Utilities Technology Council of Canada™
Conseil canadien des technologies pour les services publics™

Quality of Service in power utility telecommunications networks

White Paper

September 2017

Many power utilities are facing challenges as they are transitioning critical services from legacy TDM to IP/Ethernet core networks, thereby combining IT and OT traffic onto a single converged network. While QoS is a crucial requirement that must be considered from the beginning of such a migration, it becomes even more critical when traffic has to go through lower bandwidth links.

Contents

Quality of Service in power utility telecommunications networks

1 Overview

UTC Canada's objective is to assist its members with their ICT needs. This discussion focuses primarily on the use of traffic prioritization in a utility telecommunications network.

The convergence of OT traffic with IT traffic onto an MPLS network brings many challenges. Many utilities are facing those challenges as they are transitioning critical services from legacy TDM to IP/Ethernet core networks. While QoS is a crucial requirement that must be considered from the beginning of such a migration, it becomes even more critical when traffic has to go through lower bandwidth links, such as a fixed service point-to-point microwave links.

One crucial step is the inventory and differentiation of the services to be carried, necessitating a thorough knowledge of the utility's various internal customers and their applications. Although performance requirements can be numerous depending upon the type of service, they can often be reduced to availability, packet loss, bandwidth, latency, jitter and (sometimes) delay asymmetry.

For some utilities, microwave networks play a prominent role in their wide area network architecture. In order to make the best use of those assets, a good understanding of the active role that microwave radios can play in an IP/MPLS network is required. The

OT is defined as Operation technology. It consists of all the applications required for the control and operation of a power grid.

Typical OT applications include teleprotection, RAS and SCADA.

migration of a microwave network from legacy TDM to all-packet technologies will be very involved.

Finally, Quality of Service (QoS) is a tool that was designed to help dealing with traffic congestion issues. Hence, the use of QoS is particularly necessary in networks that contain severe bandwidth limitations (e.g. narrow-band radio networks or legacy transport technologies), as is often the case within utility private networks.

2 Key considerations

This discussion focuses on the application of QoS in a private utility network. QoS considerations of utility applications in an ISP network are not discussed.

2.1 Service requirements

The network architect should always maintain a close relationship with his customers and ensure that the customers' needs and requirements are well understood. This is of course the key input for establishing the levels of service that will lay the foundation of a good QoS architecture.

Technical requirements differ between various services. The most demanding services are generally those that are critical to the power grid operation (OT) like line relays, Remedial Action Schemes (RAS) or SCADA. These applications would be considered high availability and they are often carried on MPLS/L2-VPN (Ethernet or circuit emulation).

On the IT side, the range of corporate applications is larger, and it is less straight-forward to capture the key requirements for each and every application. Among them, VoIP probably has the most stringent performance requirements. Other applications could need low latency or high bandwidth, for example video conferencing, video broadcasting, database access, file transfer or email.

Each service has its own bandwidth, latency, packet-loss, jitter and availability requirements. The following table is an example of a number of different services a utility may require. The performance metrics listed in the table are typical, but these must be determined by each utility on a case-by-case basis.

<u>Typical stringent electrical services performance requirements – RAS and relays</u>

- Latency: <20ms
- Asymmetry/Jitter: <1ms
- Packet loss: 0.01%
- Availability: 99.9%

