

Data Center Flexible Load Initiative (DCFlex): Flex MOSAIC™ and Demonstration Summary

Executive Overview

The EPRI DCFlex research team, together with their members, are developing a technology neutral structure for describing and evaluating the flexibility capabilities of large loads, including modern data centers called Flex MOSAIC™. By establishing a shared language and transparent performance expectations, the framework will enable utilities, system operators, regulators, and developers to make faster, more consistent, and more confident decisions in designing and planning for flexibility. The framework was informed by the first cohort of demonstration projects which show that high-performance AI data centers can operate as power-flexible, grid-responsive resources without degrading service levels. These EPRI led projects demonstrate that flexibility strategies tailored to the availability, scale, and duration of grid events may enable a practical pathway to faster data-center interconnection, reduced system costs, and improved grid reliability. EPRI is continuing to collect and analyze feedback on the framework and is developing additional tools and approaches to integrate the standardized flexibility classes into the analytical approaches needed to realize flexibility’s value in grid planning.

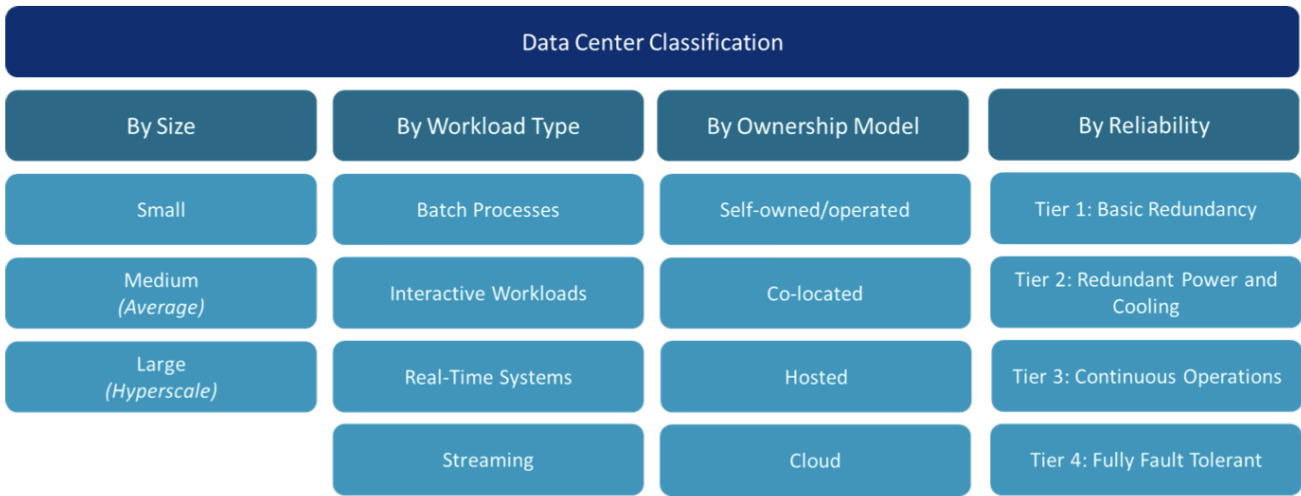
Background

The Data Center Flexible Load Initiative (DCFlex) is EPRI’s multiyear, broad stakeholder R&D initiative demonstrating how data centers can operate as flexible, grid-supportive assets, not solely as passive electricity consumers. It will drive a fundamental shift in how data centers are designed, interconnected, and operated, supporting grid reliability, resilience, affordability, and the rapid growth of AI-driven compute demand. Uniquely, DCFlex is solely focused on data centers, bringing developers, hyperscalers, IPPs, ISO/RTO, technology providers, utilities, finance, and construction under one collaborative program structure. The consortium currently has over 65 members and is continuing to expand. Members of DCFlex assign up to 2 representatives per company per workstream to Technical Advisory Groups (TAGs). These groups drive the technical scope and execution of the workstream by meetings 2–3 times per year, including an annual in-person meeting of all advisors together, plus monthly office hours with key EPRI researchers.

