DAFFODIL GROUP

EDUCATION / RESEARCH / ENTERPRISE

Daffodil AI Infrastructure

PROGRAMME IN DEVELOPMENT

Building the foundation for Bangladesh's AI future.

One governed AI platform for education, research and enterprise.

Daffodil is developing a shared foundation for GPU computing, private knowledge services, research and accountable automation. The first priority is to serve Daffodil's institutions and businesses. Proven internal delivery will create a basis for selected enterprise services across Bangladesh.



Illustrative rack concept.
Reference architecture, not installed equipment.

Institutional control

Approved knowledge, access and ownership under one governance model.

Shared capability

Allocated resources for departments, research and application teams.

Practical innovation

Supervised learning and research connected to reviewed applications.

Bangladesh AI Business Summit 2026

5 November 2026 | BCFCC, Dhaka

A complete platform designed to scale.

Planned AI compute at Daffodil Smart City, Ashulia, with protected off-site recovery at Dhanmondi. Configuration shown is a reference design.

2
GPU servers

8
initial GPUs

100TB+
protected usable

100GbE
redundant fabric

GPU SERVER A

96GB | 96GB | 96GB | 96GB

4 GPUs initially | 8-GPU-capable chassis

96+ physical CPU cores
1TB+ ECC RAM | 15-20TB local NVMe
Mirrored boot SSD

GPU SERVER B

96GB | 96GB | 96GB | 96GB

4 GPUs initially | 8-GPU-capable chassis

96+ physical CPU cores
1TB+ ECC RAM | 15-20TB local NVMe
Mirrored boot SSD

Reference GPU: NVIDIA RTX PRO 6000 Blackwell Server Edition. Eight GPUs provide 768GB distributed GPU memory; memory is not automatically pooled. Expansion to sixteen GPUs requires OEM and facility validation.

Shared storage	One HA system: 100TB+ protected usable, expandable to 300TB+.
Network	Dual 100GbE minimum fabric and separate management access.
Management	One HA platform: three control nodes or a verified HA equivalent.
Backup and recovery	One dedicated system; protected copies, retention and restore testing.

SEPARATE ACADEMIC OPTION

Two labs | 52 endpoints | Two shared H200 NVL-class GPUs. Lab server count is vendor-proposed, with mirrored boot, local NVMe and protected user data. This option is separate from the eight production GPUs.

From governed access to useful AI.

Security, monitoring and human accountability apply across every layer.

1. **People and applications**
   Students, faculty, departments and approved business systems.
2. **Identity and policy**
   Sign-in, roles, source permissions, tenant separation and quotas.
3. **Knowledge and AI services**
   Permitted knowledge retrieval, approved models and controlled tools.
4. **Infrastructure operating layer**
   Allocation, scheduling, monitoring, usage accounting and audit.
5. **Compute and data foundation**
   GPU servers, storage, network, power and recoverable backups.

A request in practice: Check access, retrieve only permitted knowledge, use an approved model and return a source-grounded answer. Any tool action needs its own authorization; sensitive actions retain human approval.

INTENDED APPLICATIONS

Learning

Supervised assistants, coding practice and student projects.

Research

GPU experiments, model evaluation and authorized tuning.

HR and accounts

Policy retrieval, summaries and draft analysis with owner review.

Daffodil Online

Service knowledge and network, server and web assistance.

DCL and Group

Product engineering, staff knowledge and workflow assistance.

Cybersecurity

Isolated research and human-supervised security analysis.

Illustrative uses. Each integration requires a responsible owner, permitted data and acceptance.

Prove value inside Daffodil. Grow with Bangladesh.

01

Develop the foundation

Build on data-center and AI pilot experience. Complete central procurement, installation and acceptance.

02

Serve Daffodil

Begin with selected knowledge, service-desk and academic uses. Measure value and operating readiness.

03

Open enterprise pathways

After internal stability, validate customer isolation, spare capacity, data terms and service support.

WHY IT MATTERS

- Institutional knowledge, access and service ownership.
- Shared capacity and reusable application integrations.
- Lower recurring external API exposure for suitable work.
- Practical skills and a reviewed path into production.

LOCAL FIRST, WITH CHOICE

Private and recurring workloads are intended for controlled local processing. External AI is available only under approved data and provider policies. A local timeout does not authorize private-data export.

Knowledge retrieval is different from training. Research and selective tuning use authorized data; this is not a claim of AGI.

FUTURE SERVICE PATHWAYS

First candidate: Private company knowledge assistants and employee productivity.
Later: Managed models and LLM APIs, AI application hosting, GPU-backed research and inference, and GPU colocation where facilities and capacity support it.

Discuss your use case

info@daffodilnet.com
+880 1847-334837