ETSI The Standards People

Security Conference

On using Containers, Virtual Machines, LightVMs and Unikernels in a Secure Operational Environment for Critical Infrastructure

Presented by: Dr. Djibrilla AMADOU KOUNTCHE

AKKODIS

19/10/2023

Context of our work: Critical Infrastructure Protection

PRECINCT

This work was carried out in PRECINCT (Preparedness and Resilience Enforcement for Critical INfrastructure Cascading Cyber physical Threats) which is part of the ECSCI Cluster aiming to :

- Provide a platform for combined safety and security for European Critical Infrastructures
- Provide European Common Platform for cascading effects on the different Critical Infrastructures
- Contribute to standards and regulations on the protection of Critical Infrastructure

Source: https://www.finsec-project.eu/ecsci

This project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 101021668.

Context of our work: PRECINCT Vision

PRECINCT

- PRECINCT aims to **connect private and public CI stakeholders** in a geographical area to **a common cyber-physical security management approach** which will yield **a protected territory** for **citizens and infrastructures**
- **Enable interdependent CIs and First Responders / Public authorities** to plan for, prevent, absorb, recover and adapt efficiently and effectively to the effects of cyber-physical and hybrid threats / attacks as well as **impede their cascading effects**.
- **PRECINCT CIs Coordination Centres** (3Cs): will provide collaboration and governance models that link CIs, first responders and other CI stakeholders **harmonising CIs emergency processes with command structures and data sharing, thus enabling the quantification and management of resilience** via identification and implementation of measures that **minimise the impact of cascading effects arising from the interdependencies between different types of critical infrastructures**
- **PRECINCT Digital Twins** will enable trusted, efficient, accurate and cost-effective operations for 3Cs by discerning and tracking events within and across system boundaries, underpinned by machine learning principles that, over time, provide self-adapting cognition based on learned behaviours, learned corrections, learned patterns and learned interventions thus incentivising automated upgrading of interdepent CIs resilience

This project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 101021668.

Context of our work: PRECINCT Outputs

PRECINCT

Digital Twin modelling the current and future behaviour of territory-based interdependent CIs in a variety of conditions and configurations, to anticipate threats, to detect anomalies, and to incentivise optimised command structure and coordinated responses

Model the dynamic interdependencies and cascading effects among (and within) the risk components in critical infrastructures and **quantify resilience (Resilience Index)** and **identify short-term and long-term resilience enhancement measures**

Directory of CIP Blueprints

Serious Games based innovative tool that that can be used to investigate vulnerabilities in complex multi-system cascading effects and support focused development of new resilience enhancement services **New CIP Blueprints**

Consolidation and Harmonisation of key assets from CIP and INFRA-01 projects

This project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 101021668.

Main objectives of our work: Re-usability and transferability

PRECINCT

- Systematically re-used results of past CIP projects:
  - Facilitating better understanding of these results
  - Ease the deployment, upgrade of the IT tool
  - Sharing knowledge and training new colleagues

- Need of a consistent approach to capitalize on previous projects for:
  - Re-using these results
  - Standardize IT assets for CIP
  - Training of new team members on existing solutions
- Question still unanswered:
  - How to share, re-use this knowledge in a readily manner ?

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under Grant Agreement No. 101021668.

Towards re-usable outcomes of CIP assets

PRECINCT

- An approach based on **Blueprints** :
  - Reference architectures (example: *Big Data Reference Architecture*)
  - Human and machine-readable description languages
  - Standards Compliance

- Ease the work of *deployment*, *orchestration* and *maintenance* **for IT Teams**

- Facilitates Re-usability, Transferability and Experimentation

This project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 101021668.

