictions

Confidentiality

Information provided as part of individual responses may be accessible under the Access to Information Act (ATIA). No commercially confidential information which is submitted to the Alliance will be disclosed unless otherwise authorized by Respondents; required to be released by law (including ATIA); or required by the Minister of Innovation, Science and Industry to be released to an international or internal trade panel for the purposes of the conduct of a dispute in which Canada is a party or a third-party intervener. Additional information on the federal government's Access to Information Act is available on the official website. Any information in the project submission that is or may be of a proprietary or confidential nature should be clearly marked as "Proprietary" or "Confidential" on each relevant item or page or in a statement covering the entire submission.

No Lobbying

Respondents to this Call for Participation shall not engage in lobbying the Government of Canada on any matter related to the Respondent's response to the Call for Participation.

13. General Inquiries

Any questions, inquiries or clarifications related to this Call for Participation *Distributed Storage and Compute Grid* must be submitted to funding-subventions@alliancecan.ca. The Alliance is committed to an open and transparent process and will share process related questions and responses with all eligible Respondents.

Inquiries must be received at least 7 calendar days before the submission due date. The Alliance is unable to guarantee a response to inquiries received after this date.

14. List of Appendices

Appendix 1. Response Form

Appendix 2. Cost Schedule Form

Appendix 3. Technical Architecture

Appendix 4. Alliance Financial Management Guide

Appendix 5. Alliance Funding Agreement Template

Appendix 1: Response Form

This form helps the Alliance assess each site's technical, operational, and infrastructure capabilities in relation to the needs of the national grid. Please provide responses to the sections below, adhering to the specified page limits.

1. **Organization Information**
    a. Name of institution
    b. Address
    c. Type of Organization
        ☐ University, college, or research hospital in Canada with active or potential research capacity
        ☐ Not-for-profit organization in Canada conducting or capable of conducting research
        ☐ Institution or not-for-profit organization in Canada supporting Canada's digital research infrastructure
    d. Primary contact name
    e. Contact title
    f. Contact email
    g. Contact phone

2. **Preferred Hosting Role (½ page)**
    Please indicate your preference and ability to host equipment for one or both of the following models. If the site is able to host both models, then please identify which model is being addressed in the sections below.
    ☐ multi-site cluster (part of a three-site geo-distributed cluster operating as one logical cluster)
    ☐ single-site cluster (self-contained cluster at one data centre)

3. **Rack Capacity (½ page)**
    The basic building block is two racks. For redundancy and balance, the Alliance expects each site to host up to six racks.
    a. How many racks could you host today?
    b. What is the maximum number of racks you could host within the Pilot window?

4. **Power and Cooling Infrastructure (2 pages)**
    a. Describe the power supply and capacity for the proposed footprint.
    b. Include the UPS system and capacity details (if available) and estimate how long the UPS can power the rack(s) at the proposed load.
    c. If a diesel generator is available, indicate whether the rack(s) are on generator-backed circuits and state its availability.
    d. Are any upgrades required to provide power to the rack(s)?
        ☐ Yes. If yes, outline the scope and lead time.
        ☐ No.

5. **Data centre management (1 page)**
   Please describe the on-site operations team and their availability. Ie) 7x24 on-site or call in? Availability for emergencies? Etc.

6. **Cybersecurity and physical security (1 page)**
   What are the cyber and physical security features available in the data centre? For instance, the availability of firewalls, and the software being used to monitor network traffic.

7. **Change management (½ page)**
   a. Describe your data-centre change management process.
   b. How will the Alliance Grid Operations Centre be notified of planned changes or maintenance that could affect the Pilot (lead time, channels, approval points, maintenance windows)?

8. **Incident management (½ page)**
   a. Describe your incident management process, including classification and escalation.
   b. How will the Alliance Grid Operations Centre be notified of unplanned outages or security events (notification method and target response times)?

9. **Monitoring (1 page)**
   a. Describe the monitoring systems in place for power, cooling, environment, racks, and network.
   b. Can your systems share telemetry or alerts with the Alliance Grid Operations Centre? Specify interfaces or methods for sharing (for example, APIs, SNMP, syslog).

10. **Tape facility (1 page)**
    a. Can you provide access to a tape library?
       ☐ Yes.
       ☐ No.
    b. If yes, provide details.
       - Is a tape librarian or equivalent support available?
       - What management software is used, and can it be made available for backup and archive of inactive data?
       - Would additional drives or media be required? Include estimates.

