

Services on the primary Admin Node

- **SSM – Server Status Monitor**

Monitors the OS and hardware state of the node (CPU, memory, disk, interfaces, etc.) and reports that health information back into the grid for display and alerting in Grid Manager.

There is an SSM instance on every grid node, not just the Admin Node.

- **AMS – Audit Management System**

Collects and writes all audited events and transactions to log files on the Admin Node (for example, S3 operations, ILM decisions, configuration changes).

Those audit logs are what you feed into external tools (Splunk, syslog, etc.) for compliance and troubleshooting.

- **CMN – Configuration Management Node**

Holds and manages grid-wide configuration state (connectivity, protocol features, grid tasks).

There is exactly one CMN in the entire grid, which is why the node that hosts it is the **primary** Admin Node and the place where many maintenance procedures must be initiated.

- **NMS – Network Management System**

Backs the Grid Manager UI: it provides the monitoring, reporting, and configuration functions that show up in the browser interface.

When you click around in Grid Manager to view node status, alerts, or configuration, you're effectively talking to NMS (plus Prometheus, mgmt-api, etc.) on an Admin Node.

Services on a Storage Node

- **SSM – Server Status Monitor**

Same role as on the Admin Node: monitors OS and hardware metrics for that Storage Node and reports them to the management layer.

- **ADC – Administrative Domain Controller**

Maintains the grid's topology and provides internal authentication services; other services (like LDR and CMN) query ADC to discover where things live and to validate connections.

At each site you must have at least three Storage Nodes running ADC to give a resilient "control-plane directory" for that site.

- **DDS – Distributed Data Store**

Acts as the interface to Cassandra for object metadata; when the grid stores or reads object metadata, it goes through DDS.

DDS is responsible for coordinating metadata reads/writes and node participation in the Cassandra cluster on that Storage Node.

- **LDR – Local Distribution Router**

Handles the actual object I/O pipeline: ingest, retrieve, replicate, erasure-coded fragment management, and internal traffic for object data across the grid.

When a client puts or gets an object (after any load balancer), the work of storing, moving, and retrieving the bits is driven by the LDR service on Storage Nodes.