evolve

Leading the Big Data Revolution

D1.1 Public Project Presentation

Deliverable type
report

Dec. 24th 2018

Work package
WP1

Leader
DDN

Dissemination level
public

Author
Contributor

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

DDN, BULL, IBM, FORTH, ONAPP, ICCS, MEMO, WLT, LOBA, TAS, SPH, CYB, NEURO, MEMEX, TIEMME, VIF, AVL, BMW, KOOLA

Contributors

The Deliverable 1.1 being based on materials presented by all partners during the project kick-off meeting the list of contributors encompasses authors from all partners. The following persons have therefore contributed to the realization of this document.

DDN	Jean-thomas Acquaviva	
ATOS	Nam Nguyen	
IBM	Christian Pinto	Michael Johnston
FORTH	Angelos Bilas	Christos Kozanitis
ONAPP	Michail Flouris	Stelios Louloudakis
ICCS	Dimitrios Soudris	Sotirios Xydis.
MEMOSCALE	Per Simonsen	Rune Erlend Jensen
WEBLYZARD	Max Goebel	Arno Scharl
LOBA	Alexandre Almeida	Catarina Pereira
SPH	Socrates Costicoglou	
CYBELETECH	Marion Carrier	Teddy debroutelle
NEUROCOM	Frédéric Hernu	Vassilis Spitadakis
MEMEX	Giorgio Ambrosino	
TIEMME		Mauro Pallari
V2C2	Alexander Stocker	Gernot Lechner
AVL	Eric Armengaud	Bernhard Peischl
BMW	Ulrich Fastenrath	
THALES ALENIA SPACE	Marc Spigai	Michelle Aubrun
KOOLA	Sergio Ermacora	

Peer reviews

Name	Organization
Jean-Thomas Acquaviva	DDN

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

Revision history

version	date	reviewer	modifications
0.1	21/12/2018		Initial template
0.2	24/12/2018		Contribution to content

Executive Summary

EVOLVE is a 3 years projects with 19 parters. The project is centered on the concept of data at scale and is an ambitious community effort to handle the looming data deluge.

Two communities are brought together, the platform providers, to be understood at the hardware and software level and the data providers.

Platform core strategy is to tackle data using heterogeneity with specialized components both from the processing sides and storage aspects. Evolve considers that data are intrinsically structured and depending on this structure the right processing unit is key to deliver speed-up. Similarly for storage, data locality, both spatial and temporal, with the right usage of the storage depth can improve performance by several orders of magnitude.

Ensuring the management of this heterogeneity requires a middle-ware upgrade and the project platform includes major efforts on the middelware layer. This aspect is key to ensure a correct convergence between the HPC and cloud applications, where one the differentiator is the dependence on specific deployment stacks.

On the data producers side, EVOLVE is bringing together partners owning technologies in markets where data capability is already the source of disruption, or which are the turn point of being disrupted. EVOLVE partners bring use cases mostly from mobility (autonomous vehicle, ground mass transportation, maritime transport), agriculture and urban planing. As these markets are socially critical from European citizens, EVOLVE is not a pure technology project but frames itself in the more global perspective of data ownership in an open society.

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

Index of Contents

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

1. Project Value Proposition

1.1. Project structure to bring value

The project is built on the convergence of two main components,

1. platform providers
2. data providers

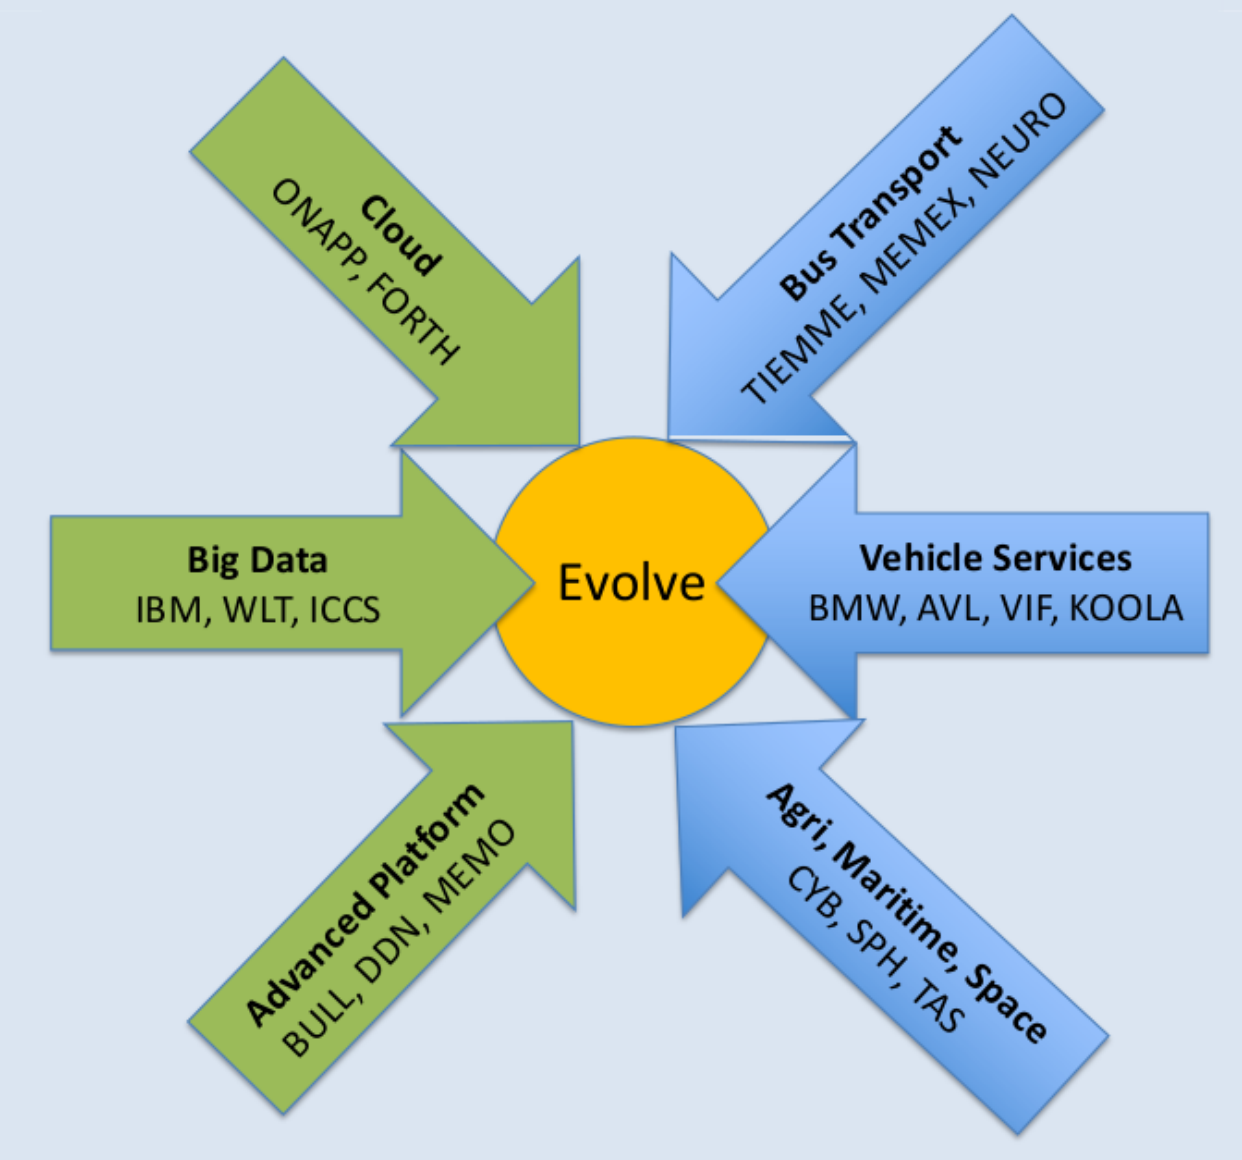

Figure 1: Evolve consortium organization

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

1.1.1. Heterogeneity is key for a successful convergence

Where platforms providers are themselves issued either form HPC or the Cloud market. While these two market are historically disjoint they are now on the verge of convergence. Evolve is therefore an opportunity to realize the many times announced but never realized convergence between Big Data and HPC.

By working jointly on a unique platform partners will mechanically ensure convergence.

The too often overlooked issues in convergence is that the two fields while exhibiting overlap are not identical. The risk of converged platform is to be efficient on the minimal common overlap between the two communities.

To ensure a better market penetration and ensure not only for the converged segments, but to support the specificities of each filed, the key enabling technology in EVOLVE is the usage of heterogeneity. Having the ability to specialized the platform is the guarantee that Cloud and HPC workload will be processed at optimal speed.

1.1.2. Data providers from heterogeneous workload

From a processing roadmap it is clear that accelerators are here to stay and the project will help applications to exploit complex architectures. However, there is an on-going revolution of usages, the pace of changes for applications is currently the innovation prescriber. EVOLVE will consequently help platform providers

→ To architect relevant solutions

→ To think workflow instead of application

As an example it could be pointed that the same data source, Sentinelle-2 satellite images are used for agricultural use case, which crop prediction, urban planing. The multi-usages a single source of data is a fact that may have been noticed by most of platform providers.

1.1.3. Lean project management

From this initial acknowledgement that several communities are represented within the consortium and that in term of data processing we are the verge of the data deluge, the project has opted for a lean management.

Evolve partners will learn from each other in a lean is earning together from each other… in a lean way

Therefore, in its life time the project has planned 3 industrial-grade HW pilots to be delivered during the project [M06, M18, M36]. Section 2, will illustrated the lean approach taken in this project with the delivery of the V0 platform.

