NetApp

White Paper

Industry 4.0 made easy

ShopFloor 4.0

Juergen Hamm, NetApp
December 2021 | WP-7339

In partnership with

Fraunhofer IESE

objective partner

intel

Abstract

This white paper provides an architecture overview of the Industry 4.0 middleware ShopFloor 4.0 on a converged infrastructure.

TABLE OF CONTENTS

LIST OF FIGURES

Overview

This document describes the main components and the base architecture of the Industry 4.0 middleware solution ShopFloor 4.0 based on Cisco Container Platform (CCP), NetApp® StorageGRID® object-based storage software, and NetApp Trident.

ShopFloor 4.0

Industry 4.0 solution ShopFloor 4.0

In today's market, manufacturing systems are designed for the mass production of identical goods. Although manufacturing systems often have a certain flexibility, they are not entirely changeable—and changes are associated with high costs. On the other hand, changeable production allows manufacturers to respond more quickly to changing demands and produce small batches more efficiently. The Digital Twin is a key concept for implementing the required changeability. But this changeability cannot be achieved by integrating individual Industry 4.0 devices into the production system. Instead, new architectural concepts are needed. ShopFloor 4.0 defines a reference architecture for production systems that enables the transformation to Industry 4.0. ShopFloor 4.0 is a commercially supported solution of Open-Source Middleware Eclipse BaSyx. This solution is delivered by three partners: Fraunhofer IESE, objective partner, and NetApp.

To go live with this industry solution, NetApp is also working with Cisco and Intel. ShopFloor 4.0 addresses use cases mainly for automotive and manufacturing. But other use cases in healthcare, the public sector (Smart Cities), and other verticals are planned as well. Customers in these industries need to run the IoT-enabled scenarios with the lowest TCO and reduced production downtimes. This is where Intel's Optane persistent memory technology plays a crucial role. This white paper is the starting point to defining the solution workflow for a specific industry solution, developing a reference architecture, and providing the customer with the information and guidance on how to run them with the highest level of performance availability. NetApp and its partners ensure that the solutions and applications can be restarted and reused as quickly as possible by retaining the data during a server downtime or power outage, which can happen more often in an industrial environment on a factory floor than in a traditional data center.

Use cases and advantages

The realization of a changeable production and the possibility to build Digital Twins of your assets, processes, or products are important. Still, they are not the only Industry 4.0 use cases that ShopFloor 4.0 can address. Others use cases include:

- Cost savings from reduced production downtimes:
  - Virtual commissioning enables users to virtually plan and test the production of a product. It is not necessary to stop production during virtual commissioning.
  - Predictive maintenance prevents unpredictable production downtimes, as well as unnecessary downtimes due to excess maintenance. Monitor the status of your production lines and benefit from optimized on-demand maintenance schedules.
- Automated documentation – digital processes:
  - Provide automatic per-workpiece and per-product documentation with the product Digital Twin. With ShopFloor 4.0, the product Digital Twin contains all relevant documentation, which can be passed to stakeholders and customers.
  - Document the invisible. Digital Twins can document more than quality. It can also demonstrate sustainability, the CO_2 footprint of a product, and everything else that is deemed relevant.

Competitive advantage – changeable production:

- Economic production of low volumes enables a quick reaction to changing market trends and the ability to process low-volume orders.
  - ShopFloor 4.0 integrates your programmable logic controllers (PLCs) into the next generation ShopFloor, enables customized products at the cost of mass production, and makes “lot size one production” a reality.
- Improve process understanding:
  - Understand your process and locate areas of improvement. ShopFloor 4.0 enables the development of dashboards and complete virtual control stations.
  - ShopFloor 4.0 provides a holistic view of your production process and enables a quick evaluation of capacity, resource use, and manufacturing quality.
  - With cycle monitoring, you can detect deviations from planned production cycles and address them online while your production is up and running.
- Added business value and new business models:
  - Improve customer experience—involve your customers in the production as never before.
  - Changing of orders—not a problem until the actual production starts.
  - New business models. ShopFloor 4.0 provides end-to-end data access and is, therefore, an enabler for new business models.
- Easily extensible and low risk to adopt:
  - ShopFloor 4.0 is easily extensible through a unified, HTTP REST-based API. Connect your existing software to ShopFloor 4.0 and access all production data.
  - ShopFloor 4.0 can be introduced on a step-by-step basis—a big bang approach is not necessary. Focus first on the most critical features, which creates the foundation for further uses.

ShopFloor 4.0 in RAMI

The reference architecture model for Industry 4.0 (RAMI 4.0) is a DIN standard and an international prestandard (IEC PAS 63088). The RAMI 4.0 structures Industry 4.0 into three dimensions and defines the relevant components such as the Asset Administration Shell (AAS), such as an equivalent to a standardized Digital Twin. The AAS concept is put into practice by ShopFloor 4.0.

Figure 1 illustrates the ShopFloor 4.0 in RAMI architecture model.

Figure 1) ShopFloor 4.0 in RAMI architecture model (graphic supplied by Fraunhofer).

