

3. ARCHITECTURE OF BUSINESS INTELLIGENCE

Amara Atif

Learning Objectives (LO)

By the end of this chapter, you will be able to:

LO1 Understand how business intelligence (BI) architecture supports decision-making by transforming raw data into actionable insights.

LO2 Identify the main components and layers of BI architecture.

LO3 Understand the difference between OLTP (online transactional systems) and OLAP (online analytical systems) and their roles in BI architecture.

LO4 Compare and contrast various BI architecture frameworks, understanding their advantages and challenges.

LO5 Explore real-world examples of BI architecture implementation to understand how organisations successfully leverage BI systems for business success.

3 Overview of business intelligence architecture

BI architecture refers to the framework that manages the flow of data through the system to provide valuable insights. It involves various components such as data sources, Extract-Transform-Load (ETL) processes, data storage, data analytics tools, reporting systems, and more. A well-designed BI architecture ensures that data (i) flows efficiently through the system, (ii) is processed correctly, and (iii) is made available to decision-makers in a usable format.

3.1 Key concepts in BI architecture

3.1.1 What is BI architecture?

BI architecture is a **blueprint** that ***outlines the structure of BI systems and their components***. It defines how data is collected from different sources, processed, stored, analysed and presented to users. The BI architecture ensures that all parts of the BI system interact seamlessly, providing a cohesive and efficient process for generating business insights.

3.1.2 The role of BI architecture in modern business

The role of BI architecture is to support organisations in converting raw data into strategic insights. This involves the integration of various technologies and processes that help collect data from different sources, process it, and then deliver it to the appropriate stakeholders. A strong BI architecture is essential for the smooth flow of information and helps organisations make quick, data-backed decisions.

3.1.3 Components of BI architecture

As mentioned in the overview, the architecture of BI systems consists of several essential components, including data sources, ETL processes, storage systems, analytics tools and reporting systems. Each component plays a vital role in transforming raw data into actionable insights for decision-making. To fully understand how BI systems work, it is important to comprehend the underlying data processing systems that power these components. The way data is processed, stored and accessed directly influences how organisations utilise their BI systems to generate valuable insights and make informed decisions.

Example scenario: A retail store using BI architecture

A retail chain, **ShopSmart**, operates physical stores and an online platform. The company aims to optimise inventory management and make data-driven business decisions.

Data collection: ShopSmart collects raw data from multiple sources:

- point-of-sale (POS) systems (sales transactions in physical stores)

- e-commerce platforms (customer orders)
- social media (customer feedback).

Data processing (ETL): The collected data is processed through an ETL pipeline:

- **Extract** – data is pulled from all sources (sales, online, social media).
- **Transform** – the data is cleaned, standardised and aggregated for applying business logic (e.g., sales totals by region, customer sentiment).
- **Load** – transformed data is stored in a central storage system such as a data warehouse for structured analysis and a data lake for unstructured data like social media reviews.

Data analysis: With BI tools such as Tableau or Power BI, ShopSmart can analyse the processed data. For example, the analysis reveals that certain products are selling fast in specific regions and that social media sentiment around new promotions is positive.

Reporting and dashboards: Real-time dashboards provide managers¹ with updated data on sales, inventory levels and customer feedback, allowing for quick decisions on restocking or adjusting marketing strategies.

Delivering insights: These insights are delivered to key stakeholders (e.g., inventory managers and marketing teams) who use the information to make informed decisions about stock levels and promotions.

In this scenario, ShopSmart's BI architecture integrates various technologies and processes to convert raw data into valuable, strategic insights. Data is collected from multiple sources, processed for analysis and delivered to the appropriate stakeholders, enabling the company to make quick, data-backed decisions that optimise its operations.

3.2 OLTP vs. OLAP

To ensure the data in BI systems is useful and actionable, organisations need to manage two fundamental types of data processing that form the backbone of BI systems.

3.2.1 OLTP (OnLine Transaction Processing)

OLTP systems are designed for managing and processing everyday business transactions that keep an organisation running. These systems are designed for fast data entry and retrieval, ensuring real-time processing of transactional data.

1. Here the managers referred to are those who oversee various operational and strategic areas of the business. Specially, these might include inventory managers, marketing managers, sales managers or operations managers.

