FP DAM · CFGS · 1r curs · FP Illes Balears · 1665



Digitalització Aplicada als Sectors Productius

Contents

Unit 1 · Digitalisation and IT/OT **1**
1.1 What is digitalisation 4
1.2 Industry 4.0 and IT systems 10
1.3 Plant, business and IT–OT 16
1.4 Deliverable 1 · IT/OT X-ray 24

Unit 2 · Enabling technologies **28**
2.1 Enabling technologies I 30
2.2 Enabling technologies II 37
2.3 Products, services and markets 44
2.4 Sustainability, business and plant 50
2.5 Writing the technology report 55
2.6 Deliverable 2 · Technology report 61

Unit 3 · Cloud systems **65**
3.1 Cloud and service models 68
3.2 Deployment and functions 75
3.3 Edge, fog and mist 80
3.4 Advantages and sovereignty 86
3.5 Deliverable 3 · Cloud architecture 92

Unit 4 · Artificial intelligence **97**
4.1 What is AI 100
4.2 Generative AI and responsible use 106
4.3 Automation and profitability 112
4.4 Sectors, languages and software 118
4.5 Deliverable 4 · AI plan 124

Unit 5 · Data and cybersecurity **129**
5.1 Data, information and lifecycle 131
5.2 Big data 137
5.3 Data science and the cloud 143
5.4 Cybersecurity and regulation 149
5.5 Deliverable 5 · Data and security brief 156

Unit 6 · Digital transformation project **160**
6.1 Goals and areas 162
6.2 Fit and technologies 168
6.3 Security, data and integration 173
6.4 People and documentation 178
6.5 Final presentation · D6 Transformation plan 183

The project **188**

UNIT 1



01

Unit 1 · Digitalisation and IT/OT

Picture a city-builder game such as *Cities: Skylines*. Your city has two halves. One half physically **works**: the water plant, the power station, the traffic lights, the factories. The other half **keeps the records**: the city hall with its registers, taxes, permits and the service desk for citizens. A real company is the same kind of city. In this unit you learn to see both halves, the office side called **IT** and the operational side called **OT**, and to understand why the most interesting work happens in the pipes that connect them.

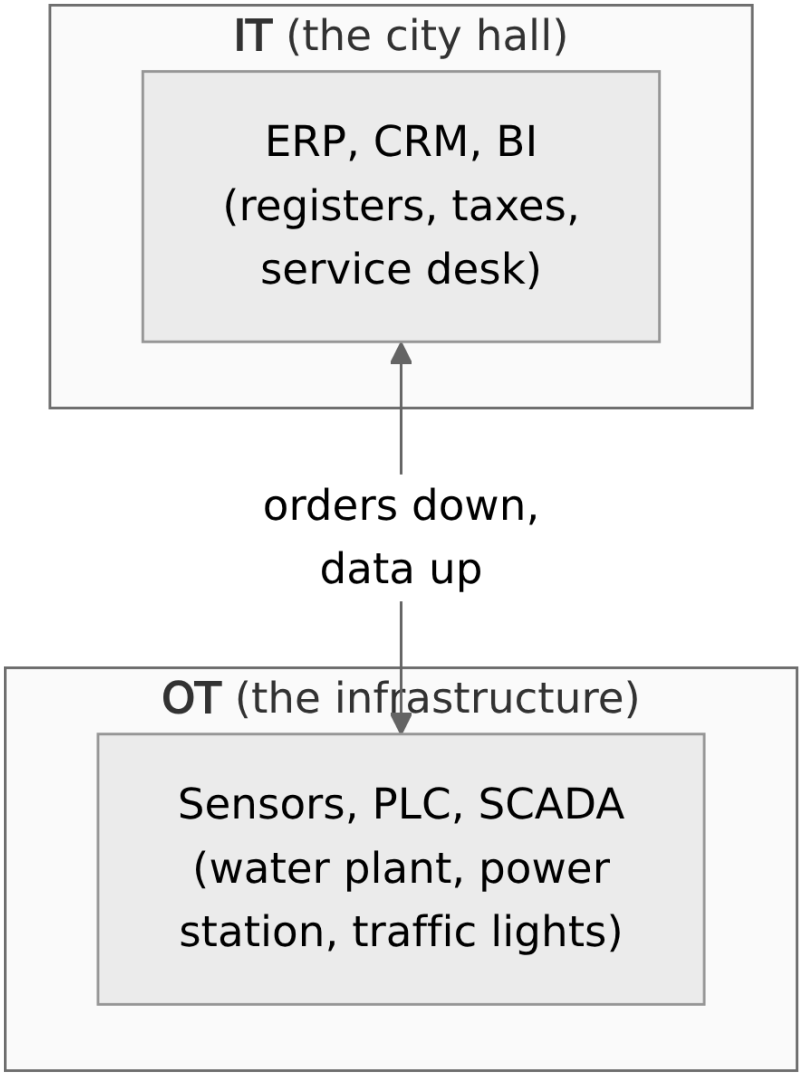
What you will learn

- Tell **digitisation**, **digitalisation** and **digital transformation** apart, and avoid the trap of thinking that digitalising means buying computers
- Explain how digital technology changes the way a company is **organised**
- Name the typical **departments** of a company and the IT systems each one uses (**ERP**, **CRM** and others)
- Compare **IT** and **OT**: their differences and, just as important, their similarities
- Place the plant technologies (sensors, PLC, SCADA, MES) on the **automation pyramid**, and choose the right plant and business technologies for a real company
- Explain why connecting IT and OT matters, what it costs, and what an **end-to-end** digital company gains

The arrow in the middle is the whole unit. Lessons 1 and 2 are about the two boxes; lesson 3 is about the arrow.

Lessons

#	Lesson	What it covers
1	What is digitalisation	Digitisation, digitalisation and digital transformation; how they change a company; IT and OT for the first time
2	Industry 4.0 and IT systems	The four industrial revolutions; departments, ERP and CRM; IT and OT in depth
3	Plant, business and IT–OT	The automation pyramid with a real plant in Mallorca, business technologies, connecting IT and OT, end-to-end digitalisation
4	Deliverable 1 · IT/OT X-ray	What you hand in, who writes what, how it is assessed, the defence


IT (the city hall)
ERP, CRM, BI (registers, taxes, service desk)
orders down, data up
OT (the infrastructure)
Sensors, PLC, SCADA (water plant, power station, traffic lights)

1.1

What is digitalisation

Think about this morning. The alarm on your phone woke you up, you checked the bus on an app, you paid a friend back with Bizum and you listened to music on the way. How many of those things would your parents have done on paper, face to face, or not at all, when they were your age?

That change has a name, and it is the subject of this whole module. It is not only about you: in 2025, people in Spain sent about **1,237 million** Bizum payments, money that twenty years ago moved in cash or at a bank counter. Every company you will write software for is going through the same change, and many of them will hire you precisely to push it forward.

What you will learn

- Tell **digitisation**, **digitalisation** and **digital transformation** apart with one example of each
- Avoid the trap: digitalising is not buying computers, it is changing **how the work is done**
- Explain how digitalising a process changes the **organisation** of a company: tasks, roles and departments
- Recognise the two faces of a company, **IT** and **OT**, and give an example of each
- See your company as a **city** you run, with a city hall and an infrastructure

1.1.1 Three words that people mix up

The three words sound alike, and even newspapers confuse them. They are three different steps.

> **Digitisation** is turning analogue information into a digital format. The medium changes; the process does not.
>
> **Digitalisation** is using digital technology to improve or automate a process that already exists.
>
> **Digital transformation** is a deep change in the business model, the organisation or the relationship with the customer, made possible by digital technology.

Follow one hotel in Calvià through the three steps. It is a made-up hotel, but every Balearic hotel has walked this road.

Step	What the hotel does	What changed
Digitisation	The paper guest cards are scanned and saved as PDF files	Only the medium: paper became a file

Step	What the hotel does	What changed
Digitalisation	Guests check in online and open their room with the phone	The process: no queue at reception, no plastic key
Digital transformation	Prices change every day according to demand, and most rooms are sold directly on the hotel's website instead of through travel agencies	The business: how the hotel earns its money

Now your morning:

Service	Which step?	Why
Photos of your class notes	Digitisation	The notes are the same; only the medium changed
Bus ticket on the phone	Digitalisation	Same trip, better process: no coins, no paper ticket
Bizum and mobile banking	Digitalisation, and a transformation of the sector	Paying is faster, and the relationship with the bank moved from the branch to the phone
Spotify or Netflix	Digital transformation	The business changed from selling a physical object (a CD, a DVD) to selling access

Careful: "digitalising means buying computers"

A company can fill its offices with new laptops and keep working exactly as before: printing emails, copying numbers from one program to another, filing paper. That is not digitalisation. Digitalising changes **the process**. The technology is the tool, not the goal.

1.1.2 Digitalising changes the organisation

When the hotel moves check-in online, the change does not stop at reception. Ask yourself: what happens inside the company?

Before	After
Several receptionists typing guest data at the desk	Fewer people on repetitive tasks; more people on customer care
The hotel learns who the guest is when they arrive	The guest's data arrive days earlier: the room can be prepared, services can be offered
Nobody in charge of the systems	New roles: systems administrator, data analyst, cybersecurity

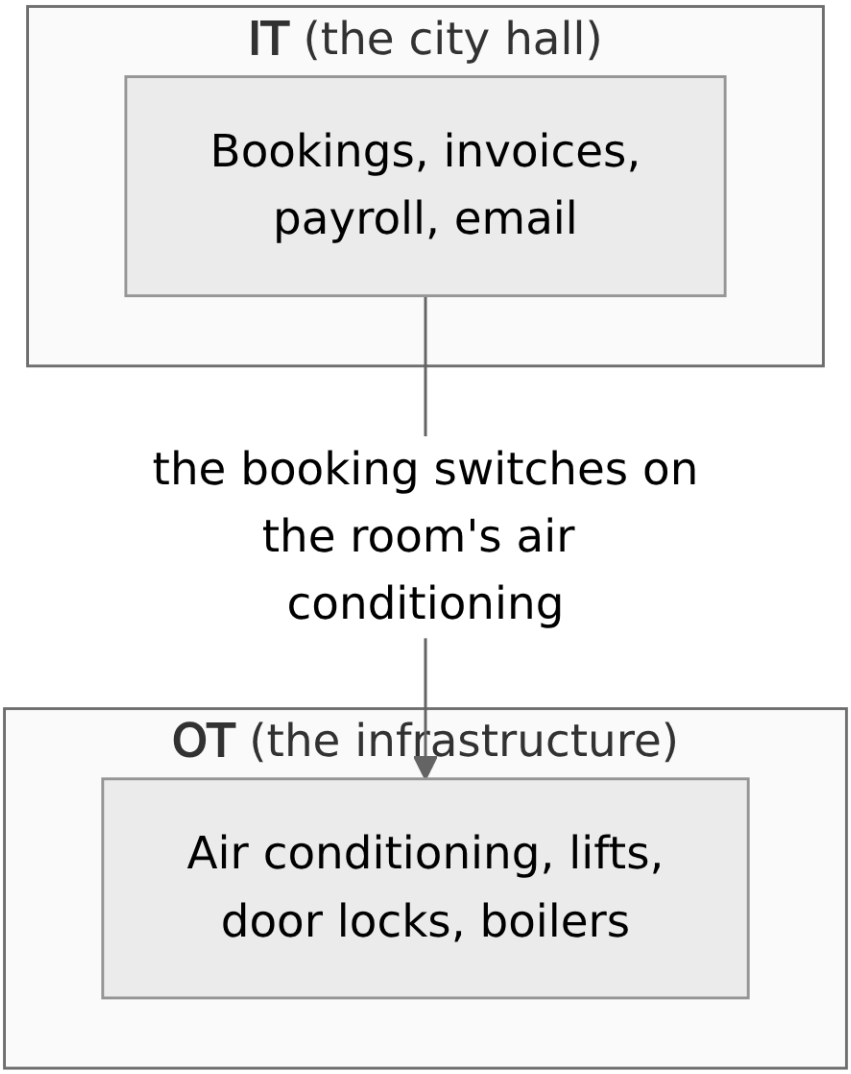
Technology moves tasks from one person to another, creates new jobs and makes others disappear. That is why the programme of this module asks you to **relate** digital technology to **the organisation of the company**: you cannot digitalise a process without reorganising the people around it.

1.1.3 IT and OT: the two faces of a company

Back to the city-builder game. A company is a city you run, and every city has two halves.

IT (*information technology*) is the technology that manages information and business data: the office side.

OT (*operational technology*) is the technology that monitors and controls physical processes, machines and installations: the operational side.



	IT · Information technology	OT · Operational technology
Manages	Information and business data	Physical processes, machines and installations
Where	Offices: administration, accounting, HR, sales	Plant, factory, installations, the field
Examples	ERP, CRM, email, website, booking system	PLC, SCADA, sensors, robots, air conditioning, access control
Priority	Confidentiality and integrity of the data	Safety of people and availability: it cannot stop
Lifespan	3–5 years, frequent updates	10–15 years or more, very controlled changes

Look at the arrow in the diagram. When the hotel's booking system (IT) switches on the room's air conditioning (OT) before the guest arrives, the two halves of the city are talking to each other. This connection is the heart of the unit, and you will study it properly in Plant, business and IT–OT.

> **Does every company have OT?**

> Almost every company has IT. Not every company has much OT. A company that sells hotel rooms online to other businesses, such as Hotelbeds in Palma, is almost all IT: its "factory" is software. A water company like EMAYA, an airport or a cheese factory has a lot of OT. For your course project you need a company with **both** halves.

1.1.4 In class: my digital life and the two faces

The session had two short activities.

Step	Do this	You should see
1	Write down three digital services you used today between waking up and arriving in class	A list that mixes apps, payments and media
2	Classify each one: digitisation, digitalisation or digital transformation	Most are digitalisation; streaming and banking are transformations
3	In groups, take one real Balearic company (Meliá, Palma airport, the Port of Palma, EMAYA, EMT Palma or Hotelbeds) and write **two IT systems** and **two OT systems** you think it has	A short table; if the company has no OT, a sentence saying why
4	Propose **one connection** between an IT system and an OT system that would add value	For example: the bus app (IT) shows the position sent by the bus GPS (OT)

These were **hypotheses**: you guessed what systems each company has. From now on, in the project, you will have to **find** them and **cite** the source.

1.1.5 For your project

Your group now works on one real company for the whole year (see the course project). This lesson feeds Deliverable 1:

- The **group introduction** says what digitalising means for *your* company, with the three words used correctly.
- **Section A** describes how digital technology is organised in the company: departments, roles, systems.
- The **group comparison** of IT and OT starts from the table on this page, filled with your company's real systems.

1.1.6 Key takeaways

- **Digitisation** changes the medium, **digitalisation** changes the process, **digital transformation** changes the business.
- Buying computers is not digitalising: what matters is **how the work is done**.
- Digitalising a process changes the **organisation**: tasks move, new roles appear, others disappear.
- **IT** manages information (the city hall); **OT** controls physical processes (the infrastructure).
- IT protects the **data**; OT protects **people and availability**, and lives much longer.
- The most valuable work happens where IT and OT **connect**.

1.1.7 Quick check

1 A dental clinic scans its old paper patient files and stores them as PDFs. Nothing else changes. What is it?

- **A** Digitisation
- **B** Digitalisation
- **C** Digital transformation
- **D** Industry 4.0

2 Which example is a digital transformation?

- **A** A shop starts accepting card payments at the till
- **B** An office scans its contracts
- **C** A software company stops selling licences and sells a monthly cloud subscription instead
- **D** A school sends the timetable by email instead of printing it

3 A company buys 50 new laptops, but everybody keeps printing the orders and typing them again into the accounting program. Has it digitalised?

- **A** Yes: it has more digital technology than before
- **B** No: the process is the same; digitalising means changing how the work is done
- **C** Yes, as long as the laptops are connected to the internet
- **D** No, because digitalising requires a factory

4 Which of these is an OT system in a hotel?

- **A** The booking system
- **B** The payroll program
- **C** The control of the air conditioning in the rooms
- **D** The hotel's website

5 What is usually the first priority in an OT environment?

- **A** Keeping the data confidential
- **B** The safety of people and keeping the process running
- **C** Installing updates as soon as they appear
- **D** Replacing the equipment every three years

6 A hotel moves its check-in online. Which TWO changes in the organisation are likely?

- **A** Fewer people doing repetitive data entry at reception
- **B** The hotel no longer needs any staff
- **C** New roles appear, such as someone in charge of the systems or the data
- **D** The accounting department disappears

7 In the city-builder analogy, which part of the city is OT?

- **A** The city hall with its registers and taxes
- **B** The citizens' service desk
- **C** The water plant, the power station and the traffic lights
- **D** The city's budget

Now you know the two halves of the city. Next, where all this comes from, who uses IT inside a company, and how similar IT and OT really are: Industry 4.0 and IT systems.

Solutions

1. A · The information moved from paper to a file, but the way the clinic works is the same. Only the **medium** changed, so it is **digitisation**.

2. C · Selling a subscription instead of a product changes **how the company earns money**: the business model. That is **digital transformation**. The other three improve a process or change a medium.

3. B · **Digitalisation** improves or automates a **process**. New hardware with the same paper-and-retyping process changes nothing. The technology is the tool, not the goal.

4. C · **OT** controls physical processes and installations. The air conditioning heats and cools real rooms; the other three manage information, so they are **IT**.

5. B · In **OT**, if the system stops, the plant stops, and a failure can hurt people. That is why **safety and availability** come first, and why OT equipment is changed rarely and carefully.

6. A, C · Digitalising **moves tasks**: repetitive work shrinks and new technical and customer-care roles appear. It does not remove people or whole departments by magic.

7. C · **OT** is the infrastructure that physically works. The city hall, the service desk and the budget handle information and records: that is the **IT** half.

1.2

Industry 4.0 and IT systems

You book a table at a restaurant on your phone. The waiter's tablet already shows your name when you arrive, the kitchen sees the order the moment it is taken, and at the end of the month the accountant does not type a single ticket. One piece of data, typed once, travels through the whole company. How did companies get here, and which programs make it possible?

This lesson answers both questions. First, where the idea comes from: a story of four industrial revolutions that ends in a word you will read everywhere, **Industry 4.0**. Then, who uses IT inside a company, and with which systems. Those systems are exactly the programs a DAM technician builds, integrates and maintains.

What you will learn

- Place **Industry 4.0** at the end of the four industrial revolutions and say what makes it different
- Explain what a **cyber-physical system** is with an everyday example
- Name the typical **departments** of a company and the IT system each one uses
- Say what an **ERP**, a **CRM** and a **PMS** do, and why typing a piece of data only once matters
- Give **differences and similarities** between IT and OT, and avoid the trap of listing only differences

1.2.1 Four industrial revolutions

Revolution	When	What made it possible	What changed
1st	Late 18th century	Steam engine, mechanisation	From handwork to the factory
2nd	Late 19th – early 20th century	Electricity, the assembly line	Mass production
3rd	Second half of the 20th century	Electronics, computers, PLCs and robots	Automated production
4th	Today	Cyber-physical systems, the internet of things, cloud, data and AI	Machines, systems and people **connected**; decisions based on data

The sequence to remember is **steam → electricity → automation → connection and data**. You do not need the exact dates of the first three.

> **Industry 4.0** is the fourth industrial revolution: machines, products and systems connected to each other, generating data that the company uses to decide and to act.

The term was born in Germany. *Industrie 4.0* was presented in public at the **Hannover trade fair in 2011**, as a project of the German government's high-tech strategy, by **Henning Kagermann**, **Wolfgang Wahlster** and **Wolf-Dieter Lukas**.

The difference with the third revolution is not that there are computers. The third revolution already had them. The difference is that now **everything is connected and produces data** that someone uses.

> A **cyber-physical system** is a physical thing (a machine, a building, a vehicle) with sensors, software and a network connection, so that it can be monitored and controlled digitally.

In the city-builder game, the third revolution is a city where the traffic lights change on a timer. The fourth is a city where every traffic light reports the traffic it sees, and you watch and adjust them all from one screen.

> **Not only factories**
> The word says *industry*, but the idea applies everywhere: a hotel whose rooms report their temperature, a port whose cranes report their position, a water network whose pipes report their pressure. A hotel with magnetic keys and bookings in a spreadsheet lives in the third revolution. A hotel that opens the door with your phone and changes prices with demand lives in the fourth.

1.2.2 Who uses IT inside a company

Every company, from a hotel to a software start-up, is organised in **departments**. Each one has its job and its systems.

Department	What it does	Typical IT systems
Management	Strategy and decisions	Dashboards, business intelligence (BI)
Administration and finance	Accounting, invoices, taxes	ERP (finance module), electronic invoicing
Human resources	Contracts, payroll, shifts	Payroll and HR software, time tracking
Sales and marketing	Sales, campaigns, customers	CRM, website, online shop
Purchasing and logistics	Suppliers, stock, shipping	ERP (purchasing and warehouse), warehouse management system
Customer service	Questions, incidents, after-sales	CRM, ticketing system
IT (systems)	Network, servers, security, support	User management, backups, monitoring

Three families of systems you must be able to name:

> An **ERP** (*Enterprise Resource Planning*) brings the company's processes (finance, purchasing, stock, sales, HR) together in **one single database**. Examples: SAP, Microsoft Dynamics 365, Odoo.
>
> A **CRM** (*Customer Relationship Management*) manages the whole relationship with customers: contacts, sales opportunities, incidents, campaigns. Examples: Salesforce, HubSpot.
>
> A **PMS** (*Property Management System*) is the management system of a hotel: bookings, rooms, check-in and billing. Every sector has its own; a shop or a bar has its point-of-sale (POS) system.

Before the ERP, each department had its own program and its own spreadsheet, and the same customer was typed three times: by sales, by the warehouse and by accounting. Three copies means three chances to get it wrong. The point of an ERP is that a piece of data is **entered once** and everybody sees the same one. In the city, it is one register at the city hall instead of a different notebook in every office.

> **This is DAM work**
> Odoo has an open-source edition (Community), and its server is written mostly in Python, with a web interface in JavaScript. Customising an ERP, connecting it to an online shop or writing a report for it is everyday work for a DAM technician.

1.2.3 IT and OT in depth

In the first lesson you compared what IT and OT manage, where they live, their priority and their lifespan. Here are more differences:

	IT	OT
Response time	Waiting a few seconds is usually acceptable	Real time: the machine cannot wait
Updates	Frequent patches, regular restarts	Very few; production must be stopped and the update planned
Networks and protocols	Internet standards: TCP/IP, HTTP, email	Industrial buses and protocols, some proprietary: Modbus, Profinet, OPC UA
If it fails	Data, time or money are lost	Production stops; there may be physical damage or danger to people
Who runs it	The IT department	Maintenance and plant engineering

And now the part most people forget. The programme asks for differences **and similarities**:

- Both are **hardware + software + networks**, and both must be programmed and maintained.
- Both **produce and consume data** that the company needs.
- Both need **cybersecurity**, backups and access control.
- More and more, they share the same technology: Ethernet, TCP/IP, standard operating systems, virtualisation.

> **IT/OT convergence** is the trend of IT and OT using more and more of the same technology and being connected to each other.

> **Careful: "IT and OT have nothing in common"**
> It is the most common mistake. They have different **priorities** (data versus people and availability) but they are built from the same materials and face the same attackers. In the city, the city hall and the water plant are different buildings, but both have doors, keys and staff who can be tricked.

1.2.4 In class: the systems map

In pairs, each pair took one case:

- **Case A**: the hotel in Calvià, 150 rooms, a restaurant, open from March to November.

- **Case B**: a software company with 25 employees at ParcBit, making custom applications for other companies.

Step	Do this	You should see
1	Make a table with **five departments**, and for each one **one IT system** and **the data** it manages	A table like the one on this page, adapted to your case
2	Find **one piece of data that two departments need at the same time** (for example, a booking: reception and accounting)	The reason an ERP exists: if each department has its own copy, they disagree
3	As a developer: which **integration** between two systems would you program first, and why?	One sentence connecting two systems

Then, individually, the whole class played **IT, OT or both?** with sentences such as "if it stops for ten minutes we lose a whole batch" (OT), "we install the security updates every month" (IT) or "it sends data over an Ethernet network" (both), and each student wrote **two differences and two similarities** between IT and OT on Classroom.

1.2.5 For your project

For Deliverable 1:

- **Section A** lists the departments of *your* company and the real IT systems it uses. Find them and cite them: the company website, its annual report, the press, or its job offers (a job offer asking for "experience with SAP" is a source).
- The **group comparison** of IT and OT must include similarities, not only differences, with examples from your company.

1.2.6 Key takeaways

- Four revolutions: **steam → electricity → automation → connection and data**. **Industry 4.0** was presented in Hannover in 2011.
- A **cyber-physical system** is a physical thing with sensors, software and a connection.
- Each **department** uses its own IT systems; the **ERP** puts them in one database, the **CRM** manages customers, the **PMS** runs a hotel.
- Enter each piece of data **once**: that is the reason integrated systems exist.
- IT and OT differ in **priorities**, response time, updates and protocols, but share hardware, software, networks, data and risks: **IT/OT convergence**.

1.2.7 Quick check

1 What makes the fourth industrial revolution different from the third?

A Computers appear for the first time
B Electricity reaches the factories
C Machines, systems and people are connected and generate data that are used to decide
D Production moves from the workshop to the factory

2 Where and when was the term "Industrie 4.0" presented in public?

A In Silicon Valley, in 1995
B At the Hannover trade fair, in Germany, in 2011
C In Barcelona, at the Mobile World Congress of 2016
D In Japan, in 1980

3 Which system puts finance, purchasing, stock and invoicing in one single database?

- **A** CRM
- **B** ERP
- **C** PMS
- **D** SCADA

4 Which department uses a CRM the most?

- **A** Human resources
- **B** Accounting
- **C** Sales and customer service
- **D** Plant maintenance

5 Reception and accounting each keep their own list of bookings in different programs. What is the most likely problem?

- **A** None: two copies are safer than one
- **B** The two lists end up being different, and someone has to type the same data twice
- **C** The hotel needs more OT
- **D** The bookings cannot be paid by card

6 Which TWO of these are similarities between IT and OT?

- **A** Both need cybersecurity, backups and access control
- **B** Both can be stopped and restarted at any time to install updates
- **C** Both are made of hardware, software and networks
- **D** Both put the confidentiality of data first

7 A city's traffic lights now report the traffic they see and can be adjusted from a control room. Which concept is this?

- **A** Digitisation
- **B** An ERP
- **C** A cyber-physical system, typical of Industry 4.0
- **D** A CRM

You now know the offices of the city. Next, you go down to the plant, where the machines are, and see what happens when both halves are connected: Plant, business and IT–OT.

Solutions

1. C · Computers and PLCs already existed in the **third** revolution. **Industry 4.0** is about **connection and data**.

2. B · It was presented at the **Hannover** fair in **2011** as part of the German government's high-tech strategy.

3. B · The **ERP** integrates the company's processes in one database. The CRM is about customers, the PMS is a hotel's management system, and SCADA belongs to the plant (OT).

4. C · A **CRM** manages the relationship with **customers**: contacts, sales opportunities and incidents. That is the daily work of sales and customer service.

5. B · Two copies of the same data drift apart and double the typing. Entering each piece of data **once** in a shared system, such as an **ERP**, is the fix.

6. A, C · IT and OT share materials and risks: hardware, software, networks and the need for **cybersecurity**. They differ in priorities: OT cannot be stopped easily and puts **people and availability** first.

7. C · A physical thing (the traffic light) with sensors, software and a connection, monitored and controlled digitally, is a **cyber-physical system**: the core idea of **Industry 4.0**.

1.3

Plant, business and IT–OT

You open the tap and water comes out. You never think about it. But in Palma, behind that tap there are more than 1,100 km of pipes, split into 600 small sectors, and a control centre where people watch pressures and flows on screens all day. In 2025 EMAYA, Palma's water company, said it had cut the water lost in leaks from 25% to 16%, after renewing pipes and installing digital meters. How do you find a leak under a street from a screen?

You find it because the **plant** (the pipes, pumps and meters) is talking to the **business** (the people who plan, pay and decide). This lesson goes down to the plant, names the technologies you find there, and then looks at the most valuable thing in a digital company: the connection between both halves.

What you will learn

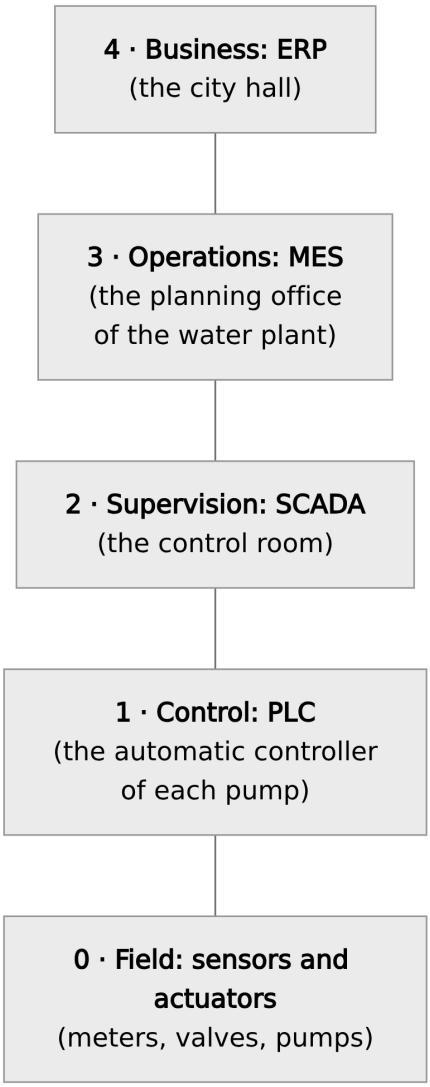
- Place sensors and actuators, **PLC**, **SCADA**, **MES** and **ERP** on the **automation pyramid**
- Recognise those levels in a real plant in Mallorca and on a real SCADA screen
- **Select** the plant and business technologies a specific company needs, and avoid the trap of "the more technology, the better"
- Explain what travels between IT and OT, what the company gains, and what risk it takes
- Describe an **end-to-end** digital company and list its advantages
- Start your **Deliverable 1** with real, cited sources

1.3.1 The automation pyramid

In the city-builder game, the water plant has levels. Down at the bottom, valves open and pumps spin. A bit higher, automatic controllers keep the pressure right. Higher still, an operator watches the whole network on a big screen. At the top, the city hall decides budgets and plans. Companies with a physical operation are organised in the same layers.

> The **automation pyramid** is a model that orders the technologies of a company in levels, from the machines at the bottom to the business at the top.

Level	Technology	What it does
4 · Business	**ERP**	Plans and manages the company: orders, purchasing, finance
3 · Operations	**MES**	*Manufacturing Execution System*: manages production orders, traceability, quality and the performance of the lines



Level	Technology	What it does
2 · Supervision	**SCADA**	*Supervisory Control And Data Acquisition*: collects data from the controllers, shows them on screens, raises alarms and lets operators give orders from a control room
1 · Control	**PLC**	*Programmable Logic Controller*: an industrial computer that reads sensors and drives actuators in real time, following a program
0 · Field	**Sensors** and **actuators**	Sensors measure (temperature, level, pressure, presence); actuators act (motors, valves, relays)

Read it **from the bottom up**: the higher you go, the less real time and the more business view. Levels 0 to 2 are **OT**; level 4 is **IT**; level 3 is the **bridge** between them.

> **Different sources, different numbers**
> This five-level pyramid is the common simplified picture of the **ISA-95** standard (IEC 62264) and the Purdue model. Some sources number the levels slightly differently, or put sensors on level 1. What matters is the order, not the numbers.

A real plant in Mallorca: TIRME at Son Reus

At Son Reus, near Palma, **TIRME** runs the plant that burns Mallorca's household waste and turns the heat into electricity. Its control system was supplied by the Japanese company Yokogawa, which published the case. According to that case, the plant has four lines, and the two newest ones generate 38 MW, enough electricity for about 53,000 households.

Here is what the published case lets you place on the pyramid:

Level	What the source says about TIRME
4 · Business	Nothing. If you wrote "TIRME uses an ERP", it would be a **hypothesis**, not a fact
3 · Operations	Exaquantum, a plant information system that collects the emissions data for reports
2 · Supervision	A central control room with operator screens, staffed 24 hours a day, and a large screen showing the environmental data
1 · Control	The CENTUM control system, which keeps the combustion above 850 °C, and a separate safety system (ProSafe-RS)
0 · Field	Field instruments connected by an industrial bus (FOUNDATION Fieldbus), feedwater pumps, a steam turbine

Two lessons hide in this table. First, real plants do not always match the textbook: TIRME uses a large control system called a DCS that does the job of the PLCs and the SCADA together, and a plant information system rather than a factory MES. Second, look at level 4: **the source does not say**. In your project, what you find you cite; what you cannot find you label as hypothesis.

What a SCADA screen looks like

Inductive Automation, a company that makes SCADA software, keeps a public demo of a water treatment plant that you can open in the browser, with no account: Ignition Water Treatment Demo.

Open it and look for: the **filters** with their mode and status, the **pumps** in automatic mode switched *On* or *Off*, the **tanks** with their level, the **flows** in and out of the plant, and a **trend** chart of the last few days. Every one of those numbers comes from a sensor (level 0), passes through a controller (level 1) and appears on this screen (level 2). It is the city's info-view overlay, for real.

1.3.2 Business technologies

You met the ERP and the CRM in the previous lesson. Here is the full catalogue of the business side:

Technology	What it is for
ERP and **CRM**	Integrated management of the company and of the relationship with customers
Business intelligence (BI)	Dashboards and reports to decide with data
Website and e-commerce	Selling and serving customers online, at any hour
Cloud collaboration tools	Email, shared documents, video calls, task management
Document management and e-signature	Contracts, invoices and files without paper
Automation of office tasks	Programs that do repetitive work: moving data from one system to another, sending reminders

To **select** technologies means to choose, for one specific company, the ones it needs and the ones it does not. A four-person carpentry workshop does not need a MES; a water company cannot live without a SCADA.

> **Careful: "the more technology, the better"**
> Every technology costs money, training and maintenance, and adds something that can be attacked. A good selection is short and justified: *this company needs this, for this reason*. In Deliverable 1 you will be marked on the reasons, not on the length of the list.

1.3.3 Why connect IT and OT

Remember the hotel in the first lesson: the booking system (IT) switches on the air conditioning (OT) before the guest arrives. Generalise it:

Direction	What travels	What the company gains
OT → IT (plant to business)	Units produced, energy used, temperatures, breakdowns, downtime	Knowing what happens in real time, real costs, **traceability**, maintenance before the machine breaks
IT → OT (business to plant)	Orders, production orders, recipes, plans	The customer's order reaches the line without anybody typing it again: fewer errors, less time

Without the connection there are **information islands**: someone writes on paper what the machine says, and someone else types it into the computer. A mechanical water meter is an island: a person has to go, read it and write the number down. A digital meter with remote reading sends the number by itself.

How are they connected? Level 3 acts as the bridge, using standards such as **OPC UA** (an industrial standard for machines to share data securely) and **MQTT** (a lightweight messaging protocol for the internet of things), plus APIs and databases. That integration layer is software: work for developers like you.

The price: cybersecurity

Many OT devices are old and were not designed to be connected: few updates, default passwords. Once they share a network with the offices, a problem in IT can reach the plant. That is why the networks are separated, with firewalls between them, and access is controlled.

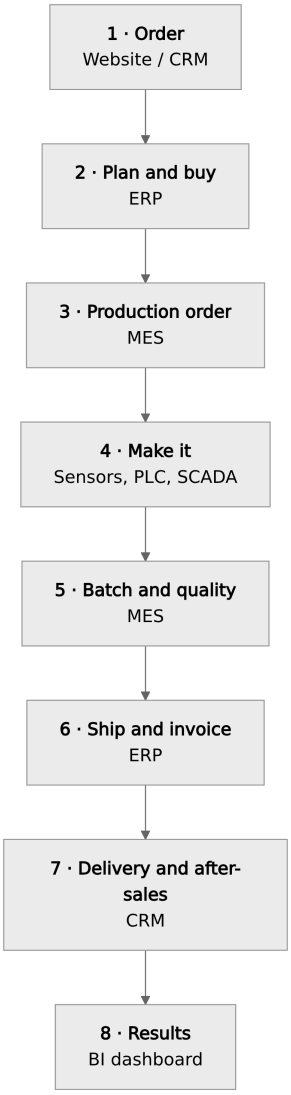
A real case: in May 2021, a ransomware attack (software that encrypts your files and asks for a ransom) hit the IT systems of **Colonial Pipeline**, which carries fuel to the east coast of the United States. There is no evidence that it reached the OT systems, but the company stopped the whole pipeline as a precaution, and it took about five days to restart. You will study cybersecurity properly in Unit 5.

1.3.4 End to end

> A company is digitalised **end to end** when the whole journey, from the customer's order to delivery and the invoice, runs through connected systems, without typing any piece of data twice.

Follow an order in a biscuit factory:

What does the company gain?



Advantage	What it means
Real-time visibility	Everybody knows where every order is
Fewer errors, less repeated work	Each piece of data is entered once
Traceability	For each packet you know the batch, the line, the day and the raw materials
Faster answers	Shorter delivery times for the customer
Less stock and less waste	You make what was ordered
Maintenance before the breakdown	Machine data warn you before it fails
Energy and material savings	What you measure, you can reduce
Decisions with data	Instead of intuition

It also has a price: investment, training for the staff, and more exposure to cyberattacks. But a company with the whole journey connected can answer in hours what costs another one days.

1.3.5 In class: express X-ray of your company

In this session your group chose its company and each member did a first piece of research for one part. The parts are the four individual sections of Deliverable 1.

Part	Question to answer	Where to look
A · Company and IT systems	What does the company do, which departments does it have, and which **real IT system** does it use (an ERP, a CRM, an app, an online shop)?	The company website (“About us”, “Work with us”), annual report, press. Tip: job offers often say “experience with SAP” or similar
B · Plant and pyramid	What physical operation does it have, and which **real plant technology** is mentioned (sensors, control, control room)? Place it on the pyramid	News about the plant, supplier case studies (like Yokogawa’s for TIRME), the company website. If you find nothing about your company, a real case **from the same sector**, cited
C · IT–OT connection	Open the Ignition SCADA demo, take a screenshot and note **three data** it shows. Which of them, or a similar one in your company, should reach an IT system? What for? What risk?	The Ignition Water Treatment Demo and your company
D · End to end	What is your company’s customer journey (order, service, invoice)? Write the first five steps and the system of each	The company website (how you buy or book), press

For each part, at least: the answer in 3–5 lines, **one source** (title, author or website, link and date you consulted it) and the word **hypothesis** in front of anything you could not find. Your name on your part. The group hands in one document on Classroom; each member is marked on their own part. **No generative AI**: this is an exercise in finding and citing.

1.3.6 For your project

This lesson launches Deliverable 1 · IT/OT X-ray, due **Sunday 8 November 2026 at 23:59**, with individual defences in session 9 (21 December 2026). Today’s four parts are the first draft of its sections:

Today’s part	Becomes	Criterion
A	Section A: company, departments and IT systems	b, d, e (business)
B	Section B: the OT side and the automation pyramid	e (plant)
C	Section C: one IT–OT connection	f
D	Section D: the end-to-end journey and its advantages	g

1.3.7 Key takeaways

- The **automation pyramid**, bottom up: **sensors and actuators**, **PLC**, **SCADA**, **MES**, **ERP**. Levels 0–2 are OT, level 4 is IT, level 3 is the bridge.
- Real plants, like **TIRME** at Son Reus, do not always match the textbook. What a source says you **cite**; what it does not say is a **hypothesis**.
- **Selecting** technologies means choosing and justifying, not listing everything.
- Connecting IT and OT sends **data up** (production, energy, breakdowns) and **orders down**; it removes **information islands** and needs **cybersecurity**.
- **End to end**: the whole journey from order to invoice through connected systems, typing each piece of data once.

1.3.8 Quick check

1 Which order of the automation pyramid, from bottom to top, is correct?

- **A** PLC → sensors → MES → SCADA → ERP
- **B** Sensors and actuators → PLC → SCADA → MES → ERP
- **C** ERP → MES → SCADA → PLC → sensors
- **D** Sensors → SCADA → PLC → ERP → MES

2 Which system lets operators watch alarms, levels and flows of the whole plant from a control room?

- **A** ERP
- **B** CRM
- **C** SCADA
- **D** A sensor

3 Which level of the pyramid acts as the bridge between the plant and the business?

- **A** Level 0, sensors and actuators
- **B** Level 1, PLC
- **C** Level 3, MES
- **D** Level 4, ERP

4 The published case of TIRME's plant says nothing about its ERP. In your deliverable, what should you write?

- **A** "TIRME uses SAP", because most large companies do
- **B** Nothing about the business side at all
- **C** A reasoned guess, clearly labelled as a hypothesis, or a different source that does say it
- **D** That TIRME has no IT

5 Which is a real advantage of connecting IT and OT?

- **A** The plant no longer needs any cybersecurity
- **B** The business decides with real plant data, almost in real time
- **C** The machines no longer need maintenance
- **D** The ERP can control the motors directly

6 Why did the Colonial Pipeline case of 2021 matter for IT–OT?

- **A** Hackers took control of the pipeline's valves
- **B** An attack on the IT systems led the company to stop the physical pipeline as a precaution
- **C** The pipeline broke because its sensors were too old
- **D** It proved that OT never needs to be connected

7 A small bakery with four workers asks you which technologies to install. Which TWO are sensible?

- **A** A simple online ordering page connected to its till
- **B** A full MES with production lines and traceability modules
- **C** A temperature sensor on the oven that sends an alarm to the phone
- **D** A control room with a SCADA staffed 24 hours a day

8 What does “digitalised end to end” mean?

- **A** Every employee has a computer
- **B** The factory has robots at both ends of the line
- **C** The whole journey, from the customer’s order to delivery and the invoice, runs through connected systems without typing data twice
- **D** The company sells only online

You have the whole picture of Unit 1. Now turn it into your group’s first chapter: Deliverable 1 · IT/OT X-ray.

Solutions

1. B · From the machine up to the office: **sensors and actuators** measure and act, the **PLC** controls, the **SCADA** supervises, the **MES** manages operations and the **ERP** runs the business.

2. C · **SCADA** collects the data from the controllers and shows them on screens with alarms and trends. The Ignition demo you opened is a SCADA screen.

3. C · Levels 0–2 are OT and level 4 is IT. The **MES**, at level 3, receives production orders from the ERP and sends back what was actually produced.

4. C · Facts about a company must be **found and cited**. If you cannot find it, you may still reason about it, but you must label it as a **hypothesis**.

5. B · Data going **up** from the plant (production, energy, breakdowns) let the business calculate real costs and react quickly. The connection **adds** cybersecurity risk; it does not remove it.

6. B · The ransomware hit **IT**, and there is no evidence it reached OT, yet the company stopped the whole pipeline for about five days. Once both worlds are connected, a problem in one affects the other.

7. A, C · **Selecting** means fitting technology to the company. A small online shop and one useful sensor solve real problems; a MES and a 24-hour control room are designed for large plants.

8. C · **End to end** is about the whole **journey** being connected: order, planning, production, shipping, invoicing, after-sales and results.

1.4

Deliverable 1 · IT/OT X-ray

Before you change anything in a city-builder game, you open the map. Where is the city hall? Where are the water plant, the power station, the depots? Which pipes and cables join them, and where does a citizen's request travel from the moment they ask until it is done? Only with that map can you plan the next district.

Deliverable 1 is that map for your company. Your group draws its **IT/OT X-ray**: the office side with its departments and systems, the operational side with its automation pyramid, one connection between both, and the whole journey of a customer's order from end to end. It is the first chapter of your digital transformation dossier, and every later deliverable builds on it.

