MOTT MACDONALD

Next-Generation Telecoms

The issues for utilities

December 2016

10 Fleet Place
London EC4M 7RB
United Kingdom

T +44 (0)20 7651 0300
F +44 (0)20 7248 2698
mottmac.com

Next-Generation Telecoms

The issues for utilities

December 2016

Information class: Standard

Contents

Executive Summary 1

1 Drivers for Telecom Technology Change 2

1.1 What's wrong with circuit-based networks? 2
1.2 What's wrong with packet-based networks? 2
1.3 What's wrong with commercial telecom services? 3
1.4 Will utilities be forced to migrate to next-generation telecoms? 3

2 Meeting Utility Requirements using Next-Generation Networks 4

2.1 The Teleprotection Challenge 4
2.2 Special Requirements for Power Utilities 6
2.3 Consolidating Operational and Business Telecom Networks 7

3 Conclusions 10

Executive Summary

For many years, power utilities have operated their own telecoms networks. These networks support a number of services that are critical to the operation of the power grid:

- Control centre monitoring of alarms and status indications from critical parts of the grid
- Automated protection systems to minimise the damage caused by faults on the grid
- Remote control of switchgear to enable rapid service restoration after a fault
- Access control for substations and other sensitive locations
- Operational voice services, including mobile communication with field staff

It is claimed that these critical services are now under threat as a result of fundamental changes that are occurring in telecoms technology. Telecoms networks are replacing legacy TDM technologies (such as PDH and SDH) with next-generation networking technologies based on IP, MPLS and Ethernet, and this has raised concerns that the very tight telecoms requirements for grid applications such as teleprotection can no longer be met.

Utilities are finding it increasingly difficult to buy TDM-based services (such as leased lines) from commercial telecoms providers, and so are seeking to migrate these services back to their in-house networks. However, the TDM equipment needed to expand their in-house networks is often no longer available because telecoms vendors are migrating their product lines to next-generation technologies. Even the spare parts and maintenance services needed to support existing networks are starting to be withdrawn.

Utilities are finding that they cannot hold back the tide of change

Power utilities operate critical national infrastructure. Mindful of their obligation to keep the lights on, these utilities have been reluctant to move away from existing TDM networks in favour of packet-based networks that seem to be less well aligned to their needs. However, they are finding that they cannot hold back the tide of change; the commercial telecoms industry has already embraced packet-based networks, and utility telecoms is a much smaller market than commercial telecoms. Whether they like it or not, utilities will soon have no option but to migrate to the new generation of networking technology.

Is this a disaster? Many in the power industry will claim that it is. However, the new networking technologies can deliver significant benefits, and the small number of utilities that have fully adopted the technology have been very positive about the experience. In this white paper, Mott MacDonald identifies the key issues and provides recommendations for how utilities should manage the migration from circuit-based to packet-based telecoms. Whilst this paper focusses primarily on the power industry, many of the comments are also applicable to other utilities with mission-critical telecoms requirements.

1 Drivers for Telecom Technology Change

1.1 What's wrong with circuit-based networks?

A telephone call that is dialled across a traditional circuit-based network results in a dedicated connection being established for the duration of the call. This circuit was originally set up as a physical connection using copper wires, but more recent circuit-based networks have carried telephone calls using dedicated streams of bits within a high-speed digital circuit. The stream of bits allocated to each telephone call is not available for sharing with other users of the network, so if the callers both fall silent then the bits are effectively wasted.

Circuit-based networks are relatively efficient for carrying voice services, but they are much less effective for carrying data. To illustrate this point, consider the demands placed on the network when you are surfing the World Wide Web. When you type the address of a web page into your browser, a significant amount of bandwidth is required to download the page quickly. However, you may then spend ten minutes reading that page – or even wander off and have lunch! As a result, data traffic tends to be much more "bursty" than voice. If a circuit was established to carry this data, it would need to have sufficient bandwidth to download a page quickly, but the bandwidth between page downloads would be wasted.

