Executive Summary

In terms of business continuity and disaster recovery planning, IDVerifact's mission is to ensure the availability of critical business functions and Information Technology (IT) operations within acceptable timeframes should a disaster or outage affect the ongoing availability of IDVerifact's Digital Identity Software as a Service offering. IDVerifact's plan is developed to meet the criteria and sound practices defined in the ISO 22301:2019 Security and Resilience – Business Continuity Management Systems – Requirements. The planning approach has the following features:

- Proactively evaluating risks and mitigating their effect through the implementation of loss control countermeasures, e.g., cyber risks & internet security provisions
- Reviewing the plan quarterly
- Frequently upgrading business continuity strategies, the supporting disaster recovery strategies and their associated action plans – part of our product development process
- Annually exercising sections of the business continuity plans and annual testing of disaster recovery plans
- Effectively providing awareness communications and training related to emergency operations, mitigation measures and recovery responsibilities for all associates

Business Continuity Plan Brief

IDVerifact Digital Identity platform is deployed on Google Cloud Platform and subscribes to the GCP high availability configuration to reduce downtime when a zone or instance becomes unavailable. IDVerifact Digital Identity Platform is distributed across multiple zones wherein the standard SLA From Google Cloud Platform commits to provide 99.99% uptime.

IDVerifact Digital Identity platform is composed of three main components: the backend application component, frontend application component & database hosted on Google Cloud SQL.

The platform is deployed on Kubernetes cluster composed of 9 nodes each having a backend and frontend application component pod. These 9 clusters are distributed across three availability zones from GCP.

The IDVerifact Digital Identity Platform data is replicated across three zones in real-time so in the event of a disaster in one specific zone instance, the other two instances will be available, and the business operations will get executed as usual.

Depending on the volume of calls during the outage on specific zones, all the processing will be handled by the remaining active nodes in respective zones which could experience a little latency.

The network diagram, figure 1 below, provides a high-level overview of the IDVerifact Digital Identity Platform. The high availability configuration, figure 2 below, provides a high-level overview of the IDVerifact Digital Identity Platform deployment architecture and figure 3 below outlines the IDVerifact Digital Identity Platform database redundancy schema.