Transaction and transactional data

A **transaction** is an event or activity that involves the exchange or transfer of goods, services or information and **transactional data** refers to the information captured during the execution of a transaction. It represents the details of an event or interaction and is typically recorded and processed by various systems (e.g., databases, accounting software, POS systems). This data is used to track the actions taken and the outcomes of those actions, often forming the basis of operational business processes.

For example, in a sales transaction, the data might include customer details (name, contact information), product information (product ID, price, quantity), transaction date and time, payment method (credit card, cash, online payment), sales tax applied and total amount paid.

Table 3.1 provides an overview of some key characteristics of OLTP.

Table 3.1: Characteristics of OLTP

Characteristic	Description
High Transaction Volume	OLTP systems handle frequent transactions, such as sales orders, inventory updates or financial transactions.
Real-time Processing	Transactions are processed immediately, allowing businesses to update and access critical information in real time.
Data Structure	OLTP systems are designed to manage normalised, relational databases, which focus on maintaining data integrity and reducing redundancy.

Examples

Relational DBMS (e.g., Microsoft SQL Server, Oracle Database, MySQL, PostgreSQL), ERP systems (e.g., SAP ERP, Microsoft Dynamics 365, Oracle ERP), CRM systems (e.g., Salesforce, HubSpot CRM), POST Systems (e.g., Square² POS, Toast³ POS), and e-commerce platforms.

3.2.2 OLAP (OnLine Analytical Processing)

OLAP systems focus on analysing large complex datasets over time to generate insights that help guide long-term strategy and decision-making. They allow users to query large datasets, often involving multi-dimensional analysis, to understand trends, patterns and business performance. OLAP is used primarily for decision support and strategic planning.

Table 3.2 provides an overview of some key characteristics of OLAP.

Table 3.2: Characteristics of OLAP

Characteristic	Description
Historical Data	OLAP systems work with aggregated and historical data rather than transactional or real-time data.
Multi-dimensional Analysis	Data in OLAP systems is typically stored in multi-dimensional formats such as OLAP cubes, allowing users to slice and dice data along various dimensions (e.g., time, geography and product categories).
Complex Queries	OLAP queries are typically more complex, involving large datasets and requiring significant computational resources to generate reports and dashboards.

Examples

Data warehouse tools (e.g., Amazon Redshift, Google BigQuery, Snowflake), business intelligence and analytics tools (e.g., Tableau, Microsoft Power BI, QlikView, IBM Cognos Analytics), OLAP Cube tools (Microsoft SQL Server Analysis Services – SSAS, Oracle OLAP), and Excel with OLAP Add-ins (e.g., Excel with Power Pivot).

Now that we have reviewed both, can you highlight the key differences between

-
- 2. Square POS is a versatile, cloud-based system suitable for various industries, including retail, food and beverage, and professional services.
 - 3. Toast POS is specifically tailored for the restaurant industry. It provides comprehensive solutions for order management, payment processing, and customer engagement. Toast POS integrates with various hardware components, such as handheld devices for tableside ordering and kitchen display systems, to streamline restaurant operations.

OLTP and OLAP?

3.2.3 How do OLTP and OLAP work together in BI architecture?

In a typical BI system, OLTP systems provide the raw transactional data, which is then extracted, transformed and loaded (ETL) into OLAP systems for deeper analysis. While OLTP systems ensure the smooth operation of business processes, OLAP systems help with reporting, trend analysis and decision-making by providing a historical view of business data.

Scenario in the health industry: OLTP and OLAP systems in a hospital

Imagine a **hospital** that uses both OLTP and OLAP systems to manage and analyse patient data. Here is how these systems work together in this healthcare environment:

OLTP system – real-time transactional data

In the hospital, OLTP systems are responsible for handling real-time transactional data. These systems record every action that happens in the hospital's day-to-day operations. For example:

- **Patient registration:** When a new patient arrives at the hospital, the registration system records their personal details (e.g., name, age, medical history, insurance).
- **Appointment scheduling:** When a patient books or reschedules an appointment with a doctor, the OLTP system updates the hospital's calendar, records the time slot and confirms availability.
- **Billing and payment:** As medical services are rendered (e.g., tests, treatments, surgeries), the billing system processes the charges, generates invoices and tracks payments in real time.

The key here is that OLTP systems handle high-frequency, operational transactions that need to be processed immediately and accurately to ensure smooth daily operations of the hospital.