<u>Typical SCADA services performance requirements</u>

- Latency: 500ms-2s
- Jitter: 500ms
- Availability: 99.9%

<u>Typical IT services performance requirements – VoIP</u>

- Latency: 150 ms
- Jitter: 30 ms
- Packet loss: 1%

Table 1 - Example of services requirements

Service Name	Bandwidth	Latency	Jitter	Packet Loss	Availability	Notes
Corporate Voice	Medium	<150-200ms	<10 to 30ms	<1%	>99%	Depends on codec and jitter buffer
Corporate Data	High	<2s	<500ms	<5%	>95%	
Security Video	Very High	<5s	<1s	<5%	>95%	
SCADA Collector	Medium	<500ms to 2s	<500ms	<0.5%	>99.998%	SCADA Concentrators to control centers
SCADA RTU	Low	<500ms to 2s	<500ms	<0.5%	>99.5%	Individual SCADA RTU to control centers
RAS All	Low	<17 to 25ms	<1ms	<0.01%	>99.95 to 98.1%	Deployed as separate PY and SY circuits
Differential protection	Low	<10ms	Asymmetry delay** <1ms	<0.01%	>99.9%	Deployed as separate PY and SY circuits
Permissive Trip (Line Relay) 735/500kV	Low	<10 to 12ms	<1ms	<0.01%	>99.95 to 99.9%	Deployed as separate PY and SY circuits
Permissive Trip (Line Relay) 315/230kV	Low	<10 to 17ms	<1ms	<0.01%	>99.9 to 99.5%	Deployed as separate PY and SY circuits when justified
Permissive Trip (Line Relay) <230kV	Low	<34ms	<1ms	<0.01%	>99.9 to 95%	
Transfer Trip 735/500kV	Low	<10 to 17ms	<1ms	<0.01%	>99.95 to 99.9%	Deployed as separate PY and SY circuits
Transfer Trip 315/230kV	Low	<10 to 25ms	<1ms	<0.01%	>99.9 to 99.5%	Deployed as separate PY and SY circuits when justified
Transfer Trip <230kV	Low	<50ms	<1ms	<0.01%	>99.9 to 95%	

** Asymmetry delay must not be confused with jitter.

2.2 Typical QoS features

The complexity of implementing a uniform QoS architecture presents a real challenge. Apart from the fact that some pieces of network equipment may be operating at different network layers from others, network equipment vendors each have various proprietary features or limiting constraints. Of course each situation is unique, but there are a few basic principles to consider.

2.2.1 Classification

The first step in the implementation of a QoS architecture is to define a limited set of forwarding classes to which the various services (or types of traffic) will be assigned.

To facilitate management and operation, and to ease the correlation with hardware queues, it is recommended to limit the set of classes to 8 or fewer. The lower the number of classes, the easier it is to manage, troubleshoot and adapt. A trade-off between priority granularity and architecture complexity must be determined. In the industry, standard sets of classes are typically made up of 3, 5 or 8 classes.

2.2.2 Marking

Marking is a fundamental part of a QoS architecture. Depending on the forwarding classes defined for each service, the packets will be assigned a specific marking.

Complexity arises from the different types of available marking protocols. Generally, in an IP/MPLS network, 3 types of marking can be used depending on the layer of the packet: CoS, DSCP and EXP. It is therefore necessary to clearly define the hierarchy of the classes of service for each level and to keep track of the correlation between them.

Typically, it is recommended that packets should be marked as close as possible to the customer equipment at network ingress. Based on that statement, CoS would be the marking mechanism of choice. However, in a private IP/MPLS network, class of service consistency ought to remain a major concern and L3 marking would be more suitable. Since the L3 domain is generally more evenly distributed through the network, and that DSCP can more easily

<ins>CoS (dot1p)</ins>

L2, 3 bits, 8 levels

<ins>DSCP</ins>

L3, 6 bits, 64 levels

<ins>EXP</ins>

MPLS, 3 bits, 8 levels

remain unaltered across the layers, using the latter as the reference marking is a good practice. CoS marking could still be used to mark voice traffic at switch level in order to differentiate it with other IT lower priority traffic.

Note that the number of available DSCP priority levels will generally be much higher than the number of forwarding classes. It is possible to increase the priority granularity by assigning different DSCP markings to individual services within a class (e.g. AFxx to affect rejection probability).

2.2.3 Queuing

In terms of number of queues, most vendors offer a basic set of 8 queues, which would be considered sufficient in most cases. Typically, each forwarding class will be allocated a specific queue.

For each queue, a CIR (Committed Information Rate) and a PIR (Peak Information Rate) must be allocated. In simple terms, CIR is typically defined as the throughput guaranteed on the link and PIR as the maximum total bandwidth that can be used by borrowing unused bandwidth from another queue.

The packets will be profiled as follows depending on their traffic rate:

- in-profile (or conforming), when within the CIR
- out-of-profile (or exceeding) when above CIR but within PIR
- violating when above PIR

Profiling allows the scheduler to double the number priorities with regards to the number of classes. Out-of-profile packets will not be dropped, but will not benefit from the same level of priority as the in-profile packets.

It should be mentioned that in a private network, one should focus on maximizing the available bandwidth so that in the absence of congestion, applications could take advantage of the full bandwidth. This idea will be further developed in the Policing section.

Real time services and strict priority scheduling are often associated to expedited forwarding classes.