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

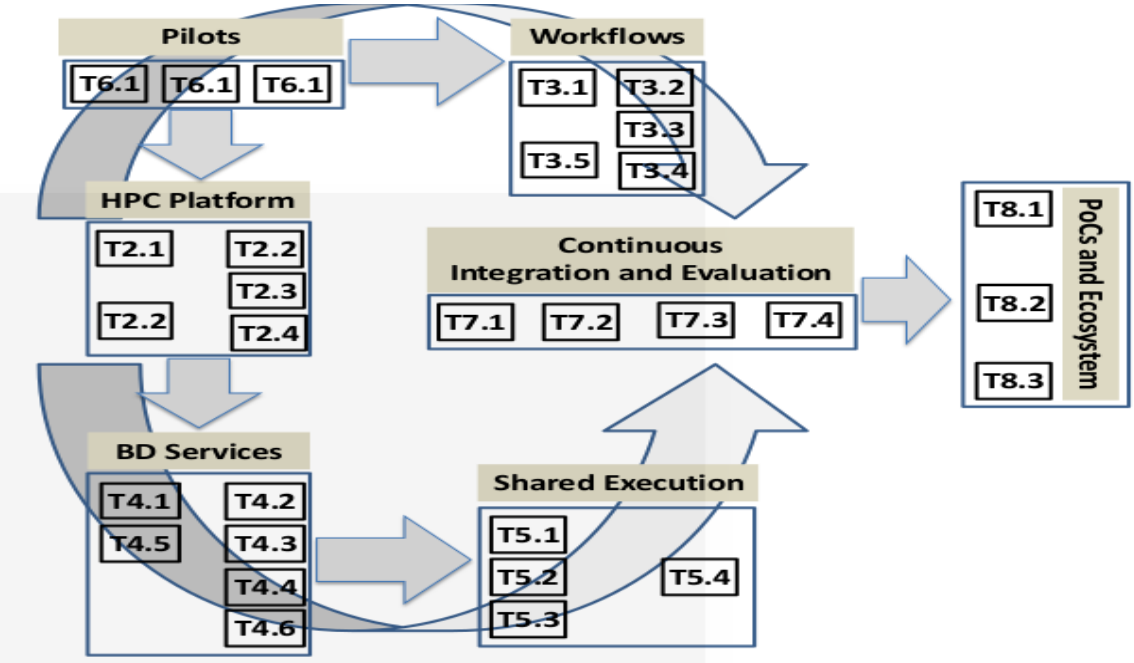

Figure 2: Circular flow and interaction within project Work-package

The lean aspect of the project management is illustrated by Figure 2 where a continuous flow of information appears across the different work-packages toward the continuous integration engine

1.2. Value proposition: 3 pillars

Evolve contribution is structured around 3 axis:

1. Increase productivity (easiness of problem expression)
   → From application to work-flow
2. 2 Computational ubiquity
   → Cloud like platform abstraction
3. 3 Efficiency
   → HPC technology Cycle and Byte optimization

For applications and data producer #1 ensure that the platform will allow a highly level of abstraction. The ability to express easily a problem is a powerful productivity enabler.

The #2 is addressing a quick often neglected aspect in project oriented toward Proof-Of-Concept. The operation side of the platform. With is strong links to Cloud community one

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

important contribution of Evolve is the guarantee that the platform will be easy to operated (which is not always a priority in pure HPC environment). At last, symmetrically to #2, the deep HPC roots of EVOLVE will allow applications and more generally use cases to be run in a efficient way where CUP cycle are not wasted around and byte movement are carefully counted.

Bridge the cultural gap and build a culture of trust

In order to deliver the vision while the consortium is built form various communities it is mandatory that some actions have to undertaken.

First, each partners has to acknowledge the complexity of each other assets. This will be ensure by the mandatory presentation during general meeting of all technical aspects. Platforms and applications work-package are presented jointly and a full work-package is dedicated to integration and testing.

Second, platform providers will set-up tigers team to engage applications. While a formal collaborative process is mandatory to deliver a full fledge solution, early engagement with a small team of highly-qualified experts is a proven scheme to build trust and ramp-up a collaboration. Evolve plan to follow this path.

1.3. EVOLVE in the European perspective

Accelerators are at the core of EU strategy form computing sovereignty. As illustrated by the launch of the EPI simultaneously with the Joint Undertaking HPC initiative

Therefore EVOLVE by addressing accelerated platforms on modern workloads is ideally poised

→ Market trend is with us

→ EU acknowledges and is funding EPI

In term of time line, the European agenda is ideally aligned with EVOLVE schedule since the EVOLVE results will be available for the second phase of the EuroHPC activities (see screen-shot bellow)

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

It is up to EVOLVE consortium to deliver and leverage our strong initial position

2. Timeline and short term schedule

2.1. Hardware platform aggressive schedule

While the platform has been initially planned to be delivered on Month 06, the consortium has successfully set-up a more aggressive schedule thanks to the commitment of ATOS and DDN.

The V0 of the heterogeneous platform will be opened to partners at Month03.

This platform will already support processing heterogeneity with CPU anf GPU nodes.

This platform will already support storage heterogeneity with a capacity hard drive based storage and an multi-Terabyte layer of storage acceleration based on SSD.

2.2. Short term deliverable and milestones: overview

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

A more detailed view of the deliverables an milestones is available in the second part of the document, where a per work-packages presentation is provided.

- D1.1 Public project presentation [M01] [DDN]

→ this document

- D1.2 Quality management plan [M03] [DDN]
- D2.1 Plan for converged hardware platform dimensioning [M6] [BULL]
- D6.1 (a,b,c) Specification of pilot workflow [M06]
- D7.1(a,b,c) Intermediate hardware prototypes [M06] [BULL]
- D7.2(a,b,c) Intermediate software stack prototypes [M06] [IBM]
- D9.1 Dissemination and communication strategy and plan [M03] [LOBA]
- D9.2(a,b,c,d) Data management plan (and revisions) [M03, M12, M24, M36] [LOBA]
- D9.3 Project web page, project stationary, and other dissemination and communication tools [M06] [LOBA]

Milestones

- MS1.1 [M3] Web page and public project presentation.

→ project splash page already available: evolve-h2020.eu

- MS9.1 [M06] Launch of Website
- MS2.1 [M6] Plan for converged hardware platform

→ project converged platform v0 anticipated to M02

Next general assembly

The next general assembly is planed for Month 05 in order to proceed to a review of the deliverables due for Month 06.

3. Dissemination and communication

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

Public communication on the project at SIG I/O UK:

“Getting ready for platform and workload heterogeneity: A technical manifesto from the EVOLVE project”

Workshop on Storage Challenges in the UK, March 6, 2019

https://hps.vi4io.org/events/2019/sig-io-uk

Poster session at the Per3S conference:

Per3S Fifth Edition -- Inria Bordeaux SO in Talence, France -- January 25th, 2019

https://per3s.sciencesconf.org/

Electronic presence

NEWSLETTER

- 2 per year (total: 6)
- Based on GDPR, Distribution: Through EVOLVE’s website subscribers + Partners contacts (as institutional or personal level
- Dissemination through social media

PRESS RELEASES

- Every time there is an important announcement
- First: M2-M3?
- Agility Tool: 800.000 contacts

Dissemination and communication materials

DISSEMINATION & COMMUNICATION MATERIALS

Whenever partners need a communication materials (e.g.: events) please contact LOBA with your requirements and proposed contents and/or structure.

STATIONARY

- Word and PPT Templates
- Folder
- Letterhead paper
- Email Signature
- Business Card

PROMOTIONAL MATERIALS

- Brochures, Leaflets, Flyers
- Posters, Roll-ups
- Videos
- Banners & Images for Social Media

GIVEAWAYS

- Mini notebooks with mini pen
- Keychain with laser pointer
- Luggage tag
- Other suggestions?

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

EVOLVE detailed view on selected work-packages

The second part of this document is based on a sub-set of the slides presented during the kick-off meeting on Dec. 1st..

For the sake of clarity only a selection of materials are reproduced. The purpose is not to present a per work-package detailed view but to provide a higher level understanding of the project and its orientation.

Namely, only the work-page presenting the platform (WP2), middle-ware (WP3); the work-flow (WP4), the use cases (WP6 pilots) and the continuous integration (WP7) are presented.

The work-packages which are not included are of the same level of importance for a successful project but remains more technicals or have been already addressed in the first part of this document.

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

1. WP2: Platform detailed view

Specifications

- **WP2 :** Advanced Computing Platform
- **Main Goals**
  - Design and Development of a HPC Platform
  - Execution of Concurrent Application Workflows
  - Support of RAS, Security
- **Main Challenges**
  - Convergence of HPC, BigData and AI workflows
  - Expansion of HPC's scope
    - Data deluge
    - Increase in compute power at affordable costs
  - Platform to unite heterogeneous computing technologies, and converged software stacks.
  - Project socket => Availability/Reliability/Accessibility/Security

