

How the Internet of Things will change the way we monitor the Railways

White paper by Tim Taberner, Eurotech UK



Introduction

In the rail industry, in common with most of the markets it serves, Eurotech sees clients and end users being increasingly frustrated with the integration issues raised by the lack of interoperability between different systems, the need to acquire data from new and 'difficult' assets, and the complexities posed by providing reliable communications to highly mobile assets operating in remote locations. We often hear comments like

'if only we could easily combine data from system A and system B we could correlate X and Y and...';

Or

'if only we could easily share this information with these other stakeholders then...';

'... it would give us a saving/increase our efficiency/provide a better customer experience'

As an example that can be easily understood, imagine you went to the facilities manager of a large station and asked him to tell you about its current status. He would say something like "Well, that's quite hard, I can go to the fire alarm system and tell you that we have no fires, and then we can go next door and I can look at the heating & ventilation system to tell you the temperature in all of the zones and that it's working properly, then we could go upstairs to look at the security system, and then in the basement we could look at the electricity & gas metering etc. etc., but I don't have just one place where a summary of all that is available"

Let's think about this example a bit more – the fire system has sensors deployed throughout a building used to detect a fire. The HVAC system similarly has deployed sensors measuring the temperature in the building. If one of these saw a temperature rising at a high rate, that's probably indicative of a fire; similarly, if there is a sudden drop in temperature in the middle of the night, that may suggest a window or door has been opened, which may be useful information for the security system. There are therefore obvious benefits in these systems being able to share common sensor data. There are similar benefits if these systems are able to interact at a higher level. If the ventilation system knew there was a fire, it could cut power to all of the heaters/air conditioning to reduce the hazard risk. Not only would we have more connected and responsive buildings, we could reduce duplication in the sensors and devices installed to feed the various systems with data.

The traditional way to get this summary information (if it's possible at all) is to introduce custom code above and in each of the relevant systems to extract the needed bits of data and then generate and present the resulting information. If, or more accurately, when at a future date the summary data requirements change, or there is a change in one of the systems sitting below it, then the custom code needs to be changed. This leads to 'software spaghetti' solutions which are extremely difficult, and therefore expensive, to support and maintain. Some solutions attempt to overcome this by requiring all of the systems to feed data into a large SQL database, only to find that this approach can lead to problems with data flooding, or with the inflexibility of the database structure once defined.

The 'Internet of Things'

The operating principles of what has become known as the 'Internet of Things' provide a potential solution to these issues. Fundamentally, it introduces the concept of data 'producers' and data 'consumers' and the notion that there is no reason why a producer of data needs to know anything about the systems that will consume that data. As an example, take the temperature sensor above. Its whole world can be summed up as 'the temperature doesn't change often, I need to report (publish) the



temperature once every x minutes UNLESS I see a rapid change in which case I need to publish it immediately'. The heating system subscribes to reports from the temperature sensor and uses them to control the temperature of the building, BUT now the fire system can also subscribe to this information to help indicate a potential fire; a summary 'building status' system can subscribe to the temperature in order to provide summary dashboards; the security system can infer intruders from unexpected drops in temperature; an external maintenance company can subscribe to the data as part of a preventative/responsive maintenance regime.....

By providing this method to share data, effectively an 'enterprise service bus for devices', we facilitate the easy propagation of data across applications, thereby enabling the interoperability outlined above. As importantly however, this model enables the enterprise level software sitting above the system to evolve away from the large and expensive monolithic systems typically deployed today, into collections of smaller components, each performing its own part of the process, subscribing to the necessary foundation data, and publishing its result for use by other applications. This dramatically increases the ability to respond to new or emerging requirements, which in turn reduces the cost, risk and deployment time associated with maintenance operations, adapting the system to integrate new technologies, or the addition of new applications.

Examples in Rail

Now think what could happen if this was applied across the rail network, and to related systems beyond it.