2.2.4 Queue management - Scheduling

Queue management (scheduling) is typically based on the three following features: First-In-First-Out (FIFO), Strict Priority and Weighted Fair Queuing (WFQ). Strict Priority is generally used for real time services that are delay sensitive, for example: voice or OT applications. With this mechanism, the buffers with the highest priority are always scheduled first and the latency and jitter are minimal.

Weighted queuing will help prevent buffer starvation and maximize bandwidth usage in lower priority queues that are not sensitive to latency.

2.2.5 Traffic conditioning - Policing

Policing is a vital function on bandwidth-limited WAN links in order to guarantee the priority of the critical applications in their queue. It is performed at ingress and consists essentially of a rate limiting tool that ensures that the traffic allowed into in each forwarding class does not exceed its allocated bandwidth.

Policers of different types are available: Single rate/two-color, Single rate/three-color and dual rate/three-color. Depending on the configuration of each specific class and queue, any mechanism may be suitable.

As mentioned before, in a private network, there are no real benefits for limiting the total usable bandwidth of the queues. So to help burst management and maximize bandwidth usage when available, a dual rate 3-color (two token buckets) policer could be chosen. CIR would be set to the class allocated bandwidth and PIR to 100% link capacity.

A Real Time class, however, should be treated differently. In order to avoid queues starvation problems (Real time classes are managed with Strict priority, as opposed to WFQ for other classes), this type of traffic should not be allowed to burst and monopolise capacity. Single rate 2-color policers are better adapted to this situation. This type of traffic is generally less bursty and CIR should be limited to the class allocated bandwidth. In other words, PIR should be limited to CIR for the Real Time traffic class. It should be noted that the lack of flexibility in this class necessitates a meticulous monitoring of the allocated bandwidth in order to adjust the parameters when necessary to avoid dropping critical traffic that bursts above the CIR.

<u>Real time classes</u>

Strict priority, CIR = PIR = Allocated bandwidth, single rate 2-color policer

<u>Other classes</u>

WFQ, CIR = Allocated bandwidth, PIR = total link bandwidth, dual rate 3-color policer

2.2.6 Traffic conditioning - Shaping

Egress traffic shaping is necessary when dealing with bandwidth restrictions. Shaping is usually realised on a per port basis. However, on newer equipment, shaping on a per VLAN basis is now possible. This feature is part of a more general concept known as H-QoS and it was designed to allow a per-customer shaping on trunks. In a utility private network, it is suitable for the transport of electrical services on restricted bandwidth links like microwave or SONET.

2.2.7 Congestion avoidance

Random Early Detection (RED) can help smooth the throughput and optimize bandwidth utilization by preventing the saw tooth throughput pattern that TCP connections can cause. DSCP (through AFxx differentiation) allows for a granular modulation of the rejection probability of the packets within the classes. However, RED comes with the drawback of early (prior to buffer filling) detection and rejection of a certain number of packets, which makes it unfit for critical Real Time applications.

2.2.8 Multicast

Some vendors allow special classification and queueing of multicast and broadcast traffic. Most utility applications do not require special multicast or broadcast capabilities. However, it may be helpful to provide a small rate-limited queue for broadcast traffic such as ARPs in critical multipoint LAN services. This will ensure that ARPs on critical services are prioritized correctly at service ingress, as well as limited to avoid broadcast storms in the case of a misconfiguration resulting in a loop.

2.3 Bandwidth allocation

Careful network planning is required to ensure that the bandwidth allocation on each link is suitable for each class. The Committed Information Rate (CIR) for each class should take into account expected utilization during both normal operations and in a potential failure scenario. Critical protection, SCADA, and operational voice traffic should always be guaranteed enough bandwidth to meet minimum functional requirements in the event of a network outage where this traffic is rerouted from its normal path through the network.

<u>SGT traffic considerations</u>

The control plane traffic (self-generated traffic, or SGT) can be a bit of a problem in a routed network since it is by default prioritized over any other application. It competes with sensitive applications (e.g. differential protection) for the highest forwarding priority and it can cause unacceptable latency and jitter, which in turn can make the application inoperable. To correct that issue, some vendors allow the user to customize the priority of the SGT. Routing protocol messages should have a higher classification, but nodal network management (FTP, SNMP, SSH etc.) should have a lower classification.

One approach to traffic engineering would be to commit a small portion of the link bandwidth to Best Effort traffic, while reserving the rest of the bandwidth for higher-priority traffic. The sum of the CIR for each traffic class should match the available bandwidth of the link.

