agenticscale.ai

AGENTICGRID · STANDARD AGT-GFM-1.0

GridFlex Methodology v1.0

Issuing, verifying and retiring GridFlex Tokens (GFT) for place-matched, work-conserving flexibility from AI compute at the grid edge.

STATUS	PILOT NETWORK	ISSUE DATE
Pilot standard — for comment	Partner network (to be announced)	9 November 2026
REGISTRY OPERATOR	**APPLIES TO**	**IMPLEMENTED IN**
AgenticGrid (AgenticScale.ai)	SAII Vault sites on flexible connections	GridFlex Tokens Register v0.1

Proposed scheme. This standard is published by AgenticGrid for pilot use and consultation. Organisation names identify the relevant network, regulator, research or technology for a region and do not imply participation, certification or endorsement. Methodology components that are patent-pending are described by function only.

Contents

0	Document control
1	Purpose and scope
2	Definitions
3	Unit of issue and serials
4	Token classes
5	Site eligibility
6	Baseline
7	Measurement and verification
8	Deferral integrity
9	Issuance and lifecycle
10	Value, shares and attribution
11	Stacking and double counting
12	Registry, signatures and data
13	Claims code
14	Pilot plan — partner network
15	Governance and versioning
16	Implementation in the Register
A	Annex A — Token record schema
B	Annex B — Worked example

0 · Document control

FIELD	VALUE
Document code	AGT-GFM-1.0
Version / status	1.0 — Pilot standard, released for comment with the first network pilot
Owner	AgenticGrid (Registry Operator), a subsidiary of AgenticScale.ai
Review cadence	End of each pilot phase (§14); minor revisions published as v1.x with a change log
Change log	v1.0 — initial pilot release (9 Nov 2026). Codifies the conventions used for shadow minting at the pilot site since 1 July 2026. Editorial revision of 8 Nov 2026: the instrument is named the GridFlex Token (GFT); serials and class codes carry the GFT- prefix. No change to issuance rules.
Related	GridFlex Tokens Register (product); GridFlex one-page overview; SAII Vault V2000 technical specification

Design principle. A GridFlex Token (GFT) must be **invoice-grade**: every figure on it reconciles to a meter and to money on a bill. Where carbon markets failed on baselines, double counting and claims drift, this standard fixes each by construction — modelled counterfactuals (§6), a single-use claim and stacking disclosure (§11), and a claims code (§13).

§1 Purpose and scope

This standard defines how flexibility delivered by an AI compute site to its local electricity network is measured, verified, recorded as a token, valued, attributed and claimed. Its purpose is to make that contribution legible to four readers at once — the network, the regulator, the compute buyer and the community — and to reduce the energy cost of compute without shifting cost to other customers.

1.1 Scope

- **Sites.** Containerised AI compute facilities (SAII Vault class, 0.2–10 MW) connected under a flexible connection agreement, dynamic operating envelope or an equivalent network instruction channel.
- **Networks.** Distribution and, where applicable, transmission network operators that signal constraints, operating envelopes or scheduled windows to the site. The pilot jurisdiction is a partner distribution network area in the NEM, to be announced. Other jurisdictions adopt by reference under §15.
- **Tokens (GFT).** Five classes (§4) covering event response, standing availability, load shaping, avoided augmentation and community resilience.

1.2 Out of scope

- Renewable origin or carbon content of electricity (LGC, REC, GO and similar energy attribute certificates). GridFlex Tokens record verified network response, not energy source.
- Wholesale market participation rules. Where a site also earns wholesale or ancillary revenue, §11 requires disclosure, not exclusion.
- Site safety, connection engineering and metering installation standards, which are governed by the network's connection requirements.