Convergence of HPC and Big Data

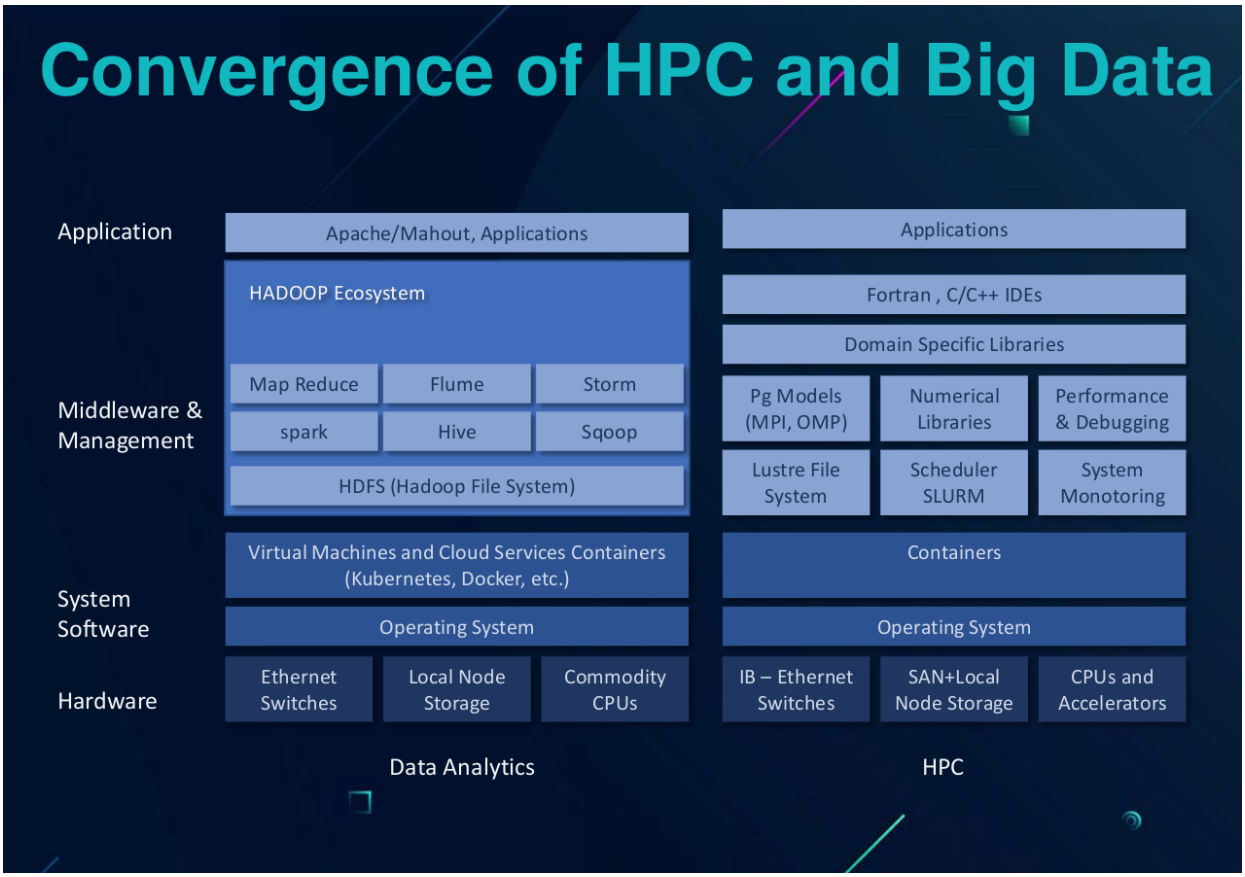

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

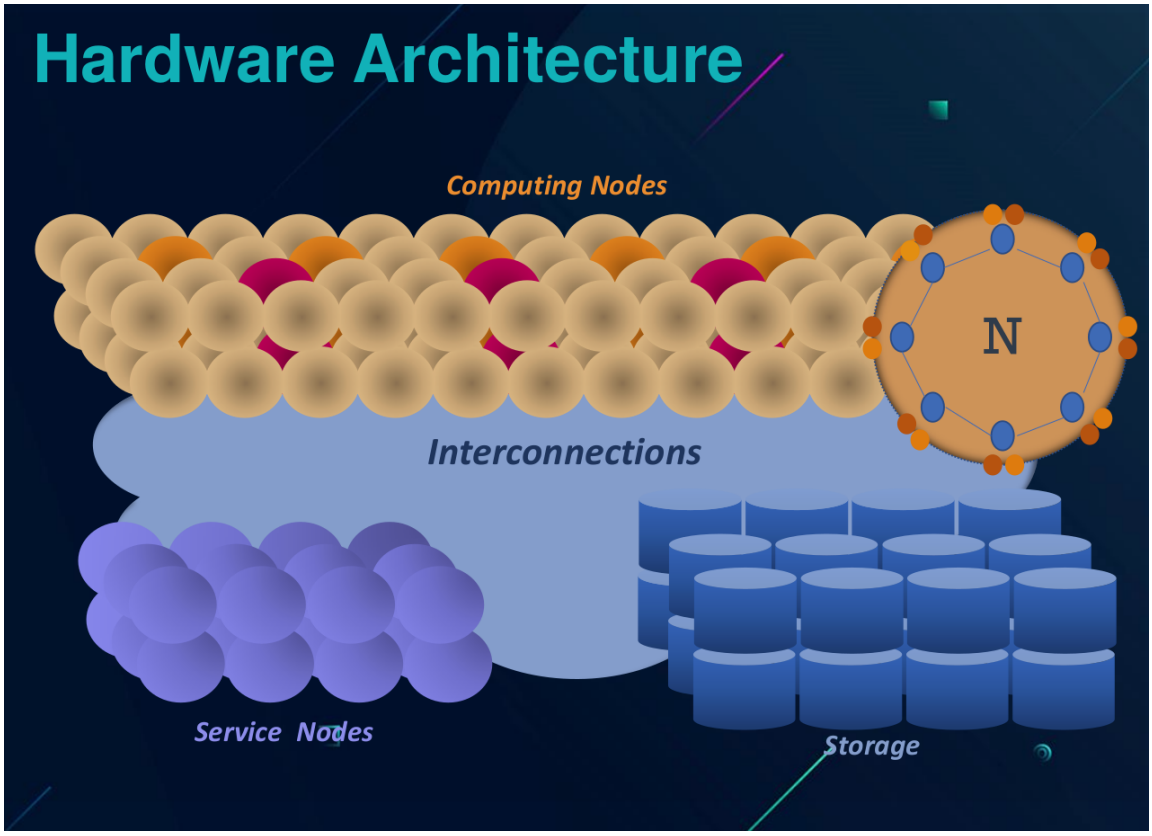

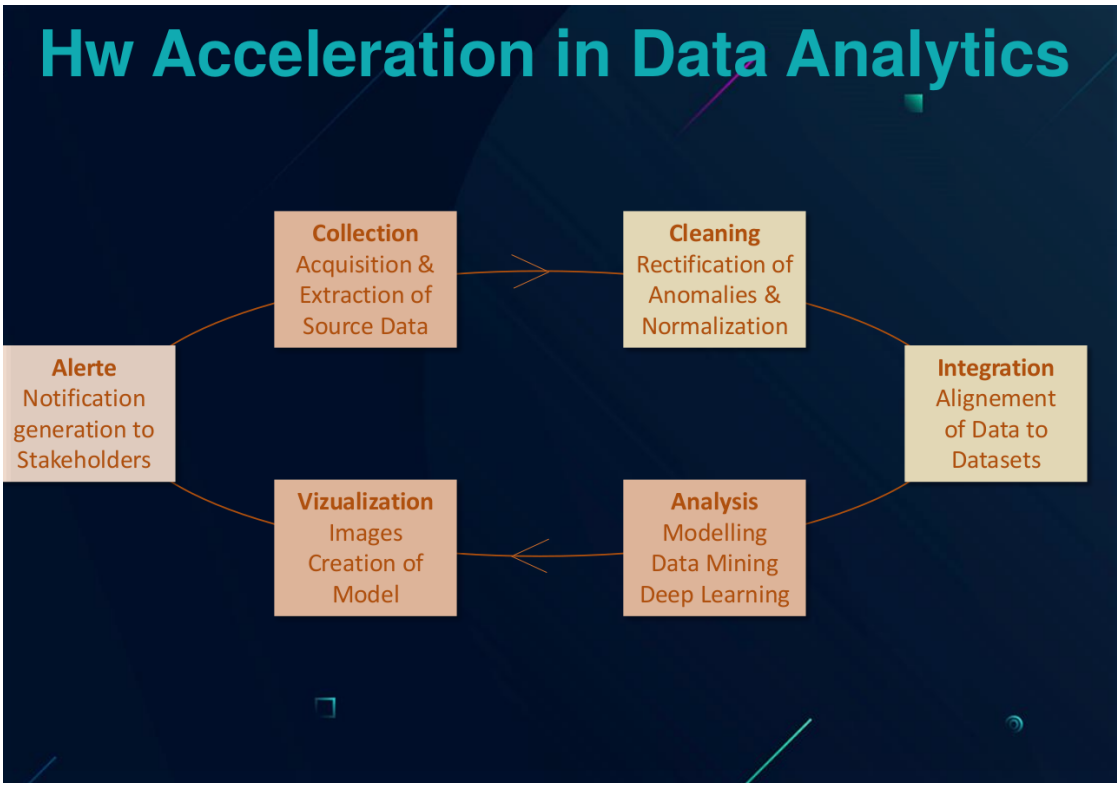

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

- **Partners : Bull**, DDN, FORTH, ONAPP, ICCS, MEMO
- **Duration :** M1-M24
- **Tasks**
  - Task 2.1 Hardware platform dimensioning for integrating HPC features with Big Data processing. [M1-M12] [**BULL**, DDN, MEMO, FORTH, ICCS, ONAPP]
  - Task 2.2 Accelerated processing subsystem. [M1-M24] [**ICCS**, BULL, DDN, MEMO, FORTH]
  - Task 2.3 Storage subsystem hierarchy based on fast devices. [M1-M24] [**DDN**, MEMO, BULL, FORTH]
  - Task 2.4 Data protection via efficient coding over fast devices. [M1-M24] [**MEMO**, DDN, BULL, ICCS]
  - Task 2.5 Support for RAS. [M1-M24] [**MEMO**, DDN, BULL, ONAPP]
  - Task 2.6 Systems software support. [M1-M24] [**ONAPP**, FORTH, ICCS, DDN, BULL, MEMO]

WPs & Tasks interdependance

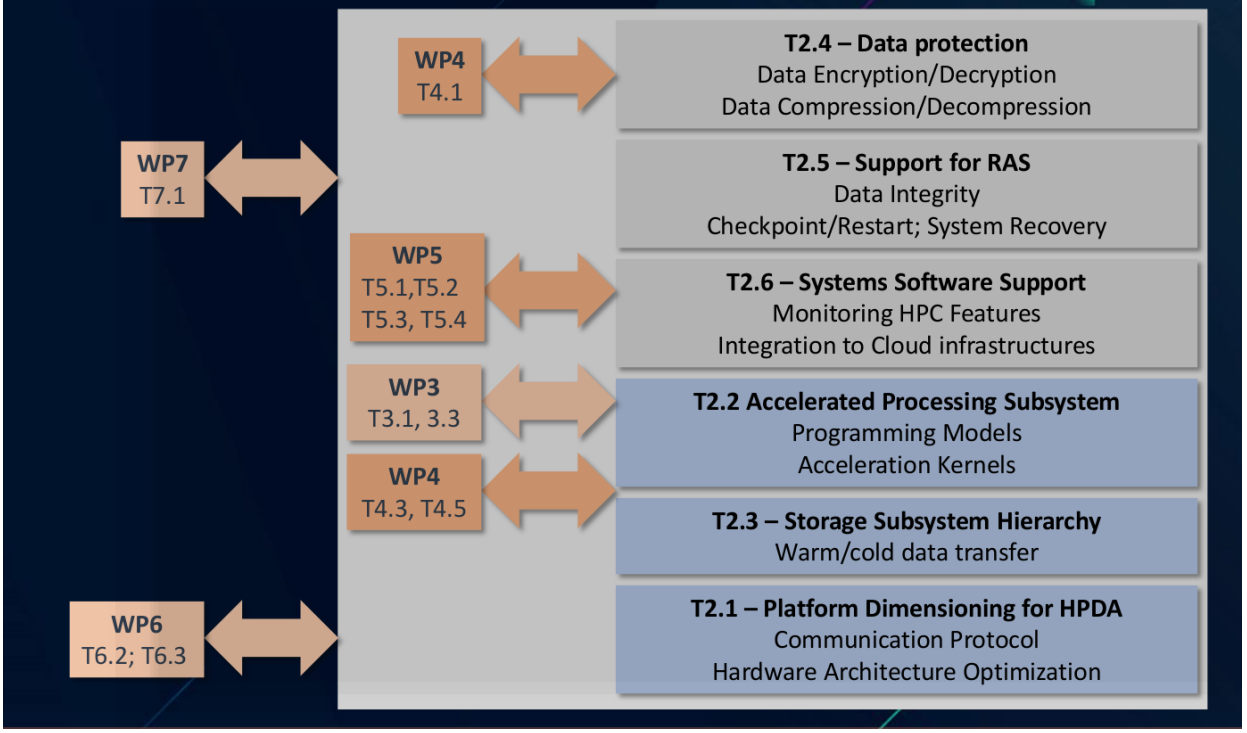

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

Tasks, Milestones & Deliverables

- Deliverables of various nature : Hw, Sw, Documentation
- Iterative delivery process (intermediate deliverables)
- Delivery of Platform V0 on January 2019

Task	M6	M12	M24
Task 2.1	MS2.1/D2.1 Platform Plan		
	MS2.2/D2.2 Converged Hw Platform		
			MS2.3 Access to refined Platform
Task 2.2	D2.3 – Accelerator Features & Capabilities		
Task 2.3	D2.4 Storage Subsystem & fast storage devices		
Task 2.4	D2.5 Data Protection efficient coding over fast devices		
Task 2.5	D2.6 RAS Features		
Task 2.6	D2.7 Systems Sw Support		

