

## **Dave Milham Interview 19<sup>th</sup> December 2025**

Firstly, many thanks for agreeing to do this talk on 14<sup>th</sup> January, it will be both face to face and on a Teams Webinar.

With this interview we hope to give prospective attendees an insight into the content of the talk and some of the background.

### **Our first question is how long have you been working in telecoms and can you describe the key organisations you've worked and with?**

To my surprise, I've worked in the industry for over 55 years, which is quite astonishing. And I've worked in the telecom sector starting in universities, through working with BT when it was a monopoly, and indeed actually, when it was the GPO.

The GPO and BT were initially monopolies and then I've worked for BT through its transformation into a private company and then into a competitive telecoms environment with many other competitors. Before and since I left BT I've been working with an International Trade Association the Tele Management Forum for about 14 or 15 years. The TM Forum is US based, which focuses on telecoms operations, standards and best practises.

I've also worked voluntarily with the organisations like the IEEE Network Operations and Maintenance Symposium and also supported and managed a number of European projects involving many European operators on R&D topics.

### **So you've been working on OSS and BSS for many years. Can you tell me what the network and operations landscape was when you first joined the GPO?**

Well, that's really a trip to the past and you realise how much has happened in 50 years, but the environment when I joined the GPO networks were based on rotary mechanical switches, electromechanical exchanges, frequency division, multiplexed transmission systems and copper baseband access. On the administration side, largely the processes were manual with limited office automation capabilities like multiple carbon copies, Advice Notes and things like that.

So when you compare the environment that I started off in the in telecoms to where

we are now with computerization, digitalization, optics, AI based technologies, and also cloud computing, it's extremely primitive.

### **What were the first attempts at automation of operations?**

Well the GPO and BT as it became, were using mechanical exchanges for much longer than most comparable Telcos, particularly Rotary mechanical exchanges [Strawger]. So the opportunities for automation were quite limited because the central network elements were mechanical and analogue, and required manual work. The real transformation took place was when BT and the UK industry decided to move towards creating digital exchange systems called System X. This basically transformed the network from analogue into a digital network and it also introduced computer controls into telephone exchanges. This meant you've got the basis for creating an IT infrastructure, integrated electronically and through open protocols into these exchanges. And indeed a lot of the work was done on operations and maintenance in BT with System X. There are still the same principles in place for much of the automation that we see even today, even say 40 years later.

### **Were the first attempts at OSS standardisation around then successful, or how successful were they?**

I'd say that certainly in BT the thinking was well ahead of the game and there was a recognition of a need to move towards automated systems, probably around about the end of the 1980s. There was a need to integrate the Operations Systems, which were that time based on many computers and to integrate them both with the networks and the front-office systems which were predominantly based on general purpose IT systems. Indeed, that split between the kinds of the IT technology and thinking in the networks, and use of IT in the front office systems is actually still there in most service providers today. It's the split between responsibility of the CTO office, which tends to look at networks, and the IT that goes with the CIO and the IT department, which tends to look at front end systems and customer experience. So those divisions are still present.

In terms of the degree of success, I'd say it was quite limited. There were some areas where we had some successes particularly around integration interfaces into networks, but networks themselves evolved and some of those standards had

relatively short-lived life cycles and new technologies came out with different views on what standards should be. So there has been quite a large churn in the actual technology of interfaces.

But the business principles of what they were trying to do from an interaction point of view has hardly changed until recently, where we're seeing the emergence of autonomous networks, rather than passive managed and controlled networks which are based on traditional automation, as opposed to autonomy.

We learned a lot from those early stages, but one of the challenges, of course, was these interfaces (APIs) getting them at the right point into the marketplace because often equipment vendors decided to create interfaces ahead of the requirement for standardisation, which made it rather difficult to retrofit. But I'd say it was patchy at the start, but there were lots of good learning points came out of it and all the principles have served well for perhaps 30 years or so.

However those principles do need changing as we move towards autonomous networks where the integration models change somewhat.

**Just as a point of detail on this one, is the concept of management information base, the MIB, an enduring one? And how do you see that as having evolved and evolving?**

Well, actually this is where BT was pretty much head of the game. The Maintenance and Control Subsystem of System X introduced the idea of resources hierarchies, relationships between entities, which is essentially the basis of a MIB. What we're seeing is the evolution of MIBs from those early days towards more formal knowledge graphs using Semantic Web concepts. That's kind of what's happening with the introduction of AI at the moment.

But I would say the idea of having a formal representation of the entities, their characteristics and the relationships between them has been a long-established principle, but the technology for representing them has changed over the years.

But that underlying principle is very long standing, and even into AI based systems

you still need those kinds of models. It's just that technology now is based on basically knowledge graphs and implemented typically in a label property graph database.