What you will learn

- Describe a real company's departments and the IT systems they use, with sources
- Place a real company's plant technologies on the **automation pyramid**
- Design and justify one **IT–OT connection**: the data, the systems, the gain and the risk
- Draw the **end-to-end** journey of a customer's order and argue its advantages
- Integrate four individual sections into one coherent document, and defend your own section alone

1.4.1 What you hand in

Format	One PDF per group, about 6–10 pages without the cover
Where	Google Classroom, assignment "D1 · IT/OT X-ray". The Coordinator hands it in
Deadline	**Sunday 8 November 2026, 23:59**
Defence	Individual, in **session 9, Monday 21 December 2026** (the same day as Deliverable 2)
Weight	70% of the RA1 mark. The other 30% are the class activities of sessions 1 to 3

Every fact about your company needs a source

Everything you state as a fact about your company (what it does, its departments, its systems, its plant) must be **found and cited**: its website, annual report, the press, job offers, supplier case studies, Ibestat, the BOIB. Each citation has a title, the author or website, the link and the date you consulted it. What you cannot find, you may still propose with reasons, but you label it **hypothesis**. A claim without source or label counts as an error.

No generative AI

Generative AI is **not allowed** in this deliverable. A section written with AI or copied scores **0**, and the defence will show it: you cannot explain a text you did not write.

1.4.2 The sections

Each member **owns and signs one section** (the individual part). With 3 members, Section D becomes part of the group part and everybody writes it together.

Section	Criterion	What to write	Owner
A · The company, its departments and IT systems	RA1 b), d), e) business	What the company does, how big it is and where it operates. Its **departments** and, for at least four of them, the **real IT system** they use (ERP, CRM, PMS, online shop, BI…) and the data it manages. How digital technology has changed the way the company is **organised** (a new department, a role, a process). Which **business technologies** it should add or drop, with one reason each.	Individual
B · The OT side and the automation pyramid	RA1 e) plant	The company's physical operation: plants, installations, vehicles, buildings. Its **automation pyramid** (a diagram) with **one concrete example per level**: which sensor or actuator, what the PLC controls, what the SCADA shows, what level 3 manages, which ERP. Mark each example as **cited** or **hypothesis**. If the company publishes nothing about its plant, use a real cited case from the same sector to support your hypotheses.	Individual
C · One IT–OT connection	RA1 f)	**One** connection between a system of section B and one of section A: **which data** travel, **from which system to which**, in which direction, how often, what the company **gains** (with a number if you can find or estimate one), and which **risk** it opens and how it would be reduced. A small diagram.	Individual
D · The end-to-end journey and its advantages	RA1 g)	The journey of one customer order or service, from the first contact to delivery, invoice and after-sales, as a flow diagram with the **system** at each step. Where data are still typed twice today (or probably are). The **advantages** of connecting the whole journey for this company, and its costs.	Individual (4th member)
Introduction	RA1 a)	What the company is, why you chose it, and what digitalisation means **for this company**, using digitisation, digitalisation and digital transformation correctly.	Group
IT vs OT comparison	RA1 c)	A table comparing the company's IT and OT: what they manage, priority, lifespan, response time, updates, who runs them, and at least **two similarities**, all with examples from this company.	Group
Integration	RA1 a)–g)	One page that joins the four sections: one diagram with the pyramid, the connection of section C and the journey of section D, and a short conclusion. The four sections must talk	Group

Section	Criterion	What to write	Owner
		about the same company, the same systems and the same names.	
Sources	All	One list of every source cited, in the same format.	Group
Group log	All	Who owned which section, who had which role, what you decided and when (template on the project page).	Group

1.4.3 Roles in this deliverable

Choose one role each and write it in the group log. In Deliverable 2 everybody changes role and section.

Role	In Deliverable 1
Coordinator	Splits the sections, sets an internal deadline (we suggest Thursday 5 November), checks the group log and hands in the PDF
Source checker	Opens every link, checks every fact about the company has a source and every guess is labelled hypothesis
Technician	Draws the diagrams (pyramid, connection, journey) and checks that each system is on the right level
Editor	Format, spelling, one voice; the same name for the same system in every section. With 3 members, the Coordinator does it

1.4.4 How it is assessed

Part	Weight	What is assessed
Your section	40%	Content correct for its criterion, real and cited sources, hypotheses labelled, reasons given, clear writing
Your defence	20%	You explain your section in your own words, answer why, and know how it connects to another member's section
Group part	40%	Introduction, IT vs OT comparison, integration diagram, coherence between sections, sources list, format and spelling, group log

Rubric

Part	Excellent	Adequate	Insufficient
Your section	Everything the table asks for, with real systems and real examples; every fact cited; hypotheses labelled and reasoned; terms used precisely	Most of what is asked; some facts without source or some generic examples; the main ideas are correct	Missing parts; generic text that could be about any company; facts without sources; levels or terms confused
Your defence	Answers fluently and with reasons; can say where each fact comes from; explains how a teammate's section links to yours	Answers the main questions with some help; knows the sources of the key facts	Cannot explain parts of the section, or does not know where its content comes from
Group part	The four sections clearly describe the same company and fit in one diagram; the comparison has differences and similarities with exam-	The sections fit with small contradictions; comparison mostly correct; minor format problems	Four separate texts glued together; comparison generic or missing; no log; many spelling or format errors

Part	Excellent	Adequate	Insufficient
	ples; clean format, one voice, complete log		

Programme rules: RA1 needs at least a 5. Late work is marked out of 5; more than 48 hours late is a 0. Plagiarism or unauthorised AI is a 0. If you fail RA1, the recovery is an individual task: you redo one section of this deliverable alone and defend it, for a maximum of 5.

1.4.5 The defence

On **Monday 21 December 2026** (session 9) the groups work on their project while the teacher calls each student for **2–3 minutes**. No paper and no computer: a conversation.

- **Two questions about your own section.** For example: “Why is this system on level 2 and not on level 3?”, “Where did you find that the company uses this ERP?”, “What would happen to your connection if the network went down?”
- **One question about another member’s section**, to check that the group really put the sections together. For example: “Which step of the journey in section D uses the connection of section C?”

1.4.6 Before you hand in

Check	Who
Every section has its owner’s name	Coordinator
Every fact about the company has a working source; every guess says hypothesis	Source checker
The pyramid has one example per level, and each is on the right level	Technician
The section C connection appears in the section D journey and in the integration diagram	Technician
The same system has the same name in every section	Editor
The IT vs OT comparison includes at least two similarities	Editor
The group log is filled in and signed	Coordinator
One PDF, uploaded to Classroom before Sunday 8 November 2026 at 23:59	Coordinator

Start from what you did in class: the four parts of the express X-ray are the first draft of sections A to D.

UNIT 2



02

Unit 2 · Enabling technologies

In Unit 1 you drew the X-ray of your company: its offices (IT), its physical operation (OT) and the pipes between them. Now comes the fun part of any city-builder game: the **tech tree**. Every upgrade you unlock (sensors on the water pipes, a control room, robots in the depot, a 3D model of the whole city) makes something possible that was impossible before. In this unit you meet the eleven upgrades that matter most today, see what each one is good for, and decide which ones your company should unlock first.

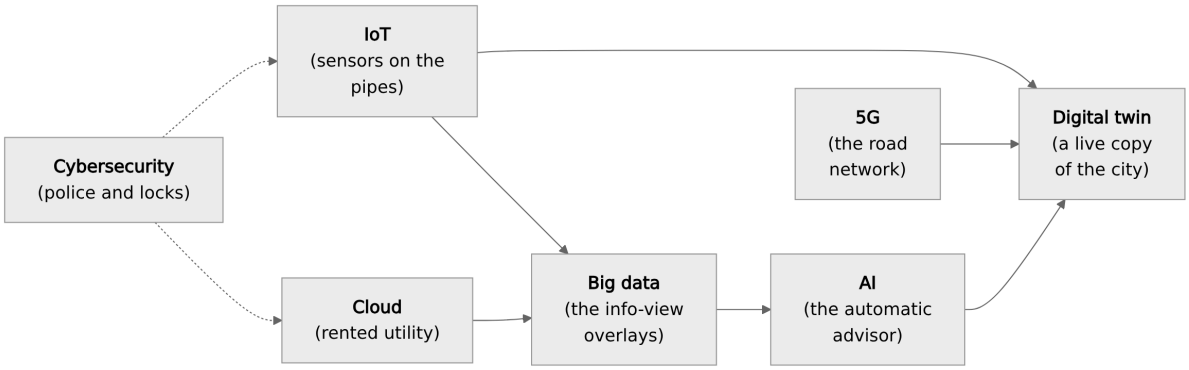
What you will learn

- Name the eleven **enabling technologies** of the unit and say in one sentence what each one does
- Tell a **characteristic** (what a technology is like) from an **application area** (where and for what it is used)
- Find a real deployment of a technology, ideally in the Balearic Islands, and cite it properly
- Explain how these technologies change products and services and create **new markets**
- Weigh their benefits and costs for a **sustainable** economy, and place each one in the business side, the plant side or both
- Write the **technology report** of your company: Deliverable 2 of the course project

The arrows are prerequisites, as in a real tech tree: a digital twin needs live data, and live data need sensors and a network. Cybersecurity protects every node.

Lessons

#	Lesson	What it covers
1	Enabling technologies I	IoT, cloud, big data, AI and cybersecurity; follow a real sensor
2	Enabling technologies II	Blockchain, robotics, additive manufacturing, AR/VR, 5G and digital twins; read real spec sheets
3	Products, services and markets	Connected products, servitisation, business models and new markets
4	Sustainability, business and plant	The two faces of sustainability, business vs plant, IT and OT improvements; measure a footprint
5	Writing the technology report	Search, evaluate and cite sources; the structure of the report
6	Deliverable 2 · Technology report	What you hand in, who writes what, how it is assessed, the defence



2.1

Enabling technologies I

You are at the bus stop and the app says: “Line 3, 4 minutes”. How does it know? Nobody is phoning the driver. Somewhere a small box on the bus measures its position, sends it over the mobile network to a server, the server compares it with months of past trips, and a program predicts when the bus will reach your stop. Four different technologies, one sentence on your screen.

Companies on the islands use the same building blocks. In February 2023 EMAYA, Palma’s municipal water company, had already replaced 30,000 mechanical water meters with smart ones, out of about 104,000 planned by 2026. A smart meter reads the consumption every day, hour by hour, without anybody visiting the house, and raises an alarm when it detects a leak. That is not one technology either: it is a sensor, a network, a data platform and an analysis, all working together.

What you will learn

- Explain what an **enabling technology** is and why it is called “enabling”
- Describe five of them, **IoT**, **cloud**, **big data**, **artificial intelligence** and **cybersecurity**, with their key characteristics
- Draw the **data chain** that links the five in one solution
- Tell a **characteristic** from an **application area**, the confusion that costs most marks in the report
- Follow a real sensor on a public map and find a real deployment of a technology, with a proper citation

2.1.1 Unlocking the tech tree

> A **digital enabling technology** is a digital technology that, on its own or combined with others, makes possible processes, products, services or business models that were not viable before. It is not the final product: it is what the product is built on.

In a city-builder game you do not buy “a better city”. You unlock upgrades in the **tech tree**: water sensors, a traffic control room, an automatic advisor. Each upgrade does nothing spectacular by itself; it **enables** the next buildings and policies. Enabling technologies work the same way, and they share three features:

Feature	What it means	Example
Transversal	The same technology serves very different sectors	Sensors are used in a hotel, a farm and a hospital

Feature	What it means	Example
Combinable	Real solutions chain several of them	The bus app needs a sensor, a network, a server and a prediction
Foundation	Industry 4.0, from Unit 1, is built on them	A smart factory is IoT + cloud + data + AI + security

There is no single official list. This module works with **eleven**, the ones in its programme. You meet five today and six in the next lesson.

This lesson	Next lesson
Internet of things (IoT) · Cloud · Big data · Artificial intelligence (AI) · Cybersecurity	Blockchain · Robotics · Additive manufacturing · Augmented and virtual reality · 5G · Digital twins

2.1.2 Internet of things (IoT)

> The **internet of things (IoT)** is the network of physical objects with **sensors** (which measure) and **actuators** (which act) that are connected to the internet, send data and can receive orders.

Characteristic	In plain words
Sensor + connectivity + platform	Something measures, something carries the data, something stores and shows it
Many small, cheap devices	Thousands of meters, not one big machine
Two directions	The thermostat sends the temperature and receives “switch off”
Industrial version	In factories it is called **IIoT** (industrial IoT)

EMAYA’s smart water meters are IoT: the meter is the sensor, the radio network carries the reading, and a platform shows it to the company and to the customer.

2.1.3 Cloud

> The **cloud** is computing (servers, storage, applications) that you use over the internet, on demand, instead of buying and running your own machines.

In the city-builder, the cloud is renting water and electricity from a utility instead of building your own power station. The US standards body NIST lists five characteristics, and they are the ones to remember:

Characteristic	In plain words
On-demand self-service	You switch it on yourself, without phoning anybody
Broad network access	It works from anywhere with internet
Resource pooling	Many customers share the same big data centres
Rapid elasticity	It grows and shrinks with the demand: more servers in August, fewer in January
Measured service	You pay for what you use, like the electricity bill

You will study the cloud properly in Unit 3. Today you only need to know what it is and what it is good for.

2.1.4 Big data

> **Big data** means data sets so large, so fast or so varied that ordinary tools (a spreadsheet, one database server) can no longer handle them.

The classic way to recognise it is the **three Vs**:

V	Question	EMAYA example
Volume	How much?	About 104,000 meters reporting every day
Velocity	How fast does it arrive?	A reading every hour, all day
Variety	How many kinds?	Meter readings, alarms, customer records, the weather

> **Careful: big data is not "a lot of data"**
> Volume is only one V. A huge file that arrives once a year in one format is just a big file. Velocity and variety are what force new tools.

2.1.5 Artificial intelligence (AI)

> **Artificial intelligence (AI)** is software that does tasks we associate with human intelligence: recognising, predicting, generating content or making decisions. Today it almost always **learns from data** instead of following rules written by a programmer.

The EU AI Act (Regulation (EU) 2024/1689) defines an AI system as a machine-based system that **infers** from the input it receives how to generate outputs such as predictions, content, recommendations or decisions. In the city-builder, AI is the automatic advisor that tunes the traffic lights.

Characteristic	Consequence
Learns from data	Without good data, it learns nothing useful
Gives probable answers	It can be wrong, and it does not warn you
Needs computing power	It usually runs in the cloud

Unit 4 is all about AI. Remember the course rule: **generative AI is not allowed** in assessed work unless the teacher authorises it for a specific activity.

2.1.6 Cybersecurity

> **Cybersecurity** is the set of technical and organisational measures that protect systems, networks and data.

It protects three properties, the **CIA triad**:

Property	Question	If it fails
Confidentiality	Can only the right people read it?	Customer data leaked
Integrity	Is it correct and unaltered?	A meter reading changed
Availability	Is it there when you need it?	Nobody can work

A real case, very close: on **13 January 2024** the Ajuntament de Calvià suffered a **ransomware** attack (a program that encrypts the files and asks for money to unlock them) by the LockBit 3.0 group. The attackers demanded 10 million euros; the council refused to pay and spent months restoring its systems. Availability failed for a whole town hall.

Every sensor you connect is one more door somebody can try to open. That is why cybersecurity is not "one more technology" at the end: it is a condition for all the others. Unit 5 goes deeper.

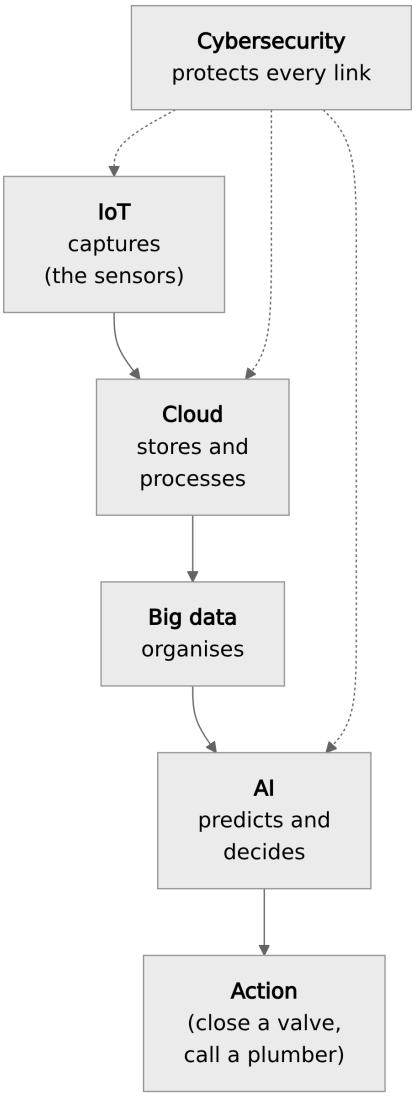
2.1.7 The data chain

The five technologies of today are rarely used alone. Most real solutions chain them like this:

A sensor with nobody analysing its data is junk; an AI with no data learns nothing. When you propose a technology for your company, always ask: what comes before it in the chain, and what comes after?

2.1.8 Characteristic or application area?

This is the distinction the report lives on, so learn it now.



> A **characteristic** says what the technology **is like**. An **application area** says **where and for what** it is used in a company.

Technology	Characteristic	Application area
IoT	Many cheap connected sensors	Water network: detect leaks at night
Cloud	Rapid elasticity	Booking system: absorb the August peak
AI	Learns from data	Sales: forecast next month's demand
Cybersecurity	Protects availability	IT department: backups and recovery plan

Test: if your sentence names a department, a process or a place in the company, it is an application area. If it would be true in any company, it is a characteristic.

2.1.9 In class: follow a real sensor

Two short tasks, done in the browser, handed in as one document on Classroom (class activity, individual). **No AI tools.**

Part 1 — A real IoT device. Sensor.Community is a citizen-science network of home-made air-quality sensors that publish their readings openly.

Step	Do this	You should see
1	Open `maps.sensor.community` and zoom into Mallorca	Hexagons coloured by air quality
2	Click a hexagon near Palma or your town	The sensors inside it, with their ID and their latest readings
3	Write down the sensor ID, the PM2.5 and PM10 values and the time of the reading	Numbers in µg/m³ and a time-stamp a few minutes old
4	Draw the data chain of that sensor: what captures, what carries, where it is stored, who could analyse it, what could go wrong	One line per link of the chain

Part 2 — Deployment hunt. Find **one real deployment** of one of today's five technologies, the one that belongs to your section of Deliverable 2 (if your section has none of today's five, take any). Prefer the Balearic Islands or your company's sector.

Write	Example (do not copy it)
What, who, where, when	EMAYA replaces mechanical water meters with smart ones in Palma, 2023–2026
Which characteristic it shows	Many cheap connected sensors; hourly readings (velocity)
Business or plant?	Plant (the water network), and business (billing, customer website)
The source, cited	Mallorcadiario. (2023, February 15). *Emaya apuesta por los contadores inteligentes*. https://www.mallorcadiario.com/emaya-apuesta-por-los-contadores-inteligentes

A company website, a public body, a newspaper or a specialised medium is fine. A forum post or an AI answer is not a source.

2.1.10 For your project

Today your group splits Deliverable 2: each member takes one **section** (a set of two or three technologies) and writes it down in the group log. The deployment you found today is field

2 of your first technology sheet, and the definition and characteristics are field 1. Keep the citation exactly as you wrote it: in the report every fact about a real company needs one.

2.1.11 Key takeaways

- An **enabling technology** makes new processes, products and services possible; it is an upgrade in the tech tree, not the final building.
- **IoT** captures data from the physical world with **sensors** and acts with **actuators**.
- The **cloud** is computing on demand over the internet: self-service, network access, pooling, elasticity, pay per use.
- **Big data** is defined by the **three Vs**: volume, velocity and variety, not by size alone.
- **AI** learns from data to predict, generate or decide, and it can be wrong. **Cybersecurity** protects confidentiality, integrity and availability, and it protects all the others.
- A **characteristic** says what a technology is like; an **application area** says where and for what the company uses it.

2.1.12 Quick check

1 Why is a technology called “enabling”?

A Because it is the newest technology on the market
B Because it makes possible processes, products or services that were not viable before
C Because it replaces the workers of a department
D Because the government subsidises it

2 EMAYA’s smart water meters send a reading every hour. Which technology captures that reading?

A The internet of things (IoT)
B Artificial intelligence
C Cybersecurity
D Big data

3 A data set of a single huge file that arrives once a year in one format. Is it big data?

A Yes: anything large is big data
B Probably not: big data also needs velocity and variety, not only volume
C Yes, as long as it is stored in the cloud
D No, because big data only exists in social networks

4 Which one is a characteristic of the cloud, and not an application area?

A Hosting the hotel’s booking system
B Storing the backups of the accounting department
C Rapid elasticity: resources grow and shrink with demand
D Running the warehouse’s stock control

5 In the Calvià attack of January 2024, the files were encrypted and nobody could work. Which property of the CIA triad failed most visibly?

A Confidentiality
B Integrity
C Availability
D Elasticity

6 In the typical data chain, what is the right order?

- **A** AI → IoT → cloud → big data
- **B** IoT captures → cloud stores → big data organises → AI predicts, with cybersecurity protecting everything
- **C** Cybersecurity → AI → IoT → cloud
- **D** Big data → IoT → AI → cloud

7 A bakery with three shops wants to know each morning how many loaves to bake in each shop. Pick the two technologies that are most directly needed.

- **A** Big data: put together past sales, the day of the week, the weather and local festivities
- **B** AI: learn from that history to predict tomorrow's demand per shop
- **C** Robotics: a robot arm that kneads the dough
- **D** Blockchain: a shared ledger among the three shops

Five technologies are about data. The next six touch the physical world, the network and trust. Next: Enabling technologies II.

Solutions

1. B · An **enabling technology** is a foundation, like an upgrade in a tech tree: it is valuable for what it lets you build on top of it, not by itself.

2. A · The meter is a **sensor** connected to a network: that is **IoT**. Big data and AI may come later in the chain, when the readings are stored and analysed.

3. B · Big data is defined by the **three Vs**. Volume alone makes a big file; **velocity** and **variety** are what make ordinary tools fail.

4. C · **Elasticity** describes what the cloud is like, in any company. The other three name a department or a process: they are **application areas**.

5. C · **Availability** means the system is there when you need it. Ransomware attacks it directly. (Elasticity is a characteristic of the cloud, not part of the triad.)

6. B · Data are born in the physical world (**IoT**), travel to where they are stored and processed (**cloud**), are organised (**big data**) and turned into a prediction or decision (**AI**). **Cybersecurity** is not a step; it covers every link.

7. A, B · The problem is a **prediction**, so you need organised data (**big data**) and something that learns from it (**AI**). A robot would change how bread is made, not how many loaves; a blockchain only makes sense with several parties that do not trust each other.

2.2

Enabling technologies II

Have you ever walked around the island of Cabrera? Since July 2026 you can do it from your phone: the Consell de Mallorca and the University of the Balearic Islands (UIB) published *Mallorca Digitalitzada*, a website and an app with 360-degree videos, 3D reconstructions and augmented reality for 50 heritage sites, from Bellver Castle to La Dragonera. The project cost 600,000 euros of EU funds. Is that virtual reality? Augmented reality? A digital twin? By the end of this page you will be able to tell.

The five technologies of the previous lesson were mostly about data. Today's six are more varied: some touch the physical world (robots, 3D printers), one connects (5G), and some change how we see or how we trust (augmented reality, digital twins, blockchain). They are the upgrades at the far end of the tech tree, and most of them need the first five as prerequisites.

What you will learn

- Describe **blockchain**, **robotics**, **additive manufacturing**, **augmented and virtual reality**, **5G** and **digital twins**, with their key characteristics and limits
- Give a real example of each one, at the plant or in the business
- Read a manufacturer's **spec sheet** and compare two real products
- Avoid three classic traps: blockchain is not cryptocurrency, 5G is not just "faster internet", and a 3D model is not a digital twin
- Decide which of the eleven technologies fit your company, which might, and which do not

2.2.1 Blockchain

A **blockchain** is a shared record (a **ledger**) copied on many computers, where entries are grouped in blocks linked by cryptography, so that nobody can change a past entry on their own without the others noticing.

Characteristic	In plain words
Distributed	Many synchronised copies, not one central database
Tamper-evident	Changing an old block breaks the links and everybody sees it
Consensus	Participants agree on new entries, usually without a central authority
Smart contracts	Small programs stored on the chain that run automatically when a condition is met

A real case: in a pilot with IBM on Hyperledger Fabric, the US supermarket chain Walmart cut the time needed to trace a packet of sliced mangoes back to its farm from about **7 days to 2.2 seconds**. Farmers, packers, transporters and the shop all wrote to the same ledger.

> **Careful: blockchain is not cryptocurrency**
> Cryptocurrencies are one application of blockchain, not the technology itself. And blockchain is usually slower and more expensive than an ordinary database. It only makes sense when **several parties that do not fully trust each other** must share one record. A single company with its own database does not need it.

2.2.2 Robotics

> **Robotics** is the design and use of programmable machines that do physical tasks, more and more often next to people.

Type	What it is like	Typical task
Industrial robot	Fast, precise, heavy loads; works behind a fence or safety barrier	Welding, painting, palletising
Cobot (collaborative robot)	Designed to work next to people without a fence; slower, lighter loads	Packing, screwing, feeding a machine
Autonomous mobile robot	Moves around a building on its own	Carrying goods in a warehouse or a hospital
Drone	Flies, usually with a camera	Inspecting roofs, solar fields, crops

Robots are described by a **spec sheet** (the manufacturer's technical data). The Universal Robots UR5e cobot, for example, lifts a **payload** of 5 kg with a **reach** of 850 mm. Closer to home, the Son Espases hospital in Palma uses a Da Vinci Xi surgical robot, and in August 2025 it performed the first robotic thoracic surgeries in the Balearic public health system.

> **Careful: software robots are not robots**
> **RPA** (robotic process automation) is software that repeats rule-based tasks on a screen, as a person would: copying invoices into the ERP, filling forms. It is very useful, and many DAM graduates will build it, but it is software, not robotics.

2.2.3 Additive manufacturing

> **Additive manufacturing** (3D printing) builds an object by adding material layer by layer from a digital 3D model, instead of cutting it out of a block or pouring it into a mould.

Characteristic	Consequence
Layer by layer from a file	You can send the part by email and print it far away
Complex shapes, no moulds	One printed piece can replace an assembly
Short runs and custom parts	Prototypes, spare parts, prostheses
Slow and expensive at scale	For a million identical parts, a mould still wins

The famous case: GE Aviation's fuel nozzle for the LEAP jet engine. Printed in metal as **one piece instead of 20**, it is 25% lighter and five times more durable than the previous model.

2.2.4 Augmented and virtual reality

> **Augmented reality (AR)** adds digital information on top of what you see (through a phone or glasses). **Virtual reality (VR)** replaces what you see with a simulated environment (through a headset). In between, **mixed reality** places digital objects that interact with the real room.

	AR	VR
What you see	The real world + digital layers	Only the simulated world
Device	Phone, tablet, smart glasses	Headset
Plant example	A technician sees repair instructions on the machine	Training operators on a dangerous procedure
Business example	A shop app that shows the sofa in your living room	A virtual visit to a hotel room before booking

Mallorca Digitalitzada mixes several things: its 360-degree videos and 3D reconstructions are virtual experiences you watch on a screen, and the project also includes augmented-reality content in an experimental app. Limits: devices cost money, 3D content is expensive to produce, and long sessions tire people.

2.2.5 5G

> **5G** is the fifth generation of mobile networks, defined by the International Telecommunication Union (ITU) as IMT-2020.

The ITU describes 5G with **three usage scenarios**, and they are its real characteristics:

Scenario	What improves	Use
eMBB (enhanced mobile broadband)	Much more speed	Video, mobile work
URLLC (ultra-reliable low-latency communications)	Very low **latency** (delay) with very high reliability	Remote control of machines, vehicles
mMTC (massive machine-type communications)	Huge numbers of devices at once	Massive IoT: meters, sensors

5G also allows **private networks** inside a factory or a port, and **network slicing**: several virtual networks on the same antennas, each with its own guarantees. In Spain the first commercial 5G service was launched by Vodafone on 15 June 2019 in 15 cities.

> **Careful: 5G is not just "faster internet on the phone"**
> For a company, the interesting parts are latency and the number of devices, not the download speed. A port that wants to control cranes remotely cares about URLLC; a water company with 100,000 meters cares about mMTC.

2.2.6 Digital twins

> A **digital twin** is a virtual replica of a real asset, process, building or city that is **kept up to date with real data** (usually from sensors), so that you can monitor it, simulate changes and predict what will happen.

Characteristic	In plain words
Connected to the real thing	Live data flow from the asset to the model
Simulate before acting	Try a change on the copy, not on the factory
Predict	Anticipate failures, peaks, risks
Needs the other technologies	IoT feeds it, the cloud hosts it, AI predicts with it

In the city-builder, the digital twin is the whole game: a live copy of the city on your screen, fed by the sensors. In November 2025 the Govern de les Illes Balears announced a *Gemelo Digital de Baleares* (about 4.6 million euros) to monitor the archipelago in real time and simulate scenarios for emergencies, water resources and tourist flows. It was an announcement: when you cite it, say so.

> **Careful: a 3D model is not a digital twin**
> A beautiful 3D reconstruction of Bellver Castle that never changes is a model. It becomes a twin only when real data keep it in sync with the real thing. That is the test to apply to *Mallorca Digitalitzada*: is anything there fed by live data?

2.2.7 The six, side by side

Technology	One characteristic	One limit
Blockchain	Tamper-evident shared ledger	Slower and costlier than a database
Robotics	Repeats physical tasks precisely	Investment, integration, safety of people
Additive manufacturing	Complex parts without moulds	Slow and expensive for large series
AR / VR	Digital layers on reality / a simulated world	Device cost, content creation, fatigue
5G	Low latency and massive numbers of devices	Uneven coverage, compatible devices
Digital twins	A live, data-fed replica	Useless without good sensors and data

2.2.8 In class: spec-sheet detective

Each member works on today's technologies of their own section of Deliverable 2, in the browser, and hands in one **technology card** per technology on Classroom (class activity, individual). **No AI tools.**

Technology	Your task	Hand in
5G	Open a public coverage map such as nPerf (`nperf.com`) and look up your company's address	Which operators offer 5G there, the date of the map data, and which of the three scenarios the company would need
Digital twins	Find a real digital twin (a city, port, building or factory)	Name, who, when, and the test: what live data feed it?
AR / VR	Open `mallorcadigitalizada.com` and try two resources	For each one: AR, VR, 360° video or 3D model, and why
Robotics	Open the spec sheets of two real robots of the same type (two cobots, two warehouse robots…) from the manufacturers' websites	A comparison table: payload, reach or speed, weight, typical task; which one could work at your company

Technology	Your task	Hand in
Additive manufacturing	Find a real part made by 3D printing in industry, and one part of your company that could be printed	The case with its source; the part and why
Blockchain	Read a real traceability case (the Walmart case study at `lfdecentralizedtrust.org` is one)	Who the parties are, why they need a shared ledger, and whether your company has several parties that do not trust each other

Each card has three fields: **what it is and three characteristics**, **a real deployment with its source cited**, and **verdict for our company** (fits / maybe / does not fit, with half a line of reason). With the minutes left, the group fills a table of the eleven technologies with the verdict of each one.

2.2.9 For your project

The cards of today and of the previous lesson are fields 1, 2 and 7 of your technology sheets in Deliverable 2. The group table of eleven verdicts is the first draft of the **summary matrix** of the group part. “Does not fit” is a good answer if it is reasoned: a software company hardly needs additive manufacturing, and saying so shows judgement.

2.2.10 Key takeaways

- **Blockchain** is a distributed, tamper-evident ledger; it only pays off when several parties that do not trust each other share one record.
- **Robotics** ranges from fenced **industrial robots** to **cobots** that work next to people; **RPA** is software, not robotics.
- **Additive manufacturing** builds layer by layer from a file: complex, custom and short-run parts, but slow for big series.
- **AR** adds layers to the real world; **VR** replaces it.
- **5G** is defined by three scenarios: more speed, very low **latency** with high reliability, and massive numbers of devices.
- A **digital twin** is a replica kept in sync with real data; a 3D model without live data is not one.

2.2.11 Quick check

1 A winery, its grape growers, the bottling plant and the distributor want one shared record of where every bottle comes from, and none of them wants another to control it. Which technology fits best?

- **A** A spreadsheet in the winery’s cloud
- **B** A blockchain shared by all the parties
- **C** A digital twin of the winery
- **D** A cobot in the bottling line

2 What is the main difference between an industrial robot and a cobot?

- **A** The cobot is always faster and carries more weight
- **B** The cobot is designed to work next to people without a fence, usually slower and with lighter loads
- **C** The cobot is software that automates screen tasks
- **D** There is no difference: it is a marketing name

3 A boatyard needs a spare plastic part for an old engine that is no longer manufactured. Which technology helps most?

- **A** 5G
- **B** Blockchain
- **C** Additive manufacturing
- **D** Virtual reality

4 A technician wears glasses that show the next repair step on top of the real machine. What is it?

- **A** Augmented reality
- **B** Virtual reality
- **C** A digital twin
- **D** Robotics

5 For a port that wants to control cranes remotely, which 5G characteristic matters most?

- **A** The download speed for videos
- **B** Very low latency with very high reliability (URLLC)
- **C** Massive numbers of cheap devices (mMTC)
- **D** The new phone models

6 A city publishes a detailed 3D model of its old town, built in 2024 and never updated. Is it a digital twin?

- **A** Yes, any 3D model of a real place is a digital twin
- **B** No: a digital twin must be kept up to date with real data from the real place
- **C** Yes, as long as it can be viewed in virtual reality
- **D** No, because digital twins only exist in factories

7 Which statements about blockchain are true? Pick two.

- **A** Cryptocurrencies are one application of blockchain, not the technology itself
- **B** A blockchain is always faster than an ordinary database
- **C** Changing an old entry is detected because the blocks are linked by cryptography
- **D** A blockchain needs one central authority to validate every entry

You now know the eleven upgrades of the tech tree. The next question is what comes out of them: new products, new services and markets that did not exist. Next: Products, services and markets.

Solutions

1. B · Several parties that do not fully trust each other, one record that nobody should be able to change alone: that is exactly the case where a **blockchain** is worth its cost.

2. B · **Industrial robots** are fast and strong but work behind barriers; **cobots** trade speed and payload for the ability to share the space with people. Software that automates screens is **RPA**.

3. C · One custom part, no mould, no factory making it any more: a typical job for **additive manufacturing**. Somebody models the part in 3D (or scans it) and prints it.

4. A · The technician still sees the real machine, with digital information added: **augmented reality**. In virtual reality the real world would disappear.

5. B · Remote control needs the order to arrive almost instantly and always: **URLLC**. Speed (eMBB) and massive devices (mMTC) are the other two scenarios, useful for other jobs.

6. B · The test is the **live data**. Without sensors keeping it in sync, it is a model, however beautiful. Digital twins exist for machines, buildings, ports and whole cities.

7. A, C · Blockchain is **tamper-evident** and usually works by **consensus** without a central authority. It is slower and more expensive than a database, which is why it only pays off in specific cases.

2.3

Products, services and markets

How many songs do you own? Not “have on your phone”: own, as in bought once and kept forever. Probably none. You pay a monthly fee, or you listen with ads, and the music is never really yours. Twenty-five years ago people bought CDs. The product became a service, and a whole industry changed its way of making money.

That shift was not a fashion. It was made possible by the technologies of the last two lessons: the cloud stores the catalogue, big data and AI learn what you like, and your phone is a connected product that the company improves after you have bought it. This page is about what comes out of the tech tree: new products, new services and markets that would not exist without it. Most of them are software, which is where you come in.

What you will learn

- Explain four ways enabling technologies change products and services: **connected products**, **personalisation**, faster development and **servitisation**
- Recognise the most common digital **business models**: subscription, pay per use, platform and freemium
- Name markets that would not exist without enabling technologies, including some born in the Balearic Islands
- Read a real pricing page, identify its model and calculate what it would cost your company
- Avoid the trap: a business model is how a company earns money, not the technology it uses

2.3.1 Four ways products change

Change	What it means	Example	Technologies
Connected product	The product carries sensors and software, sends usage data and can be improved after it is sold	A smartwatch; a car that gets new functions with an update	IoT, cloud, 5G
Personalisation	Each customer gets a version adapted to them, without the cost exploding	Music recommendations; a hotel offer based on your past stays	Big data, AI
Faster development	You test before you build	A 3D-printed prototype in a day; a line simulated in a digital twin	Additive manufacturing, digital twins, AR/VR

Change	What it means	Example	Technologies
Servitisation	The company stops selling the product and sells its use or its result	Engines paid per flying hour; software by subscription instead of a licence	IoT, cloud, AI

> **Servitisation** is the shift from selling a product to selling it as a service: the customer pays for the use or for the result, and the maker keeps the product working.

Two real cases. Since 1962 Rolls-Royce has offered aircraft engine contracts called *Power-by-the-Hour* (today *TotalCare*): the airline pays a fixed price per flying hour that includes maintenance, so the maker only earns while the engine works. That is why Rolls-Royce fills its engines with sensors: when you are paid for hours of flight, a breakdown is your problem. And in May 2013 Adobe announced it would stop developing the boxed Creative Suite and sell its software only by subscription, through Creative Cloud.

In the city-builder, servitisation is like the utility company you rent water from: you do not buy the pipes, you pay for the litres.

2.3.2 Business models

> A **business model** is the way a company creates value for its customers and earns money from it.

Enabling technologies made four models very common:

Model	How it earns money	Examples
Subscription	A regular fee to access the product or service	**SaaS** (software as a service), music and video streaming
Pay per use	You pay for what you actually consume: hours, requests, kilometres, gigabytes	Cloud services, shared bikes and cars
Platform (marketplace)	It connects supply and demand and takes a commission or a fee	App stores, holiday rentals, online hotel distribution
Freemium	A basic version is free; advanced functions or higher limits are paid	Many mobile apps, online collaboration tools

On top of all four there is a fifth idea: **data as an asset**. Whoever has the usage data of its customers can improve the product, personalise it and decide better than its competitors.

> **Careful: a business model is not a technology**
> "Our business model is the cloud" is a wrong sentence. The cloud is a technology; **pay per use** or **subscription** is the business model it makes possible. In the report, always say both: which technology, and how the company would charge.

2.3.3 Markets that did not exist

Some markets were born with these technologies: cloud services; mobile apps and app stores; wearables and digital health; the connected home; shared mobility; mobile payments and fintech; cybersecurity as a service; and, very close to you, **travel tech**.

Mallorca has a travel-tech sector made of platform businesses. Two real examples:

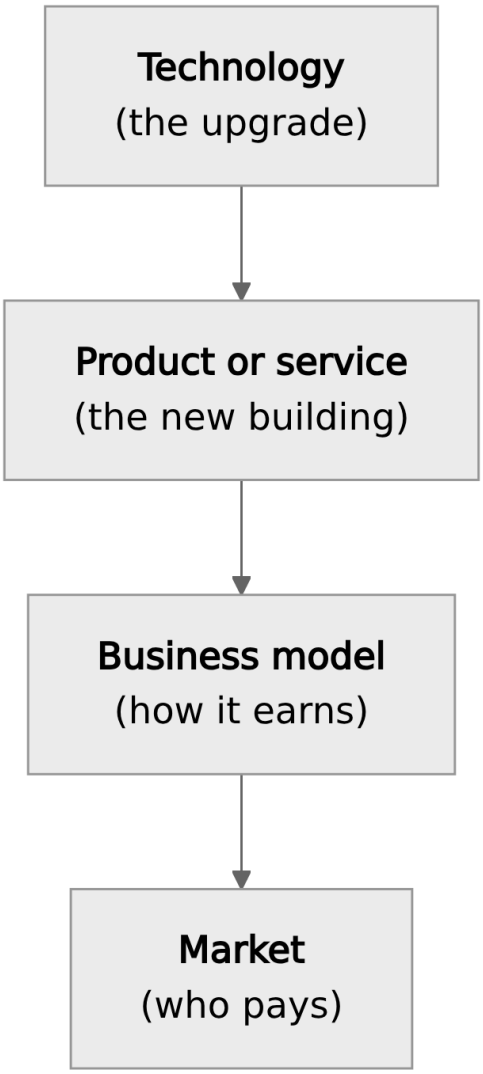
Company	What it does	Model
Hotelbeds (HBX Group, headquartered in Palma)	Connects hotels and other travel providers with travel agencies and tour operators, business to business	Platform
Habitissimo (founded in 2009 in the Balearic Islands; part of the HomeServe group since 2017)	Connects private customers with renovation professionals	Platform

Neither is a factory. Their product is a piece of software that puts two sides of a market in touch, and their market did not exist before the internet.

2.3.4 In class: business model detective

Two parts, handed in on Classroom as one document (class activity, individual). **No AI tools.**

Part 1 — Read a real pricing page. Pick one real digital product from your company's sector, or one of these: Canva, Google Workspace, Holded (a Spanish cloud ERP), Habitissimo, Hotelbeds.



Step	Do this	You should see
1	Open its official website and find the pricing or “how it works” page	Plans, fees or commissions
2	Name its business model (one or two of the four) and say why	“Freemium + subscription: free plan with limits, paid plan per user”
3	Name two enabling technologies the product needs	“Cloud (it runs in the browser), AI (it suggests designs)”
4	Calculate what it would cost your company in a year, with a number of users you justify	“8 office workers × price per user per month × 12 = …” with the date you read the price
5	Cite the page	Author or company. (Year). *Title*. URL, and the date you read it

Prices change: what matters is that your number comes from the real page and says when you read it.

Part 2 — A new product for your company. For each technology of your section of Deliverable 2, write one sentence: which product or service it would enable or improve, and which of the four changes it is. Then, in the last minutes, the group picks the most promising one and writes its pitch in this fixed format:

> *For* (customer), *who* (need), *we offer* (product or service), *thanks to* (technologies), *and we charge* (business model).

2.3.5 For your project

Part 2 gives field 4 of each of your technology sheets in Deliverable 2 (criterion b). The group pitch is the seed of the **new market** section of the group part (criterion d): does it open a new market for the company, or improve its position in the one it already has? Facts about the company need a source; the product is your proposal, so say clearly that it is one.

2.3.6 Key takeaways

- Enabling technologies change products in four ways: **connected products**, **personalisation**, faster development and **servitisation**.
- **Servitisation** sells the use or the result instead of the product: Rolls-Royce’s engines paid per flying hour, Adobe’s move to subscription.
- A **business model** is how a company earns money: **subscription**, **pay per use**, **platform** or **freemium**, often combined.
- Data about how customers use a product are an asset in themselves.
- New markets were born with these technologies, such as travel-tech **platforms** in Mallorca.
- Behind almost every new digital business model there is software somebody has to write.

2.3.7 Quick check

1 What does servitisation mean?

A Hiring more people for customer service
B Moving from selling a product to selling its use or its result as a service
C Moving the company’s servers to the cloud
D Giving the product away for free

2 Under Power-by-the-Hour, why does the engine maker fill its engines with sensors?

- **A** Because the law forces it
- **B** Because it is paid only while the engine works, so it needs to prevent breakdowns
- **C** To sell the sensor data to other airlines
- **D** Because sensors make engines lighter

3 A design app has a free version with limits and a paid version per user per month. Which models is it using?

- **A** Platform only
- **B** Pay per use only
- **C** Freemium combined with subscription
- **D** Servitisation of hardware

4 What does a platform (marketplace) business do?

- **A** It manufactures the products it sells
- **B** It connects supply and demand and earns a commission or a fee
- **C** It rents servers by the hour
- **D** It only sells advertising

5 A classmate writes in the report: “The business model of our proposal is artificial intelligence.” What is wrong?

- **A** Nothing: AI is a valid business model
- **B** AI is a technology; the business model is how the company would charge for the product AI enables
- **C** Business models only exist for big companies
- **D** The report should not mention business models

6 Which of these markets exists thanks to enabling technologies?

- **A** Selling fresh fish at the town market
- **B** Wearables for digital health, such as smartwatches that measure your heart rate
- **C** Renting deckchairs on the beach
- **D** Selling paper road maps

7 A water company installs smart meters and offers businesses a paid service that sends an alert to their phone when it detects a leak at night. Pick the two true statements.