<p>Workstream 1</p> <p>Flexible Data Center Designs</p> <p>Reference architectures enabling DCs to operate flexibly as grid resources</p>	<p>Workstream 2</p> <p>Transformational Utility Programs</p> <p>Market structures and incentive frameworks capturing flexibility value</p>	<p>Workstream 3</p> <p>Grid Planning & Operational Flexibility</p> <p>Standardized taxonomy, load forecasting, and interconnection processes</p>	<p>Workstream 4</p> <p>Data Center Informed Energy Supply</p> <p>Technology analysis and modeling for viable energy supply options</p>	<p>Workstream 5</p> <p>Distribution Data Centers</p> <p>Frameworks for distribution-connected, grid-interactive data centers</p>
--	---	---	---	---

In December 2025, Work Stream 1 EPRI team published a paper called Grid Flexibility Needs and Data Center Characteristics which provided definitions and classification system for data centers.

Table 1: Data center characteristic classification



Source: EPRI, Grid Flexibility Needs and Data Center Characteristics, 2025, available here: <https://www.epri.com/research/products/3002031504>

Once a data center has been characterized, there are a number of capabilities to provide flexibility. Data centers comprise common components and systems that may each be flexible; however, similar to the grid, the constraints within which they operate may influence their potential to respond. A high level view of various types of flexibility capabilities is available in Table 2.

Table 2: Flexibility Potential from Data Center Assets and Subsystems

Underlying Asset	Potential Capability	Response Time	Duration Sustained	Availability
On-site generation				
On-site battery storage				
On-site thermal storage				
Cooling systems				
Standby generation				
UPS				
Compute				

Source: EPRI, Grid Flexibility Needs and Data Center Characteristics, 2025, available here: [\[Link\]](#)

Flex MOSAIC™






In March 2026, the DCFlex team released the Flex MOSAIC™ Framework. It provides a unified, technology-neutral structure for describing and evaluating the flexibility capabilities of large loads, including modern data centers based on power system requirements—such as congestion relief, peak reduction, balancing, and frequency response. Rather than focusing on specific technologies or assets, the framework defines flexibility in terms of performance characteristics, including Notification time, Duration, Frequency of use, Depth of load

adjustment, Ramp behavior, and Availability. These characteristics are grouped into five flexibility classes, each representing a clear, consistent set of capabilities that represent needs that are aligned with most of the high value use cases where flexibility can unlock issues that constrain grid connection and that can be understood and used by utilities, operators, and regulators across regions.

- What flexibility is available?
- What flexibility is needed?
- How flexibility can be relied upon in interconnection, planning, and operations

The goal of introducing Flex MOSAIC™ is to increase transparency and predictability around flexibility capabilities to operationalize flexibility and facilitate large load interconnection processes.

Table 3: Flex MOSAIC classification

Class	Description
	<i>Critical Peaking</i> Responds to <i>rare</i> scarcity events of 5 hours or less
	<i>Peaking</i> Responds to <i>frequent</i> scarcity events 5 hours or less
	<i>Prolonged</i> A + B + responds to <i>prolonged</i> events, up to 24+ hrs
	<i>Fast</i> A + B + provide fast response with short notice
	Fully <i>grid responsive</i>

Source: EPRI, Flex MOSAIC Briefing Note, 2026, available here: <https://restservice.epri.com/publicattachment/97236>

The Flex MOSAIC™ framework was publicly released on March 24, 2026 with an Open Letter of support from over 30 Reliability & Standards Authorities, Grid & Market Operators, Electric Utilities, Technology, Data Centers, OEM’s and Financial & Advisory organizations. Today, over 50 organizations, including NERC, MISO, SPP, Exelon, Rayburn Electric Cooperative, Google, Meta, NVIDIA, Siemens and many others have added their names to the [Open Letter](#).

In-production Tests

AI-driven data centers are increasing electricity demand faster than grid infrastructure can expand. In many areas of the world, this mismatch can lead to multi-year interconnection queues, especially in constrained regions. This challenge is often compounded by planning frameworks that treat data centers as inflexible “firm loads,” requiring the grid to be built for rare worst-case peaks. This approach often drives up consumer costs and slows the deployment of AI infrastructure.

In December 2025 EPRI, National Grid (NG), Nebius, and Emerald AI conducted a series of tests to evaluate the ability of data centers to respond to real-world grid events as a part of DCFlex demonstrations. The tests were designed using the Flex MOSAIC™ framework developed under EPRI’s DCFlex initiative. Several tests also included unannounced, back-to-back grid events, demonstrating that flexible data centers can operate as a

continuously dispatchable grid resource, at least for short periods of time, rather than a one-time demand response asset. Once defined, the test events were entered into a utility-accessible cloud interface that allowed for the control of the size, duration, ramp rate, and timing of load reductions. Details of the developed test events can be seen in Table 4.