### **What difference did the Internet make to telecoms OSS?**

Well, I'd say that would probably be around about the late 80s. The Internet was coming out of academia. These standards were becoming more accepted as the basis of commercial offers in the marketplace.

And certainly the latest standardisation activities, I'm thinking now in the late 1990s, most of the protocols were based on Internet thinking and particularly standards coming out of the IETF, and they were very influential. The IETF, were creating standards which were MIB based, actually for managing network devices things like routers, firewalls, those sorts of things.

However, those early standards from the IETF didn't really have particularly strong concepts around MIBs, and that's one of the reasons why all of that is becoming obsolescent, and why the IETF is moving towards model and MIB based solutions which includes some of the properties we mentioned, formal modelling, resources, attributes and relationships amongst them.

### **What would you say have been the key successes of the delivery of end-to-end processing automation enabled by the kind of architectures and standards that you will have been working on with the Telemanagement Forum, perhaps just ahead of the AI era?**

So recently and the majority of the time, we have had a very successful API programme. We seem to have hit the market at exactly the right point. There were leading companies like Telefonica and others deeply involved. Indeed, BT had quite a lot of leading thinking about moving towards a set of formal APIs and an Integration model.

What they've basically done and what's happened is to create a family of APIs which are based on the idea of automating transactions between systems that sit within an automation model.

Here we have been defining an end-to-end process, allocating some of the process steps to systems and where there are boundaries between systems using formal APIs that are transaction based. In those kinds of interfaces, we've been extremely successful. The adoption rate for the standards that we've got for APIs has been extremely high.

There are a few areas for improvement. One is that the APIs tend to be fairly generic and they do require service providers to add their own additional information models to them. As an example, the product catalogue has a definition of what a product is, but you usually have to extend the product definition to suit each individual service provider and of course the instances have to be created within the service providers' systems.

But for things like ordering, pre-order qualification, basic trouble ticketing, alarm management, all sorts of things that will be used and talked about within a process centric traditional automation integration approach. Those APIs have been particularly successful.

And actually, what's interesting is that success can be measured by the number of comments and change requests we get on APIs. So we've been quite lucky on this sort of API because we get lots of feedback from members, lots of discussions about practises for extending them for specific deployment circumstances. This demonstrates that people are really using this stuff and I'd say that TM Forum APIs and the certification programme are pretty much the core of what most OSS and BSS vendors are delivering into the marketplace at the moment and it's been quite successful.

**Moving on to AI or more properly machine learning with large language models, that seems set to revolutionise IT across many industries. How are you using those capabilities and which companies are you working with to develop autonomous networks and operations?**

So this is a very interesting topic actually, because around about the year 2000 or 2001, maybe even a bit earlier, the TM Forum kicked off a programme of work called Autonomous Networks.

And the thinking there was to create networks that were self-healing, self-maintaining, self-remediating and those kinds of concepts. And we created a fairly simple reference model which is based on the idea of autonomous domains, based on intent based interfaces and also based on the use of closed control loops. So what that allows you to do is that for an autonomous domain, say for example SDN controller, you define an intent so you describe *not* the procedural steps for setting up the SDN across multiple routers. You go to an SDN Controller and you say this is the outcome I need, I need a number of ingress and egress points,: these are the characteristics of them; these are the capacities that I need between them; these are the SLAs that I need. So you describe the end state to the controller rather than the intermediate steps, which is what you do with the traditional integration mechanisms.

Now the companies that are working heavily in this area, Ericsson and Huawei, who are probably the two leading companies. Ericsson recently ran a conference in London called the OSS/BSS Summit, which is available on the Internet if you're prepared to look for it. There was some really good input from lots of service providers picking up these kinds of concepts and running with them. In particular, our Australian colleagues, NBN Co, which is the National Broadband Network and Telstra are pretty well advanced with introducing the ideas of self-healing networks, controllers, the use of control loops, and to support these you need to have in knowledge graphs that are formally defined. A lot of that work is moving towards the use of standards coming out of the Semantic Web, and those standards are Resource Description Framework, Ontology Web Languages. I'm not going to go into the details of these, but that's the kind of the technology base that's being picked up.

Behind those leadership companies, we're seeing massive numbers of companies getting into this discussion. Within the discussions we've got, we're seeing ZTE very active in this area. We've got tremendous inputs from Orange in France, Telefonica

in Spain, and Deutsche Telekom are also putting in contributions in this space. In our last face to face technical meeting in China, the number of contributions coming in from the Chinese operators, that's China Mobile, China Unicom and China Telecom was quite remarkable. Their intended use of AI covers virtually the whole of telecoms operations, right from customer experience, chat bots, agents working on behalf of customers. And that's AI agents working on behalf of customers, right the way through to operational efficiency where networks are self-healing and self-optimising.