- **A** It is a new service built on IoT and data analysis
- **B** It is a freemium model because the water is not free
- **C** It could be charged as a subscription added to the water bill
- **D** It is servitisation because the company stops selling water

Technologies make new products possible, but at what cost for the planet, and where in the company do they act? Next: Sustainability, business and plant.

Solutions

1. B · With **servitisation** the customer pays for the use (hours, litres, results) and the maker keeps the product working, like Rolls-Royce’s engines paid per flying hour.

2. B · When you are paid per flying hour, every hour on the ground costs **you**. IoT sensors and predictive maintenance protect the maker’s own income.

3. C · The free basic version is **freemium**; the monthly fee for the full version is a **subscription**. Real products often combine models.

4. B · Like Hotelbeds (hotels and travel agencies) or Habitissimo (customers and renovation professionals), a **platform** puts two sides of a market in touch.

5. B · A **business model** answers "how do we earn money?": subscription, pay per use, platform, freemium. The technology answers "what makes it possible?". The report needs both.

6. B · **Wearables** need sensors (IoT), the cloud and data analysis. The other three markets existed long before any of these technologies.

7. A, C · The meter data (**IoT**) analysed for anomalies create a new **service**; a monthly fee would make it a **subscription**. The company still sells water, so it is not servitisation of its main product.

2.4

Sustainability, business and plant

Last night you watched a match on your phone. No paper, no petrol, no plastic: digital is clean, right? Somewhere a data centre sent you that video, its servers ran all night, and they were cooled by machines that use electricity and sometimes water. Your phone was made from metals dug out of a mine, and one day it will become electronic waste.

So, are enabling technologies good or bad for the planet? The honest answer is "it depends, and you have to measure it". On islands with limited water and land, and a population that swells every summer, measuring and adjusting consumption matters a lot. This page gives you both faces of the coin, then answers a second question your report needs: where does each technology act, in the offices or on the plant floor, and what does it improve there?

What you will learn

- Explain how enabling technologies help an **efficient** and **sustainable** economy (face A) and what environmental cost they have (face B)
- Recognise the **rebound effect**, the trap of "more efficient means less consumption"
- Place each technology in the **business** side, the **plant** side or both, with an example
- Identify the improvements they bring to **IT** and **OT** environments and to the connection between them
- Measure two real footprints in the browser: the carbon intensity of Mallorca's electricity and the estimated emissions of a website

2.4.1 Efficient and sustainable

> An **efficient** economy gets the same result with fewer resources: energy, water, raw materials, time. A **sustainable** economy does it without compromising the future: environmental, economic and social.

The European Union talks about the **twin transition**: the digital transition and the green transition must move together, because they can reinforce each other, and they can also clash.

2.4.2 Face A: how the technologies help

Technology	What it allows	Effect
IoT + big data	Measure energy and water use in real time; spot leaks and waste	You cannot reduce what you do not measure
AI	**Predictive maintenance** (anticipating failures from sensor data); optimising routes, loads, air conditioning, production	Fewer breakdowns, less fuel, less energy

Technology	What it allows	Effect
Digital twins	Simulate a change before making it	Fewer physical tests, fewer costly mistakes
Additive manufacturing	Use material only where needed; print spare parts locally, on demand	Can reduce waste, transport and stock
Cloud	Share infrastructure among many customers in large data centres	Usually better used than many half-empty servers of your own
Blockchain	Trace the origin of products along the supply chain	Supports the circular economy and transparency

2.4.3 Face B: the cost of digital

Cost	What happens	A number to remember
Energy	Data centres use a lot of electricity, and AI is the main driver of its growth	About 415 TWh, around 1.5% of the world's electricity in 2024 (International Energy Agency)
Water	Depending on the cooling system, data centres also consume water	It varies a lot by site and climate
Materials and waste	Devices need critical raw materials, and constant renewal generates **e-waste**	62 million tonnes of e-waste in 2022, only 22.3% documented as properly collected and recycled (ITU and UNITAR)

> The **rebound effect** happens when making something more efficient and cheaper leads people to use more of it, so total consumption can end up rising.

> **Careful: "digital is always greener"**
> An app that saves paper but runs on servers switched on all day is not automatically more sustainable. Neither is it automatically worse. The only honest answer is to **measure** both sides, which is exactly what you will do in class today.

2.4.4 Business or plant?

Remember Unit 1: the **business** side is the offices, management and customer relations (IT systems); the **plant** side is where the product is made or the service is physically delivered (OT systems). In the city-builder, the city hall and the infrastructure.

Technology	In the business	On the plant
IoT	Product usage data for sales and after-sales	Sensors on machines, installations, vehicles
Cloud	ERP, CRM, email and collaboration tools as a service	Platform where plant data are sent and analysed
AI	Sales forecasts, chatbots, fraud detection	Quality control by camera, predictive maintenance
Big data	Customer and market analysis	Analysis of process and production data
Cybersecurity	Protecting customer data and management systems	Protecting control systems: stopping them stops the plant
Blockchain	Contracts and certificates	Traceability of batches and raw materials
Robotics	Software robots (RPA) for admin tasks	Industrial robots, cobots, mobile robots
Additive manufacturing	Prototypes for design and marketing	Spare parts, tools, short runs

Technology	In the business	On the plant
AR / VR	Sales demos, virtual visits	Maintenance assistance, operator training
5G	Mobile work	Wireless connection of machines, vehicles, sensors
Digital twins	Simulating scenarios to decide investments	Replica of the line, the machine or the building

2.4.5 What they improve in IT and OT

Improvements in IT	Improvements in OT	Improvements from connecting IT and OT
Scale resources with demand; automate repetitive tasks; decide with data instead of intuition; work from anywhere	Fewer unplanned stops; more quality and less scrap; safer work for people; remote supervision and control	The business sees the plant in real time; the plant produces according to real demand; traceability from the order to the product

An improvement only counts if you can check it. That is why every improvement in the report needs a **KPI** (key performance indicator): one number you would measure before and after. "Fewer stops" is vague; "unplanned stop hours per month" is a KPI.

2.4.6 In class: measure the footprint

Three short tasks in the browser, handed in as one document on Classroom (class activity, individual). These tools only look at public information, as any visitor does: **passive only**. **No AI tools.**

Step	Do this	You should see
1	Open Electricity Maps (`app.electricitymaps.com`), find the zone **Mallorca** and then **Spain** (mainland)	The carbon intensity of each, in gCO_2eq/kWh, and the mix of sources right now
2	Write both numbers, the time, and the main source in each	Two different numbers: an island grid and a mainland grid are not the same
3	Open the Website Carbon Calculator (`websitecarbon.com`) and test your company's home page	Grams of CO_2 per visit and whether the hosting appears as green
4	Check the same domain with the Green Web Foundation's Green Web Check	Whether the hosting provider is listed as running on green energy
5	Write two sentences: what the numbers suggest, and why they are estimates and not exact measurements	The calculator estimates from the page weight and average grid figures

Then, for each technology of your section of Deliverable 2, write:

Field	What to write
Business or plant	Business, plant or both; the exact department or process
Sustainability	One benefit and one cost, and what you would measure
IT / OT improvement	Which environment (IT, OT or the connection) and one KPI

2.4.7 For your project

The three fields you just wrote are fields 3, 5 and 6 of your technology sheets in Deliverable 2 (criteria e, c and f). The two measurements of today are good evidence to quote in field 5: a real number with its source and date is worth more than a paragraph of good intentions. If

your company publishes a sustainability report, cite it; if not, write that you looked for it and did not find it.

2.4.8 Key takeaways

- **Efficient** means the same result with fewer resources; **sustainable** means without compromising the future. The EU calls doing digital and green together the **twin transition**.
- Face A: IoT, data and AI measure and optimise consumption; **predictive maintenance**, digital twins and additive manufacturing reduce waste.
- Face B: data centres used about 1.5% of the world's electricity in 2024, and only 22.3% of the 62 million tonnes of **e-waste** of 2022 was documented as recycled.
- The **rebound effect** can turn efficiency into more total consumption: measure, do not assume.
- Every technology acts on the **business** side, the **plant** side or both, and improves IT, OT or the connection between them.
- An improvement needs a **KPI**: one number you would measure before and after.

2.4.9 Quick check

1 What is predictive maintenance?

A Repairing machines only after they break
B Replacing every part once a year, whatever its state
C Anticipating failures from sensor data before they happen
D Buying a new machine every five years

2 A company makes its delivery app so efficient that each delivery costs much less, and customers order three times as often. Total fuel use goes up. What is this called?

A Twin transition
B Rebound effect
C Servitisation
D Predictive maintenance

3 Which statement about the environmental cost of digital technologies is correct?

A Digital technologies have no environmental cost because they are virtual
B Data centres use electricity and sometimes water, and devices become e-waste
C Only cryptocurrencies consume energy
D All e-waste is recycled in Europe

4 Which one is an enabling technology on the business side?

A A cloud CRM with AI that ranks which customers to call first
B IoT sensors watching the temperature of an oven
C A cobot that palletises boxes
D A 5G network connecting cranes in a port

5 Which improvement belongs to the OT environment?

A Scaling servers with demand without buying hardware
B Fewer unplanned machine stops
C Working from home with cloud tools
D Faster approval of invoices

6 A student writes as the improvement of a technology: “It will make the company more efficient.” What is the best advice?

- **A** It is perfect: efficiency is the goal
- **B** Name the process and add a KPI you would measure before and after, such as litres lost per day
- **C** Remove it: reports should not talk about improvements
- **D** Add more technologies so the improvement is bigger

7 Your company’s website shows a high CO_2 figure per visit in the Website Carbon Calculator. Pick the two sensible conclusions.

- **A** It is an estimate based on the page weight and average figures, not an exact measurement
- **B** It proves the company pollutes more than all its competitors
- **C** Making the page lighter (smaller images, fewer scripts) would probably lower the estimate
- **D** The company must be breaking the law

You now have every piece of the technology report. The last lesson is about putting it together well: finding sources you can trust, citing them and writing it as a team. Next: Writing the technology report.

Solutions

1. C · **Predictive maintenance** uses IoT data and analysis (often AI) to fix a machine just before it would fail: fewer breakdowns and less waste than repairing late or replacing early.

2. B · When efficiency makes something cheaper and people use much more of it, total consumption can rise: the **rebound effect**.

3. B · Face B is real: about 1.5% of the world’s electricity went to data centres in 2024, and most **e-waste** is not documented as properly recycled.

4. A · The **CRM** serves sales and customer relations: offices, IT. The other three act on the physical operation: the **plant**.

5. B · Unplanned stops happen on the plant floor, so it is an **OT** improvement. The other three are IT improvements.

6. B · An improvement only counts if it can be checked. A **KPI** turns a vague promise into something the company can measure.

7. A, C · The calculator **estimates** emissions from data transferred, so a heavier page scores worse. It is useful for comparing and improving, not as proof of anything.

2.5

Writing the technology report

Imagine you are the manager of your company and two reports land on your desk. Both say "we should install smart sensors on the water network". The first one says it because "it is the future". The second one says it because the company lost a known amount of water last year (link to its annual report), because EMAYA did something similar in Palma (link to the news), and because a sensor of that type costs a given amount (link to the manufacturer). Which one do you believe?

The second one, obviously. Not because it is longer, but because you can **check** it. That is the whole difference between an opinion and a report. This page is about the craft of the second kind: finding sources, deciding whether to trust them, citing them, and putting four people's work together into one document that reads as if one person wrote it.

What you will learn

- Search the web with **search operators** to find primary sources quickly
- Evaluate a source with the **SIFT** method and **lateral reading**, and tell a **primary source** from a secondary one
- Separate **facts** (with a source) from **hypotheses** (your reasoned guesses), and label each one
- Cite web sources in **APA style**, in the text and in the list of sources
- Know the structure of Deliverable 2 and how the group turns four sections into one report
- Avoid the trap: "I found it on Google" is not a source

2.5.1 Search like a professional

A search engine gives you what most people click, not what is most reliable. **Search operators** are short codes that tell it exactly what you want.

Operator	What it does	Example
`"..."`	Exact phrase	`"contadores inteligentes" Emaya`
`site:`	Only one website or domain	`site:caib.es gemelo digital`
`filetype:pdf`	Only PDF files (annual reports, official documents)	`Iberostar sostenibilidad filetype:pdf`
`-`	Exclude a word	`blockchain -bitcoin -crypto`
`OR`	Either word	`cobot OR "robot colaborativo" Mallorca`

Where to look, depending on the claim:

You need	Look in
What the company is and does	Its own website, "about us", annual or sustainability report
Official figures about the islands	Ibestat, the Govern (`caib.es`), the BOIB, the Spanish INE
What a technology is	Standards bodies and public agencies: ITU, NIST, the European Commission
A real deployment	The company's press room, regional and specialised press
A product's characteristics	The manufacturer's spec sheet

Most search engines also let you filter by date (in Google: *Tools* → *Any time*). Use it: a "latest" article from 2017 is not the latest.

2.5.2 Can I trust this source?

> A **primary source** is the original: the company's own report, the law in the BOIB, the manufacturer's spec sheet. A **secondary source** reports on it: a newspaper article, a blog post.

Primary sources are best for facts; secondary sources help you find them and add context. Before you use any source, apply **SIFT**, a four-step method proposed by the digital-literacy researcher Mike Caulfield:

Step	Question	What you do
S · Stop	Do I know this website? Why am I trusting it?	Pause before you copy anything
I · Investigate the source	Who is behind it, and what do they want?	**Lateral reading**: open new tabs and search what others say about that website, instead of reading only its own "about us"
F · Find better coverage	Do other reliable sources say the same?	Search the claim itself and compare
T · Trace to the original	Where does this figure or quote come from?	Follow the links back to the primary source and cite that one

Warning signs: no author and no date, a page that sells the product it praises, figures with no origin, a forum or social-media post, and any answer written by an AI chatbot. An AI answer is not a source: it can invent references that look real, and in this course it is not allowed in assessed work unless the teacher authorises it.

2.5.3 Facts and hypotheses

Your company will not publish everything you would like to know. That is normal, and the course rule is simple:

> A **fact** is a statement you can support with a source, and you cite it. A **hypothesis** is a reasoned guess you could not verify, and you label it as such.

Statement	Type	How to write it
The company has 3 hotels in Mallorca	Fact	"…has three hotels in Mallorca (Company, 2026)."
The company probably uses a cloud booking system	Hypothesis	"*Hypothesis:* given its size and its online booking page, it probably uses…"
A cobot UR5e lifts 5 kg	Fact	"…a payload of 5 kg (Universal Robots, n.d.)."

A hypothesis is not a weakness. Presenting a guess as a fact is.

2.5.4 Cite it

The report uses a simplified **APA style** (7th edition), the most common in technical and social-science writing. Each citation has two parts.

In the text, right after the claim: (Author, year). The author can be a company or a public body.

> EMAYA expects about 104,000 smart meters installed by 2026 (Mallorcadiario, 2023).

In the list of sources, at the end, in alphabetical order:

Kind of source	Pattern
Web page or news	Author or organisation. (Year, Month Day). *Title*. Site name. URL
Report (PDF)	Organisation. (Year). *Title of the report*. URL
Page that changes over time (a live map, a price list)	Organisation. (n.d.). *Title*. Retrieved Month Day, Year, from URL

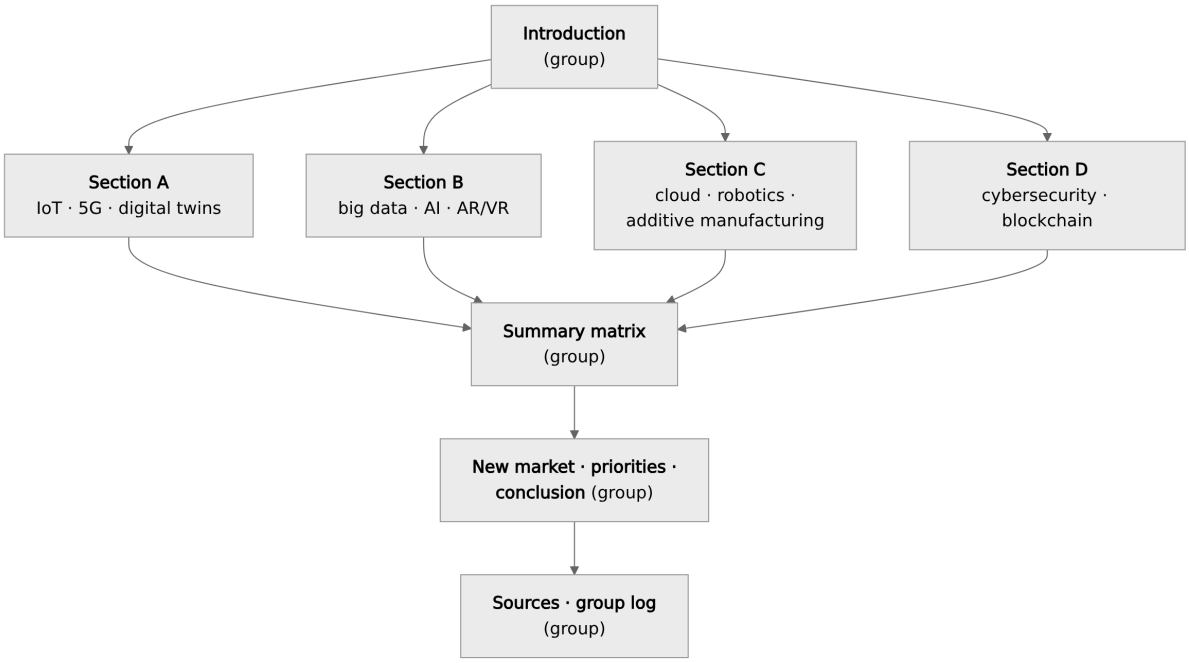
`n.d.` means “no date”. If the author and the site name are the same, write it once. Every source in the list must be cited in the text, and every citation in the text must be in the list.

2.5.5 The structure of Deliverable 2

The report is built from the work of the last four sessions. Each member owns a section of two or three technologies; the group writes the parts that tie them together. The full brief is on the Deliverable 2 page.

Each technology gets the same **technology sheet** with seven fields, so that the four sections can be compared:

Field	Content	Lesson
1	What it is and three characteristics	Enabling technologies I and II



Field	Content	Lesson
2	A real deployment, with its source	Enabling technologies I and II
3	Where in the company: process, business or plant	Sustainability, business and plant
4	Product or service it would enable	Products, services and markets
5	Sustainability: a benefit, a cost, what to measure	Sustainability, business and plant
6	IT or OT improvement, with a KPI	Sustainability, business and plant
7	Verdict: fits, maybe, does not fit, and why	Enabling technologies II

2.5.6 Writing as a team

Habit	Why
One shared document, everybody writes in it	The version history shows who did what; it is part of the group log
The same sheet format for every technology	The reader can compare; the summary matrix almost writes itself
The **Source checker** opens every link before the hand-in	A dead link is a missing source
The **Editor** reads the whole report aloud once	Four voices become one
Cite while you write	Adding sources at the end means losing half of them

2.5.7 In class: source audit

Individual first, then as a group. Handed in on Classroom (class activity). **No AI tools.**

Part 1 — Audit three claims of your own section.

Step	Do this	You should see
1	Pick three statements from your technology sheets: one about a technology, one about a real deployment, one about your company	Three sentences
2	For each, search with at least one operator and write the query you used	e.g. `site:emaya.es contadores`
3	Apply SIFT to the best result: who is behind it, primary or secondary, date	Three short notes
4	Write the APA reference and the in-text citation	Three references in the pattern above
5	If one claim cannot be verified, rewrite it as a labelled hypothesis	"*Hypothesis: …*"

Part 2 — Assemble the skeleton. In the group's shared document: put the headings of the report in place, paste each member's sheets in their section, start the summary matrix and the group log, and let the Source checker list the sources that are still missing.

2.5.8 For your project

Today you leave the session with Deliverable 2 assembled as a draft and a list of what is missing. The report is due on **Monday 21 December 2026 at 23:59**, and in the session of that same day each of you will defend your own sections of Deliverable 1 and Deliverable 2 in a short conversation with the teacher.

2.5.9 Key takeaways

- **Search operators** (`"..."`, `site:`, `filetype:`, `-`, `OR`) take you straight to primary sources.
- A **primary source** is the original; a **secondary source** reports on it. Cite the primary one whenever you can.
- **SIFT**: Stop, Investigate the source, Find better coverage, Trace to the original. Investigate with **lateral reading**.
- A **fact** carries a citation; a **hypothesis** carries a label. Never present a guess as a fact.
- **APA style**: (Author, year) in the text, the full reference in the alphabetical list of sources.
- An AI chatbot's answer is not a source, and AI is not allowed in assessed work without the teacher's authorisation.

2.5.10 Quick check

1 You want only PDF documents from the Govern's website about digital twins. Which query does that?

A `gemelo digital PDF Govern`
B `site:caib.es filetype:pdf gemelo digital`
C `"caib.es" OR "pdf" gemelo digital`
D `gemelo digital -caib.es -pdf`

2 Which one is a primary source for the number of smart meters EMAYA plans to install?

A A forum thread where a neighbour comments on his bill
B A chatbot's answer
C EMAYA's own press release or report
D A blog that summarises Mallorca news

3 In SIFT, what is lateral reading?

A Reading the whole website carefully before judging it
B Opening other tabs to find what other sources say about the website you are evaluating
C Reading only the headlines of an article
D Translating a source into another language

4 You could not find whether your company uses an ERP, but it is large and has several warehouses. How do you write it?

A "The company uses an ERP." with no source
B As a labelled hypothesis, with the reasoning behind it
C Leave out anything you cannot verify, even if it is important
D Ask a chatbot and cite its answer

5 Which in-text citation follows the APA pattern of this page?

A "...installed by 2026 [link]."
B "...installed by 2026 (found on Google)."
C "...installed by 2026 (Mallorcadiario, 2023)."
D "...installed by 2026 (https://www.mallorcadiario.com, page 1)."

6 Why does every technology get the same seven-field sheet in the report?

A To make the report longer
B So the four sections can be compared and integrated into the summary matrix
C Because the teacher only reads the first field
D So that each member can copy another member's sheet

7 A news article says a hotel chain "cut its water use by 30%" and links to the chain's sustainability report. What should you cite?

- **A** The news article only, because it is shorter
- **B** The sustainability report, after checking that the 30% is really there
- **C** Nothing: figures from companies are never reliable
- **D** The search engine where you found the article

Everything is ready: the technologies, their applications and impact, and the craft of the report. The brief, the rubric and the defence are on the last page of the unit: Deliverable 2 · Technology report.

Solutions

1. B · `site:` limits the search to one domain and `filetype:pdf` to PDF files. The last option would exclude exactly what you want.

2. C · The **primary source** is the original, issued by the organisation itself. A newspaper that reports it is a good secondary source; a forum or a chatbot is not a source at all.

3. B · Instead of trusting a website's own description, you **leave it** and check what independent sources say about it. Professional fact-checkers work this way.

4. B · A reasoned guess is welcome if it is labelled: "*Hypothesis:* given its size and its warehouses, it probably uses an ERP". Presenting it as a fact would be the mistake.

5. C · In the text: **(Author, year)**. The full reference with the URL goes in the alphabetical list of sources at the end.

6. B · A common format turns four people's work into one report: the reader compares technologies field by field, and the group builds the **summary matrix** from it.

7. B · This is the **T** of SIFT: trace the claim to the **original** and cite it there. If the figure is not in the report, you have learned something important about the article.

2.6

Deliverable 2 · Technology report

The management of your company asks your team one question: *which enabling technologies suit us, what for, and where would we apply them?* Deliverable 2 is your answer. It is the second chapter of the company's digital transformation dossier, and it builds directly on the IT/OT X-ray of Deliverable 1: you already know the company's departments, systems and plant; now you decide which upgrades of the tech tree it should unlock.

The report relates the eleven technologies of the unit to their **characteristics** and to their **application areas** in your company (criterion g). Each member writes and signs one section; the group writes the parts that make it one report.

> **Facts must be cited**
> Everything you state as a fact about a real company, a real deployment or a real product must have a source you **found and cite** (company website, annual or sustainability report, press, Ibestat, BOIB, the manufacturer's spec sheet). What you could not find is written as a labelled **hypothesis**. Invented data presented as facts is the most serious mistake in this report. How to search and cite: Writing the technology report.

2.6.1 What you hand in

One PDF per group, on Google Classroom, by **Monday 21 December 2026 at 23:59**. Between 8 and 12 pages without the cover and the list of sources.

2.6.2 The sections

The eleven technologies are split into four **individual sections**. Each member owns one, writes it and signs it with their name. With three members, section D is written by the whole group and counts as group part.

Section	Criterion	What to write	Owner
Introduction	—	The company in half a page (from Deliverable 1) and three needs or problems that technology could address	Group
A · Sense and connect	a, b, c, e, f	A technology sheet for **IoT**, **5G** and **digital twins**	Individual
B · Data and intelligence	a, b, c, e, f	A technology sheet for **big data**, **AI** and **AR/VR**	Individual
C · Platforms and machines	a, b, c, e, f	A technology sheet for **cloud**, **robotics** and **additive manufacturing**	Individual

Section	Criterion	What to write	Owner
D · Trust (4th member)	a, b, c, e, f	A technology sheet for **cybersecurity** and **blockchain**	Individual (group with 3 members)
Summary matrix	g	One table: the eleven technologies × the company's areas (departments or processes), with the verdict and the main characteristic of each	Group
New market	d	One new product or service built on the technologies, its business model, and whether it opens a new market or strengthens an existing one	Group
Priorities and conclusion	g	Which two or three technologies first, in what order, and why; ten lines	Group
Sources	—	Every source cited in the text, in APA style, in alphabetical order	Group
Group log	—	Who owned which section, who had which role, what was decided and when	Group

Members take a different kind of section in every deliverable, so choose one you did not have in Deliverable 1.

The technology sheet

Every technology of an individual section gets the same sheet, about two thirds of a page:

Field	Content	Criterion
1 · What it is	Definition in your own words and three characteristics specific to this technology	a
2 · A real deployment	One real case, ideally in the Balearic Islands or in the company's sector: who, what, when, and the source	a, g
3 · Where in our company	The exact process or department; business, plant or both	e
4 · Product or service	What new or better product or service it would enable, and which change it is (connected product, personalisation, faster development, servitisation)	b
5 · Sustainability	One benefit and one cost, and what you would measure	c
6 · IT / OT improvement	Which environment it improves (IT, OT or the connection) and one KPI	f
7 · Verdict	Fits / maybe / does not fit, with the reason. "Does not fit" is a good answer when it is reasoned	a, g

2.6.3 Roles

One role each, rotated from Deliverable 1. With three members, the Editor's job merges into the Coordinator's.

Role	Job in this deliverable
Coordinator	Plans the sessions, keeps the deadline, hands in the PDF
Source checker	Opens every link before the hand-in; every fact has a working, reliable source; hypotheses are labelled
Technician	Owns the hands-on parts: the measurements of the sessions (sensor map, coverage map, Electricity Maps, Website Carbon) are correct and dated
Editor	One format, one voice, spelling, the same sheet layout for every technology

2.6.4 How it is assessed

Deliverable 2 is 70% of the RA2 mark; the class activities of the sessions are the other 30%.

Part	Weight	What is assessed
Own section (individual)	40%	Your technology sheets: correctness, specificity, link to the company, sources
Defence (individual)	20%	Your answers in the defence of Monday 21 December
Group part	40%	Introduction, summary matrix, new market, priorities and conclusion, sources, format and spelling, group log, coherence

Rubric: own section (40%)

Criterion	Excellent (9–10)	Adequate (5–6)	Insufficient (0–4)
Characteristics (field 1)	Correct, specific to each technology, in your own words	Correct but generic (“it is fast”, “it is modern”)	Missing or wrong
Real deployment (field 2)	A real, relevant case with a reliable, cited source	A real case, but weak source or loosely related	Missing, invented or not cited
Application to the company (fields 3, 4, 7)	A concrete process or department, a credible product or service, a reasoned verdict	Area named but vague; verdict without real reasons	Not linked to the company
Impact (fields 5, 6)	Benefit and cost; IT/OT improvement with a measurable KPI	Only benefits, or improvement without a KPI	Missing

Intermediate marks (7–8) sit between Excellent and Adequate.

Rubric: defence (20%)

Excellent (9–10)	Adequate (5–6)	Insufficient (0–4)
Explains your own sections in your own words, with details of the company and of your sources; answers the question about a teammate’s section	Explains the main idea but hesitates on details or sources; partial answer about the teammate’s section	Cannot explain what you signed

Rubric: group part (40%)

Criterion	Excellent (9–10)	Adequate (5–6)	Insufficient (0–4)
Summary matrix (g)	All eleven technologies against the company’s real areas, consistent with the sheets	Complete but generic, or not quite matching the sheets	Missing or incomplete
New market (d)	A coherent product or service, a clear business model, a reasoned market	An idea without a clear model or market	Missing or unrelated to the technologies
Priorities and conclusion	A clear order with reasons tied to the company’s needs	An order without real reasons	Missing
Sources and rigour	Reliable, complete APA list; facts and hypotheses clearly separated	Some sources incomplete or weak	No sources, or invented data
Format, coherence and group log	One voice, same sheet format, no spelling mistakes, complete log	Some inconsistencies or mistakes; log incomplete	Disordered, many mistakes, no log

2.6.5 Rules

Rule	What it means
No AI	Generative AI is not allowed in this deliverable. Unauthorised AI or plagiarism = 0
Late work	Up to 48 h late: maximum 5. More than 48 h late: 0
Minimum	RA2 needs at least 5
Recovery	If RA2 is failed: redo one section alone, plus a defence; maximum mark 5

2.6.6 The defence

On **Monday 21 December 2026**, in class, the groups finish the report while the teacher calls each student for **2–3 minutes**. No paper, no computer: a conversation. You will be asked:

Question	About
1	Your own section of Deliverable 1
2	Your own section of Deliverable 2
3	A section of Deliverable 2 written by a teammate

There are no trick questions. The examples are of the kind “why did you say that technology does not fit?”, “where did you find that deployment?”, “what KPI would you measure?”, “what does your teammate propose for the cloud?”. If you wrote your section yourself and read your teammates’ sections, you are ready. If you are absent with a justified reason, your defence is moved to another day.

2.6.7 Final checklist

Check	Done?
Every member’s name is on the section they signed	
Every technology has its seven fields	
Every fact about a real company or deployment has an in-text citation	
Every hypothesis is labelled as such	
The summary matrix has the eleven technologies and matches the sheets	
The new market has a product, a business model and a market	
The list of sources is in APA style, alphabetical, with working links	
The group log says who did which section and which role	
No AI was used	
One PDF, uploaded on Classroom by Monday 21 December 2026, 23:59	



03

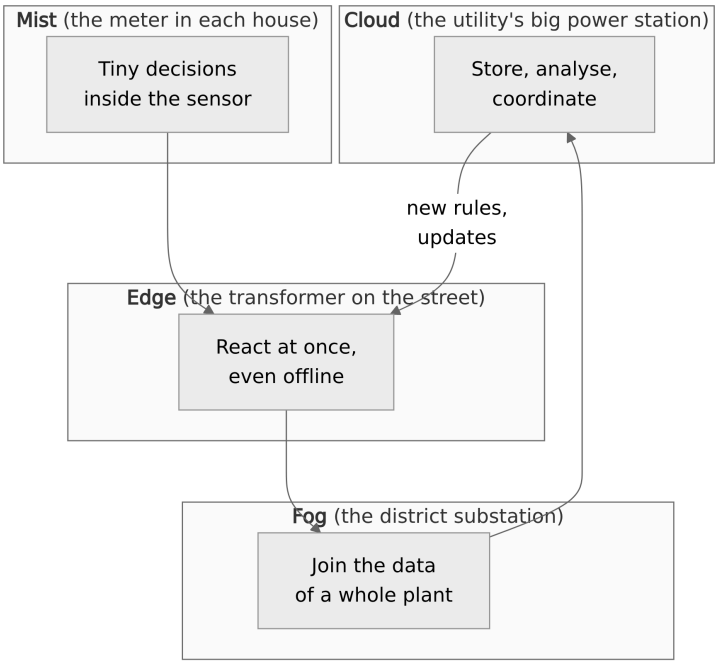
Unit 3 · Cloud systems

Back to the city-builder game. Your city needs water and electricity. You could build your own water plant and your own power station, pay for them up front and repair them yourself. Or you could sign a contract with a utility company: it runs the plants, you open the tap and pay the bill at the end of the month. Most companies today do the second thing with their computers. That is the **cloud**. In this unit you learn what exactly you can rent, who owns it, what it is used for, how much it costs, why some decisions still have to be taken in local **substations** near the neighbourhoods (**edge**, **fog** and **mist**), and why it matters in which country the utility keeps your data.

What you will learn

- Explain what the **cloud** is and tell the three **service models** apart: **IaaS**, **PaaS** and **SaaS**
- Publish a real web page with **GitHub Pages** and say which service model it is
- Tell the **deployment models** apart (public, private, hybrid, multicloud) and name the main **functions of the cloud**
- Read a provider's **pricing calculator** and produce a dated monthly cost estimate for a real system
- Decide what a real plant processes at the **edge**, in the **fog**, in the **mist** and in the cloud, and why
- Find out where a real service keeps its data and what that means for **data sovereignty**

Lessons 1 and 2 are about the top box: what you rent and what it costs. Lesson 3 is about the boxes below it. Lesson 4 asks whose power station it is.



Lessons

#	Lesson	What it covers
1	Cloud and service models	What the cloud is, IaaS, PaaS and SaaS, who manages what; your first site on GitHub Pages
2	Deployment and functions	Public, private, hybrid and multicloud; what the cloud is used for; a real estimate with a pricing calculator
3	Edge, fog and mist	Latency measured in the browser; edge computing and the cloud; fog and mist in a wind farm, a water utility and a hotel
4	Advantages and sovereignty	What the cloud brings to connected systems, its risks; where data live and which law applies
5	Deliverable 3 · Cloud architecture	What you hand in, who writes what, the group's website, how it is assessed, the defence

3.1

Cloud and service models

This morning you probably played a song on Spotify, opened Classroom and saved something to Google Drive. None of it ran on your phone, not really. Where was the song? Where is your homework right now? And who keeps the machine that stores it switched on at three in the morning?

The answer is “in the cloud”, and it is less mysterious than it sounds. A hotel in Calvià that receives ten times more bookings in August than in January cannot buy ten servers for August and throw nine away in September. It rents. As a DAM developer, almost everything you build will end up running on somebody else’s computer, so you need to know exactly what you are renting.

What you will learn

- Define **cloud computing** with its five characteristics, and tell it apart from **on-premise**
- Tell the three **service models** apart: **IaaS**, **PaaS** and **SaaS**
- Say who manages what in each model, and which part is **always** yours
- Classify real services you use, with proof from the provider’s own website
- Publish your own web page with **GitHub Pages**, and say which model it is
- Avoid the trap: “the cloud” is not in the sky, it is a building full of somebody else’s servers

3.1.1 The cloud is renting from the utility

> **Cloud computing** is on-demand access, over the network, to computing resources (servers, storage, databases, applications) that a provider shares among many customers, that can grow or shrink quickly, and that you pay for according to use.

In the city-builder game of this module, the cloud is **renting water and electricity from a utility** instead of building your own plant. You open the tap and pay the bill. The alternative is called **on-premise**: the company buys its own servers and keeps them in its own building.

The US standards institute NIST wrote the definition everyone quotes (SP 800-145, 2011). It lists five characteristics. A service that lacks them is not really cloud:

Characteristic	What it means	Example
On-demand self-service	You sign up and create resources yourself, without waiting for anyone to install anything.	Creating a server from a web form in a couple of minutes

Characteristic	What it means	Example
Network access	You reach it from any place and device with a connection.	Opening Classroom at school, at home or on your phone
Shared resources	The same hardware serves many customers, kept separate logically.	Thousands of companies in the same data centres
Elasticity	Resources can be added or removed quickly as demand changes.	The hotel's website in August and in January
Pay per use	Use is measured, and you pay for what you consume.	Like the electricity bill

Back to the hotel. On-premise, it buys a server, puts it in an air-conditioned room and pays somebody to look after it: a big investment on day one, and if it falls short in August, bad luck. In the cloud there is no initial investment, it pays monthly according to use, and in August it asks for more capacity.

> **Careful: "the cloud is somewhere up there"**
> The cloud is physical. It is a **data centre**: a building full of servers, cables, power supplies and cooling, owned by a provider such as Amazon, Microsoft, Google or Hetzner. An external hard disk or a server in the office cupboard is not cloud: nobody else shares it, it does not grow on demand and nobody bills you by use.

3.1.2 Three service models: IaaS, PaaS and SaaS

You can rent at three different heights. The higher you go, the more the provider does for you.

> A **service model** says how much of the system the provider gives you ready-made, and how much is left for you to manage.

	IaaS · infrastructure	PaaS · platform	SaaS · software
What you get	Virtual machines, network and storage	A ready environment where you deploy your code	A finished application, ready to use
What you manage	Operating system, software, application and data	The application and the data	Only your data and settings
Who it is for	System administrators	Developers	End users
Examples	Amazon EC2, Azure Virtual Machines, Google Compute Engine, Hetzner Cloud	Google App Engine, Heroku, Azure App Service, GitHub Pages	Gmail, Google Classroom, Microsoft 365, Spotify, Kahoot

- **IaaS** (*Infrastructure as a Service*): you rent a **virtual machine**, a computer simulated by software on the provider's hardware. It is like renting an empty plot with water and power connected: you build the house.
- **PaaS** (*Platform as a Service*): you bring your code and the platform runs it. Like renting a furnished flat: you bring your things.
- **SaaS** (*Software as a Service*): you just use the application. Like staying in a hotel.

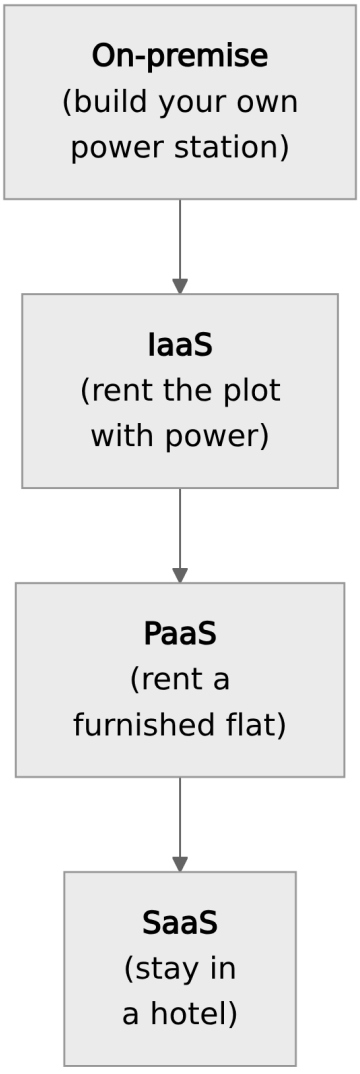
3.1.3 Who manages what

Stack the layers of any system from the building up to your data. For each model, mark what is the **customer's** job (C) and what is the **provider's** (P):

Layer	On-premise	IaaS	PaaS	SaaS
Data	C	C	C	C
Application	C	C	C	P
Runtime (Java, Node, database…)	C	C	P	P
Operating system	C	C	P	P
Virtualisation	C	P	P	P
Servers, disks and network	C	P	P	P
Building, power and cooling	C	P	P	P

Look at the first row. **Your data are always your responsibility**, whatever the model. The higher you climb, the less administration you do, but the less control you have. Providers call this the **shared responsibility model**.

> **Careful: the borders are not walls**



> Some services sit between two models. **Firebase**, for example, is a platform you build on (PaaS), but it also offers finished pieces such as user login. What is assessed is the **reason**: say what the customer manages and what the provider manages, and the label follows.

3.1.4 GitHub Pages: a platform you can use today

> **GitHub Pages** is a free GitHub service that publishes a **static website** (HTML, CSS and images that do not change per visitor) straight from a code repository.

You put the files in a repository; GitHub stores them, serves them to the world and keeps the servers running. You never see a server. That is the definition of **PaaS**, and you will use it all unit: in Deliverable 3 your group publishes its dossier as a GitHub Pages site.

Three facts to know before you start (checked on 25/10/2026 in GitHub's documentation):

Fact	What it means for you
GitHub's terms require you to be **13 or older**	Everyone in this course can open an account; it is free and needs no card
With a free account, the repository must be **public**	Everything you upload can be seen by anybody: no personal data, no phone numbers; your first name is enough
Pages is not for shops or for collecting passwords	A dossier, a portfolio or a project page is fine

3.1.5 In class: what level does it play at?

Part 1 · Classify real services. In the Classroom template:

Step	Do this	You should get
1	Write down **four services you used this week** (phone or computer) and add these four: Amazon EC2, Heroku, Google App Engine, Hetzner Cloud	A list of eight services
2	Classify each one as **IaaS, PaaS or SaaS** and write one sentence on why: what does the customer manage, and what does the provider?	Eight labels with a reason
3	For each, paste a link to the **official** page (product or pricing page, not a blog or an AI) and quote the sentence that proves it	Eight pieces of evidence
4	For your project company, list **three systems** it has or probably has (website, bookings, ERP, email, plant control…) and, for each, the model you would propose and a **real provider** that offers it	The draft of Section A of D3

Part 2 · Your first site on GitHub Pages.

Step	Do this	You should see
1	Sign in to GitHub, or create a free account at github.com/signup (email, password, username, a check and a code sent to your email)	Your GitHub home page
2	Click **+** (top right) → **New repository**. Name it exactly `yourusername.github.io`. Visibility **Public**. Turn on **Add README**. **Create repository**	Your new repository with a `README.md`
3	**Add file** → **Create new file**. Name it `index.html`. Write a minimal page: a heading with your first name and the sentence "This website is a … service because …". **Commit changes**	`index.html` in the repository

Step	Do this	You should see
4	**Settings** → **Pages**. Under **Build and deployment**, **Source**: **Deploy from a branch**; **Branch**: `main` and `/ (root)`. **Save**	A message that your site is being published
5	Wait (it can take up to 10 minutes) and open `https://yourusername.github.io`	Your page, visible from any computer in the world

If you cannot or do not want to open a GitHub account, make the page with **Google Sites** and your school account, and say in the sentence that Google Sites is SaaS, not PaaS. It does not affect your individual mark: in the deliverable, the website is group work.

Hand in on Classroom: the table of part 1 with its evidence, your three company systems and the link to your site. **No generative AI**: this activity is assessed (class work, 30%).

3.1.6 For your project

This lesson starts Deliverable 3 · Cloud architecture, due **Monday 8 February 2027 at 23:59**, with individual defences in session 14 that same day.

Today's work	Becomes	Criterion
Three company systems with a model and a real provider	Section A: service models	a
Your GitHub account and first site	The group's website, published by the Technician	Group part

3.1.7 Key takeaways

- **Cloud computing** is renting computing resources on demand, over the network, shared, elastic and paid by use. It runs in real **data centres**. The opposite is **on-premise**.
- **IaaS** gives you machines, **PaaS** gives you a place to deploy your code, **SaaS** gives you the finished application.
- The higher the model, the less you administer and the less you control. **Your data are always your responsibility** (the **shared responsibility model**).
- End users mostly meet **SaaS**; developers live mostly in **PaaS** and **IaaS**.
- **GitHub Pages** publishes a **static website** from a repository: a PaaS you can use for free, with a public repository.

3.1.8 Quick check

1 Which of these is cloud computing?

A An external hard disk where the company copies its files every Friday
B A server bought by the company and kept in its own office
C Storage rented from a provider, reached over the internet and billed by the GB used
D A USB stick shared by the whole office

2 A company rents a virtual machine and installs the operating system, the database and its application on it. Which model is it?

