Design Principles for Cloud native Egeria

Principle	Description	
Run in Kubernetes	Egeria should run in Kubernetes as a set of Deployments.	
Microservices	Each deployments represents only a single server. A cohort can consist of multiple deployments which are as indepedent as possible.	
Stateless	The deployments of Egeria are stateless in themselves. That means the state (e. g. registration of server, metadata) is stored elsewhere and accessed by the individual pods.	
Store config in the environment	The server configuration is stored in the environment or an external config file. Configuration is applied on startup. There is no need to trigger configuration after startup	
Logs	The audit logs of the servers are written to console by default to enable Kubernetes to capture and display them. If an additional audit log is required, it must be part of the configuration.	
Observable	The services can be observed on runtime. tbd.	
Disposable	A server is disposable, that means it can be started up and shut down in few seconds. Shutdown is done gracefully on a SIGTERM signal. This allows flexible movement of servers between nodes in a Kubernetes environment.	
Concurrency / horizontal scaling	As servers are stateless they can be scaled horizontally by running multiple instances of the same server.	