11. **Network connectivity (1 page)**
    a. Describe the current cross-connect to CANARIE/NREN, including port speeds and redundancy, and the bandwidth you can commit to the Pilot racks.
    b. For inter-site operation of a distributed Ceph cluster, can you provide 100 Gbps cross-connect capacity to an upgraded or new NREN/RAN connection? If not, state the maximum available capacity and any upgrades that would be required, including lead times.
    c. Describe any commercial internet connectivity that would be used, and any required cross-connects.

Appendix 2: Cost Schedule

As part of the Digital Research Alliance of Canada's Call for Participation in the Distributed Storage and Compute Grid, please complete this Cost Schedule to describe recurring and one-time costs associated with hosting a Pilot node.

Notes:

- Provide N/A where an item does not apply at your site.
- Rates should remain fixed for the Pilot term unless otherwise agreed in writing.
- Invoices will be submitted quarterly in arrears with supporting details where charges vary by consumption.
- Title to all Pilot hardware remains with the Alliance; the Alliance carries equipment insurance.

Administrative details

Organization	
Billing address	
Finance/contract contact name	
Contact title	
Contact email	
Contact phone	

Costing Breakdown

Type	Category	Subcategory	Rate	Amount FY 2026-2027	Amount FY 2027-2028
Operating Costs		Power and cooling monthly rental cost per rack			

	Space and Rental Expenses	CANARIE cross-connect monthly cost			
		Commercial connection monthly cost			
		Data centre operator hourly cost			
		Tape system monthly operating cost (if available)			
Subtotal					
Initial One-time Costs	Installation Costs	Approximate cost for taking delivery of the equipment, racking the equipment, and connecting to suitable power and network.			
	Infrastructure Costs	Approximate cost for additional cabling and cooling infrastructure if not included in the above rent costs. Note that only minor upgrades are eligible.			
	NREN/RAN cross-connect	Approximate cost if any for upgrading existing network infrastructure to be able to accept a 100 Gbps connection.			
	Rack provisioning costs	If the site has specialized rack requirements or approved suppliers, please provide an estimate of the cost per rack.			
	Tapes	If the site can provide tape library facilities, please provide the type and cost of additional tapes (per tape).			

Subtotal			
Total			

Appendix 3: Technical Architecture

The following information is provided so that the sites have the background information for the technical requirements.

Stretched Multi-site Cluster

A stretch cluster is a three site, six rack deployment of a DSCG cluster. It consists of two racks per site with all three sites working in concert. The goal of this architecture is extreme fault tolerance and high uptime. This architecture is designed to tolerate a single site going offline while the remaining two sites continue to allow reading and writing to the storage cluster. Compute applications impacted by an offline site will be automatically restarted and rescheduled on available compute nodes at the two remaining online sites. A high speed and low latency <10ms RTT 100Gbps link with VPLS between the sites are required. A 40Gbps connection to CANARIE is recommended for North-South client traffic, 10Gbps minimum. There should not be any common points of failure between the three sites, such as a common campus core network or power feed.

Logical Architecture

Users access the service through common API's provided by Ceph and Kubernetes and are load balanced across the three sites. User applications provided by the Alliance and/or run by researchers will be accessed through name-based hosting and Kubernetes Ingress. The back end Ceph replication network will leverage a 100Gbps VPLS between the sites so the Ceph OSD's can communicate and distribute data using Erasure coding.

