Clustered compute and storage support in EVE

This proposal is to add clustered compute and storage functionality to EVE.

Motivation:

EVE software today runs on single device and that device is orchestrated using the REST API which EVE exposes to the controller. Since it is a single device the HA and scalability of the EVE-OS is an issue if the customer applications demands such requirements.

The motivation of this proposal is to eliminate such limitation and make EVE-OS a distributed storage and compute operating system at the edge.

The following picture is high level overview of EVE-OS today.



Proposed EVE architecture:

There are many open source distributed clustering solutions and one of those is Kubernetes based system. The proposal here is to bring Kubernetes based eco system into EVE and convert EVE storage into a truly distributed storage system. Please note I am not proposing EVE to expose Kubernetes API, EVE will still expose its own REST API, the Kubernetes components compile into or ship with EVE and are accessed internally by EVE micro services and a new translation layer to provide the clustered storage.

The following are the new proposed features as part of this architecture:

- •True scale out EVE-OS based on Kubernetes components (k3s)
- Native container support
- Virtualization support using Kubevirt
- •VM and container failover support (HA)
- •Storage replication using CSI (ex. longhorn)
- •Kubernetes based CNI
- •Orchestrated by controller using REST API
- •API compatibility to support previous EVE versions
- •New microservice and a Translation layer to connect EVE to Kubernetes components

The following picture is high level overview of proposed EVE architecture.

A new micro service will talk to a new translation layer which interacts with Kubernetes eco system. External communication from EVE is still through the EVE REST API.



A sample 3 node EVE-OS cluster



Building new EVE:

Since EVE already has mechanism to build for various hypervisor versions using Makefile flag HV="kvm|xen|acrn" the proposal is to compile the new EVE with HV=kubevirt. The base EVE micro services will be common for all type of hypervisors. Kubevirt specific EVE will include Kubernetes components like k3s, CSI (longhorn) and kubevirt (ofcourse).

This mechanism will ensure that EVE can still be upgraded and supported on kvm based systems.

Conclusion:

With this approach we will be bringing Kubernetes based components to the edge and make EVE-OS a scalable and distributed system.