- A journey planner application could recommend the fastest or most comfortable current trip allowing for road conditions to the station, live train times, available car parking capacities, passenger loading etc., allowing passengers to make informed choices about what option will provide them with the best experience according to their personal circumstances, for example whether it is more important to have the shortest journey time, or to be guaranteed a seat. Allowing the inclusion of historic data will enable evaluation not only for a current trip, but also in a predictive way for a trip planned at a future date, based upon what is normal for the planned day and time of travel.
- Combining passenger loading information from trains with social networking apps will help to spread demand peaks. The same base information shared at a terminus can help to select the destination platform which offers the most efficient passenger egress considering the loadings of other inbound trains, whilst sharing the same information on the train can produce a more even distribution of passengers within the carriages, potentially allowing standing passengers to find a seat.
- Combining status information from diverse on-board public facing assets such as toilets, food car
 chillers and ovens, and presenting this to service organisations with current positional information
 can improve the customer experience and reduce the penalty costs associated with having these
 assets out of service.
- Intelligent CCTV cameras not only provide a record of events in case of an incident, they actively
 provide real time alarms of the occurrence of potential problems, allowing more timely
 intervention responses and potentially reducing service outages.
- Information concerning categorisation of faults can be analysed across multiple assets, even multiple operators, to spot trends and identify areas for preventative maintenance.



The fabric of the rail network starts to communicate with itself in a similar way to how the human population interacts via social networking, and also communicates with the human population via the same social networking tools they already use —so a train could tweet passengers to tell them it will be 5 minutes late at their station. All of this becomes possible once the interdependent link between devices and their host systems is broken via a middleware layer.

Other Monitoring Challenges - Cost per Point

Of course, no one will dispute that having the ability to monitor something will produce a benefit – normally the question is whether the cost of monitoring is justified by the benefit it produces. Historically, monitoring projects have been considered in isolation, resulting in 'silo' systems installed as specific solutions to a particular monitoring problem where the cost/benefit analysis has justified the investment. No consideration is typically given to whether the infrastructure provided by these systems can be reused to provide monitoring for other devices, and so the cost of entry to provide monitoring of new devices is often prohibitive.

These concepts indicate that a single architecture can be deployed and then be accessed by multiple devices and systems, reducing the cost of monitoring new devices and extending their usefulness. The communications technology used (and therefore the communications cost) can be optimised for each individual class of device. The incremental cost of adding a new class of device can often be measured in pence per device per month, allowing a huge number of currently unmonitored assets to be included, and for their data to be made available to new, often unforeseen, applications, processes and stakeholders.

M2M Integration Architecture Description

So what components are needed to make up an effective platform for the integration of Machine to Machine (M2M) data?

Real Time Publish and Subscribe (pub/sub) Messaging – the Broker

The heart of the system is the notion of data producers that **publish** data, and data consumers that **subscribe** to data. A physical device or application may be a data producer, data consumer or both. In between the two is a middleware layer known as a data broker. It keeps a record of all of the data subscriptions to a particular message **topic**, and when it receives a message on that topic, it immediately passes the message on to all subscribers.

The format of the topic is itself an important part of the principle of operation. It is split into a topic namespace and a data payload. The namespace is a hierarchical structure that mirrors the physical infrastructure being monitored. In rail for a specific Train Operating Company (TOC) one could imagine constructs such as:

TOC/region/train/car/passenger count TOC/region/train/car/toilet status



Infrastructure/region/station/platform/passenger count Infrastructure/region/point location/heater status

.. and so on.

Each of these reports would have a data payload associated with it, the format of which is understood by applications subscribing to that data, but about which the broker is agnostic. It simply passes the payload through to subscribers who then interpret it. Wildcards are allowed within subscription declarations, and so it is easy to get information about all the passenger counters in one region, or on one train etc., simply by substituting the appropriate wildcards in a single subscription. As an example, subscribing to TOC/region/*/*/passenger count would result in all passenger counter messages from any TOC (Train Operating Company) owned train within the region being immediately passed to the subscribing application.