OLAP system – in-depth analysis and decision-making

Once the OLTP systems capture the transactional data, it is extracted, transformed and loaded (ETL) into OLAP systems for deeper, multi-dimensional analysis. For example:

- **Trend analysis:** OLAP systems allow the hospital to analyse patient outcomes over time, such as identifying trends in the recovery rates for certain procedures or diseases. By looking at historical data (e.g., treatment results over the past five years), hospital administrators can identify which treatments are most effective.

- **Resource allocation:** OLAP systems aggregate patient admission data, medical staff workloads and equipment usage to identify patterns in hospital resource utilisation. For example, they can reveal peak times for patient admissions, allowing the hospital to allocate resources (e.g., beds, staff) more efficiently.
- **Cost management:** Using historical financial data, OLAP tools can generate detailed reports on hospital spending, identifying trends in operational costs, patient billing and insurance payments. This helps the finance department make informed decisions about budgeting and cost control.

The OLAP systems provide strategic insights by analysing large volumes of aggregated and historical data, helping hospital administrators and decision-makers make long-term decisions on resource allocation, patient care strategies and operational improvements.

How do OLTP and OLAP work together in a hospital scenario?

In this healthcare setting, OLTP systems ensure the smooth running of day-to-day tasks such as patient registration, billing and appointment scheduling. These systems provide real-time transactional data that is crucial for the hospital to function efficiently.

Meanwhile, OLAP systems enable healthcare administrators to analyse historical data over months or years to identify patterns, trends and insights that support strategic decision-making. For example, they might identify seasonal spikes in flu cases, forecast the need for additional staff or resources, or track the effectiveness of specific treatments over time.

3.3 BI architecture components and layers

The components and layers of BI architecture are often divided into tiers to organise and structure the BI system effectively. This tiered approach helps break down the architecture into manageable sections, ensuring that each component serves a specific function in the process of collecting, storing, analysing and presenting data. Typically, these tiers are referred to as three-tier architecture, four-tier architecture or even five-tier architecture, depending on the complexity and needs of the organisation.

In the context of this book, our focus will be limited to the three-tier architecture.

3.3.1 Three-tier BI architecture

This is the most common and simplest form of BI architecture, and it consists of three primary tiers:

3.3.1.1 Tier 1: Data source layer (data collection layer)

This tier is where raw data is generated and stored. It includes transactional data (e.g., from OLTP systems) and other types of data (e.g., external data feeds, social media data, unstructured data).

Components:

- Operational databases – these include *OLTP systems* that manage transactional data (e.g., sales transactions, customer interactions).
- External data sources – this includes *APIs*, third-party data and external datasets that are integrated into the system for a broader view of data.
- Flat files – includes **CSV**, **Excel** and other structured files used for data transfer or manual uploads.
- Unstructured data – data from sources such as *social media*, *text data* and other formats that don't have a predefined structure.
- IoT devices and sensors – devices that capture *real-time data* (e.g., environmental sensors, patient monitoring devices) that feed into the system for immediate analysis.

3.3.1.2 Tier 2: Data integration, storage and modelling layer (ETL, data warehouse and data modelling)

This tier handles the process of extracting, transforming and loading (ETL) data from various sources into a centralised storage system, which could be a data warehouse, data lake, or both. It also includes data modelling to ensure that the data is structured within the storage system in a way that facilitates easy analysis and efficient querying and reporting.

Components:

- ETL tools – these tools (e.g., **Talend**, **Informatica**, **Apache Nifi**) handle the process of extracting data from multiple sources, transforming it (cleaning, standardising and enriching) and loading it into storage systems.
- Data warehouse – a system for storing structured data that is ready for analysis, typically based on relational databases (e.g., **Amazon Redshift**, **Google BigQuery**, **Microsoft SQL Server**).
- Data lake – a storage system designed to handle unstructured or semi-structured data, often cloud-based (e.g., **Hadoop**, **AWS S3**, **Azure Data Lake**, **Apache Spark**).
- Data marts – subsets or smaller, specialised versions of a data warehouse, designed for specific departmental or business unit (e.g., finance, marketing or sales) needs.
- For data modelling – common data models include *star schema* and *snowflake schema*.
- Star schema – a simplified model where a central fact table is connected to multiple dimension tables, facilitating fast queries and reporting.
- Snowflake schema – an extension of the star schema that normalises the dimension tables, reducing redundancy and enhancing data integrity.