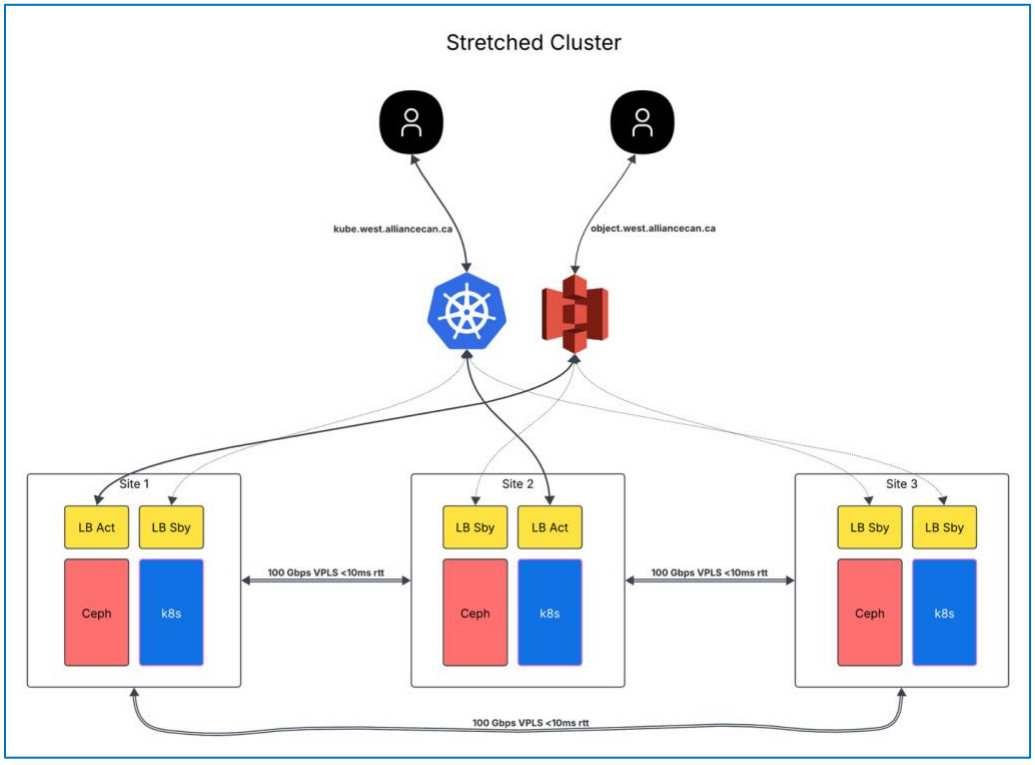

Physical Architecture

The rack layout diagram below is for example purposes however based on this design the data centre needs to provide standard 42U 19” racks and be able to accommodate 25 kilowatts per rack and 775 kilograms per rack.

In this architecture each site consists of:

- 20 storage nodes
- 6 GPU compute nodes
- 6 CPU compute nodes
- 5 Control plane servers
- A load balancer
- 25/100G Top of Rack switch
- 1G Top of Rack switch for Out of band connectivity

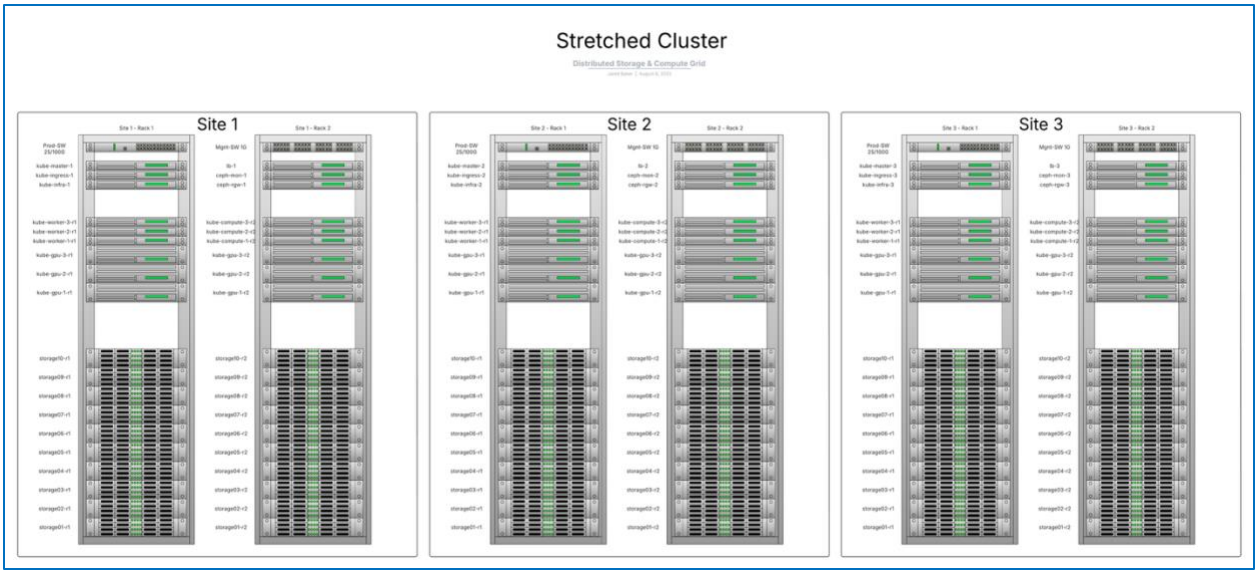

Traditional Single Site Cluster

A traditional cluster is a single site, four or more rack deployment of a DSCG cluster. The goal of this architecture is site fitment with a reduced set of requirements by the data centre provider. This architecture offers fault tolerance at the rack level, tolerating an entire rack failure while continuing to allow reading and writing to the storage cluster. Compute applications impacted by a rack failure will be automatically restarted and rescheduled on available compute nodes in the remaining racks. A 40Gbps connection to CANARIE is recommended for North-South client traffic, 10Gbps minimum.