The case of the PRECINCT Ecosystem Platform

PRECINCT

WP2 - PRECINCT Ecosystem Platform and Blueprints Directory [Months: 1-20]
AKKA, ICP, UCD, AIT, BSC, MON, NURO, TNCL, ENG, KNT, VLTN , ATTD, AIA, AMETRO, LIST, LEPIDA
Task 2.1 Semantic CIs Connectivity and Dynamic Integration Tools (Lead: KNT; Partners: ICP, AIT, AKKA, MON, TCNL, ENG, VLTN, ATTD, AIA, LEPIDA) (M1-M17)
Deliver the high-performance, highly scalable and distributed message-based semantic components for message exchanges between the federated staged systems on existing infrastructures. Introduce Knowledge Graphs for semantic data exchanges and deploy connectivity and communication-integration via Pub-Sub services components. Provide secure integration functions, ensuring privacy and blockchain based security, encryption and provenance. Integrate Complex Event Processing capabilities. Build and customize the Big Data Analytics Infrastructural services in the form of PRECINCT Blueprints for generic cloud infrastructures. Design and implement the scheduler to analyse, optimize and schedule the execution of distributed dataflows over distributed data sources.
ST2.1.1 Semantic Data Exchanges and Knowledge Graphs (ENG, KNT). Analyse, specify, integrate and configure the connectivity enabling software components for exchanging messages between CIs, validating the message structures, content, and sequences. Define and design the data exchanges protocols using a semantic approach, supported by Knowledge Graphs implemented mainly in Neo4J. Enable Organizational Interoperability supporting Semantics decomposition (OWL) integrating Graph database extensions and tools to efficiently work with RDF/semantic data, as well as publishing open-data for data sharing and exchanges between CI Systems.
ST2.1.2 Connectivity, Integration and Message Based Exchanges, Pub-Sub (KNT, ENG). Implement the Streaming and Queuing messaging exchanges functions of the layered and distributed highly scalable high availability architecture, including the components and preconfigured templates for easy distributed deployment and operation. Provide full end-to-end encryption, accommodate different data schemas and perform semantic data validation with Knowledge Graph extensions. Integrate available open source components (Apache Kafka, Apache Pulsar) for the message Broker, the Message Storage, and the semantic data registries for the CI systems and the Message Schemas. Automate deployment (Kubernetes) and cluster configurations for increased message loads and efficiency in processing. Embed monitoring and visualisation via a Management and configuration UI (e.g. Grafana). Integrate CHARIoT blockchain-based PKI for sensor and gateway authentication and blockchain aided encryption of IoT endpoints.

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under Grant Agreement No. 101021668.

PRECINCT Ecosystem Platform: Capabilities

PRECINCT

Objective: From the description of the Ecosystem Platform, identify the capabilities using a composition and decomposition mechanism. These capabilities are linked to the Reference architecture.

1 Data Acquisition & Ingestion	9 Synthetic Data Generation	17 Enterprise System Integration	23 Edge AI & Intelligence	29 Prediction		39 Basic Visualization	45 Dashboards
2 Data Streaming	10 Ontology Management	18 Eng. System Integration	24 Command & Control	30 Machine Learning ML		40 Advanced Visualization	46 Continuous Intelligence
3 Data Transformation	11 Digital Twin (DT) Model Repository	19 OT/IoT System Integration	25 Orchestration	31 Artificial Intelligence AI	35 Prescriptive Recommendations	41 Real-time Monitoring	47 Business Intelligence
4 Data Contextualization	12 DT Instance Repository	20 Digital Twin Integration	26 Alerts & Notifications	32 Federated Learning	36 Business Rules	42 Entity Relationship Visualization	48 BPM & Workflow
5 Batch Processing	13 Temporal Data Store	21 Collab Platform Integration	27 Reporting	33 Simulation	37 Distributed Ledger & Smart Contracts	43 Augmented Reality AR	49 Gaming Engine Visualization
6 Real-time Processing	14 Data Storage & Archive Services	22 API Services	28 Data Analysis & Analytics	34 Mathematical Analytics	38 Composition	44 Virtual Reality VR	50 3D Rendering
7 Data PubSub Push	15 Simulation Model Repository	52 Device Management	54 Event Logging	56 Data Encryption	58 Security	60 Safety	51 Gamification
8 Data Aggregation	16 AI Model Repository	53 System Monitoring	54 Data Governance	57 Device Security	59 Privacy	61 Reliability	62 Resilience

Data Services · Integration · Intelligence · UX · Management · Trustworthiness

Periodic table of capabilities (from Digital Twin Consortium)

This project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 101021668.

