



# Blueprint Submission for Time-critical Edge Compute

Intel Corporation, Inc.

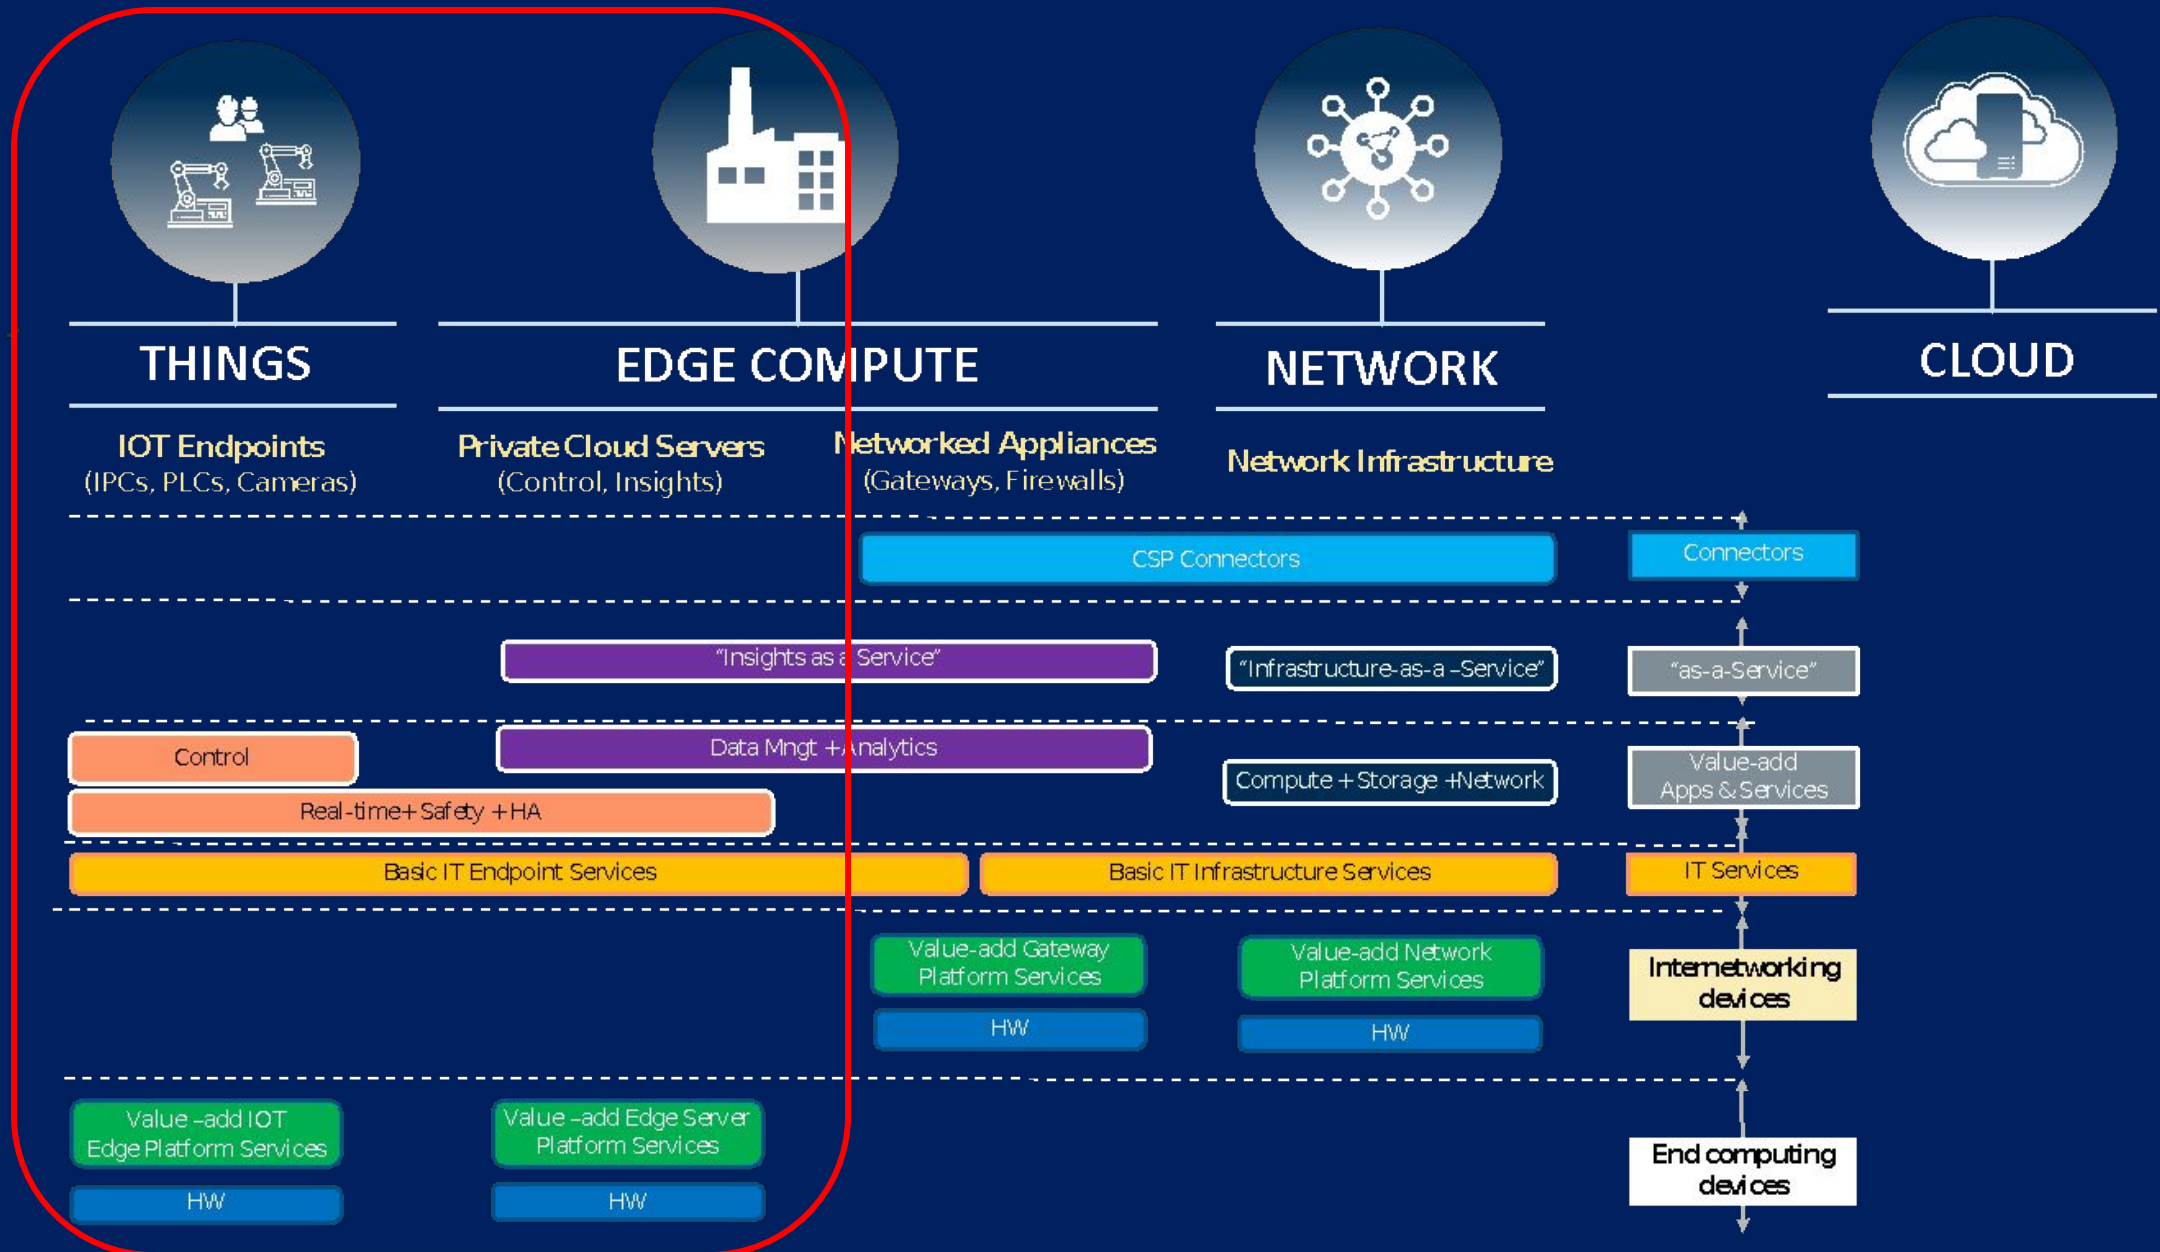
# Time-critical Edge Compute Blueprint: Use Cases