§2 Definitions

TERM	MEANING
GridFlex Token (GFT)	The unit of issue under this standard: a signed record of one verified flexibility event — or, by class, one availability period, shape window or annual aggregate — at one connection node (§3). “Token” and “GFT” are used interchangeably in this document.
Site	A compute facility at one connection point, identified by a site ID and a connection node.
Connection node	The register identifier for the connection point (network code · zone substation · voltage · feeder). A register key, not the network's asset name.
Network signal	An instruction or forecast from the network with a signal ID: a constraint event, an operating-envelope change, or a scheduled shape window.
Window	The start time and duration over which response is measured.
Baseline	The modelled counterfactual load of the site had no signal been received (§6).
Metered load	Interval load recorded by revenue-grade metering at the connection point.

Response	Baseline minus metered load, averaged over the window (kW), and its energy over the window (kWh). For the Shape class, metered minus baseline (load absorbed).
Deferral	Compute work paused or shifted during the window and completed later.
Deferral integrity (DI)	The ratio of deferred work completed within the deferral horizon to the energy reduced (§8).
Registry Operator	AgenticGrid, which operates the register, verifies records and mints tokens.
Co-signatory	The network that confirms the signal and window for a token (§7, §12).
Holder	The party to which a token is attributed and who may retire it (§10).
Retirement	The single, final act of claiming a token; a retired token cannot be transferred or claimed again.

§3 Unit of issue and serials

One GridFlex Token (GFT) records one verified event at one connection node — or, by class, one availability period, one shape window, or one annual aggregate. A token carries the response (kW), the energy (kWh), the window, the node, the class, the deferral integrity, the value shares and the holder attribution. The minimum issue size is **50 kWh** of verified response; smaller events are logged but not tokenised.

```
Serial: GFT-{country}-{network code}-{YYMM}-{sequence}
Example: GFT-AU-PILOT-2609-0003 → Australia · Partner network · September 2026 · third token that month
```

Each token also carries a record hash (SHA-256 of the canonical record), a meter-data hash, the identifier of the baseline model version used, and — for Response tokens — a deferral proof hash (§8).

§4 Token classes

CLASS	NAME	WHAT THE TOKEN RECORDS	TYPICAL BUYER
GFT-R	Response	Verified reduction during a network-signalled constraint window. Carries DI and may carry the `work-conserving` attribute.	Hyperscalers, enterprises, startups (per-workload attribution)
GFT-S	Standby	Standing availability to curtail a declared kW for a declared number of hours per month, whether or not called. Aligns with capacity-style instruments (e.g. NSW Peak Reduction Certificates; ERCOT Controllable Load Resource).	Investors, lenders, networks (planning)

GFT-H	Shape	Verified load *increase* in a scheduled window to absorb local surplus — minimum-demand and rooftop-solar conditions. DI = 1.0 by construction.	Networks, government, buyers with batch workloads
GFT-A	Avoided	An annual aggregate of Response tokens at one node demonstrably deferring a network augmentation, valued against the network's published avoided cost where one exists.	Networks, investors, regulators
GFT-X	Resilience	Hours of backup supplied from site storage to a nominated community load during an outage or planned event.	Insurers, councils, networks

Why Shape matters where rooftop solar is high. Networks with high rooftop solar experience some of the lowest midday operational demand in the world. A compute site that reliably absorbs load at 11:00 is doing the network a service that most demand-response instruments cannot describe. GFT-H is the class built for that.

§5 Site eligibility

A site may mint tokens when the Registry Operator has confirmed all of the following and recorded them in the site register:

- ☐ **Connection instrument.** A flexible connection agreement, dynamic operating envelope or an equivalent documented channel by which the network can signal the site, with a named signal interface (API, DOE feed, dispatch message or scheduled window).
- ☐ **Metering.** Revenue-grade interval metering at the connection point at 5-minute resolution or finer, with read access for the Registry Operator and the network.
- ☐ **Workload telemetry.** Per-job start, pause, resume and completion timestamps; GPU utilisation; the dispatchable share of load (the load that can be deferred without breaching a customer deadline).
- ☐ **Local model.** A calibrated model of the site's load against its own schedule and of the connection node, versioned and referenced on every token (§6). Model internals are not disclosed; model version and calibration date are.
- ☐ **Flex tier.** A declared tier (Flex-4h, Flex-day, Shape) that sets the maximum window and minimum DI the site offers (§10.4).
- ☐ **Priority protection.** A documented rule set for the workloads that are never deferred (priority queue, deadlines), so that response never breaches an agreed service level.

