Oracle Cloud Services maintains a redundant and resilient infrastructure designed to maintain high levels of availability and to recover services in the event of a significant disaster or disruption. Oracle designs its Cloud Services using principles of redundancy and fault-tolerance with a goal of fault-tolerance of a single node hardware failure.

Oracle has two separate data centers that function as primary and secondary sites for Oracle Cloud Services. Oracle will commence the disaster recovery plan under the policy upon its declaration of a disaster, and will target to recover the production data and use reasonable efforts to re-establish the production environment at the secondary site.

City of St. Charles production standby (secondary site) environment will reside in a separate data centre from City of St. Charles primary site. The secondary site is built with the exact same infrastructure than the primary site; no low cost data center is used for the standby server.

For major regional jurisdictional areas (e.g. the United States or the European Union) Oracle operates both a production and secondary site within that region. City of St. Charles data is replicated in physically separate facilities in order to restore full services in the event of a disaster at a primary site. Backups are for Oracle's sole use in the event of a disaster.

Oracle provides for the recovery and reconstitution of its Cloud Services to the most recent available state following a disaster.

Oracle has established alternate processing sites to accommodate full operating capability in the event of loss of service at a primary facility. For each Cloud Service, Oracle maintains separate Disaster Recovery Plans that describe recovery procedures. Disaster recovery operations apply to the physical loss of infrastructure at Oracle facilities. Oracle reserves the right to determine when to activate the Disaster Recovery Plan. During the execution of the Disaster Recovery Plan, Oracle provides regular status updates to customers.

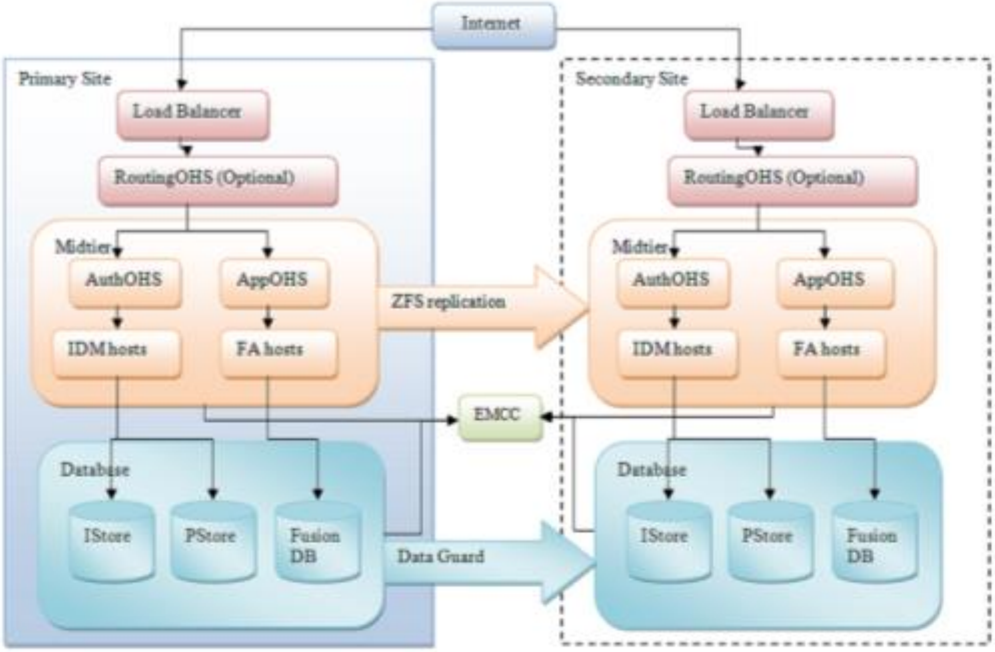

Overview of our Disaster Recovery topology

Our DR solution has two sites, primary and secondary. Under normal operating conditions, the primary site is running and is in active mode, while the secondary site, which replicates the primary site, is in passive mode.

All activities at the primary site are replicated to the secondary site at both the middle-tier and database levels. The secondary site remains in passive mode until switchover or failover occurs. City of St. Charles requests from the Internet are routed via a DNS hosting service to the primary site. After switchover or failover, the DNS service is updated to transparently route new user requests to the secondary site. Our disaster recovery solution for Oracle Cloud services is based on the following Oracle technologies:

- **Oracle Data Guard**: Replicates data from a primary database (DB) to a secondary DB.
- **ZFS Storage Replication**: Replicates file systems for middle-tier machines from a primary site to a secondary site.
- **Oracle Data Guard Broker**: Provides coordinated management of DB pairs in a DR environment.
- **Enterprise Manager Cloud Control** (EMCC): Displays status and statistics collected from EM agents and executes tasks scheduled from the console.
- **Oracle Site Guard**: Manages the DR switchover operations for the entire application stack. Oracle Site Guard, an Enterprise Manager plug-in, is a GUI-based and command line interface-based tool.

The primary and secondary sites are connected to the same EMCC instance, which monitors and orchestrates the switchover activity through the Oracle Site Guard EM plug-in. Oracle Site Guard uses EM agent jobs as a tool to perform the necessary health checks and actual switchover on the hosts where the agents are running.

Recovery time objective (RTO)

RTO is Oracle's objective for the maximum period of time between Oracle's decision to activate the recovery processes under this Policy to failover the service to the secondary site due to a declared disaster, and the point at which Customer can resume production operations in the standby production environment. If the decision to failover is made during the period in which an upgrade is in process, the RTO extends to include the time required to complete the upgrade. The RTO does not apply if any data loads are underway when the disaster occurs.

The RTO objective is **12 hours** from the declaration of a disaster.

Recovery point objective (RPO)

RPO is Oracle's objective for the maximum period of data loss measured as the time from which the first transaction is lost until Oracle's declaration of the disaster.

The RPO objective is **1 hour** from the point of service loss.

Note: the RTO and RPO do not apply to Customer customizations that depend on external components or third-party software. During an active failover event, non-critical fixes and enhancement requests are not supported. Customer will be solely responsible for issues arising from third party software and customizations (CEMLIs) to Oracle programs.