- Use cases in Manufacturing, Smart Buildings, general IIOT
  - Virtualized PLC
  - Computer vision inference
  - Machine, sensor data inference
  - Process or discrete manufacturing closed loop control
  - Ethernet TSN
- Functional Safety capable use cases
  - Discrete manufacturing soft PLC
- Onramp for 5G-URLLC

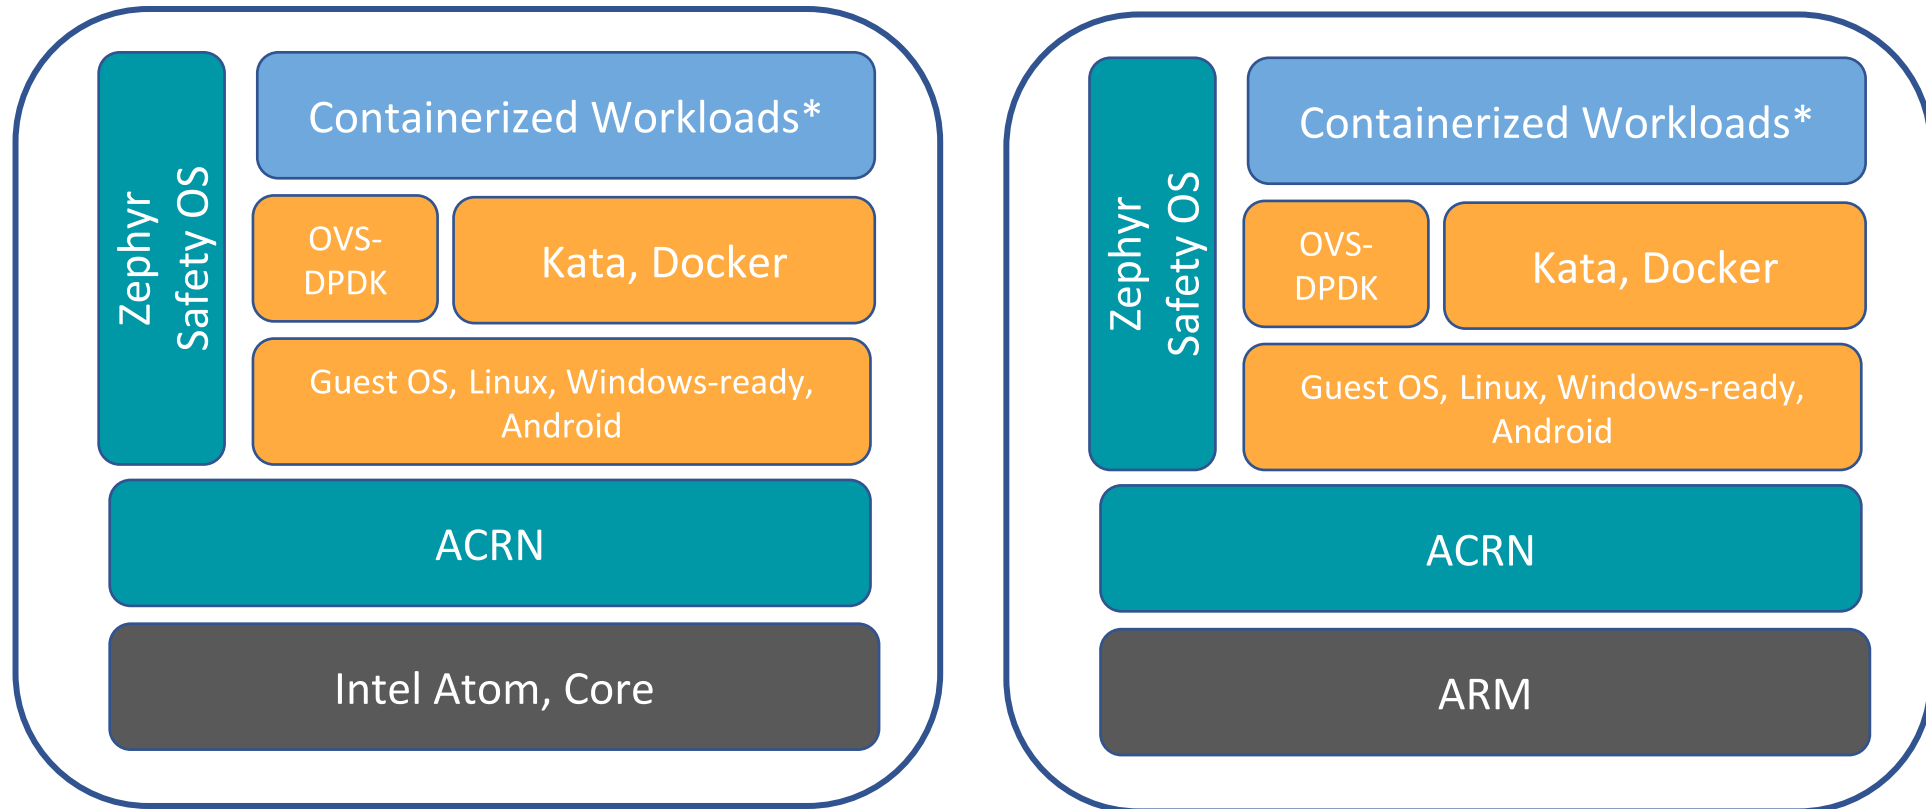
# Time-critical Edge Compute Blueprint: Hardware and Partners