METERING RESOLUTION
≤ 5 min
Connection-point interval data

DEFERRAL HORIZON
24 h
For deferred work to complete

MINIMUM ISSUE
50 kWh
Per token

§6 Baseline

The baseline is the load the site would have drawn had no signal arrived. It is **modelled, never declared**. The Registry Operator computes it from the workload schedule the site had committed to before the signal, the site's metered load in the intervals preceding the window, and the calibrated local model. The model adjusts for temperature-driven cooling load and for any scheduled maintenance.

```
B(t) = f( schedule(t), L(t−10Δ … t−Δ), T(t), model v )   for each interval t in the window
Reported with an uncertainty band u (kW): the 90 % interval of the model's residuals at the node over the calibration period
```

- A token is issued only if the average response exceeds **2 × u**. Smaller responses are logged as “below threshold”.
- The baseline model version and calibration date are recorded on every token. Re-calibration is required at least quarterly and after any change to site capacity.
- For the Shape class, the baseline is the site's committed firm schedule for the window; absorbed load is metered minus baseline.
- Self-declared baselines, “average of last 10 days” baselines without model adjustment, and baselines set after the window are not permitted.

§7 Measurement and verification

```
Response kW   R = mean over window of ( B(t) − L(t) )   [Shape: L(t) − B(t)]
Energy kWh    E = Σ ( B(t) − L(t) ) × Δt
Standby         A = declared dispatchable kW available for H hours in the period, verified by test event or by actual response
```

ROLE	VERIFIES	EVIDENCE
Registry Operator	Meter data integrity, baseline computation, response and energy, DI, value shares, serial uniqueness	Meter hash, model version, record hash
Network co-signatory	That the signal was issued, its ID and window; optionally the connection-point telemetry	Network signature on the record
Site operator	Workload telemetry and deferral proof; priority-protection compliance	Operator signature; job proof hash
Independent verifier (from Phase 3)	Annual sample audit of baselines, DI and stacking disclosures	Audit report in the public register

§8 Deferral integrity

Compute is the only large load class that can prove its flexibility was a *shift* rather than a *shed*: a paused job finishes later. Deferral integrity makes that proof part of the token.

```
DI = ( energy of deferred work completed within the horizon H = 24 h ) ÷ ( energy reduced during the window E )
Work-conserving attribute: DI ≥ 0.85  ·  Shed flag: DI < 0.85  ·  Shape class: DI = 1.0 by construction
```

- Deferred work is identified by job IDs paused or re-scheduled during the window; completion is evidenced by workload telemetry and summarised as a job proof hash on the token.
- A Response token with the shed flag remains valid and may be retired, but may not carry the work-conserving attribute or the corresponding claim (§13).
- Sites report a rolling 30-day average DI. A site whose average falls below 0.80 is moved to “review” and may not mint Response tokens until re-calibrated.

§9 Issuance and lifecycle



STEP	RULE	TIMING
Mint	Registry Operator mints once baseline, response and DI are computed and the record hash is signed by the site operator and the registry.	≤ 5 business days after the window
Co-sign	The network confirms the signal ID and window. A token is “Verified” only when co-signed. Standby tokens are co-signed monthly.	≤ 15 business days after mint
Retire	The holder retires the token to make its claim. Retirement is final, timestamped and public (holder may be anonymised by segment).	Any time before expiry
Expire	Unretired tokens expire and are marked “Expired” in the public register; they cannot be retired thereafter.	24 months after mint
Void	A token is voided if meter data is withdrawn, the signal is disputed by the network, or a stacking breach is found. Voids are public with a reason code.	At any time, with audit entry

§10 Value, shares and attribution

10.1 Gross value

```
V = network tariff avoided (demand, critical-peak or flexible-connection rebate) +
avoided-augmentation value (where the network publishes an avoided cost, $/kVA
deferred) + market revenue attributable to the window (wholesale demand response,
ancillary services) + attribute premium (what a holder pays for the claim)
```

For the pilot, V is computed from the network's flexible-connection tariff terms and an illustrative avoided-cost schedule agreed with the network; market revenue and attribute premium are reported as zero until Phase 2.

