AKF HPC SYSTEM CONFIGURATION & OPERATION POLICY

ASEAN-Korea HPC Facility in Indonesia (AKF)

Field	Detail
Document	AKF HPC System Configuration & Operation Policy
Version/Status	Draft v1.0
Supersedes/Merges	KISTI System Policy V3.8 + BRIN reconciliation (Consolidated Policy)
Owners	ASEAN HPC Laboratory Management (AHLM), with the AHSC
Basis	Assumes the AKF Consolidated Policy decisions are accepted: username <country code><affiliation code><5-digit>; V3.8 queues + CPU correction; V3.8 SRU factors with BRIN allocations; /home per-user + scratch per-project quotas & retention; granular administrators
Date	21 August 2026

Status. This operational policy is based on the discussion on 19 August 2026. The General Policy sets governance rules (eligibility, categories, allocation, review); this document covers system configuration, metering, storage, identity, lifecycle, administration, and operations. Some detailed partitioning and configuration are yet to be implemented in the system.

1 System Overview

Component	Configuration
GPU nodes	16 nodes, 64 NVIDIA H200 SXM (4/node), AMD EPYC 9335 (64 cores/node), 1,536 GB/node
NPU nodes	7 nodes, 28 Furiosa RNGD (4/node), AMD EPYC 9355 (32→64 cores/node), 1,536 GB/node
Interconnect	400 Gb/s InfiniBand per node
Storage	GPFS / IBM Storage Scale ~4.1 PiB usable (NVMe flash + HDD tiers)
Scheduler/OS	Slurm; Ubuntu baseline; Lmod modules; Apptainer containers

2 Computing Resources and Queues

2.1 Slurm User Queues

Res.	Queue	Nodes	GPU/ NPU/ Core	Wall (h)	Sub/ run	Max resource (user)
GPU	gpu_trial	gpu01–gpu02	8 GPU	12	4/ 2	1 GPU
GPU	gpu_junior	gpu03–gpu04	8 GPU	24	4/ 2	2 GPU
GPU	gpu_general	gpu05–gpu08	16 GPU	48	5/ 3	1 nodes (4 GPU)
GPU	gpu_collaborative	gpu09–gpu11	12 GPU	48	5/ 3	3 nodes (44 GPU)
GPU	gpu_container	gpu12–gpu16	20 GPU	72	6/ 4	1 node (4 GPU)
NPU	npu_trial	npu01–npu01	4 NPU	12	6/ 4	1 NPU
NPU	npu_junior	npu01–npu02	8 NPU	24	6/ 4	2 NPU
NPU	npu_general	npu03–npu05	12 NPU	48	6/ 4	1 nodes (4 NPU)
NPU	npu_collaborative	npu06–npu07	8 NPU	48	6/ 4	2 nodes (8 NPU)
CPU	cpu_general	NPU worker nodes (7)	32 cores/node	224	6/ 4	see note

- cpu_general draws CPU cores from the 7 NPU worker nodes. Each node has 64 cores; 1 core per accelerator (4) + 2 cores for OS services are reserved → 58 allocatable cores per node (7 × 58 = 406 cores).
- A higher-priority special partition may be created for a limited duration when required; it is named after the project and has a defined start and end time (Slurm reservation).
- Additional user queues may be proposed by users after a trial period.
- Additional queues such as gpu_all, npu_all, or all can be created when full-node execution is required.

2.2 Module and Container Naming and Grouping Convention

Item	Convention	Example
Modulefile	name/version[-toolchain][-variant]	gromacs/2024-foss2023a-cuda12

Item	Convention	Example
Default version	one default per name (.modulerc)	default = 2024
Grouping	Lmod hierarchy Core -> Compiler -> MPI; family() for exclusive tools	family("mpi"), family("compiler")
Category label	compiler/ mpi/ lib/ app/ ai/ tool	whatis: Category: ai
Container image	app-version[-variant].sif	pytorch-24.06-cuda12.sif
Container module	wrapper module points at the SIF	apps/pytorch/24.06
Restricted software	group-gated MODULEPATH per project/category	/apps/modulefiles/restricted/<group>

Rules. lowercase names; semantic versions; never overwrite a published version (add a new one); record the toolchain in the version string; announce additions/deprecations via change control.