Until relatively recently, telephone networks were voice networks that could carry data, and users of dial-up modems could vouch for the fact that the data capabilities of these circuit-based networks were rather limited. Modern telecoms networks carry far more data than voice, so "next-generation networks" are designed to be data networks that carry voice – rather than the other way around. As a result, they have discarded the circuit-based technologies used by traditional telephone networks in favour of the packet-based technologies used on the internet. Within the telecommunications industry, Time Division Multiplexing (TDM) technologies such as PDH, SDH and SONET are being replaced by routers and switches based on IP/MPLS and Ethernet.

1.2 What's wrong with packet-based networks?

Whilst next-generation telecoms technologies have been adopted enthusiastically within the telecoms industry, utilities have been much more cautious. Some reasons for this caution are listed below:

- Service performance requirements. Packet-based networks are often unable to commit to the service levels achieved by existing TDM networks in areas such as delay and delay variation. This is a major issue for critical grid applications such as teleprotection.
- Legacy interfaces. The equipment in power networks tends to have a working life that is measured in decades, so it is not surprising that power grids around the world contain large amounts of elderly plant. Although this equipment is still operating perfectly satisfactorily, the monitoring and control ports were designed to suit an earlier generation of telecoms equipment; instead of Ethernet interfaces, this equipment typically uses legacy interfaces such as X.21, RS232 or V.35 which are not normally supported on modern networking equipment.

As a result, the telecoms networks operated by utilities are often based on TDM networking technologies. However, the TDM equipment used to construct existing networks may no longer be available as telecom equipment vendors migrate their product lines from circuits to packets. Furthermore, existing TDM equipment is becoming increasingly difficult to operate and maintain as spare parts become unavailable and vendor support for network management systems is withdrawn. As a result, the equipment needed to expand utility telecoms networks or "refresh" obsolete equipment is becoming increasingly difficult to obtain.

1.3 What's wrong with commercial telecom services?

Power companies have traditionally used commercial telecom services such as leased lines to provide connectivity in parts of their network where a self-build solution was commercially unattractive. However, utilities are becoming increasingly reluctant to use commercial telecom services for the following reasons:

- Service performance. As the telcos migrate their services from circuits to packets, leased lines are being replaced by packet-based services with less stringent performance commitments.
- Service interfaces. Telcos typically deliver packet-based services via Ethernet interfaces rather than the traditional service interfaces required by utilities.
- Availability targets. Telcos have always struggled to meet the very high service availability targets demanded by utilities.
- Mains Power Independence. Power utilities depend on telecommunications to restore power quickly and safely after a power outage, but many nodes in commercial telecoms networks are not equipped to continue operating during a prolonged power outage.
- Dependability. Power utility applications depend on the performance of a circuit (eg latency) remaining repeatable throughout the life of the circuit and not constantly changing as the network is reconfigured to meet changing telco requirements.
- Increasing reliance on telecoms. As electricity grids become "smarter", power utilities are increasingly regarding telecoms as part of their core business. They therefore see the provision of telecoms services as something that needs to be under their direct control.

As a result, many utilities have been forced to extend their telecoms networks and bring these services in-house.

1.4 Will utilities be forced to migrate to next-generation telecoms?

As we saw in the previous section, many power companies are migrating services that were previously provided by commercial telecoms companies back to their in-house networks. This migration, coupled with growing demand for electricity and the need for smarter grids, is placing increasing pressure on the in-house networks operated by power utilities. However, enhancing utility telecoms networks is becoming increasingly difficult as TDM technology disappears from the market.

Utilities may be reluctant to abandon their TDM technologies in favour of packet-based technologies that seem to be less well aligned to their requirements, but they are finding that they cannot hold back the tide of technological change. The following section will show how utilities can become beneficiaries rather than victims of next-generation telecoms.

2 Meeting Utility Requirements using Next-Generation Networks

This section reviews some of the key concerns that power utilities have in relation to next-generation telecoms networks, and considers ways in which these concerns can be addressed. This discussion is an essential pre-cursor to the development of a strategy for deploying next-generation telecoms technology.

2.1 The Teleprotection Challenge

Since Current Differential Teleprotection systems represent one of the most challenging applications that a utility telecoms network will be required to support, it is reasonable to start by considering the requirements of this application. Any networking technology that can support teleprotection properly is likely to be able to support other time-critical applications needed by power utilities.