The PRECINCT Ecosystem Platform: Reference architecture

PRECINCT

Objective: Re-use existing reference architectures or specify a architecture during the project derived from past EU funded projects, the state of the art, etc.

"A reference architecture are abstractions of concretes architectures from a certain domain and serves to design these concrete architectures in multiple domains. Reference architectures facilitate system design and development in multiple projects. Their design and application take place in a broader and hence less defined context with a larger and less defined stakeholder's base."

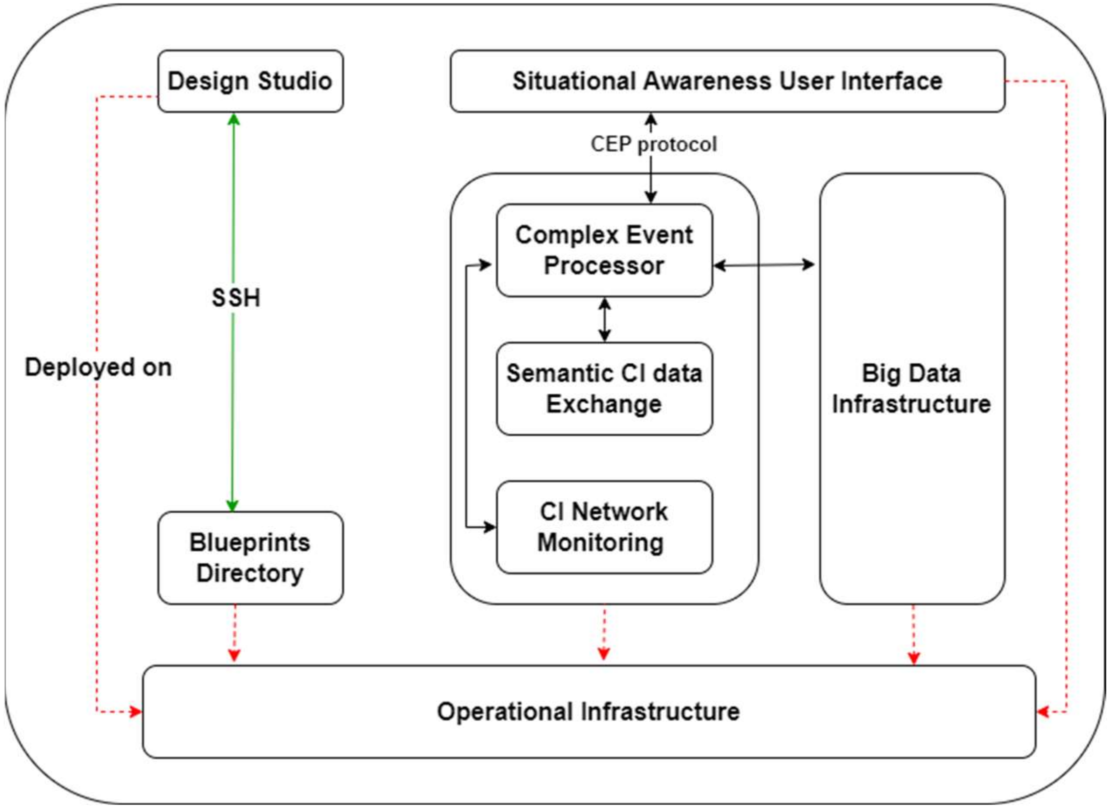

PRECINCT Ecosystem Platform

This project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 101021668.

PRECINCT Ecosystem Platform: Concrete implementations

Objective: From the reference architecture, derive a concrete architecture to be use for implementation by composing a set of building block which provides the identified capabilities..