A IaaS
B PaaS
C SaaS
D On-premise

3 In which model does the customer manage only their data and settings?

- **A** IaaS
- **B** PaaS
- **C** SaaS
- **D** In none: the customer always manages the operating system

4 You upload the files of your website to a repository and GitHub Pages publishes it without you managing any server. Which model is it?

- **A** IaaS
- **B** PaaS
- **C** SaaS
- **D** It is not cloud, because it is free

5 Which statement about the shared responsibility model is true?

- **A** In SaaS, the provider is responsible for the customer's data
- **B** In IaaS, the provider updates the operating system of your virtual machines
- **C** Whatever the model, the customer is responsible for its data
- **D** On-premise, the provider manages the building and the power

6 A classmate says: "The cloud is just the internet; my files are in the air." What is the best answer?

- **A** True: the cloud has no physical location
- **B** The files are in real data centres owned by a provider; the internet is only the road to them
- **C** True, but only for SaaS services
- **D** False: cloud files are always stored on your own device

7 A bakery in Inca wants online ordering without hiring a programmer, and a start-up in Palma wants to deploy its own app without managing servers. Which models fit each one best?

- **A** IaaS for the bakery, SaaS for the start-up
- **B** SaaS for the bakery, PaaS for the start-up
- **C** PaaS for both
- **D** On-premise for both, to save money

You now know what you can rent. Next you will see who owns the power station, what the cloud is used for and how much the bill is: Deployment and functions.

Solutions

1. C · The cloud is **on demand**, **over the network**, **shared**, **elastic** and **paid by use**. The disk, the office server and the stick have none of those features; the office server is **on-premise**.

2. A · The provider gives only the **infrastructure** (the virtual machine, network and storage). Everything from the operating system up is the customer's job: that is **IaaS**.

3. C · In **SaaS** the application is finished and the provider runs everything below it. The customer still owns the **data**, which is their responsibility in every model.

4. B · You bring the code; the platform stores and serves it. That is **PaaS**. Being free does not stop it being cloud: the provider still runs the data centre.

5. C · The **data** row is always the customer's. In IaaS the operating system is also yours; on-premise there is no provider at all.

6. B · The cloud is physical: buildings full of servers, cables and cooling. The **network** is how you reach them. That is why, in lesson 4, it will matter **in which country** they are.

7. B · The bakery needs a **finished application** (SaaS). The start-up has its own code and wants somebody else to run the servers: that is **PaaS**.

3.2

Deployment and functions

Your phone contract has a price per month, but the real bill depends on what you did: the extra gigabytes, the calls abroad, the month you forgot to switch off roaming. Cloud bills work the same way. Before a company moves a system to the cloud, somebody has to answer three questions: whose infrastructure will it be, what will it do there, and how much will it cost each month?

In the last lesson you learnt **what** you can rent. Today you learn **from whom** and **for what**, and you do what a real IT team does before signing anything: you build a cost estimate with a provider's own **pricing calculator**, and you learn to read every line of it.

What you will learn

- Tell the four **deployment models** apart: **public**, **private**, **hybrid** and **multicloud**
- Name the main **functions of the cloud**: processing data, exchanging information and running applications
- Read a **pricing calculator**: region, instance, hours, GB, data transfer, taxes
- Build and compare a dated monthly estimate for a real system of your project company
- Avoid two traps: "private means in our building", and "cloud functions" meaning only the product called *Cloud Functions*

3.2.1 Who owns the cloud: deployment models

A **deployment model** says whose infrastructure the cloud runs on and who can use it.

Model	What it is	When it makes sense
Public cloud	The infrastructure belongs to a provider that offers it to any customer; many customers share the same hardware.	Small and medium companies, new projects, changing demand. No initial investment.
Private cloud	Cloud infrastructure for the exclusive use of one organisation. It can be in its own building or hosted by a third party.	Very sensitive data, strict legal requirements, or a need for full control.
Hybrid cloud	A combination of private (or on-premise) infrastructure and public cloud, connected and working together.	Old systems that cannot move, plus peaks of demand covered by the public cloud.
Multicloud	Using services from two or more cloud providers at the same time.	Not depending on one provider, using the best of each, or higher availability.

In the city-builder game: the **public** cloud is the city's utility, shared by every neighbourhood; a **private** cloud is a power station only for your industrial estate; **hybrid** is your own small

plant plus the utility for the peaks; **multicloud** is buying from two utilities so a strike at one does not leave you in the dark.

Take the hotel in Calvià again. Its booking website, with summer peaks, runs in a public cloud. The customer records and the old management program stay on its own servers. Both are connected: **hybrid**. If it also keeps its backups with a second provider, it is **multicloud** as well.

> **Careful: three common mix-ups**
> **Private does not mean "in our building"**: what defines it is exclusive use, not location. **Hybrid and multicloud are not the same**: hybrid mixes private and public; multicloud mixes providers, and a company can be both. **Public does not mean "open to everyone"**: each customer's data are kept apart; what is shared is the hardware. NIST also defines a fourth model, the **community cloud**, shared by organisations with a common interest, such as several public administrations.

3.2.2 What the cloud is for: its functions

> The **functions of the cloud** are the jobs a company uses it for. The three main ones are processing data, exchanging information and running applications.

Function	What the cloud does	Examples
Data processing	Provides computing power to transform, analyse or train on large volumes of data.	Analysing five years of bookings to predict occupancy; processing the readings of thousands of sensors
Information exchange	Acts as the common point where people, applications and devices pass data to each other.	Shared documents, file sync, **APIs** (agreed ways for two programs to talk) between companies, IoT devices sending readings
Running applications	Hosts and runs programs so users reach them over the network.	A booking website, the server of a mobile app, an online ERP, your GitHub Pages site
Others	Storage and backups, managed databases, development and test environments.	Phone photos that copy themselves; code repositories

One tool usually does several at once. Google Docs runs an application, stores your file and lets three classmates edit it together.

> **Careful: "cloud functions" has two meanings**
> In this module, the *functions of the cloud* are **what it is used for** (the table above). Some providers also sell a product with a similar name, such as *Google Cloud Functions* or *AWS Lambda*: small pieces of code that run when something happens and are billed per execution. That is one specific service (often called **FaaS**, *Function as a Service*), not the criterion.

3.2.3 How much does it cost: reading a pricing calculator

> A **pricing calculator** is a provider's free web tool where you describe what you would rent and it estimates what you would pay.

The big providers publish one, and none needs an account to make an estimate: the AWS Pricing Calculator, the Azure Pricing Calculator and the Google Cloud Pricing Calculator. Smaller providers such as Hetzner Cloud simply publish a price list per plan. Every estimate is built from the same fields:

Field	What it means	Typical trap
Region	Where the data centre physically is. Prices change from region to region, and so does the law (lesson 4).	Leaving the default region, often in the United States
Instance type	The size of the virtual machine: how many virtual processors (**vCPU**) and how much memory.	Choosing a huge machine "just in case"
Hours of use	A month has about **730 hours**. Always on = 730. Office hours only = far fewer.	730 hours for a system that only works by day
Storage (GB-month)	The GB of disk or storage you rent, paid per month.	Forgetting the disk of the machine
Data transfer	What goes out to the internet is usually paid per GB; what comes in, often not.	Ignoring it for a website with many visitors
Payment option	On demand (no commitment) or with a one- or three-year commitment (cheaper).	Comparing a committed price with an on-demand one
Taxes	The AWS calculator does not include taxes; in Spain you add VAT (IVA).	Presenting the total as the final price

Prices change, so **an estimate is only valid with its date**. As an example of reading, on 25/10/2026 Hetzner's page listed its smallest shared cloud servers from €5.99 and €6.49 a month, without VAT. Do not memorise that figure: open the page on the day you need it and write down the date.

> **Careful: an estimate is not a price**
> The calculator multiplies what **you** typed by the provider's current prices. If your hypothesis is wrong (twice the visitors, a machine left on all night), the bill will be different. And it never includes the people: somebody still has to configure, watch and update the system.

3.2.4 In class: how much does our system cost?

Individually or in pairs, in the browser. **No generative AI**: this activity is assessed (class work, 30%).

Step	Do this	You should get
1	Choose one system of your project company that makes sense in the cloud (website, bookings, online shop, backups, sensor data…). Write which **function of the cloud** it fulfils	One system and one function
2	Decide what it needs, as a reasoned hypothesis: how many machines and of what size, how many hours a month, how many GB	A short list of needs, labelled hypothesis
3	**Estimate 1:** build it in the AWS Pricing Calculator (**Create estimate** → search the service → **Configure**) in the region *Europe (Spain)*. Note the monthly total and **today's date**	One dated figure
4	**Estimate 2:** the same in another region (for example *Europe (Frankfurt)* or *US East (N. Virginia)*) **or** with another provider (Azure, Google Cloud or Hetzner's price list)	A second dated figure

Step	Do this	You should get
5	**Evidence:** export the estimate as PDF (or take a screenshot where the configuration and the total are visible) and copy the **Share** link	A file and a link
6	**Conclusion** in three sentences: which you would choose and why; what the estimate does **not** include; which deployment model you would propose for this system	Your reasoning

Hand in on Classroom the document with the six answers, the PDF or screenshot and the link.

3.2.5 For your project

This lesson feeds Deliverable 3 · Cloud architecture:

Today's work	Becomes	Criterion
System, function and the two dated estimates with evidence	Section B: cloud functions and real cost estimate	b, e
The deployment model you proposed	Section D: deployment model and sovereignty	e

3.2.6 Key takeaways

- **Deployment models**: **public** (a provider's, shared), **private** (exclusive to one organisation, wherever it is), **hybrid** (private plus public, connected) and **multicloud** (several providers).
- The main **functions of the cloud** are **processing data**, **exchanging information** and **running applications**, plus storage, backups and more.
- A **pricing calculator** turns your hypothesis (region, instance, hours, GB, transfer) into a monthly estimate. About **730 hours** is a whole month.
- An estimate is only valid **with its date**, does not include **VAT** in the AWS calculator, and never includes the people's work.
- Changing the **region** changes the price and, as you will see, the law.

3.2.7 Quick check

1 A hotel keeps its customer records on its own servers and its booking website with a public provider, and both are connected. Which deployment model is it?

A Public
B Private
C Hybrid
D Community

2 Which statement about a private cloud is true?

A It must always be inside the company's own building
B It is for the exclusive use of one organisation, wherever it is located
C It is shared by many customers who do not know each other
D It always means using two providers

3 A travel company uses one provider for its website and another for its backups. What is it doing?

A Private cloud
B On-premise
C Multicloud
D SaaS

4 An API connects a hotel's booking system with an online travel agency. Which function of the cloud is that?

- **A** Data processing
- **B** Information exchange
- **C** Running applications
- **D** Backups

5 A virtual machine must be on all month. Roughly how many hours do you enter in the calculator?

- **A** 24
- **B** 160
- **C** 730
- **D** 8,760

6 Your AWS estimate says €40 a month. Pick the two statements that are true.

- **A** VAT has to be added, because the AWS calculator does not include taxes
- **B** The figure is guaranteed for the next three years
- **C** The estimate should carry the date, because prices change
- **D** The figure already includes the salary of whoever manages the system

7 Two estimates for the same machine give different totals in Spain and in the United States. What is the best conclusion?

- **A** One of the two calculators is wrong
- **B** Always pick the cheaper region, nothing else matters
- **C** The region changes the price, and it also changes where the data are and which law applies, so both must be weighed
- **D** Regions only affect speed, never price

The bill tells you how much the power station costs. But what if the answer has to arrive in milliseconds, or the line goes down? Next: Edge, fog and mist.

Solutions

1. C · Private (or on-premise) infrastructure plus public cloud, connected and working together, is **hybrid**.

2. B · What defines a **private cloud** is **exclusive use**. It can be hosted by a third party in another building and still be private.

3. C · Using services from **two or more providers** at the same time is **multicloud**. It may also be hybrid if it keeps private infrastructure, but the two ideas are different.

4. B · The cloud acts as the meeting point where two systems pass data to each other: **information exchange**.

5. C · A month has about **730 hours** (8,760 hours a year divided by 12). 160 would be office hours; 8,760 is a whole year.

6. A, C · The calculator gives an **estimate without taxes**, valid on **the day you made it**, for exactly what you typed. People's work is never in it.

7. C · Prices vary by **region**, and so do latency and the legal framework. A cheaper region abroad may be a worse choice for personal data, as you will see in lesson 4.

3.3

Edge, fog and mist

You tap your card on the bus validator and it beeps at once, even in a tunnel with no signal. Your phone unlocks with your face in a fraction of a second, also in airplane mode. Neither of them asked a data centre in Madrid or Frankfurt for permission. Why not, if the cloud is so powerful?

Because the cloud is far away, and distance costs time. A wind turbine that must stop when a gust becomes dangerous, a pump that must stop before a pipe bursts, a hotel door that must open when the internet is down: none of them can wait. Today you measure that distance with your own browser, and you learn where the decisions that cannot wait are taken.

What you will learn

- Measure **latency** to real data centres and explain why it limits what the cloud alone can do
- Describe **edge computing** and how it works **with** the cloud, not against it
- Define **fog** and **mist** and place the four layers in order, from the sensor to the cloud
- Decide, in a real plant, which task goes in which layer and why
- Avoid the trap: “edge replaces the cloud”

3.3.1 How far away is the cloud?

> **Latency** is the delay between sending a request and receiving its answer. It depends mainly on the distance and on the number of networks the data cross. It is measured in milliseconds (ms).

You can measure it right now. GCPing opens in the browser, with no account, and measures how long your computer takes to get an answer from each Google Cloud region. Look for Madrid (`europe-southwest1`), then for a region in the United States and one in Asia. The figures change every day and from one network to another, so trust the ones you see, not a figure in a book. The pattern is always the same: the further away, the longer it takes.

For a film, a few dozen milliseconds do not matter. For a machine that has to react, they can. And if the connection drops, the cloud is not slow: it is simply unreachable.

3.3.2 Edge computing: decide close to the data

> **Edge computing** means processing data **close to where they are produced** (in the device itself, in the same room or in the same building) instead of sending everything to a distant data centre.

In the city-builder game, the cloud is the utility's big power station at the other end of the map. You would never ask it to switch each traffic light. So you install **local substations** near the neighbourhoods, and they handle what must happen now.

	Edge	Cloud
Where	Next to the devices: a gateway, a small local server, the device itself	The provider's remote data centres
Latency	Very low	Higher: depends on distance and network
Power and capacity	Limited	Very large and expandable
Without connection	Can keep working	Cannot be reached
Data	Filters and summarises them; sends only what is needed	Receives the summary, stores it and analyses it in depth
Typical tasks	React at once, filter, decide locally	History, global analysis, training models, coordinating many sites

Why not send everything to the cloud? Four reasons: **latency**, **bandwidth** (sending everything costs money and network), **working offline**, and **privacy** (some data are better never leaving the building).

You already carry edge computing in your pocket. Apple says the Face ID data never leave the iPhone: the face is checked inside the phone. Your photos, on the other hand, sync with the cloud. Websites use edge too: a **CDN** (*content delivery network*) keeps copies of a site on servers spread around the world, so each visitor gets it from a nearby point. Your GitHub Pages site is served that way.

> **Careful: "edge replaces the cloud"**
> It does not. They share the work. The edge **reacts** and **filters**; the cloud **remembers**, **analyses** and **coordinates**, and sends back down what it has learnt: new rules, settings, updates. Fast decisions at the edge; smart decisions in the cloud.

3.3.3 Fog and mist: the layers in between

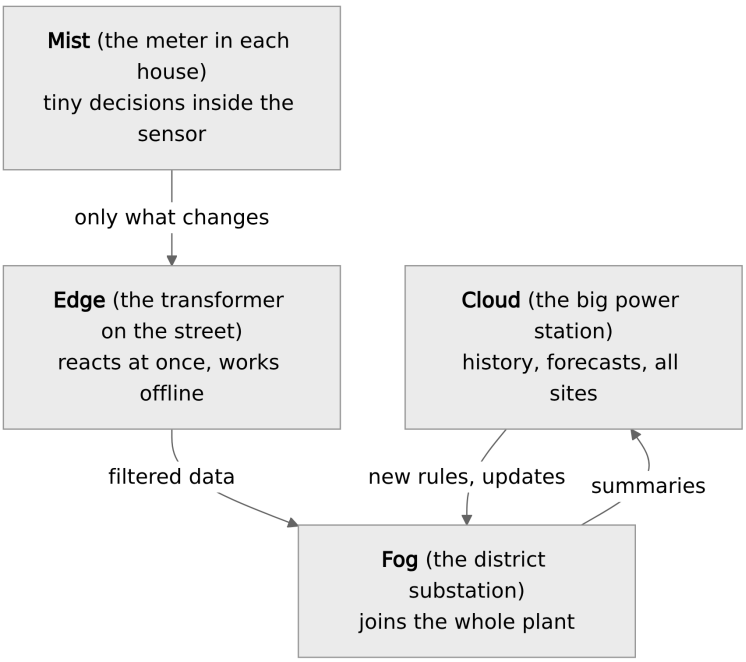
Between the sensor and the cloud there can be more than one step.

> **Fog computing** is an intermediate layer on the local network, between the edge devices and the cloud, that joins and processes the data of many points before deciding what goes up. **Mist computing** is the lightest layer, inside the sensors and actuators themselves, on very low-power microcontrollers.

A way to remember it: the cloud is high in the sky; **fog** is a cloud touching the ground; **mist** is even thinner and right at ground level. The lower you go, the closer to the sensor, the less power and the faster the answer.

The same four layers in three real kinds of plant (these are teaching examples, not the description of one particular installation):

Layer	Where it lives	Wind farm	Water utility	Hotel
Mist	In the sensors and actuators	A vibration sensor on a turbine only reports	A flow meter sends a reading every few min-	A room temperature sensor only reports changes



Layer	Where it lives	Wind farm	Water utility	Hotel
		when it crosses a threshold	utes, or at once if it jumps	
Edge	Next to the devices	Each turbine's controller turns the blades and stops it in dangerous wind	The pumping station's controller stops the pump if pressure spikes	Each floor's controller opens doors and regulates air conditioning
Fog	On the local network of the site	The farm's substation server joins all turbines	The control centre joins all stations of an area	The hotel's server joins all floors and manages the building's energy
Cloud	Remote data centres	History of all farms, production forecasts, predictive maintenance	Consumption of the whole city, leak detection by area, reports	The chain compares all its hotels and predicts occupancy

Different sources, different borders

The order **mist → edge → fog → cloud** is the convention of this module. Not all sources agree. NIST's *Fog Computing Conceptual Model* (SP 500-325, 2018) says there is no consensus: it places fog between the end devices and the cloud and mist as a lighter form of fog close to the devices, but it calls **edge** the layer of the end devices themselves. If you read it that way somewhere, you are not wrong. What matters is that you can order the layers and say **what each one does**.

3.3.4 In class: where do we process it?

Individually or in pairs. Use the plant of your project company (the OT side of D1). If it has none, choose one of the three in the table: a wind or solar farm, a town's water utility (for example EMAYA in Palma) or a hotel. **No generative AI**: this activity is assessed (class work, 30%).

Step	Do this	You should get
1	Measure latency with GCPing: note the fastest and the slowest region, with their milliseconds	Two figures and a date
2	**Guided research:** find **one real source** that says which equipment, sensors or control systems this plant (or one like it) has: company website, press release, annual report, public tender, manufacturer sheet	One citation: title, organisation, link, date consulted, and the sentence that says it
3	Write **six tasks** done with the plant's data (stop a machine, filter readings, join the whole plant's data, monthly report, predict breakdowns…)	Six tasks
4	Place each task in **mist, edge, fog or cloud** and give the **reason**: latency, working offline, data volume, privacy or computing power	A table of six rows
5	Name **one piece of data you would never send to the cloud**, and why	One sentence
6	Draw the four layers and the arrows of the data flow with diagrams.net or Mermaid Live (no account) and export it as an image	A diagram

Hand in on Classroom the citation, the table and the diagram.

3.3.5 For your project

This lesson feeds Deliverable 3 · Cloud architecture:

Today's work	Becomes	Criterion
The six tasks in their layers, with reasons and the real source	Section C: edge, fog and mist in the plant	c, d
The diagram of the four layers	The base of the group's architecture diagram	Group part

3.3.6 Key takeaways

- **Latency** is the delay between a request and its answer; it grows with distance, and without a connection the cloud cannot be reached at all.
- **Edge computing** processes data close to where they are produced, to react fast, work offline, save bandwidth and protect privacy.
- Edge and cloud **share the work**: the edge reacts and filters; the cloud remembers, analyses and coordinates.
- **Fog** is the intermediate layer that joins a whole site; **mist** lives inside the sensors themselves.
- In this module, from the sensor up: **mist → edge → fog → cloud**. Sources draw the borders differently; what counts is what each layer does.

3.3.7 Quick check

1 What is latency?

A The amount of data a connection can carry per second
B The delay between sending a request and receiving its answer
C The time a server stays switched on
D The distance in kilometres to the data centre

2 What is edge computing?

- **A** Moving all processing to the biggest data centre available
- **B** Processing data close to where they are produced
- **C** A type of SaaS for small companies
- **D** Storing backups on the border between two countries

3 A wind turbine must stop at once if a gust becomes dangerous, even if the internet is down. Where should that decision be taken?

- **A** In the cloud, which has more computing power
- **B** At the edge, in the turbine's own controller
- **C** In the provider's data centre closest to the farm
- **D** It does not matter, as long as the data are saved

4 Following this module's convention, which is the order from closest to the sensor to furthest away?

- **A** Edge → mist → cloud → fog
- **B** Fog → mist → edge → cloud
- **C** Mist → edge → fog → cloud
- **D** Cloud → fog → edge → mist

5 Which statement is true?

- **A** Edge computing replaces the cloud completely
- **B** Filtering data at the edge reduces the volume sent to the cloud
- **C** Mist computing runs in the provider's data centre
- **D** Fog computing is the same as a CDN

6 Which task fits best in the cloud?

- **A** Opening a hotel room door when the guest taps the card
- **B** Stopping a pump when pressure spikes
- **C** Comparing the energy use of all the hotels of a chain over five years
- **D** Sending a temperature reading only when it changes

7 A supermarket chain's cash registers must keep selling when the connection fails, and head office wants to know the sales of every shop each night. Which design fits?

- **A** Everything in the cloud, because it is cheaper
- **B** Everything in each shop, with no cloud at all
- **C** Each shop sells and records locally (edge) and sends the day's summary to the cloud for head office
- **D** Every sale waits for the cloud to approve it

You now know where each decision is taken. One question is left: whose power station is it, and which law applies to what it stores? Next: Advantages and sovereignty.

Solutions

1. B · **Latency** is a delay, measured in milliseconds. It depends mostly on distance and the networks crossed; how much data fits per second is **bandwidth**.

2. B · **Edge computing** processes data near their source (device, room or building) so the answer arrives fast and works without a connection.

3. B · The decision cannot wait for the network, and must work **offline**: it belongs at the **edge**.

4. C · **Mist** inside the sensor, **edge** next to the devices, **fog** on the site's local network, **cloud** far away.

5. B · The edge **filters** and **summarises**, so far less data travel up. The cloud is still needed to store, analyse and coordinate.

6. C · Long history and many sites need **storage and computing power**: that is the cloud's job. The other three are edge or mist tasks.

7. C · The same split as the wind farm: **react locally** so the shop never stops, and **send summaries** up so the cloud can compare all shops.

3.4

Advantages and sovereignty

Where is your Classroom homework right now: in Madrid, in Ireland, in Iowa? Who decided it, you or the school? And if a judge in another country asked the provider for it, would the provider have to hand it over?

Most people never ask. A company has to. When a hotel chain in Palma puts its guests' passports and card details in the cloud, it is still legally responsible for them, wherever they end up. Today you weigh what the cloud really brings to a connected company, what it risks, and you do the research a careful IT team does before signing: find out where a real service keeps its data and which law applies.

What you will learn

- Identify the **advantages** the cloud brings to **connected systems**, and its main risks
- Explain **data sovereignty**: knowing and controlling where data are, who can reach them and which law applies
- Recognise the rules that decide it: the **RGPD** and its transfer rules, the US **CLOUD Act** and the EU **Data Act**
- Find out, with public sources, where a real service keeps its data and whether they leave the EU
- Avoid the trap: "the IP address of a website tells you where the data are stored"

3.4.1 What the cloud brings to connected systems

> A **connected system** is a set of devices, applications and people that exchange data over a network: the sensors and SCADA of Unit 1, the IoT of Unit 2, the wind turbines of the last lesson.

Advantage	What it means	Example
One common point for data	Every device and site sends its data to the same place.	The data of all the hotels of a chain, together and comparable
Scalability and elasticity	Grow from ten devices to ten thousand without buying servers; shrink when they are not needed.	More capacity in August, less in January
Access from anywhere	Data and applications are reachable from any device with a connection.	The manager checks the wind farm's production on her phone
Cost that follows use	No big initial investment: spending follows activity (your estimate from lesson 2).	Opening a new hotel without a server room
Availability and copies	Providers have several data centres and let you replicate data between them.	If one server fails, bookings continue from another

Advantage	What it means	Example
Remote management	Devices are controlled and updated from a distance, all at once.	Updating the software of every door lock without visiting them one by one
Power to analyse	Computing capacity for big data and artificial intelligence.	Predicting breakdowns from years of vibration data
Integration	Services from different companies connect easily through APIs.	Bookings linked to travel agencies and a payment gateway

And the risks to keep an eye on:

Risk	What it means	How it is reduced
Depending on the connection	No internet, no cloud.	Edge processing for what cannot stop (last lesson)
Vendor lock-in	Moving to another provider becomes so hard or expensive that you are stuck.	Open formats, multicloud, and the EU **Data Act** (below)
Costs that grow	Paying per use also means paying for what is left switched on.	Budgets, alerts, reviewing the bill every month
Shared security	The provider protects the infrastructure; the configuration and the data are the customer's job.	Know which layers are yours (lesson 1)
Where the data are and who can reach them	Another country, another law.	The rest of this page

3.4.2 Data sovereignty

> **Data sovereignty** means that data are subject to the laws of the place where they are stored and of the company that holds them. Having sovereignty over your data means knowing and controlling **where they are, who can reach them and which law applies**.

In the city-builder game, renting electricity is comfortable. But before signing you want to know where the power station is, which country the company belongs to, and whether you can switch to another one without paying a fine. Four questions, four answers:

Question	What you need to know
Where are they?	Providers group their data centres in **regions**, and the customer chooses the region. The three big ones have regions in Spain: AWS in Aragon (`eu-south-2`, opened in 2022), Google Cloud in Madrid (`europe-southwest1`, 2022) and Microsoft Azure in Madrid (`Spain Central`, 2024). Hetzner, a German provider, has its data centres in Germany and Finland, plus the USA and Singapore.
Can they leave the EU?	The **RGPD** (Regulation (EU) 2016/679, Chapter V, articles 44 to 49) only lets personal data leave the European Economic Area with safeguards: an **adequacy decision** of the European Commission (the destination protects data as well as the EU), **standard contractual clauses**, or specific exceptions. For the United States there is the **EU–US Data Privacy Framework**, an adequacy decision of 10 July 2023, upheld by the EU General Court on 3 September 2025; an appeal is pending before the Court of Justice.
Whose is the provider?	A US company is subject to US law even if its data centre is in Europe. The **CLOUD Act** (2018) obliges it to hand over data under its control when US authorities request them through legal process, wherever the data are stored. That is why "sovereign" offers exist, such as the **AWS European Sovereign Cloud**, whose first region opened in Brandenburg, Germany, in January 2026, run by companies set up in the EU.

Question	What you need to know
Can I switch provider?	The EU **Data Act** (Regulation (EU) 2023/2854) has applied since 12 September 2025. It obliges cloud providers to make switching to another provider possible, and from **12 January 2027** they may no longer charge **switching charges**. It is Europe's answer to vendor lock-in.

The company that collects personal data stays **responsible** for them even if a provider stores them. Choosing a provider and a region is a **legal** decision, not only a technical one. The rules on personal data are studied in depth in Unit 5.

> **Careful: "switching is now free"**
> The Data Act bans the **switching charges** themselves. The provider can still bill the normal service while you use it, and some costs of running two clouds in parallel. Read it as "they cannot fine you for leaving", not "everything is free".

3.4.3 Where does a website answer from?

There is a quick technical check you can do from the browser, with no account: look up a domain's IP address and ask who owns that address.

Tool	What it does
Google Admin Toolbox Dig	Asks the **DNS** (the internet's phone book) for a domain's address. Choose the `A` record and copy the IP
ipinfo.io	Opening `ipinfo.io/<IP>` shows the organisation that owns the address and an approximate city

Try it with your own GitHub Pages site. On 25/10/2026, a `github.io` address showed the organisation **Fastly** (a CDN), the city **San Francisco**, and the label `anycast: true`. Yet the page reached Mallorca in milliseconds. **Anycast** means the same address is announced from many places at once, and you are served by the closest copy: the CDN of the last lesson, which is edge.

> **Careful: the IP address is not where the data live**
> A website is often served from a CDN or an anycast address, so the IP points to a nearby cache or to the CDN company's head office, not to where the data are stored. If ipinfo says "Madrid", it does **not** prove the data are in Madrid. What the provider **writes** in its privacy policy and contracts is what counts; the IP check only tells you who delivers the page.

3.4.4 In class: where do the data live?

Individually. Research **two services**: one you use in class (Kahoot, Google Classroom or GitHub) and one your company uses or that you proposed in D3 (the provider of your estimate, the company's website...). **No generative AI**: this activity is assessed (class work, 30%). **Passive only:** you only read what any visitor can see; never test passwords or scan anything.

Step	Do this	Where to look
1	**Who is the provider**, and in which country is its head office?	The legal notice or privacy policy

Step	Do this	Where to look
2	**Where does it say it stores the data**: country, region or data centre?	Privacy policy, “data location” or “data residency” page, list of **subprocessors** (other companies that handle data for it)
3	**Do data leave the EEA**, and with which safeguard (adequacy, Data Privacy Framework, standard clauses)?	The “international transfers” section of the privacy policy
4	**Where does the website answer from?** Look up the domain in Dig (`A` record), copy the IP and open `ipinfo.io/<IP>`. Note the organisation and city	Dig and ipinfo.io
5	**Conclusion** in three sentences: is it a good option for your company’s data? Which region or provider would you choose, and why?	What you found

For every step, the citation: link, date consulted, and the sentence copied in quotation marks. If the policy does not say, write “it does not say” and where you looked: that is also a finding.

Then: your group’s website

Then, in your project group, the **Technician** sets up the site where the group will publish Deliverable 3. Everybody watches and takes part.

Step	Do this	You should see
1	The Technician creates a **public** repository called `dossier-<company>` (no spaces or accents), with **Add README**	The new repository
2	**Settings** → **Collaborators** → **Add people**: invite the other members by username. Each one accepts the invitation (email or GitHub notification)	Everybody can edit
3	Create `index.md` with the skeleton of D3: title, company name, one heading per section (A, B, C, D) and one for the group part. Each member writes one line under their heading and commits it	The history shows who wrote what
4	**Settings** → **Pages** → **Deploy from a branch** → `main`, `/ (root)` → **Save**	The site at `https://<technician-username>.github.io/dossier-<company>/`

GitHub Pages turns `.md` files into web pages by itself, so you can write in Markdown. Mermaid diagrams are **not** drawn on the published site: add the architecture diagram as an **image** (PNG). The repository is public: no personal data, first names only if you prefer. The Coordinator hands in the two links on Classroom. If nobody in the group has a GitHub account, make a Google Sites site with your school accounts.

3.4.5 For your project

This lesson closes the research for Deliverable 3 · Cloud architecture, due **Monday 8 February 2027 at 23:59**, the same day as your defence in session 14.

Today's work	Becomes	Criterion
Where the data of your company's provider live, and your choice of region	Section D: deployment model and sovereignty (with 3 members, group part)	e
Advantages and risks for your company	Sections B and D, and the group's conclusions	e
The group's repository and site	The published dossier (group part)	Group part

3.4.6 Key takeaways

- For **connected systems**, the cloud brings one common point for data, elasticity, access from anywhere, cost that follows use, availability, remote management, analysis power and integration.
- Its risks: depending on the connection, **vendor lock-in**, growing costs, shared security, and where the data are.
- **Data sovereignty** is knowing and controlling where data are, who can reach them and which law applies. The customer stays responsible for its data.
- The **RGPD** only lets personal data leave the EEA with safeguards such as an **adequacy decision**; the US **CLOUD Act** reaches US providers wherever their data centres are.
- The EU **Data Act** obliges providers to allow switching, with no switching charges from 12 January 2027.
- An IP lookup tells you who **delivers** a page (often a CDN with **anycast**), not where the data are **stored**.

3.4.7 Quick check

1 What is data sovereignty?

A The provider's right to use its customers' data
B Knowing and controlling where data are, who can reach them and which law applies
C Storing all data on the company's own servers
D Encrypting data so that nobody can read them

2 A hotel chain moves its guest records to a cloud provider. Who is responsible for those personal data?

A Only the provider, from the moment it stores them
B The hotel chain, which remains responsible even though a provider stores them
C Nobody, because the data are in the cloud
D The guests themselves

3 Under the RGPD, when can personal data be sent to a country outside the European Economic Area?

A Never, in any case
B Always, if the provider is a large company
C Only with safeguards, such as an adequacy decision of the European Commission or standard contractual clauses
D Only if the data are less than one year old

4 A US provider keeps a company's data in a data centre in Madrid. Which statement is true?

- **A** Only Spanish law applies, because the data centre is in Spain
- **B** US law can also reach those data, because the CLOUD Act applies to US providers wherever the data are stored
- **C** No law applies to data in the cloud
- **D** The data automatically move to the United States

5 What does the EU Data Act change for a company that wants to leave its cloud provider?

- **A** Nothing: it only covers mobile phones
- **B** It forbids companies from changing provider more than once a year
- **C** Providers must make switching possible, and from 12 January 2027 they cannot charge switching charges
- **D** It makes every cloud service free of charge

6 ipinfo.io says a website's IP address is in Madrid. What can you conclude?

- **A** The website's data are stored in Madrid
- **B** The company is Spanish
- **C** Only who delivers the page and roughly from where; the data could be stored somewhere else, so you must read the privacy policy
- **D** The website does not use a CDN

7 A small dental clinic in Manacor wants online appointments for its patients. Which choice protects its patients' data best, with the least effort?

- **A** The cheapest service found, whatever the country, without reading its policy
- **B** A SaaS whose policy says it stores data in the EU, after checking its transfers section
- **C** Its own server in the waiting room, so that nobody else sees the data
- **D** A spreadsheet shared publicly so patients can book themselves

You now have everything for your cloud architecture: what to rent, from whom, at what cost, what stays local and where the data live. Put it together in Deliverable 3 · Cloud architecture.

Solutions

1. B · **Data sovereignty** is about **location, access and law**. Encryption and own servers can help, but they are tools, not the definition.

2. B · Under the **RGPD** the company that collects the data stays **responsible**. The provider processes them on its behalf, under contract.

3. C · **Chapter V** of the RGPD allows transfers outside the EEA only with safeguards: an **adequacy decision**, **standard contractual clauses** or specific exceptions.

4. B · Location matters, but so does **who the provider is**. The **CLOUD Act** obliges US providers to hand over data under their control when legally requested, wherever they are stored.

5. C · The **Data Act** (Regulation (EU) 2023/2854) attacks **vendor lock-in**: switching must be possible and, from 12 January 2027, without switching charges. The normal service is still paid.

6. C · Many sites are served through a **CDN** or an **anycast** address. The IP tells you who **delivers** the page, not where data are **stored**: the provider's written policy is the evidence that counts.

7. B · A **SaaS** saves the clinic administration work, and checking the **region** and **transfers** in the provider's policy is the sovereignty homework. A server in the waiting room is not safer just because it is close.

3.5

Deliverable 3 · Cloud architecture

In a city-builder game, the moment comes when you stop building every plant yourself. You sign contracts with the utilities, decide which neighbourhoods get their own substation, check the monthly bill and read the small print: where is the power station, and can you switch company if the service gets worse? That is a cloud architecture.

Deliverable 3 is that plan for your company. Your group decides which of its systems go to the cloud and at which **service model**, what the cloud will do for it and **how much it will cost** with a real provider's calculator, what the plant must keep deciding locally at the **edge**, **fog** and **mist**, and **where the data will live**. And this time the result is not only a PDF: your group publishes its dossier as a real website on **GitHub Pages**. It is the third chapter of your digital transformation dossier.

What you will learn

- Choose and justify a **service model** and a real provider for the systems of a real company
- Produce a dated, reproducible **cost estimate** with a provider's pricing calculator, and read its limits
- Split a real plant's work between **mist**, **edge**, **fog** and **cloud**, with reasons
- Choose a **deployment model** and decide where the data live, with the law that applies
- Publish the group's dossier as a static website, and defend your own section alone

3.5.1 What you hand in

Format	The group's website on GitHub Pages and one PDF of the same content, about 6–10 pages without the cover. The PDF includes the link to the website
Where	Google Classroom, assignment "D3 · Cloud architecture". The Coordinator hands it in
Deadline	**Monday 8 February 2027, 23:59**
Defence	Individual, in **session 14, Monday 8 February 2027** (the same day)
Weight	70% of the RA3 mark. The other 30% are the class activities of sessions 10 to 13

Every fact about your company needs a source

Everything you state as a fact about your company (its systems, its plant, the providers it already uses) must be **found and cited**: its website, annual report, the press, job offers, public tenders, Ibestat, the BOIB. Each citation has a title, the author or website,

the link and the date you consulted it. What you cannot find, you may still propose with reasons, but you label it **hypothesis**. The same applies to providers: every claim about a service (what it offers, where it stores data) cites the provider's own page. Prices always carry the **date** of the estimate.

No generative AI

Generative AI is **not allowed** in this deliverable. A section written with AI or copied scores **0**, and the defence will show it: you cannot explain a text you did not write.

The website is public

With a free GitHub account the repository must be public, so anybody can read your site. Publish no personal data: no phone numbers, no addresses, and first names only if you prefer (your full names go in the PDF on Classroom). Nothing confidential about the company: only what you found in public sources.

3.5.2 The sections

Each member **owns and signs one section** (the individual part). With 3 members, Section D becomes part of the group part and everybody writes it together.

Section	Criterion	What to write	Owner
A · Service models for three systems	RA3 a)	**Three systems** of the company (for example the website, the bookings or orders, the ERP, email, backups, plant supervision). For each: what it does, the **service model** you propose (IaaS, PaaS or SaaS) and **why**, saying what the company would manage and what the provider would. A **real provider and product** for each, with a link to its official page and the sentence that proves the model. At least **two different models** among the three.	Individual
B · Cloud functions and real cost estimate	RA3 b), e)	The **functions of the cloud** in this company: at least one example each of processing data, exchanging information and running applications. Then a **real estimate** for one system with a provider's pricing calculator (AWS, Azure, Google Cloud) or price list (Hetzner): the configuration (region, size, hours, GB, transfer) as a labelled hypothesis, the monthly total, the **date**, and what it does **not** include (VAT, people's work...). **Evidence**: the exported PDF or a screenshot with configuration and total, and the share link if there is one. A **second estimate** in another region or with another provider, and a short conclusion: which one and why, and what advantage the pay-per-use model gives this company.	Individual
C · Edge, fog and mist in the plant	RA3 c), d)	The company's plant (the OT side of D1). A **real cited source** on its equipment or control systems, or on a plant of the same kind.	Individual

Section	Criterion	What to write	Owner
		A latency measurement to two cloud regions (GCPing or similar) with its date. A table of at least **six tasks** done with the plant's data, each placed in **mist, edge, fog or cloud** with its **reason** (latency, working offline, data volume, privacy, computing power). What happens if the internet connection drops. **One piece of data** that would never go to the cloud, and why. A diagram of the four layers.	
D · Deployment model and data sovereignty	RA3 e)	The **deployment model** (public, private, hybrid, multicloud) for the whole architecture, justified with the company's size, data, existing systems and budget. **Where the data will live**: provider, country of its head office, region, and what its privacy policy says about **transfers outside the EEA** (cited, with date). Whether the **CLOUD Act** could reach them, and how the **Data Act** protects the company if it wants to switch. **Three advantages** of the cloud for this company's connected systems and **two risks** with how you would reduce them.	Individual (4th member)
Introduction	RA3 a)–e)	The company in five lines (from D1 and D2) and what this architecture changes for it: which systems move, which stay.	Group
Architecture diagram	RA3 a)–e)	One diagram that joins the four sections: the plant's devices, the mist, edge and fog layers of C, the cloud services of A with their provider and region, and the arrows of the data flow. Exported as an image (PNG), so that it shows on the website.	Group
The published website	RA3 a), b)	The whole dossier published on **GitHub Pages** (or Google Sites if nobody in the group can use GitHub): one page or several, with the four sections, the diagram and the sources, the same version as the PDF. In the PDF, one paragraph on how you published it and **which service model GitHub Pages is**.	Group
Integration and conclusions	RA3 a)–e)	The four sections must agree: the provider of A is the one of B's estimate and D's region; the plant of C is the one in the diagram. Half a page of conclusions.	Group
Sources	All	One list of every source cited, in the same format, with consultation dates.	Group
Group log	All	Who owned which section, who had which role, what you decided and when (template on the project page).	Group

3.5.3 Roles in this deliverable

Everybody changes role and section from Deliverable 2. Write them in the group log.

Role	In Deliverable 3
Coordinator	Splits the sections, sets an internal deadline (we suggest Thursday 4 February 2027), checks the group log and hands in the PDF and the link

Role	In Deliverable 3
Source checker	Opens every link, checks every fact has a source and every guess is labelled hypothesis, and that every price has its date and its evidence
Technician	Creates the repository, invites the others, publishes the website, exports the diagram as an image and checks the site opens from another computer
Editor	Format, spelling, one voice; the same name for the same system and provider everywhere; the PDF matches the website. With 3 members, the Coordinator does it

3.5.4 How it is assessed

Part	Weight	What is assessed
Your section	40%	Content correct for its criterion, real and cited sources, dated evidence, hypotheses labelled, reasons given, clear writing
Your defence	20%	You explain your section in your own words, answer why, and know how it connects to another member's section
Group part	40%	Introduction, architecture diagram, the published website, coherence between sections, conclusions, sources list, format and spelling, group log

Rubric

Part	Excellent	Adequate	Insufficient
Your section	Everything the table asks for, about this company; every choice justified; real providers and sources cited; estimates dated and reproducible; terms used precisely	Most of what is asked; some generic reasons, a missing date or a source without the exact sentence; the main ideas are correct	Missing parts; generic text that could be about any company; models or layers confused; prices without date or evidence
Your defence	Answers fluently and with reasons; can say where each fact and figure comes from; explains how a teammate's section links to yours	Answers the main questions with some help; knows the sources of the key facts	Cannot explain parts of the section, or does not know where its content comes from
Group part	One coherent architecture: same providers, region and plant in every section and in the diagram; the website works and matches the PDF; clean format, one voice, complete log	Small contradictions between sections; the website works but is incomplete or differs from the PDF; minor format problems	Four separate texts glued together; no diagram or no website; no log; many spelling or format errors

Programme rules: RA3 needs at least a 5. Late work is marked out of 5; more than 48 hours late is a 0. Plagiarism or unauthorised AI is a 0. If you fail RA3, the recovery is an individual task: you redo one section of this deliverable alone and defend it, for a maximum of 5.

3.5.5 The defence

On **Monday 8 February 2027** (session 14) the groups finish their deliverable while the teacher calls each student for **2–3 minutes**. No paper and no computer: a conversation about your section as it is that day.

- **Two questions about your own section.** For example: "Why SaaS and not PaaS for this system?", "Which is the most expensive line of your estimate, and how would you reduce it?", "What happens in your plant if the internet goes down?", "Where would the data be stored, and which law would apply?"

- **One question about another member's section** or the group part. For example: "Follow one piece of data through your diagram, from the sensor to the cloud", or "Which service model is GitHub Pages?"

If you miss the session with a justified absence, your defence is moved to the start of session 15; tell the teacher as soon as you can.