Table 4. Flex MOSAIC Aligned Demonstration Test Events

Flex Moasic Alignment	Name	Notification time	Load reduction (% of instantaneous)	Ramp down time	Duration of load reduction	Ramp up time
A	Short but strong peak	15 min	40%	10 min	1 h	15 mins
B	Ramp support	15 min	25%	10 min	2 h	15 min
C	Extended heat wave or wind drought	3 h	10%	15 mins	10 h	15 mins
D	Less aggressive fast response	<1 min	30%	<1 min	30 mins	2 min
D	Aggressive fast response	<1 min	40%	<1 min	15 mins	<1 min
E	DA peak	3 h	30%	20 min	3 h	20 min
E	TX maintenance	3 h	25%	15 mins	4 h	15 mins

Source: Power-Flexible AI Factories: A UK-First Demonstration of Grid-Responsive AI Infrastructure, 2026. Available here: <https://restservice.epri.com/publicattachment/96969>

The controlled AI cluster, which was a high-compute production cluster that was taken out of active production operation for this testing, achieved 100% compliance with more than 200 requested power targets and ramp-rate constraints. ***Load reductions of up to 40% were delivered with sub-minute response times, while service level agreements for high-priority workloads were preserved throughout all tests.*** Test results also included second-to-minute controllability, autonomous operation, explicit workload prioritization, and measurable, verifiable compliance. Apart from power flexibility, the UK test also included responses to five-minute carbon intensity signals, the cluster dynamically adjusted load, and reduced associated emissions by more than 10% during the test period.

Project impact

This UK demonstration tests fed into the design of the Flex MOSAIC™ framework and showed that flexible data centers can dynamically adjust data-center compute loads and electricity consumption in response to grid conditions. This flexibility is enabled via the use of AI-trained software that controls the scheduling of training jobs, natural pause points, GPU-level power capping, and the ability to shift compute temporally or geographically. Once completed, demonstration results showed that high AI performance and grid responsiveness are compatible, and that AI-enabled data centers may provide a practical pathway to faster grid connections, lower system costs, and improved reliability—supporting both economic competitiveness and energy system affordability.

Findings from this small-scale demonstration showed that AI data centers can be treated as flexible, dispatchable resources rather than inflexible loads. Thus, this approach may enable earlier interconnections using existing grid headroom, reduce the need for costly overbuilding of infrastructure, and improve system resilience during both short-term disturbances and prolonged stress events.

If implemented at a larger scale, the findings suggest multiple benefits, including flexible (non-firm) interconnection agreements that allow limited curtailment and market structures that compensate compute loads for providing grid flexibility may be possible. In addition, dynamic tariffs could reward load shifting that aligns electricity consumption with periods of high renewable generation availability.

Next Steps

The second cohort of DCFlex demonstration projects was [announced](#) in February 2026 and includes 6 additional data center demonstration projects in the US and Europe. Building on these results and to evaluate the benefits at scale, DCFlex plans to conduct flexibility testing at additional demonstration projects, such as the Aurora Project. In late 2026, this U.S.-based demonstration will target nearly 100 MW of power-flexible AI load, further validating that software-defined flexibility can scale from kilowatt-level pilots to hyperscale, gigawatt-class AI facilities.

The EPRI DCFlex Workstream 1 team is collecting feedback on the Flex MOSAIC™ framework from DCFlex members and select external organizations. The team is developing tools for base features and reference design support elements. By the end of the year, the team intends to rollout the final reference design as well as structural incentive program designs and tools for adoption of the framework.

EPRI's DCFlex Workstream 4 team introduced a practical framework to help power system planners evaluate how much additional load, particularly from rapidly growing data centers (DCs), can be integrated without expanding generation, storage, or transmission infrastructure. Unlike existing studies, this framework provides a how to approach for planners and will allow planners to evaluate how the Flex MOSAIC™ flexibility classes may be deployed to help realize a system's headroom potential.

By establishing a shared language and transparent performance expectations, the Flex MOSAIC™ framework will enable utilities, system operators, regulators, and developers to make faster, more consistent, and more confident decisions in designing and planning for flexibility.

For More Information

[DCFlex.epri.com](https://dcflex.epri.com)

Irene Danti Lopez, DCFlex Workstream 1 Research lead, Frank Sharp, Technical Executive for Demonstrations, dcflex@epri.com

Jessica Lin, Global Marketing and Communications Lead, jlin@epri.com