Describe the implementation using OASIS TOSCA.

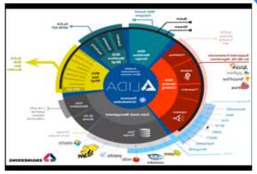

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under Grant Agreement No. 101021668.

OASIS TOSCA as Blueprint Description Language

PRECINCT

TOSCA is a domain-specific language for designing services and for defining the deployment and run-time management aspects of these services with the goal of enabling fully automated service management to support all three phases of the service lifecycle:

1. **Day 0—Service Design**: Service designers use TOSCA to model services as topology graphs that consist of nodes and relationships. Nodes model the components of which a service is composed, and relationships model dependencies between these service components.
2. ***Day 1*—Service *Deployment***: TOSCA can also be used to define mechanisms for deploying TOSCA service topologies on external platforms.
3. **Day 2—Service Management**: TOSCA can enable run-time management of services by providing support for updating and/or upgrading deployed services and by providing service assurance functionality.

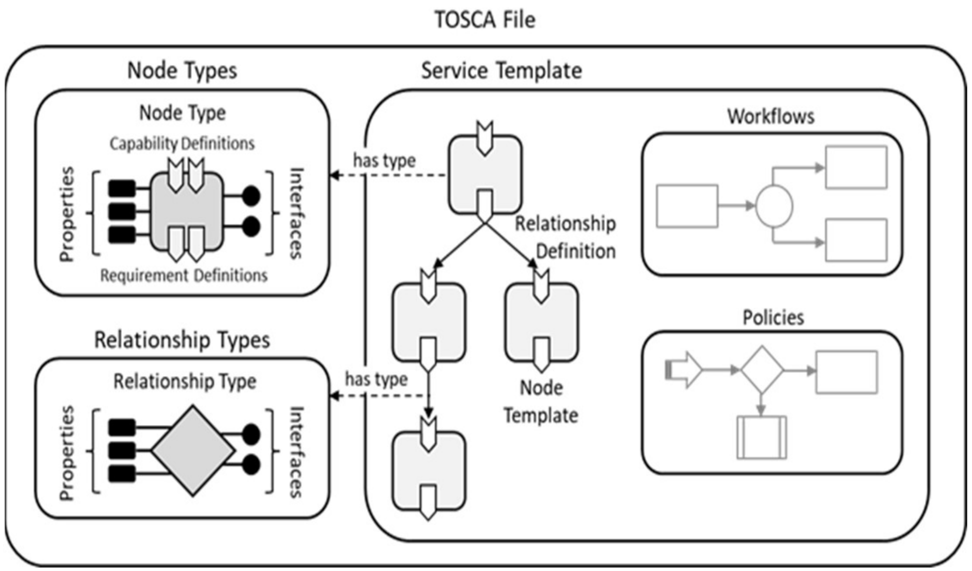

OASIS TOSCA Service Template: https://docs.oasis-open.org/tosca/TOSCA/v2.0/TOSCA-v2.0.html

This project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 101021668.