Database

Of course, whilst the broker is essentially an exchange for real time information, it is often necessary to include historical data for analysis within an application. The platform therefore needs to include some form of historical archive that reproduces the flexibility inherent in the pub/sub message topic construct. This flexibility can be achieved by using a schemaless 'NoSQL' database that essentially subscribes to everything. The advantage of this approach is that the database can use exactly the same topic space naming convention to store its data. Therefore, the database does not need to be designed in advance of deployment. As new topics are published, the database simply files the data payload away in a new location table defined and accessed by the topic name. Again, the database is completely agnostic about the data payload, and leaves it to any application looking at the data to interpret it. Also, like with the broker, enquiries can be made to the database using wildcards within the topic name, meaning it is easy to aggregate data across multiple sources. Finally, the database engine is optimised for data entry and egress, unlike older SQL-based approaches, making it uniquely suited to receive massive amounts of information from a large number of sources in real time and at the same time.

Rules Engine for Real Time Data

We now have a way to pass data from large numbers of devices to applications in real time, and a way to store historical data from all of these sources for offline analysis. Whilst deriving actions from this data could be performed at application level, there are advantages in certain circumstances if this can be done within the integration platform itself. The final component of the M2M integration platform is therefore a complex events processing rules engine. This component allows users to enter queries based not upon the historical data contained in the database, but upon the real time messages themselves, with the result of a match being an action, for example sending an email, SMS or tweet, or publishing data on a new topic for consumption and action elsewhere. This means that without any external applications, the integration platform can handle notifications such as 'if this asset is not working (or is showing symptoms of an impending fault), then notify the service desk', or 'if the number of passengers exiting all trains that arrive at a terminus within a 10-minute rolling window exceeds a threshold, then slow down the arrival of the next inbound train to allow time for the platforms to clear'.



Device Management

Of course, providing the integration platform for the exchange of M2M data is only part of the story. When considering a large population of remote devices, it is also crucial to take into account how those devices will be managed through their lifetime. Device management covers issues such as authentication onto the M2M integration platform, the provisioning of remote device configuration and applications, over the air updates for upgrades and bug fixes, as well as more specialised considerations such as remote device engineer access for investigative or debug work.

Device Software Framework

These considerations are enabled by a standards based software framework within the connected edge gateways and devices. This framework abstracts the physical characteristics of the remote unit to provide a common platform onto which user business logic may be installed. The framework takes care of issues such as communications session maintenance, application download management and application start/stop, as well as providing a set of APIs (application programming interfaces) and industry-specific protocols which simplify the user's development process, allowing users to focus on their core competencies and resulting in faster, lower-risk deployment lifecycles.

Project Specific Requirements

With the integration platform and device management pieces in place, implementation of a project becomes an exercise in topic space definition, remote business logic and enterprise dashboard and data-mining applications. The use of the integration platform, providing standard, open API based access to the data from the device fabric with a flexible but standardised access mechanism means that these processes are greatly simplified. Better still is the fact that new devices, applications and dashboards can be added in the future without any impact on the existing processes and architecture, and that any new data provided by these elements is immediately available for use in 'mashups' both in the new systems and the existing applications.

Impact on System Considerations

As the 'Internet of Things' revolution permeates into rail applications, it will affect some of the design considerations when implementing a new technology deployment.

Trend to Multi-Service Gateways

One of the first elements is the notion that (e.g.) a train will need a data concentrator that acts as a gateway between the on-board devices and the global integration platform. Such devices already exist, but are typically dedicated to one functional system on board the train, such as a ticketing, CCTV, PIS or maintenance monitoring. Whilst these could remain as isolated systems, each communicating independently with the integration platform, more efficient use of communications space, and more cost



effective maintenance regimes can be realised if a single device capable of accepting remotely downloaded applications is deployed. In this case, a smaller number of un-configured spares can be held, replaced in the field by electrical technicians, and then provisioned over the air to perform the functions required in that particular installation. The gateway becomes a shared resource, a **Multi-Service Gateway**, hosting apps for potentially different stakeholders, each exchanging data over a single unified communications path to the integration platform and ultimately to the enterprise level systems interacting with that data. Implicit in this model is that the gateways deployed need to be sized not only for the immediate applications being considered, but also to have capacity to accept and run undefined applications at some time in the future.