Logical Architecture

Functionally the same as the Stretched Cluster however the back end Ceph replication network will remain local and leverage the 100Gbps Top of Rack switch topology so the Ceph OSD's can communicate and distribute data using Erasure coding.

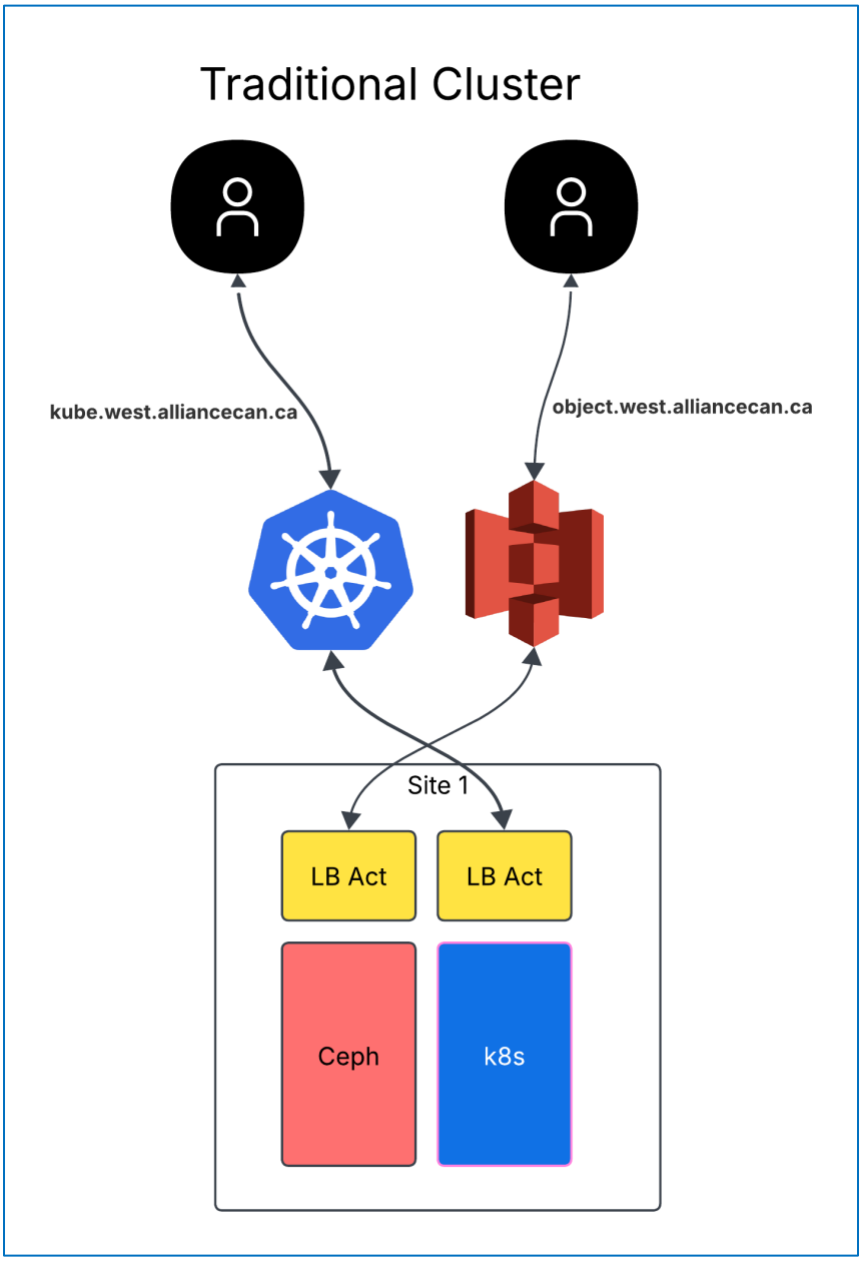

Physical Architecture

The rack layout diagram below is for example purposes however based on this design the data centre needs to provide standard 42U 19" racks and be able to accommodate 25 kilowatts per rack and 789 kilograms per rack. This architecture lends itself to some flexibility in number of racks and could potentially be bigger but this needs to be done carefully and in collaboration

with the colocation provider as Erasure coding settings and failure domains need to be carefully considered.