- Low power, ruggedized hardware
  - Dell 3000, 5000 IPC
  - Huawei XXX industrial gateway
- Potential to attract new members to Akraino project
  - Industrial ODM's e.g. Advantech, Adlink
  - Industrial OEM's/ISV's e.g. TTTech, Nebbiolo, IOTech
  - Industrial end-users e.g. ExxonMobil

# Time-critical Edge Compute Blueprint: Deployment Focus




# Time-critical Edge Compute Blueprint: Deployment Scenarios



Virtualized, Functionally Safe workloads in addition to others  
Easily Extensible and Expandable, by just adding more systems  
Evaluating Airship for ZTP and deployment

*\* See next page for some sample targets*



# Containerized edge workloads

- Containerized workloads orchestrated via Kubernetes tuned for lightweight, time-critical embedded deployments
- Sample workloads include
  - Tensorflow via Kubeflow
  - OpenVINO for Video and Inference
  - Closed loop control (e.g. IEC 61131)
  - EdgeX Foundry
    - Building controller

# Demo

This stack is largely functional today.

Work ahead is in hardware software validation and validation of the workloads described.

Demo Link: <https://youtu.be/1qkRJIuUSY>

# Backup: Deep dives for underlying technology

- Zephyr OS
- ACRN Hypervisor
- Kata Containers
- Celadon - A fully Open Source Android Stack
- OVS-DPDK







A scalable real-time operating system (RTOS) supporting multiple hardware architectures, optimized for resource constrained devices, and built with security in mind.

<https://www.zephyrproject.org/>

# Overview – A Fully Featured Open Source RTOS (since 2016)

## Safety

- Thread Isolation
- Stack Protection (HW/SW)
- Quality Managed (QM)
- Build time configuration
- No dynamic memory allocation
- FuSA (2019)

## Security

- User-space support
- Crypto Support
- Software Updates

## Configurable & Modular

- Zephyr Kernel can be configured to run in as little as 8k RAM
- Enables application code to scale
- Configurable and Modular

## Cross Platform

- Support for multiple architectures
- Native Port
- Developed on Linux, Windows and MacOS

## Open Source

- Licensed under Apache II License
- Managed by the Linux Foundation\*
- Transparent development
- Fork it on Github!