So a good example of this is radio access networks, where based on the situation in the radio access network you might have for example have fuses on power supplies fail. You might have a period of limited power so now you need to optimise the way in which the RAN is being used and what traffic you're accepting on the RAN. Based on that situation, and when it remediates, you'll completely redo the thinking.

There's a lot of work going on in planning and design and the use of AI to support planning and design activities. We've seen several TM Forum Catalyst demonstrations over the last year, all looking at how you can augment and improve AI based network planning. We saw planning of 5G networks on long major high speed rail networks in China, where information has been collected from drones, from line plant records, from geospatial information, from satellite information and also on planning practises that are embedded into the AI Agents.

They come up with an optimal placing of masts, antenna positions and what the antenna power level should be. There's a big emphasis on energy efficiency, particularly in radio access networks, and that's an area where there's a lot of work going on in AI.

So I'd say it's quite pervasive actually and in the background, some of the Japanese companies are doing quite a bit of work in this area, and also some of the Korean companies are quite active, particularly SK Telecom, who are looking at the use of AI not only for their own internal operations but as the basis of services that they deliver to customers, so it's remarkably active, but the challenges are that the operational model for managing autonomous agents looks to be quite different from the traditional automation model.

So the idea of defining an end-to-end process, and then at design time automating each of those steps statically, isn't what's happening with AI agents. Some of them are supporting process steps, but some of the agents are effectively replacing all this sequential analysis by an approach which is more like a delegation model. So for example, if you're running a shop, you could be supervising people and telling them do this, do that, do the other. But in a delegation model you're saying achieve this outcome, which is more like a managerial model of delegation, supervision, goals, targets and that's a different integration model; and these models have to co-exist, which makes things a little bit more complicated. But it does mean that some of the practises that have been involved over the last 40 or so years do have to be adjusted quite a bit.

The software components that TM Forum is defining do have to be able to support embedded AI decision making in them. And at the moment, most components are regarded in terms of they do something quite static. You send them a task with some parameters and then it does what it's told and it's deterministic. But we're moving to a position where some of these components are non-deterministic. They're stochastical, so there is quite an important change.

**Are there any other points that you wanted to leave us with which might encourage people to attend your your lecture on the 14th of January?**

Well, what I would say is that things are changing quite rapidly and we are getting impact and progress, certainly in my activities in the TM Forum, almost on a daily basis and there will be more to talk about on the day.

Across the TM Forum community, which shall I say, consists of operators in especially Australasia and Europe with Telefonica and Orange contributions quite substantial as is Vodafone's contribution, particularly in the AI area. But I think where most operators are, is a lot of them are in proof-of-concept stages, because they realise they need to understand what the capabilities of the technology are, how they would actually interact and manage with those kinds of tools in an operational environment.

There's certainly some evidence that scaling these things up is not just simply a question of pushing a button on a cloud computing server and say, make this run

500 times bigger, faster or, you know, 500 times the scale it was. For example, in knowledge bases knowledge graphs, you typically need about 30 billion elements in these graphs for a reasonable size operator area, which is really approaching the limits of what can be done with current technology, labelled property graphs. But it does mean that you need to be able to introduce federated models into telecoms, which we haven't seen in other industries as yet, so there's lots of good stuff coming out of practical demos which we can share.

**30 billion elements in these graphs? Could you just position that? What might these elements be?**

They would be things like fibre splitters, fibre sections, cable sections, cross connect cables, frame plug-in points, routers, DWDM type things because I'm thinking now of people who've got real physical networks as opposed to running things in computer servers. In the UK there might be 30 million home endpoints. So each of those might have 1000 other elements associated with it to get to 30 billion.

And why? Why do you need to do that? Because when you get failures in the networks, you need to be able to resolve down to where the root cause is because you only want to send one person out to the fault, not lots of people. And the economics and the Return-on-Investment opportunities around physical networks and the use of active networks to provide insights to where faults exist in passive networks is actually quite large. I mean huge. It's measured in billions of dollars per annum for quite average size telco.

So for operators it's a very interesting area, but I think we're at the experimental stage at the moment and there's a lot of hype out there, some of which is equating AI with witchcraft.

But there are some useful things that are going on, and we see lots of things around contact centres, for example, where it's very difficult to maintain and motivate contact centre staff and therefore it's difficult to maintain consistency. But the use of AI both in conjunction with human staff and independently as chatbots and to drive up the consistent experience sounds paradoxical because they're stochastic systems. But we can actually drive up the consistency of the way which an organisation

interacts with this customer, so it's all rather fascinating because we're at one of those Infexion points where I think a lot is going to change quite rapidly.

**That might be a good place for us to finish, Dave. Well, thank you very much.  
We are very much looking forward to your talk on 14<sup>th</sup> January.**