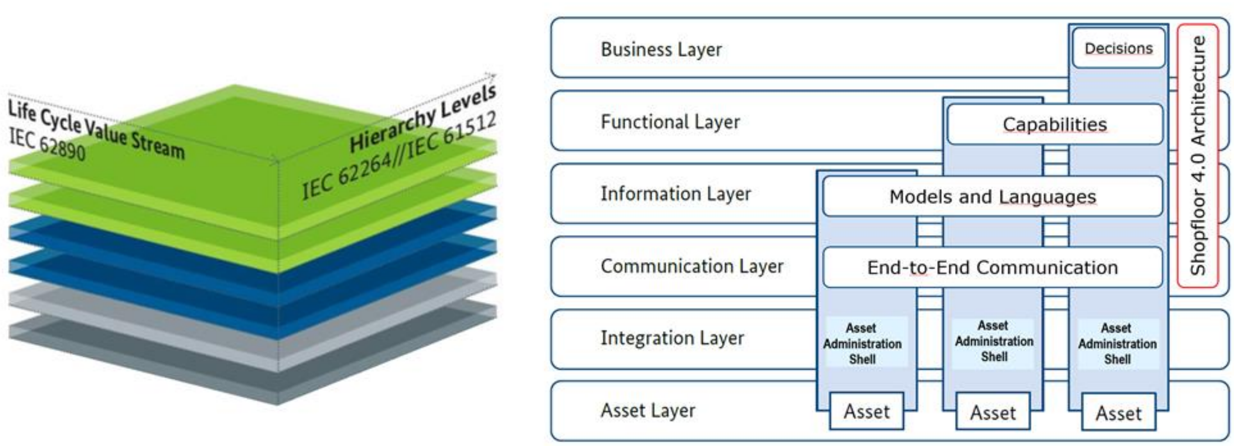

BaSys 4 middleware can be classified in the function, information, and communication layers of the RAMI metamodel of the Industry 4.0 platform. BaSys 4 defines the interfaces both to process components and applications in IT. These include enterprise resource planning (ERP) systems, for example, but also include analysis applications. Various adapters are available for the device connections. These adapters map the native protocols of the devices to the BaSys 4 data models. In addition, BaSys 4 defines the interfaces and submodels that map central aspects of the manufacturing processes.

Regarding the hierarchies defined by RAMI, the BaSys 4 scope can be seen in the control areas (control device) up to the work center. Likewise, BaSys 4 covers the area types and asset instances, as well as the development and runtime views.

BaSys 4

The BaSys 4 architecture is an Industry 4.0 reference architecture driven by Industry 4.0 key use cases. These uses cases include, but are not limited to, process documentation, the monitoring of manufacturing, and the enabling easier changes to plant systems. A central concept of the BaSys 4 architecture is the concept of the AAS.

The AAS is specified by Platform Industry 4.0—it provides access to both properties and functions through standardized interfaces. Additionally, a defined meta-model specifies a common minimum set of information that all Digital Twins must export. Each AAS has multiple submodels attached. In each submodel, different functionalities of the Digital Twin can be realized. Furthermore, views enable the splitting of information by type, but also by the role (for example, type of person or connecting machine) that might be interested in it. For example, one submodel might describe the physical dimensions of the asset, while another submodel might provide access to sensor data, and other submodels might provide information on the availability of spare parts.

By adding specialized submodels, AAS can be used to implement different versions of the Digital Twin. For example, submodels can provide static data, such as a digital data sheet. Submodels can also provide access to live data, enable the controlling of a machine, and implement their own behavior to monitor data streams or the availability of spare parts. By using AAS and submodels, it is possible to provide a harmonized, service-based interface to the whole production system in which physical and virtual devices, products, workers, and other assets are represented by their own object instance.

Figure 2 illustrates the BaSys 4 architecture.

Figure 2) BaSys 4 key architecture (graphic supplied by Fraunhofer).