A TOSCA Service Template in PRECINCT

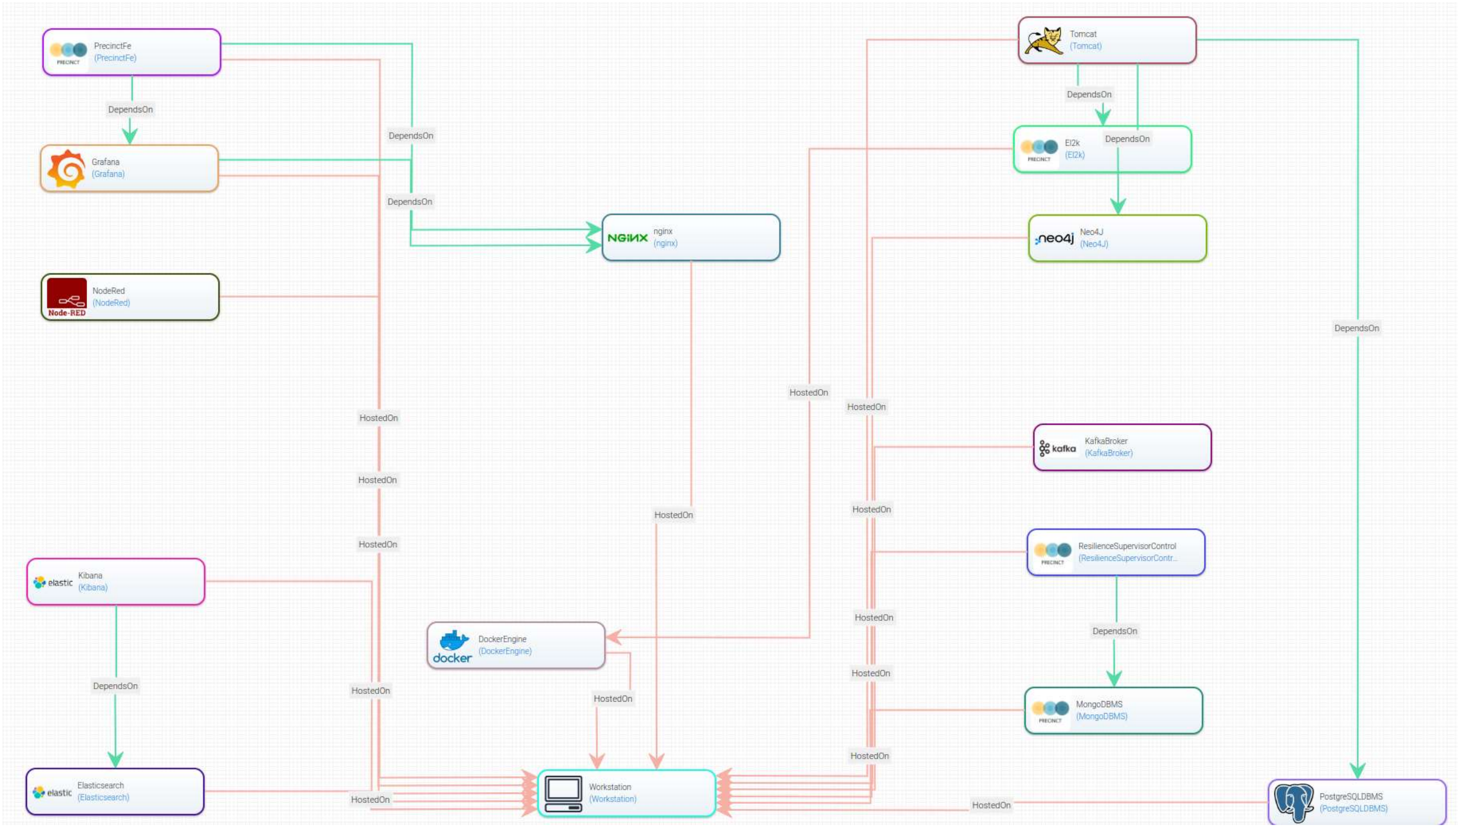

Exemple of an OASIS TOSCA Service Template used to deploy a PRECINCT Living Lab assets

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under Grant Agreement No. 101021668.

Our exeprimenting approach for the deployment of the applications

PRECINCT

Objective: Test the deployment of the CIP assets in an operational infrastructure and test several virtualisation techniques and recommends their usage for CIP operators