## Connected

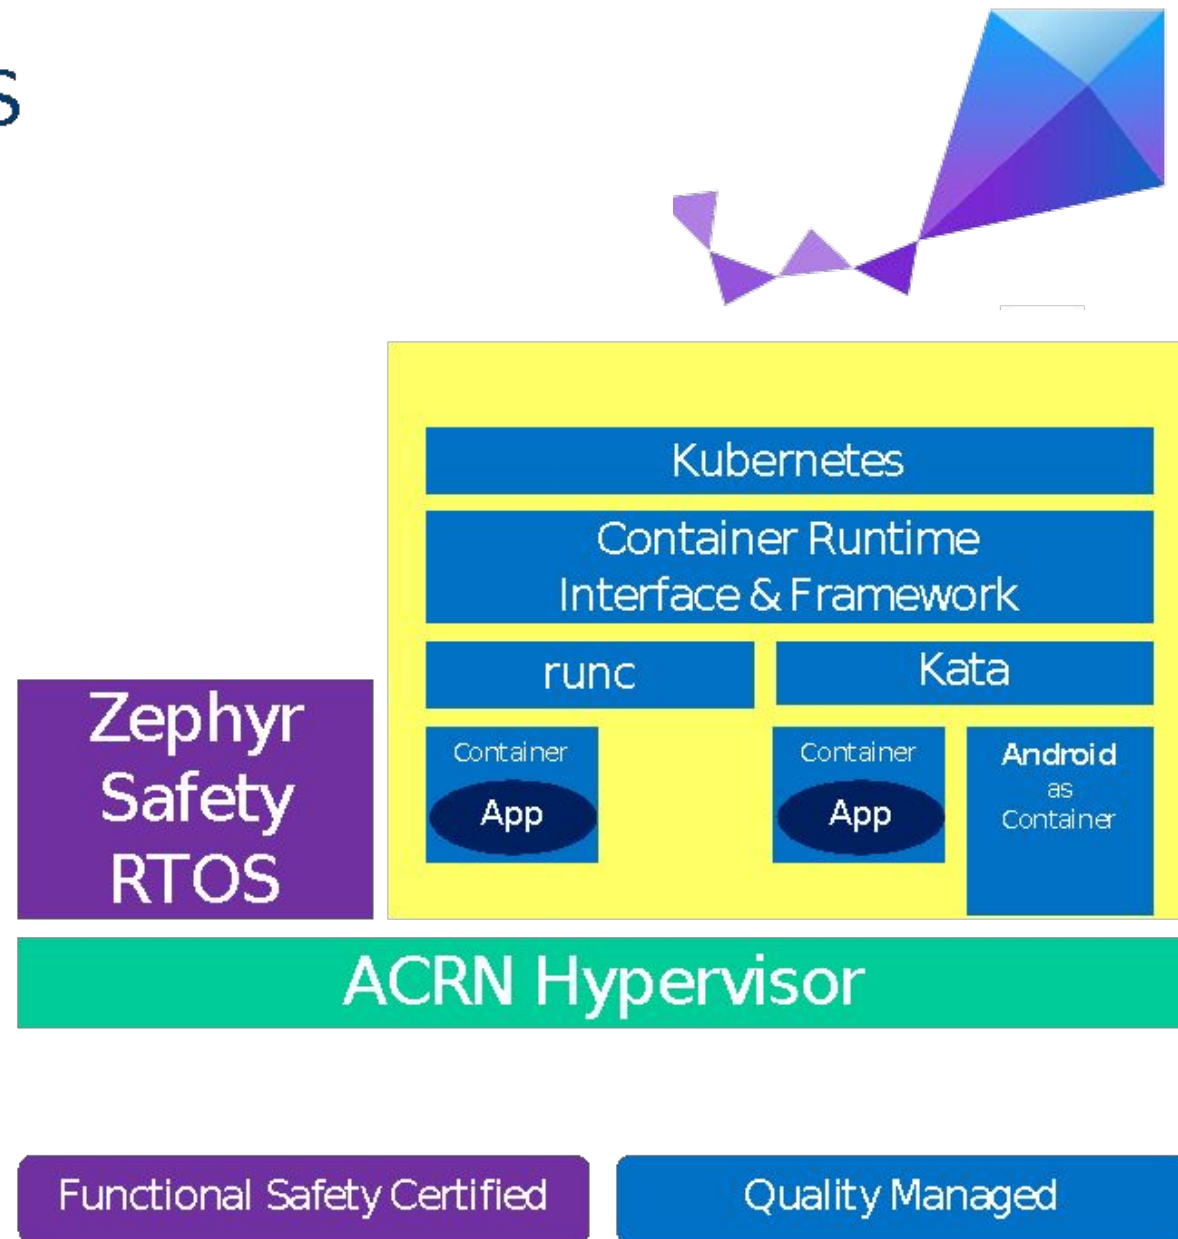
- Full Bluetooth 5.0 Support
- Bluetooth Controller
- BLE Mesh
- Thread Support
- Full featured native networking stack
- DFU (IP+BLE)

Zephyr™ is not an ingredient, Zephyr™ provides a complete solution.

# Zephyr Enabled as a Safety Critical OS

- Runs on a custom hypervisor that is safety critical capable
- Security updates with the latest fixes
- Similar to Cloud Software Defined Infrastructure (SDI)

- Zephyr =FuSa (2019)
- Linux =Quality Managed
- Android =non-FuSa



# Zephyr™ OS Direction



## Safety & Security

- Functional Safety (FuSa) core OS certification: secure & harden kernel (IEC61508 SIL3).
- Development model & process with safety and security in mind.
- Trusted Execution Environments.

## E2E Platform

- Bootloader.
- Device firmware updates.
- Cloud connectivity.
- Development tools.

## Expanded Use Cases

- Industrial, safety, and security features.
- Deep embedded usages (BLE, 802.15.4 (zigbee), BT Mesh).
- Advanced configurations and use cases: Multicore, SMP, AMP.

## Ecosystem & Portability

- Improve support on Mac and Windows.
- IDE integration.
- 3rd party tools: tracing, profiling, debugging.
- LLVM, commercial compilers.
- Standard APIs and portability: POSIX layer (PSE54), BSD socket, and CMSIS RTOS.



ACRN

---

**A Big Little Hypervisor for IoT Development**

---

# What is ACRN™?



ACRN is a flexible, lightweight **reference hypervisor**, built with real-time and safety-criticality in mind, **optimized** to streamline **embedded development** through an open source platform.

