

Video Example

[EdgeLake Overview Video](#)

[EdgeLake in 1 minute](#) (no sound)

[EdgeLake with Open Horizon](#)

Text for EdgeLake with OH Demo

```
<speack>
  <p>
    Hello, I'm Joanna, and I will be guiding you through a deployment of EdgeLake using Open Horizon.
    EdgeLake is a platform that transforms the edge to a virtual data lake.
    Open Horizon is a platform for managing the service software lifecycle of containerized workloads.
  </p>
  <p>
    On the left side you can see data managed today by transferring the data to the cloud.
    On the right side we show an EdgeLake Network. The green nodes are EdgeLake instances at the edge, and
these are the nodes that host the data.
    When an application needs data, it is services from the edge, from a virtual database, replacing the
need to centralize the data.
  </p>
  <p>
    Here is the process with more details:

    At the bottom, sensors and applications generate data.
    The second layer shows the edge nodes where EdgeLake is deployed using Open Horizon.
    Open Horizon deploys the EdgeLake instances and the data is streamed from the sensors and apps to the
EdgeLake nodes.
    This is all that is needed, from here the data is available to the applications:

    An app issues a query to one of the EdgeLake instances and this instance identifies where the needed
data resides and process the query like MapReduce

  </p>
  <p>
    The outcome is a data lake where sensors and apps generate the data, the data remains in place, but
serviced to the apps as if it is centralized.
  </p>
  <p>
    Here is a live demo. The green dots show the EdgeLake instances of and EdgeLake network.
  </p>
  <p>
    We have a few examples of Grafana that visualize the data. Grafana is issuing the query to a single
node in the network without knowing which are the nodes that host the data.
    The query is satisfied as if the data is centralized.
  </p>
  <p>
    This Grafana example, shows KubeArmor data, this is event data that is monitored to alert on security
issues with Kubernetes Pods. EdgeLake satisfies the queries without centralizing the data.
  </p>
  <p>
    This is a WEB GUI that we use to query the network.

    In this first example I issue a SQL query to the network. We get the result set back as if we are
working with a centralized database.
    But if I'll scroll to the bottom, you can see from the summary section that this query was
satisfied from 6 EdgeLake nodes.
    In this second example, show that with EdgeLake, users can monitor all their distributed resources from
a single point, as if all the edge resources are a single machine.

  </p>
  <p>
    In this last part of the demo, we will deploy a new EdgeLake instance using Open Horizon.

    Here is the Open Horizon Policy and I will start the deployment of the EdgeLake instance.
    and here is a Grafana instance visualizing the data.
    This Grafana instance represents data from 2 edge nodes, the deployment not only deploys EdgeLake
instance but also integrates EdgeLake with a data source.
    You will see that after the deployment, a 3rd node will contribute data to the table.

    Note, this entire EdgeLake and Open Horizon demo transformed the edge to a virtual data lake.
    As you have seen the entire process required only deployment of EdgeLake instances but without a
single line of code.
  </p>
</speack>
```

Add your comments.