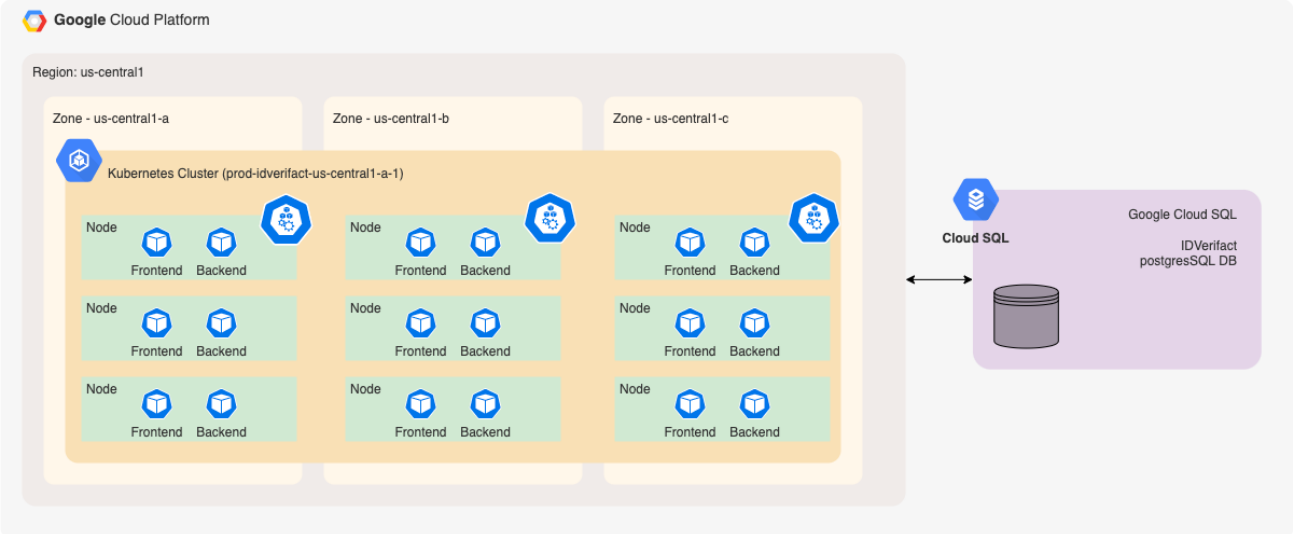

Figure 1 IDVerifact Digital Identity Platform Network Diagram

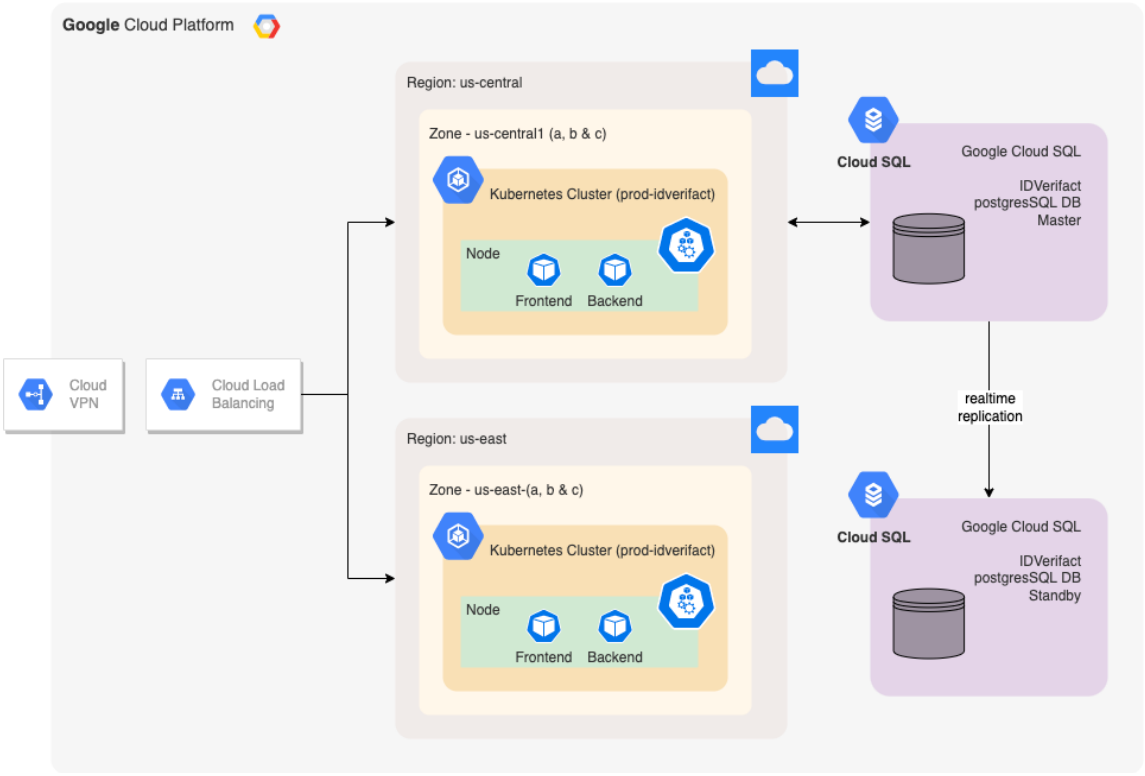

Figure 2 IDVerifact Digital Identity Platform Deployment Architecture

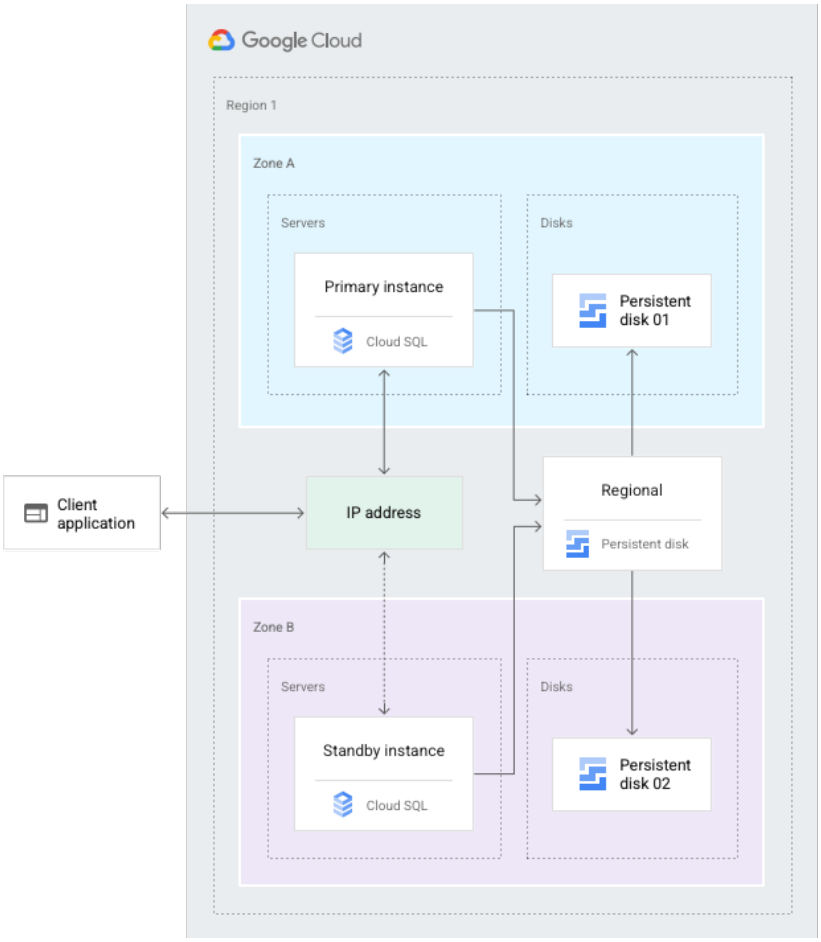

Figure 3 IDVerifact Database Redundancy

The IDVerifact Digital Identity Platform does not persist any transaction data other than logging the call control data (i.e. call unique identifier, composite API ID, call timestamp and call status) which is used for monitory the call volumes for a given API.

IDVerifact platform persists the information pertaining to the organization and its users that accesses the platform to compose a Digital Identity API to ensure that digital identity APIs are accessible to authorized organizations with a unique API key generated by the authorized user of a given organization.

Kubernetes Node failure

In case a specific node is not available the Google Cloud Platform Kubernetes Engine would automatically retry to bring back the node. Following are the steps that are executed during the node failure and recovery:

- Post node failure, in about a minute, Kubernetes engine will report “NotReady” state.

- In about 5 minutes, the states of all the pods running on the “NotReady” node will change to either “Unknown” or “NodeLost”. This is based on pod eviction timeout settings, the default duration is five minutes.
- Irrespective of deployments (“StatefuleSet” or “Deployment”), Kubernetes will automatically evict the pod on the failed node and then try to recreate a new one with old volumes.
- If the node is back online within 5 – 6 minutes of the failure, Kubernetes will restart pods, unmount, and re-mount volumes.
- If in case an evicted pod gets stuck in “Terminating” state and the attached volumes cannot be released/reused, the newly created pod(s) will get stuck in “ContainerCreating” state. There are 2 options now:
    - Either to forcefully delete the stuck pods manually (or)
    - Kubernetes will take about another 6 minutes to delete the Volume Attachment objects associated with the Pod and then finally detach the volume from the lost Node and allow it to be used by the new pod(s).

In summary, if the failed node is recovered later, Kubernetes will restart those terminating pods, detach the volumes, wait for the old Volume Attachment cleanup, and reuse (re-attach & re-mount) the volumes. Typically these steps would take about 1 ~ 7 minutes.

Any calls to Digital Identity APIs are replicated to Google Cloud SQL instance in a different availability zone in real-time which minimizes loss of data in the event of a disaster occurred in the primary region.

Google Cloud Platform – Security and Data Protection

