

### PRODUCT SPEC SHEET · AI COMPUTE BUFFER

Power buffer for AI & GPU clusters. Sized for modern compute transients.

*DRAFT SPEC SHEET — engineering review pending. Specific specifications may change before final release.*

CETUS is XEYAR's dedicated power buffer for AI and GPU compute environments. Modern hyperscale AI workloads produce millisecond power transients that collapse conventional Li-ion UPS cycle life. CETUS uses the same electrostatic SuperCap architecture as DRACO, tuned for the burst-heavy draw profiles of GPU clusters.

SUPERCAPACITOR OPTION	SODIUM-ION OPTION
GPU transient dispatch · 500,000+ cycles	Sustained compute backup

**Built for the AI Boom** · A single rack of H100-class GPUs can swing hundreds of kilowatts in milliseconds. CETUS is designed to handle that load profile natively — absorbing spikes that would otherwise destabilize upstream PDUs.

## SPECIFICATIONS

Attribute	Value
Response Time	<1 ms
Discharge Rate	5C continuous, higher peak
Cycle Life (SuperCap)	500,000+
Form Factor	Rack or row-integrated (TBD)
Deployment Pattern	Per-rack, per-pod, or row-level
Integration	Upstream of GPU PDU
Typical Sizing	Configured to GPU draw profile
Certifications	UL 9540, UL 1973 (confirm scope)

## IDEAL FOR

- Hyperscale AI training clusters
- HPC research facilities
- Pod-level power conditioning
- GPU inference data centers
- Crypto & high-density compute
- Retrofit of legacy UPS-limited rooms

**Ready to Specify Your CETUS Project?**

[Request a quote](#) · [Book a demo](#) · [Talk to engineering](#)

**Quote:** [xeyar.com/contact/?topic=quote-cetus](https://xeyar.com/contact/?topic=quote-cetus) **Demo:** [xeyar.com/contact/?topic=demo-cetus](https://xeyar.com/contact/?topic=demo-cetus) **Sales:** [sales@xeyar.com](mailto:sales@xeyar.com)