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

2. WP3:Middleware

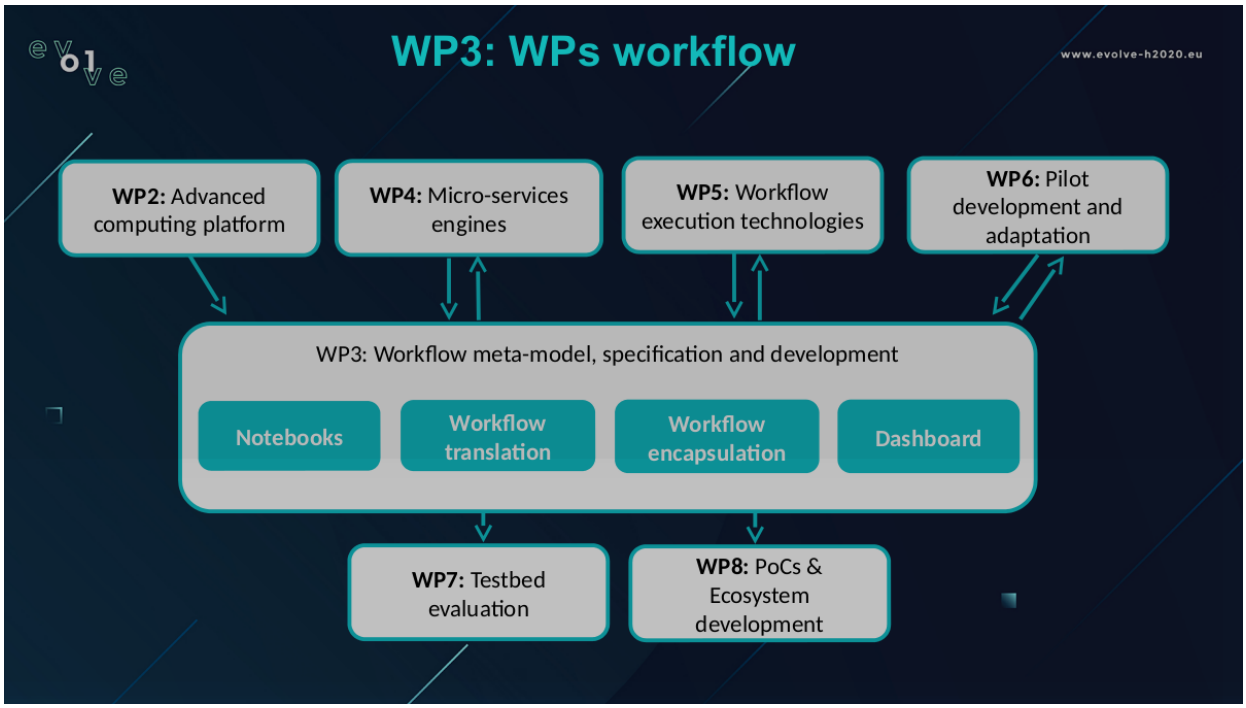

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

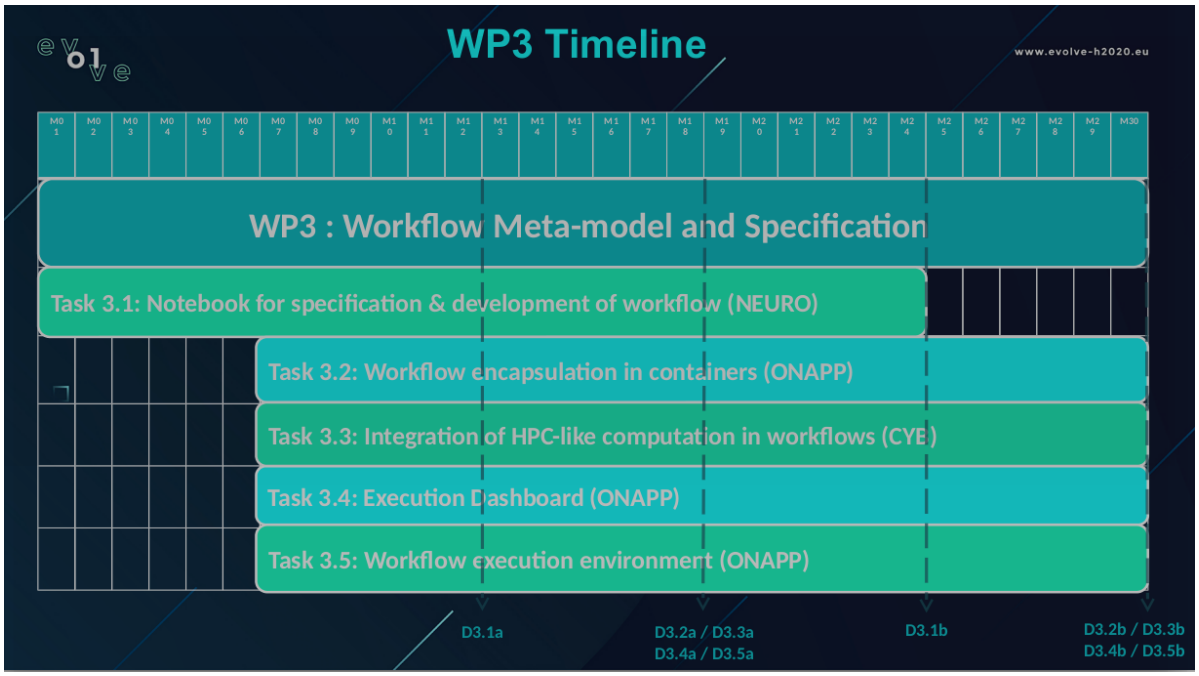

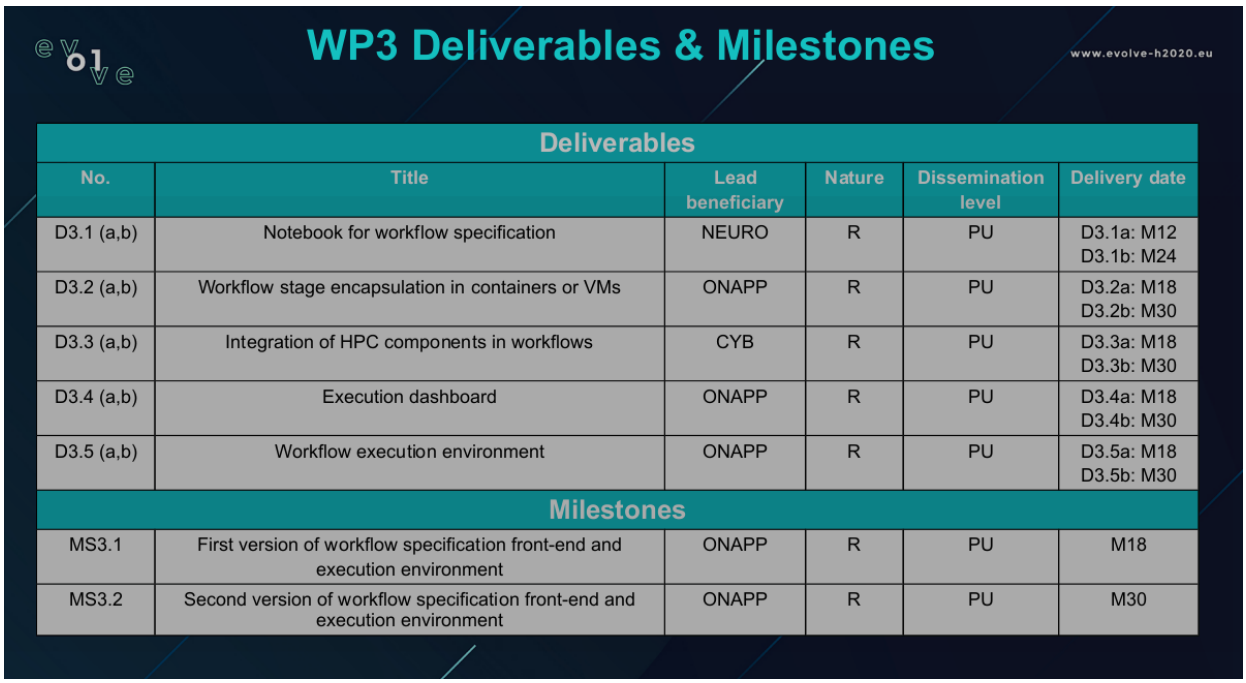

Deliverables					
No.	**Title**	**Lead beneficiary**	**Nature**	**Dissemination level**	**Delivery date**
D3.1 (a,b)	Notebook for workflow specification	NEURO	R	PU	D3.1a: M12 D3.1b: M24
D3.2 (a,b)	Workflow stage encapsulation in containers or VMs	ONAPP	R	PU	D3.2a: M18 D3.2b: M30
D3.3 (a,b)	Integration of HPC components in workflows	CYB	R	PU	D3.3a: M18 D3.3b: M30
D3.4 (a,b)	Execution dashboard	ONAPP	R	PU	D3.4a: M18 D3.4b: M30
D3.5 (a,b)	Workflow execution environment	ONAPP	R	PU	D3.5a: M18 D3.5b: M30
Milestones					
MS3.1	First version of workflow specification front-end and execution environment	ONAPP	R	PU	M18
MS3.2	Second version of workflow specification front-end and execution environment	ONAPP	R	PU	M30