3.3.1.3 Tier 3: Analytics and presentation layer (BI tools and dashboards)

This tier is where data is analysed and presented. BI tools are used to generate reports, perform data analysis and provide visualisations. This layer allows users to interact with the data and access actionable insights.

It bridges the gap between raw data stored in previous layers and business decision-makers who need clear, understandable insights.

Components:

- BI tools – these tools allow users to perform ad-hoc queries, generate reports and explore data visually. Examples include **Tableau**, **Power BI** and **QlikView/Qlik Sense**, which offer powerful data exploration and reporting capabilities.
- Reporting tools – these tools are specifically designed to create formal reports, often with advanced features for scheduling, distribution and formatting. Examples include **IBM Cognos** and **SAP BusinessObjects**.
- Data visualisation – data is presented in the form of *charts*, *graphs* and *dashboards*, allowing users to visually interpret trends, patterns and insights quickly.
- OLAP tools – tools such as **Microsoft SSAS** and **Oracle OLAP** allow for multi-dimensional data analysis, where users can slice and dice data across various dimensions (e.g., time, geography, product) for deeper insights.
- Advanced analytics – this includes tools and techniques such as *predictive analytics* and *machine learning models* that help identify trends, forecast future outcomes and make data-driven predictions.

Misconceptions

01 Microsoft Excel as an ETL tool

Excel itself is not traditionally considered a fully-fledged **ETL (Extract, Transform, Load)** tool, but it can be used for simple ETL processes, especially on a smaller scale. Excel's built-in features allow users to extract, transform and load data, but it lacks the advanced automation and scalability of dedicated ETL tools. Excel is not built for handling very large datasets (millions of rows), which can cause performance issues. While you can set up some automation with **VBA (Visual Basic for Applications)**, it's not as robust or scalable as dedicated ETL tools. ***Excel can serve as a useful, albeit limited, ETL tool.***

02 Microsoft Excel as a data warehouse or tool for creating a data warehouse

While **Excel** is an excellent tool for data analysis, calculations and reporting, it **cannot** be considered a **data warehouse**, nor is it typically used to **create a data warehouse** in a traditional sense.

Excel is not designed to handle the massive data volumes that a typical data warehouse manages. Excel files have row limits (e.g., 1,048,576 rows in Office 365 Excel 2024), and

performance reduces significantly when dealing with large datasets. Excel lacks the ability to automatically integrate and consolidate data from multiple sources in real-time or on a scheduled basis, which is a fundamental characteristic of a data warehouse. **ETL processes** are usually required to populate a data warehouse, something Excel cannot handle efficiently without manual intervention.

Excel is not built for concurrent access by multiple users and is not optimised for collaborative work on a large scale, which is often required for data warehouse systems. Data warehouses are optimised for running complex queries across large datasets efficiently, which Excel is not designed for. Running complex queries on a large dataset in Excel can lead to poor performance or crashes.

While Excel is **not suitable as a data warehouse** or **for creating a data warehouse**, it can still play a role in conjunction with a data warehouse. After data is loaded into a data warehouse, Excel can be used as a **front-end tool for data analysis, creating reports and visualisations**. Excel can connect to data warehouses using tools like **Power Query** or **Power Pivot**. For small projects or teams with limited data needs, Excel can be used to store and manage data for analysis, but this is not the same as a full data warehouse solution. Data extracted from a data warehouse can be exported into Excel for detailed analysis, reporting or ad-hoc querying.

03 Excel as a data mart: a limited use case

Having said that, **Excel can be used as a data mart** in specific scenarios, but it has many limitations compared to traditional data mart solutions. While Excel is not typically used as a fully-fledged **data mart** for large-scale enterprise environments, it can serve as a **simple, lightweight data mart** for smaller or less complex needs.

3.4 BI architecture frameworks

The architecture framework for BI provides various approaches to structuring and implementing BI systems. These frameworks help organisations determine the best strategies for managing data, performing analysis and making informed decisions. In this section, we will explore key BI architecture approaches.

3.4.1 Top-down vs. bottom-up approach

The top-down and bottom-up approaches are two strategies for implementing BI architecture, specifically focusing on how to organise and structure the data storage and analysis layers (such as data warehouses and data marts). Both approaches involve creating a system that consolidates data and makes it available for reporting and decision-making, but they do so in different ways.