3.5.6 Before you hand in

Check	Who
Every section has its owner's name	Coordinator
Every fact about the company has a working source; every guess says hypothesis	Source checker
Every provider claim cites the provider's own page	Source checker
Every estimate has its date, its configuration and its evidence (PDF, screenshot or link)	Source checker
Section A uses at least two different service models	Editor
The provider of A, the estimate of B and the region of D agree	Editor
The diagram shows the four layers and the cloud services, and is an image on the website	Technician
The website opens from a computer outside the group, and matches the PDF	Technician
The group log is filled in and signed	Coordinator
One PDF with the website link, on Classroom before Monday 8 February 2027 at 23:59	Coordinator

Start from what you did in class: the activities of Cloud and service models, Deployment and functions, Edge, fog and mist and Advantages and sovereignty are the first drafts of sections A to D.



04

Unit 4 · Artificial intelligence

In a city-builder game such as *Cities: Skylines*, there comes a point where you stop changing every traffic light by hand. You switch on an **automatic advisor**: it watches the traffic overlay, learns when the jams happen and tunes the lights by itself. It is not magic and it is not always right. It is only as good as the data it watches, it costs money to run, and somebody still has to check what it does. That advisor is **artificial intelligence**. In this unit you train one in your browser, test a generative AI critically, calculate whether an AI project pays for itself, and find out what AI is really made of, and what it changes in your own future job.

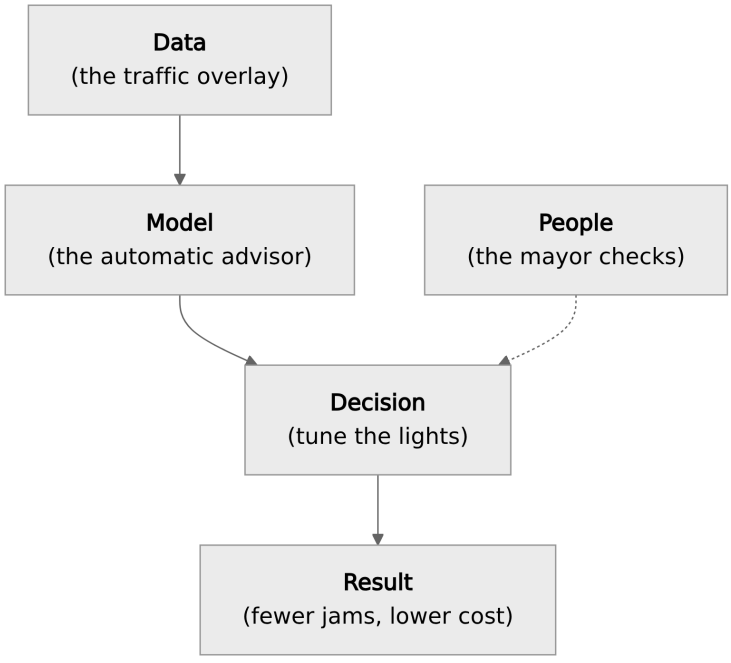
What you will learn

- Explain what **AI**, **machine learning** and **deep learning** are, and avoid the trap of calling everything automatic "AI"
- Train an image classifier in the browser and explain why it fails
- Use a **generative AI** critically: check its answers against sources and spot **hallucinations**
- Classify AI uses by the risk levels of the European **AI Act**
- Find real AI uses in Balearic companies and estimate, in a spreadsheet, whether an AI project is **profitable**
- Name the languages and tools AI is built with, and describe how AI is changing software development

The model is the part everybody talks about. This unit is just as much about the other boxes: where the data come from, who checks the decision, and whether the result is worth the cost.

Lessons

#	Lesson	What it covers
1	What is AI	AI, machine learning and deep learning; training and inference; your first model with Teachable Machine



#	Lesson	What it covers
2	Generative AI and responsible use	Large language models, hallucinations, rules for responsible use and the risk levels of the AI Act
3	Automation and profitability	Rules vs AI, automating vs optimising, a real Balearic case, and a profitability spreadsheet
4	Sectors, languages and software	Which sectors use AI most (Eurostat), the languages behind real AI projects on GitHub, AI in software development
5	Deliverable 4 · AI plan	What you hand in, who writes what, how it is assessed, the defence

4.1

What is AI

This morning your phone unlocked when it saw your face. Your keyboard guessed the next word before you typed it. Your music app put on a song you had never heard, and you liked it. Nobody at Apple, Google or Spotify wrote a rule that says “if it is Marta with her hair up, unlock”. So how does the phone know?

That question is the start of this unit. Companies in Mallorca use the same kind of technology to guess how full a hotel will be next month or to sort thousands of support requests. Before you can say where AI helps a company, and where it is only hype, you need to know what is actually inside. Today you will not only read about it: you will train one yourself.

What you will learn

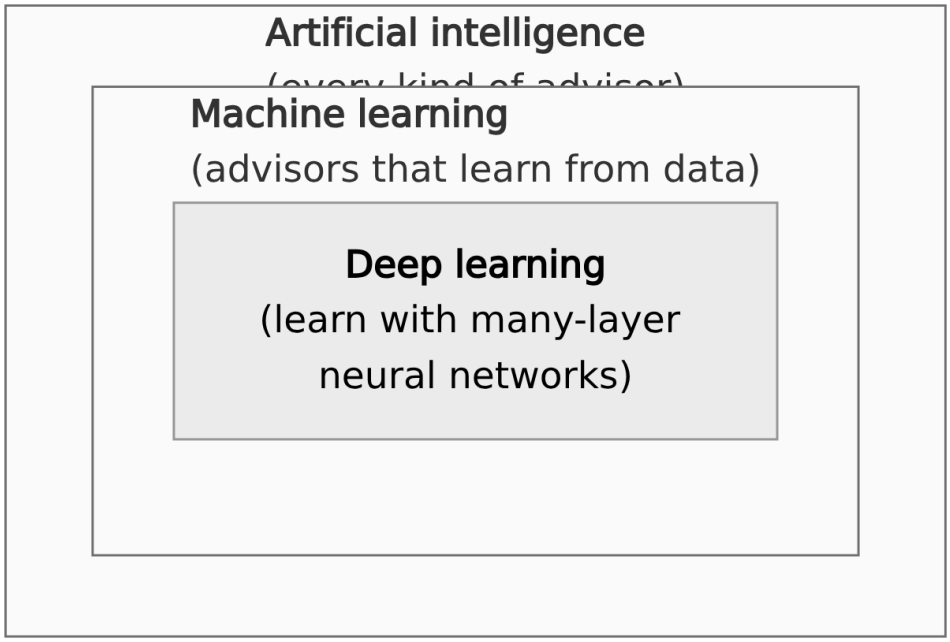
- Define **artificial intelligence** and tell it apart from a program that follows fixed rules
- Place **machine learning** and **deep learning** inside AI, and name the three ways a machine can learn
- Tell **training** from **inference**
- Train an image classifier in the browser with Teachable Machine and explain why it fails
- Avoid the trap of the day: not everything automatic is AI, and a model is only as good as its data

4.1.1 Three circles, one inside the other

> **Artificial intelligence (AI)** is the field of building systems that do tasks we associate with human intelligence: perceiving, reasoning, deciding, generating.

In the city-builder game of this course, AI is the **automatic advisor** that tunes the traffic lights. Some advisors follow rules written by a person; others learn from what they see. That difference gives three circles:

Concept	What it is	Example
Artificial intelligence (AI)	Any system that does an “intelligent” task. Includes rule systems that **do not learn**	A route planner; a medical expert system with rules written by doctors
Machine learning (ML)	Part of AI: the system **learns patterns from data** instead of receiving the rules	A spam filter trained with emails already marked as spam or not
Deep learning (DL)	Part of ML: learning with **neural networks** that have many layers. Good with images, sound and text; needs a lot of data and computing power	Face unlock, automatic translation, chatbots



> A **neural network** is a program made of many small connected units arranged in layers. Each connection has a number (a weight) that gets adjusted while it learns. “Deep” means many layers.

4.1.2 Rules or examples?

This is the key difference, and it is worth drawing:

In classic programming, you write `if temperature < 20 then heating on`. In machine learning, you give thousands of examples with their answers, and the algorithm finds the rules by itself. The result is a **model**: a file full of numbers that captures what it learned.

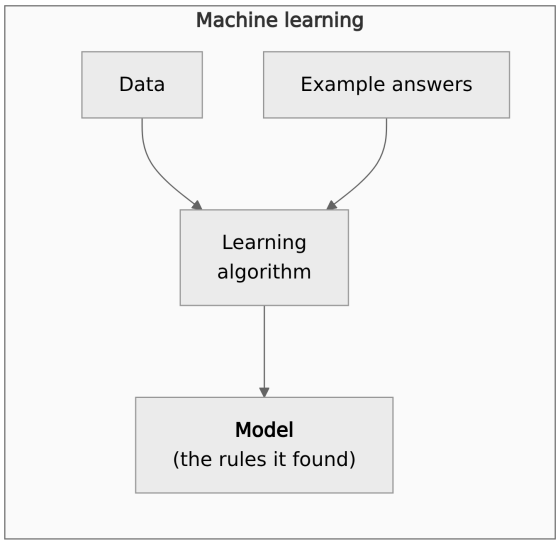
> **Training** is adjusting the model with example data. It is slow and expensive, and it happens a few times. **Inference** is using the trained model on new data to get an answer. It is fast, and it happens millions of times.

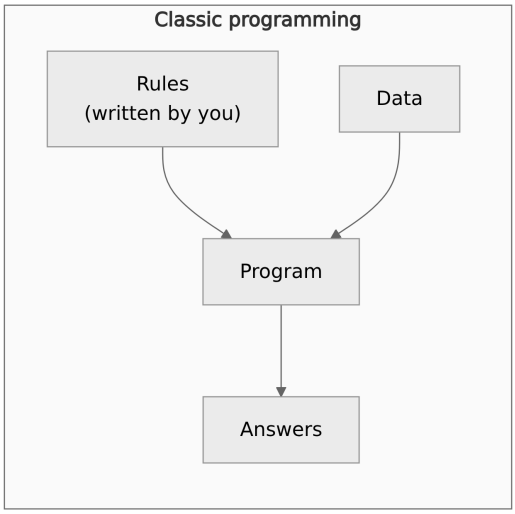
When your phone learned your face, that was training. Every time it unlocks, that is inference.

4.1.3 Three ways to learn

Type	How it learns	Example
Supervised learning	From **labelled** examples: input plus the correct answer. It classifies or predicts a number	Spam yes/no; next month’s occupancy of a hotel in Calvià
Unsupervised learning	**No labels**: it looks for groups or structure in the data	Grouping the customers of an online shop into profiles
Reinforcement learning	It tries actions and receives rewards or penalties	A program that learns to play chess; a robot that learns to walk

> **Careful: “everything automatic is AI”**





> A thermostat that turns on the heating below 20 °C is **automation**, not AI: a person wrote the rule and it never learns. A program that calculates VAT is classic programming. Calling these "AI" is the most common mistake in company brochures, and in student projects. Ask: *did it learn from data, or did someone write the rule?*

4.1.4 What AI is not

Two ideas will come back in every lesson of this unit.

A model is only as good as its data. If the examples are few, unbalanced or unfair, the model learns that. This is called **bias**: a model trained mostly with photos of one kind of person recognises that kind of person better. It is not malice; it is the data.

Today's AI is narrow. A model that recognises faces cannot translate a sentence; a translator cannot drive. Each model does one task very well and nothing else. It does not "think" in general.

Why is AI everywhere now, if the ideas are decades old? Three things came together: **lots of data** (everything is digital), **lots of computing power** (graphics cards and the cloud you saw in Unit 3) and **better algorithms** (deep neural networks).

4.1.5 In class: train your first model

You will train an **image classifier** with **Google Teachable Machine**. It runs in the browser, it needs no installation, and you do not need to sign in to train or test. The training happens on your own computer: the images do not leave it unless you choose to save the project to Google Drive.

The task: a quality check in a factory. Is a part **correct** or **defective**? (A pen with or without its cap, a mug clean or with a paper mark.)

Step	Do this	You should see
1	Open teachablemachine.withgoogle.com → *Get Started* → *Image Project* → *Standard image model*. Name the two classes "correct" and "defective"	Two empty classes
2	Add about **30 images** to each class: hold *Webcam* pressed, or *Upload* the photos your teacher shared on Classroom	A counter of examples in each class
3	Click *Train Model* and do not close the tab	A progress bar, then the *Preview* panel
4	**Inference**: test the model with **10 new cases** (5 per class) that were not used for training. Write down the class it says and its confidence	A table of 10 tests with hits and misses
5	**Break it on purpose**: change the background, the light or the distance, or show an object that is in neither class	Wrong answers given with high confidence
6	**Improve it**: add 20 more varied examples (other backgrounds and light) per class, train again and repeat the 10 tests	More hits, or not: the data decide

Do not record faces or people: this is a practice with objects.

Hand in on Classroom (individual, "S15 · My first model", by 23:59 the same day): the classes and number of examples with a screenshot; the table of 10 tests and the hit rate before and after improving; the clearest mistake and **why** you think it happened; one thing this model could **never** do even with more training; and two lines on a process of your project company where an image or sound classifier could help. **No generative AI** in this hand-in.

4.1.6 For your project

This lesson feeds Deliverable 4 · AI plan, due **Monday 22 March 2027 at 23:59**:

Today's work	Becomes	Criterion
Your trained model and its test table	Section D: the working in-browser demo	e
The process where a classifier could help	A first idea for Section A: processes to automate	a
Present and future: why AI is everywhere now, and its limits	Section C: the AI landscape of the sector	c

4.1.7 Key takeaways

- **AI** is the big circle; **machine learning** learns from data; **deep learning** learns with many-layer **neural networks**.
- Classic programming: you write the rules. Machine learning: you give examples and the algorithm finds the rules, stored in a **model**.

- **Training** adjusts the model (slow, rare); **inference** uses it on new data (fast, constant).
- **Supervised** (labelled examples), **unsupervised** (find groups), **reinforcement** (rewards and penalties).
- A model is only as good as its data: **bias** in, bias out. Today's AI is **narrow**.
- Not everything automatic is AI: a fixed rule is automation.

4.1.8 Quick check

1 Which statement describes the relationship between AI, machine learning and deep learning?

A They are three different technologies with nothing in common
B AI is part of machine learning, which is part of deep learning
C Deep learning is part of machine learning, and machine learning is part of AI
D Machine learning and deep learning are two names for the same thing

2 A thermostat turns on the heating whenever the temperature drops below 20 °C. Is it machine learning?

A Yes, because it decides by itself
B No: a person wrote the rule and it never learns from data
C Yes, it is reinforcement learning
D Only if it is connected to the internet

3 Your Teachable Machine model, already trained, says a new photo is "defective" with 92% confidence. What is happening at that moment?

A Training
B Inference
C Unsupervised learning
D Labelling

4 A shop wants to group its customers into profiles, but nobody has told the system what the profiles are. Which kind of learning fits?

A Supervised
B Unsupervised
C Reinforcement
D None: this cannot be done with AI

5 You trained a classifier with all photos taken on the same white desk. On a wooden desk it fails. What is the most likely reason?

A The model is broken and must be reinstalled
B Deep learning does not work with photos
C The model may have learned the background as well as the object, because all the examples shared it
D The model needs an internet connection to see colours

6 Which two factors explain why deep learning took off in recent years?

A Huge amounts of digital data
B Programmers stopped writing rules
C Much more computing power, such as graphics cards rented in the cloud
D AI became able to think in general, like a person

7 A hotel in Calvià wants to predict how many rooms will be occupied each day next month from its past bookings, flights and holidays. How would you describe the system?

- **A** A fixed rule written by the manager
- **B** Supervised machine learning that predicts a number, trained on historical data
- **C** Reinforcement learning, because it gets a reward
- **D** Not AI, because it uses numbers instead of images

You now know what is inside the advisor and you have trained one. Next, the kind of AI everybody is talking about: models that write. Next: Generative AI and responsible use.

Solutions

1. C · Three circles, one inside the other: **AI** is the largest, **machine learning** is the part that learns from data, and **deep learning** is the part of machine learning that uses neural networks with many layers.

2. B · It is **automation** with a fixed rule. Machine learning needs a model that learned its rules from **data**. Deciding "by itself" is not enough.

3. B · Using a trained model on **new data** to get an answer is **inference**. Training was when you clicked *Train Model* with your labelled examples.

4. B · With **no labels**, the system looks for groups by itself: that is **unsupervised learning**. Supervised learning would need examples already labelled with their profile.

5. C · A model learns whatever pattern separates the examples, including the ones you did not intend. The fix is **more varied data**. A model is only as good as its data.

6. A, C · Data, computing power and better algorithms came together. AI is still **narrow**: it does specific tasks, not general thinking.

7. B · Past days with their **real occupancy** are labelled examples; the model learns from them and **predicts a number** for new days. That is supervised learning, and its quality depends on the history it is trained on.

4.2

Generative AI and responsible use

Ask a chatbot who founded your village's football club, and it answers in a second: a name, a year, a nice sentence about the founders. It sounds certain. Is it true?

Sometimes it is. Sometimes the name, the year and even the "source" it quotes were invented on the spot, written just as fluently as the truth. Companies in the Balearic Islands already use these tools to answer guests, write descriptions and draft code. A hotel chatbot that invents a cancellation policy is not a funny mistake: it is a complaint, maybe a lawsuit. Today you learn how these models work, how to catch them when they are wrong, and which rules Europe has set for them.

What you will learn

- Explain what **generative AI** is and how a **large language model** writes, without magic
- Catch a **hallucination** by checking a model's answers against real sources
- Write rules for the **responsible use** of AI in a company
- Classify AI systems by the four risk levels of the **AI Act**, and know which parts already apply
- Avoid the trap of the day: a model that writes well is not a model that knows

4.2.1 Models that create

> **Generative AI** is AI that **produces new content** (text, code, images, audio, video) from an instruction, instead of only classifying or predicting a value.

Your Teachable Machine model from the last lesson **classified**: correct or defective. A generative model **creates**: a paragraph, an image, a function.

> A **large language model (LLM)** is a very large neural network trained on enormous amounts of text to **predict the next piece of text**. Repeating that prediction again and again, it writes whole answers.

That is the whole trick. It does not look the answer up in a database; it writes the most likely continuation of your text, piece by piece. In the city-builder game it is an advisor that talks with great confidence, even when it has not looked at the map.

Term	What it means
Prompt	The instruction or input text you give the model. More context and more precision give better answers

Term	What it means
Token	A piece of text (a word or part of a word) that the model works with. Prices and limits are counted in tokens
Hallucination	A false or invented answer, well written and convincing
Knowledge cutoff	The model only "knows" what was in its training data; anything later it does not know unless you give it, or it searches the web

What it does well	Its limits
Write, summarise, translate and rephrase	On its own it does not check facts: it can invent data, quotes and sources
Explain and generate code, tests and documentation	It can repeat the **biases** in its training data
Pull information out of messy texts	The same question can give different answers
Give you a fast first draft	What you type may leave the company (privacy); it uses a lot of energy

> **Careful: "it sounds sure, so it is right"**
> An LLM writes a correct answer and a wrong answer in exactly the same confident style. The problem is not that it fails: every tool fails. The problem is that it **fails without warning**. That is why every output is a draft until a person checks it.

4.2.2 In class: error hunters

This is the **only activity of the unit where generative AI is authorised**, and only on these terms:

> **What is authorised, and what is not**
> **Authorised:** using the AI assistant your teacher names, **in class, during this activity only**, to ask the questions below and analyse the answers.
>
> **Not authorised:** writing your conclusions with AI, any other class activity, and any part of Deliverable 4. Never type personal data (yours or anyone else's) or internal company data into the assistant.

In pairs. Every answer must be checked against a **real source**: an official website, the press, the language documentation, or by running the code.

Step	Ask the assistant	Check it by
1	Three questions about things you know very well: your village or neighbourhood, your school, a sport or a game you master	Your own knowledge, confirmed with a website
2	Two facts about your project company that you **already cited** in your dossier (year founded, number of hotels or staff…)	Comparing with the source in your dossier
3	A short function in the language you use in Programming	Running it, with a normal case and an edge case
4	The sources for one of its answers	Opening the links: do they exist? Do they say what the model claims?

For each answer, mark **correct**, **wrong** or **cannot be checked**, and paste the link you used to check it. Then name the **type** of error: invented fact, non-existent source, code that does not work, outdated answer, bias.

4.2.3 Rules for responsible use

From the errors you catch, five rules follow. Most companies that use AI write something similar.

Rule	What it means
Check it	Everything a model produces is a draft until someone checks it
Protect the data	No personal data, and no confidential code or documents, in tools the company has not approved
Be transparent	Say when you used AI and what for
You are responsible	Whoever signs the work or publishes the code answers for it, not the tool
Respect authorship	Respect the authorship and licences of content and code

Your group will turn these into rules for **your** company, with at least one rule specific to its sector. In a hotel, for example: “the chatbot never confirms a booking or a price without the booking system”.

This is not only good manners. Since **2 February 2025**, the AI Act requires companies that use AI to make sure their staff have enough knowledge of it, called **AI literacy**.

4.2.4 The AI Act: rules by risk

> The **AI Act** is **Regulation (EU) 2024/1689**, the European law on artificial intelligence. It entered into force on **1 August 2024** and applies in stages.

It mostly affects **providers** (whoever develops or sells an AI system) and **deployers** (whoever uses one professionally) in the European Union. Its idea is a **risk-based approach**: the more harm a system can do to people, the more obligations it has. The European Commission summarises it in four levels.

Level	What it means	Examples
Unacceptable	Prohibited practices	Social scoring of citizens; manipulation that exploits vulnerabilities; emotion recognition at work or at school (except for medical or safety reasons)
High	Allowed with strict obligations: risk management, data quality, documentation, logging, human oversight	Sorting CVs to hire people, assessing students, deciding on loans, critical infrastructure
Limited (transparency)	Transparency obligations	Chatbots must say you are talking to a machine; generated or manipulated content (*deepfakes*) must be labelled
Minimal	No specific obligations	Spam filters, AI in video games

Large language models also have their own obligations, as **general-purpose AI models**: documentation, transparency and respect for copyright, with more obligations for the biggest ones.

What applies, and since when

Date	What starts to apply
1 Aug 2024	Entry into force
2 Feb 2025	Prohibited practices and AI literacy
2 Aug 2025	Governance and the obligations of general-purpose AI models
2 Aug 2026	General application, including transparency obligations
2 Dec 2027	High-risk systems such as hiring, credit or education (postponed by the *AI Omnibus*)
2 Aug 2028	High-risk systems built into regulated products

> **The calendar has already changed once**
> The high-risk dates were first set for 2026 and 2027. An amendment, the **AI Omnibus**, entered into force on **27 July 2026** and pushed them back. When you write about the AI Act, check the current dates on the European Commission's page and cite it: laws move. And do not confuse it with the **GDPR** (Regulation (EU) 2016/679), the data protection law you will see in Unit 5.

4.2.5 In class: which risk level?

In your project groups, classify each case and justify it in one sentence. Then classify the two AI ideas you noted for your own company in the last lesson.

Case	Your level
A chatbot that answers customers' questions on a hotel's website	
A program that sorts CVs and rejects candidates	
The spam filter of a transport company	
A public administration system that scores citizens by their behaviour	
An AI that marks exams and decides who gets into a course	
A promotional video with a person's voice generated by AI	
Cameras that detect the mood of office workers	
A model that decides whether a bank grants a loan	

Hand in on Classroom (in pairs, "S16 · Error hunters and risk levels", by 23:59 the same day): your error table with the link for every check, the five rules adapted to your company, and the classification of the eight cases and of your two ideas. Remember: the AI authorisation ended with the error-hunting activity.

4.2.6 For your project

This lesson feeds the **group part** of Deliverable 4 · AI plan:

Today's work	Becomes	Criterion
Your five rules, adapted to the company	Group part: responsible-use rules	c
The risk level of your own ideas	Group part: the AI Act risk level of each proposal	c
What generative AI can and cannot do	Section C: present and future of AI in the sector	c, f

4.2.7 Key takeaways

- **Generative AI** creates content; an **LLM** writes by predicting the next piece of text, again and again.
- It writes right and wrong answers equally well: **hallucinations** come without warning. Always check against sources.
- Responsible use: **check it**, **protect the data**, **be transparent**, **you are responsible**, **respect authorship**.
- The **AI Act** (Regulation (EU) 2024/1689) follows a **risk-based approach**: unacceptable, high, limited (transparency), minimal.
- Prohibitions and AI literacy apply since 2 February 2025; transparency since 2 August 2026; high-risk obligations were postponed to December 2027 and August 2028.
- In this module, AI is allowed **only** when the teacher authorises it for a specific activity.

4.2.8 Quick check

1 What does a large language model basically do when it answers?

A It looks the answer up in a database of verified facts
B It predicts, piece by piece, the most likely continuation of the text
C It searches the internet and copies the best page
D It runs rules written by experts for every question

2 A chatbot gives you a quote from a report, with the title and the year. You search for the report and it does not exist. What is this called?

A A token
B A prompt
C A hallucination
D A knowledge cutoff

3 Which two actions follow the rules of responsible use?

A Pasting a list of customers with their phone numbers into a public chatbot to write emails faster
B Telling your team that a first draft was written with an AI tool the company approved
C Testing a function generated by an assistant before adding it to the project
D Publishing an AI answer as it is, because the tool is responsible for it

4 Under the AI Act, which level is a program that sorts CVs and rejects candidates?

A Unacceptable: it is prohibited
B High risk: allowed with strict obligations such as human oversight
C Limited: it only has to say it is an AI
D Minimal: no obligations

5 A hotel puts a chatbot on its website. What is the minimum the AI Act requires?

A Nothing: chatbots are minimal risk
B A licence from the European Commission
C Telling users that they are talking to an AI
D Switching it off at night

6 Which of these is a prohibited practice under the AI Act?

- **A** A spam filter
- **B** A video game with AI-controlled enemies
- **C** Cameras that recognise the emotions of workers in an office
- **D** A recommender in an online shop

7 Your group proposes an AI that predicts which machines in a bakery will break down next week, from their sensor data. Which risk level fits best, and what must you still do?

- **A** Unacceptable; you must drop the idea
- **B** Probably minimal: it decides nothing about people, but a person should still check its warnings
- **C** High, because it uses sensors
- **D** Limited, so it must announce it is an AI to the machines

You know how a generative model writes, why you check it, and what Europe asks of it. Now switch sides: you are the company, and you want to know whether AI pays. Next: Automation and profitability.

Solutions

1. B · An **LLM** is a predictor of text. That is why it can write a convincing answer that is false: nothing in the prediction checks the facts.

2. C · An invented but convincing answer is a **hallucination**. Invented sources are one of its most common forms, which is why you always open the links.

3. B, C · **Be transparent** and **check it**. Personal data never go into unapproved tools, and the person who publishes is the one responsible, not the tool.

4. B · Hiring decisions affect people's lives, so the AI Act lists them as **high risk**: risk management, data quality, documentation and **human oversight**.

5. C · Chatbots are in the **limited (transparency)** level: people must know they are talking to a machine. Those transparency obligations apply since **2 August 2026**.

6. C · Emotion recognition **at work or at school** is an **unacceptable** practice (except for medical or safety reasons), prohibited since 2 February 2025.

7. B · A maintenance prediction does not decide about people's rights, so it fits **minimal** risk. Responsible use still applies: someone **checks** before acting, and the data must be good.

4.3

Automation and profitability

You email your internet provider because the router has stopped working. Two minutes later an answer arrives: it has understood the problem, it has opened a ticket with the right technician, and it suggests restarting the router while you wait. Did a person read your email? Probably not.

Somebody decided to pay for that system. Not because it is modern, but because it saves hours of work, and those hours cost money. In September 2026, Telefónica announced that Meliá's technical support, which includes 37 hotels in the Balearic Islands, handles around 240,000 requests a year and already automates more than 70% of them. Today you learn to ask the question every company asks before it buys AI: **does it pay?** And you will answer it with numbers, in a spreadsheet.

What you will learn

- Tell apart three steps of automation: manual, by **rules** and with **AI**
- Tell **automating** a task from **optimising** it, and name the number that shows the improvement
- Find and cite a real AI use in a Balearic company, separating what the source says from your **hypothesis**
- Explain the chain from **data** to money, and why **big data** matters
- Build a **profitability** estimate in a spreadsheet: savings, net benefit, **payback period** and **ROI**
- Avoid the trap of the day: an AI project that saves 800 euros and costs 1,000 is not innovation, it is a loss

4.3.1 Three steps of automation

> **Automation** is doing a task with less human work. **Optimisation** is doing it **better**: the best price, the best route, the best moment.

Step	How it works	Example: a company's technical support
Manual	A person does every step	A technician reads every support email and forwards it to the right person
By rules (scripts, **RPA**)	A program repeats fixed steps defined by people. Works when the input is always the same	If the subject contains "invoice", the email goes to accounts
With AI	A model handles variable, unstructured input (free text, images, voice), predicts and proposes decisions	A model understands what the email is about, classifies it, estimates its urgency and drafts a reply

> **RPA** (*robotic process automation*) is software that copies the clicks and keystrokes a person would make, following fixed rules.

AI does not replace rules; it covers the cases rules cannot reach. In the city-builder game, a timer on each traffic light is a rule; the advisor that watches the traffic overlay and changes the timings is AI.

Where companies use it

Use	What AI does	The number that shows the improvement
Classification and extraction	Reads invoices, emails or forms and pulls out the data	Time per document; typing errors
Customer service	Answers frequent questions and passes the rest to a person	Response time; questions solved without a person
Demand forecasting	Estimates sales, occupancy, consumption or workload	Forecast error; leftover stock or staff
Dynamic pricing	Adjusts the price to the predicted demand	Revenue per unit sold
Predictive maintenance	Spots signs of a breakdown in sensor data (the OT side) before it happens	Hours of unplanned downtime
Routes and planning	Calculates more efficient routes or shifts	Kilometres, fuel, hours

4.3.2 From data to money

Without data there is no AI, and with bad data the AI decides badly: garbage in, garbage out. When the data are huge, varied and arrive fast, we call it **big data**: too much for a spreadsheet, so it needs the cloud from Unit 3 and specific tools. You will study it in Unit 5.

AI can make a company more profitable in three ways: **more revenue** (sell more, or better), **lower costs** (time, errors, energy) and **lower risk** (fraud, breakdowns, fines). But it also costs: collecting and cleaning the data, the cloud service, setting it up and connecting it to existing systems, maintaining the model, and the people who supervise it.

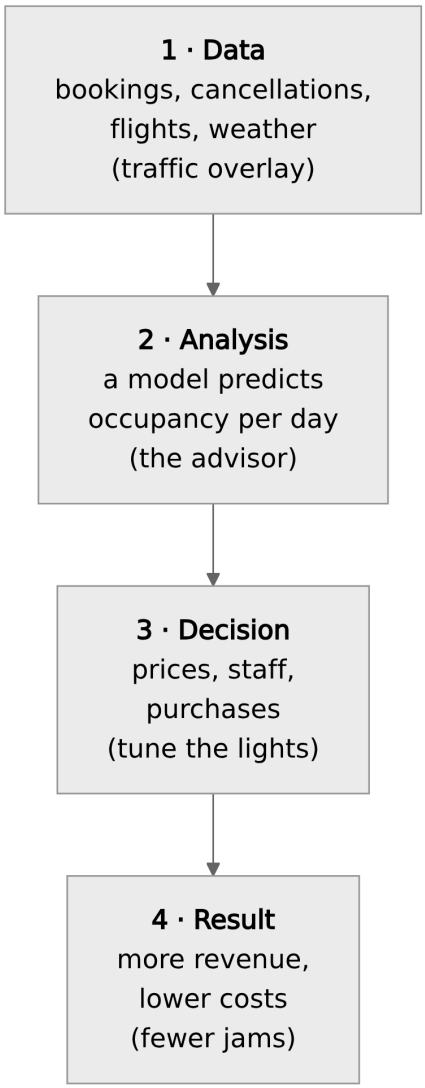
> **Return on investment (ROI)** = (what you gain – what it cost) / what it cost. The **payback period** is how many months it takes to recover the initial investment.

> **Careful: “it saves time, so it pays”**
> Saved hours are only real savings if those people spend them on other valuable work. And every mistake the model makes costs a person’s time to fix. A profitability estimate that ignores supervision and errors is too optimistic.

4.3.3 In class: a real Balearic case

In pairs. The question: *which company or organisation operating in the Balearic Islands uses AI, or automatic forecasting, in a real process, what exactly does it do, and how do we know?* Pick one starting point, or find a case from your own company’s sector.

Sector	Starting point	What to find
Tourism	Telefónica’s press release about Meliá (10 September 2026), on telefonica.com	How many Balearic hotels are in the deal, how many support requests a year, what share is automated, and what AI does



Sector	Starting point	What to find
Energy	Red Eléctrica: “Real-time electricity demand in the Balearic Islands”, on ree.es	The three curves (actual, forecast and scheduled). How does REE say it calculates the forecast? Is that AI, statistics, or can we not tell?
Ports and transport	The innovation news of Ports de Balears (portsdebalears.com), or EMT Palma / TIB	One concrete digitalisation or AI project, with a date. If there is no AI, say so

For your case write: company and sector; the process; what the AI or the model does; the number the source gives; the full citation (title, website, date, link); and one line that separates **what the source says** from **what is your hypothesis**.

> **Does the source say “AI”?**
> Red Eléctrica explains that its forecast is built from consumption in similar past periods, corrected for working days, weather and economic activity. It does not say it is AI. Saying “REE uses AI” would be your guess, not a fact. That distinction is exactly what the project asks of you.

4.3.4 In class: does it pay?

In pairs, in Google Sheets (or LibreOffice Calc). The figures are **invented for the exercise**.

A small software company in Mallorca receives **400 support tickets a month**. Each ticket takes a technician **15 minutes** on average, and an hour of a technician costs the company **€20**. They want an AI cloud service that classifies tickets and drafts a reply; they estimate the time would drop to **9 minutes** per ticket. The service costs **€300 a month**, and setting it up (integration and preparing the data) costs **€2,000** once.

Step 1 — Build the sheet with **formulas**, not with results typed by hand. Use only operators (`*`, `/`, `-`), because function names change with the language of the spreadsheet.

Row	A (concept)	B (value or formula)	You should see
2	Tickets per month	400	
3	Minutes per ticket now	15	
4	Minutes per ticket with AI	9	
5	Cost of one technician hour (€)	20	
6	AI service (€ per month)	300	
7	Set-up (€, once)	2000	
9	Hours saved per month	`=B2*(B3-B4)/60`	40
10	Savings per month (€)	`=B9*B5`	800
11	Net benefit per month (€)	`=B10-B6`	500
12	Payback period (months)	`=B7/B11`	4
13	Net benefit, first year (€)	`=B11*12-B7`	4,000
14	ROI, first year	`=B13/(B6*12+B7)`	0.71 (71%)

Step 2 — Sensitivity. Suppose one reply in ten has to be rewritten, and the real time with AI is **11 minutes**. Change only `B4`. What happens to the payback period?

Step 3 — Your process. Duplicate the tab and model a process of your project company: the one from your Teachable Machine idea, or the Balearic case you found. Every number that does not come from a source goes in a cell with the word **hypothesis** in column C.

Hand in on Classroom (in pairs, "S17 · A real case and a sheet that pays", by 23:59 the same day): the table of your real case with its citation, and the link to your spreadsheet (shared with the teacher, view only) with both tabs. **No generative AI** in this activity.

4.3.5 For your project

This lesson feeds Deliverable 4 · AI plan:

Today's work	Becomes	Criterion
The process in your second tab	Section A: processes to automate or optimise	a
The data that process needs, and your spreadsheet	Section B: data and profitability estimate	b
Your real Balearic case, cited	Section C: the AI landscape of the sector	d

4.3.6 Key takeaways

- Three steps: **manual**, by **rules** (scripts, **RPA**) and with **AI**, which handles variable, unstructured input.
- **Automating** means less human work; **optimising** means better decisions. Each improvement needs a number that measures it.
- The chain: **data → analysis → decision → result**. Without good data there is no useful AI; **big data** needs the cloud.
- **ROI** = (gain − cost) / cost; the **payback period** says when the investment is recovered. Supervision and errors are costs too.
- A source that says "forecast" is not a source that says "AI". What you cannot cite is a **hypothesis**.

4.3.7 Quick check

1 What does AI add compared with automation by fixed rules?

A It is always cheaper
B It can handle variable and unstructured input, such as free text or images
C It needs no data
D It removes the need for any supervision

2 A hotel uses a model's demand forecast to adjust its room prices every day. Is it automating or optimising, and what number shows the improvement?

A Automating; the number of emails answered
B Optimising; the revenue per room
C Automating; the number of rooms in the hotel
D Neither; prices cannot be improved with data

3 An AI tool costs €300 a month and saves €800 a month of work. Setting it up cost €2,000. When is the investment recovered?

A In 2.5 months
B In 4 months
C In 7 months
D Never

4 In the support example, the real time with AI turns out to be 11 minutes instead of 9. What happens?

A Nothing: the saving is the same
B The monthly saving drops and the payback period roughly doubles
C The project becomes more profitable
D The ROI cannot be calculated any more

5 Red Eléctrica's website says the Balearic demand forecast is built from similar past periods, corrected for weather and working days. Your teammate writes "REE uses AI to forecast demand". What is the right correction?

A It is fine: every forecast is AI
B Write what the source says, and if you think it may be AI, label it as a hypothesis
C Remove the case: forecasts are not part of RA4
D Replace the source with a chatbot's answer

6 Which two costs are often forgotten in an AI profitability estimate?

- **A** The time people spend checking and correcting the model's output
- **B** The colour of the dashboard
- **C** Preparing and cleaning the data the model needs
- **D** The number of letters in the company name

7 A water company wants to know whether to buy a predictive maintenance system for its pumps. Which data would be essential?

- **A** The customers' favourite songs
- **B** The history of sensor readings from the pumps and the dates of past breakdowns
- **C** Only the price of the system
- **D** The company's social media followers

You can now say where AI helps a company and whether it pays. Last question of the unit: which sectors use it most, what is it built with, and what does it change for you? Next: Sectors, languages and software.

Solutions

1. B · Rules work when the input is always the same. A model can deal with **free text, images or voice** and propose decisions, which is where rules stop working.

2. B · Choosing a **better** price is **optimisation**, and the indicator is the revenue per room. Automating would be doing the same task with less human work.

3. B · The net benefit is 800 − 300 = **€500** a month, and 2,000 / 500 = **4 months** of **payback period**.

4. B · With 11 minutes, 26.7 hours are saved (about €533), the net benefit falls to about €233 a month and the payback period grows to about **8.6 months**. Profitability is very **sensitive** to the real time of review.

5. B · The source describes a **forecast**, not AI. In the project, what a source says is a fact with a citation; what you deduce is a **hypothesis** and must be labelled.

6. A, C · **Supervision** and **data** are real costs. Leaving them out makes any AI project look more profitable than it is.

7. B · **Predictive maintenance** learns from the sensors (the OT side) and from past breakdowns, which are the labelled examples. Without that history, the model has nothing to learn from.

4.4

Sectors, languages and software

In Programming you write Java. Then you open the news and every AI headline seems to mention Python. Did you pick the wrong language? Will an AI assistant write your code before you finish the course anyway?

Good questions, and they deserve data, not opinions. Today you look at three kinds of evidence: what European statistics say about which companies use AI, what real AI projects on GitHub are actually written in, and what developers themselves report in the largest surveys. Then you look at your own future job: what AI changes in software development, and what it does not.

What you will learn

- Name the sectors and kinds of company where AI is most used, with **Eurostat** figures
- Name the main **programming languages** and **libraries** for AI and what each one is for
- Check the languages of real AI projects on **GitHub** and cite developer surveys
- Describe how AI is changing **software development** and the developer's responsibility
- Avoid the trap of the day: "AI is Python" is only half true; the engine is often C++, and many apps just call a model through an **API**

4.4.1 Where AI is used most

> **Eurostat** is the statistical office of the European Union. Every year it asks companies with 10 or more people which digital technologies they use.

Its 2025 figures:

Figure (EU, 2025)	Value
Companies that use at least one AI technology	19.95% (about 6.5 points more than in 2024)
Large / medium / small companies	55.03% / 30.36% / 17%
Sector with most use: information and communication	62.52%
One of the sectors with least use: construction	10.79%
Country with most use / least use	Denmark 42.03% / Romania 5.21%
Most used AI technology: text mining	11.75%

Figure (EU, 2025)	Value
Most common purpose: marketing or sales	34.70%

So roughly **one company in five** in the EU uses AI. It takes root first where there are **lots of digital data**, **repeated decisions with economic value**, and **staff who know how to use it**: that is why IT companies and large companies lead.

Beyond the ranking, some sectors have well-known uses:

Sector	Typical uses
IT and software	Code assistants, search, cybersecurity, content moderation
Finance and insurance	Fraud detection, credit risk, automated customer service
Health	Help with diagnosis from medical images, prioritising cases, drug research
Retail and marketing	Recommenders, customer segmentation, sales forecasts
Industry	Predictive maintenance, visual quality control, robotics
Transport and logistics	Routes, demand forecasting, driver assistance
Tourism	Dynamic pricing, occupancy forecasting, multilingual chatbots

Tourism is not at the top of Eurostat's table, but it is **our** sector, and you already found a real Balearic case in the last lesson. Remember too that hiring, credit and health are **high risk** under the AI Act.

4.4.2 Languages and tools

Language	Role in AI	Libraries and tools
Python	The reference language for training and experimenting: simple, with a huge ecosystem	NumPy, pandas, scikit-learn, PyTorch, TensorFlow/Keras, Hugging Face Transformers
C++ (and CUDA)	Speed: the engine under many libraries, and running models on devices and graphics cards	The cores of PyTorch and TensorFlow, llama.cpp
SQL	Getting and preparing the data	Databases and data warehouses
R	Statistics and data analysis	Statistical and visualisation packages
Java and Kotlin	Adding AI to business and Android apps, by calling models	LangChain4j, Deeplearning4j, API clients
JavaScript and TypeScript	Running or calling models in the browser and on the web	TensorFlow.js (the engine of Teachable Machine), API calls

> A **library** is code that somebody else wrote and published so you can reuse it. Nobody writes a neural network from zero: you import PyTorch.

Around the language there are three more tools you will meet:

- **Notebooks** such as Jupyter or Google Colab: one document that mixes code, text and results.
- **Graphics cards (GPUs)**, usually rented in the cloud from Unit 3, for training.
- **Ready-made models and AI cloud services** that you call through an **API**: a web address your program sends a request to and gets an answer from. Very often a company trains nothing; it **consumes** a model someone else trained.

> **Careful: “to do AI I must switch to Python”**
> To **train** and experiment, Python is the reference. To **use** AI inside an app, which is what most developers do, you call a model through an API from whatever language the app is written in: Java, Kotlin, JavaScript. Your Java is not wasted.

4.4.3 In class: what are real AI projects made of?

In pairs. On the page of every GitHub repository, the right-hand column shows a **Languages** bar with the share of each language. You do not need an account.

Step 1 — Choose **four** repositories from this list (at least one that is not mostly Python). Write down the **top three languages** with their percentage, the date you looked, and what the project is for, in one sentence from its README.

Repository	What it is
`pytorch/pytorch`	A deep learning library
`tensorflow/tensorflow`	Google’s machine learning library
`scikit-learn/scikit-learn`	“Classic” machine learning in Python
`huggingface/transformers`	Ready-made models for text, images and sound
`ggml-org/llama.cpp`	Running language models on an ordinary computer
`tensorflow/tfjs`	TensorFlow for the browser
`langchain4j/langchain4j`	Adding language models to Java applications

Step 2 — Find and cite two figures in primary sources: **(a)** in Stack Overflow’s 2026 Developer Survey (on stackoverflow.blog), the share of developers who use Python and the language in first place; **(b)** in GitHub’s Octoverse 2025 report (on github.blog), the most used language on GitHub and the language that still leads AI and data science.