Under normal operations, the link utilization should be much lower than 100%, and BE traffic may expand to consume un-used bandwidth committed to other classes. In a failure scenario, where other high-priority traffic is routed onto the link, the reserved bandwidth would then be used by higher priority applications. BE traffic would then be throttled to ensure critical applications continue to operate.

While it is possible to over-subscribe CIR in some equipment (i.e. sum of CIR for all the traffic classes exceeds 100% of the bandwidth available on the link), this is not recommended, as it is difficult to predict how the bandwidth would be used in a failure scenario when many high-priority traffic flows are competing for limited bandwidth resources.

A set of default CIR allocations may be generated for classes of links (e.g. nxT1, DS3, Gigabit Ethernet, etc.). These allocations may be changed over time on a per-link basis to reflect the actual requirements of each link during normal operations and failure scenarios.

2.4 Packet radio considerations

The first fundamental question to answer when considering the migration of a microwave radio network to all-packet technology is the OSI layer of operation of the radio. Choices between L1, L2, L3 and MPLS will have an impact on network performance and the OPEX. Once again, a trade-off must be found. Operation at L1 will be quite simple to configure and operate (reduced OPEX), but the limitations on flexibility and future-proofing capability will be greatly affected.

Increasing bandwidth of IP traffic on a microwave link is expected to cause a reduction of receive threshold levels (IP Error Rate substituted for Bit Error Rate, ITU-T Y.1541), leading to a decrease in system gain margins. To compensate for that loss, more sophisticated mechanisms like adaptive modulation and channel bonding might be required. These features must be considered in the effort to minimize CAPEX and reuse existing sites and infrastructure. However, these mechanisms have a reliance on the network's QoS scheme, and it may prove complicated to integrate the QoS scheme with advanced radio features.

Considering these limitations, L1 radios might prove to be sub-optimal. The features described above are unlikely to be available, and they might significantly impair the latency. Indeed, in a multi-link configuration (cascade of radio repeater sites), the inability of L1 radios to deal with VLANs might force all the traffic to be dropped at every site and be processed by the site router. This significantly increases the total number of nodes to transit and increases the latency significantly. By operating the radios at L2, those drawbacks can be overcome.

Regarding L3 and MPLS, the supported features in the radios are very limited at the moment. Microwave radios cannot be considered as a substitute for specifically designed routers in a near future.

Since the microwave network is part of the IP network backhaul, the radio equipment does not interface with the customer applications directly. Thus, they do not change or impose any traffic marking. Radios should be configured to take action on pre-marked packets. Scheduling and control in the radios should be carried out according to the priority tag available. The type of priority tag to be used between the router and the radio should be chosen between CoS priority bits, DSCP bits or EXP bits, depending on those available. The number of queues and the bandwidth allocation should be coherent as well.

One major impediment of using bandwidth varying technologies such as adaptive modulation in microwave radio links is the inability of the IP/MPLS network to be aware of the changes and adapt accordingly. A variable link throughput might have an impact on QoS health as well as the efficiency of routing protocols.

A possible solution to explore is the use of a hierarchical shaper with the outer shaper having a CIR equivalent to the minimum or worst-case bandwidth achievable on that link. Within that parent shaper, a child shaper could provide a CIR guarantee to the most critical classes at least.

Another solution to consider is the Bandwidth Notification Message protocol (ITU-T G.8013/Y.1731). The practical network applications for this feature are: routing protocols metric adjustments, network QoS congestion management adjustments (traffic shaping), load balancing adjustments and MPLS fast reroute (FRR) operation. Further study is required, as few network equipment vendors support this protocol.

3 QoS architecture examples

The recipe to a well-designed QoS architecture is unique for every situation. It is far from one size fits all. Guidelines can provide a good start, but lab testing and ongoing monitoring and analysis are necessary to fix potential issues. A utility's QoS architecture will have to be continuously adapted as technology and customer needs evolve.

To help understand how the recommendations presented in the previous section can be applied practically, this section will present two example QoS architectures, namely a 5-class model and an 8-class/16-priority model. The steps of service classification and configuration design will be treated independently.

3.1 Services classification and marking

In order to assign a forwarding class to each service, it is necessary to first bundle them into larger groups of similar network behavior. However, routers typically have a limited number of classification and queuing resources, so assigning each service its own dedicated queue is not feasible. Main groups (classes of service) could be defined as follows, each with an associated queue and specific marking.