10.2 Default shares



Shares are defaults for the pilot and are set per site in the site register. The community dividend is paid into a network tariff rebate or a community energy fund nominated by the network, and reported per customer per year (§14).

10.3 Attribution

Response and Shape tokens are attributed to holders **pro-rata to the GPU-hours actually flexed** in the window, from workload telemetry. A holder's **GF%** is the share of its GPU-hours that were flexed over the reporting period. Unattributed tokens remain in the pool and may be attributed later or retired by the operator.

10.4 Flex tiers

TIER	WINDOW	MINIMUM DI	PRICING BASIS
Firm	—	n/a	Base price; no tokens
Flex-4h	≤ 4 h per event	≥ 0.85	Base less the customer share of verified savings, expressed as a price band
Flex-day	≤ 24 h	≥ 0.95	As above; larger band
Shape	Scheduled windows	1.0	As above; largest band; runs when the network wants load

§11 Stacking and double counting

- **Stacking is permitted, concealment is not.** A token may coexist with payments for the same response — a network flexibility contract, a wholesale demand response mechanism, a capacity-style certificate — but every concurrent payment or scheme must be listed in the token's stacking disclosure.

- **The claim is single-use.** A token's claim (§13) may be made once, by the holder who retires it. Where the same response is also recognised under another attribute scheme, the GridFlex claim must say so.
- **Energy attributes are separate.** GridFlex Tokens record verified network response, not renewable origin. It may be presented alongside hourly energy attribute certificates as the “place-matched” leg of a source-, time- and place-matched report, but never as a substitute for them.
- **Duplicate detection.** The registry rejects any record whose node and window overlap an existing token of the same class.

§12 Registry, signatures and data

The register is operated by AgenticGrid. Each token is a signed data object carrying three signatures — site operator, network co-signatory and registry — plus the holder's retirement signature. The audit log is append-only.

REGISTER VIEW	VISIBLE
Public register	Serial, class, status, network, region, node ID, window, response kW/kWh, DI attribute, holder (or segment if anonymised), methodology version, signatures (hashes only), retirement timestamp and claim statement.
Holder portal	The holder's tokens, GF%, savings credited, statements for download, retire and claim.
Network portal	Signals issued, co-sign queue, delivered vs requested, avoided-augmentation record, community dividend accrued.
Regulator view	Issuance vs retirement, verification lag, void and duplicate checks, dividend by region, full audit log, methodology change log.
Restricted	Meter interval data, workload telemetry, model internals and commercial terms — available to the network and an appointed verifier under agreement; never published.

§13 Claims code

Permitted claims

- "*N* kWh of verified [class] flexibility, place-matched to [network area], under GridFlex Methodology v1.0."
- "Work-conserving (DI ≥ 0.85)" — Response tokens with the attribute only.
- "GF% of our GPU-hours were flexed in [period]."
- "Community dividend of [amount] accrued to [network area] customers."

Not permitted

- Carbon-neutral, net-zero, renewable or "offset" claims on the strength of a GridFlex Token.
- Claiming a token that has not been retired, or has been retired by another holder.
- Describing a shed-flagged token as work-conserving.
- Aggregating tokens across networks into a single "grid" claim without naming each network area.

The Registry Operator generates the claim statement at retirement and prints it on the token; holders use that text or a shortened form of it.

§14 Pilot plan — partner network

The methodology is proven first in a partner network area where a flexible-connection framework already exists, operating envelopes are dynamic, and minimum-demand conditions make the Shape class valuable from day one. The partner network is named when the pilot is agreed. The pilot runs on one site and one node, and is monitored weekly in the Register's pilot view.