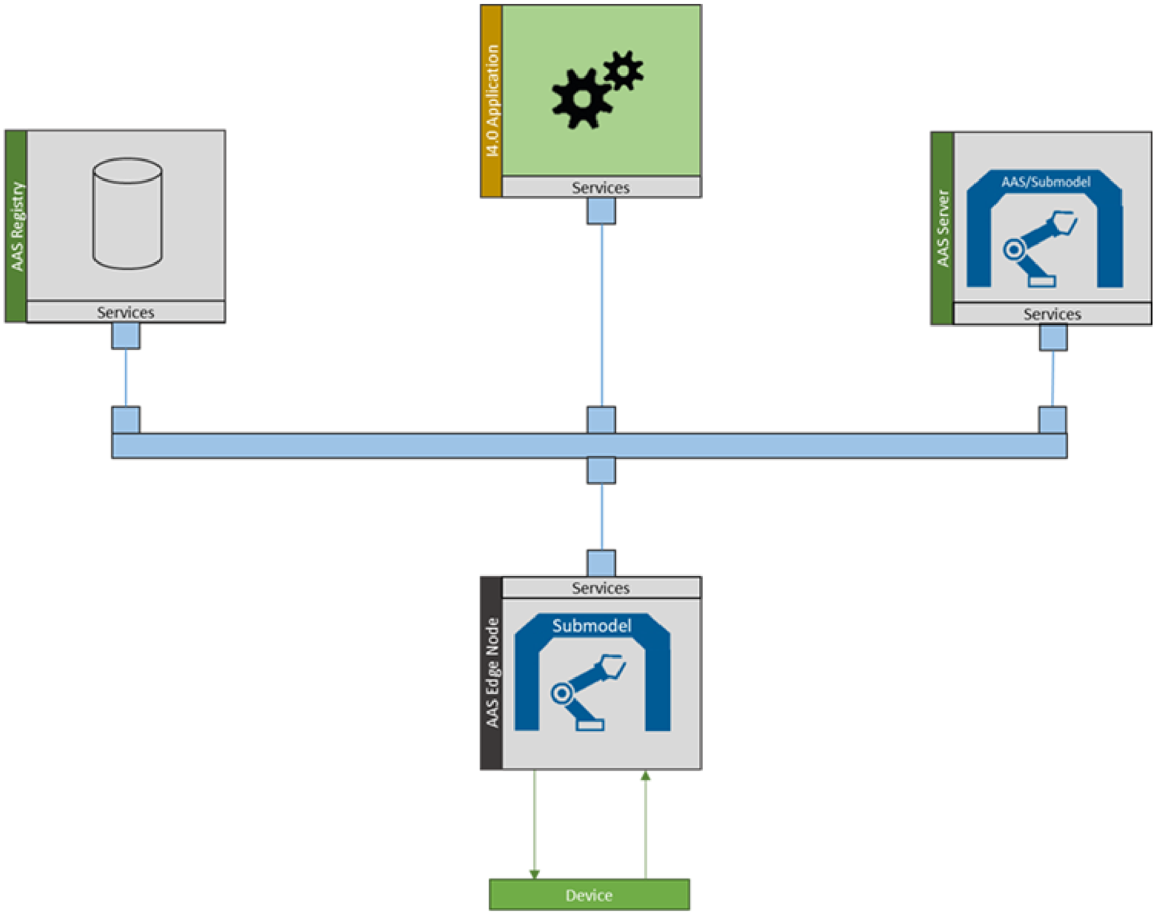

The BaSys 4 architecture consists of several key components:

- **AAS Registry.** The AAS Registry acts as a central entry point into the AAS system. This registry is used by Industry 4.0 applications to look up available AAS and to retrieve access information: for example, network endpoints and metadata.
- **AAS Cloud Server.** The AAS Cloud Server provides the capability to upload and store AAS at runtime and in a persistent manner, in contrast to hosting AAS in-memory on a device. This feature increases the availability of an AAS in the event of a device failure. This feature also enables efficient retrieval of information on devices; for example, ordering its replacement in case of failure. This AAS Server can be part of a private or a public cloud.
- **AAS Edge Node.** Similar to the AAS Cloud Server, an Edge Node is also able to host an AAS. In contrast to the AAS Cloud Server, the AAS Edge Node does not support the upload of AAS at runtime. Instead, it provides preconfigured AAS, for example, with sensor data and digital nameplate. A device AAS is provided by the device manufacturer for modern smart devices, or by a third party.

ShopFloor 4.0 architecture

Architecture considerations

The ShopFloor 4.0 architecture supports important quality criteria required by the central Industry 4.0 use cases. This section describes these architecture considerations.

Scalability

Scalability enables the user of an architecture to quickly react to quantitative requirement changes; for example, numbers of users or models to be provided by the architecture. The ShopFloor 4.0 architecture was built with scalability in mind. Thus, the key architecture components and setup support easy scalability. Each component of the architecture is available as a Docker image. Thus, scaling the number of devices and products (such as deploying a new AAS) is as easy as deploying a new container. Additionally, the AAS Cloud Server supports uploading and deleting AAS and its accompanying submodels with easy-to-use API calls. Furthermore, by using converged infrastructure (see the section titled "Converged infrastructure"), scaling up the application through Docker images is made easy.

Finally, the AAS Registry can be scaled by using sub-registries and a cascading access.

Note: Scalability is referred to as Adaptability in ISO 25010.

Maintainability

Maintainability, such as the ease of changing a system, is a key consideration of the ShopFloor 4.0 architecture. Thus, deploying or redeploying AAS or updating existing models is supported by various means. The various off-the-shelf images provided by the ShopFloor 4.0 architecture support the AAS configuration and infrastructure with little to no implementation effort. Instead, the off-the-shelf images can easily be configured through text-based property files, as documented in the Eclipse BaSyx wiki. Additionally, changing or deleting an already existing AAS is supported by a comprehensive API. In the same manner, uploading a new AAS is also possible.

By leveraging the concept of an AAS Registry, it is possible for applications to discover existing AAS and incorporate their data into their algorithms. Additionally, based on several filter criteria, applications can tailor search queries to their specific needs. Through an event mechanism, applications are always up-to-date and informed about changes of registered AAS.

In addition to these benefits, the AAS Registry offers another key advantage—ease of redeployment. If a requirement changes, moving an AAS or submodel might be necessary. Because the AAS Registry keeps track of endpoints, this redeployment is completely transparent to applications.

Lastly, the structure of the AAS with its submodels enables an easy extension of existing models. If a new application requires additional aspects of assets that are currently not covered by existing submodels, the AAS and AAS Registry dynamically support supplementing existing models with new data.

Network considerations

In contrast to the Internet of Things (IoT), Industry 4.0 has specific network requirements. Distributed control loops, for example, realized by a Supervisory Control and Data Acquisition (SCADA) system, typically have hard real-time constraints. Thus, data must be available with only milliseconds (or even faster) of a delay. Additionally, devices create massive amounts of data, estimated to be in the range of several terabytes (TBs) per day per plant. This number is expected to multiply with a factor of 5 to 10 within the next few years.

Implementing a public cloud-based solution would require first uploading the data to the public cloud before being able to use it for control. However, this will most likely introduce unpredictable latency, which is not acceptable for real-time control of manufacturing. In other extremes, deploying all applications in the edge will not be feasible due to resource constraints of Edge Nodes standing in harsh contrast to the requirements of big data analytics and AI solutions.