3.1.1 Example 1: 5-class model (classification and marking)

Shown here is an example of services grouping and marking of a 5-class model.

Table 2 – 5-class model classification and marking

	Group / Class name	Service / Type of traffic	Packet loss Tolerance	Latency Tolerance	Jitter Tolerance	DSCP	L2-CoS IP-Prec	MPLS-EXP
5	**Real time**	OT – Relays and RAS	Very low	Very low	Very low	EF	5	5
		OT – SCADA Corporate Voice	Very low	Very low	Very low	CS5		
		Signaling	Very low	Very low	Very low	CS4	4	
4	**Multimedia**	Control plane traffic (SGT)	Very low	Very low	Low	CS6	6	4
		Security video	Very low	Very low	Low	CS3	3	
		Corporate interactive	Very low	Very low	Low	AF32		
3	**Sensitive**	Device management	Low	Medium-Low	Yes	CS2	2	3
		Sensitive apps (Database access, trading services, instant messaging, privileged applications)	Low	Medium-Low	Yes	AF22		
2	**Volume**	Heavy apps (File transfers, database synchronisation, emails)	Low	Medium-High	Yes	AF12	1	2
1	**Best effort**	Default	n/a	n/a	n/a	BE	0	1

3.1.2 Example 2: 8-class model (classification and marking)

Shown here is an example of service grouping and marking of an 8-class, 16-priority model.

Table 3 – 8-class model classification and marking

Traffic Class	Profile	Usage	Packet Loss Tolerance	Latency Tolerance	Jitter Tolerance	EXP	Dot1p	DSCP
7	In Profile	Network Control: Routing Protocols, ARP, STP	Very low	Very low	Very low	7	7	56 (nc2)
6		Service Traffic: Protection and RAS	Very low	Very low	Very low	6	6	48 (nc1)
5		Service Traffic: Critical Voice Service Traffic : High-impact SCADA	Very low Very low	Very low Low	Very low Medium	5	5	46 (ef)
4		Service Traffic: EMS Database Synchronization	Very low	Medium-Low	Medium	4	4	34 (af41)
3		Service Traffic: Telecom Device Management Service Traffic : Low-impact SCADA	Low Low	Medium Medium	Medium Medium	3	3	18 (af21)
2		Service Traffic: Corporate Voice	Very low	Low	Very low	2	2	10 (af11)
1		Service Traffic: Corporate Interactive	Medium	Medium	Medium	1	1	8 (cs1)
0		Service Traffic: Security Video	N/A	N/A	N/A	0	0	2 (cp2)
7	Out of Profile	Future Use				N/A		52 (cp52)
6		Future Use						44 (cp44)
5		Network Control: SNMP, SNMP-notification	Low	Medium	High			32 (cs4)
4		Future Use						36 (af42)
3		Network Control: SSH	Low	Medium	Medium			20 (af22)
2		Network Control: DHCP, DNS, ICMP, VRRP	Medium-Low	Medium	Medium			12 (af12)
1		Future Use						6 (cp6)
0		Future Use						0 (be)

3.2 Configuration design

The QoS architecture design is pursued with the same examples. The different QoS features are associated to the classes defined in the previous section.

3.2.1 Example 1: 5-class model (configuration)

In this example, all the traffic is considered in-profile at ingress and egress. It can be argued that specific actions on out-of-profile traffic are of no real benefit in private corporate network. In the absence of congestion, the lower priority classes in this example can benefit of all the available bandwidth. The action to take on the traffic that exceeds CIR in each class is left to the treatment of the equipment scheduler.

Table 4 – 5-class model configuration

	Class name	Marking	Queue	Profile	CIR %	PIR %	Scheduling	RED	Policing (ingress)
5	**Real time**	EXP 5 EF, CS5, CS4	Q5	In	35	CIR	Expedite / Strict priority	No	Single Rate 2-color
4	**Multimedia**	EXP 4 CS6, CS3, AF3x	Q4	In	15	100	WFQ	No	Dual rate 3-color
3	**Sensitive**	EXP 3 CS2, AF2x	Q3	In	20	100	WFQ	Yes	Dual rate 3-color
2	**Volume**	EXP 2 AF1x	Q2	In	5	100	WFQ	Yes	Dual rate 3-color
1	**Best effort**	EXP 1 BE	Q1	In	25	100	WFQ	Yes	Dual rate 3-color