3 SRU Accounting

SRU is a usage-accounting unit (not a billing unit). A fixed initial grant is deducted per running job, proportional to resources held × wall-clock time, using the factors below (from V3.8 Section 3).

Resource	SRU factor (per hour)	Note
H200 GPU	1	1 SRU = 1 GPU-hour
RNGD NPU	0.1	0.1 SRU = 1 NPU-hour
AMD CPU core	0.01	0.01 SRU = 1 CPU-core-hour

3.1 Initial SRU Allocation per Project Category

Sizing basis: Initial SRU ≈ A × D × 24 × U (A = max concurrent accelerators, D = duration days, U = planning utilization). Adopted allocations:

Category	Max compute	Duration	Max users	Scratch h/ f	Initial SRU
Trial	1 GPU	2 weeks	1 (cohort)	small	100
Junior	up to 2 GPU	1 month	2	20 TB / 5 TB	500
General	up to 1 node (4 GPU)	3 months	5	40 TB / 10 TB	5,000
Collaborative	by negotiation	up to 1 year	many	by negotiation	12,500 (base)

Notes. NPU and CPU time draw from the same SRU pool at the factors above, so accelerator-light work consumes far less. Grants are shared team pools and can be topped up on extension. Per-project-type limits derive from the General Policy; values here may be revised after a user trial.

4 Storage and Quotas

4.1 Filesystem Partitioning

Mount	Total (TiB)	NVMe	HDD	Notes
/home	1,000	0	700	user home directories
/apps	100	0	100	NVMe reserved for container images
/scratch_f	350	350	–	NVMe flash scratch
/scratch_h	2,950	0	2,950	HDD scratch

4.2 Quota and retention

Space	Quota	Scope	Retention/ Purge
/home	30 GB	per user	retained 90 days post-project
/scratch_h (HDD)	40 TB	per project	purge 15 days after last access
/scratch_f (flash)	10 TB	per project	purge 5 days after last access

Not an archive. AKF storage is for active computation. At the end of the project, the account retains data management access for 90 days (Section 6); after that, the project data is deleted, and the quota is released.

5 Identity and naming

5.1 Username

Element	Value	Notes
country code (2)	ISO-3166 alpha-2	bn kh id la my mm ph sg th

Element	Value	Notes
		vn tl
affiliation code/ type (1)	r / e / i	research/ educational/ industrial
sequence (5)	00001–99999	assigned once; retained for continuous status

Convention: <cc><inst><5-digit> → e.g. idr00142 (Indonesia, research, #00142). 8 characters, POSIX UID from a per-country range. Regular user starts from 00101

5.2 Project/Group Code

Convention: <cat><yy><apasti><nnnn> → e.g. g26a0031 (General, 2026, APASTI area a, #0031). The group (gid) name equals the project ID; enables per-project accounting and quotas.

Format	Example	Fields
<cat><yy><apasti><nnnn>	g26a0031	cat = t/j/g/c/d; yy = 2-digit year; apasti = area letter a–j; nnnn = 4-digit serial; gid = project ID

Element	Value	Notes
Project category (1)	t/ j/ g/ c/ d	trial / junior / general / collaboration / development
Year (2)	yy	Two last digit of year (20yy)
APASTI area (1)	a/ b/c/ d/ e/ f/ g/ h/ i/ j/ x	See below
sequence (4)	0001-9999	assigned once;

Code	APASTI area	Code	APASTI area
a	AI, Data Science & Advanced Computing	f	Energy Transition & Net-Zero
b	Climate, Disaster Risk Reduction & Digital Twins	g	Ocean, Marine Science & Blue Economy
c	Health, Bioinformatics & Precision Medicine	h	Cybersecurity & Trusted Digital Infrastructure
d	Food Security, Smart Agriculture & Water Systems	i	Innovation, Tech Commercialisation & Digital Economy
e	Advanced Materials,	j	Regional HPC Infrastructure &