3.4.1.1 Top-down approach

The top-down approach begins by first building a central data warehouse that acts as a single, comprehensive repository for all organisational data. The data warehouse is populated with data from various operational systems (e.g., sales, finance, marketing) and external sources. The idea is to create a unified source of truth where data from across the organisation is integrated and cleaned. Once the data warehouse is established, data marts and reporting systems are built around it.

Table 3.3 lists the advantages and challenges of the top-down approach.

Table 3.3: Advantages and challenges of the top-down approach

Advantages	Challenges
Centralised data management ensures consistency across the organisation.	The upfront effort to build a central data warehouse can be complex and time-consuming.
Reporting and analytics are based on a unified data source, reducing duplication and errors.	Requires significant planning and a clear understanding of data integration needs.
Scalable for long-term, large-scale data analysis needs.	

3.4.1.2 Bottom-up approach

The bottom-up approach, in contrast, starts by building small, departmental data marts rather than a centralised data warehouse. These data marts are designed to serve specific business functions (e.g., finance, sales or HR), and they store data tailored to the needs of the department. Once these individual data marts are established and working, they are gradually integrated into a larger, enterprise-wide data warehouse. This approach allows for quicker wins because each department can start using its data mart immediately, and they can later consolidate the data as part of a broader system.

Table 3.4 lists the advantages and challenges of the bottom-up approach.

Table 3.4: Advantages and challenges of the bottom-up approach

Advantages	Challenges
Quicker implementation since departments can work independently and don't have to wait for a centralised warehouse.	Potential for inconsistency between departments, as each data mart may have different data models.
Business units can directly control and tailor their data to meet specific needs.	As the system grows, integrating data marts into a central warehouse can be complex and require significant rework.

3.4.2 Real-time BI vs. batch processing

In BI systems, real-time BI and batch processing represent two different methods for managing how and when data is updated, processed and analysed.

3.4.2.1 Real-time BI

Real-time BI focuses on continuous, live data feeds. It involves processing and analysing data as it becomes available. This means that business intelligence reports, dashboards and insights are updated in real time (or near real time), providing organisations with the most up-to-date information. For instance, in a retail environment, real-time BI could track customer purchases as they happen, enabling immediate inventory adjustments or personalised offers.

Table 3.5 lists the advantages and challenges of the real-time BI approach.

Table 3.5: Advantages and challenges of real-time BI

Advantages	Challenges
Provides up-to-the-minute insights that can lead to quick decision-making.	Complex to implement and maintain, as it requires handling continuous data streams.
Essential for environments where timely information is critical, such as stock trading, e-commerce or customer service.	Often needs specialised tools and infrastructure (e.g., message brokers, streaming data platforms such as Apache Kafka or AWS Kinesis).

3.4.2.2 Batch processing

Batch processing involves collecting and processing data in large batches at regular intervals (e.g., daily, weekly or monthly). The data is typically extracted from various sources, transformed and loaded into a data warehouse or reporting system, often during off-peak hours to minimise system load. This approach is common for historical analysis or reporting purposes, where immediate updates are not required.

Table 3.6 lists the advantages and challenges of the batch processing approach.

Table 3.6: Advantages and challenges of batch processing

Advantages	Challenges
Simpler to implement than real-time BI, as data can be processed in bulk at scheduled intervals.	Does not provide the most up-to-date information, which can be a drawback for real-time decision-making.
Less resource-intensive during the processing phase since updates are less frequent.	Delays in data processing can lead to stale insights, especially in fast-paced industries.

3.4.3 Cloud BI architecture

Cloud BI architecture refers to the use of cloud-based technologies to manage, store and analyse data for business intelligence. Traditional BI often relied on on-premises infrastructure (refers to IT hardware and software systems that are physically located within an organisation’s facilities or data centres, rather than being hosted remotely or in the cloud), but cloud BI solutions have grown in popularity due to their flexibility, scalability and cost-effectiveness.