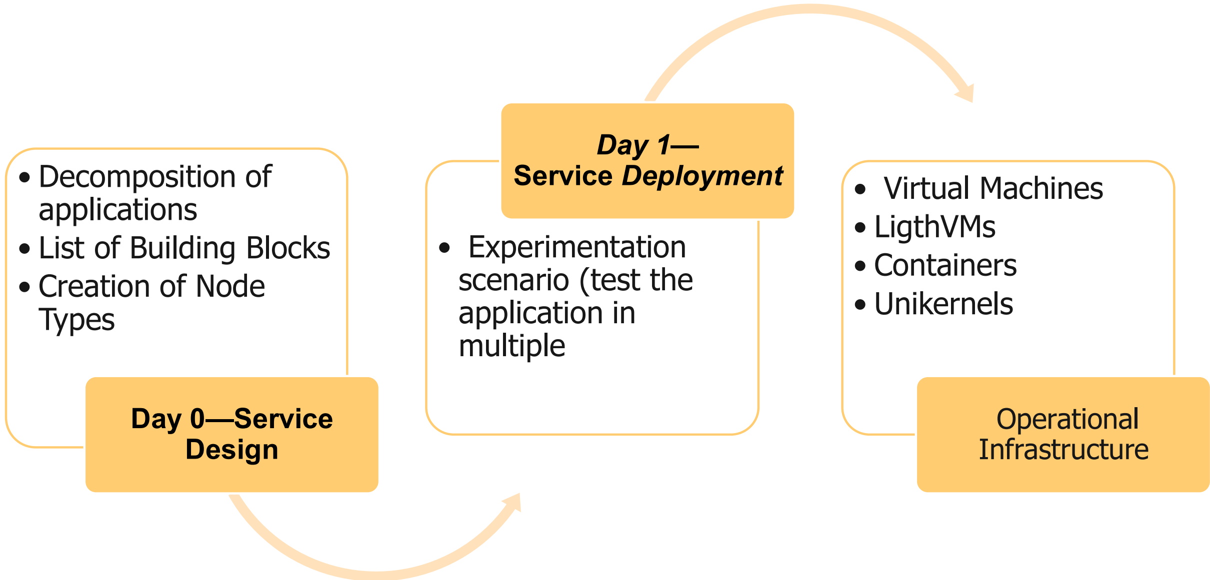

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under Grant Agreement No. 101021668.

Deployment using VM

PRECINCT

- Type 1 and type 2 hypervisor virtualization virtualize of the entire operating system resulting into instance being designated as virtual machine (VM) .
- Type 1 hypervisors are found in data centers to provide different users or services on the same machine or cluster.
- Type 2 are used in personal workstations for sandboxing different applications, tests environments, etc. and exploit different Operation Systems
- Example of Type hypervisor for SCADA systems virtualization Virtual Machine in SCADA systems. Source.
- In our experiments, we used VMs to host the Docker, etc.