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

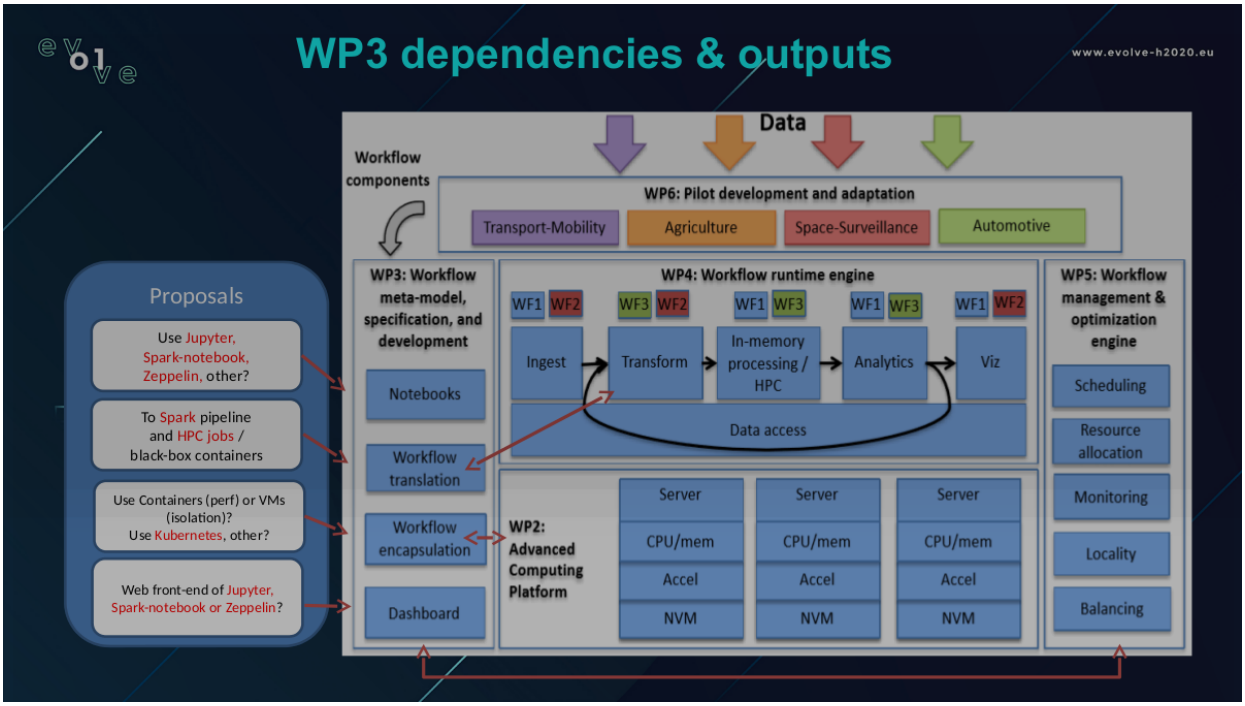

WP3: Discussion issues

www.evolve-h2020.eu

- Which computational notebook is more prominent for use within the EVOLVE project? (Jupyter – Zeppelin – Spark_Notebook)
- Workflow encapsulation in containers, VMs
  - Workflow integration with Spark? Define templates, image store and encapsulation
- HPC-like processing (black-box containers / MPI, etc. integration with notebooks?)
  - Need input from applications / pilots on the type & size of data, type and stages of processing, libraries used, etc.
- Execution Dashboard (based on web front-end of notebook?)
- Workflow execution environment (Use Kubernetes, Mesos or Ansible? VMs?)

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

WP3: Implementation timeline

www.evolve-h2020.eu

Month	Implementation
M06	Version 0 – First implementation of testbed and mock-up stack
M12	First integrated SW&HW testbed prototype
M18	First version of workflow specification front-end and execution environment (MS:3.1)
M18	Intermediate iteration (Prototype implementation with features that allow running full applications)
M24	Second integrated version of the testbed (All project pilot workflows running)
M30	Second and final version of workflow specification front end and execution environment (MS:3.2)
M30	Full prototype ready and technology evaluation initiation
M36	Final results of full deployment and evaluation (results from final workflows and large datasets)

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

3. Evolve Workflow

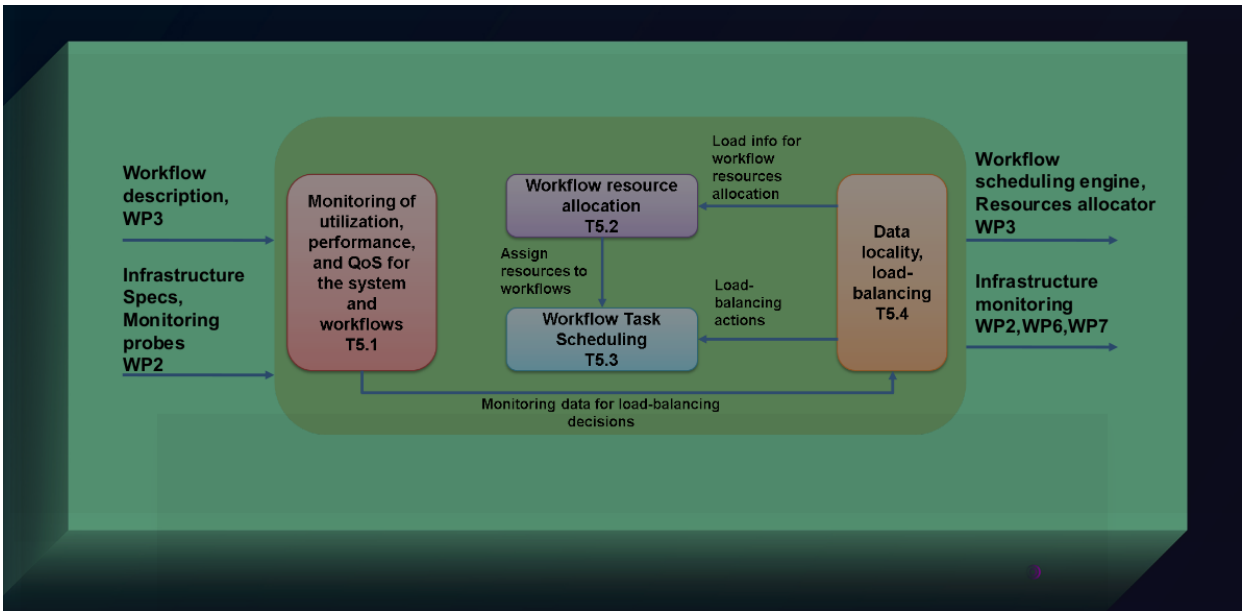

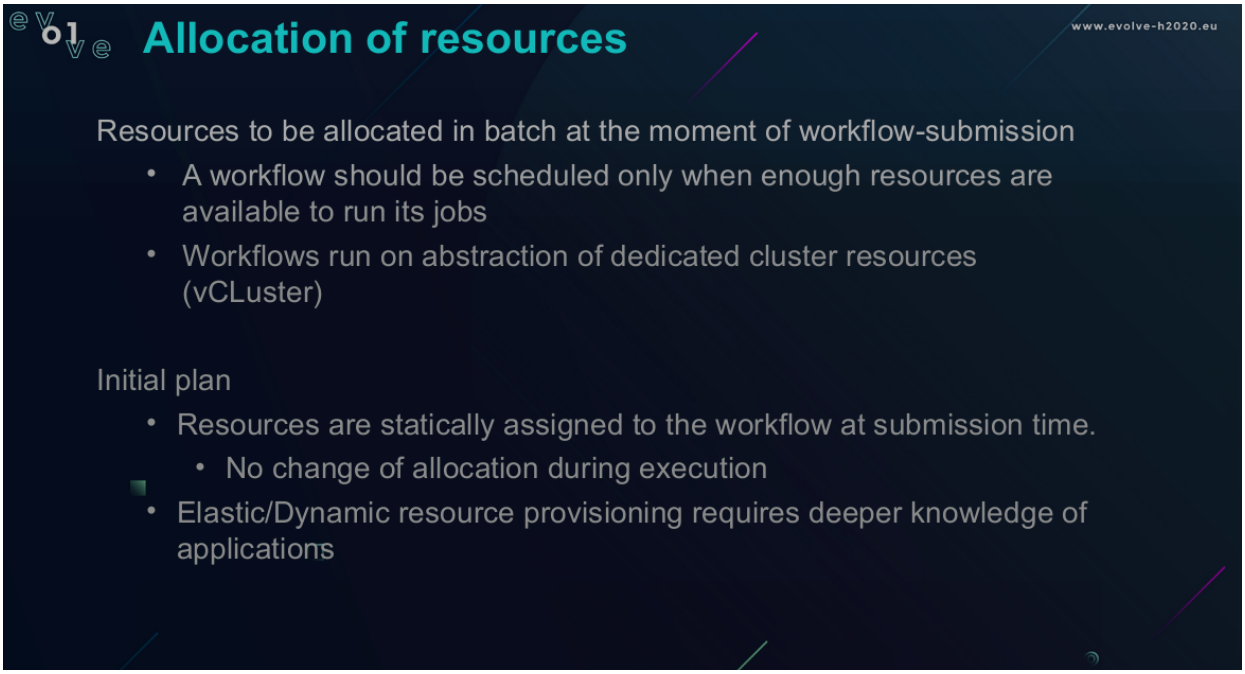

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

Scheduling of tasks and Load Balancing

www.evolve-h2020.eu

- Tasks are submitted to the scheduler by workflow orchestrator
- Scheduler has information on the workflow and takes informed decisions
  - Colocation of jobs (e.g., frequent interactions)
  - Deployment of jobs on specific infrastructure nodes for hardware needs (e.g., network requirements, connection to storage, cached data present on site)
  - Optimize usage of vCLuster resources (e.g., minimize interference between applications)
  - Create call-back mechanisms to monitor intermediate or exit job conditions
- Use monitoring information to perform load balancing
  - Adapt to applications dynamic behaviour
  - Avoid hot-spots

Monitoring of resources

www.evolve-h2020.eu

- Enable monitoring to guarantee efficient usage of hardware resources and final user QoS
- Monitoring ad system level and application level
- Monitoring data to be collected on per-workflow basis to guarantee privacy
- How do we monitor custom hardware? (e.g., FPGAs)
- What are the probes available in the hardware infrastructure?
- Reuse existing monitoring frameworks:
  - Prometheus: https://prometheus.io/
  - Telegraf: https://www.influxdata.com/time-series-platform/telegraf/