Step 3 — Answer: which project is “more C++ than Python”, and why do you think so?

> **What you will probably find**
> The **interface** of most AI libraries is Python, but the part that has to be fast (the **engine**) is C++ or CUDA. In TensorFlow, C++ is more than half the code. A Java library like LangChain4j is almost entirely Java. Percentages change a little every day, so always write the date.

Hand in on Classroom (individual, “S18 · What are AI projects made of?”, by 23:59 the same day): the table of your four repositories; the two cited survey figures; and **your own paragraph** (5–8 lines) on how AI would change the work of your project company’s IT department. **No generative AI** in this activity.

4.4.4 AI in software development: your sector

Your qualification’s sector is software development, and AI is changing it fast.

Area	What changes
Writing code	Assistants and agents complete and generate functions, tests and documentation, explain old code and help debug

Area	What changes
Products	Apps add AI features (search, summaries, recommendations, chat) by calling models through APIs
Jobs	New roles: machine learning engineer, data engineer, model operations (**MLOps**)
Risks	Code with bugs or vulnerabilities that looks correct, unclear licences, confidential code leaking, losing your skills if you accept what you do not understand
What does not change	Understanding the customer's problem, designing, reviewing, testing and **answering for** what you publish

Look at this code. An assistant proposed it to find a user. It compiles and it works. Would you accept it?

```
String sql = "SELECT * FROM usuaris WHERE nom = '" + nom + "'";
Statement st = con.createStatement();
ResultSet rs = st.executeQuery(sql);
```

It is vulnerable to **SQL injection**: whatever the user types is glued into the query, so a name containing a quote can change the query itself. The correct version uses a parameterised query (`PreparedStatement` with `?`). Generated code must be reviewed like the code of a new teammate.

> **SQL injection** is an attack where text typed by a user is executed as part of a database query, because the program pasted it in without separating data from code.

AI will most likely not take your job away; it will change it. Whoever can only copy what a tool proposes is replaceable. Whoever can explain **why** that code is wrong is not.

4.4.5 For your project

This lesson feeds Deliverable 4 · AI plan. It is also the session where your group shares out the sections and roles of the deliverable:

Today's work	Becomes	Criterion
Eurostat figures and the sector table	Section C: the AI landscape of the sector	d
Your paragraph on the IT department	Section C: how AI changes software work	f
Your four repositories and the survey figures	Section D: languages and tools	e

4.4.6 Key takeaways

- About **one EU company in five** used AI in 2025 (**Eurostat**); **information and communication** leads, large companies use it far more than small ones.
- **Python** is the reference for training; **C++** is often the engine; **SQL** prepares the data; Java, Kotlin and JavaScript integrate AI into apps.
- Many apps train nothing: they call a ready-made model through an **API**.
- Real projects on **GitHub** show it: Python on the outside, C++ on the inside. Cite the date, the numbers change.
- AI changes **software development** (assistants, new products, new roles) but you still **review**, **test** and **answer for** the code.

4.4.7 Quick check

1 According to Eurostat, which sector of the EU had the highest share of companies using AI in 2025?

- **A** Construction
- **B** Information and communication
- **C** Accommodation and food services
- **D** Agriculture

2 Why is C++ so important in AI if most people write AI code in Python?

- **A** Because C++ is easier to learn than Python
- **B** Because the fast engine under many Python libraries is written in C++
- **C** Because Python cannot do mathematics
- **D** Because GitHub only accepts C++

3 Your company's Java app needs to summarise customer reviews with a language model. What is the most common approach?

- **A** Rewrite the whole app in Python
- **B** Train a new language model from zero
- **C** Call a ready-made model through an API from the Java app
- **D** It cannot be done without a GPU in every office

4 Where do you see which languages a GitHub repository is written in?

- **A** In the name of the repository
- **B** In the Languages bar in the right-hand column of the repository page
- **C** Only after downloading all the code
- **D** In the list of contributors

5 An assistant proposes code that builds an SQL query by gluing in what the user typed. It compiles and works. What should you do?

- **A** Accept it: if it works, it is correct
- **B** Reject it and use a parameterised query, because it is open to SQL injection
- **C** Accept it, because the assistant is responsible for its code
- **D** Delete the database

6 Which two statements about AI and the developer's job are true?

- **A** Reviewing, testing and answering for the code remain the developer's job
- **B** Code assistants make knowing how to program unnecessary
- **C** New roles such as machine learning engineer and MLOps have appeared
- **D** AI is only used by developers who work in Python

7 A small hotel company in Menorca, with 15 staff, wonders why it uses less AI than a large chain. Which explanation fits the Eurostat data best?

- **A** Small companies are not allowed to use AI under the AI Act
- **B** AI use grows with size: large companies have more data, more repeated decisions and more specialised staff
- **C** AI only works in information and communication companies
- **D** Menorca is not part of the EU statistics

You now know where AI is used, what it is made of and what it changes for you. Time to put it all together for your company: Deliverable 4 · AI plan.

Solutions

1. B · **Information and communication** led with 62.52%. Construction was among the lowest, at 10.79%.

2. B · Python is the comfortable **interface**; the heavy calculations run in a **C++** (and CUDA) engine underneath. That is what the Languages bar of PyTorch or TensorFlow shows.

3. C · Most apps **consume** a model through an **API**. Libraries such as LangChain4j exist exactly for this in Java.

4. B · The **Languages** bar shows the share of each language. It changes over time, so always write the date you looked.

5. B · Working is not the same as safe. Gluing user text into a query allows **SQL injection**. The developer who publishes the code is the one responsible.

6. A, C · AI changes the job: new tools, new roles. The responsibility for what is published, and the skill to judge it, stay with you.

7. B · Eurostat shows a big gap by size: 55% of large companies against 17% of small ones. **Data**, **scale** and **skills** explain much of it, not the law.

4.5

Deliverable 4 · AI plan

Your city is running: the map is drawn, the tech tree is unlocked, the utilities are rented from the cloud. Now the game offers you an **automatic advisor**. Before you switch it on, a good mayor asks four questions. Which jobs would it do? What does it need to watch, and does it pay for itself? Who else uses one, and what does it change? And what could go wrong?

Deliverable 4 is that decision for your company: an **AI plan**. Your group proposes where AI would really help, which data it needs, whether it pays, what the sector is already doing, what it would be built with, and the legal risk level of each proposal. One member also builds a **working demo** in the browser. It is the fourth chapter of your digital transformation dossier.

What you will learn

- Propose concrete processes of a real company that AI could **automate** or **optimise**, and measure the improvement
- Say which data each proposal needs and estimate its **profitability** with a spreadsheet
- Describe the AI landscape of the company's sector, present and future, with cited sources
- Choose languages and tools for the proposals and show a model that works
- Classify each proposal under the **AI Act** and write rules for its responsible use

4.5.1 What you hand in

Format	One PDF per group, about 6–10 pages without the cover
Where	Google Classroom, assignment "D4 · AI plan". The Coordinator hands it in
Deadline	**Monday 22 March 2027, 23:59**
Defence	Individual, in **session 19, Monday 22 March 2027** (the same day)
Weight	70% of the RA4 mark. The other 30% are the class activities of sessions 15 to 18

Every fact about your company needs a source

Everything you state as a fact about your company or its sector (what it does, what systems it has, what AI it or its competitors already use, any figure) must be **found and cited**: the company website, annual report, press releases, the press, Eurostat, Ibestat. Each citation has a title, the author or website, the link and the date you consulted it. What you cannot find you may still propose with reasons, but you label it **hypothesis**. A

claim without a source or a label counts as an error. A source that says “forecast” or “automated” does not say “AI”: do not add the word yourself.

No generative AI

Generative AI is **not allowed** in this deliverable. The only AI authorisation of the unit was the error-hunting activity in session 16, and it ended in class. A section written with AI or copied scores **0**, and the defence will show it: you cannot explain a text you did not write. The Teachable Machine model in Section D is not an exception to this rule: it is the object you build and test, not a tool that writes for you.

4.5.2 The sections

Each member **owns and signs one section** (the individual part). With 3 members, Section D becomes part of the group part and everybody builds it together.

Section	Criterion	What to write	Owner
A · Processes to automate or optimise	RA4 a)	**Two** processes of the company that AI could improve. For each: what it is, who does it today and how (manual, by rules, already with AI); what the AI would do (classify, extract, predict, recommend, generate) and whether it **automates** or **optimises**; why a fixed rule is not enough; the **indicator** that would show the improvement; who supervises the AI and what happens when it is wrong.	Individual
B · Data and profitability	RA4 b)	For each process of Section A: which **data** the model needs, whether the company already has them (and in which system, from your earlier deliverables), whether any are **personal data**, and how much data there would be. A **profitability estimate** in a spreadsheet for at least one process: monthly saving or extra revenue, costs (service, set-up, supervision), net benefit, **payback period** and first-year **ROI**, plus one sensitivity case. Every number without a source is labelled **hypothesis**. Link to the sheet.	Individual
C · The sector’s AI landscape	RA4 c), d), f)	Where the company’s sector stands: **Eurostat** figures and at least **two real cases** of AI in the sector, one of them in the Balearic Islands if possible, cited. What AI can and cannot do today in this sector, and what is likely to change in the next years (separating facts from forecasts). How AI changes the work of the company’s **IT department and its developers**: tools, new roles, risks.	Individual
D · Languages, tools and a working demo	RA4 e)	Which **languages, libraries and services** the proposals of Section A would be built with, and why (training or calling a model through an API?), supported by what you found on GitHub and in developer surveys. A **working demo** in the browser related to the company, normally a **Teachable Machine** model (image or sound): its classes and	Individual (4th member)

Section	Criterion	What to write	Owner
		number of examples, a table of at least 10 tests with new data, the hit rate, its clearest error and why, and its limits. Screenshots and, if you saved or exported it, the link.	
Introduction	RA4 c)	The company in two lines (link to your earlier chapters), what an AI plan is, and why it makes sense, or not, for this company now.	Group
Risk level and responsible use	RA4 c)	For **every proposal**, its risk level under the **AI Act** (Regulation (EU) 2024/1689) with a one-sentence reason and what that level obliges the company to do. Five **rules of responsible use** for this company, at least one specific to its sector.	Group
Integration	RA4 a)–f)	One table that joins the sections: proposal · process · data · profitability · risk level · owner. One paragraph saying **which proposal you would do first and why**.	Group
Sources	All	One list of every source cited, in the same format, with the date consulted.	Group
Group log	All	Who owned which section, who had which role, what you decided and when (template on the project page).	Group

4.5.3 Roles in this deliverable

Choose one role each and write it in the group log. Choose a section and a role you have not had before, if you can.

Role	In Deliverable 4
Coordinator	Shares out the sections in session 18, sets an internal deadline (we suggest Friday 19 March), checks the group log and hands in the PDF
Source checker	Opens every link; checks that every claim of “this company uses AI” has a source that really says so; labels hypotheses
Technician	Checks the spreadsheet formulas and that the demo works; makes sure the sheet is shared with view access
Editor	Format, spelling, one voice; the same process has the same name in every section and in the integration table. With 3 members, the Coordinator does it

4.5.4 How it is assessed

Part	Weight	What is assessed
Your section	40%	Content correct for its criterion, real and cited sources, hypotheses labelled, reasons and numbers given, clear writing
Your defence	20%	You explain your section in your own words, answer why, and know how it connects to another member’s section and to the risk levels
Group part	40%	Introduction, risk level and responsible-use rules, integration table and priority, coherence between sections, sources list, format and spelling, group log

Rubric

Part	Excellent	Adequate	Insufficient
Your section	Everything the table asks for, about this company; real cases and figures cited; numbers that add up (the sheet works); hypotheses labelled and reasoned; terms used precisely	Most of what is asked; some generic examples or a missing source; the main ideas and numbers are correct	Missing parts; text that could be about any company ("AI will improve efficiency"); facts without sources; "AI" claimed where the source does not say it; numbers that do not add up
Your defence	Answers fluently with reasons and numbers; knows where each fact comes from; explains a teammate's proposal and its risk level	Answers the main questions with some help; knows the sources of the key facts	Cannot explain parts of the section, or does not know where its content or numbers come from
Group part	Every proposal has a justified risk level and its obligations; rules are concrete and fit the sector; the integration table joins all sections; the priority is argued; clean format, one voice, complete log	Risk levels given with weak reasons; generic rules; integration with small contradictions	Risk levels missing or wrong; no rules; four separate texts glued together; no log; many spelling or format errors

Programme rules: RA4 needs at least a 5. Late work is marked out of 5; more than 48 hours late is a 0. Plagiarism or unauthorised AI is a 0. If you fail RA4, the recovery is an individual task: you redo one section of this deliverable alone and defend it, for a maximum of 5.

4.5.5 The defence

On **Monday 22 March 2027** (session 19) the groups finish the deliverable while the teacher calls each student for **2–3 minutes**. No paper and no computer: a conversation.

- **Two questions about your own section.** For example: "Why is a fixed rule not enough for this process?", "Where does the monthly saving in your sheet come from, and which numbers are hypotheses?", "Does your source really say the company uses AI?", "Show me, in words, one error of your model and why it happened."
- **One question about another member's section or the group part**, to check that the group really put it together. For example: "What risk level does your teammate's proposal have, and why?", "Which proposal would you do first?"

If you cannot come on 22 March for a justified reason, tell the teacher in advance: your defence is moved to the start of the next session.

4.5.6 Before you hand in

Check	Who
Every section has its owner's name	Coordinator
Every fact about the company and the sector has a working source; every guess says hypothesis	Source checker
No source is quoted as saying "AI" when it does not	Source checker
The spreadsheet opens with the link, its formulas work and the numbers match the text	Technician
The demo has at least 10 tests with new data, and screenshots	Technician

Check	Who
Every proposal of Section A appears in Section B, in the risk table and in the integration table, with the same name	Editor
Every proposal has a risk level with a reason; there are five rules of responsible use	Editor
The group log is filled in and signed	Coordinator
One PDF, uploaded to Classroom before Monday 22 March 2027 at 23:59	Coordinator

Start from what you did in class: your Teachable Machine report from What is AI, your rules and risk levels from Generative AI and responsible use, your real case and spreadsheet from Automation and profitability, and your GitHub table from Sectors, languages and software.



05

Unit 5 · Data and cybersecurity

In a city-builder game you never steer the city by looking at the buildings. You switch on the **info views**: the traffic overlay, the pollution map, the happiness bars. Those overlays are **data**, and without them you are flying blind. But a city with great overlays and no police, no walls and no locks on the doors will not last long either. This unit is about both halves: what data are worth to a company, how they live and die, how to get answers out of them, where to keep them safe, and which laws say how they must be protected.

Every session is hands-on. You will download a real open dataset and take it from a question to a chart, design a backup plan, run a **passive** security check of a real public website, and decide which data-protection and cybersecurity laws apply to your project company. Everything feeds **Deliverable 5**.

What you will learn

- Tell **data** from **information**, and follow one dataset of your company through its **lifecycle**, from capture to deletion
- Explain what makes something **big data** (the V's) and how big data, data analysis, machine learning, deep learning and AI fit together
- Follow the stages of **data science** on a real open dataset, from a question to a chart and a conclusion
- Choose where a company keeps its data in the **cloud** and design a **3-2-1** backup plan
- Run a **passive** security check of a public website: security headers, cookies before consent, privacy policy
- Decide whether the **RGPD**, the **LOPDGDD** and **NIS2** apply to a company, and why

Lessons

#	Lesson	What it covers
1	Data, information and lifecycle	Data vs information, the six stages of the data lifecycle, reading a real privacy policy
2	Big data	The five V's, structured and unstructured data, big data vs data analysis vs ML vs AI
3	Data science and the cloud	The stages of data science on a real Ibestat dataset, cloud storage, the 3-2-1 backup rule
4	Cybersecurity and regulation	Confidentiality, integrity, availability; device, system and global levels; a passive website audit; RGPD, LOPDGDD and NIS2
5	Deliverable 5 · Data and security brief	The brief, the sections, the rubric and the individual defence

5.1

Data, information and lifecycle

This morning you probably unlocked your phone, checked the bus time and paid for a coffee with a card. Each of those moments left a small trace somewhere: a timestamp, a card number, a bus stop code. None of those traces means much on its own. So why do companies care so much about keeping them?

Think of a hotel in Calvià in July. It receives thousands of bookings: names, dates, room numbers, prices. Stored in a database, they are just rows. But the moment somebody asks "is this July better or worse than last July?", those rows turn into something the manager can act on: hire more staff, change prices, open the second restaurant. Rows become decisions. This lesson is about that jump, and about the path every piece of data follows from the day it is captured to the day it should disappear.

What you will learn

- Tell **data** from **information**, and say what turns one into the other
- Place data, information, knowledge and decisions on the **DIKW pyramid**
- Describe the six stages of the **data lifecycle** and recognise them in a real company
- Find the data a real company says it collects by reading its **privacy policy**
- Avoid the trap: a full database is not the same as an informed company

5.1.1 Data and information

> A **datum** (plural **data**) is a raw fact, value or symbol with no context. **Information** is data that have been processed and put in context so that they mean something to somebody.

Write the number **38.5** on the board. What does it mean? A fever? A shoe size? A price? Nobody knows. Now add context, one piece at a time: "38.5 °C", "temperature of the server room", "today at 3 a.m.", "the manufacturer's limit is 35 °C". Same number. Now it is an alarm.

	Data	Information
What it is	A raw fact, value or symbol	Data processed and put in context
Answers	Nothing on its own	Who, what, when, where, how much
Hotel example	"214", "2027-07-14", "2"	"Room 214 is booked for 2 people from 14 July"
Good for	Being stored and processed	Making decisions

In the city-builder game, data are every single reading coming from every sensor in the city. Information is the **info-view overlay**: the colour map that tells you, at a glance, that the north

district is jammed. The game does the processing for you. In a real company, somebody has to program it. That somebody is often a DAM developer.

> **Careful: information is not a different kind of file**
> A spreadsheet full of numbers is not information yet. Information appears when somebody, or some program, asks the data a question and gets a useful answer. The same datum can give different information to different people: 38.5 means one thing to a doctor and another to a system administrator.

5.1.2 The DIKW pyramid

A classic way to draw the climb is the **DIKW pyramid**: data → information → knowledge → wisdom. It is usually credited to the systems thinker Russell Ackoff (1989). It is a teaching model, not a law: different authors define the levels differently. Use it as a ladder, not as a definition to memorise.

Level	Hotel example	What was added
Data	"214 booked", "215 free"	Nothing: raw records
Information	"Occupancy this week is lower than the same week last year"	Data organised and compared
Knowledge	"When there is a congress in Palma, weekday occupancy goes up"	A pattern learned from a lot of information
Wisdom (decision)	"On congress weeks, raise prices and add staff"	Somebody acting on the knowledge

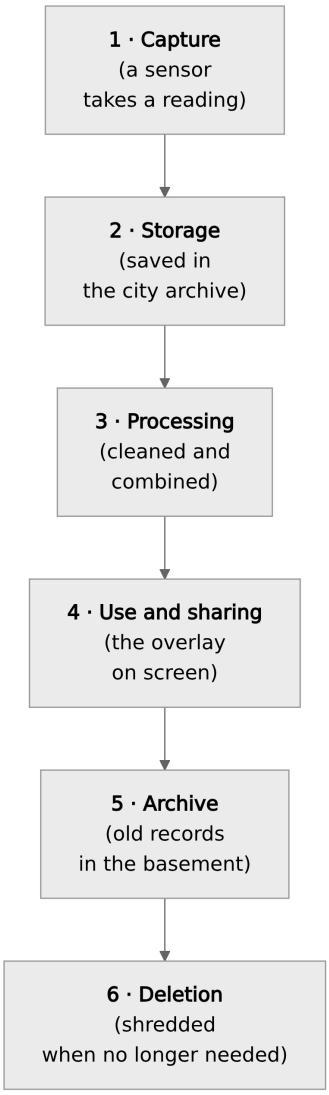
5.1.3 The data lifecycle

Data are not eternal. They are born, they are used, and at some point they should die.

> The **data lifecycle** is the sequence of stages a piece of data goes through, from the moment it is captured or created until it is archived or deleted.

Stage	What happens	At the hotel in Calvià
1 · Capture or creation	Data enter the system: forms, sensors, logs, imports	The guest fills in the booking form; the electronic lock records every opening
2 · Storage	Data are saved somewhere: a database, files, the cloud, with backups	The booking is saved in the hotel management system
3 · Processing	Data are validated, cleaned, transformed and combined	Availability is checked, the price is calculated, a room is assigned
4 · Use and sharing	Data are analysed, shown and sent to whoever needs them	Reception sees today's arrivals; management reads the occupancy report
5 · Archive	Data are no longer used every day but are kept, for legal reasons or history	Bookings from past seasons go to an archive with restricted access
6 · Deletion	Data are erased or anonymised safely when there is no reason to keep them	At the end of the retention period, personal data are deleted or anonymised

There is no single official model: depending on the source you will find five, six or eight stages with slightly different names. What matters is the **order** and being able to give an example of each stage.



Careful: the stage everybody forgets

The lifecycle does not end when data are stored. Stage 6 is the one companies forget most, and it matters: data you no longer need still cost money to store and can still be stolen. Each stage also has its own security risk (a form that leaks, a backup next to the original, an archive anybody can open). You will come back to this in Cybersecurity and regulation.

5.1.4 Where to find a company's real data

You cannot see inside your project company's database. But there is a public document that has to say, by law, which personal data the company collects, what for, who receives them and how long they are kept: its **privacy policy** (*política de privacidad*). The European data-protection regulation (RGPD, Articles 13 and 14) requires companies to give people that information when they collect their data.

Usually in a privacy policy	Which lifecycle stage it tells you about
Which data are collected (name, email, booking dates, card…) and how (forms, cookies, app)	1 · Capture
Purposes: what the data are used for	3 · Processing and 4 · Use
Recipients: other companies that receive the data (payment, cloud, marketing)	4 · Sharing
International transfers: whether data leave the European Union	2 · Storage
Retention period: how long the data are kept	5 · Archive and 6 · Deletion
Your rights and how to exercise them	All stages

5.1.5 In class: inventory and journey of one dataset

Work **individually**, on your group's project company. Use the company's own website. **No AI tools** in this activity.

Step	Do this	You should see
1	Open the company's website and find the privacy policy (usually a link in the footer: *Privacidad*, *Aviso legal*, *Privacy*). Copy its URL.	A long legal page, often in Spanish
2	Write a table with **at least five kinds of data** the company handles. For each one: what it is, where it comes from, whether it is **personal** data, and which information or decision it allows. Mark each row as *found* (with the source) or *hypothesis*.	Some rows found in the policy, some hypotheses about the OT side (sensor readings, machine logs)
3	Pick the **most important dataset** (bookings, orders, water-meter readings…) and follow it through the six stages: who captures it, where it is stored, how it is processed, who uses it, when it is archived, when it should be deleted.	A six-row table
4	Mark the stage where the data become **information** and for which decision. Mark the stage with the **highest risk** and say why.	Two short sentences
5	Hand in the document on Classroom with the privacy-policy URL and the date you read it.	

If the company has no privacy policy, or it does not list retention periods, write that down: it is a finding too.

5.1.6 For your project

This activity is the first draft of **section A (data inventory and lifecycle)** of Deliverable 5. Whoever owns section A in your group will build on everybody's tables, so keep yours in the group's shared folder. Remember the rule: what you found, you cite; what you could not find, you label *hypothesis*.

5.1.7 Key takeaways

- A **datum** is a raw fact; **information** is data processed and put in context to answer a question or support a decision.
- The **DIKW pyramid** (data → information → knowledge → wisdom) is a useful teaching ladder, not an exact definition.
- The **data lifecycle** has six stages: **capture**, **storage**, **processing**, **use and sharing**, **archive** and **deletion**.
- The last stage, deletion, is the one companies forget most, and every stage carries its own security risk.

- A company's **privacy policy** is a public, legally required source that tells you which personal data it collects, why, who receives them and for how long.

5.1.8 Quick check

1 Which of these is information, not just data?

- **A** `07014`
- **B** `2027-03-01T08:15:00Z`
- **C** "60% of this month's incidents were opened on a Monday"
- **D** `404`

2 What turns data into information?

- **A** Saving them in a database
- **B** Processing them and putting them in context so they answer a question
- **C** Collecting a very large amount of them
- **D** Converting them to a PDF

3 Which is the correct order of the data lifecycle?

- **A** Storage → capture → use → processing → deletion → archive
- **B** Capture → storage → processing → use and sharing → archive → deletion
- **C** Capture → processing → deletion → storage → use → archive
- **D** Use → capture → storage → archive → processing → deletion

4 A hotel keeps the bookings of past seasons in a restricted archive because the law requires it. Which stage is that?

- **A** Processing
- **B** Use and sharing
- **C** Archive
- **D** Capture

5 Where can you find, legally and publicly, which personal data a company collects and how long it keeps them?

- **A** In its internal database, by asking an employee for the password
- **B** In its privacy policy, which the RGPD requires it to provide
- **C** Nowhere: companies never publish that
- **D** Only in its annual accounts

6 In the city-builder analogy, what is the info-view overlay that shows the jammed north district?

- **A** Raw data
- **B** Information
- **C** A backup
- **D** The deletion stage

7 A classmate says: "Our company has a huge database, so it is already well informed." Pick the two answers that are true.

- **A** A full database is not information until somebody asks it a question and gets a useful answer
- **B** Correct: the more rows, the better the decisions
- **C** Data that are kept with no purpose still cost money to store and can still be stolen
- **D** Data can never be deleted once they are in a database

You can now follow one piece of data from birth to deletion. Next: what happens when the data are so many, so fast and so varied that a normal database is not enough? Big data.

Solutions

1. C · The other three are raw values with no context: a code, a timestamp, a number. The sentence about incidents is data that have been **grouped, counted and put in context**, so it can support a decision.

2. B · Storing or accumulating data does not make them meaningful. **Information** appears when the data are processed and given context for somebody who needs to decide.

3. B · Data are first **captured**, then **stored**, **processed**, **used and shared**, then **archived** when they are no longer used every day, and finally **deleted** or anonymised.

4. C · The data are not used day to day any more, but there is a reason to keep them: that is the **archive** stage. Deletion comes when that reason ends.

5. B · The **privacy policy** must tell people which data are collected, for which purposes, who receives them and the retention period. It is the right public source for your data inventory.

6. B · The overlay is what you get after the game **processes** thousands of sensor readings (data) and draws them in context. That is **information**: ready for a decision.

7. A, C · Same lesson in a new situation. **Accumulating data** is not the same as being informed, and data with no purpose are a cost and a **risk**. That is why the lifecycle ends with deletion.

5.2

Big data

Is the spreadsheet with your class marks big data? Of course not. And every tap of every travel card on every bus in Mallorca for a whole year? Or every search for a flight and a hotel that a Palma travel website receives in August? Somewhere between those examples there is a line. Where is it?

Big data is one of the enabling technologies you met in Unit 2, and in Unit 4 you saw that AI learns from data. This lesson explains what makes data "big", why "a lot of data" is not the right answer, and how big data, data analysis, machine learning, deep learning and AI fit together, so that you can stop using the five terms as if they were synonyms.

What you will learn

- Describe the five V's that define **big data**: volume, velocity, variety, veracity and value
- Tell **structured**, **semi-structured** and **unstructured** data apart
- Decide whether a case is a big data problem, and which V weighs most
- Explain how big data, **data analysis**, **machine learning**, **deep learning** and AI are related
- Name the four types of analysis by the question they answer
- Avoid the trap: big data does not just mean "lots of data"

5.2.1 What makes data "big"

> **Big data** are datasets that, because of their volume, the speed at which they arrive or their variety, cannot be stored, processed or analysed efficiently with traditional tools (a spreadsheet, a single database on one server) and need a **scalable** architecture.

In the city-builder game, a small town has a few dozen sensors and you could read them one by one. A city of a million inhabitants has traffic counters, water meters, bus GPS units and citizens complaining on social media, all at once, every second. The overlays still work, but behind them there is no longer one computer: there is a whole data centre.

The first three V's were proposed by the analyst Doug Laney in a 2001 research note for the META Group (now part of Gartner). Most textbooks today add two more.

V	What it means	Example
Volume	The amount of data is very large and keeps growing	Every flight and hotel search a travel website receives each day
Velocity	Data arrive fast and must be handled almost in real time	The GPS position of every bus in a network, sent every few seconds

V	What it means	Example
Variety	Many formats and sources mixed together	Bookings (tables), reviews (text), photos, server logs (JSON), sensors
Veracity	Quality and reliability: incomplete, duplicated, wrong or biased data	Fake reviews, badly calibrated sensors, the same customer typed twice
Value	Data are only worth keeping if they bring a benefit	Forecasting occupancy to adjust prices and staff

Some sources add more V's (variability, validity, visualisation...). The US National Institute of Standards and Technology (NIST), for example, adds **variability**. You do not need to memorise them: volume, velocity and variety are in almost every definition, and veracity and value complete the usual five.

5.2.2 Structured, semi-structured, unstructured

Variety becomes clearer when you sort data by how much structure they have.

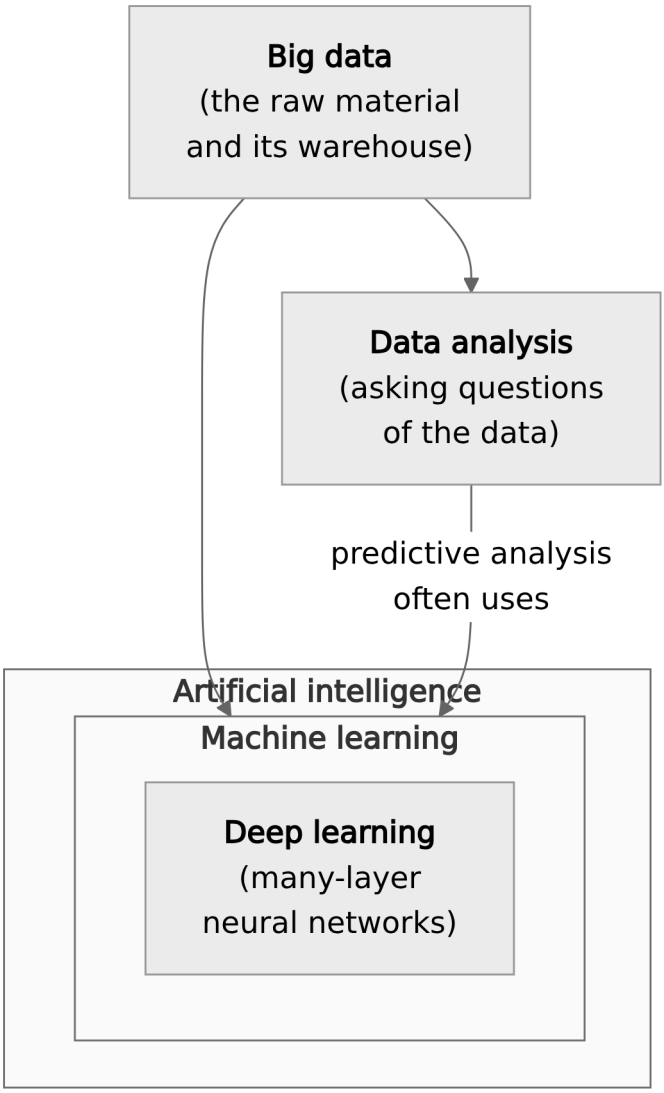
Kind	What it looks like	Example
Structured	Rows and columns with a fixed schema	A SQL table of bookings
Semi-structured	Labels or keys, but no rigid schema	JSON from an API, XML, a server log
Unstructured	Free content with no predefined fields	Review text, photos, audio, video

> **Careful: big data is not "a lot of data"**
> A huge file that is uniform and never changes is not the same problem as a continuous stream of mixed data. And the other way round: **not every company needs big data**. Many small and medium companies are perfectly served by a good database and somebody who knows how to ask it questions. The test is not "how many gigabytes?" but "does the traditional tool stop working here, and why?"

5.2.3 Who is who: big data, analysis, ML, DL and AI

Term	What it is
Artificial intelligence (AI)	Systems that do tasks we associate with human intelligence (you saw it in Unit 4)
Machine learning (ML)	The part of AI where the system **learns from data** instead of following rules written by hand
Deep learning (DL)	The part of ML based on neural networks with many layers; it usually needs a lot of data and computing power
Big data	Not AI. It is the **raw material** and the infrastructure to store and move it
Data analysis	The process of examining data to get information out of them; it can use simple statistics or ML

Big data without analysis is a full warehouse that nobody looks at. Analysis without good data gives bad answers. And a lot of the AI that companies actually use is predictive analysis done at scale. That is why the three usually go together.



5.2.4 Four questions, four types of analysis

Type	Question	Hotel example	Typical tool
Descriptive	What happened?	Average occupancy in July	Reports, charts
Diagnostic	Why did it happen?	Occupancy fell the week an event was cancelled	Crossing data sources
Predictive	What will happen?	Expected bookings for next week	ML
Prescriptive	What should we do?	Recommended price per room and day	ML plus optimisation

5.2.5 In class: big data in your company's sector

Work **in pairs**, on the sector of your group's project company. This is guided web research: every claim needs a working source. **No AI tools.**

Step	Do this	You should see
1	Search for one **real, documented** case of big data or data analytics in your company's sector, ideally in the Balearic Islands or in Spain (a company's	Two or three candidate articles

Step	Do this	You should see
	press release, a news article, a public project, a case study). Try queries like `big data turismo Baleares`, `analítica de datos agua contadores inteligentes`, `big data puerto`	
2	Pick the most reliable one: who publishes it, when, and does it name the company and the data? Copy the URL and write the full citation.	One source you trust
3	Fill in a table: which **V's** are present and why (at least two), which kinds of data (structured, semi-structured, unstructured), and which **type of analysis** it is (descriptive, diagnostic, predictive, prescriptive).	A five-row table
4	Say whether the case uses **ML** or **DL**, or only statistics, and how you know from the source. If the source does not say, write "not stated".	One sentence
5	Now your own company: does it have, or could it have, a big data problem? Justify it with at least two V's. Write one **predictive question** it would like to answer with its data.	Two short paragraphs
6	Hand in the document on Classroom (one per pair, both names).	

5.2.6 For your project

Steps 3 to 5 are raw material for **section B (mini data-science pipeline)** of Deliverable 5: section B starts from a question, and the question you write today is a good candidate. The real case you found can also go into the group part as context for the sector. Cite it.

5.2.7 Key takeaways

- **Big data** are data that traditional tools cannot handle efficiently because of their **volume**, **velocity** or **variety**; **veracity** and **value** complete the usual five V's.
- Data can be **structured**, **semi-structured** or **unstructured**; variety means mixing them.
- Big data is not "lots of data", and not every company needs it.
- **Deep learning** is inside **machine learning**, which is inside **AI**; big data is the raw material, and **data analysis** is asking the data questions.
- The four types of analysis answer four questions: what happened, why, what will happen, what should we do.

5.2.8 Quick check

1 Which three V's appear in almost every definition of big data?

A Value, veracity and visualisation
B Volume, velocity and variety
C Volume, value and validity
D Velocity, veracity and variability

2 Review text, photos, sensor readings and booking tables, all mixed in one system. Which V is that?

A Volume
B Velocity
C Variety
D Value

3 A gestoria keeps its 800 clients in one spreadsheet. Is it big data?

- **A** Yes, because it contains personal data
- **B** No: few, structured, stable data that a normal database handles well
- **C** Yes, because any customer list is big data
- **D** It depends only on the size of the file in megabytes

4 Which statement about AI, ML and DL is correct?

- **A** Machine learning is inside deep learning, which is inside AI
- **B** Deep learning is inside machine learning, which is inside AI
- **C** Big data is a kind of AI
- **D** AI, ML and DL are three names for the same thing

5 "How many bookings will we get next week?" Which type of analysis answers it?

- **A** Descriptive
- **B** Diagnostic
- **C** Predictive
- **D** Prescriptive

6 Which V is about whether the data can be trusted?

- **A** Velocity
- **B** Veracity
- **C** Volume
- **D** Variety

7 A water company in Mallorca reads thousands of smart meters every hour and also receives phone complaints about leaks. Pick the two statements that are true.

- **A** Velocity and variety weigh heavily: frequent readings plus voice or text complaints
- **B** It cannot be big data because water is an OT system
- **C** A model that warns of a likely leak before it is reported would be predictive analysis
- **D** It is big data only if it uses deep learning

You know what big data is and how it feeds analysis and AI. Next you will do the analysis yourself, step by step, on a real dataset, and decide where the data should live: Data science and the cloud.

Solutions

1. B · The classic three, proposed by Doug Laney in 2001, are **volume**, **velocity** and **variety**. Veracity and value were added later to make the usual five.

2. C · Many **formats and sources** at once, structured and unstructured: that is **variety**.

3. B · No traditional tool stops working here. The data are few, **structured** and stable. Big data is defined by the problem, not by the word "data".

4. B · **DL** (many-layer neural networks) is a part of **ML** (learning from data), which is a part of **AI**. Big data is the raw material, not AI.

5. C · "What will happen?" is the **predictive** question. "What should we do about it?" would be prescriptive.

6. B · **Veracity** is about quality and reliability: fake reviews, duplicated customers and badly calibrated sensors all lower it.

7. A, C · New sector, same tools. Frequent meter readings mean **velocity**, readings plus complaints mean **variety**, and warning before it happens is **predictive analysis**. Big data does not require deep learning.

5.3

Data science and the cloud

Everybody "knows" that August is the busiest month in the Balearic Islands and that most tourists come from Germany and the United Kingdom. But how many, exactly? Is the UK ahead of the rest of Spain? How much quieter is January? You could guess, or you could download the official figures and find out in twenty minutes with a spreadsheet.

That is what data science looks like from the inside: a clear question, the right data, a lot of cleaning, a chart and one honest sentence at the end. In this lesson you will do it with real data from **Ibestat**, the Balearic statistics institute. Then you will look at where companies keep their data in the cloud, and at the one rule every backup plan starts with.

What you will learn

- Describe the stages of **data science** and how each one depends on the previous one
- Name the main **objectives** of data science in a company
- Take a real **open dataset** from a question to a chart and a conclusion in Google Sheets
- Tell the main kinds of **cloud storage** apart and the procedures that protect data in them
- Design a backup plan with the **3-2-1 rule**
- Avoid two traps: skipping the cleaning, and thinking a synced folder is a backup

5.3.1 The stages of data science

> **Data science** combines statistics, programming and knowledge of the business to extract information and knowledge from data. It follows a process.

Stage	What you do	Hotel example
1 · Ask the question	Understand the business problem and what you want to know or decide	Which sales channel should we invest in?
2 · Get the data	Find and collect them: databases, files, APIs, sensors, open data	Export the bookings from the hotel system
3 · Clean and prepare	Remove duplicates, handle empty or impossible values, unify formats	Delete a repeated booking; unify "ES", "spain", "Spain"
4 · Explore and analyse	Summarise, group, chart, look for patterns	Bookings and average rating per channel
5 · Model	If needed, train a model (ML) to predict or classify	A model that estimates cancellations
6 · Communicate and deploy	Explain the results to whoever decides and put them to use	A dashboard for the management team

The process is **iterative**: exploring the data often sends you back to look for more data or to rewrite the question. And the stages are chained: **bad cleaning ruins everything after it**. One of the best-known reference models, **CRISP-DM** (*Cross-Industry Standard Process for Data Mining*, version 1.0 published in 2000), also has six phases and also loops back. Not every project reaches stage 5: many questions are answered with a good descriptive analysis.

In the city-builder game this is the moment you stop admiring the overlays and ask one precise thing: "why is the north district always jammed at 8 a.m.?" Then you go and get exactly the data that can answer it.

What companies use data science for

Objective	Example
Decide with data, not by intuition	Open the second restaurant only on the weeks that need it
Know the customer and personalise the offer	Different emails for families and for couples
Optimise processes and cut costs	Plan cleaning shifts from the arrivals forecast
Forecast demand, breakdowns or customers leaving	Predict which pump will fail next month
Detect anomalies: fraud, errors, intrusions	A card used in two countries in ten minutes
Create new products and services from data	Sell anonymised mobility statistics to a town council

> **Careful: "data science is building models"**
> Most of the time goes into stages 2 and 3, getting and cleaning the data. IBM's documentation of CRISP-DM estimates that data preparation usually takes between 50% and 70% of a project's time and effort. Stages 2, 3 and 6 are, to a large extent, development work: connectors, cleaning scripts, dashboards and APIs. That is your future job.

5.3.2 In class: from question to chart with real open data

Work **in pairs**, one computer, in **Google Sheets**. **No AI tools.**

The question: *Which three places of residence sent the most tourists to the Balearic Islands in 2025, and how much does the number of tourists change between the busiest and the quietest month?*

Step	Do this	You should see
1 · Get	Go to ibestat.es, search for *Turistes amb destinació principal les Illes Balears per lloc de residència. Illes Balears i mes* and open it. Use the **Descarregar** button to download it as CSV. If the download fails, use `ibestat-tourists-by-residence.csv` from Resources, which is the same table (2023 to August 2026).	A CSV file with periods, places of residence and numbers of tourists
2 · Get	In Google Sheets: *File* → *Import* → *Upload*, choose the file, *Replace spreadsheet*. Write down the source and the download date in a cell.	About 660 rows and three columns
3 · Clean	Find the problems before you count anything. Look for: rows that are **yearly totals** mixed with monthly rows; rows that are **aggregates** (*Món…*, *Altres…*) that would be counted twice; **empty** values; the **order** of the rows. Write down every problem and what you decided.	At least four problems, each with a decision
4 · Analyse	Keep only the 12 months of 2025. Make a **pivot table** (*Insert* → *Pivot table*): rows = place of residence, values = SUM of tourists. Sort it from largest to smallest.	Germany first, then the United Kingdom and the rest of Spain

Step	Do this	You should see
5 · Analyse	For the row *Món (exclosa Illes Balears)* (all tourists), list the 12 months of 2025 and find the busiest and the quietest month.	A peak in summer and a valley in winter
6 · Represent	Insert a **bar chart** of the top places and a **line chart** of the 12 months. Give each a title and the source.	Two charts
7 · Communicate	Write **one sentence** of conclusion for a hotel manager and **one limitation** of the analysis.	Two sentences
8	Hand in the spreadsheet link (shared with the teacher) on Classroom.	

Stage 5, the model, is skipped on purpose: this question does not need one. Keep empty cells empty: replacing a missing value with 0 would make the average lie.

5.3.3 For your project

In **section B of Deliverable 5** you repeat this pipeline on your own: a question that matters to your project company, a **real open dataset related to its sector** (Ibestat, the Balearic Government's open data catalogue, datos.gob.es or the INE), the cleaning decisions, one chart and a conclusion with its limitation. Today's sheet is your model.

5.3.4 Where the data live: cloud storage

In Unit 3 you saw the cloud as a service model (IaaS, PaaS, SaaS): renting the water and electricity from a utility instead of building your own plant. Now look at it as the place where the data are kept.

Kind	How it stores data	Used for
Object storage	Whole files with metadata inside containers (**buckets**), reached through an API	Images, videos, backups, logs, data lakes
Block storage	Volumes that a virtual machine mounts as a disk	The system disk and database disk of a server
File storage	A shared folder system over the network	Folders shared between teams or applications
Managed database	A SQL or NoSQL database that the provider administers	The operational data of an application
Data warehouse	Structured data already prepared for analysis	Business reports and dashboards
Data lake	Raw data in any format, at large scale	Big data and training models

The usual procedures when data are stored in the cloud:

Procedure	What it protects against
Replication in several locations	A data centre going down (availability)
Backups and versioning	Deleting or overwriting by mistake, ransomware
Encryption at rest and in transit	Somebody reading the data who should not
Access control: identities, roles, least privilege, access logs	Too many people able to touch too much
Storage tiers: hot (frequent access, more expensive) and cold or archive (rare access, cheaper)	Paying "hot" prices for data you never read; it matches the archive stage of the lifecycle
Choice of region: where the data physically are	Legal problems with data leaving the EU (sovereignty, Unit 3)

Why the cloud matters for data: capacity on demand (scale up and down), pay per use without buying hardware, access from anywhere, high availability and analysis and AI services ready to use. Without the cloud, working with big data and AI would need an infrastructure that most small companies cannot afford.