As a result, a hybrid approach was chosen. This approach allows placing preprocessing in the edge and storing data with high requirements on latency in an on-premises cloud while still leveraging the benefits of a public cloud.

Thus, the ShopFloor 4.0 architecture supports a distributed deployment of data and Digital Twins. The AAS Server can be placed in a public or private cloud. The AAS Edge Node enables the handling of time-critical data in the edge, thus enabling distributed control loops.

Converged infrastructure

Experimenting with platform integration often results in complex systems that require expertise that is difficult to find. For this reason, NetApp and its partners worked together to deliver a converged infrastructure for ShopFloor 4.0, which doesn't require endless maintenance. This solution quickly and easily delivers your IoT applications with outstanding control, speed, and scale.

This converged infrastructure for ShopFloor 4.0 solution is designed to simplify your environment for today and tomorrow. With this architecture, you can accelerate every application without disruption, lower license costs, dramatically improve TCO, leverage built-in hybrid cloud capabilities, and prepare your infrastructure for the future with confidence. These solutions continue to evolve with innovative technology, delivering the proven performance, agility, and value required for a changing enterprise.

Hardware architecture

This document describes the Cisco and NetApp HW ShopFloor 4.0 solution, which is a tested approach for deploying NetApp and partner technologies as a shared infrastructure. This design provides a framework for deploying VMware vSphere, the most popular virtualization platform in enterprise-class data centers.

Figure 3 is an overview of the hardware architecture.

Figure 3) Hardware architecture.

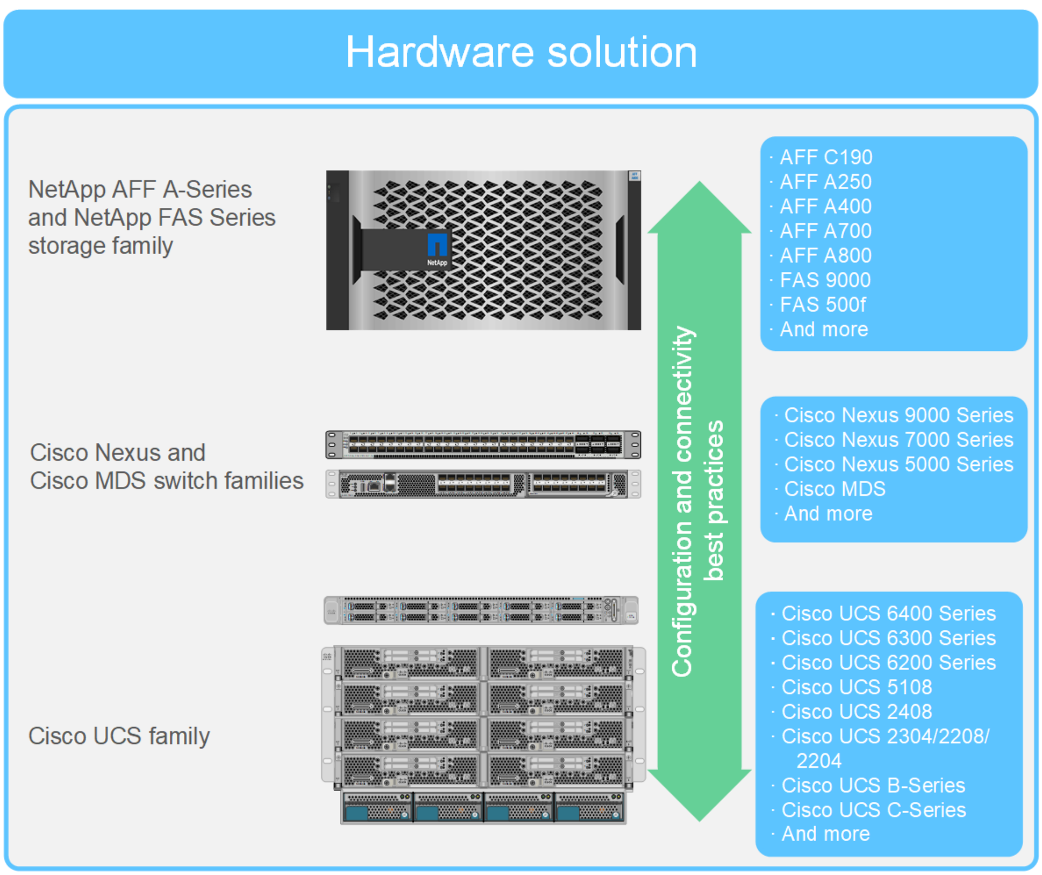

Computing

Cisco Unified Computing System (UCS) blade and rack servers deliver enterprise-class performance, reliability, and expandability to your ShopFloor 4.0 environment. These servers harness the power of the second-generation Intel Xeon Scalable processor family. Intel Optane™ DC persistent memory enables you to increase memory density per server to better serve your ShopFloor 4.0 middleware, lower costs, and achieve a much faster restart if required.

The Intel Scalable Xeon processor allows the required scalability in a larger industry environment with thousands of IoT devices, which results in much higher throughput for solutions such as ShopFloor 4.0. It is essential to have a backbone with the highest performance level (using Intel Scalable Xeon processors in Cisco UCS) to ensure real-time transaction and interconnectivity between the industry devices such as robots, machines, sensors, and PLC devices, and the downstream workflow through the BaSys 4 framework to the backend system.