	Semiconductors & Manufacturing		Human Capacity Dev.
		x	No specific area

6 Account lifecycle and retention

A person has one persistent identity that may join several time-limited, extendable projects. At the end of the project, the account retains data management access for 90 days, then is purged and set to dormant, and can be reactivated within two years; otherwise, it is closed and the person must re-register.

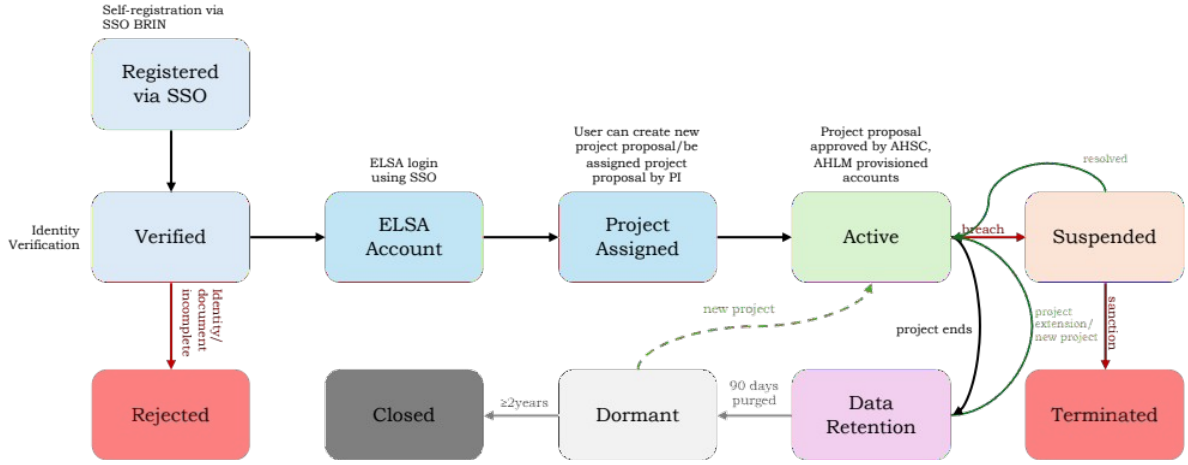

Figure 1. Account state machine.