> **Careful: the data are still yours**
> Moving data to the cloud does not move the responsibility. Under the **shared responsibility model**, the provider is responsible for the security **of** the cloud (buildings, hardware, its network); the company is still responsible for what it puts **in** the cloud: who has access, what is encrypted and how long it is kept.

5.3.5 The 3-2-1 backup rule

> **3-2-1 rule**: keep **3** copies of any important data (the original and two backups), on **2** different types of media, with **1** copy off-site.

It was popularised by the photographer Peter Krogh in his books on digital asset management, and the US government's computer emergency team (US-CERT, now part of CISA) recommended it in a 2012 paper, *Data Backup Options*. In the city-builder game it is keeping the city plans in the town hall, a copy on a different kind of storage in the fire station, and another copy in a vault in a different town.

Copy	Example for a small hotel company
1 · Original	The hotel management system's database in its cloud region
2 · Second copy, different medium	A nightly backup to object storage, or a NAS (a network disk) in the office
3 · Off-site copy	A copy in another region or another provider, ideally one that cannot be changed for 30 days (immutable)

> **Careful: a synced folder is not a backup**
> If a file is deleted or encrypted by ransomware in a folder synced with Google Drive or OneDrive, the sync faithfully copies the damage everywhere. A backup is a **separate** copy you can go back to, and you only know it works if you have **tested a restore**. The classic small-company disaster, where the only backup sat on the same server that got encrypted, breaks the "1 off-site" part of the rule.

5.3.6 Key takeaways

- **Data science** follows chained stages: question, data, cleaning, analysis, model (only if needed) and communication; it is **iterative**.
- Cleaning takes most of the time, and bad cleaning ruins every result after it.
- Companies use data science to decide with data, know customers, optimise, forecast, detect anomalies and create new services.
- The main kinds of **cloud storage** are object, block, file, managed databases, **data warehouses** and **data lakes**, protected by replication, backups, encryption, access control, storage tiers and the choice of region.

- The **3-2-1 rule**: three copies, two types of media, one off-site. A synced folder is not a backup; a backup you have never restored is a hope.

5.3.7 Quick check

1 What is the first stage of a data-science project?

- **A** Making a chart
- **B** Asking the question: understanding the business problem
- **C** Training a model
- **D** Buying more storage

2 In the Ibestat table, yearly totals ("2025") are mixed with monthly rows. What happens if you add them all up for 2025?

- **A** Nothing: the result is the same
- **B** You count every tourist twice, because the yearly row already adds up the months
- **C** The spreadsheet refuses to add them
- **D** You get the average instead of the total

3 A value is missing in the data. What is the best decision?

- **A** Replace it with 0 so the formulas work
- **B** Leave it empty and note it as a limitation, so it is not counted as a real zero
- **C** Delete the whole column
- **D** Invent a reasonable number

4 Which kind of cloud storage keeps whole files in buckets, reachable through an API?

- **A** Block storage
- **B** Object storage
- **C** A data warehouse
- **D** File storage

5 What does the 3-2-1 rule ask for?

- **A** 3 passwords, 2 users, 1 administrator
- **B** 3 copies of the data, on 2 types of media, with 1 copy off-site
- **C** 3 cloud providers, 2 regions, 1 data centre
- **D** Backups every 3 days, kept 2 weeks, tested once a year

6 A company keeps all its documents in a folder synced with a cloud drive and says it is backed up. Ransomware encrypts the laptop. Pick the two true statements.

- **A** The sync can copy the encrypted files to the cloud too, so it is not a real backup
- **B** The cloud provider is responsible for recovering the files under the shared responsibility model
- **C** A separate copy that cannot be changed, and a tested restore, would have protected it
- **D** It is safe because the cloud is always encrypted

7 A water company wants to know which pumps are likely to fail next month. Which stage of data science might it need that today's tourism question did not?

- **A** Asking the question
- **B** Cleaning
- **C** Modelling: training a model that predicts failures
- **D** Communicating

Your data now have a home and a backup plan. The last step of the unit is protecting them from people who should not get in, and knowing which laws require it: Cybersecurity and regulation.

Solutions

1. B · Everything else depends on the **question**: it decides which data to get, how to clean them and which chart answers it.

2. B · The yearly row is an **aggregate** of the monthly ones. Counting both doubles the total. Spotting aggregates is part of **cleaning**, stage 3.

3. B · A 0 is a real value and would pull totals and averages down. Leaving it **empty** and saying so is honest cleaning.

4. B · **Object storage** keeps files with their metadata in **buckets**. It is the usual place for images, logs and backups.

5. B · Three copies (the original and two backups), two different media, and one copy **off-site**, so that one fire, theft or ransomware attack cannot destroy all of them.

6. A, C · A synced folder mirrors **damage** as well as changes. Under the **shared responsibility model** the data are still the company's responsibility. A separate, ideally immutable, copy and a tested restore are what make a backup.

7. C · Today's question was **descriptive**, so no model was needed. Predicting failures is a **predictive** question, and that is where stage 5, the model, usually comes in.

5.4

Cybersecurity and regulation

Have you ever received an email saying "your package is waiting, click here to pay the customs fee"? Or a message from "your bank" asking you to confirm your password? Now imagine that the person who clicks is not you but the receptionist of a hotel, and the computer behind the click holds the passports of every guest of the season.

Every company you have studied this year depends on its data: bookings, orders, meter readings, patient records. In the city-builder game, all those info views are worthless if anybody can walk into the town hall and change the records, or switch off the power station. Cybersecurity is the police, the walls and the locks of the city. And just like a real city, there are laws that say which locks are compulsory. This lesson covers both, and you will check the locks of a real website yourself, from the outside, without touching anything.

What you will learn

- Explain what security protects: **confidentiality**, **integrity** and **availability**
- Sort security measures into **device**, **system** and **global** level
- Run a **passive** security check of a public website: security headers, cookies before consent and the privacy policy
- Check whether your own email address appears in a known **data breach**
- Explain what the **RGPD**, the **LOPDGDD** and **NIS2** regulate, and decide whether each one applies to a company
- Avoid the trap: securing every laptop is not enough if there is no tested backup

5.4.1 What we protect: confidentiality, integrity, availability

> **Information security** protects information in any form, including paper and people. **Cybersecurity** is the part that protects networks, systems and digital data against attacks and incidents.

Both protect the same three things, often called the **CIA triad** (in Catalan and Spanish, *CID*):

Pillar	What it guarantees	It breaks when…	In the city
Confidentiality	Only authorised people can access the data	A customer database leaks	Strangers reading the town hall's files
Integrity	The data are correct and nobody changed them without permission	Somebody alters prices or marks	Somebody forging the land registry
Availability	Data and systems work when they are needed	An attack or a failure takes the service down	The power station switched off

Remember the Unit 1 lesson on IT and OT: in OT (the plant, the pumps, the machines) **availability** and physical safety usually come first, because a stopped water plant is worse than a leaked spreadsheet. In office IT, confidentiality often gets most attention. It is a general rule, not a law.

5.4.2 Three levels of security

Level	Scope	Typical measures	In the city
Device	Each device: computer, phone, server, sensor	Updates, antimalware, disk encryption, screen lock, strong passwords and **two-factor authentication** (a second proof, such as a code on your phone)	A lock on every door
System	The whole organisation: network, applications, people, processes	Firewall and network segmentation, least privilege, **tested** backups, monitoring, an incident response plan, staff training	The police, the walls, the fire brigade's plan
Global	Sectors, countries and the supply chain	Laws and directives, standards, public incident-response teams, international cooperation, security requirements for suppliers	The laws of the country and the cooperation between cities

Examples of the global level: **ENISA**, the EU Agency for Cybersecurity, and in Spain **INCIBE**, the National Cybersecurity Institute, whose free and confidential helpline **017** serves citizens, companies and professionals.

> **Careful: the device level is not enough**
> A perfectly protected laptop does not save a company with no backups. A firewall does not stop an employee from typing their password into a fake page. The **human factor** (**phishing**: fake messages that trick you into clicking or giving your credentials) crosses all three levels. Security is a system, not a product.

5.4.3 In class: a passive security check of a real website

You will look at your project company's public website the way any visitor's browser does, and nothing more.

> **Passive only**
> You only **look** at what the website already shows to any visitor: its response headers, the cookies it sets in your own browser, its cookie banner and its privacy policy. **Never** scan ports, try passwords, test forms with strange input or use any "hacking" tool against a real company. Doing that without written permission can be a crime under the Spanish Criminal Code. The two online tools below only make an ordinary visit to the home page, like your browser does.

Work **individually**. **No AI tools.**

Step	Do this	You should see
1 · HTTPS	Open the company's website. Click the icon left of the address and look at the certificate: who issued it and until when is it valid?	A padlock or "connection secure", an issuer and an expiry date

Step	Do this	You should see
2 · Headers	Go to securityheaders.com, paste the address and scan. Write down the grade and which headers are **missing**.	A grade from A+ to F and a list in red and green
3 · Observatory	Do the same in Mozilla HTTP Observatory. Write down the score and the two most important failed tests.	A score and a list of tests
4 · Cookies before consent	Open a **private window**, open DevTools (`F12`) → *Application* (*Storage* in Firefox) → *Cookies*, and load the site. **Do not click the banner.** List the cookies already set. Then click *Reject* and reload: did any analytics or advertising cookies (names like `_ga`, `_fbp`) stay?	A list of cookies with their domains
5 · Banner	Is there a cookie banner? Is *Reject* as easy as *Accept* (same level, same size)?	Yes or no, with a screenshot
6 · Privacy policy	In the privacy policy, find: who is responsible, the contact of the data protection officer (if any), the purposes, the retention periods, the rights and how to exercise them, and the right to complain to the AEPD.	A checklist with found / not found
7	Hand in the audit sheet on Classroom: URL, date and time, screenshots, and a table of findings with one recommended measure for each.	

What the most common headers do:

Header	What it does, in one line
`Strict-Transport-Security` (HSTS)	Tells the browser to always use HTTPS for this site
`Content-Security-Policy` (CSP)	Lists where scripts and other content may be loaded from, which limits injected scripts
`X-Frame-Options`	Stops other sites from showing this one inside a frame to trick clicks
`X-Content-Type-Options`	Stops the browser from guessing file types (`nosniff`)
`Referrer-Policy`	Controls how much of the address is sent to other sites when you click a link
`Permissions-Policy`	Switches off browser features the site does not need (camera, microphone, location)

A bad grade is a finding, not proof that the site has been hacked. Write findings politely and with evidence: "on 26/04/2027 at 10:15 the site did not send `Content-Security-Policy`".

Your own email: Have I Been Pwned

Have I Been Pwned is a free service run since 2013 by the security researcher Troy Hunt. It collects the lists of accounts from known data breaches and lets you check whether **your own** email address appears in them. It asks only for the email, **never** for a password.

Check your own address (only yours). If it appears: change the password of that service and of any other where you reused it, use a different password per site (a password manager helps) and switch on two-factor authentication. **Do not** hand in your result: it is personal data. In the Classroom activity just answer "what would a company do if an employee's email appeared in a breach?".

5.4.4 The laws: RGPD, LOPDGDD and NIS2

	RGPD (GDPR in English)	LOPDGDD	NIS2
Official name	Regulation (EU) 2016/679	*Ley Orgánica 3/2018*, of 5 December	Directive (EU) 2022/2555
Type	EU regulation: applies directly in every member state, since 25 May 2018	Spanish organic law, in force since 7 December 2018	EU directive: each state must **transpose** it into its own law; deadline 17 October 2024
Protects	People, regarding the processing of their **personal data**	The same: adapts and completes the RGPD in Spain and adds **digital rights** (for example, digital disconnection at work)	The **networks and information systems** of key sectors
Who it binds	Anybody who processes personal data of people in the EU, whatever their size	Processing subject to Spanish law	**Essential** and **important** entities in the sectors of its annexes, usually from medium size up
Key ideas	Principles, legal basis, rights, privacy by design, breach notification within 72 hours	The authority is the **AEPD** (*Agencia Española de Protección de Datos*)	Risk management, incident reporting (24 h early warning, 72 h notification, final report within one month), management accountability, supply-chain security
Maximum fines	Up to €20 million or 4% of worldwide annual turnover, whichever is higher	Classifies RGPD infringements as very serious, serious or minor	Up to €10 million or 2% of turnover for essential entities; €7 million or 1.4% for important ones

RGPD: what a developer needs to know

- **Personal data**: any information about an identified or identifiable person: name, email, ID number, IP address, location. **Special categories** (health, biometric and genetic data, political opinions, religion, ethnic origin, sexual orientation…) have extra protection.
- **Principles**: lawfulness, fairness and transparency; purpose limitation; **data minimisation** (collect only what you need); accuracy; storage limitation (the deletion stage of the lifecycle); integrity and confidentiality; accountability.
- **Rights**: access, rectification, erasure, restriction, portability and objection.
- **Roles**: the **controller** decides what the data are for; the **processor** handles them on the controller's behalf (a cloud provider, a software company).
- **Privacy by design and by default**: privacy is planned when the application is designed, not patched at the end.
- **Breaches**: a breach of personal data must be notified to the AEPD without undue delay and, where feasible, within **72 hours**, unless it is unlikely to be a risk to people.

Cookies have their own Spanish rule: Article 22.2 of the information society services law (*Ley 34/2002*, LSSI) requires informed consent before storing cookies that are not strictly necessary. The AEPD's cookie guide (updated in May 2024) adds that rejecting must be as easy as accepting. That is what you checked in steps 4 and 5.

NIS2 in Spain

NIS2 replaced the first NIS directive (2016/1148). As of October 2026, **Spain has not yet transposed it**: the government approved the draft *Ley de Coordinación y Gobernanza de la Ciberseguridad* on 14 January 2025, and on 8 July 2026 the European Commission referred Spain (with Ireland, France and the Netherlands) to the EU Court of Justice for not transposing it. Meanwhile, *Real Decreto-ley 12/2018*, which transposed the first NIS directive, is still in force. **Check the BOE** when you write your deliverable: this may have changed.

Does it apply to my company?

Question	If yes
Does the company process personal data of people in the EU (customers, employees, patients)?	**RGPD** and, in Spain, **LOPDGDD** apply. Almost every company answers yes.
Does its website use cookies that are not strictly necessary?	**LSSI** Article 22.2: consent first
Is it in an NIS2 Annex I sector? Energy, transport (air, rail, water, road; ports and airports), banking, financial market infrastructure, health, drinking water, waste water, digital infrastructure, ICT service management, public administration, space	Possibly **essential** (if large) or **important** entity
Is it in an Annex II sector? Postal services, waste management, chemicals, food (industrial production, processing and wholesale distribution), manufacturing of some products, digital providers, research	Possibly an **important** entity
Is it at least medium-sized: 50 employees or more, or annual turnover and balance sheet above €10 million?	Needed, in general, for NIS2 to apply; smaller companies can still be affected as **suppliers** of an entity that is covered

A hotel chain: RGPD and LOPDGDD yes; tourism is not an NIS2 sector. A water company: RGPD and LOPDGDD yes, and drinking water is in NIS2 Annex I.

5.4.5 For your project

The passive audit is **section D** of Deliverable 5 (the fourth member's section). The question "which of RGPD, LOPDGDD and NIS2 apply, and why" is the **group part**, together with the risk table: each risk, which pillar of the triad it threatens, its level and the measure that reduces it.

5.4.6 Key takeaways

- Security protects **confidentiality**, **integrity** and **availability** (the CIA triad).
- Measures act at **device**, **system** and **global** level; the human factor (**phishing**) crosses all three.
- A **passive** audit only looks at what any visitor sees: headers, cookies, banner, privacy policy. Never scan or try logins on a real company.
- The **RGPD** (Regulation (EU) 2016/679) and the **LOPDGDD** (Ley Orgánica 3/2018) protect people's personal data; **NIS2** (Directive (EU) 2022/2555) protects the networks and systems of key sectors.
- As of October 2026 Spain has not transposed NIS2; Royal Decree-law 12/2018 is still in force.
- Cookies that are not strictly necessary need consent first (LSSI Article 22.2).

5.4.7 Quick check

1 A denial-of-service attack takes the booking website offline for a day. Which pillar is hit hardest?

- **A** Confidentiality
- **B** Integrity
- **C** Availability
- **D** None: no data were stolen

2 Encrypting a laptop's disk is a measure at which level?

- **A** Device
- **B** System
- **C** Global
- **D** None: encryption is not a security measure

3 Which of these is allowed in this course's security activity?

- **A** Scanning the company's server ports to see what is open
- **B** Trying common passwords on the staff login page
- **C** Reading the response headers and the cookies the site sets in your own browser
- **D** Sending strange input to the contact form to see if it breaks

4 Which law applies directly in every EU state without being transposed?

- **A** The RGPD, Regulation (EU) 2016/679
- **B** NIS2, Directive (EU) 2022/2555
- **C** The LOPDGDD, Ley Orgánica 3/2018
- **D** Real Decreto-ley 12/2018

5 A booking app asks for date of birth, ID number and food allergies just to book a table. Which RGPD principle does it break most clearly?

- **A** Portability
- **B** Data minimisation
- **C** Accuracy
- **D** Availability

6 When you load a site in a private window, before touching the banner, DevTools already shows a `_ga` analytics cookie. What does it suggest?

- **A** Nothing: all cookies are allowed until you click
- **B** A cookie that is not strictly necessary may have been set before consent, which LSSI Article 22.2 does not allow
- **C** The website has been hacked
- **D** The RGPD does not cover websites

7 Your project company is a medium-sized drinking-water utility in Mallorca with an online customer area. Pick the two true statements.

- **A** RGPD and LOPDGDD apply, because it processes customers' personal data
- **B** Only NIS2 applies, because water is infrastructure
- **C** Drinking water is an NIS2 Annex I sector, so it is likely an entity under NIS2 once Spain transposes it
- **D** No law applies until Spain transposes NIS2

You have now covered the whole unit. Put it together for your company in Deliverable 5 · Data and security brief.

Solutions

1. C · Nobody read or changed the data, but the service did not work when it was needed. That is **availability**.

2. A · It protects **one device**. Network segmentation or tested backups for the whole company would be **system** level.

3. C · Only **passive** checks: looking at what the site already shows to any visitor. The other three probe the company's systems without permission.

4. A · An EU **regulation** applies directly. A **directive** such as NIS2 must be transposed by each state, and the LOPDGDD and the royal decree-law are Spanish laws.

5. B · **Data minimisation**: collect only what the purpose needs. An allergy is also health data, a special category.

6. B · Analytics cookies are not strictly necessary, so consent should come **first**. It is a finding to report politely with evidence, not proof of an attack.

7. A, C · The laws **add up**: personal data bring in the **RGPD** and **LOPDGDD**, and the sector and size bring in **NIS2**. The RGPD applies today regardless of the NIS2 transposition.

5.5

Deliverable 5 · Data and security brief

Your project company already has a map of its IT and OT systems, a technology report, a cloud architecture and an AI plan. All of them run on data. Deliverable 5 answers four questions about those data: which data the company has and how they live, what it could learn from them, where they are kept safe, and how well its public front door is locked. The group then decides which laws apply and puts everything in one risk table.

In the city-builder game, this is the chapter where you list every info view the city depends on, check that the archive has copies in another town, walk around the walls looking for open gates (without opening any), and read the laws the city must follow.

What you will hand in

- One PDF per group on Google Classroom, by **Monday 03/05/2027 at 23:59**
- Four individual, signed sections (A, B, C and D) and a group part
- An individual **defence** of 2–3 minutes in session 24 (03/05/2027)

5.5.1 The rules

- **Facts must be cited.** Everything you state as a fact about the company (what data it collects, which cloud it uses, what its website sends) must have a source: the company's website or privacy policy, an annual report, the press, Ibestat, the BOE, a screenshot with date and time. What you cannot find is labelled **hypothesis**, and you explain why it is reasonable.
- **Passive only.** The security audit only **looks** at what the public website shows to any visitor: headers, certificate, cookies, banner, privacy policy. Never scan ports, try logins or send strange input to a real company's systems.
- **No generative AI** in any part of this deliverable, unless the teacher authorises it for a specific task. Unauthorised AI or plagiarism means a 0.
- **Real data only** in section B: a real open dataset, cited, with its download date. No invented numbers.

5.5.2 Sections

Each member owns and signs **one** section (put your name under its title). Pick one you have not done before in other deliverables. With 3 members, section D goes to the group part.

Section	Criteria	What to write	Owner
A · Data inventory and lifecycle	a, b	A table of at least **six kinds of data** the company handles (office and plant): what it is, where it comes from, personal or not, which **information** or decision it allows, found or hypothesis.	Individual

Section	Criteria	What to write	Owner
		Then the **six stages of the lifecycle** of the most important dataset, from capture to deletion, with the retention period if the privacy policy gives it. Mark where data become information and the riskiest stage.	
B · Mini data-science pipeline	c, d, e, h	A question that matters to the company, answered with a **real open dataset related to its sector** (Ibestat, Balearic Government open data, datos.gob.es, INE…). Show each stage: question, source and download date, cleaning problems and decisions, the analysis (a pivot table or formulas), **one chart**, one conclusion and one limitation. Then: which **objective** of data science this serves, whether the company has a **big data** problem (justified with at least two V's) and whether ML would add anything. Link the spreadsheet.	Individual
C · Cloud storage and backups	f, g	For three of the company's datasets (from section A): which **kind of storage** (object, block, file, managed database, warehouse, lake), which deployment model (Unit 3) and **region**, and why. A **3-2-1 backup plan**: where each copy lives, how often it is made, how long it is kept, who can restore it and how a restore is tested. Explain why the cloud matters for this company, with one real price or feature from a provider's page as evidence.	Individual
D · Passive security audit of the website (4th member)	i	The audit of the company's public website: certificate, **securityheaders.com** grade and missing headers, **Mozilla HTTP Observatory** score and two failed tests, cookies set **before consent**, the banner, and the privacy-policy checklist. Screenshots with date and time. For each finding: which pillar (confidentiality, integrity, availability) it affects and one recommended measure. Polite wording: findings, not accusations.	Individual (group if 3 members)
Group · Regulation	i	Which of **RGPD**, **LOPDGDD** and **NIS2** apply to the company and why (sector, size, personal data), with the official references. Spain's NIS2 transposition status, checked in the BOE on the date you write it. Two concrete consequences for the company's applications.	Group
Group · Risk table	i, b	At least **five risks** across sections A–D: the risk, the asset (data or system), the pillar it threatens, the level (device, system, global), the measure, and who in the company would own it.	Group
Group · Integration	all	A short introduction, one diagram or table that connects the four sections (the dataset from A, analysed in B, stored in C, exposed through the website in D), conclusions, the sources list, format and spelling.	Group
Group · Group log	—	Who owned which section, who had which role, what was decided and when.	Group

5.5.3 Roles

Rotate them: nobody repeats the role they had in Deliverable 4. With 3 members, the Editor role merges into the Coordinator's.

Role	In this deliverable
Coordinator	Plans the work back from 03/05, chases deadlines, hands in the PDF
Source checker	Opens every link, checks that every fact has a source and every hypothesis is labelled; checks the BOE for NIS2
Technician	Helps with the hands-on parts: the spreadsheet of section B, the DevTools steps of section D
Editor	One format and one voice, spelling, numbered figures, the sources list in a single style

5.5.4 Assessment

Deliverable 5 is the 70% of the RA5 mark. The other 30% is the class activities of sessions 20 to 24.

Part	Weight	What is assessed
Your own section	40%	The section you signed, using the rubric below
Your defence	20%	Your answers in the 2–3 minute conversation in session 24
Group part	40%	Regulation, risk table, integration and coherence, sources, format, group log

Rubric for the sections

Section	Excellent (9–10)	Adequate (5–6)	Insufficient (0–4)
A	Six or more kinds of data, clearly separated from the information they give, mostly found and cited; a lifecycle specific to the company, with deletion and retention resolved	The inventory and the six stages are there, but generic or with few sources	Confuses data with information, misses stages, or does not talk about this company
B	A real business question, a cited dataset, cleaning decisions explained, a correct chart, an honest conclusion and limitation; V's and objective well argued	The pipeline is complete but the question is vague or the cleaning is not explained	Invented or uncited data, missing stages, or a chart that does not answer the question
C	A reasoned storage choice per dataset, a complete 3-2-1 plan with frequency, retention and restore test, cloud importance backed by evidence	Correct choices with little justification; the 3-2-1 plan has gaps	No backup plan, or storage choices that do not fit the data
D	All checks done passively, with dated evidence, findings linked to the triad and a sensible measure for each	Most checks done, evidence or measures partial	Checks missing, no evidence, or anything other than a passive check
Group part	Correct and justified application of the three laws, a risk table that uses all four sections, one coherent docu-	Laws identified but weakly justified, risk table generic, some inconsistencies	Laws wrong or missing, no risk table, sections that do not fit together

Section	Excellent (9–10)	Adequate (5–6)	Insufficient (0–4)
	ment, every fact cited, complete log		

The intermediate marks (7–8 and 3–4) are for work between two columns.

5.5.5 Deadline and how to hand in

- **One PDF per group**, uploaded to the Classroom assignment by the Coordinator, by **Monday 03/05/2027 at 23:59**. Include the link to the section B spreadsheet, shared with the teacher.
- Late work: maximum mark 5; more than 48 hours late, 0.
- Keep the shared document's **version history**: it shows who wrote what.

5.5.6 The defence

In session 24 the group works on the final version while the teacher calls each of you, one at a time, for **2–3 minutes**. No paper, no computer: a conversation. You will get two questions about **your own section** (why you chose that dataset, what a header does, why that copy is off-site) and one about **another member's** section, to check that the group really put the document together. If you wrote your section yourself, you will be fine.

5.5.7 Final check before you hand in

Check	Done?
Every section has its owner's name	
Every fact about the company has a source; every hypothesis is labelled	
Section B uses a real open dataset, cited with its download date, and the spreadsheet link works	
Section C has all three copies of the 3-2-1 plan and a restore test	
Section D has dated screenshots and only passive checks	
The group part says which laws apply and why, with the NIS2 status checked in the BOE	
The risk table has at least five risks drawn from the four sections	
The group log says who did what and which role each person had	
No generative AI has been used	



06

Unit 6 · Digital transformation project

In a city-builder game, there is a moment when you stop placing one road at a time and open the master plan: which neighbourhoods get renovated first, which pipes have to connect them, who will run the new metro, and in what order the works happen. Your group has spent the whole year building the chapters of a real company's dossier. This unit is the master plan: you put D1 to D5 together behind the company's own goals, and you defend the result in front of the class.

What you will learn

- Find the **strategic goals** of a real company in its own public sources and write them so they can be measured
- Measure where the company starts with a real **digital maturity** self-assessment and decide which areas go first
- Check how the digitalised areas **fit** with the rest, and choose technologies from needs, not from fashion
- Find the **security gaps** of each area, define the data and their indicators, and draw how everything **integrates**
- Compare the skills the plan needs with the skills the staff have, and document every change in a **roadmap**
- Present your own part of the plan in 10 minutes as a group, and answer questions about it

Lessons

#	Lesson	What it covers
1	Goals and areas	Strategic goals with sources, the Open DMAT maturity test, the priority matrix
2	Fit and technologies	Interfaces between areas, present and future needs, the area × technology matrix
3	Security, data and integration	Security gaps per area, indicators, the integration diagram in diagrams.net
4	People and documentation	Skills gaps against DigComp 3.0, the change log and the Gantt roadmap
5	Final presentation	The D6 brief, sections and owners, presentation format, assessment

6.1

Goals and areas

You open a city-builder game and the tutorial says: “build something”. You place a road, a few houses, a shop. Two hours later the traffic is a mess and the money has run out. What was missing was not more buildings. It was a goal.

Companies have the same problem. A company that “goes digital” without knowing why buys tools that nobody uses. So every transformation plan starts in the same place: **where does the company want to be in two or three years?** In March 2025, Meliá, the hotel group with its head office in Palma, published its 2024 results and said that Melia.com and its other own sales channels already brought in **50% of its centralised sales**. That sentence is not about technology. It is about strategy, and everything digital Meliá does has to serve it.

What you will learn

- Tell a **strategic goal** from a tool, and write goals that can be measured
- Find a company’s goals in its **public sources** and cite them
- Measure where a company starts with a real **digital maturity** self-assessment
- Map the production, business and communication **areas** and align each one with a goal
- Rank the areas with a **priority matrix**, and avoid the trap of digitalising whatever looks most old-fashioned

6.1.1 Goals, not tools

> A **strategic goal** says where the company wants to be, by when, and how you will know it got there.

In the city-builder, “build a metro” is not a goal. “Halve the commuting time to the city centre within five years” is. The metro is one way to get there; buses might be another.

The same test works for companies. The most widespread checklist is **SMART**, first published by George T. Doran in 1981: a goal should be **S**pecific, **M**easurable, **A**chievable, **R**elevant and **T**ime-bound.

Badly written	Why it fails	Well written
“Reduce waste.”	Not measurable, no deadline	Cut rejected or unsold product from 6% to 4% of production in 18 months
“Install a B2B ordering portal.”	It is a tool, not a goal	Receive 70% of hotel orders online within 12 months

Badly written	Why it fails	Well written
“Be more sustainable.”	Too vague to check	Cut oven energy per kilo of product by 10% in 2 years

The right-hand column belongs to the example company of this unit: a **fictitious** industrial bakery in Marratxí with 80 people, which bakes bread and pastries for hotels and supermarkets. It has an office (IT) and a bakery floor with ovens, cold rooms and a packing line (OT).

> **Careful: “our goal is to implement a CRM”**
> It is the most common mistake in this unit. A CRM, an ERP or an AI model are tools. If your “goal” names a product, ask “and what for?” until the answer is something the company’s management would put in a report.

6.1.2 Where real goals are written

You are not inside the company, so you cannot ask the management. You can read what the company publishes about itself.

Source	Where to find it	Who publishes it
Annual report or **sustainability report**	“Investors”, “Shareholders” or “Sustainability” on the corporate website	Large companies; listed ones also file them with the CNMV
Results press release	“Press room” or “News”	Large companies, every quarter or year
News in the local or business press	Última Hora, Diario de Mallorca, Ara Balears, Expansión, Cinco Días	Any company that makes news
“About us” or “Our strategy” pages	Corporate website	Almost every company
Job offers	LinkedIn, InfoJobs, the company’s careers page	A company hiring a “data analyst” is telling you something

Two rules from the course project apply here with full force:

- Every fact about the company carries its **source**: title, author or outlet, date and a working link.
- What you cannot find is written as a **hypothesis**, and labelled as such.

6.1.3 Where does the company start?

Before deciding what to digitalise, you need to know how digital the company already is.

> **Digital maturity** is how far a company already uses digital technology in its strategy, its processes and its people.

The European Union has a free tool for this: the **Open DMAT** (Digital Maturity Assessment tool) of the European Digital Innovation Hubs network. It works in the browser, in Spanish and the other EU languages, and it scores a company in six dimensions:

Dimension	The question behind it
Digital business strategy	Does the company have a plan for digital, and money for it?
Digital readiness	Which technologies does it already use?

Dimension	The question behind it
Human-centric digitalisation	Are the staff trained and involved?
Data governance	Are the data managed and protected?
Automation and artificial intelligence	Are machines and software doing repetitive work?
Green digitalisation	Does digital help the company use less energy and material?

In the city-builder, this is the info view that colours each neighbourhood from red to green before you start renovating.

> **You are not the company**
> The test is designed for the company's own managers. You answer it with what you learned in D1 to D5, and every answer you had to guess is a hypothesis. Write in the dossier which ones they were.

6.1.4 Areas, aligned with goals

> An **area** is a part of the company that does one kind of work: the bakery floor, purchasing, sales, administration. **Communication areas** are the channels that connect people inside and outside the company: website, phone, messaging, email, and the network underneath.

Aligning means writing, for each area, which goal it serves. In the example bakery:

Area	Type	Serves goal
Hotel orders (phone and WhatsApp today)	Business + communication	Online orders
Ovens	Production (OT)	Energy
Quality and food safety (paper temperature logs)	Production (OT)	Waste
Administration (ERP already in place)	Business (IT)	Online orders
Website and messaging with customers	Communication	Online orders

6.1.5 The priority matrix

You cannot renovate every neighbourhood at once. Two questions decide the order: **how much does this area move a goal?** and **how digital is it already?**

> A **priority matrix** scores each area on two axes and multiplies them, so the areas with high impact and low maturity come first.

Score impact from 1 (low) to 3 (high), and maturity from 1 (low) to 3 (high). Then: priority = impact × (4 – maturity).

Area	Impact	Maturity	Priority
Hotel orders	3	1	**9**
Ovens	3	1	**9**
Quality	2	1	6
Administration	2	3	2

Orders and ovens go first. Administration already has an ERP: it is not a bad area, it just has less to gain.

> **Careful: “digitalise whatever looks most old-fashioned”**
> A paper notebook that nobody depends on is not a priority, however analogue it looks. An area climbs the matrix only if it moves a goal. That is why goals come first.

6.1.6 In class: goals, maturity and matrix

You work on your own company, partly alone and partly with your group.

Step	Do this	You should see
1 · alone	Find one strategic source of your company that no one else in your group has used (annual report, results release, news, job offer)	A document with a date and a working link
2 · alone	Copy one sentence or one figure that says where the company wants to go, and turn it into a SMART goal	A goal you could check in a year
3 · group	The Technician opens the Open DMAT, chooses Spanish, and the group answers together with D1–D5 open	A result per dimension; a screenshot or PDF of it
4 · group	In your shared spreadsheet, one row per area (at least four production or business areas and one communication area): goal, impact, maturity, priority formula	Two or three areas at the top
5 · group	Decide who signs each section of D6 and the roles, and write it in the group log	The group log for D6

Hand in on Classroom (individual class activity): your source card (source, sentence or figure, SMART goal, source or hypothesis), with your group’s Open DMAT result and matrix attached. Generative AI is not allowed in this activity.

6.1.7 For your project

This lesson builds **section 1 · Strategy and areas** of D6 · Transformation plan (criteria a, b and c). The areas come from the departments you described in D1; the goals and the matrix are new.

6.1.8 Key takeaways

- A **strategic goal** says where the company wants to be and by when; a tool is only a way to get there. Use **SMART** to check it.
- Real goals are in the company’s **public sources**: annual and sustainability reports, results releases, press, job offers. Cite them; label the rest as **hypothesis**.
- **Digital maturity** is how digital the company already is. The EU’s **Open DMAT** scores it in six dimensions.
- Map the production, business and **communication areas**, and **align** each one with a goal.
- A **priority matrix** (impact × room to improve) decides which areas go first. Old-fashioned is not the same as important.

6.1.9 Quick check

1 Which of these is a well-written strategic goal?

A Implement a CRM in the sales department
B Become a more digital company
C Receive 70% of hotel orders online within 12 months
D Improve customer service

2 Your company does not publish an annual report. Where can you still find signs of its strategy?

- **A** Nowhere: without an annual report you must invent the goals
- **B** Press news, the “about us” page and its job offers
- **C** By phoning the company and asking for internal documents
- **D** In the Open DMAT results of other companies

3 What does a digital maturity self-assessment such as the Open DMAT measure?

- **A** How fast the company’s internet connection is
- **B** How much the company spends on computers
- **C** How far the company already uses digital technology, dimension by dimension
- **D** Whether the company complies with the RGPD

4 In the priority matrix, which area goes first?

- **A** Low impact on the goals, high maturity
- **B** High impact on the goals, low maturity
- **C** The area with the most employees
- **D** The area with the oldest computers

5 Which of these is a communication area?

- **A** The bakery ovens
- **B** The warehouse
- **C** The phone and messaging channels the hotels use to place orders
- **D** The accounting software

6 A classmate wants to put “digitalise the paper visitors’ book at reception” at the top of the matrix because it is the most old-fashioned thing in the company. What is the best answer?

- **A** Agree: the more analogue an area is, the higher its priority
- **B** First check which goal it serves; if it serves none, it is not a priority
- **C** Agree, as long as the Open DMAT says the company is not mature
- **D** Remove it: paper areas can never be digitalised

7 You answer the Open DMAT for a real hotel chain and guess how many staff have had digital training. What must you do in the dossier?

- **A** Nothing: the test result is official once submitted
- **B** Leave that question blank in the dossier and say nothing
- **C** Say which answers were guesses and label them as hypotheses
- **D** Ask an AI to estimate the real figure

You know where the company wants to go and which areas go first. Next, check whether those areas fit with the rest, and choose their technologies: Fit and technologies.

Solutions

1. C · Only the third one is **SMART**: specific, measurable (70%), relevant and time-bound (12 months). A CRM is a **tool**, and the other two cannot be checked.

2. B · Small companies rarely publish reports, but they appear in the **press**, describe themselves on their **website** and reveal plans in their **job offers**. The course project uses only **public information**; whatever you deduce is a hypothesis.

3. C · **Digital maturity** covers strategy, technology in use, people, data, automation and green digitalisation. The result tells you where the company **starts**, which is what the matrix needs.

4. B · Priority = impact × (4 – maturity). An area that **moves a goal a lot** and is **barely digital** has the most to gain.

5. C · **Communication areas** are the channels between people, inside and outside the company, and the network that carries them. The ovens and the warehouse are production; accounting is business.

6. B · An area climbs the matrix only if it has **impact on a goal**. Low maturity alone is not enough. Goals first, then areas.

7. C · You are not the company, so some answers are **hypotheses**. The project rule is to say so. And generative AI is not allowed in assessed work.

6.2

Fit and technologies

You order a pizza in an app. The app works perfectly: you pay, you get a confirmation. Twenty minutes later the restaurant calls you because the order never reached the kitchen. The app was digital. The kitchen was a printer that had run out of paper.

Digitalising one area does not help if the next area cannot receive what it sends. And choosing a technology because it is fashionable, instead of because an area needs it, can be very expensive. Between 2011 and 2018 the supermarket chain Lidl tried to replace its inventory system with SAP. According to the business press, it stopped the project in July 2018 after about **500 million euros**: Lidl valued its stock at purchase prices, the standard software at retail prices, and instead of changing the process, the software was bent to fit it.

What you will learn

- Describe how a digitalised area **fits** with the others using an **interface table**
- Spot the **frictions** where a digital area meets one that still works on paper
- Tell **present** needs from **future** ones, and back a future need with a real figure
- Build an **area × technology matrix** that starts from needs and reuses D2, D3 and D4
- Avoid the trap of choosing the technology first and then looking for somewhere to put it

6.2.1 The roads between neighbourhoods

In the city-builder, you can renovate a neighbourhood beautifully. If the road that leaves it is still a dirt track, the traffic jams right at the border.

> An **interface** is the point where two areas hand something to each other: data, a document, a product, an order.

> A **friction** is an interface that wastes time or causes errors: someone retypes data, a format does not match, the same customer exists twice with different addresses.

For each digitalised area, ask four questions: **what does it receive, from whom, in what format, and what does it hand to whom?**

From	To	What they exchange	Friction
Orders (digital: B2B portal)	Bakery floor (not digital)	Quantities per product and day	The floor manager prints the list and copies it onto the night-shift whiteboard
Orders (digital)	Administration (ERP)	Delivered orders to invoice	A manual export every week: errors and late invoices

From	To	What they exchange	Friction
Ovens (digital: IoT meters)	Maintenance (not digital)	Alerts about abnormal consumption	The alert goes by email to a person who is not on shift

Typical frictions: retyping data by hand, incompatible formats, duplicated data that contradict each other, and depending on a single person who "knows how it is done".

> **Careful: "we digitalise the area and that's it"**
> Digitalising an isolated area only moves the problem to its border. Every friction you find must have an answer in your plan: digitalise the other side too, connect the two systems, or accept it and say why.

6.2.2 Today and tomorrow

> A **present need** hurts today. A **future need** will appear in one to three years, because of the company's goals or because the sector is moving.

Present (hurts today)	Future (1–3 years)
Orders by phone get lost or misunderstood	Serve the new hotels of the season without hiring more office staff
Nobody knows how much each oven consumes	Large customers will ask suppliers for emissions data
Temperature logs on paper	Health inspections may ask for digital records and traceability

A future need must be backed by something outside your head. The **INE** (Spain's national statistics institute) surveys every year how companies use technology. Its press release of **22 October 2026** says that **33.6%** of Spanish companies with 10 or more employees used artificial intelligence in the first quarter of 2026, and **39.4%** paid for cloud services. If your company's sector is moving there, keeping up is a future need.

A choice that solves today but breaks tomorrow must be justified or changed. This is where scalability and pay-per-use, from D3, come in.

6.2.3 From need to technology

In the city-builder, the tech tree shows what you have unlocked. You have unlocked a lot this year: the enabling technologies of D2, the cloud of D3 and the AI of D4. Now you pick from that menu.

> An **area × technology matrix** lists, for each priority area, the need, the technology that answers it, the goal it serves and the chapter of the dossier it comes from.

The rule is one line: **from need to technology, never the other way round.**

Area	Need	Technology	Goal	Comes from
Orders	Receive orders in writing, without errors	B2B ordering portal as a service (SaaS), connected to the ERP	Online orders	D2, D3
Ovens	Know the consumption of each oven and shift	IoT energy meters and a dashboard	Energy	D2, D3

Area	Need	Technology	Goal	Comes from
Orders	Predict how much bread the hotels will need	Demand forecasting model on the order history	Waste	D4
Quality	Reliable temperature records	Sensors in the cold rooms with alarms (edge at the bakery)	Waste	D2, D3

Name the **type** of solution ("a B2B portal as a service"). If you name a specific product, say why that one.

> **Discarding is a decision too**
> Your D2 probably covered technologies that do not fit any goal: a bakery hardly needs blockchain or augmented reality. Write one discarded technology and the reason. It shows that you chose, instead of copying a catalogue.

6.2.4 In class: sector figure, interfaces and matrix

Step	Do this	You should see
1 · alone	Find **one** external figure about your company's sector that backs a future need: the INE press release on ICT in companies, Ibestat, a sector report or a news item	A figure, its source with date and link, and the goal it connects to
2 · group	In the shared spreadsheet, the interface table: for each priority area, what it exchanges with at least two other areas (one digital, one not), in what format, and the friction	Three or more rows, each with a friction
3 · alone	Fill **two rows** of the area × technology matrix: area, need, type of technology, goal, chapter it comes from	Two rows that start from a need
4 · group	Add one discarded technology and why	One row with a reason

Hand in on Classroom (individual class activity): your future-need card and your two matrix rows, with the group's interface table attached. Generative AI is not allowed, even if your plan proposes AI for the company.

6.2.5 For your project

This lesson builds **section 2 · Fit, needs and technologies** of D6 · Transformation plan (criteria d, e and f). The technologies come from your D2, D3 and D4; this section is where you choose among them.

6.2.6 Key takeaways

- An **interface** is where two areas hand something to each other; a **friction** is an interface that wastes time or causes errors.
- Digitalising one area alone moves the problem to its border. Every friction needs an answer.
- **Present needs** hurt today; **future needs** come from the goals and the sector. Back them with a real figure, such as the INE survey.
- The **area × technology matrix** goes from need to technology, never the other way round, and reuses D2, D3 and D4.
- Bending a process to a tool, or a tool to a process, without checking the **fit** can cost a fortune: Lidl's SAP project was stopped in 2018 after about 500 million euros.