For a smaller starting environment, this platform allows customers to consolidate other customer applications, such as SAP/4HANA and SAP Data Intelligence, and establish improved OLAP- and OLTP-based processes.

Storage

Named a leader for its AFF A-Series in the 2019 Gartner Magic Quadrant for Primary Storage, NetApp has a 25-year history of innovative technology. Based on NetApp ONTAP® data management software, these systems enable end-to-end solutions to integrate tightly with standard enterprise applications for seamless operations.

Networking

Cisco UCS and NetApp AFF storage systems are connected through Cisco Nexus 9000 Series Switches, providing 100 and 40 Gigabit Ethernet (100GbE and 40GbE) connectivity together with Cisco Application-Centric Infrastructure (ACI). These high-performance, low-latency switches offer exceptional power efficiency at cloud scale to your ShopFloor 4.0 environments. Cisco SingleConnect technology delivers an easy, intelligent, and efficient way to connect and manage your ShopFloor 4.0 computing resource.

Cisco Container Platform

CCP is a fully curated, lightweight container management platform for production-class environments, powered by Kubernetes and delivered with Cisco enterprise-class support (Figure 4). It reduces the complexity of configuring, deploying, securing, scaling, and managing containers through automation coupled with Cisco's best practices for security and networking.

Figure 4) Cisco Container Platform.

CCP is built with an open architecture (Figure 5) using open-source components. It works across both on-premises and public cloud environments.

Figure 5) Architecture of CCP.

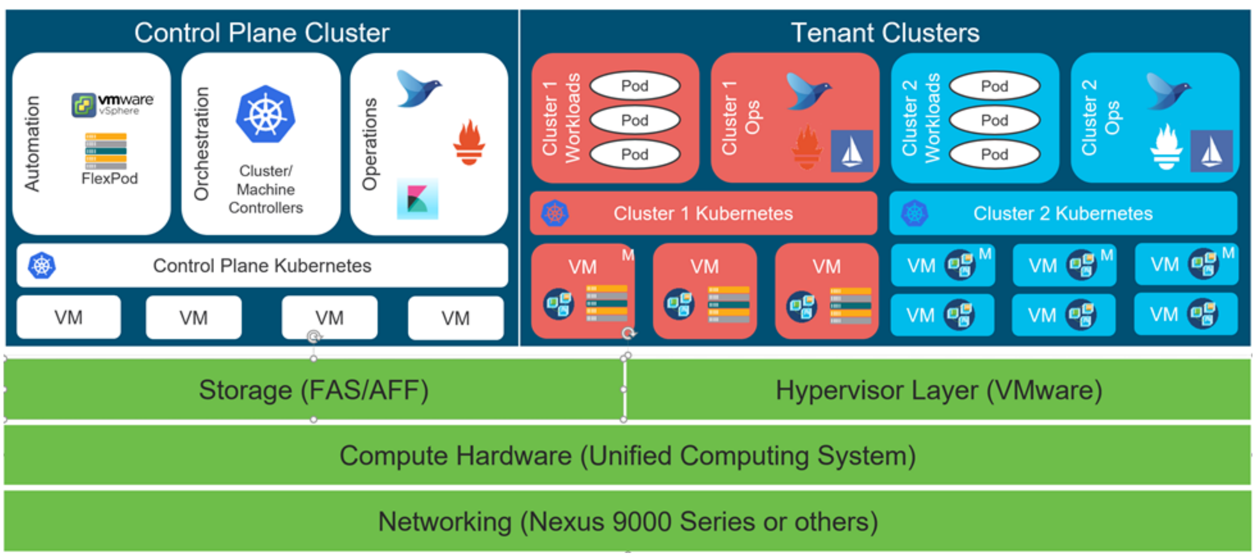

Cisco Container Platform architecture

At the bottom of the stack is level 1, the networking layer, which can consist of Cisco Nexus switches, Cisco Application Policy Infrastructure Controllers (APICs), and fabric interconnects.

Note: CCP can run on top of a Cisco ACI networking fabric as well as on a networking fabric other than Cisco ACI that performs standard L3 switching.

Level 2 is the computing layer, which consists of Cisco UCS, or third-party servers that provide virtualized computing resources through VMware and distributed storage resources.

Level 3 is the hypervisor layer, which is implemented using VMware ESXi.

Level 4 consists of the CCP control plane and data plane (or tenant clusters). In Figure 5, you can see that CCP has two types of clusters: control-plane clusters and tenant clusters. Tenant clusters can support a single master or multiple-master Kubernetes cluster for your container applications. These tenant clusters are preconfigured to support persistent volumes using the VMware vSphere Cloud Provider and Container Storage Interface (CSI) plug-in. Tenant clusters can be configured with multiple masters (master high availability) to maintain a steady state across all the worker nodes in the event of a single master failure.

In addition to being easy to deploy, CCP provides the following benefits:

- **Reduced risk.** CCP is a full stack solution built and tested. It provides automated updates and enterprise-class support for the entire stack. It is built to handle production workloads.
- **Greater efficiency.** CCP provides your IT operations team a ready-to-use, preconfigured solution that automates repetitive tasks and removes pressure on staff to update people, processes, and skill sets in house. It provides developers with the flexibility and speed they need to be innovative and to respond to market requirements quickly.
- **Remarkable flexibility.** CCP gives you a choice of deployment on virtual environments: from hyperconverged infrastructure to vSphere ESXi clusters. And because it is based on open-source components, you are free from vendor lock-in.