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

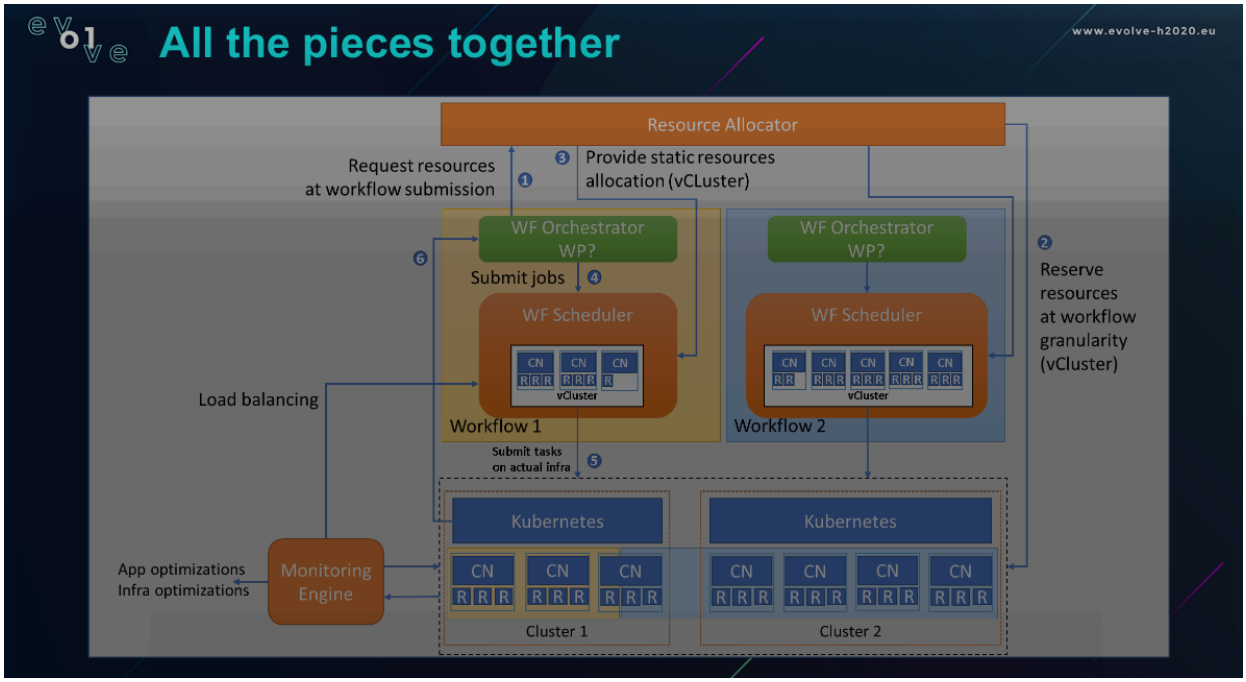

3.3. Deliverable for Work-flow

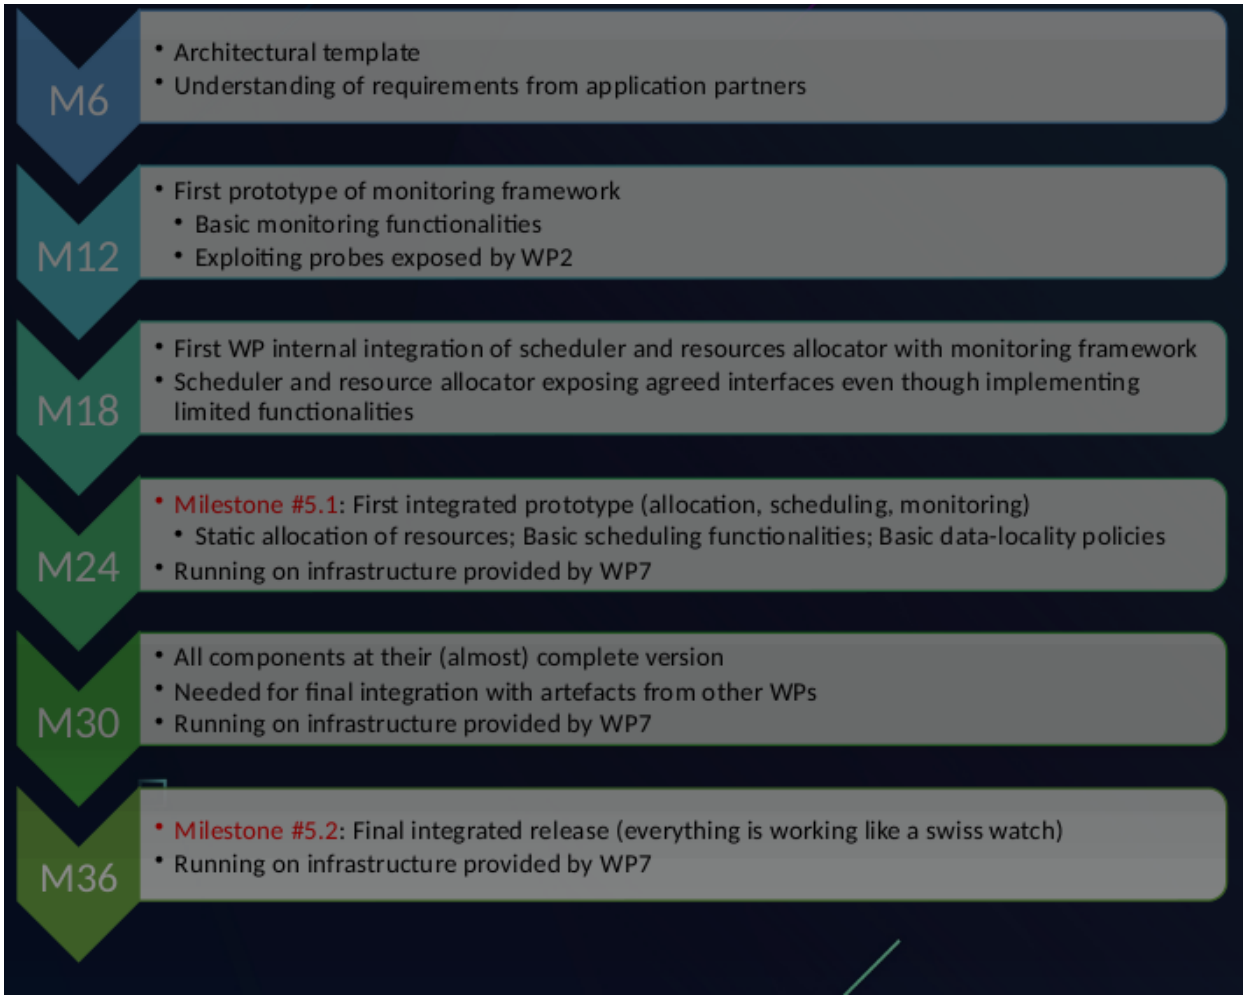

4. Pilots and use cases

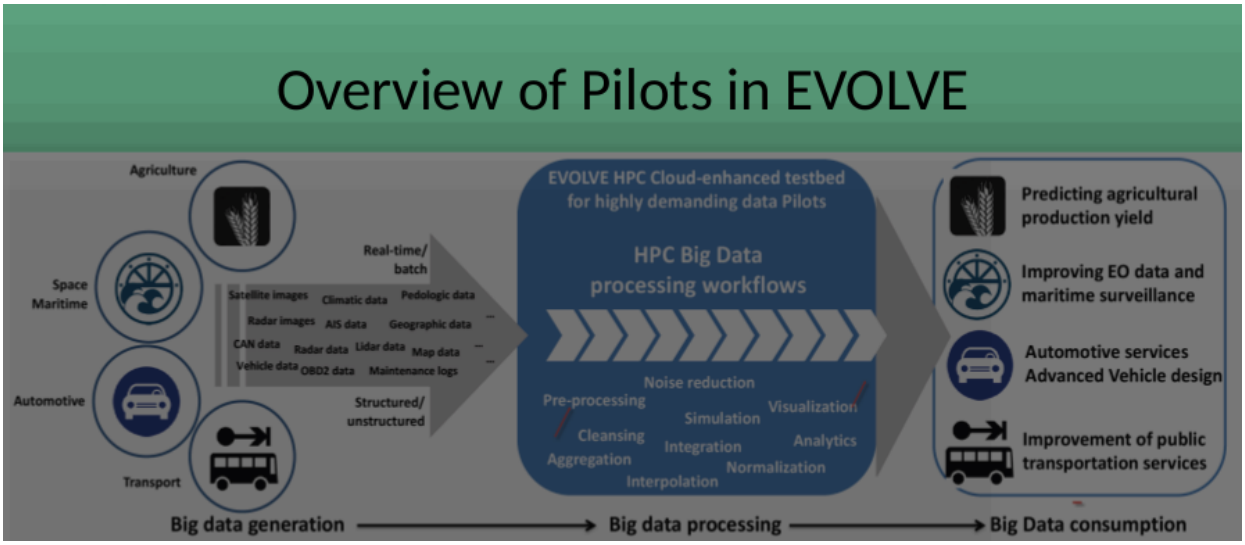

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

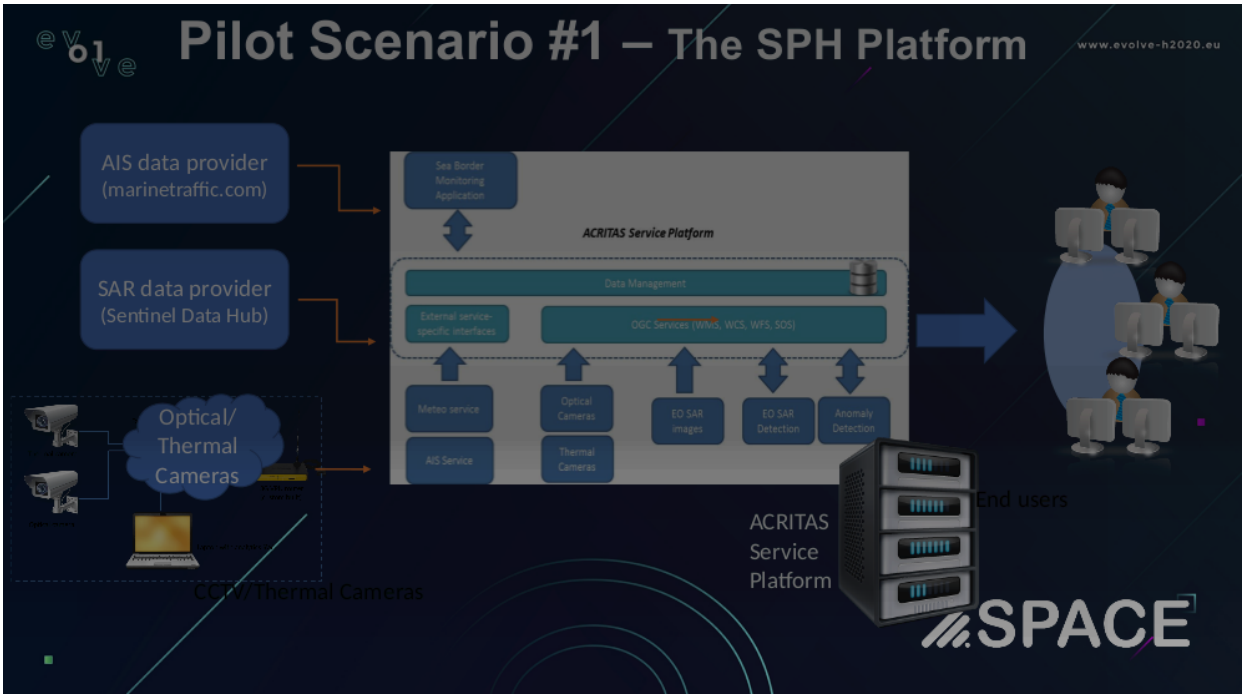

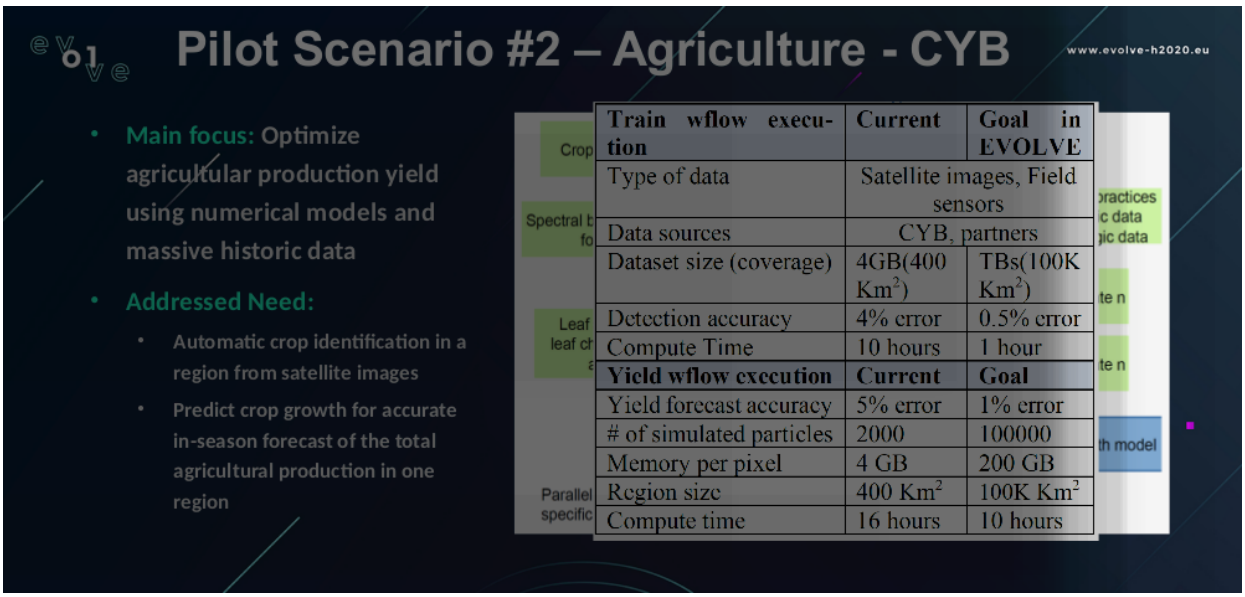

Train wflow execution	Current	Goal in EVOLVE
Type of data	Satellite images, Field sensors	
Data sources	CYB, partners	
Dataset size (coverage)	4GB(400 Km^2)	TBs(100K Km^2)
Detection accuracy	4% error	0.5% error
Compute Time	10 hours	1 hour
Yield wflow execution	**Current**	**Goal**
Yield forecast accuracy	5% error	1% error
# of simulated particles	2000	100000
Memory per pixel	4 GB	200 GB
Region size	400 Km^2	100K Km^2
Compute time	16 hours	10 hours

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

Pilot Scenario #3 – Satellite Images - TAS

www.evolve-h2020.eu

- **Main focus:** Radiometric correction and change detection on Sentinel-2 images
- **Addressed Need:**
  - Enhanced security, emergency, marine and land surveillance by examining changes of a designated area over a period of time sing these high-resolution images (10m per pixel)

Single execution	Current	Goal in EVOLVE
Type of data	Images and sensors	
Source of data	Private network, AIS	
Dataset size	10 GB	10 TB
Detection accuracy	~80%	>95%
Compute Time (1xSAR scene)	~20 mins	< 30 secs
Area coverage	1K Km^2	50K Km^2
Temporal window	1 month	10 years
Overall improvement	Improve time, spatio-temp window 10x-50x	

Pilot Scenario #4 – Public Transportation – NEURO, MEMEX, TIEMME

www.evolve-h2020.eu

- **Main focus:** Improvement of bus transportation services and city planning using observation and historic operational data
- **Addressed Need:**
  - Monitoring, controlling, validati bus services in a medium-size urban and rural areas
  - Assessment of the quality of bus services to improve service planning
  - Assessment of the effectiveness of activities related to service operation

AVM workflow	Current	Goal in EVOLVE
Type of data	10.000 trips; 160.000 transits; 900.000 bus events;	1.800.000 bus events;
Dataset size	• AVM[13] data f 1 GB/day x 30 days, and ~3-5x more than AVM from other sources (Traffic sensors, parking real-time availability, O/D matrix, sat maps, social media) • Not processed today automatically	Combine, process automatically
Source of data	TIEMME, Available earth observation and other services (for all workflows)	
Detection accuracy	10% detected trips as "suspected anomaly" or unreported	2%
Compute Time:	5 hours (manual procedures)	2 hours (auto+manual)
Overall improvement	5% of daily trips not operated	1%
	Lack of regularity of bus service from users' point of view	60% increase
	Poor attractiveness of PT (in particular small, medium cities)	2% passeng. increase
AFC workflow	**Current**	**Goal in EVOLVE**
Type of data	500 validation transactions daily with 50.000 transits through bus stops and 100.000 passengers	2.500 validation transactions
Dataset size	AFC[13] data: 300 MBytes/day x 30 days, not possible to combine today with data coming from other sources	Combine, process automatically