## A Big Little Hypervisor for IoT Development



# ACRN™ Features



Small Footprint



Built for IoT




Adaptability



Built for Real-Time



Safety Criticality



Truly Open Source



# Features Roadmap - Proposal


*	Limited to specific HW
PT	Pass through



Dates below are for reference only and subject to change					
Area	v0.1@Q2'18	v0.2@Q3'18	V0.5@Q4'18	V1.0@Q1'19	V1.x@2019
<b>HW</b>	<ul style="list-style-type: none"> <li>• APL NUC (UEFI)</li> <li>• APL UP2 (UEFI)</li> </ul>	<ul style="list-style-type: none"> <li>• APL NUC (UEFI)</li> <li>• APL UP2 (UEFI)</li> </ul>	<ul style="list-style-type: none"> <li>• APL NUC (UEFI)</li> <li>• KBL NUC (UEFI)</li> <li>• APL UP2 (UEFI)</li> </ul>	<ul style="list-style-type: none"> <li>• APL NUC (UEFI)</li> <li>• KBL NUC (UEFI)</li> <li>• APL UP2 (UEFI)</li> </ul>	<ul style="list-style-type: none"> <li>• APL NUC (UEFI)</li> <li>• KBL NUC (UEFI)</li> <li>• APL UP2 (UEFI)</li> <li>• ARM</li> </ul>
<b>Hypervisor</b>	<ul style="list-style-type: none"> <li>• VT-x</li> <li>• VT-d</li> <li>• CPU static-partitioning</li> <li>• memory partitioning</li> <li>• Virtio (v0.95)</li> <li>• VHM</li> <li>• EFI boot</li> <li>• Clear Linux as guest</li> </ul>	<ul style="list-style-type: none"> <li>• Virtio (v1.0)</li> <li>• Power Management (Px/Cx)</li> <li>• VM management</li> <li>• ACRN debugging tool</li> <li>• vSBL</li> </ul>	<ul style="list-style-type: none"> <li>• Android as guest</li> <li>• AliOS as guest</li> <li>• Zephyr as guest</li> <li>• MISRA C compliance</li> <li>• <b>Logical partitioning without SOS</b></li> <li>• Trusty (Security)</li> <li>• SBL boot</li> </ul>	<ul style="list-style-type: none"> <li>• vHost</li> <li>• Power Management (S3/S5)</li> <li>• <b>Hybrid Mode (Privilege VM loaded by SOS)</b></li> <li>• Real Time phase I</li> </ul>	<ul style="list-style-type: none"> <li>• Real Time phase II</li> <li>• <b>Hybrid Mode (Privilege VM loaded by hypervisor)</b></li> <li>• Windows as guest</li> <li>• VxWorks as guest</li> <li>• Functional Safety capable</li> <li>• CPU sharing</li> <li>• OVMF</li> <li>• ARM</li> </ul>
<b>I/O virtualization</b>	<ul style="list-style-type: none"> <li>• Storage</li> <li>• Ethernet</li> <li>• USB host controller (PT)</li> <li>• USB device controller (PT)</li> <li>• Audio (PT)</li> <li>• WiFi (PT)*</li> <li>• Touch (PT)</li> </ul>	<ul style="list-style-type: none"> <li>• GPU Sharing:</li> <li>• GPU Surface Sharing</li> <li>• IPU Sharing*</li> </ul>	<ul style="list-style-type: none"> <li>• GPU Prioritized Rendering</li> <li>• Touch sharing</li> <li>• IOC sharing*</li> <li>• Audio sharing</li> <li>• USB host controller Sharing</li> <li>• USB DRD virtualization</li> </ul>	<ul style="list-style-type: none"> <li>• GPIO virtualization</li> </ul>	<ul style="list-style-type: none"> <li>• HECI sharing (Security)</li> <li>• CSME/DAL sharing (Security)</li> <li>• TPM Sharing (Security)</li> <li>• eAVB/TSN Sharing</li> <li>• SR-IOV*</li> </ul>