3.2.2 Example 2: 8-class model (configuration)

The following is an example of an 8-class 16-priority model, where traffic is prioritized by whether it is in-profile or out-of-profile. In this example, it is assumed that all MPLS-encapsulated service traffic is marked as in-profile, while DSCP is used to mark network control traffic as always being out-of-profile. On network egress, the scheduler allocates traffic to each queue in descending order of priority Q8-Q1, first for queues that are within CIR, then in a weighted manner for queues with traffic above CIR or marked as out-of-profile. Within each queue, traffic marked as in-profile may exceed the CIR on egress, so the scheduler may treat excess traffic as out-of-profile, despite its marking. On service ingress, all traffic is marked as in-profile, but if it exceeds the service ingress CIR, it is treated as out-of-profile on the network egress scheduler on the first hop, regardless of whether the traffic exceeds the CIR for that class on that network interface. On subsequent hops, all service traffic is treated as described above, since the dot1p and EXP markings do not have enough granularity to retain knowledge of whether particular packets were in-profile or out-of-profile at service ingress. RED is not applied at this time.

Table 5 – 8-class model configuration

	Class name	Profile	DSCP Marking	Dot1p Marking	Queue	CIR %	PIR %	RED	Policing (ingress)
	7	In Out	56 (nc2) 52 (cp52)	7 -	Q8	3	3	No	Single Rate 2-color
	6	In Out	48 (nc1) 44 (cp44)	6 -	Q7	10	100	No	Single Rate 2-color
	5	In Out	46 (ef) 32 (cs4)	5 -	Q6	21	100	No	Single Rate 2-color
	4	In Out	34 (af41) 36 (af42)	4 -	Q5	20	100	No	Single Rate 2-color
	3	In Out	18 (af21) 20 (af22)	3 -	Q4	16	100	No	Dual rate 3-color
	2	In Out	10 (af11) 12 (af12)	2 -	Q3	15	100	No	Single Rate 2-color
	1	In Out	8 (cs1) 6 (cp6)	1 -	Q2	10	100	No	Dual rate 3-color
	0	In Out	2 (cp2) 0 (be)	0 -	Q1	4	100	No	Dual rate 3-color

4 What's next

There is a lot of uncertainty currently regarding all-packet OT applications and the coming of IEC 61850. It would of great interest to see how the present recommendations hold within the network architectures that will be developed to fulfill those emergent needs.

Following that trend, microwave networks are also migrating to all-packet technologies and QoS will certainly be of capital importance in the design of those future super robust networks, especially if infrastructure investments are expected to be kept as low as possible. Lessons learned and more detailed recommendations for cross-domain IP/MPLS-to-packet-radio QoS schemes are anticipated.

All UTC members are invited to share any test results, proofs of concept, specific case studies or even QoS architectures relating to the above-mentioned topics for publication to the UTC community. Recommendations for other topics for future white papers are also welcome.

Please contact Stephen LaRoy (stephen.laroy@bchydro.com) to submit ideas and content for publication.

About the authors

Principal authors

Chrystian Paquet, B. Sc., P. Eng [paquet.chrystian@hydro.qc.ca] – Chrystian is currently telecommunications engineer for Hydro-Quebec. Chrystian has the role of microwave network designer and architect as well as technical lead for the microwave radios technology. He is currently responsible for developing the microwave network evolution architecture and leading the migration of this network to all-packet microwave technology. Prior to 2010 Chrystian was employed by MDA Corporation (Montreal, Canada) as a technical lead engineer for the RF design, analysis and testing of satellite antennas. Chrystian holds a Bachelor of Science in Electrical Engineering with an orientation in space technology from the Ecole Polytechnique (Montreal, Canada).

Stephen LaRoy, BASc, P. Eng [stephen.laroy@bchydro.com] – Stephen is currently a telecommunications planner for BC Hydro. He is responsible for planning the telecom portion of IT and OT projects, evaluating new networking technologies, and setting the vision and direction for the expansion and maintenance of BC Hydro's MPLS network. He has previously worked on the deployment of BC Hydro's Smart Metering and Infrastructure program. Before coming to BC Hydro in 2008, Stephen worked as a product engineer at PMC-Sierra. Stephen holds a Bachelor of Applied Science in Engineering Physics from the University of British Columbia (Vancouver, BC).

Contributing Authors

Richard Spaeth, P. Eng, SaskPower [Canada] <rspaeth@saskpower.com>