Current Differential Teleprotection systems work on the principle that the electrical current entering a transmission line at one end should be the same as the current emerging at the other end. If there is a significant discrepancy between the two measurements, then current must be leaking to ground somewhere along the line. Since fault currents on a transmission network can be extremely high, the protection scheme needs to react very fast (typically within a few 50Hz AC cycles) in order to minimise any damage.

Fault currents can be detected by sampling the value of the current at both ends of the line several times per cycle. The samples are then transmitted over the telecoms network to the opposite end where the two sets of samples are compared. If it is found that the currents at the two ends of the line are not the same, then it is assumed that there is a fault on the line and a circuit breaker is opened. Since this comparison takes place simultaneously at both ends of the line, circuit breakers will be opened simultaneously at both ends and the line will be completely isolated.

The current carried by an electricity grid is constantly fluctuating in response to changing loads, so it is essential that the comparison uses a pair of samples from the equivalent time period. Since the samples from the remote end of the line will encounter delay in the telecoms network, the two sets of samples must be aligned to compensate for this, and this means that the teleprotection relays need to be able to measure the delay introduced by the telecoms network. They do this by sending a message which generates an immediate response from the far end of the line. They then measure the round-trip delay encountered by the message, and divide this value in half to obtain the one-way delay.

Obviously, if the two directions of transmission do not introduce the same amount of delay, then this calculation will yield the wrong answer and the system will fail to operate correctly. Path asymmetry in the telecoms network is therefore a major issue for this form of teleprotection. Furthermore, the transmission delays need to remain stable - if they change following a network reconfiguration then this will interfere with the correct operation of the protection system.

As shown in Table 1, Current Differential Teleprotection places some demanding requirements on telecoms networks. These requirements are difficult for most IP-based networks to meet, but they can be met relatively easily by traditional TDM-based networks:

Table 1: The Teleprotection Challenge

Requirement	IP-based Networks	TDM-based Networks
Delay (Latency) The end-end delay across a telecoms network linking the teleprotection relays generally needs to be less than 10mSecs.	The delays caused by collecting enough samples to fill an IP packet, coupled with the delays caused by packet queues within the network, can easily exceed 10mSecs.	Time division multiplexers do not put the data into packets, so there are no packetisation delays and no queuing delays, although a small amount of framing and buffering is required.
Delay Variation (Jitter) Latency variations in the network need to be managed within very tight limits (eg 500µSecs).	The statistical nature of packet queuing delays on an IP network means that delay variations are very likely to exceed 500µSecs.	Time division multiplexing does not depend upon statistics. There should be no variations in delay unless the network has synchronisation problems or is reconfigured.
Path Asymmetry Some protection schemes require the delay across a telecoms network to be the same in both directions.	IP networks may not constrain all packets to follow the same path, so path asymmetry can occur.	TDM multiplexers normally use bi-directional circuits that follow a fixed route across the network. There is therefore no asymmetry between the two directions of transmission.

Source: Mott MacDonald

In the United Kingdom, the Electricity Networks Association (a representative body for transmission and distribution network operators) has established a specification for the performance of communications circuits in terms of transmission delay, mains autonomy, repair time, separation and performance. It defines three categories for communications circuits for use in different voltage levels in the power system:

- Transmission grid (250kV and above)
- Distribution grid (66kV and 110kV)
- Distribution network (below 66kV)

TDM technologies such as PDH, SDH and SONET have been used in utility networks for many years, and are quite capable of meeting the stringent performance requirements of applications such as teleprotection. Since teleprotection systems are fundamentally important for maintaining the safe operation of electricity grids, it is not surprising that most power utilities have no particular wish to experiment with IP-based alternatives.

However, the real-time performance of IP networks can be dramatically improved by the introduction of MPLS. MPLS is a circuit-based technology that can be used in a packet-based network, and it re-introduces some of the more valuable features of TDM networks that were lost in the original transition to IP. For example:

- MPLS can be used to constrain traffic to follow a particular path across the network.
- MPLS can be used to provide quality of service guarantees to sensitive traffic such as teleprotection.
- MPLS can support high-speed circuit protection schemes that are similar to SDH path protection.
- Some IP/MPLS networks can distribute accurate synchronisation to the network nodes (as was the case in SDH networks). This helps to minimise delay variation.