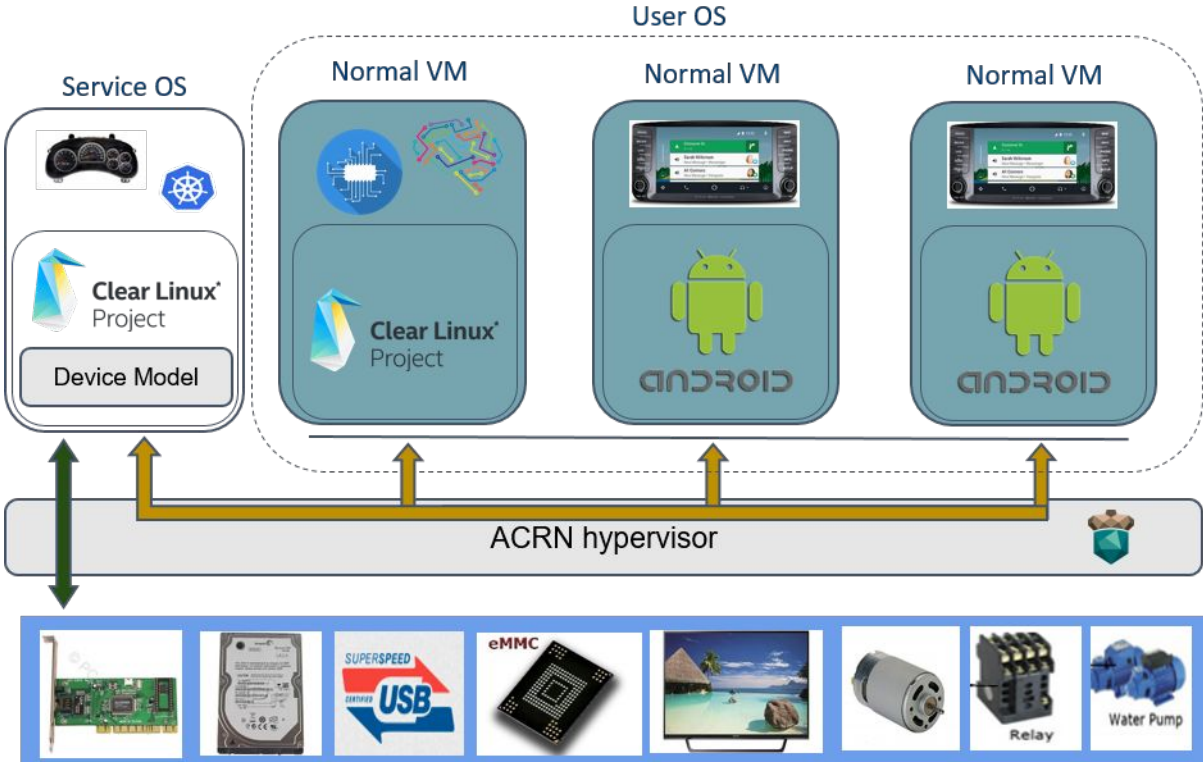
# Towards MISRA-C Compliance



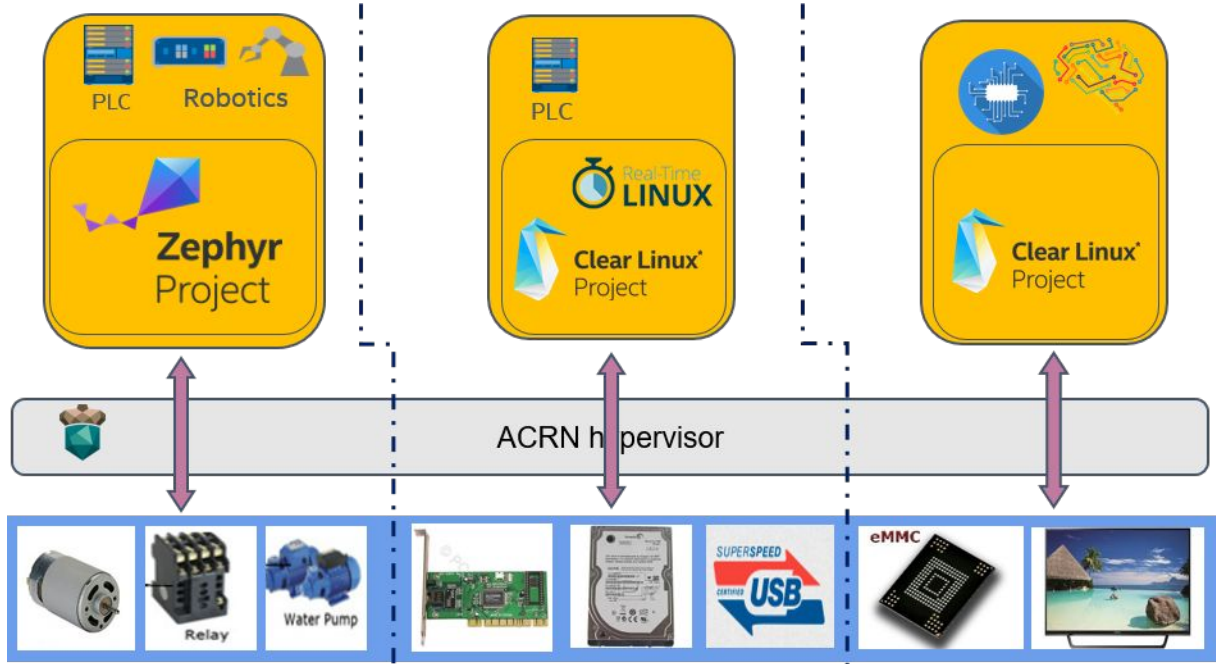
- Statistics from commercial safety-qualified checker.
- False positives and intended deviations tracked in weekly-updated sheets.
- Pull requests are scanned hunting for new violations.



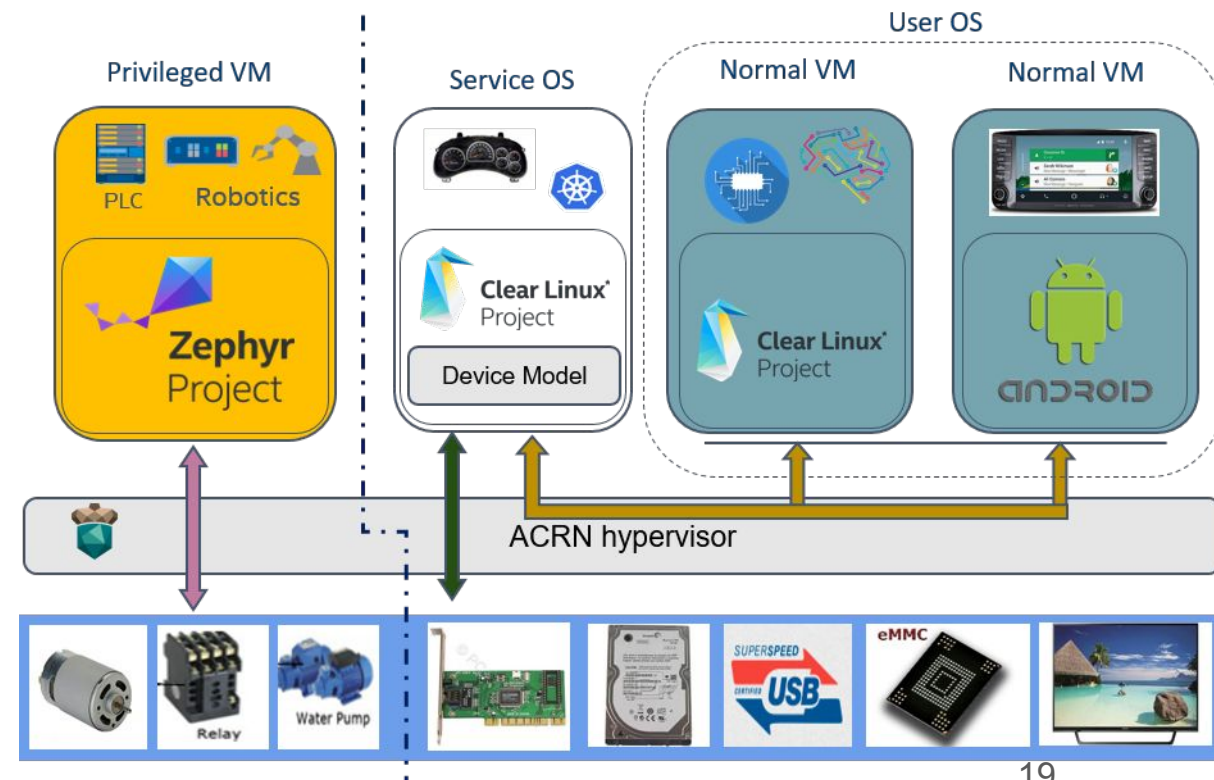
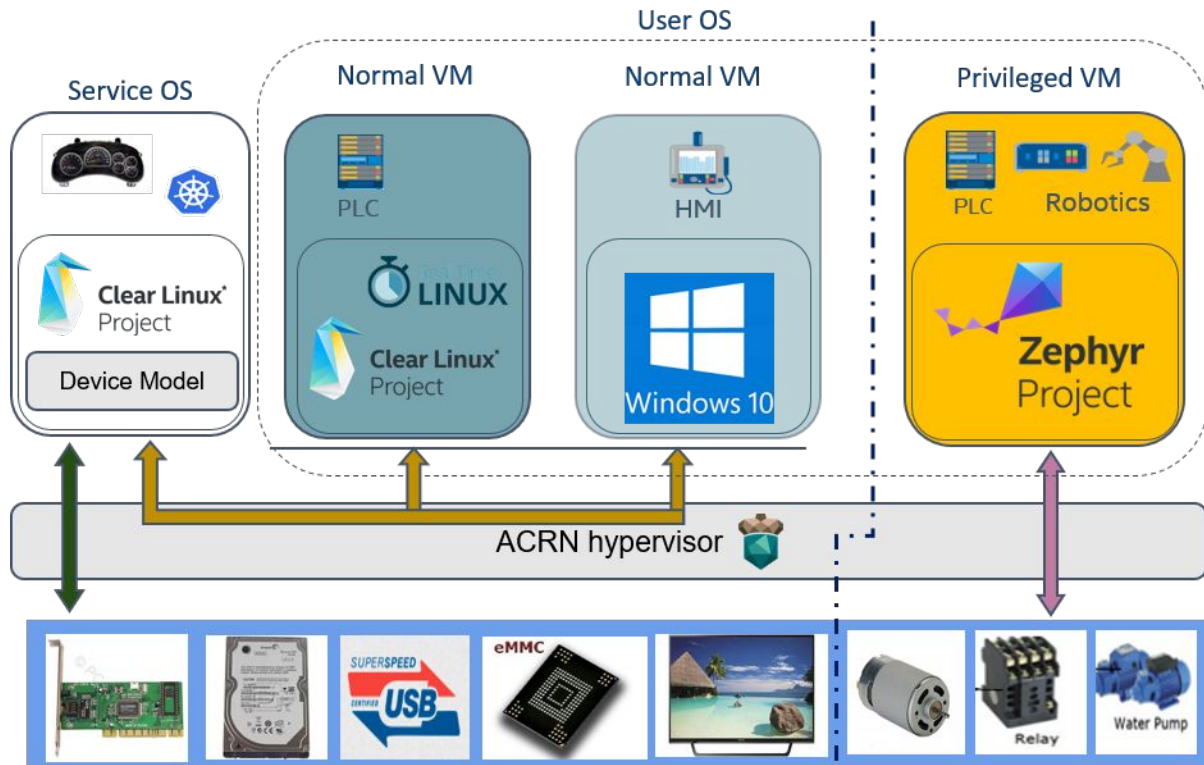
# Sharing Mode