CCP is a turnkey, production grade, extensible platform to deploy and manage multiple Kubernetes clusters. It runs on 100% upstream Kubernetes. CCP offers seamless container networking, enterprise-grade persistent storage, built-in production-grade security, integrated logging, monitoring and load balancing.

Monitoring solution

To achieve the highest availability with the infrastructure, and to detect problems and bottlenecks as early as possible, it is imperative that all components of the solution are continuously monitored. The infrastructure monitoring approach is based on an open monitoring system that combines best-of-class, open-source components with Cisco and NetApp modules acting as the interface to the hardware components building this stack. Figure 6 shows the layered monitoring architecture.

Figure 6) Monitoring architecture.

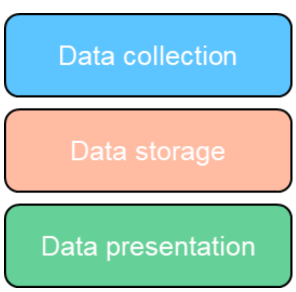

- **Data collection.** This layer is where Cisco and NetApp provide modules that are designed to collect information from their infrastructure components using appropriate APIs (such as SNMP and NetApp ZAPI). Additional modules for various software layers are provided by open-source projects (for example, hypervisor and container platform). There is an easy-to-use interface to add additional application-specific metrics to the system if needed. The amount of information and time granularity can be customized at any time.
- **Data storage.** This layer is a time-series database (in this example, InfluxDB) that provides the performance and capabilities to support a high-ingest data rate of datapoints with queries showing information in real-time (Figure 7).

Figure 7) Data flow monitoring.

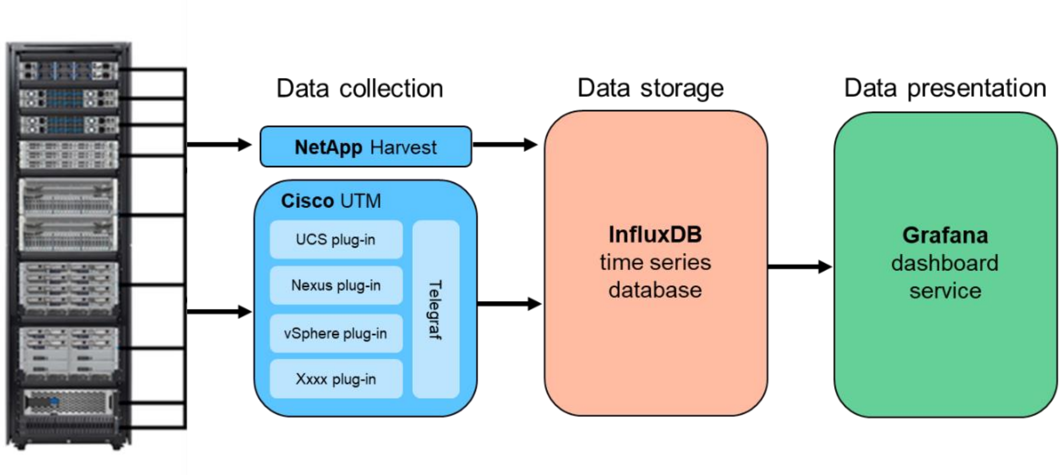

- **Data presentation.** This layer uses the Grafana dashboard system (Figure 8), which comes with a rich set of predefined dashboard views that can be easily extended to match the operational requirements of the customer.

Figure 8) Monitoring dashboard.

Note: There are dashboards that provide a quick overview of the complete stack with the option to dive deep into any detail as needed. Individual components can be directly retrieved from the appropriate download pages.

NetApp Trident

Trident enables microservices and containerized applications to leverage enterprise-class storage services (such as QoS, storage efficiencies, and cloning) to meet an application's persistent storage demands. The Kubernetes volumes are static. Depending on an application's requirements, dynamic provisioning of Kubernetes volumes can be performed with the help of Trident.

Trident is used to dynamically provision storage from the following sources:

- NetApp ONTAP data management software running on NetApp AFF or FAS systems

Trident uses the storage class object that was introduced in Kubernetes 1.4 to dynamically provision persistent volumes (PVs) when a persistent volume claim (PVC) object is created. A storage class allows administrators to describe the classes of storage that they offer. A storage class might map to different QoS levels, backup policies, or other storage characteristics. For more information, see the Trident documentation.

You can automate the management of datasets and the inference results by using Trident to dynamically provision PVs. Compared to manual data management by a storage administrator, this automation is inexpensive, simple, convenient, and time saving.

NetApp StorageGRID object storage

NetApp StorageGRID is a software-defined object storage solution for large archives, media repositories, and web data stores. StorageGRID uses intelligent policy-driven data management to store, protect, and preserve fixed-content data over long periods of time, seamlessly moving data between on-premises and public cloud storage to optimize its availability, protection, performance, and cost.

Content is placed in the right location, at the right time, and on the right storage tier—optimizing workflows and reducing costs for globally distributed rich media. In Industry 4.0 scenarios, the data management infrastructure must store an almost infinite amount of data types in distributed locations and it must allow centralized access. With the S3 interface, StorageGRID offers the possibility to integrate the IoT data into your ERP and other business applications as well as into public clouds.

The basic building block of a StorageGRID system is the grid node. Nodes contain services, which are software modules that provide a set of capabilities to a grid node. The values and statuses for all the functions of the StorageGRID system are reported through attributes.