Detection accuracy	12% missed transaction/validation data	2%
Compute Time	1 hour to update population O/D matrix and verify coherence with scheduled service, transport network, selling volume	20 mins
Overall improvement	Not optimized selling network, poor attractiveness of PT	5% increase of tickets
APC workflow	**Current**	**Goal in EVOLVE**
Type of data (detection coverage)	100.000 passengers to join with 50.000 transits through bus stops and 500 validation transactions	200.000 passengers, 2.500 transactions
Dataset size	APC[13] data: not combined today with data coming from other ITS/data sources, 250 GBytes/day x 90 days	Combine, process automatically
Detection accuracy	6% "errors" in the total # of passengers counted at bus stops	4%
Compute Time	30 mins to estimate ticket pay evasion	10 mins
Overall improv.	10% of users having not paid tickets	8%

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

Pilot Scenario #5 – Mobility Services – BMW

- **Main focus:** Advanced vehicle routing algorithms and mobility services optimization
- **Addressed Need:**
  - Improve the efficiency of on demand mobility services, shared car services, parking
  - Increase fleet efficiency by improving the rate of capacity utilization
  - Optimized revenue management,
  - Coordination of demand and supply for transportation service,
  - Improved estimated time to pick-up and geo-spatial take over optimization of shared vehicles

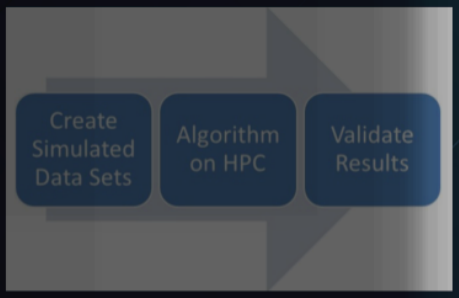

Pilot Scenario #6 – Automotive Services – KOOLA - VIF

- **Main focus:** Automotive services for predictive vehicle maintenance optimization
- **Addressed Need:**
  - Customer driven always up-to-date predictions of vehicle heath
  - High-quality predictive maintenance services

Single workflow execution	Current	Goal in EVOLVE
Type of data	Vehicles, workshops, observation	
Source of data	VIF/KOOLA, existing EO service	
Dataset size	5 GB	50 GB
Detection accuracy	~30%	50%
Response time	hours	minutes
# Vehicles	5000	50000
Overall improvement	Increase accuracy by ~2x, #vehicles by 10x, decrease response by 10-100x	

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

Pilot Scenario #7 – Vehicle Design – AVL VIF

- **Main focus:** Data-assisted automotive service development
- **Addressed Need:**
  - Testing of autonomous driving functionalities with novel scenarios captured by a fleet of vehicles operated in daily traffic. Examples for typical well-known driving scenarios are "Free Driving", "Long Stop", "Following Driving", or "Overtaking"

Single workflow execution	Current	Goal in EVOLVE
Type of data	Vehicles, workshops, observation	
Source of data	VIF/KOOLA, existing EO service	
Dataset size	5 GB	50 GB
Detection accuracy	~30%	50%
Response time	hours	minutes
# Vehicles	5000	50000
Overall improvement	Increase accuracy by ~2x, #vehicles by 10x, decrease response by 10-100x	

WP6 Deliverables and Milestones

List of Deliverables

- D6.1 Specification of Pilot Workflows – **M6**
- D6.2 First Implementation Workflows – **M18**
- D6.3 Second / Optimized Implementation of Pilot Workflows – **M30**

Milestones

- MS 6.1 First Implementation Workflows - **M18**
- MS 6.2 Second / Optimized Implementation of Pilot Workflows – **M30**