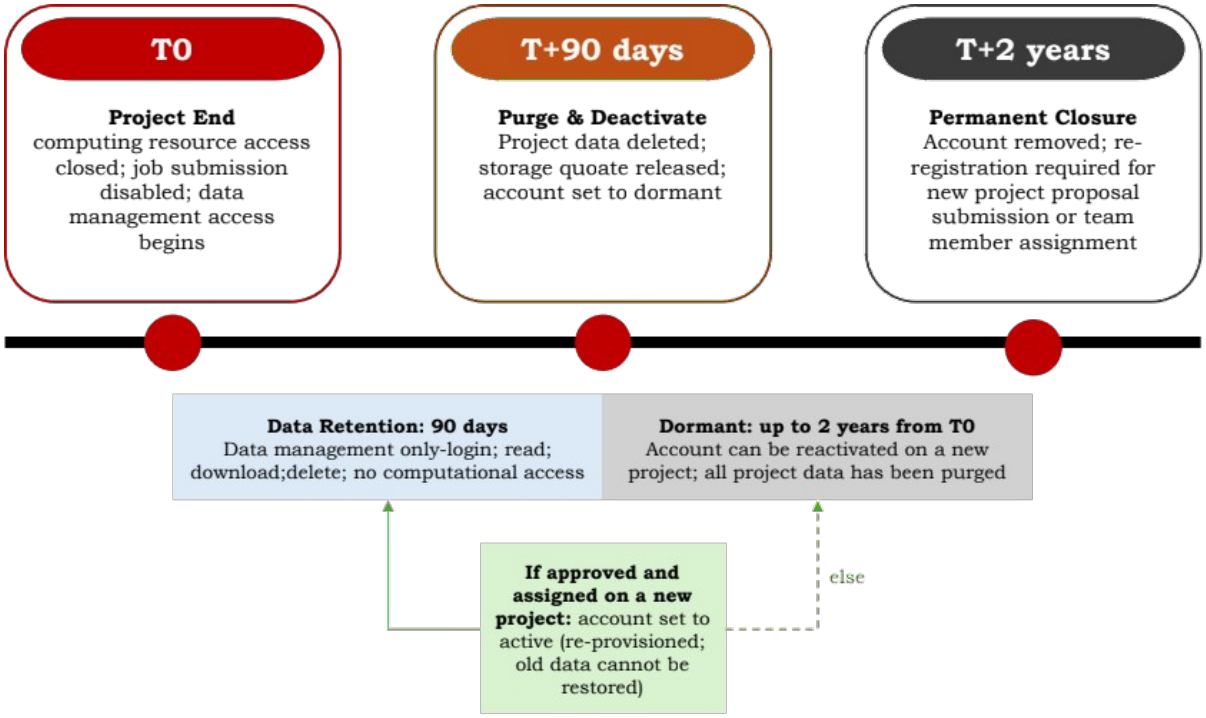

Figure 2. Post-project timeline: 90-day retention, dormancy, reactivation within 2 years, closure.

7 Administration and Access Control

The AKF uses the general system administrator model with privileged access management. Separation of duties and specific roles are managed within the system administrator.

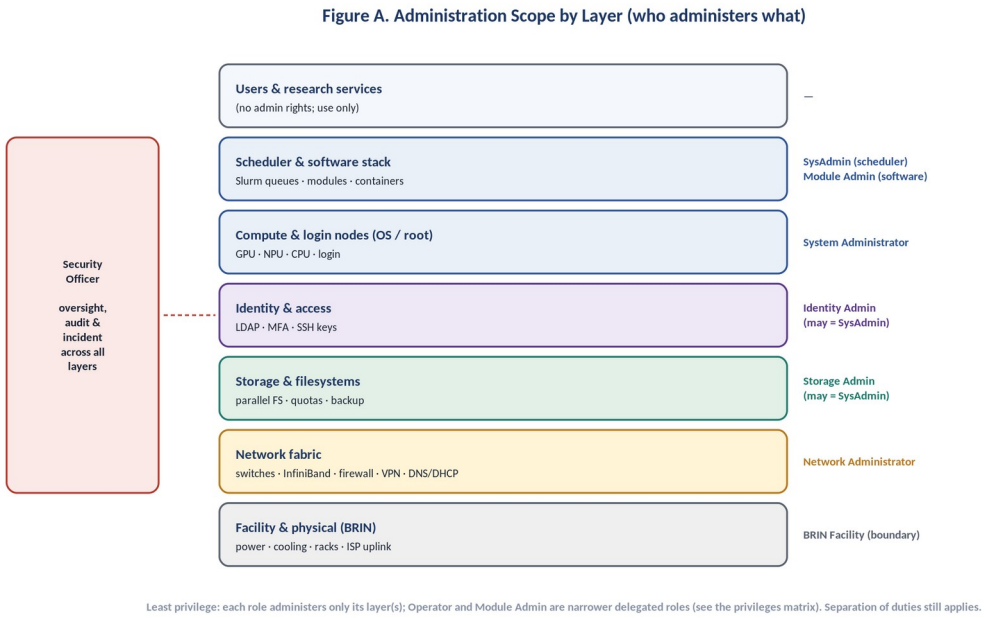

Figure 3. Administration scope by layer.