# Partition Mode



# Hybrid Mode



# Kata Container Project

<https://katacontainers.io/>

## Project Overview, Status



# What is Kata?

- kata-runtime, an OCI (Open Containers Initiative) compliant runtime
  - *Seamless* integration into cloud native ecosystem
- “Providing the speed of containers with the security of virtual machines”
  - Light-weight enough to be used with micro-services design patterns
  - More than just security of virtual machines, it is an additional layer on top of existing container security primitives.
  - Each container/pod is created within its own virtual machine



# Who is Kata?

- Open source, open governance project with original contributions from Intel's Clear Containers and Hyper.sh's runV
- Under the Openstack Foundation Umbrella (not managed by openstack)
- Architecture Committee: Google, Huawei, Hyper.sh, Intel
- Contributors include: AMD, ARM, Branch, IBM, Intel, Google, Huawei, Hyper.sh, Microsoft, Nvidia, Openstack Fountain, Redhat, Suse, ZTE, 99Cloud ...



# Where does Kata make sense?

- Regulated and sensitive production environments
- Too many capabilities required which increase attack surface
- Desire to easily run on multiple or custom kernel versions
  - Legacy applications on older kernels in containerized environment
  - Custom kernel features required
  - Testing on cutting edge kernels





# Where else does Kata make sense?

- Bare-metal infrastructure
- Mixed levels of trust
  - Multiple tenants
  - Untrusted workloads





# Kata Updates since release

V1.0 (May 2018)	V1.2 (August 2018)
<ul style="list-style-type: none"><li>• Seamless integration with Kubernetes (CRI), Docker</li><li>• Hardware isolation using KVM/QEMU</li><li>• Optimizations for minimal footprint and boot-time</li><li>• Seamless integration with major networking plugins<ul style="list-style-type: none"><li>◦ Advanced networking available through DPDK (VPP/OVS and SR-IOV)</li><li>- High bandwidth, low latency networking<ul style="list-style-type: none"><li>▪ Ability to run custom kernels at the container or pod level</li></ul></li></ul></li><li>• Direct device assignment (GPU, RDMA, QAT, etc.)</li></ul>	<ul style="list-style-type: none"><li>• Support multiple architectures</li><li>• VM-Factory support [1]</li><li>• Vsock support [2]</li><li>• K8S deployment through container based daemonset [3]</li><li>• Bug fixes, enhancements</li></ul> <p>[1] - <a href="https://github.com/kata-containers/runtime/pull/303">https://github.com/kata-containers/runtime/pull/303</a> [2] - <a href="https://github.com/kata-containers/runtime/issues/383">https://github.com/kata-containers/runtime/issues/383</a> [3] - <a href="https://github.com/kata-containers/packaging/pull/65">https://github.com/kata-containers/packaging/pull/65</a></p>



# Kata Roadmap

V1.3 (September 2018)	Looking forward
<ul style="list-style-type: none"><li>• Full network hotplug</li><li>• Full storage hotplug</li><li>• Open-tracing support (Jaeger)</li><li>• CNI-Macvlan support</li><li>• Containerd v2 shim</li></ul>	<ul style="list-style-type: none"><li>• Runtimeclass</li><li>• More native integration with CRI (containerdv2 for CRIO)</li><li>• Security Enhancements</li><li>• Live upgrade</li><li>• Performance optimizations</li></ul> <p>See <a href="https://github.com/orgs/kata-containers/projects/12">https://github.com/orgs/kata-containers/projects/12</a></p>