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

5. Continuous testbed integration

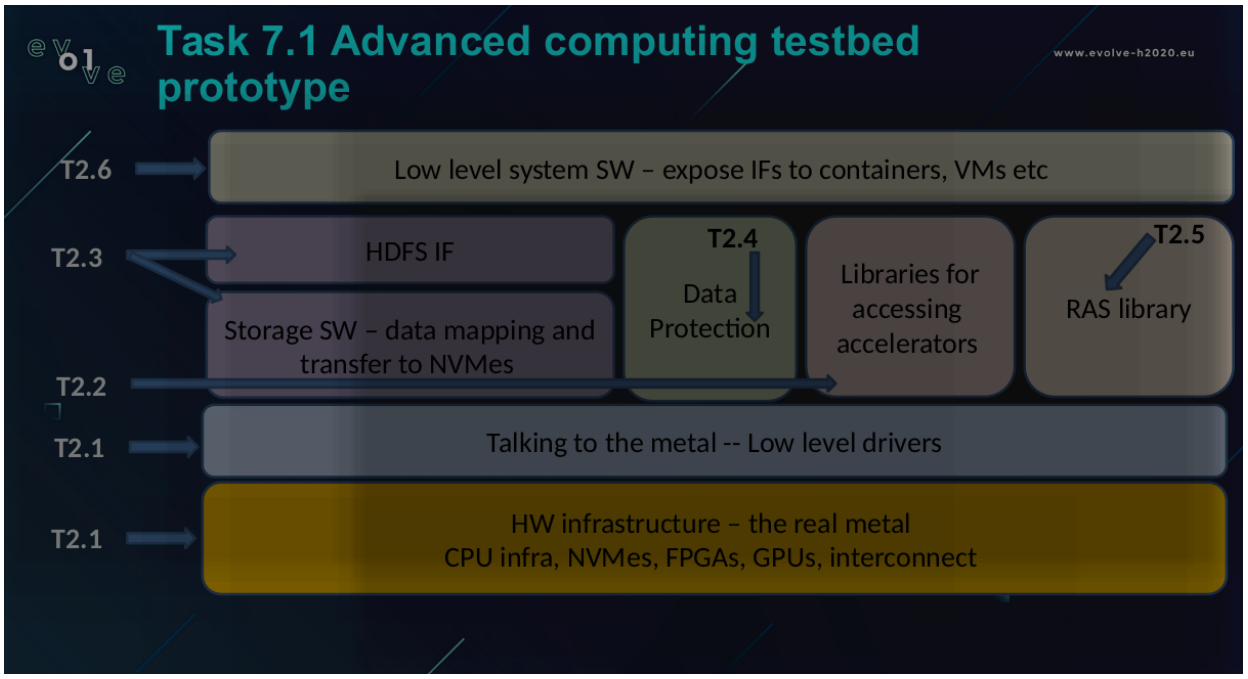

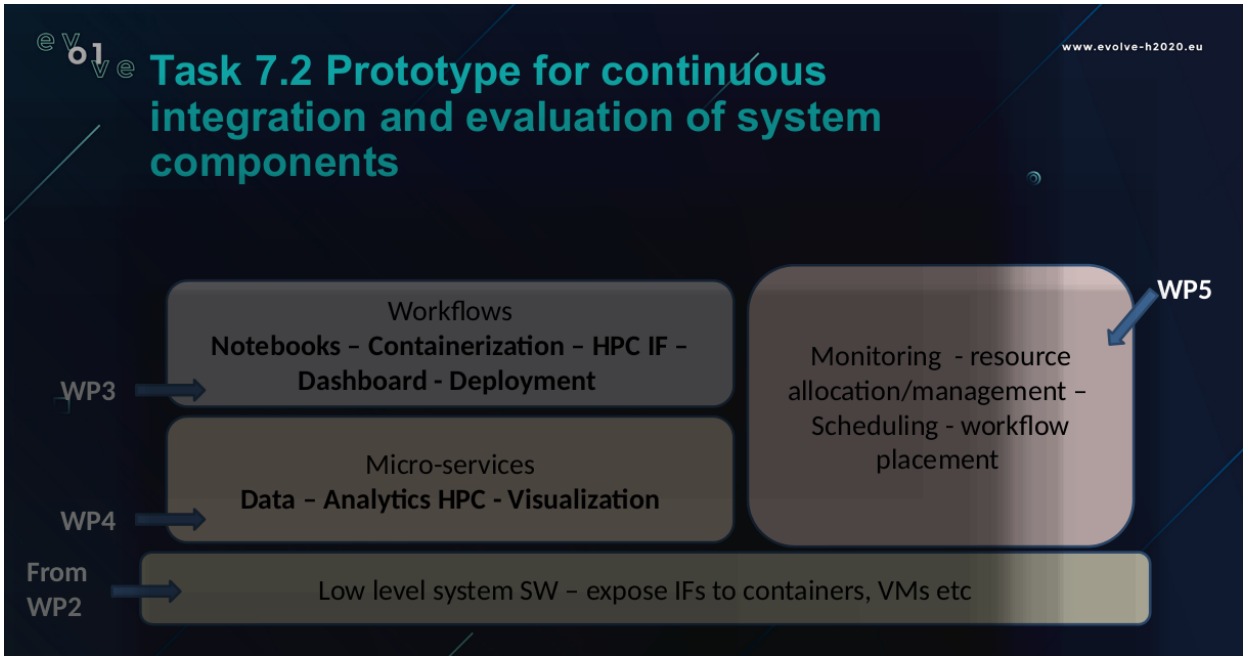

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

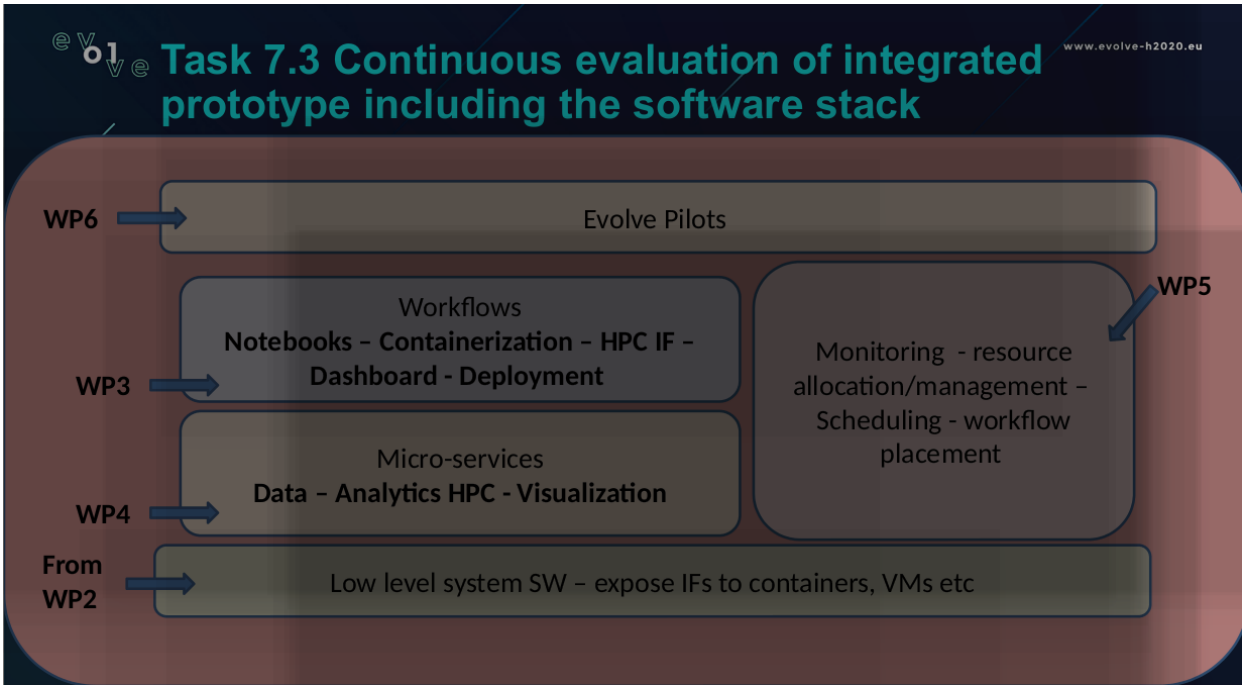

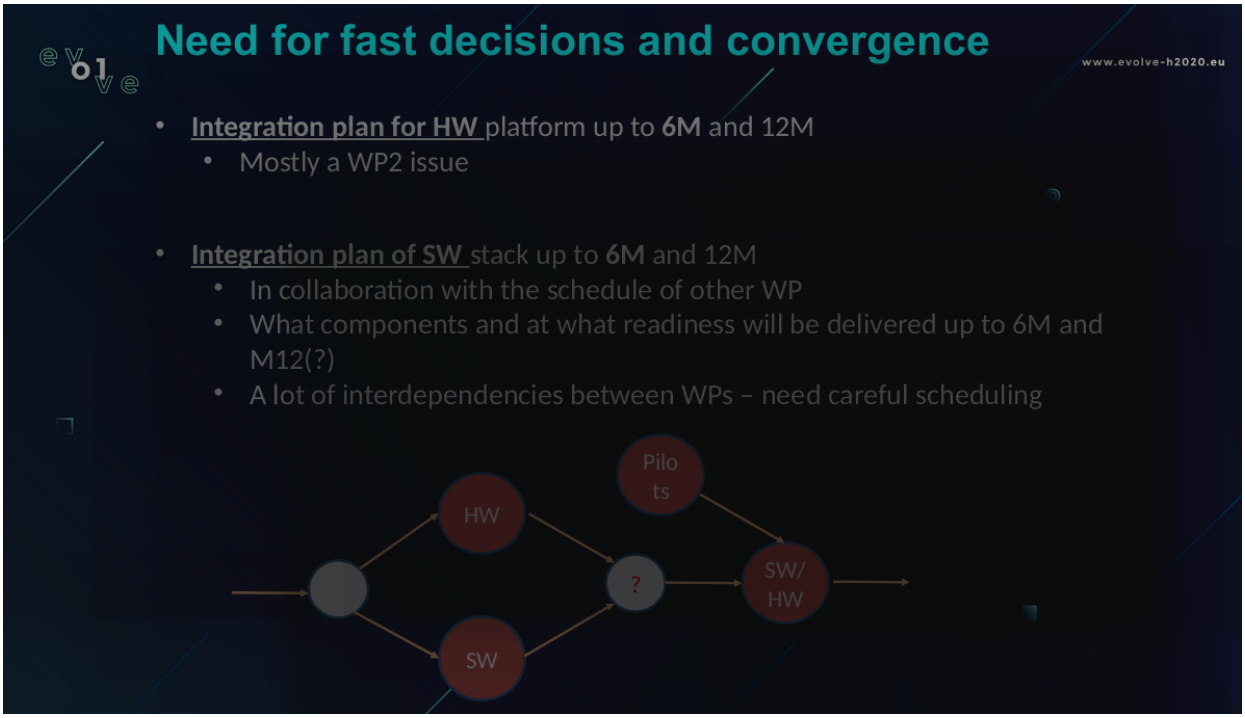

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

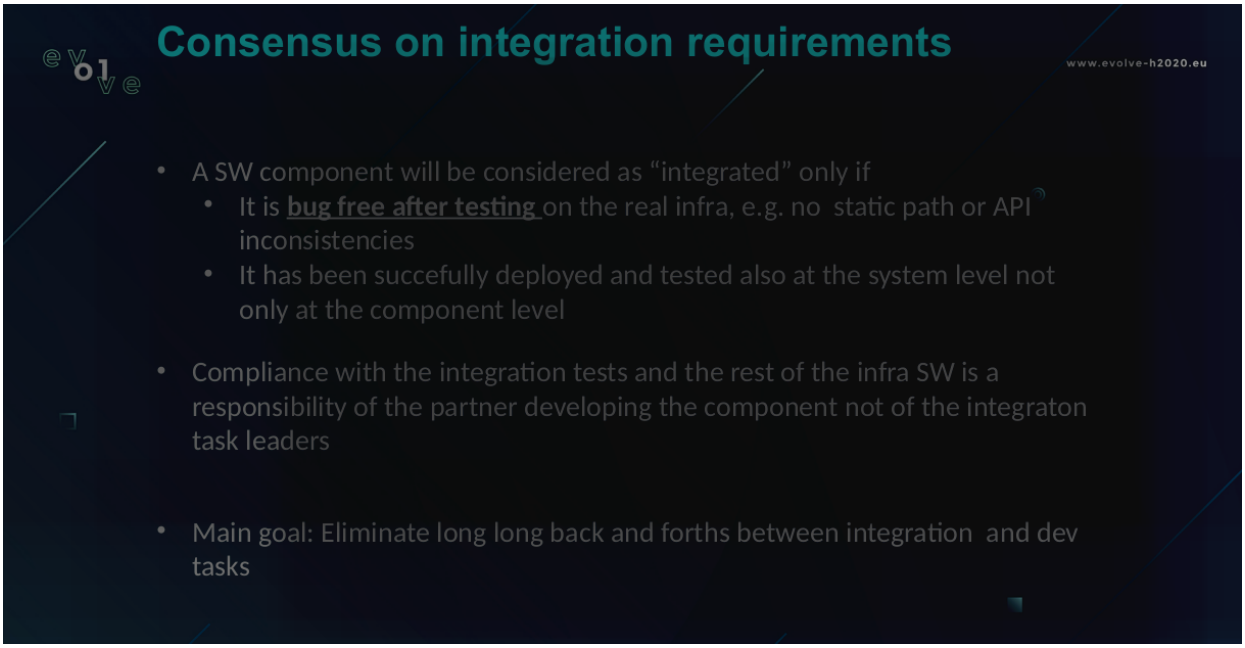

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061

Consortium as seen from organizations

Consortium as seen from Kick-Off meeting

This project has received funding from the European Union's Horizon H2020 research and innovation programme under grant agreement No 825061