In this architecture each site consists of:

- 40 storage nodes
- 12 GPU compute nodes
- 12 CPU compute nodes
- 15 Control plane servers
- A pair of load balancers
- 4 25/100G Top of Rack switches
- 4 1G Top of Rack switches for Out of band connectivity

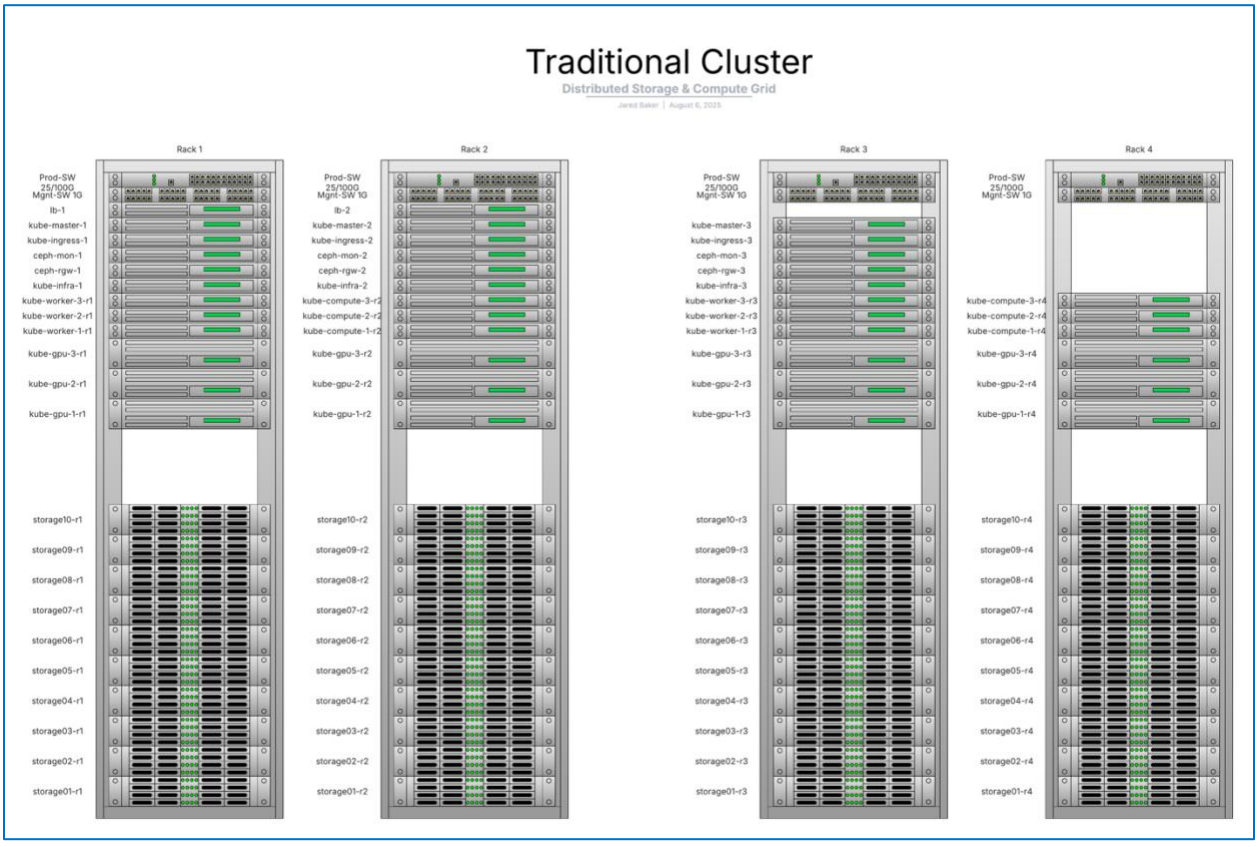

Appendix 4: Alliance Financial Management Guide

Access the *Financial Management Guide*

Appendix 5: Alliance Funding Agreement Template

Access the *Alliance Funding Agreement*