PHASE	NAME · WINDOW	SCOPE	EXIT CRITERIA
0	Design & shadow minting · Jul–Oct 2026	Signal interface and metering confirmed; local model calibrated; events minted under draft conventions to tune the baseline.	Model uncertainty u ≤ 6 % of site load; signal interface tested end to end
1	90-day minting pilot · Nov 2026–Feb 2027	Every signalled event at the pilot site minted under v1.0; network co-signs; weekly monitoring; Standby token issued monthly.	KPIs below met for the phase
2	Attribution & dividend · Mar–May 2027	Holder attribution on invoices; GF% reporting; community dividend paid through the network's nominated channel; claims code live.	Dividend paid and reported; ≥ 3 holders retiring tokens
3	Second network · from Jun 2027	Adoption at a second network area; cross-jurisdiction serials; independent verifier appointed; v1.1 published.	Second co-signatory live; verifier's first audit

14.1 Phase 1 KPIs

KPI	TARGET	MEASURED AS
Verified events	≥ 20	Response, Shape and Resilience tokens co-signed by the network in the phase
Average deferral integrity	≥ 0.90	Mean DI of Response tokens minted in the phase
Verification lag	≤ 15 days	Mean days from mint to network co-signature
Double-claims and duplicate serials	0	Registry duplicate check; retirement uniqueness
Community dividend accrued	≥ A$25,000	Sum of the community share across pilot tokens
Midday surplus absorbed	≥ 60 MWh	Energy of Shape tokens in the phase

14.2 Monitoring cadence

PER EVENT

Signal → response → baseline and uncertainty → mint within 5 business days → deferral proof attached when deferred work completes.

WEEKLY

Pilot view reviewed with the network: events, DI distribution, co-sign queue, voids. Shape windows scheduled against the minimum-demand forecast.

MONTHLY

Standby token; savings shares settled to invoices; community dividend accrued; stacking disclosures checked.

END OF PHASE

KPI report; methodology issues logged for v1.1; go / no-go for the next phase agreed with the network.

14.3 Roles in the pilot

- **Partner network** — signals and windows; co-signature; nomination of the community dividend channel; participation in weekly review.
- **AgenticGrid** — registry operation, baseline modelling, minting, monitoring view, KPI reporting.
- **Site operator (AgenticScale)** — response delivery, workload telemetry, deferral proofs, priority protection.
- **Research partners** — review of baseline method and DI definition; publication of pilot findings.
- **Regulators (state regulator, AER, AEMO)** — informed at Phase 1 start; invited to the Phase 1 KPI review; regulator view available throughout.

§15 Governance and versioning

- **Registry Operator** — AgenticGrid owns the methodology and the register, and is accountable for verification and for the public change log.

- **Technical Advisory** — research partners review the baseline method, DI definition and any change that affects issuance; their review is noted in the change log.
- **Network co-signatories** — each adopting network signs an adoption note recording its signal interface, tariff basis and dividend channel.
- **Independent verifier** — appointed before Phase 3; samples at least 10 % of tokens annually and publishes a summary in the register.
- **Versioning** — minor versions (v1.x) clarify without changing issuance; a major version (v2) changes issuance and applies only to windows after its effective date. Every token names the version it was issued under.
- **Alignment** — definitions are kept consistent with industry flexible-data-centre metrics as they are standardised, and with the deliverability concepts in evolving corporate energy accounting guidance, so that GridFlex can serve as the place-matched leg of a buyer's report.

§16 Implementation in the Register

The GridFlex Tokens Register is the product that implements this standard. Each clause maps to a feature; the mapping is also shown on the Register's Methodology page.