7.1 Administrator assignments

Role	Members	System group
System Administrator	Andria, Ihsan, Wiarno, Rizki, Arida, Lina, Iftitah	akf-sysadmin
Network Administrator	Ditya	akf-netadmin

Separation of duties. Within the System Administrator team, duties are shared but no single person both makes and solely approves the same production change, or is the sole reviewer of logs they can write. Change control and immutable audit logging apply to both roles

7.2 Responsibilities, Privileges, and Boundaries

7.2.1 System Administrator

- **Responsibilities**: OS install/patch/harden on compute and login nodes; Slurm (queues, partitions, fair-share, accounting, jobs); software, modules and containers (Lmod, Spack, container registry); storage — GPFS filesets, /home per-user and scratch per-project quotas, snapshots, backup/restore, scratch purge; identity and accounts (LDAP, MFA, SSH keys, user/project lifecycle); monitoring, user support, and incident response; owns the immutable audit logs.
- **Privileges**: root on compute/login nodes (named account + time-boxed elevation); full administration of Slurm, software, storage, identity, and accounts.
- **Boundaries**: no network-device/firewall/VPN/DNS administration (Network Administrator); production changes via change control; user research data on a need-to-know basis.

7.2.2 Network Administrator

- **Responsibilities**: Ethernet/VLANs; InfiniBand fabric (subnet manager, pkeys); firewall and segmentation; VPN/remote access; DNS/DHCP/NTP; network monitoring; ISP-uplink coordination with BRIN.
- **Privileges**: administer network devices, InfiniBand, firewall, VPN, DNS/DHCP (akf-netadmin group); firewall changes under change control.

- **Boundaries**: no compute root; no Slurm/software/account/storage administration; no user data; firewall changes with security impact need Security Officer sign-off; cannot disable security logging.

7.3 Privileges matrix

A = administer, O = operate/monitor, R = read, n = need-to-know, — = none.

Privilege / domain	Sys	Net
Root on compute / login nodes	A	—
Slurm config (partitions, QOS, fair-share)	A	—
Slurm job ops (drain/resume, requeue, cancel)	A	—
Software / modules / containers	A	—
User / project accounts	A	—
Identity (LDAP / MFA / SSH keys)	A	—
Storage filesets & quotas	A	—
Backup & restore	A	—
Network, InfiniBand, firewall, VPN, DNS	—	A
System monitoring (view)	O	O
Security / audit logs	A	R
Suspend account / isolate host (incident)	A	O
User research data	n	—