3.4.3.1 Cloud-based BI solutions

Cloud BI platforms allow organisations to store and analyse large amounts of data on remote servers (the “cloud”) rather than maintaining physical data centres and hardware. These platforms offer various tools and services for data integration, reporting and analytics. Popular examples of cloud BI solutions include:

- **Google BigQuery** – a serverless data warehouse for handling large-scale data analytics
- **Amazon Web Services (AWS)** – a suite of cloud computing services that includes analytics platforms such as Amazon Redshift (data warehouse) and Amazon QuickSight (BI and analytics service)
- **Microsoft Azure** – offers a variety of cloud analytics services, including Azure Synapse Analytics for data integration and Power BI for visualisation and reporting.

Table 3.7 lists the advantages and challenges of the cloud BI architecture.

Table 3.7: Advantages and challenges of cloud BI architecture

Advantages	Challenges
Cloud solutions can easily scale up or down based on the needs of the organisation, accommodating fluctuating data volumes.	Storing sensitive data in the cloud raises concerns about data privacy and security, though cloud providers often offer robust security features.
Pay-as-you-go pricing models allow organisations to only pay for the storage and computing resources they actually use.	Cloud-based BI relies on consistent internet access, and connectivity issues can hinder performance or accessibility.
Cloud platforms support a wide range of tools, including data storage, analysis and visualisation, and can be integrated with various other cloud services.	
Since the platform is hosted in the cloud, authorised users can access the data and tools from anywhere with an internet connection.	

3.5 Case studies and real-world applications of BI architecture

Several Australian companies/institutions have successfully utilised BI solutions to drive business success.

3.5.1 Australian examples of BI architecture in action

All these examples are from [Telstra Purple](#), which was launched by Telstra in 2019. Telstra Purple is the largest Australian-owned technology services business, bringing together Telstra Enterprise's business technology services capabilities and a number of its recently acquired companies focused on outcome-based, transformative tech solutions.

3.5.1.1 InfraBuild's procurement transformation

This case study exemplifies how leveraging modern digital solutions can transform procurement processes, leading to significant efficiency gains and improved operational effectiveness.

The problem

InfraBuild, Australia's only fully vertically integrated sustainable steel manufacturing and recycling business, faced challenges in its device procurement process. Employees experienced a cumbersome and time-consuming ordering system, leading to inefficiencies and a lack of visibility into device inventory and procurement status. To address these issues, InfraBuild partnered with Telstra Purple to develop an automated procurement platform.

The solution

The solution, built on Microsoft's Power Platform, offers a user-friendly interface where employees can select and order devices from a centralised catalogue. This integration streamlined the ordering process, reducing order processing time from two weeks to approximately four hours and enabling faster device deployment. The platform also provided real-time visibility into orders, invoices and device inventory, enhancing control over approved devices and preventing unauthorised purchases.

The results

The automated procurement system not only improved operational efficiency but also enhanced the onboarding experience for new employees by ensuring timely access to necessary devices and support services. Encouraged by these benefits, InfraBuild plans to expand the platform's use to include mobile

services, telecommunications, software procurement and printers, aiming to create a comprehensive procurement solution that maintains consistent control and visibility across various services.

3.5.1.2 Victoria University's digital transformation

This case study illustrates how adopting advanced digital infrastructure and smart campus solutions can enhance connectivity, security and operational efficiency, creating a seamless and secure environment for both students and staff at Victoria University.

The problem

Victoria University (VU), established in 1916, is a public institution offering both higher education and Technical and Further Education (TAFE) to over 40,000 students across multiple campuses in Melbourne, Sydney and Brisbane. As part of a multimillion-dollar project to consolidate five locations in Melbourne's CBD, VU developed the state-of-the-art VU City Tower at its City Campus.

To enhance its digital infrastructure, VU partnered with Telstra Purple and Cisco to implement a connected, digitised and smart campus. The objectives included automating IT troubleshooting, delivering next-generation network services, improving network resilience and security, maximising student engagement, and supporting a range of digital services for students, staff and industry partners.

The solution

The solution involved deploying Cisco's Software-Defined Access (SDA) network architecture, which utilises Software-Defined Networking (SDN) with physical network devices. This approach enabled policy-based automation from edge to cloud, including automated user and device policies, end-to-end segmentation, and zero-trust security across all users and devices. The implementation simplified IT operations by automating processes like device provisioning, issue identification and policy updates, resulting in better security, transparency and visibility across the campus.