6.2.7 Quick check

1 What is a friction between two areas?

- **A** Any disagreement between two department heads
- **B** An exchange between areas that wastes time or causes errors, such as retyping data
- **C** A security gap in the network between two offices
- **D** The cost of a new technology

2 The orders area is now digital, but the bakery floor still works from a printed list. What happens?

- **A** Nothing: the orders area is already digital
- **B** Someone has to copy the data by hand, so there is a bottleneck and a risk of errors at the border
- **C** The bakery floor becomes digital automatically
- **D** The orders portal stops working

3 Which of these is a future need, rather than a present one?

- **A** Orders by phone are getting lost
- **B** Nobody knows how much each oven consumes
- **C** Large customers will start asking suppliers for emissions data
- **D** Paper temperature logs are hard to read

4 A group writes: “We will use blockchain. Now let’s find an area for it.” What is wrong?

- **A** Nothing: every enabling technology must appear in the plan
- **B** Blockchain is illegal in the European Union
- **C** They start from the technology instead of the need; the matrix goes from need to technology
- **D** Blockchain can only be used by banks

5 Where should the technologies of the matrix come from?

- **A** From the chapters you already wrote: D2 (enabling technologies), D3 (cloud) and D4 (AI)
- **B** From a list generated with an AI chatbot
- **C** From the most popular products of the year
- **D** Only from D5

6 According to the business press, why did Lidl’s SAP project run into trouble?

- **A** SAP refused to sell the software to Lidl
- **B** Lidl’s process (stock at purchase prices) did not fit the standard software, and the software was adapted instead of the process
- **C** The servers were stolen
- **D** The project was finished early and under budget

7 A hotel chain wants to receive 30% more guests in two years. Which two choices are better justified by that future need?

- **A** A booking and check-in system that scales with demand, paid per use
- **B** A server bought today with exactly today’s capacity
- **C** Automating repetitive guest questions so the same staff can serve more people
- **D** Keeping check-in on paper because it works today

The areas fit and have their technologies. Every new technology brings new data and new doors to lock. Next: Security, data and integration.

Solutions

1. B · A **friction** sits at an **interface**: the point where two areas hand data, documents or products to each other. Retyping, mismatched formats and duplicated data are the classic ones.

2. B · Digitalising an isolated area **moves the problem to its border**. Your plan has to answer that friction.

3. C · The first, second and fourth hurt **today**. The third will appear in the coming years: it is a **future need**, and it should be backed by a source.

4. C · The rule is **from need to technology, never the other way round**. A technology without a need is a discarded technology, and saying so is a good justification.

5. A · You unlocked the tech tree during the year: **D2, D3 and D4** are the menu. And generative AI is not allowed in assessed work.

6. B · The process and the tool did not **fit**, and adapting the software made it slower and more expensive until the project was stopped in July 2018. Fit is a decision criterion, not a detail.

7. A, C · A **future need** justifies choosing now something that **scales**. A fixed server or paper solve today and break tomorrow.

6.3

Security, data and integration

On 19 July 2024, thousands of passengers at Palma airport queued for hours at the check-in desks. There was no attack. A faulty update of CrowdStrike Falcon, a security program, had crashed about **8.5 million Windows computers** around the world, according to Microsoft. Aena switched to manual procedures. One piece that everything depended on fell, and everything that was connected to it stopped.

Every technology you added to your plan in the last lesson brings two gifts: **new data**, which you have to manage, and **new doors**, which someone has to lock. And all of it has to be connected, so you also have to know what happens when one piece fails.

What you will learn

- Find the possible **security gaps** of each area of your plan, starting from your D5 passive audit, and give each one a measure
- Define how each new kind of data is **treated** along its lifecycle
- Give each strategic goal an **indicator** with a formula and a data source
- Draw the **integration diagram** of data, applications and platforms in diagrams.net, with the **source of truth** of each piece of data
- Avoid the trap of generic measures (“install an antivirus”) that answer no specific gap

6.3.1 Locks for every neighbourhood

In the city-builder, cybersecurity is the police, the walls and the locks. A new neighbourhood without them is an invitation.

> A **security gap**, in this unit, is a weak point of an area: something that could go wrong with its data or systems, by attack, by mistake or by accident.

Do not confuse it with a **personal data breach**: that is an incident that has already happened, which the RGPD (Regulation (EU) 2016/679, article 33) says must be notified to the authority “without undue delay and, where feasible, not later than 72 hours” after the company becomes aware of it. Your plan analyses the gaps so that the breaches do not happen.

For each area: **which asset or data is there, what could go wrong, what measure you put in, and which law applies.**

Area	Asset or data	Possible gap	Measure	Law
Orders	B2B portal with hotel accounts	Hotel reception staff share one password	One account per person, MFA, access log	RGPD (contact data)

Area	Asset or data	Possible gap	Measure	Law
Ovens	IoT meters on the bakery network	Devices with factory passwords on the same network as the office	Separate network for OT, change credentials	—
Quality	Temperature records	A sensor fails and nobody notices that the records stopped	Alarm when data stop arriving, daily copy	Food safety rules
All	Third-party cloud services	Dependence on the provider	Own backups, contract and data location checked	RGPD, NIS2 if it applies

The starting point for your company is the **passive audit** of D5: the security headers, cookies and privacy policy of its public website.

> **Passive only**
> Everything you say about the real company's security comes from what its website shows to any visitor: securityheaders.com, Mozilla HTTP Observatory, the browser's DevTools, the cookie banner, the privacy policy. Never scan it, never probe ports, never try a login.

> **Careful: "install an antivirus"**
> A measure that does not answer a specific gap is decoration. If you cannot say which gap a measure closes, delete it.

6.3.2 Data in, decisions out

Each new technology produces new data: portal orders, oven consumption, cold-room temperatures. Two things must be defined for each.

Treatment. The lifecycle you used in D5: how it is collected, where it is stored, who can access it, how long it is kept and how it is deleted. If it is personal data, the RGPD and the LOPDGDD (Ley Orgánica 3/2018) apply.

Analysis. The data must tell you whether the goals are being met.

> An **indicator** (or **KPI**, key performance indicator) is a number, calculated with a fixed formula from data you have, that shows whether a goal is being reached.

In the city-builder, indicators are the info-view overlays: traffic, pollution, happiness. Each goal of section 1 needs at least one.

Goal	Indicator and formula	Where the data come from	How often
Waste	kg rejected or unsold ÷ kg produced	ERP (production) and returns log	Weekly
Online orders	Hotel orders via the portal ÷ all hotel orders	Portal and ERP	Monthly
Energy	Oven kWh ÷ kg produced	IoT meters and ERP	Monthly

6.3.3 The pipes between buildings

> **Integration** is how data travel between applications, and on which platforms those applications run.

Think in three layers:

Layer	What it holds	In the example bakery
Data	Customers, orders, consumption, temperatures	Each piece of data has **one source of truth**: the customer "lives" in the ERP
Applications	The programs that use the data	Portal, ERP, dashboard, forecasting model. They talk through **APIs**; manual exports are the last resort
Platforms	Where the applications run	Cloud services from D3, an edge device at the bakery for the sensors

> The **source of truth** of a piece of data is the one system where it is created and corrected; every other system copies it from there.

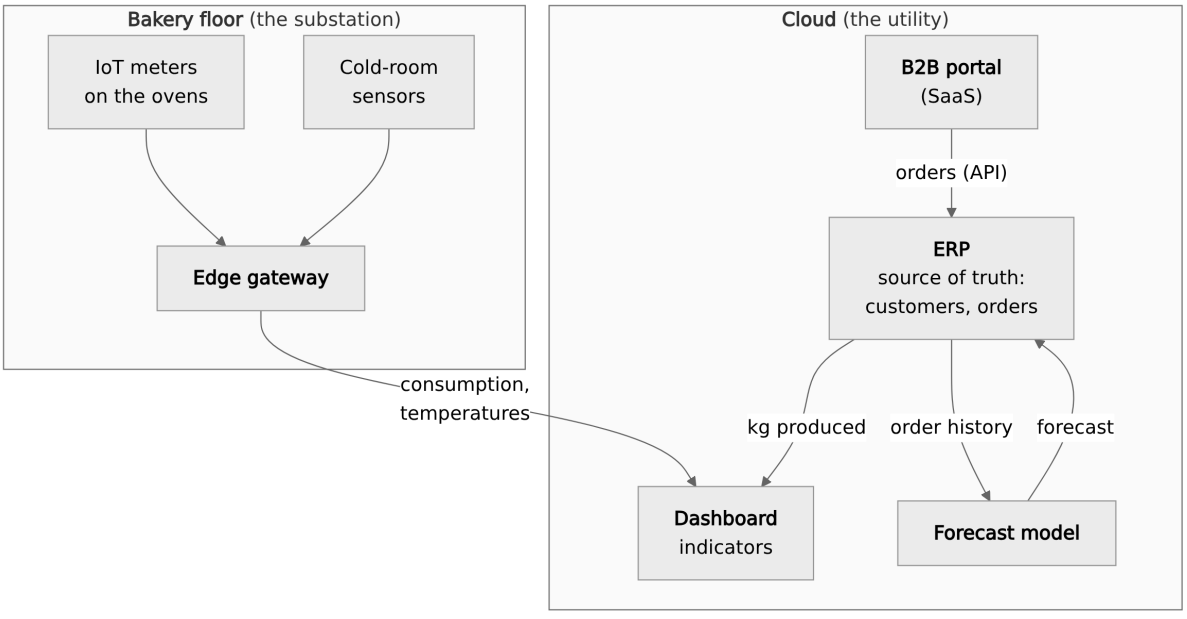
> An **API** (application programming interface) is a door that one program opens so another program can ask it for data or send it data, without a person in between.

Here is the example bakery's integration diagram, the same kind you will draw:

Check every diagram with three questions: **where is the source of truth? how does the data travel? what happens if one piece falls?** The Palma airport story is the answer to the third question when nobody asked it.

6.3.4 In class: diagram, gaps and indicators

Step	Do this	You should see
1 · group	The Technician opens app.diagrams.net, chooses **Device** (save to this computer) and creates a blank diagram. The others dictate from D3 and section 2	An empty canvas



Step	Do this	You should see
2 · group	One box per application, grouped by platform (cloud, edge, local); arrows labelled with **what** travels and **how** (API, file, by hand); the source of truth of each piece of data marked	A diagram someone outside the group could read
3 · group	**File › Export as › PNG**, and also save the `.drawio` file	Two files: the picture and the editable diagram
4 · alone	Take **one area** of the plan and write two possible gaps, the asset affected, one measure for each and the law involved. At least one must come from your D5 passive audit	Two rows in the shared gaps table
5 · alone	Write the indicator of **one** goal: name, formula, where the data come from in the diagram, how often	One indicator you could actually calculate

Hand in on Classroom (individual class activity): your two gap rows and your indicator, with the group's diagram as PNG attached. Passive only, and no generative AI.

6.3.5 For your project

This lesson builds **section 3 · Security and data** of D6 · Transformation plan (criteria g and h) and the integration diagram of **section 4** (criterion i). The gaps start from your D5 audit; the platforms come from your D3 architecture. If D3 no longer fits the technologies of section 2, update it and write the change in the change log.

6.3.6 Key takeaways

- A **security gap** is a weak point of an area; a **personal data breach** is an incident that already happened, notifiable within 72 hours under the RGPD. Plans analyse gaps to prevent breaches.
- For each area: asset, possible gap, a specific measure, the law involved. Start from your **passive** D5 audit.
- Every new kind of data needs its **treatment** defined along its lifecycle.
- Every goal needs an **indicator** (KPI) with a formula and a data source you actually have.
- **Integration** has three layers: data, applications, platforms. Each piece of data has one **source of truth**, and applications talk through **APIs**.
- Always ask what happens if one piece falls: on 19 July 2024, one faulty update stopped check-in at Palma airport.

6.3.7 Quick check

1 What is the difference between a security gap and a personal data breach?

- **A** None: they are two names for the same thing
- **B** A gap is a weak point that could be exploited; a breach is an incident that has already happened
- **C** A gap only affects OT and a breach only affects IT
- **D** A breach is always caused by an external attacker

2 Which of these measures is well matched to a specific gap?

- **A** "Install an antivirus on every computer", for every area
- **B** "One account per person and MFA", for a portal where hotel staff share one password
- **C** "Buy a faster router", for paper temperature logs
- **D** "Hire a lawyer", for IoT devices with factory passwords

3 For your company's real website, which of these is allowed in the course project?

- **A** Scanning its open ports to find vulnerable services
- **B** Trying common passwords on its login page
- **C** Reading its security headers in securityheaders.com and its cookie banner
- **D** Sending it a test phishing email

4 What makes an indicator usable in the plan?

- **A** It sounds impressive in the presentation
- **B** It measures something unrelated to the goals, so it is objective
- **C** It has a fixed formula and can be calculated with data the plan actually collects
- **D** It is calculated only once, at the end of the project

5 What is the source of truth of a piece of data?

- **A** The most expensive application in the company
- **B** The one system where that data is created and corrected; the others copy it from there
- **C** The backup copy in the cloud
- **D** Any spreadsheet where the data appears

6 In an integration diagram, what should each arrow say?

- **A** Nothing: arrows only show that two boxes are related
- **B** What travels and how: for example "orders, via API"
- **C** The price of each application
- **D** The name of the person who installed it

7 A hotel uses a cloud booking system connected to its door-key system. Which two questions should its integration diagram help answer?

- **A** If the booking system is down, can guests still open their rooms?
- **B** What colour should the boxes of the diagram be?
- **C** Which system holds the correct guest list, and how does the key system get it?
- **D** How many floors does the hotel have?

Data, doors and pipes are in place. The last pieces are the people who will run the new city, and the written plan of what changes and when: People and documentation.

Solutions

1. B · The plan analyses **gaps** to prevent **breaches**. And a breach can be caused by a mistake too, such as an email sent to the wrong person.

2. B · A measure must **close a specific gap**. Shared passwords are closed with personal accounts and multi-factor authentication.

3. C · Only **passive** observation: what the website already shows to any visitor. Never scan, probe or try logins on a real company.

4. C · An **indicator** turns data into a check on a goal. If the data are not in your diagram, you cannot calculate it.

5. B · With one **source of truth**, the customer's address is corrected in one place and every other system gets the right one. With several, they contradict each other.

6. B · An arrow without content hides the important part. Label **what** travels and **how** (API, file, by hand), so you can see the frictions.

7. A, C · The control questions of every diagram: **what happens if one piece falls**, and **where is the source of truth and how does the data travel**.

6.4

People and documentation

Your family's group chat moves to a new app. Your cousin set it up in five minutes. Your grandmother, three weeks later, still sends photos to the old one, and nobody sees them. The app works. The change did not.

Most transformation projects that fail do not fail because of the technology. They fail because the people who have to use it were not trained, not asked, or not told why. In 2023, only **56%** of people aged 16 to 74 in the European Union had at least basic digital skills, according to Eurostat. The EU's target for 2030 is 80%. If your plan needs people to do something new, you have to check whether they can, and say what you will do if they cannot. Then you write down every change, so anyone in the company can follow it.

What you will learn

- Describe a **skills gap** between what the plan needs and what the staff can do today
- Name the needed skills with the European framework **DigComp 3.0**
- Choose one of four answers to each gap: train, reassign, hire or outsource
- Explain why **change management** is part of the plan, not an extra
- Write a **change log** where every change points to a goal and an indicator, and turn it into a **Gantt chart** roadmap
- Avoid the trap of treating resistance to change as people's fault

6.4.1 Who will run the new city?

In the city-builder, you can build a metro line in one click. In a real city, someone has to drive the trains, sell the tickets and fix the escalators, and the neighbours have to learn to use it.

> A **skills gap** is the distance between what a job will need once the plan is in place and what the people in that job can do today.

For each area of the plan, ask: **who works there, what will they need to know that they do not know today, and how do you cover the difference?**

6.4.2 DigComp 3.0: naming the skills

"They need digital skills" is too vague to plan anything. The European Commission has a common vocabulary for it.

> **DigComp** is the European Digital Competence Framework, published by the Commission's Joint Research Centre (JRC). Its current version, **DigComp 3.0**, came out on 27 November 2025.

It has five competence areas, 21 numbered competences and four proficiency levels: **basic**, **intermediate**, **advanced** and **highly advanced**. It replaces DigComp 2.2 (2022), which had eight levels, so check the version of any source you read.

Area	Example competences
1 Information search, evaluation and management	1.3 Managing information
2 Communication and collaboration	2.4 Collaborating through digital technologies
3 Content creation	3.4 Computational thinking and programming
4 Safety, wellbeing and responsible use	4.1 Protecting devices · 4.2 Protecting personal data and privacy
5 Problem identification and solving	5.1 Identifying and solving technical problems

6.4.3 Four answers to a gap

Answer	When it fits	Example bakery
Train	The people are there and can learn in time	The order-desk staff learn the B2B portal from its provider
Reassign	Someone else in the company already has the skill	A person in administration who knows the ERP well supports the order desk
Hire	The skill is new, permanent and central	A part-time profile to run integrations and accesses
Outsource	The skill is specialised or needed only for a while	Maintenance of the oven meters by the installer for the first year

A row of the skills table looks like this:

Job	DigComp competence	Needed	Now (hypothesis)	Answer
Order desk	2.4 Collaborating through digital technologies	Intermediate	Basic	Train, with phone orders in parallel for one month
Bakery floor manager	1.3 Managing information	Intermediate	Basic	Train internally to read the dashboard
All staff	4.1 Protecting devices	Basic	Uneven	Awareness session and reminders
Maintenance	5.1 Identifying and solving technical problems	Advanced	Intermediate	Outsource for the first year

> **Write about jobs, not about people**
> What you say about the staff of a real company is always a hypothesis. Talk about jobs and competences, never about named people, and write it with respect.

6.4.4 Change management

> **Change management** is everything you do so that people adopt a change: explain why, train before switching on, and support them during the first days.

Your grandmother did not refuse the new app because she is stubborn. Nobody explained why it changed, nobody showed her how, and the old one still worked. **Resistance to change** is a risk of the plan, like a server failing. You plan for it.

> **Careful: “if they don’t want to learn, that’s their problem”**
> If the people who must use the system do not use it, the plan has failed, whoever is “to blame”. A plan that ignores the staff is an incomplete plan, and criterion k) is graded.

6.4.5 Writing down what changes

In the city-builder, the master plan is a map with the works in order. In a company, it is two tables.

> A **change log** is a table with one row per change: what was there before, what there will be after, which goal it serves, how it will be measured and when it happens.

ID	Before	After	Goal	Indicator	When
C1	Orders by phone and WhatsApp	B2B portal connected to the ERP	Online orders	% of orders via the portal	Months 1–4
C2	No energy measurement	IoT meters on the ovens, dashboard	Energy	kWh per kg	Months 3–6
C3	Paper temperature logs	Cold-room sensors with alarms	Waste	% of batches rejected	Months 5–8
C4	Production by eye	AI demand forecast	Waste	kg unsold	Months 9–14

Every row must answer “why are we doing this?” with a goal. A row without a goal is a change nobody needs.

> A **roadmap** places those changes in time. A **Gantt chart** draws it: one horizontal bar per task, from its start to its end, along a timeline.

Order matters. C4 needs the order history that C1 starts to collect, so C4 cannot start first. What unblocks the rest, or gives quick results, goes first.

6.4.6 In class: skills gap, change log and Gantt

Step	Do this	You should see
1 · alone	Take **one job** affected by your plan. On the DigComp page, pick the competence that best describes what it will need. Fill one row: job, competence (number and name), level needed, current level (hypothesis, justified), answer and why	One row of the skills table
2 · group	Collect every change of sections 1–3 in the change log: ID, area, before, after, goal, indicator, responsible job, start and end month	One row per change, none without a goal
3 · group	The Technician builds the Gantt in the spreadsheet: columns for change, start, end, then one column per month; select the month cells and add conditional formatting with the custom formula `=AND(D$1>=$B2; D$1<=$C2)` (if your sheet is set to an English locale, use commas instead of semicolons)	Coloured bars, one per change
4 · group	Mark at least one dependency (which change needs another first)	An arrow or a note on the Gantt
5 · group	Split the 10-minute presentation: who presents which section, who opens, who closes with the Gantt	A running order with names

Hand in on Classroom (individual class activity): your skills-gap row, with your group's Gantt as a screenshot attached. Generative AI is not allowed.

6.4.7 For your project

This lesson builds the people half of **section 4 · Integration and people** of D6 · Transformation plan (criterion k) and the **group part**: change log and roadmap (criterion j). It is also the last workshop before the dossier is due on Sunday 6 June 2027.

6.4.8 Key takeaways

- A **skills gap** is the distance between what a job will need and what its people can do today.
- **DigComp 3.0** (JRC, 27 November 2025) names the skills: 5 areas, 21 competences, 4 levels (basic, intermediate, advanced, highly advanced).
- Four answers to a gap: **train**, **reassign**, **hire** or **outsource**. Pick one and say why.
- **Change management** (explain, train before switching on, support) is part of the plan; **resistance to change** is a risk you plan for, not a fault.
- The **change log** has one row per change, each with a goal and an indicator; the **roadmap** puts them in order as a **Gantt chart**, respecting dependencies.

6.4.9 Quick check

1 What is a skills gap?

- **A** The difference in salary between two jobs
- **B** The distance between what a job will need after the plan and what its people can do today
- **C** A security hole caused by an employee
- **D** The time between hiring someone and their first day

2 What is DigComp?

- **A** A European law on digital services
- **B** A certificate that every employee must pass
- **C** The European framework that describes digital competences, with areas, competences and levels
- **D** A tool to measure a company's digital maturity

3 The staff do not know how to use a new tool. Is hiring new people the only option?

- **A** Yes: it is always cheaper than training
- **B** No: you can also train, reassign someone who already knows it, or outsource
- **C** Yes, because DigComp requires it
- **D** No: the best option is to cancel the tool

4 A department keeps using spreadsheets after the new system is switched on. What does a good plan say about this?

- **A** It is the department's fault, and the plan does not need to mention it
- **B** Resistance to change is a risk the plan foresees, with explanation, training before the switch and support
- **C** The spreadsheets should be deleted at night without warning
- **D** It proves the new system was a bad choice

5 Which row belongs in a change log?

- **A** “We like the new portal.”
- **B** “Buy laptops.”
- **C** “C1 · Orders by phone → B2B portal · goal: online orders · indicator: % via portal · months 1–4”
- **D** “Meeting with the provider on Tuesday.”

6 In a Gantt chart, what does each bar represent?

- **A** A person in the company
- **B** A strategic goal
- **C** A task or change, from the moment it starts to the moment it ends
- **D** The budget of an area

7 A hospital wants to move its appointment system online for patients over 70. Which two actions answer the people side of the plan?

- **A** Keep phone booking for a while and help patients at reception during the first weeks
- **B** Switch off phone booking on the first day so everyone gets used to it
- **C** Train reception staff before the switch so they can guide patients
- **D** Leave the change out of the change log because it is only a website

You have every piece of the plan. Now it has to become one dossier and one 10-minute presentation: Final presentation.

Solutions

1. B · The **skills gap** is measured job by job, and the plan must say how it is covered.

2. C · **DigComp** is a shared vocabulary of digital skills from the European Commission’s JRC. Its current version, DigComp 3.0, has 5 areas, 21 competences and 4 levels. Company maturity was the Open DMAT.

3. B · There are four answers: **train, reassign, hire, outsource**. Each fits a different situation, and the plan must justify its choice.

4. B · **Change management** treats resistance as a **risk to plan for**, not a fault of the people.

5. C · A **change log** row says before, after, the **goal** it serves, the **indicator** and when. Without a goal, the change has no reason in the plan.

6. C · A **Gantt chart** draws the **roadmap** on a timeline: one bar per task. Dependencies show which bar must finish before another can start.

7. A, C · Same lesson, new place: **change management** means supporting users and **training** staff before switching on. And every change goes in the change log.

6.5

Final presentation · D6 Transformation plan

All year you have been building one city. D1 drew the map of its town hall and its infrastructure. D2 unlocked the tech tree, D3 rented water and power from the utility, D4 installed the automatic advisor, and D5 put up the info views and the locks. **D6 is the master plan**: one document that puts every piece behind the company's own goals and says what changes, in what order, with what risks and with which people. Then your group defends it in 10 minutes in front of the class.

D6 is the deliverable of RA6, which weighs **25%** of the module mark.

What you will learn

- What D6 must contain and which RA6 criterion each section answers
- Who signs which section, and what the group writes together
- How D1 to D5 feed the plan
- The format of the 10-minute presentation and how the questions work
- How D6 is assessed, the deadlines and how to recover it

6.5.1 What D6 is

A **transformation plan** for the real company your group chose in October. It is not a sixth chapter next to the others: it is the document that **integrates** them. Each earlier deliverable becomes evidence for a decision in the plan.

Two rules of the project apply to every line:

- Every fact about the company is **found and cited**: title, author or outlet, date, working link (company website, annual report, press, Ibestat, INE, BOIB...).
- What you cannot find is written as a **hypothesis** and labelled as such. Never contact the company or use internal information.

Generative AI is **not allowed** in D6, in the slides or in the class activities, unless the teacher authorises it for a specific task. Plagiarism or unauthorised AI means 0.

6.5.2 Sections and owners

Each member owns and **signs** one section. With three members, section 4 goes to the group part. Choose a section you did not own in D1–D5 if you can.

Section	Criteria	What to write	Feeds from	Owner
1 · Strategy and areas	a, b, c	Three SMART goals with their sources; the Open DMAT result by dimension, with hy-	D1 (departments, IT systems, OT)	Individual

Section	Criteria	What to write	Feeds from	Owner
		potheses marked; the map of production, business and communication areas aligned with the goals; the priority matrix		
2 · Fit, needs and technologies	d, e, f	The interface table with frictions and answers; three present and three future needs, each future one backed by a real figure; the area × technology matrix; one discarded technology and why	D2 (enabling technologies), D3 (cloud), D4 (AI)	Individual
3 · Security and data	g, h	Two or more security gaps per area with asset, measure and law, starting from the passive audit; the treatment of each new kind of data along its lifecycle; one indicator per goal with formula, source and frequency	D5 (data inventory, audit, regulation)	Individual
4 · Integration and people	i, k	The integration diagram (diagrams.net) with platforms, labelled flows and sources of truth, and what happens if one piece fails; the skills table against DigComp 3.0 with the answer to each gap; the change-management measures	D3, D5	Individual (4th member); with three members, group
Group part	j	Introduction (the company in half a page); the **change log** (every change with goal, indicator and dates); the **roadmap** as a Gantt chart with dependencies; a one-page summary that links goals, changes and indicators; conclusions; sources list; group log; format and spelling	All	Group

The criteria are listed in full on the unit page.

6.5.3 How D1–D5 feed the plan

Earlier deliverable	What you reuse in D6	Where
D1 · IT/OT X-ray	Departments, IT systems, OT side, the IT–OT connection	Section 1 (areas)
D2 · Technology report	The enabling technologies that fit the company	Section 2 (matrix)
D3 · Cloud architecture	Service and deployment models, edge/fog/mist, costs	Sections 2 and 4 (platforms)
D4 · AI plan	Processes to automate, data needed, AI Act risk level	Section 2 (matrix)
D5 · Data and security	Data inventory and lifecycle, passive audit, applicable laws	Section 3

Reuse means **choose and update**, not paste. If an earlier decision no longer fits, change it and record the change in the change log.

6.5.4 Roles and group log

New roles for D6, one each: **Coordinator** (plan, deadline, hands it in), **Source checker** (every claim has a working, reliable source), **Technician** (Open DMAT, diagram, Gantt), **Editor** (format, spelling, one voice). With three members, the Editor merges into the Coordinator.

The **group log** (*acta de grup*) is a short table at the end of the dossier: who owns which section, who had which role, what was decided in each session. Contributions are visible in the shared document's version history.

6.5.5 The presentation

Monday 7 June 2027, in class, in the order drawn on 31 May.

Part	Time	Who
Opening: the company and its three goals	0:30	Coordinator
Section 1	about 1:30	Its owner
Section 2	about 1:30	Its owner
Section 3	about 1:30	Its owner
Section 4 or group part	about 1:30	Its owner, or the group
Closing: the Gantt and how you will know it worked	0:30	Any member
Questions	3:00	The teacher asks **each member** at least one question about their own section, and maybe one about another member's

- Each group has a **10-minute slot**: 7 minutes of presentation and 3 minutes of questions. Time is cut at the limit.
- Every member presents the section they signed. If you only read, it counts as reading.
- Slides: six to eight, as a PDF, with your matrix, diagram and Gantt. Little text. You present from the teacher's computer.
- The questions are of four kinds: **why** (why this area first?), **where from** (is this figure a source or a hypothesis?), **what if** (what happens if the cloud provider goes down?) and **cross** (which gap does your colleague's technology open?).
- While other groups present, you fill a short **feedback form** on Classroom (one strength and one question per group). It is the class activity of the session.

If you are absent on 7 June with a justified reason, you present your part and answer the questions on another day, in front of the teacher.

6.5.6 Assessment

D6 follows the same split as every deliverable of the project.

Part of the RA6 mark	Weight	What is assessed
D6 · individual	60% of the deliverable	Your signed section (**40%**) + your part of the presentation and your answers to the questions (**20%**)

Part of the RA6 mark	Weight	What is assessed
D6 · group	40% of the deliverable	The integrated dossier, change log and roadmap, sources, format, group log, and the coherence of the presentation
The deliverable	**70%** of RA6	
Class activities	**30%** of RA6	The activity of each session from 10 May to 31 May, and the feedback form of 7 June, on Classroom

Rubric · your signed section (40%)

Level	What it looks like
Excellent	Every item of the section is there; each decision is linked to a goal; facts are cited, hypotheses are labelled; it reuses and updates the earlier deliverables
Adequate	The items are there, but some decisions are not justified or some links to the goals are weak; a few facts lack a source
Insufficient	Items are missing, decisions are listed without reasons, or the section is pasted from an earlier deliverable or from the company's website

Rubric · your part of the presentation and answers (20%)

Level	What it looks like
Excellent	You explain your decisions without reading, in your time; you answer with data from the dossier, say what is a source and what is a hypothesis, and handle the cross question
Adequate	You explain what you did but read part of it or run over a little; your answers are correct but vague
Insufficient	You read all the time or do not finish; you do not answer, or your answer contradicts the dossier

Rubric · group part (40%)

Level	What it looks like
Excellent	One story from goal to change to indicator; every change in the log has a goal, an indicator and dates; the Gantt shows dependencies; sources complete; clean format; the presentation flows from member to member and reaches the Gantt on time
Adequate	Sections correct but loosely connected; some rows without indicator; minor format or spelling problems; the presentation is a little over or a little uneven
Insufficient	Four sections that talk about different decisions or different companies; no change log or roadmap; missing sources; the presentation is cut before the end

Programme rules: RA6 needs at least **5**. Late work counts for at most 5, and more than **48 hours** late is 0.

6.5.7 Deadlines and how to hand in

What	When	Where
Class activities S25–S28	At the end of each session, or by 23:59 the same day	Classroom, individual
D6 dossier, one PDF per group	**Sunday 6 June 2027, 23:59**	Classroom, handed in by the Co-ordinator

What	When	Where
Slides, one PDF per group	**Sunday 6 June 2027, 23:59**, with the dossier	Classroom
Presentation	**Monday 7 June 2027**, in class	Classroom computer

6.5.8 If you fail RA6

The recovery is an **individual task**: you redo one section of D6 on your own (the one the teacher tells you, usually the one you signed) and you defend it in a 2–3 minute conversation with the teacher. The maximum mark is **5**. The dates are the ones the school sets for June.

6.5.9 Final check before you hand in

Check	Done?
Cover with company, group and members; who signs each section	
The four sections (three with three members) plus the group part, in the order above	
Every fact with a source; every guess labelled as hypothesis	
Matrix, diagram (PNG) and Gantt legible and titled	
Every row of the change log has a goal and an indicator	
Group log with roles and decisions	
Spelling checked; one PDF; slides as a separate PDF	
Nothing written with generative AI	



07

The project

In a city-builder game you do not win by reading the manual. You win by running a city: you lay the pipes, you fix what breaks, you plan the next district. This module works the same way. There are no written exams. Instead, from October to June, your group runs one real company on paper and writes its **digital transformation plan**, one chapter per unit.

This page is the rulebook of that project: how groups work, how to choose the company, what you hand in and when, and how each person is marked. Come back to it before every deliverable.

What you will learn

- Explain what the **digital transformation dossier** is and how the six deliverables build it
- Choose a company that will last a whole year of work
- Tell your **individual part** from the **group part**, and the four rotating roles
- Keep the **group log** and prepare for the **individual defence**
- Avoid the two traps that cost the most marks: facts without sources, and a member who disappears

What the project is

> The **digital transformation dossier** is the document where your group analyses one real company and plans how it should go digital. It grows by one chapter per unit and ends as a complete plan in June.

Each unit gives you one tool, and each deliverable applies that tool to **your** company:

D6 is not written from scratch: it integrates and improves D1 to D5. Every comment the teacher makes on an earlier deliverable is a free improvement for June.

Why a cooperative project

Real software is built by teams, and real teams have two classic problems: the person who does everything, and the person who does nothing. The project is designed to avoid both.

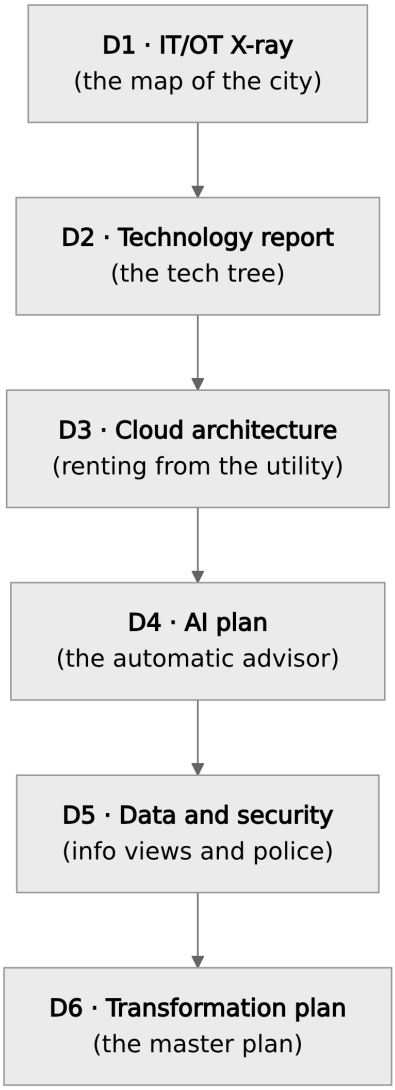
> **Cooperative learning** is a way of working in groups where each member is responsible for a piece nobody else can do for them, and the group only succeeds if every piece fits.

Two ideas make it work:

Idea	What it means	How the project does it
Individual accountability	Your mark depends on your own work, and everyone can see what you did.	Each member owns and signs one section, defends it alone with the teacher, and the shared document's version history shows every change.
Positive interdependence	The group needs everyone's piece to finish.	The group part (introduction, summary table, conclusions) cannot be written until every section is in, and it must make them all talk about the same company.

Forming the group

- Groups of **3 or 4** people, formed in session 3, on **26 October 2026**.



- The **same group all year**, with the same company. Changing group is only possible with the teacher's agreement, for a serious reason.
- With **3 members**, the section marked "4th member" in each brief goes to the group part, and the Editor role merges into the Coordinator.

Choosing a good company

You will spend eight months on this company. Choose it well: a bad choice in October makes every deliverable harder.

Criterion	Why it matters	Quick test
It is real and still operating	Every fact must be found and cited.	It has a website, and it appears in the press or in an official register.
It has an **office side (IT)** and a **physical operation (OT)**	D1 needs both, and so does the rest of the year.	Can you name one machine, plant, vehicle or building system it runs?
There is **public information** about it	Without sources, everything becomes hypothesis.	Website with "about us", annual report, news, job offers, public tenders.
Preferably from the Balearic Islands	Local sources (Ibestat, BOIB, local press) and cases you can visit.	Head office or main plant on the islands.

Criterion	Why it matters	Quick test
Not too tiny	A three-person shop has almost no IT or OT to analyse.	More than a handful of employees, or several sites.

Agree your choice with the teacher in session 3, so that no two groups take the same company.

Some real Balearic companies with both sides, as a starting point. The IT and OT columns are what a company of that kind **usually** has: until you find a source, treat them as hypotheses.

Company	Sector	Typical IT side	Typical OT side	Start here
EMAYA (Palma, owned by the city council)	Water, sewage, waste collection, street cleaning	Customer billing, incident reports, fleet planning	Water treatment plants, pumping stations, pipe network, collection lorries	emaya.es
Quely (Inca, family company, biscuits since 1853)	Food industry	Orders from supermarkets, ERP, export	Mixers, ovens, packaging lines	quely.es
Coinga (Alaior, Menorca: Cooperativa Insular Ganadera de Menorca, 1966)	Dairy: Mahón-Menorca cheese, milk, kefir	Cooperative members, sales, traceability	Milk collection from the farms, dairy plant, curing rooms	coinga.com
Bodegues José L. Ferrer (Binissalem)	Wine	Online shop, distribution, visits	Vineyards, fermentation tanks, bottling	vinosferrer.com
Ports de Balears (Autoritat Portuària de Balears: Palma, Alcúdia, Maó, Eivissa, La Savina)	Ports and logistics	Port services management, ship and passenger records	Berths, cranes, access control, lighting and port security	portsdebalears.com
A **hotel chain** based in Palma, such as RIU, Iberostar or Meliá	Tourism	Bookings, reception (PMS), loyalty, revenue	Air conditioning, laundry, kitchens, pools, building management	Each chain's corporate website

The six deliverables

Each deliverable is one PDF per group, handed in on **Google Classroom** by **23:59** on the due date. The full brief of each one (sections, rubric, checklist) is on its own page.

Deliverable	RA	What it contains	Due	Defence
D1 · IT/OT X-ray	RA1	The company, its departments and IT systems; the OT side with the automation pyramid; one IT–OT connection; the end-to-end flow and its advantages	Sun 08/11/2026	S9 · 21/12/2026
D2 · Technology report	RA2	The enabling technologies applied to the company: characteristics, areas, products and services, new markets, sustainability, plant and business, IT/OT improvements	Mon 21/12/2026	S9 · 21/12/2026

Deliverable	RA	What it contains	Due	Defence
D3 · Cloud architecture	RA3	Service and deployment models for the company's systems, cloud functions, a real cost estimate with a pricing calculator, edge/fog/mist for the plant, sovereignty; the dossier published as a static website on GitHub Pages	Mon 08/02/2027	S14 · 08/02/2027
D4 · AI plan	RA4	Processes to automate with AI, the data needed and a profitability estimate, AI in the sector, languages and tools, a working demo in the browser, the AI Act risk level of each proposal	Mon 22/03/2027	S19 · 22/03/2027
D5 · Data and security	RA5	Data inventory and lifecycle, a small data-science pipeline on a real open dataset, cloud storage and backups, a passive security audit of the company's public website, which of RGPD, LOPDGDD and NIS2 apply	Mon 03/05/2027	S24 · 03/05/2027
D6 · Transformation plan	RA6	The integrated dossier (objectives, areas, fit, technologies, security gaps, data, integration, people, documented roadmap) and a 10-minute group slot (7 min presenting, 3 min questions) where every member presents their own part	Dossier Sun 06/06/2027 · presentation Mon 07/06/2027	The presentation itself

D1 and D2 are defended together, in the workshop of 21 December. The full timetable is in the Calendar.

Individual part and group part

	Individual part	Group part
What	One **section** of the deliverable, owned and signed by one member. The sections are listed in each brief.	The common pieces: introduction, the summary table or diagram that joins all sections, conclusions, sources list, format and spelling, coherence.
Who	You alone. You choose a **different section** in each deliverable, so by June you have done every kind of work.	Everyone, coordinated by the roles.
Weight in the deliverable	**60%**: your section 40% + your defence 20%	**40%**, the same mark for every member
What ruins it	A section copied, unsourced, or that you cannot explain in the defence	Four sections that describe four different companies, or contradict each other

The four rotating roles

On top of your section, you have one **role** in each deliverable. Roles rotate: nobody is Coordinator twice in a row.

Role	What you do	You have done it well if...
Coordinator	Splits the sections, sets internal deadlines, writes the group log and hands in the PDF.	The PDF is on Classroom before 23:59 and every member knew what to do.
Source checker	Opens every link and checks every fact has a working, reliable source, and that hypotheses are labelled.	No claim about the company is left without a source or a "hypothesis" label.
Technician	Leads the hands-on IT part: the pricing calculator, the GitHub Pages site, the demo, the data analysis, the passive audit.	The practical part works and the others can repeat it.
Editor	Format, spelling, one voice, the same names for the same things.	It reads like one document, not four.

With 3 members, the Coordinator also does the Editor's job.

The group log

> The **group log** (*acta de grup*) is a short table at the end of every deliverable that records who owned which section, who had which role, and what the group decided.

It protects you: if something goes wrong, the log and the version history show who did what. Copy these two tables into your shared document and fill them in.

Who did what

Member	Section owned	Role	What I did (one line)	Signature
		Coordinator		
		Source checker		
		Technician		
		Editor		

What we decided

Date	Decision	Why	Who was there

The individual defence

> The **individual defence** is a short one-to-one conversation with the teacher about your own section of the deliverable.

It takes place in the workshop sessions (S9, S14, S19 and S24). While the groups work, the teacher calls each student for **2 to 3 minutes**. No paper, no computer: just a conversation.

- **Two questions about your own section.** "Why did you choose this technology?" "Where did you find this number?" "What would happen if...?"
- **One question about another member's section**, to check the group really integrated its work.

If you wrote your section yourself, you will find it easy. It is not a trap: it is the reason nobody can sign a section somebody else wrote. In D6 the defence is your own part of the group presentation and your answers to the questions.

How each person is marked

Part of the RA mark	Weight	What it is
Deliverable	**70%**	**Individual 60%**: your own section 40% + your individual defence 20%. **Group 40%**: integration, coherence, sources, format, group log.
Class activities	**30%**	The practical activity of each session, handed in on Classroom. Many of them feed a section of the next deliverable.

Every RA needs at least a 5. The worked examples, the recovery and the rules for late work are on Learning outcomes.

Sources and hypotheses

Everything you state as a **fact** about your company must be **found and cited**: the company's website, its annual report, the press, Ibestat, the BOIB. What you cannot find, you may still propose, but you label it.

> A **hypothesis** is a claim you believe is reasonable but have not been able to check. You write it as such: "*Hypothesis:* the plant probably uses a SCADA to supervise its ovens."

Write this	Not this
Coinga is a cooperative founded in 1966 (Coinga, *website*, consulted 26/10/2026).	Coinga is a big company with many workers.
Hypothesis: the curing rooms have temperature and humidity sensors connected to a PLC.	The curing rooms use Siemens PLCs. (Where does it say so?)

The sources we use most, and a short guide to citing them, are on Resources.

Rules that apply to every deliverable

No generative AI

Generative AI is **not allowed** in assessed work unless the teacher authorises it for a specific activity. A section written with unauthorised AI, or copied, scores **0**. The defence makes this easy to spot: you cannot explain a text you did not write.

Passive only

The security tasks only *look* at what the company's public website already shows to any visitor: securityheaders.com, Mozilla HTTP Observatory, the browser's DevTools, the cookie banner, the privacy policy. Never scan, probe ports or try to log in to a real company's systems. It can be a crime, and it scores 0.

If a member does not work

It happens. The project is built so that it does not sink the others:

- The **individual part** (60% of every deliverable) is yours alone. A member who does not write their section gets a low mark **in their section**, not in yours.
- The **group part** suffers if a section is missing. So **tell the teacher early**, the same week, not on the deadline. The teacher can then redistribute the work or adapt the group part.

- Write it in the **group log**. Together with the version history of the shared document, it shows who did what.

Key takeaways

- The project is a **digital transformation dossier** of one real company, built in **six deliverables**, one per unit; D6 integrates the other five.
- **Cooperative learning**: everyone owns one **section** (individual part) and the group joins them (group part); the four **roles** rotate in every deliverable.
- Each deliverable is **70%** of its RA: **individual 60%** (section 40% + defence 20%) and **group 40%**; class activities are the other **30%**.
- In every workshop session you have a 2–3 minute **individual defence** of your own section.
- Facts about the company need a **source**; what you cannot find is labelled **hypothesis**.
- No generative AI unless authorised, security tasks **passive only**, and if a member does not work, tell the teacher early.

Ready to start? The first chapter of your dossier begins in Unit 1.