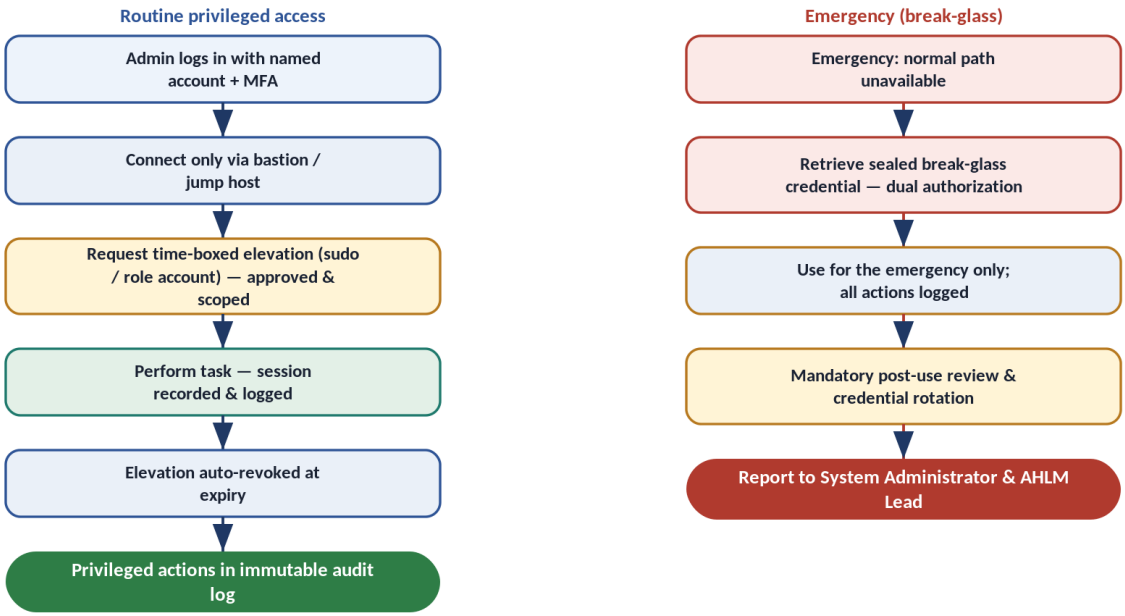

Figure 4. Privileged access (named account + MFA → bastion → time-boxed elevation) and break-glass.

No shared admin accounts. Every privileged action is attributable to a named person and recorded in a log admins cannot alter.

8 Access and Login

Users reach the AKF two independent ways — the myAKF portal and direct SSH — so that a failure of the portal does not block access. Every path requires two-factor authentication (2FA), and SSH uses the user's public key (registered at account creation) so no password is needed.

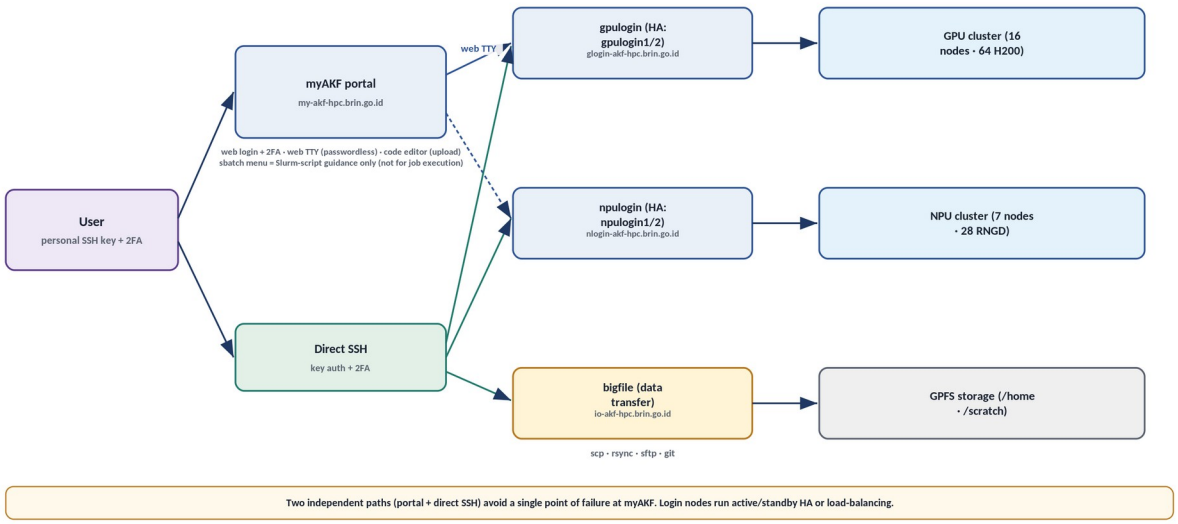

Figure 5. AKF access architecture: myAKF portal and direct SSH, both with 2FA.

8.1 myAKF portal

- Web login with 2FA at the portal.

- myAKF for system administrators to manage users, configure, and monitor the cluster system.
- Web TTY to the gpulogin and npulogin nodes — opened directly as the logged-in user (their registered SSH key), so no username/password re-entry.
- Code editor menu for document/file upload.
- The sbatch menu is used to guide Slurm-script writing only — not to execute or submit jobs.