Best of Breed Devices for Each Application

The abstraction of the data sources implicit in the architecture allows users to choose current 'best of breed' devices for each project deployment, without risking any previous investment in similar but different devices. Going back to the topic space definitions, 'TOC/region/train/car/passenger count' is a generic search term for which the subscriber doesn't need to know if the primary source is a Eurotech passenger counter or one made by another vendor. The differences in protocols between devices are handled by protocol agents on the Multi-Service Gateway, and so it becomes possible to operate mixed populations of devices, each chosen as a point in time best fit to the project requirements. Serviceable devices from a previous generation no longer need to be replaced to maintain compatibility with newer units being installed elsewhere.

How Eurotech is Serving the Rail Market

So what does Eurotech offer the rail industry that helps to realise these new technologies and capabilities?

Eurotech is an embedded computer and communications specialist working across a range of industry sectors, offering a wide technology base spanning from wireless sensors through to true supercomputers. From this portfolio, it supplies the rail industry in Europe and the US with ruggedized computer, networking and storage solutions, pre-qualified to standards such as EN50155, together with more specialised rail devices such as in-cab HMIs, intelligent signage, CCTV and passenger counters. It is now bringing these products and its expertise to the UK market, building upon the presence it has already established in other sectors

M2M Integration Platform

The Everyware Cloud is an M2M integration platform which integrates the broker, database and complex event processing rules engine elements. The Everyware Cloud is available as a public cloud service that reduces the initial investment and upon request, can be deployed in the customer's own data center. It offers standard interfaces using MQTT protocol (co-developed and open sourced by Eurotech & IBM) for pub/sub messaging, and REST APIs for database access, making integration with current enterprise and web services systems simple.

Eurotech offers software clients in a variety of languages, and for a variety of operating systems for embedding within a user's device to allow access to the Everyware Cloud, and to provide the session Page 7 of 9



awareness required to properly manage the communications channel between the device and the broker.

Multi-Service Gateways

- A range of EN50155 qualified on-board and static computers, whose modular construction makes them easy to configure for different physical interfaces and communications infrastructures.
- The Everyware Software Framework (ESF) is available on these devices, providing the device abstraction, connectivity, device management and applications management features required by Multi-Service Gateways. ESF offers native connectivity to the Everyware Cloud.
- For different applications, on-board computers feature embedded Windows or embedded Linux Operating Systems as well.

Infrastructure and Networking

 A range of EN50155-qualified infrastructure components such as routers, switches, NAS servers etc.

Passenger Counting

 Autonomous Stereoscopic vision based devices are installed over the door to count passenger entering or leaving the vehicles Eurotech also engages in the supply of turnkey Passenger Counting solutions.

PIS

- A range of EN50155 qualified devices used to make up PIS (Passenger Information System) systems including screens, high brightness LED display panels, media servers.
- Eurotech also engages in the supply of turnkey PIS solutions.

CCTV

- EN50155 qualified DVRs and video application engines providing functionality such as left object detection, crowd alerts etc.
- Intelligent cameras for trackside infrastructure applications such as boundary protection, number plate recognition, car park space availability, etc.
- Cloud-based management console for CCTV applications across a variety of sites and assets.
- Eurotech also engages in the supply of turnkey CCTV solutions.



Summary

The revolution underway in M2M communications will have profound implications for all sectors, including the rail industry. Sensors, devices, systems and enterprise applications must work collaboratively, interacting with each other, and with the supply chain, regulators and customers in ways that produce lower CAPEX and OPEX project costs, improved customer experience, and more efficient running of the network.

Eurotech, whilst a relatively new entrant to the UK rail market, can demonstrate a long history of supply to the rail sector in other regions, offering a range of products and services addressing both on board and trackside deployments in a variety of applications. Uniquely, Eurotech is able to combine its rail qualified hardware with field-proven M2M enabling technologies, evolved from mission critical deployments over many years. This combination makes it an ideal partner for organisations wishing to exploit the opportunities emerging as devices and the enterprise become more connected.

For further information on Eurotech Solutions, please visit www.eurotech.com or email sales.emea@eurotech.com