Google’s focus on security and protection of data is a key design criterion for the IDVerifact Digital Identity Platform business continuity planning. Google’s physical security features a layered security model, including safeguards like custom-designed electronic access cards, alarms, vehicle access barriers, perimeter fencing, metal detectors and biometrics. The data centers floors feature laser beam intrusion detection. Google’s data centers are monitored 24/7 by high-resolution interior and exterior cameras that can detect and track intruders. Access logs, activity records and camera footage are available in case an incident occurs.

Google’s data centers are also routinely patrolled by experienced security guards who have undergone rigorous background checks and training. On the data centers’ floors, security measures increase. Access to Google’s data centers’ floors is only possible via a security corridor that implements multi-factor access control using security badges and biometrics. Only approved employees with specific roles may enter.

Recovery Point Objective (RPO)

Based on the current design, IDVerifact is deployed on multiple availability zones for both the application component and database component. The application is changed only when new deployments are executed, and they are deployed to both the primary and secondary location. The database is replicated from primary zone to secondary zone in real-time (secondary zone being the slave) resulting to minimal loss of data in case the primary zone experiences an outage because of a disaster.
On top of that database are backed up periodically as follows:

- Every 4 hours during the business operational time (I.e. 8am ET, 12pm ET, 4pm ET and 8, pm ET)
- Once over night at 2am ET

Recovery Time Objective (RTO)

We anticipate that, should a catastrophe strike, the IDVerifact Digital Identity Platform can be recovered within 1-3 hours with minimal to no loss of transactional logs. This estimate is based upon our current design. During this period the IDVerifact team will execute the following tasks to bring back the platform:

- Based on the design the Load Balancer will automatically switch to the secondary zone and make the platform available.
- The secondary platform will be configured to connect with the Google SQL cloud in the secondary zone (where the data is getting replicated)
- The current state of the database will be checked to ensure that the last replication was executed a few minutes before the primary zone was unavailable.
- The Google SQL Cloud on the secondary zone will be made master going forward.
- The team will perform a quick test to ensure that traffic to the application component in the secondary zone is connected with the Google SQL and once all the tests are executed the system will be available to the users for continuing their business operations.