HA-DO - SIMPLE Clustering Engine

Superstes

Apr 08, 2022

CONTENTS

The goal of this project is it to provide a **simple base-engine for building high-available systems**.

Simplicity is a feature.

CHAPTER

ONE

BASE

1.1 Base

1.1.1 Intro

HA-DO Clustering Engine

The goal of this project is it to provide a **simple base-engine for building high-available systems**.

Simplicity is a feature.

Goals

KEEP IT SIMPLE

Admins should be able to understand and troubleshoot the system after a few minutes.

Functionality can be added using plugins. Allowing for advanced usage without increasing the core engines complexity.

Platform

Initially this solution will only target **systems that use systemd** as init.

Functionality

CORE

MONITORING

Checking the local-systems and peer states

RESOURCES

Interact with a resource needed by your HA-APP (*start, stop, promote, demote, alive, other*)

STATE-API

Lightweight http server to publish it's state

Used to exchange state-info between nodes.

Also - allows easy:

- monitoring
- debugging (*by simply reading the json output in your browser*)

Modular extensibility

The core uses PLUGINS to achieve its functionality.

Keeping the footprint of the base-engine small - you can add plugins as needed.

Example plugins:

- Resource for systemd services
- Resource for ip-addresses
- Monitoring of a running process
- Monitoring if some remote port is reachable

  The structure of such a plugin should be logic.

Transparency

The functionality of the core should be logic and easily understandable and to troubleshoot.