Some modern IP/MPLS networks can support demanding real-time applications

As a result of these developments, some modern IP/MPLS networks can support demanding real-time applications such as teleprotection. Table 2 shows how the teleprotection challenge can be met by IP/MPLS networks that have been designed to meet utility requirements:

Table 2: The Teleprotection Challenge

Requirement	Standard IP Networks without MPLS	Specialist IP/MPLS Networks
Delay (Latency) The end-end delay across a telecoms network linking the teleprotection relays generally needs to be less than 10mSecs.	The delays caused by collecting enough samples to fill an IP packet, coupled with the delays caused by packet queues within the network, can easily exceed this value.	Packetisation delays can be reduced by reducing the size of the IP packets. Packet queue delays can be minimised by assigning a higher priority to teleprotection packets than to other forms of traffic. The use of MPLS switching rather than IP routing at intermediate nodes in the network helps to minimise transit delays.
Delay Variation (Jitter) Latency variations in the network need to be managed within very tight limits (eg 500µSecs).	The statistical nature of packet queuing delays on an IP network means that delay variations are common.	Packet delay variation is minimised by assigning a higher priority to teleprotection packets than to other forms of traffic. If the IP/MPLS network is emulating a TDM circuit, then a jitter buffer driven by a highly-accurate network clock can be used at the receiver to eliminate any remaining delay variation.
Path Asymmetry Some protection schemes require the delay across a telecoms network to be the same in both directions.	IP networks may not constrain all packets to follow the same path, so path asymmetry can occur.	IP/MPLS networks can use label-switched paths to ensure that all packets associated with a particular teleprotection scheme follow the same path.

Source: Mott MacDonald

2.2 Special Requirements for Power Utilities

Although the real-time requirements of teleprotection represent a particular concern when power utilities are migrating from a TDM-based telecoms network to one based on IP/MPLS, there are a number of other challenges that are likely to arise. These challenges are set out in Table 3 below:

Table 3: Other network migration challenges

Issue	Requirement	Possible Solution
Multi-drop SCADA circuits	Some SCADA systems are designed to operate over multi-drop leased lines.	Some IP/MPLS equipment can emulate a multi-drop leased line.
Equipment designed for substation environments	Equipment deployed in electricity substations needs to be suitable for that environment. For example, it needs to be capable of operating from DC power supplies and to be screened against strong electromagnetic fields. It also needs to be physically rugged and must operate over an extended temperature and humidity range.	Although standard telecoms equipment is generally not suitable for use in an electricity substation, IP/MPLS equipment is available that has been designed specifically for such environments.
Legacy interfaces	Power grids around the world contain large amounts of elderly plant. Although this equipment is still operating perfectly satisfactorily, the monitoring and control ports typically use legacy interfaces such as X.21, RS232 or V.35, and these interfaces are not normally supported on modern routers and switches.	Next-generation telecoms equipment that has been specifically designed for utility applications is likely to support these legacy interfaces.

Source: Mott MacDonald

2.3 Consolidating Operational and Business Telecom Networks

In common with most large businesses, power utilities have a significant requirement for telecoms services to support day-to-day business activities such as office telephony and broadband Internet access. Within the utility world, services to support business applications are often referred to as Information Technology (IT), while services to support the operation of a power grid are referred to as Operational Telecoms (OT).

If the utility already has their own OT network to support the operation of the power grid, then it is natural that they will seek to use the same network to support their IT requirements as well. They may even wish to set up a utility telco (utelco) to provide commercial telecoms services to external customers. However, there are likely to be people within the organisation who will argue that mission-critical power grid applications should not be forced to share the same telecoms network as business applications. Table 4 lists some key objections that are likely to arise if OT and IT - and possibly external customers - need to share the same physical network. The table also suggests some possible solutions:

Table 4: OT/IT Network Sharing Issues

Issue	Objection	Possible Solution
Bandwidth contention	IT network applications are often bandwidth-intensive, whilst most OT network applications require surprisingly little bandwidth. There are concerns within the OT community that their critical applications could start to run slowly as a result of very heavy bandwidth demand from the IT network.	If logical separation is maintained between the IT and OT networks, then it is possible to ensure that the performance of the OT network is unaffected by traffic levels on the IT network.
Logical security	IT data applications often involve access to the Internet, and so represent a potential route via which various forms of cyber attack can enter the network. It will be claimed that the OT network needs to be “air-gapped” from the internet to protect the power grid from cyber threats.	In reality, OT networks are very rarely “air-gapped” from the internet because VPN technology is used by vendors and remote operations staff to diagnose faults and fix software problems. However, there are good security arguments for maintaining complete logical separation between the IT and OT networks.
Different availability requirements	The power grid is required to operate continuously, and any power outages can have very serious consequences. In contrast to this, most office applications are normally only required during office hours, and a failure in the middle of the night might be of little concern. This leads to different availability targets for OT and IT services, so different levels of network resilience and different fault repair times are needed.	The need to support services with different availability requirements is something that commercial telecoms networks have to handle all the time. Whilst having to support different availability targets does make network operation more complicated, it certainly does not make it impossible, and it is likely to be a more cost-effective solution than providing IT and OT with physically-separate networks.
Cultural differences	IT network operators tend to apply software patches regularly to eliminate potential vulnerabilities in their networks. OT network operators, on the other hand, generally want to keep their networks as stable as possible, and so regard software patches to fix problems that they have not encountered as introducing unnecessary risk. Without wishing to emphasise cultural stereotypes, there is often a difference in approach between the two teams, and this can lead to reluctance to share network infrastructure.	Maintaining logical separation between the IT and OT networks does not prevent the sharing of physical assets such as fibre and backbone routers, but it does enable the two networks to be operated completely separately. Each network can be protected from any technical vulnerabilities or operational weaknesses that may exist in the other network.

Source: Mott MacDonald

One way of achieving a high level of physical infrastructure sharing between the OT and IT networks, while maintaining complete logical separation, is illustrated in Figure 1 below:

Figure 1: Logical Separation between OT and IT Networks

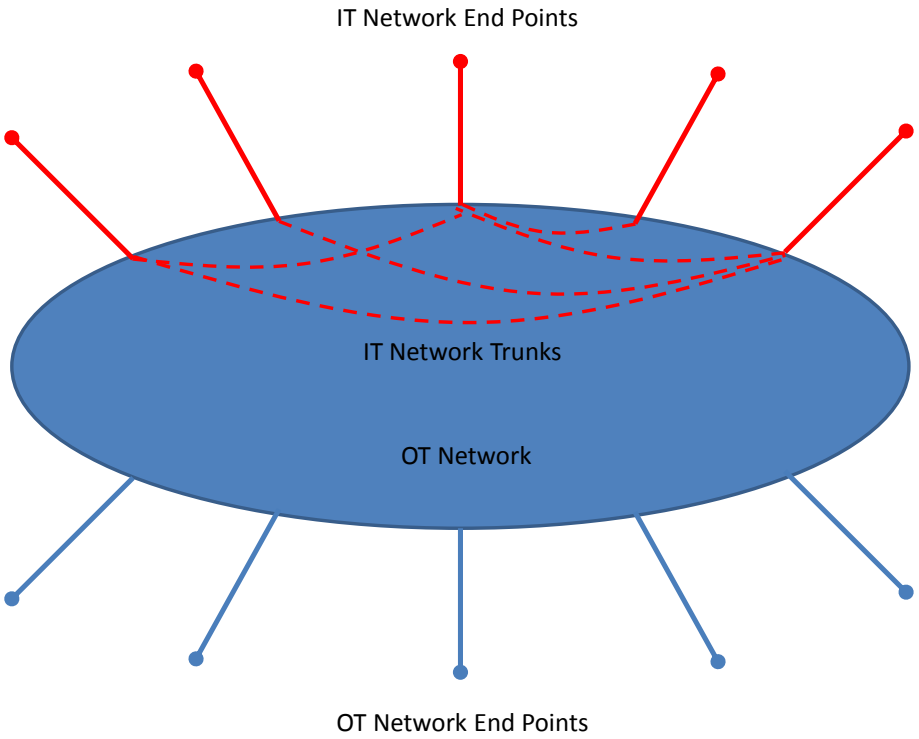

Source: Mott MacDonald

The IT network is constructed using network trunks supplied by the OT network. There is nothing to stop the IT network from also using trunks provided by commercial telecom operators if the OT network cannot provide capacity on a particular route, or if additional resilience is required. (As discussed in Section 1.3, commercial telecoms services can cause problems in OT networks, but they are likely to be acceptable for IT networks.)