# Project Celadon

<https://01.org/projectceladon/>

# Project Celadon: Elements & Benefits



## Code transparency

open source code **provides freedom and flexibility to customize and accelerate development**

## Turnkey system

supports a wide range of hardware components optimized for Intel architecture making it **easy for rapid prototyping and building new applications**

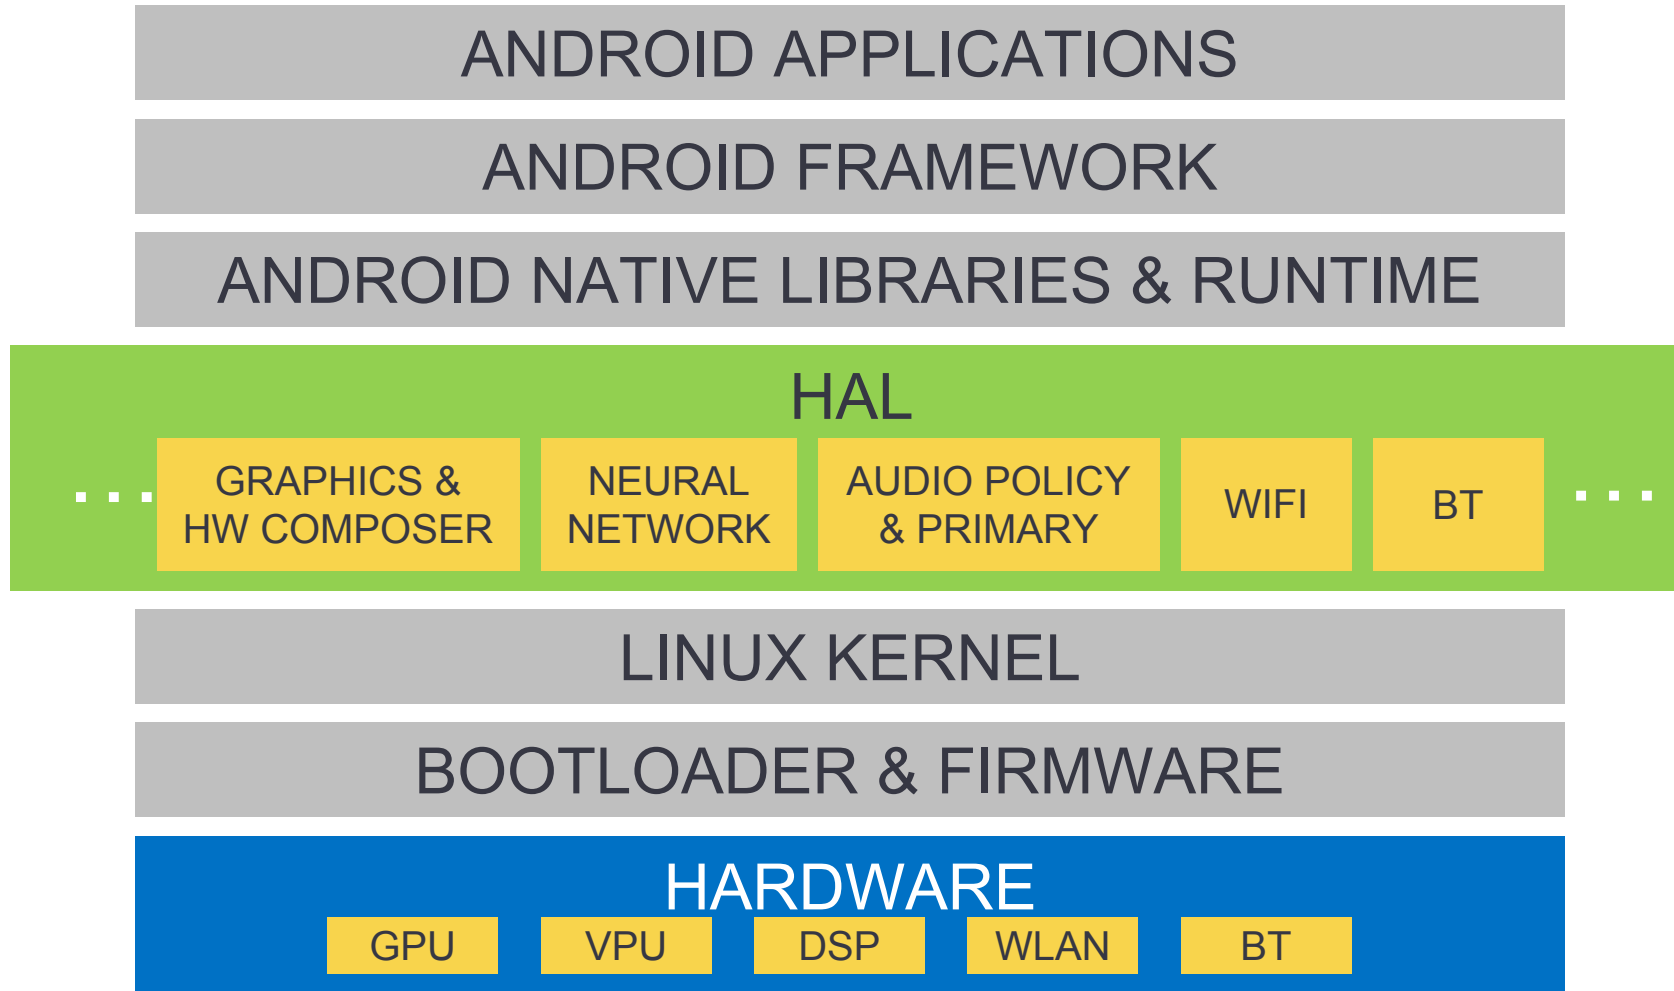
## Regularly updated

**opportunity to realize new features and improvements** by developing on the latest hardware implementations and Android software updates

## Verified compatibility

basic Android compatibility **ensures consistent application and hardware environment and experience**

# Architecture



Built on standard and familiar android stack architecture



<https://www.dpdk.org/>