8.2 Direct SSH access

Direct SSH is opened (by IP allow-list) to three entry points, each via SSH with public-key authentication and 2FA:

Entry point	Purpose	Domain / IP Address
gpulogin	log in to the GPU cluster	glogin-akf-hpc.brin.go.id / 103.224.136.239
npulogin	log in to the NPU cluster	nlogin-akf-hpc.brin.go.id / 103.224.136.240
bigfile	data transfer: scp / rsync / sftp, and git clone/push	io-akf-hpc.brin.go.id / 103.224.136.233
myAKF portal	web access (Section 8.1)	my-akf-hpc.brin.go.id / 103.224.136.241

- gpulogin and npulogin run in high availability — one server active at a time; gpulogin1/gpulogin2 and npulogin1/npulogin2 may both be activated (active/standby or load-balanced) to spread user load.
- SSH public-key authentication means users do not enter a password; 2FA is still enforced.
- The bigfile server is the recommended path for large files or many files, and for code (git).

8.3 Login banner

On login (portal and SSH), the banner shows the user their context:

- User and project group membership.
- Link to the "how to use" documentation.
- Scratch file locations (/scratch_h/<project>/<user>, /scratch_f/<project>/<user>).

- Accessible resources (queues/partitions the user may use).
- Running/queueing jobs.
- Account / project expiration date.

9 Security, Data Protection, and Continuity

- **Acceptable Use Policy**. Every user digitally accepts the AUP before activation and re-accepts it upon material revision, for each new project, and upon reactivation.
- **Data protection**. Data classification (public/internal/confidential/restricted); compliance with Indonesia's PDP Law and the user's jurisdiction; minimize/anonymize personal data; IP in results stays with the researchers.
- **Security operations**. Named credentials + MFA, least privilege, export/sanctions screening, immutable audit logs, incident response.

Business Continuity & Disaster Recovery — life-safety first, controlled shutdown, then recovery in dependency order:

