Solve the operational problem of deploying and managing distributed networks

StarlingX is a complete cloud infrastructure software stack for the edge used by the most demanding applications in industrial IOT, telecom, video delivery and other ultra-low latency use cases. With deterministic low latency required by edge applications, and tools that make distributed edge manageable, StarlingX provides a container-based infrastructure for edge implementations in scalable solutions that is ready for production now.

The StarlingX virtualization platform focuses on easy deployment, low touch manageability, rapid response to events and fast recovery -- think control at the edge, control between IoT and cloud, and control over your virtual machines. Rather than reference platforms and gap definition for edge use cases, StarlingX provides a deployment-ready, scalable and highly reliable edge infrastructure software platform to build mission critical edge clouds. Tested and released as a complete stack, StarlingX ensures compatibility among diverse open source components, such as OpenStack, Kubernetes, Ceph, Linux and more. Its unique project components provide fault management and service management among others to ensure high availability of user applications. The StarlingX community has optimized the solution for security, ultra-low latency, extremely high service uptime, and streamlined operation.

STARLINGX

Join the Community

StarlingX is an independent open source community collaboratively developing code under the Apache 2 license. Anyone is welcome to join and contribute code, documentation, and use cases. The project is supported by the OpenStack Foundation.

Get Involved

Website: starlingx.io
Git: git.starlingx.io
Docs: docs.starlingx.io
Freenode IRC: #starlingx
Mailing Lists: lists.starlingx.io
E-mail: info@starlingx.io

Ready to explore StarlingX? Try it at starlingx.io

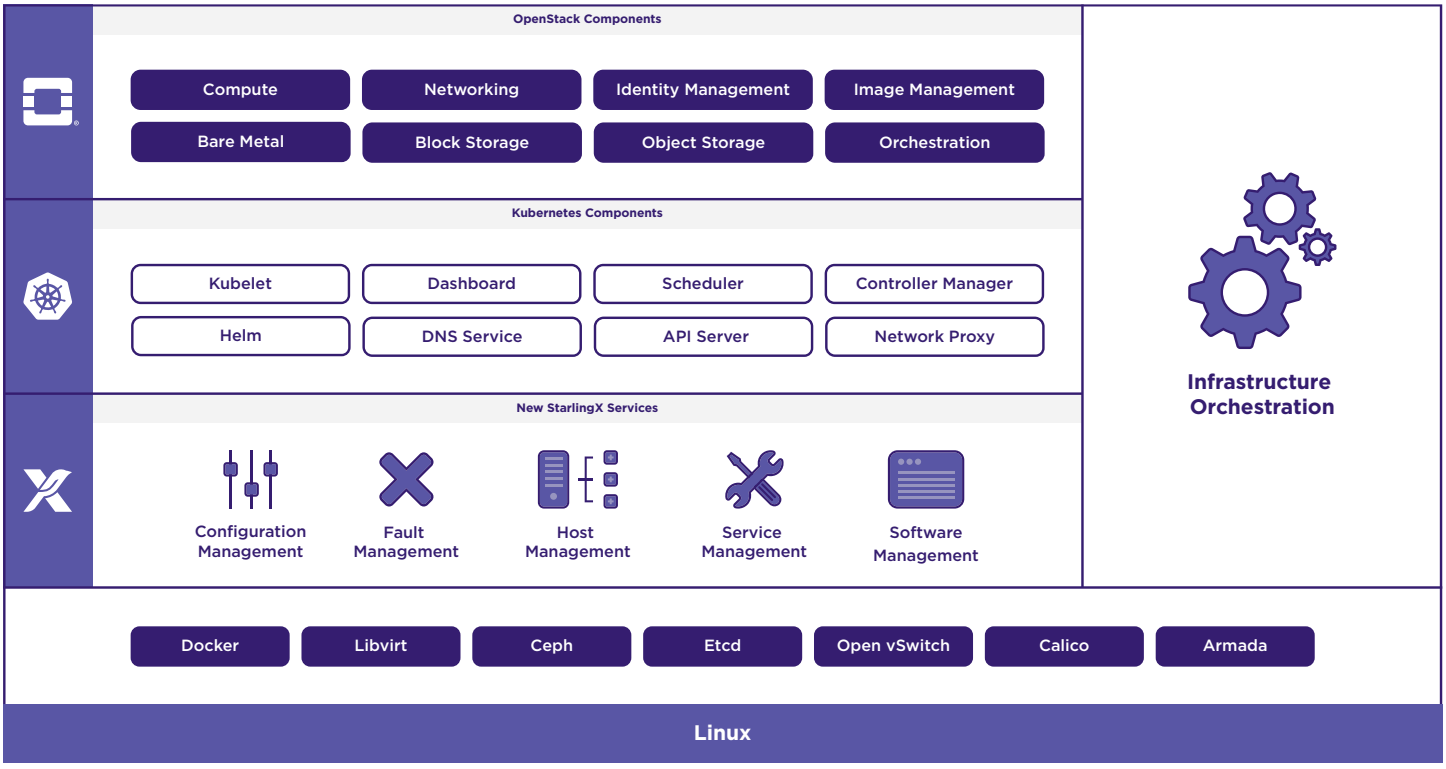

There are more OpenStack and Kubernetes components used than represented in this diagram.

Features

Reliability

Fault management, fast secure VM failover and live migration minimizes downtime

Ultra-low latency

Deterministic, tunable performance optimized for the use case

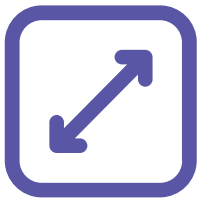

Scalability

Deployable on one to thousands of distributed nodes allowing for a single system to be used from edge to core

Edge security

Software security to avoid tampering at the edge, where physical security may be limited

Small footprint

Providing a platform for edge and IoT use cases even for environments with tight resource constraints

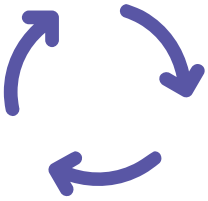

Lifecycle management

Simplified deployment and operations with full system management through comprehensive orchestration suited for the edge

Use Cases

Ultra Low-latency 5G and Industrial IoT (IIoT)

- Autonomous vehicles (drones, cars and trucks)
- Industrial automation, such as robotics and virtual Programmable Logic Controller (vPLC)
- Cloud/virtual Radio Access Network (cRAN/vRAN)
- Smart city/buildings (metering and monitoring)

High Bandwidth, Large Volume Applications

- Mobile HD video
- Content delivery
- Healthcare (imaging and diagnostics)
- Caching and surveillance

Multi-access Edge Computing (MEC)

- Augmented and virtual reality (AR/VR)
- Enterprise focused small cell services for stadiums and high-density locations
- Universal Customer Premise Equipment (uCPE) applications
- Retail