We refer here to a project TL Consulting recently undertook with a NSW government agency as a case study. The existing setup of Oracle SOA Suite (https://www.oracle.com/au/middleware/technologies/soasuite.html) suffered from critical issues, which were:

- Manual High Availability: Oracle's HA reference model is an Active-Standby that relies on two identical groups of nodes: one Active Node for serving workloads; the other Standby Node ready to take over in case the Active Node malfunctions. The switch from the Active to the Standby nodes is a manual switch which requires a Standard Operating Procedure to be run by an engineer on duty
- High Cost: Since the platform was moved to the cloud using "Lift and Shift", the cost of provisioning and maintaining the different environments for the platform exceeded that of running it on VMs. Having two groups of identical nodes with only one of them actively serving the workloads was a substantial waste of expensive infrastructure resources and compute power
- Scalability: Adding another node group to the platform to handle extra load was an onerous process, which required extensive planning prior to implementation

FIGURE 2: ORACLE SOA HA MODEL

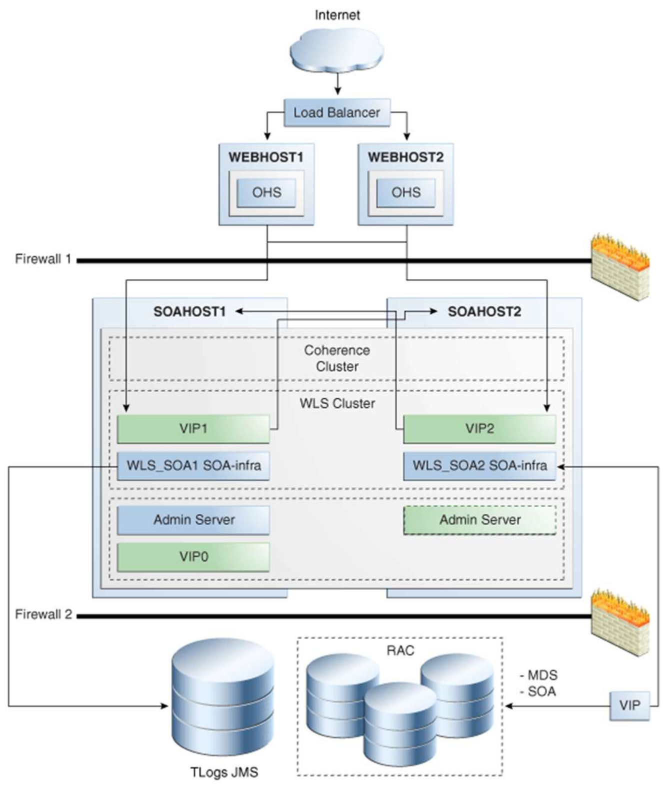

Such was the problem statement that needed to be resolved using application modernisation. In addition to the three criteria for success described, other criteria were mandated, such as:

- ✓ Compliance with the government security and data privacy policies
- ✓ Compliance with the government managed data centre models

Solution Design Stage

Firstly, the application was evaluated in a form of a POC or an MVP, intended to assess the feasibility to decouple the application components and explore different architecture design options.

Secondly, the application's data needed to be evaluated. Persistent data goes against the principle of mutable docker containers which are the basic building blocks of cloud native. Therefore, application data needs to be classified into the following categories:

- Category 1: Data that can be included into the docker image and is updated only when the corresponding image is updated. Data in this category is usually data that can be created during the application start-up such as temporary data or automatic configurations.
- Category 2: Data that is stored on external volumes that are mounted into the docker containers which can get updated at any time during container runtime. This comprises mostly configuration data.
- Category 3: Application data that is stored in external database systems which can get updated frequently and needs to be preserved. This includes master and transactional data.

Since proper categorisation and handling of data directly impact the application update workflows, careful consideration needs to be taken during the architecture design to evaluate how this aspect of the solution is going to be implemented and maintained.

The following diagram shows the architecture of the proposed design that TL Consulting proffered to the government agency, highlighting. The design shows how the different components of the system were decoupled and built as separate docker containers. These containers were executed in Kubernetes pods to maximise infrastructure cost efficiency.

FIGURE 3: SOA SOLUTION ARCHITECTURE