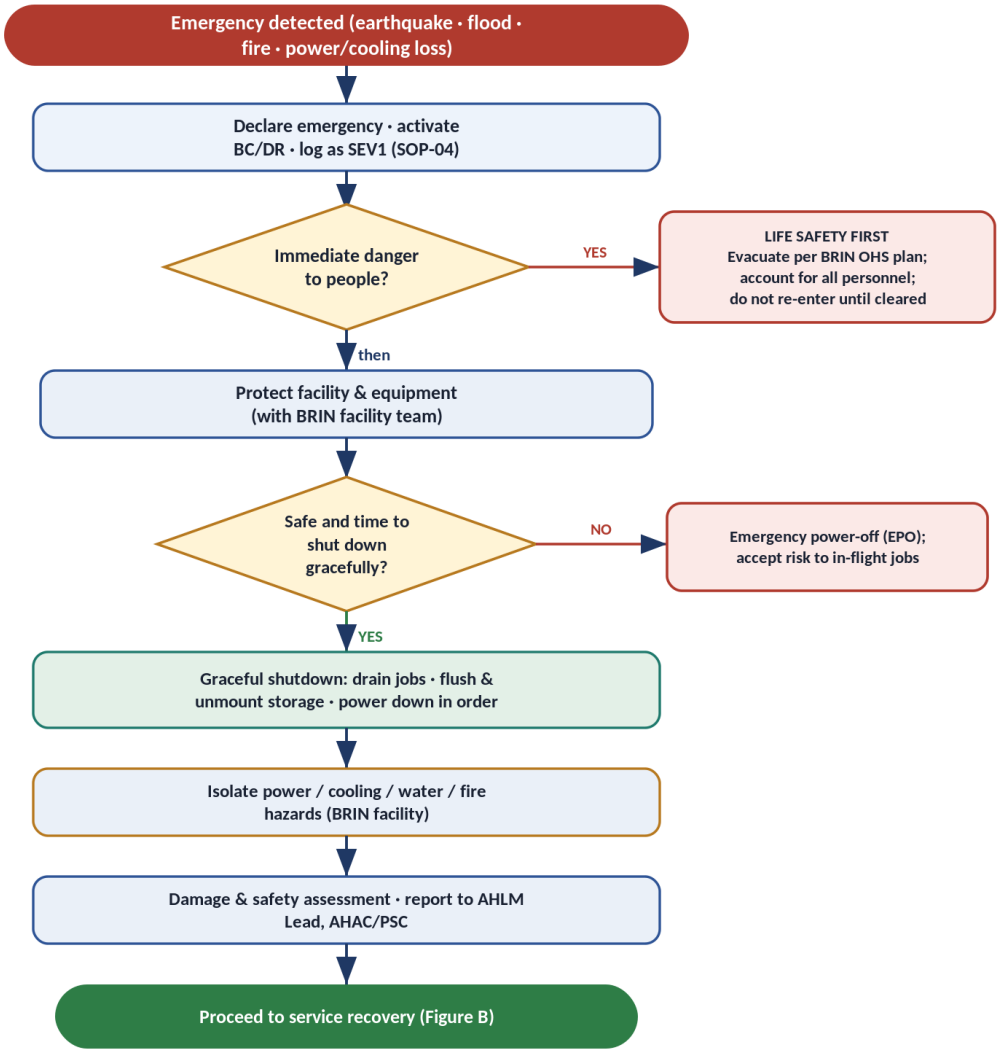

Figure 6. Emergency response and shutdown decision.

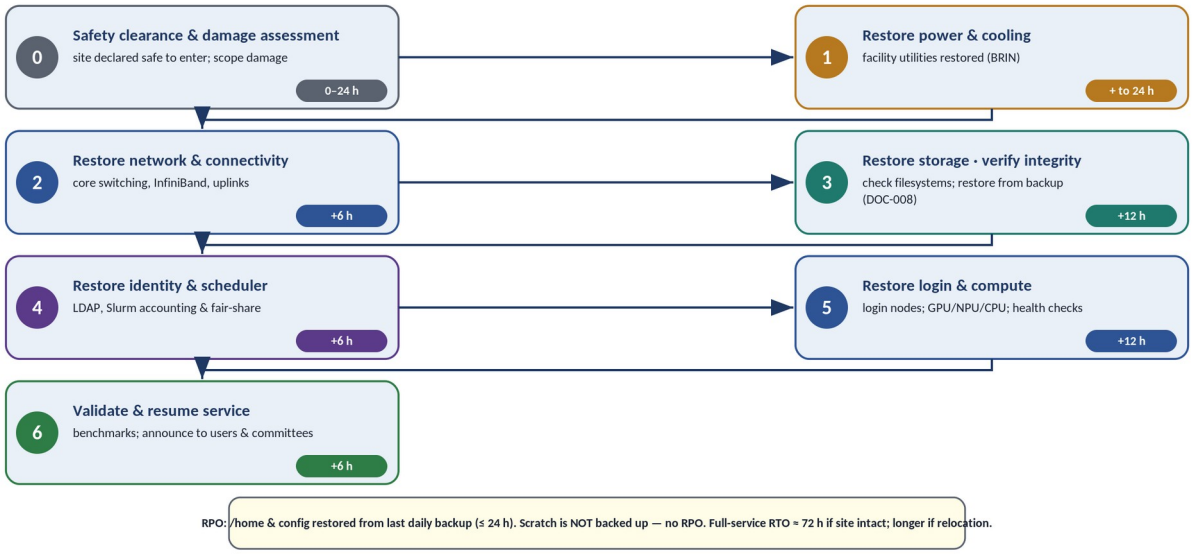

Figure 7. Service recovery sequence with indicative RTO; RPO ≤ 24 h for /home & config.

10 Change Control and Review

- All production changes (configuration, software, quotas, queues, firmware) are subject to change control; emergency changes are documented retrospectively.
- This policy is reviewed at least annually and after any major change; SRU allocations, utilization factors, and additional queues are revised after user trials.
- Indicative values (SRU allocations, host-nation reserve, RTO/RPO) are confirmed by the AHSC and against the as-built (DOC-001).

End of policy (Draft v1.0, endorsed baseline). Some items are still under development.