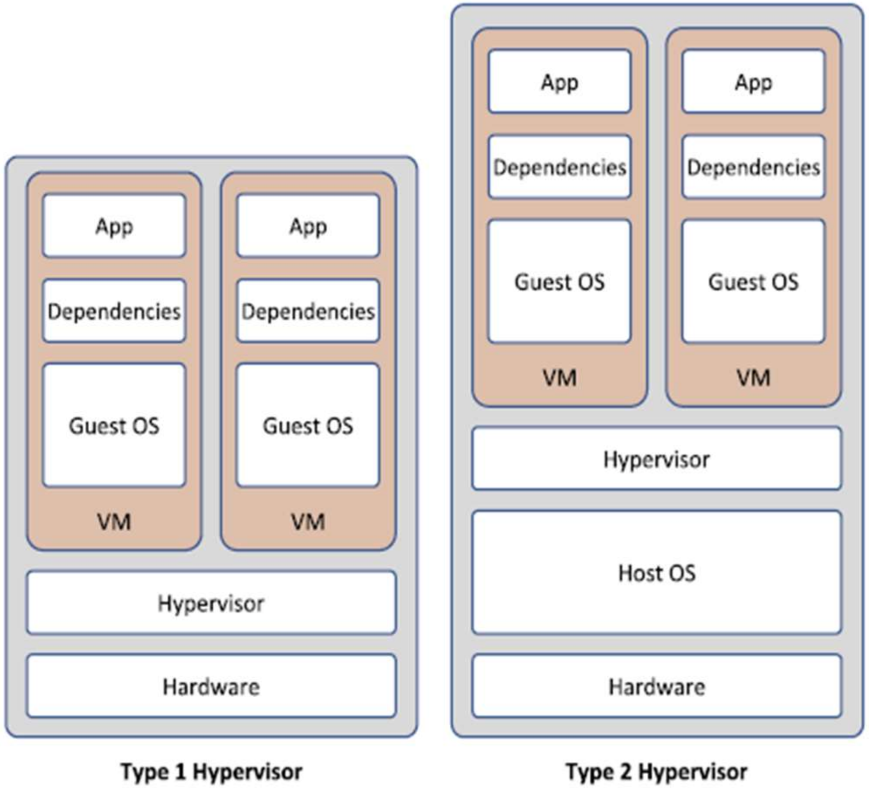

Source: Rui Queiroz, Tiago Cruz, Jérôme Mendes, Pedro Sousa, and Paulo Simões. 2023. Container-based Virtualization for Real-time Industrial Systems—A Systematic Review. ACM Comput. Surv. 56, 3, Article 59 (March 2024), 38 pages. https://doi.org/10.1145/361

This project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 101021668.

Deployment using Containers

PRECINCT

"While hypervisors provide a complete hardware platform abstraction layer, container-based virtualization is supported by a thin layer provided by kernel-level mechanisms to host a wrapped package agglomerating code and all the dependencies for its execution."

Most of the PRECINCT CIP software assets have been provided as Docker containers. (cf. slides 14).

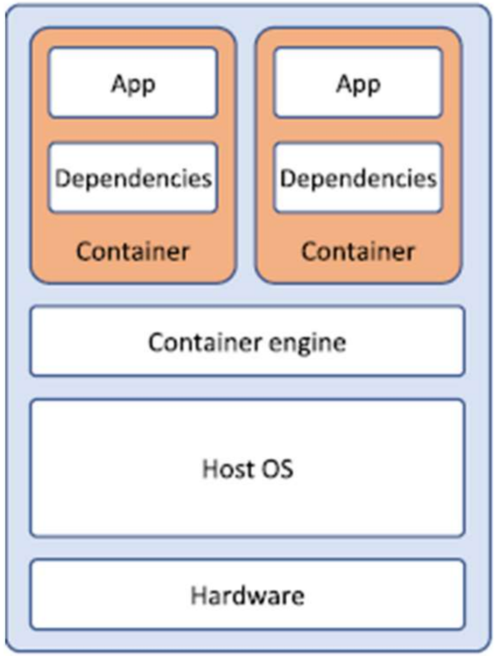

Container

Source: Rui Queiroz, Tiago Cruz, Jérôme Mendes, Pedro Sousa, and Paulo Simões. 2023. Container-based Virtualization for Real-time Industrial Systems—A Systematic Review. ACM Comput. Surv. 56, 3, Article 59 (March 2024), 38 pages. https://doi.org/10.1145/361

This project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 101021668.

Deployment using LightVMs

PRECINCT

LightVMs to replace containers but retains:
• Fast Instantiation: to reduce the boot time of virtual machines closer to container boot times
• High Instance Density: being able to instantiate many virtual machines as containers
• Pause/unpause: Along with short instantiation times, containers can be paused and unpaused quickly.