CLAUSE	REQUIREMENT	REGISTER FEATURE
§3	Unit, serial, hashes	Token record with serial, record hash, meter hash, model version, job proof; minimum-issue rule enforced at mint.
§4	Classes	Class on every record with class-specific fields; filters and charts by class.
§5	Eligibility	Site register (node, vault class, dispatchable MW, tier, status); only operating sites can mint.
§6	Baseline	Baseline, metered and uncertainty on every event token; below-threshold events not issued.
§7	M&V and roles	Lifecycle states; network co-sign queue in the Network portal; operator mint queue.
§8	Deferral integrity	DI with work-conserving / shed badge; DI distribution on the pilot page; deferral proof hash.
§9	Lifecycle and timing	Status badges, signature timestamps, verification-lag KPI, expiry note; actions for mint, co-sign and retire.
§10	Value and attribution	Shares per token; holder attribution, GF% and savings by month in the Buyer portal; community dividend by region.
§11	Stacking	Stacking disclosure on every token; duplicate check in the Regulator view.
§12	Registry views	Public register, holder portal, network portal, regulator view with audit log and CSV export; anonymised holders by segment.
§13	Claims code	Claim statement generated at retirement and printed on the token.

§14	Pilot plan	pilot page: KPIs vs targets, weekly response, DI distribution, phases, monitoring cadence, event log.

Current release. Register v0.1 runs on a specimen (synthetic) dataset for concept review. Production wiring — signal interface, metering feed, workload telemetry, signatures and the public API — is scoped for Phase 1 of the pilot.

A Token record schema

FIELD	TYPE	DESCRIPTION
serial	string	GFT-{country}-{network}-{YYMM}-{seq}; unique
hash	hex	SHA-256 of the canonical record (first 16 bytes shown)
class	enum	GFT-R · GFT-S · GFT-H · GFT-A · GFT-X
site · node · network · region	ids	Site ID; connection node ID; network code; region code
signal_id	string	Network's signal reference (event, envelope change, schedule or availability period)
window_start · window_h	datetime · hours	Local time; duration (hours available for Standby)
baseline_kw · metered_kw · uncertainty_kw	kW	Modelled counterfactual, metered load, 90 % uncertainty band (§6)
response_kw · kwh	kW · kWh	Average response and energy over the window (§7)
di · deferred_completed_h · job_proof	ratio · hours · hex	Deferral integrity, completion time of deferred work, proof hash (§8)
status · retired_at	enum · datetime	Lifecycle state (§9)
gross_value · customer_share · community_dividend · operator_share	currency	Value and shares in local currency (§10)
buyer · gpu_hours_flexed	id · number	Holder attribution and basis (§10.3)
stacking[]	list	Concurrent payments and schemes (§11)
signatures{operator, network, registry}	objects	Signer, timestamp, signature (§12)
methodology · twin_model · meter_hash	strings	Version references and data hashes

B Worked example (specimen)

Token `GFT-AU-PILOT-2609-0003` — Response at the pilot node, 27 September 2026, 13:30, 0.5 h window.

```
Baseline B = 3,790 kW (u = ±120 kW)
Metered L = 2,290 kW
Response R = 3,790 − 2,290 = 1,500 kW
 (> 2u ✓)
Energy E = 1,500 × 0.5 = 750 kWh  (≥ 50
kWh ✓)
```

```
Deferred work completed within 24 h =
652 kWh-eq
DI = 652 ÷ 750 = 0.87 ≥ 0.85 → work-
conserving ✓
Completed 19.5 h after the window
```

```
Gross value V = 750 kWh × A$0.72 =
A$540.00
Customer 50 % = A$270.00
Community 15 % = A$81.00
Operator 35 % = A$189.00
```

```
Attribution: 1,070 GPU-hours flexed →
holder Meridian Logistics AI
Status: Verified (co-signed 15 Oct
2026) → retired with claim:
“750 kWh of verified Response
flexibility, place-matched to the pilot
network area (Pilot Precinct), work-
conserving, retired by Meridian
Logistics AI under GridFlex Methodology
v1.0.”
```

All names and figures in this example are synthetic and taken from the specimen register. The A$/kWh network value is illustrative and is replaced by the network's tariff and avoided-cost terms in the pilot.

Proposed service. Savings depend on site eligibility, tariff terms and verified delivery. Token scheme proposed by AgenticScale; organisation names identify the relevant network, regulator, research or technology for a region and do not imply participation, certification or endorsement. © AgenticScale.ai 2026.