Capacity Allocation Marketplace

by Nick Clark | Published April 25, 2026

What It Specifies

Capacity-allocation transactions carry: capacity-holder identity, capacity description (location, time, quantity, class), capacity-user identity, allocation-authority signature. The transaction enters the architecture's audit retention.

Allocation rules can be class-specific: priority-based, market-based, hybrid. The architecture admits the rules through commodity-class declaration; implementations apply the rules within the architectural framework.

Why It Matters Structurally

Current capacity-allocation systems face structural challenges: cross-holder coordination, allocation-rule complexity, audit complexity for regulators.

Architectural capacity allocation produces structural support. The architecture handles the structural primitives; participants transact within the framework; allocation authorities participate as credentialed observers.

How It Composes With Mesh Operation

The architecture defines the capacity-allocation taxonomy, the allocation-rule specification format, and the federation across capacity-holders. Implementations apply the architecture; capacity participants transact structurally.

Composition with other features. Cross-jurisdictional capacity allocation, byzantine-robust capacity coordination, and dispute mechanism for allocation conflicts all build on the capacity allocation primitive.

What This Enables

Infrastructure operators (utilities, data centers, transportation networks) gain structurally-supported capacity allocation. Civilian shared-infrastructure operations gain the same.

The architecture also supports emerging capacity classes. As new capacity-allocation patterns emerge (real-time grid balancing, dynamic data-center capacity, dynamic transportation capacity), the architecture admits the new classes through declared specification.