The core network is operated by the OT team because the OT network needs to meet far more stringent service performance targets – including some that are safety-critical. Although the OT network is likely to require less capacity than the IT network, it will almost certainly provide coverage to a much larger number of sites, and that coverage will need to expand further as the grid gets “smarter”. The boundaries between the two operational teams should be entirely clear, and the OT and IT network management systems should each present the operator with a clear view of the logical network that is under their control.

One apparent disadvantage of this arrangement is that sites requiring both OT and IT services will need two separate routers to deliver them. However, this is much less of a disadvantage than it might at first appear:

- On sites where both OT and IT services are required, the two types of services are normally required in different parts of the site and so it is convenient to deliver them using separate routers.

- OT services typically have to be delivered in secure environments with highly restricted access, while IT services are normally required in offices. Having two separate routers to deliver the services means that IT technicians do not need to be trained to work in potentially-dangerous locations such as electricity substations.
- OT network services are typically delivered via legacy network interfaces, while IT network services are typically delivered over Ethernet.
- OT network services often have to be delivered in harsh environments where temperature, humidity and electromagnetic fields can all be a problem. Specialist networking equipment is required to deliver OT services, whilst standard equipment can be used to deliver IT services.
- Some utilities choose to outsource their IT network while retaining their OT network in-house. This reflects the view that the IT network requirements of a utility are pretty similar to those of any other large business, and so can safely be outsourced; OT requirements, on the other hand, are much more specialised and much closer to the utility's core business. Delivering IT and OT services on separate routers makes it much simpler to outsource the IT network if required.

The architecture shown in Figure 1 illustrates a way in which a utility can achieve economies of scale in the core network whilst maintaining complete logical and operational separation between IT and OT. It also indicates a way in which TDM and IP/MPLS networks can co-exist, because the IT network is likely to make the transition from circuits to packets considerably earlier than the OT network.

3 Conclusions

This paper has demonstrated that it is possible for next-generation telecom technologies to meet the requirements of mission-critical utility applications. Furthermore, it is possible to maintain complete logical and operational separation between the OT and IT networks without sacrificing economies of scale. On the other hand, packet-based networks are different from traditional circuit-based networks so utility staff will face a steep learning curve, and migrating network applications to next-generation telecoms technology will require very careful planning.

So how should utilities manage the issues raised in this white paper? Some key messages are listed below:

- Embrace the new technology rather than fighting against it. There is no point in ignoring these developments or hoping they will go away - that simply stores up bigger problems for the future. In our experience, utilities that have rolled-out IP/MPLS networks have become enthusiastic proponents of the new technology.
- Develop a network technology strategy. This document should provide a roadmap to show how the existing telecoms network will evolve to IP/MPLS.
- Try to avoid introducing interim technologies into the network. They add costs, make the network more difficult to manage and slow down progress towards the target architecture.
- Develop a migration plan. Migrating network applications from a TDM network to IP/MPLS cannot happen overnight. Develop a phased roll-out plan to ensure that the network delivers the full required functionality throughout the migration process.

Utilities that have rolled-out IP/MPLS networks have become enthusiastic proponents of the new technology

Source: Wikimedia Commons

Dr Andrew Wheen is a Project Director within Mott MacDonald's Digital Infrastructure practice. He is part of an integrated team of strategic advisers, technical experts and project managers that are delivering major telecommunications and IT projects around the world.

For further information, contact:

E-mail: Andrew.Wheen@mottmac.com

Office: +44 (0)20 7651 0656

Mott MacDonald Limited, 10 Fleet Place, London, EC4M 7RB, United Kingdom

www.mottmac.com

Until recently, Andrew Thomas was a Divisional Director within Mott MacDonald's Technology & Communications Consultancy. Andrew has over 25 years' experience of telecommunications in the electricity supply industry. He started out in a distribution company and his considerable experience bridges the power and telecommunications sectors. Today, Andrew is active in the smart grid field working with utilities, presenting at conferences and is secretary of a CIGRE working group.