The StorageGRID system uses four types of grid nodes:

- **Admin nodes** provide management services such as system configurations, monitoring, and logging. Each grid has one primary admin node and might have any number of additional admin nodes for redundancy.
- **Storage nodes** manage object data and metadata storage, including data protection.
- **API gateway nodes** provide a load-balancing interface to the StorageGRID system through which applications can connect to the system. Alternatively, a third-party HTTP load balancer can be used.
- **Archive nodes** provide an interface through which object data can be archived to tape.

Architecture example

Architecture diagram

The architecture diagram shown in Figure 9 is based on a BaSyx introductory example by Fraunhofer. The objective of this use case is to control a simulated edge device and obtain access to its data.

This example was adopted to demonstrate the various possibilities to distribute the AAS Edge Node depicted in Figure 2. In other scenarios, there might not be a need to distribute the AAS, submodel, and other components over more than one Docker container. If various components with different aspects and multiple functions and operations deployed within one Docker container or multiple is depending on the customer environment and use case. For example, for an AAS that contains generic descriptions, extensive CAD documentation or video tutorials should be deployed as an AASX file (a package file format extension for the serialized AAS) on the AAS Server.

Submodels, control components, or even the AAS itself can be deployed onto edge hardware that is connected through a bus connection and running closely to the asset.

The example scenario depicts efficient communication between the asset and the asset control component based on BaSyxNative. As soon as the transaction data reaches a node with a larger computing and networking capacity, it is preprocessed and sent to StorageGRID for persistence.

Compared to the high-frequency but immutable transaction data that is stored in an object storage, metadata is modifiable and stored in a database.

Each of the boxes depicted in the Kubernetes area of Figure 9 represents a container image. The AAS Server and the AAS Registry are off-the-shelf-components that are provided either as an executable jar or a Docker image.

Interfaces that can be implemented exist for each component in order to realize new backends. All of them are seamlessly integrating into the REST API provider.

The configuration and integration of the deployed application images are automated through Helm charts. When deployed as ShopFloor 4.0, the required Helm charts are included.

Figure 9) Architecture diagram (graphic supplied by objective partner).

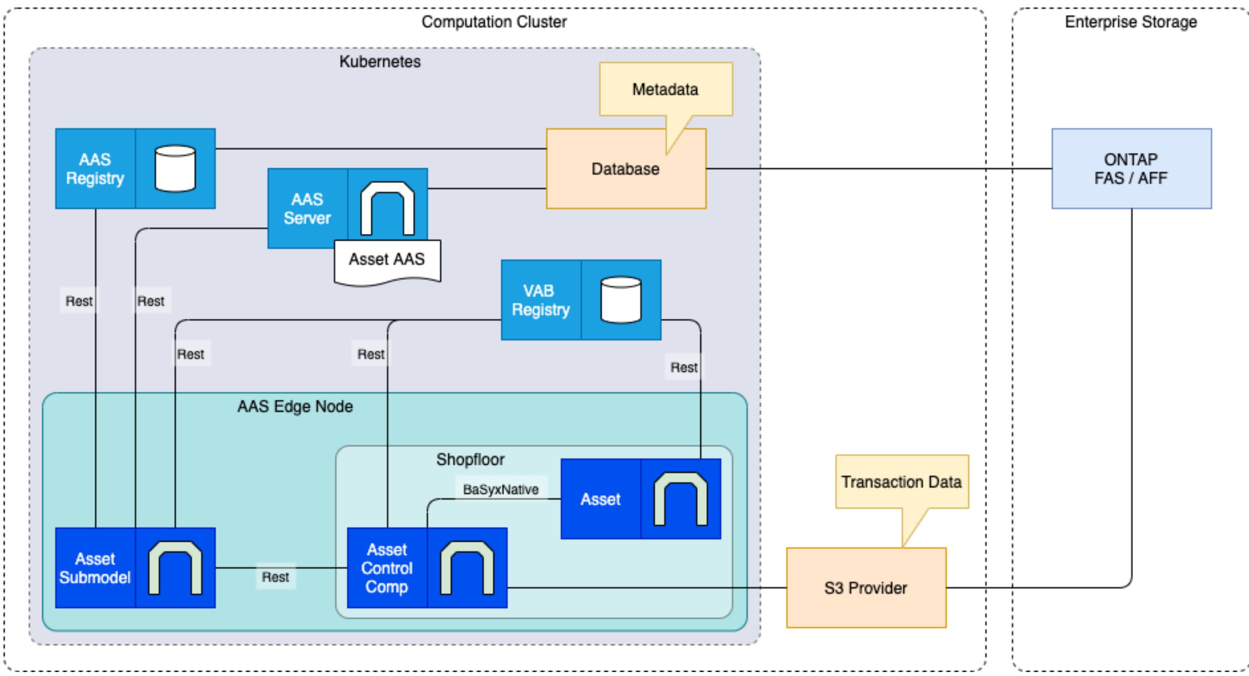

Software infrastructure

The software infrastructure shown in Figure 10 illustrates the data management.

Transaction data is transmitted from the Docker containers through Trident and NFS onto NetApp AFF and FAS storage systems.

Figure 10) Software infrastructure (graphic supplied by objective partner).