The results

The successful deployment of Cisco SDN led by Telstra Purple improved and simplified VU's IT operations, offering better device security and a consistent user experience across the City Campus and from remote locations. It also reduced network management complexity and enhanced cyber threat management and visibility. The project was completed within budget and has been recognised as a significant achievement in creating a digitally enabled campus environment.

3.5.1.3 Jemena's IoT-driven gas metering transformation

This case study demonstrates how leveraging IoT-enabled smart monitoring and data analytics can transform operational efficiency and customer service, providing real-time insights and enhancing the management of gas metering systems at Jemena.

The problem

Jemena, a leading Australian utilities provider, faced challenges with its outdated gas metering systems, which impacted operational efficiency and customer service. The existing 3G and 2G meters did not provide real-time data collection or remote monitoring capabilities, limiting the company's ability to accurately track and manage gas usage across its network.

The solution

To address these challenges, Jemena partnered with Telstra Purple and Nucleus3 to develop an Internet of Things (IoT)-enabled platform. This solution involved replacing the old gas meters with advanced IoT-enabled devices, which transmit data through Telstra's Narrowband IoT network. The project began with a successful proof-of-concept deployment of 50 Eden Worth modems in high-rise locations, including areas with poor network coverage. The solution uses the Cumulocity connected device platform for centralised management, offering dashboards and analytics tools for real-time visibility of gas usage and network health.

The results

The IoT solution enabled Jemena to achieve better billing accuracy and operational efficiency. With real-time data, the company could remotely monitor gas usage, improving customer service and providing insights into consumption patterns. The project has expanded from an initial 50 units to 600, with plans to deploy around 5,000 units, covering 200,000 customers. In the future, Jemena plans to utilise this IoT infrastructure for additional applications, such as pressure monitoring and gas leakage detection, further enhancing its digital capabilities and service offerings.

3.5.1.4 University of Tasmania's campus navigation

This case study illustrates how implementing innovative digital solutions like UniMaps and

UniNav can transform the campus navigation experience, enhancing accessibility, improving user engagement and streamlining operations at the University of Tasmania.

The problem

The University of Tasmania (UTAS), with its multiple campuses across the state, faced significant challenges in helping students, staff and visitors navigate its expansive and often complex campus infrastructure. Despite offering induction sessions and support from frontline staff, these measures were not enough to improve the overall campus navigation experience. Students and visitors had difficulty finding buildings, classrooms and public transport information, resulting in frustration and inefficiencies, particularly for those new to the campus.

The solution

To address these challenges, UTAS partnered with Telstra Purple to develop an innovative digital campus navigation tool. The project began with a discovery phase, involving workshops with students, staff and other stakeholders to identify pain points and develop effective solutions. The primary objective was to provide an accessible, user-friendly digital platform that would allow students to navigate the campus easily, improve accessibility, enhance spatial data and integrate information such as timetables and transport schedules.

The solution involved creating two interconnected tools:

- **UniMaps:** A digital wayfinding platform that allows users to explore the campus from their mobile, tablet or desktop devices. This platform was integrated into the UTAS website and the UniApp mobile app.
- **UniNav:** A solution embedded into physical **QR code signage** around the campus, enabling students and visitors to scan codes and access navigation tools, making the entire campus experience more interactive and accessible.

The results

The implementation of **UniMaps** and **UniNav** greatly enhanced the campus navigation experience for students, staff and visitors. With **UniMaps**, users can now easily explore the campus and access key information on their devices, whether on mobile, tablet or desktop. **UniNav** connected the physical campus to the digital experience through strategically placed QR codes, further simplifying navigation.

The tools were seamlessly integrated into UTAS's public website, the UniApp mobile app, and physical signage, ensuring easy access to campus information and real-time updates. This transformation not only improved the overall user experience but also boosted operational efficiency, as it streamlined campus

services and made vital information more accessible. The project has contributed to a more connected, modern and accessible campus, enhancing student satisfaction and engagement.

3.5.1.5 Queensland Police Service's data-driven decisions

This case study illustrates how leveraging data-driven solutions, such as the Mobile Business Intelligence dashboard, can enhance resource allocation, improve operational efficiency and enable informed decision-making at the Queensland Police Service.

The problem

The Queensland Police Service (QPS) faced challenges in effectively managing and monitoring the usage of over 8,000 iPads deployed statewide. Limited visibility into device utilisation hindered their ability to maximise return on investment and ensure officers had the necessary tools for efficient operations. Disparate data sources provided fragmented insights, making it difficult to obtain a comprehensive view of device performance and application uptake.