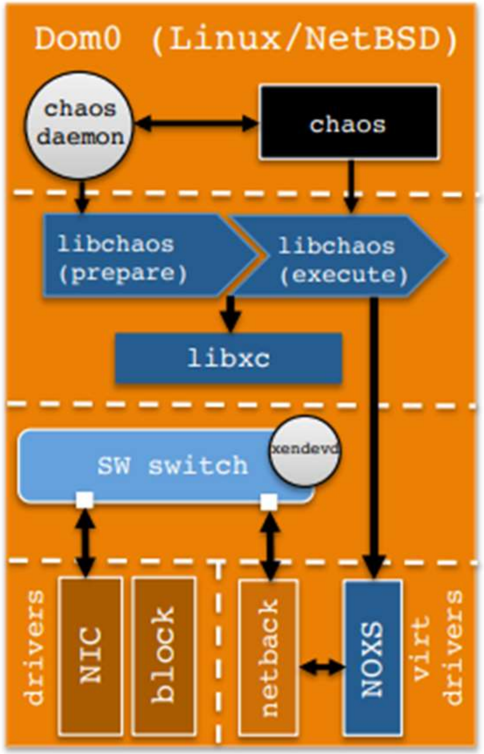

Example of Implementations:
- Firecracker (https://firecracker-microvm.github.io/)
- Katacontainers (https://firecracker-microvm.github.io/)

Source: *Filipe Manco, Costin Lupu, Florian Schmidt, Jose Mendes, Simon Kuenzer, Sumit Sati, Kenichi Yasukata, Costin Raiciu, and Felipe Huici. 2017. My VM is Lighter (and Safer) than your Container. In Proceedings of the 26th Symposium on Operating Systems Principles (SOSP '17). Association for Computing Machinery, New York, NY, USA, 218–233. https://doi.org/10.1145/3132747.3132763*

This project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 101021668.

Deployment using Unikernels

PRECINCT

"Unikernels are specialized, single-address-space machine images constructed by using library operating systems."

- **Specialized:** a unikernel holds a single application.
- **Single-address space:** unikernel does not have separate user and kernel address space (more on this later).
- **Library operating systems** are the core of unikernel systems. The following sections will explain these concepts in more details.

- Exemple:
- Unikraft: https://unikraft.org/

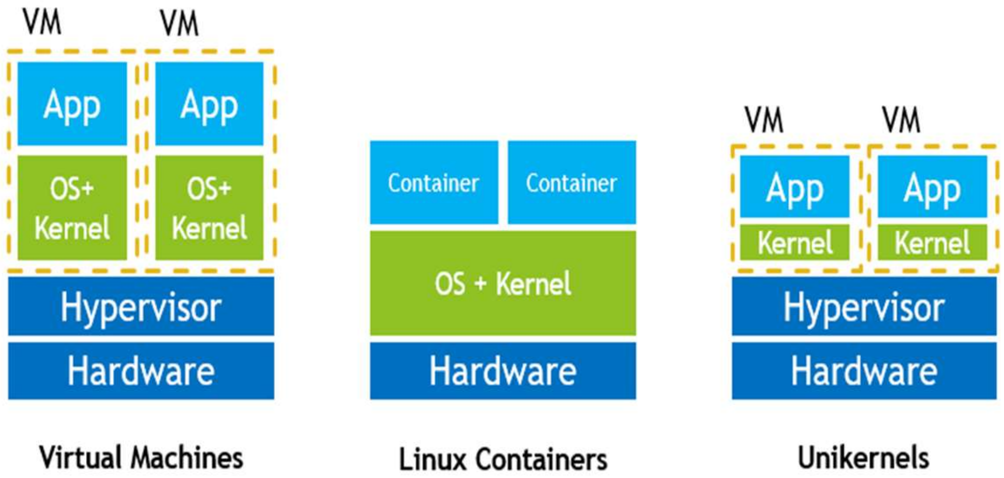

Comparison of Unikernels, Containers and VMs. Source: https://github.com/cetic/unikernels#unikernel

This project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 101021668.

Conclusions and perspectives

- Lessons Learned from CIP projects
  - Significant results are achieved during each project separately
  - These results are not always re-usable/shared since tailored for a specific project
  - IT assets are deployed manually or using specific deployment tools/scripts

- Used approach in PRECINCT WP2:
  - Define or re-use reference architectures in CIP projects
  - Describe the implementation using OASIS TOSCA for deployment
  - Experiment on the deployment using

This project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 101021668.

PRECINCT

Thank you for you attention!

Akkodis

Djibrilla Amadou Kountche

Djibrilla.amadou-kountche@akkdis.com

https://www.akkodos.com

This project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 101021668.