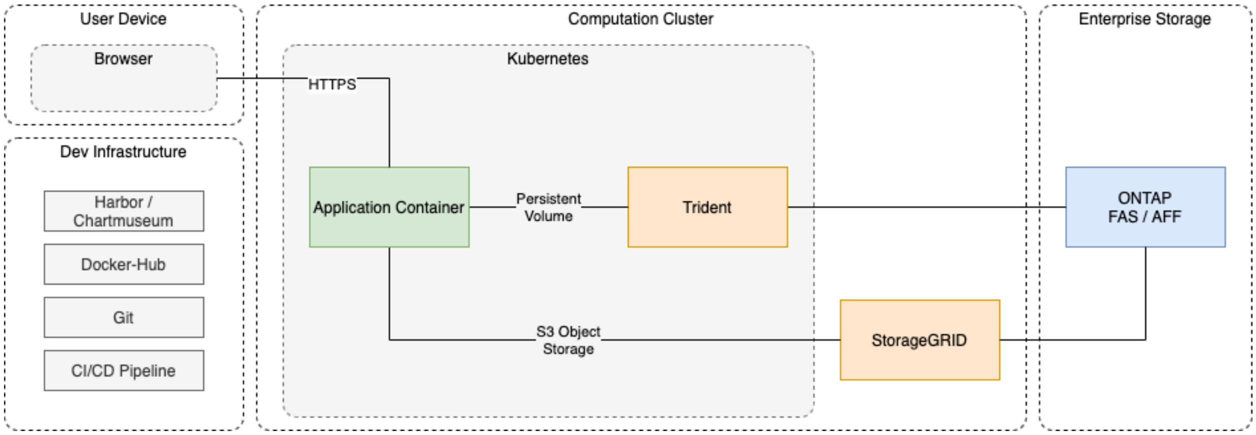

The depicted software infrastructure is either based on the example described in “Architecture diagram,” or it could be any custom case when deployed with Shopfloor 4.0.

The development infrastructure shown in Figure 10 is optional. The application is accessed through HTTP by using a browser. The configured ingress in Kubernetes is connected to the application container.

Docker images

To support an easy deployment, ShopFloor 4.0 provides commercial off-the-shelf components as Docker images ready to be deployed in the HW infrastructure. For each component (AAS Registry, AAS Cloud Server, and so on) rich configuration files exist. Thus, you can set up the ShopFloor 4.0 infrastructure with little configuration effort and without needing programming experience.

The configuration of the images allows you to select backends (for example, SQL and MongoDB) as well as support relevant AAS exchange formats (JSON, XML, and AASX). Therefore, the ShopFloor 4.0 components integrate out-of-the-box with other open-source tools such as the AASX Package Explorer. The documentation of these off-the-shelf components is described here.

Conclusion

The tested ShopFloor 4.0 architecture combines the service-based Industry 4.0 middleware BaSys 4 from Fraunhofer with the well-tested converged infrastructure using components from Cisco and NetApp. The architecture is completed by the CCP and data management infrastructure from NetApp. This ShopFloor 4.0 solution is designed to simplify your Industry 4.0 environments for today and tomorrow. This pretested stack can accelerate your journey to an Industry 4.0 ready production environment. It allows you to upgrade your applications without disruption, lower license costs, dramatically improve TCO, leverage built-in hybrid cloud capabilities, and prepare your OT infrastructure for the future with confidence. This solution continues to evolve with innovative technology—delivering the proven performance, agility, and value required for a changing IT and OT to facilitate your changeable production.

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

NetApp Trident

- Deploy NetApp Trident CSI Plug-in on Cisco Container Platform
https://www.cisco.com/c/dam/en/us/solutions/collateral/data-center-virtualization/unified-computing/trident-on-container-platform-with-flexpod.pdf

NetApp StorageGRID

- StorageGRID documentation resources
https://www.netapp.com/data-storage/storagegrid/documentation/)
- Hardware Installation and Maintenance Guide
https://docs.netapp.com/sgws-110/topic/com.netapp.doc.sga-install-sg5700/SG5700%20appliance%20installation%20and%20maintenance.pdf

Cisco Container Platform

- CCP Installation Guide
https://www.cisco.com/c/en/us/td/docs/net_mgmt/cisco_container_platform/6-1/Installation_Guide/ccp-installation-guide-6-1-0.pdf
- Cisco Container Platform for SAP Data Hub
https://www.cisco.com/c/dam/en_us/solutions/global-partners/ccp-data-hub.pdf
- Elastic Cloud on Kubernetes with Cisco HyperFlex and Cisco Container Platform - Cisco
https://www.cisco.com/c/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/elastic-cloud-on-kubernetes-on-cisco-hx.html

Intel

- Cloud Storage Solutions from NetApp and Intel
https://partner-connect.netapp.com/us/partner-directory/intel-corp/cloud-storage-solutions-from-netapp-and-intel

Other

- BaSys/BaSyx
https://wiki.eclipse.org/BaSyx
- Monitoring Harvest NetApp Harvest download page
https://mysupport.netapp.com/site/tools/tool-eula/netapp-harvest/download
- Grafana
https://grafana.com
- InfluxDB
https://www.influxdata.com

Acknowledgements

We would like to thank the following people for their invaluable contributions to the creation of this document:

- Jürgen Hamm, NetApp
- Georg Mey, NetApp
- Alexander Kern, NetApp
- Andreas Bader, objective partner
- Dr. Thomas Kuhn, Fraunhofer IESE
- Frank Schnicke, Fraunhofer IESE
- Ulrich Kleidon, Cisco

Version history

Version	Date	Document Version History
Version 1.0	February 2021	Initial version written by Jürgen Hamm.
Version 1.1	September 2021	Second version written by Jürgen Hamm.

Refer to the [Interoperability Matrix Tool (IMT)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

WP-7339-1221