The solution

To address these issues, QPS partnered with Telstra Purple to develop a Mobile Business Intelligence (MBI) dashboard powered by Microsoft Power BI. This data visualisation tool consolidated multiple data sources, offering actionable insights into device usage, application performance and user engagement. The dashboard was designed to be intuitive and user-friendly, enabling Officers-In-Charge (OIC) and Mobile Capability Centre (MCC) personnel to monitor and manage device performance effectively.

The results

The implementation of the MBI dashboard provided the QPS with enhanced visibility into device and application usage, offering several key benefits. By identifying underutilised devices, QPS was able to reallocate resources to areas with higher demand, optimising resource distribution. Additionally, monitoring application usage and device performance helped reduce costs associated with managed services and device maintenance. The dashboard also ensured that officers had access to the most effective applications and tools, thereby improving their operational efficiency. Furthermore, the dashboard empowered leadership to make data-driven decisions by providing real-time insights, ultimately enhancing the overall effectiveness of QPS operations.

3.6 Test your knowledge



An interactive H5P element has been excluded from this version of the text. You can view it online here:

<https://oercollective.caul.edu.au/decoding-data/?p=269#h5p-3>

Additional readings for Australian case studies

- Agile Analytics. (n.d.). *Tourism Australia: Azure Data Platform, Synapse, Power BI case study*. Agile Analytics. Retrieved May 28, 2024, from <https://www.agile-analytics.com.au/stories/tourism-australia-azure-power-bi-dax-solution-case-study/>
- BizData. (n.d.). *Customer case studies*. BizData. Retrieved June 2, 2024, from <https://www.bizdata.com.au/case-studies>
- CSIRO. (n.d.). *Case studies – Analytics and Decision Sciences*. CSIRO. Retrieved June 2, 2024, from <https://research.csiro.au/ads/case-studies/>
- Dear Watson Consulting. (n.d.). *Power BI case studies*. Dear Watson Consulting. Retrieved June 2, 2024, from <https://www.dearwatson.net.au/power-bi-case-studies/>
- dVT Group. (n.d.). *Case studies*. dVT Group. Retrieved June 2, 2024, from <https://dvtgroup.com.au/case-studies/>
- Expose Data. (n.d.). *Case studies*. Expose Data. Retrieved June 2, 2024, from <https://exposedata.com.au/case-studies/>
- NetSuite. (n.d.). *18 real-world examples of business intelligence*. NetSuite. Retrieved June 2, 2024, from <https://www.netsuite.com.au/portal/au/resource/articles/business-strategy/business-intelligence-examples.shtml>
- SAP Concur Australia. (n.d.). *Case studies*. SAP Concur Australia. Retrieved June 2, 2024, from <https://www.concur.com.au/casestudy>
- Telstra Purple. (n.d.). *Case studies*. Telstra Purple. Retrieved June 2, 2024, from <https://purple.telstra.com/insights/case-studies>
- WingArc Australia. (n.d.). *Customer case studies*. WingArc Australia. Retrieved June 2, 2024, from <https://wingarc.com.au/customers/case-studies/>

About the author



Dr Amara Atif

UNIVERSITY OF TECHNOLOGY SYDNEY

<https://profiles.uts.edu.au/Amara.Atif>

Dr Amara Atif is a Senior Lecturer in the School of Computer Science, Faculty of Engineering and IT (FEIT), at the University of Technology Sydney (UTS), Australia. Amara is an expert in educational technology, with a comprehensive skill set that includes research and development in the design of technology-enhanced learning environments, human-computer interaction, artificial intelligence in education, and learning analytics/educational data mining.

At UTS, she developed the undergraduate subject “Business Intelligence”, integrating culturally informed Indigenous knowledge systems through the FEIT Indigenous Graduate Attribute. She has received multiple UTS Vice Chancellor’s Learning and Teaching Awards, including recognition for data-informed teaching practices and inclusive assessment. Passionate about the accessibility and sustainability of educational resources, Amara is an advocate for Open Educational Resources (OER). *Decoding Data* continues this mission, bringing rigorous Business Intelligence concepts to learners through an open, practitioner-oriented resource grounded in Australian business contexts.