21, rue d'Artois, F-75008 PARIS
http : //www.cigre.org

CIGRE US National Committee
2018 Grid of the Future Symposium

Active Management of Communication Resources in a Modern Power System Environment

B. R. WARD, P. J. ZAWADA
American Electric Power
USA

SUMMARY

Early in the migration to packet network technology, substation WAN connectivity was rather simple because it often supported SCADA traffic and perhaps one or two other low-bandwidth applications. The simplicity of these networks meant they were often straightforward to implement and did not require a great deal of upfront planning or long-term maintenance. However, as more packet-enabled substation applications have emerged, competing demands have been placed on limited WAN capacity. This is true despite the fact that technology advances have made more bandwidth available for WAN packet network. Furthermore, even when sufficient network capacity appears to be available, applications interact in such a way that more latency-sensitive applications, like remote synchrophasor measurements, can be negatively impacted by higher-bandwidth applications such as video surveillance. These impacts can be addressed by using network throughput modeling and Internet Protocol (IP) Quality of Service (QoS) mechanisms. These tools are utilized through an active management process, whereby WAN links are sized appropriately for the planned application mix according to a network throughput model as well as QoS marking and queuing policies applied on routers and switches, thus ensuring applications operate as intended. Once the network is constructed and these measures are implemented, long-term performance monitoring of the network should be executed to maintain optimal operation as well as to ensure traffic from new applications does not encroach on previously-modeled traffic.

KEYWORDS

Quality of Service, Class Based Queuing, Traffic throughput modeling, Internet Protocol, Time Division Multiplexing, Ethernet, Smart Grid Networks, Substation Wide Area Network(WAN), Local Area Network (LAN), Supervisory Control and Data Acquisition (SCADA), Phasor Measurement Unit (PMU), Syncrophasor, Substation Bandwidth, IPFIX.

1. Recent History of the data communication network in power systems. The move from TDM to Packet

The need for data communication networks in power systems evolved from the desire to remotely control and monitor the power system from a central location. Supervisory Control and Data Acquisition (SCADA) systems require information to be transported between the substation and the control center, conveying the state of the power system as well as assert desired control operations. A variety of analog technologies were used as the communications means for early relatively simple SCADA systems, often leveraging electric utility-specific technology such as Power Line Carrier (PLC). However, as microprocessor technologies allowed for more sophisticated SCADA applications that outpaced these primitive communications approaches, the utility industry leveraged telecommunications industry standards (which were benefitting from the same microelectronics revolution), often leading to the use of Time Division Multiplexing (TDM) technology for transport of substation communications. TDM technology worked well enough for utilities; however, after many years of successful deployment, the telecom industry was moving on to a more flexible and efficient packet-based approach.

TDM technology had been designed with the inherent assumption of carrying voice telephone calls at a fixed data rate of 64 kbps each, whereas nascent data services required a variety of transmission rates. Packet networking approaches, on the other hand, are adaptable to communication rates that vary over a wide range of speeds. As telecommunications carriers began serving more data needs instead of voice, they began deploying more and more packet-native networks and the decline of TDM technology commenced. Some of the relative inefficiencies of using TDM can be seen in the utility communications network, such as when multiple applications including SCADA, revenue metering, or telephony were all being used at the station. With TDM, a separate circuit with unique provisioning through the network is required for each of these allocations. Each separate circuit utilizes a channel in the TDM backhaul connection that has a static allocation of 64 kbps and that cannot be shared between the applications. Often, the application's throughput is less than this channel size, leaving the rest of the bandwidth wasted. All of this has been eliminated as the substation communication connection moved from TDM to a packet-based network.

Another advantage of modern Internet Protocol (IP) packet-based networks is that system intelligence is pushed to the end point. The TDM/analog telecommunications model was based on the assumption that the intelligence (primarily telephone call circuit switching) was built into the core of the network, hence the reason TDM-based networks require extensive provisioning of each connection through the network. So while packet-based networks do require an incremental amount of intelligence at the end point, the result is the network can be greatly simplified, allowing it to simultaneously transport data for a variety of applications without a great deal of administrative overhead. The network is designed to simply switch packets of data based on their destination, regardless of size or the type of data they contain. Furthermore, since the intelligence is moved to the end point, there are fewer restrictive assumptions (like capacity needs to be doled out in 64 kbps block) built into the network infrastructure, meaning new applications (i.e. end point types) can be easily added to take advantage of the network.[1]

AEP's use of a packet network for the Wide Area Network (WAN) connection started about 10 to 15 years ago with the addition of applications on the Local Area Network (LAN). The change started by connecting a router over one channel of the TDM backhaul, allowing the

packet network to be built as an overlay on top of the existing TDM network. The router could then provide communication to all the packet-enabled applications on the LAN. This was done primarily as a means of migration because, over time, all applications in the substation became packet-capable and moved to the LAN. The router and WAN connection to a station replaces the need for a separate circuit for each application. With a packet-based system, all communications outside the network would move through the router; the router would then direct the traffic to the correct location. Figure 1 illustrates the evolution from the TDM based network to a packet based network. In the first few years of deployment at AEP, packet networks were deployed primarily to support IP-connected RTUs as well as a few other low throughput applications such as remote engineering access and revenue meter gathering.

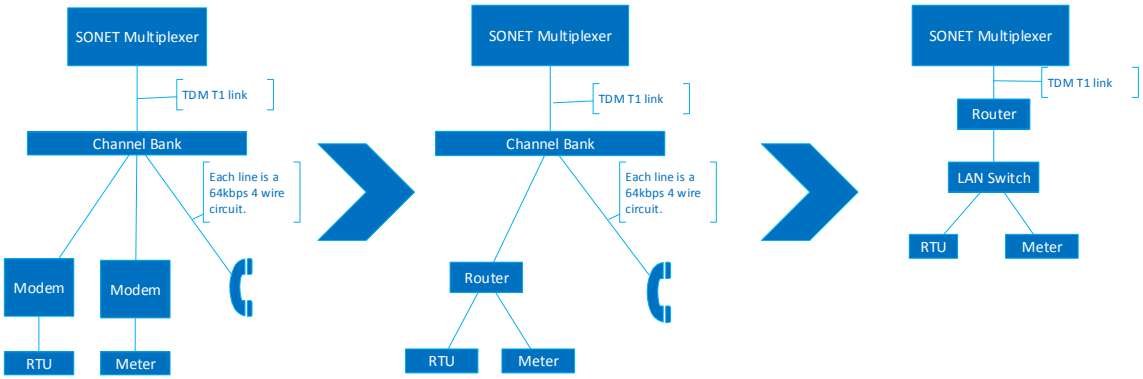

Figure 1: Evolution of the communication system from TDM to a packet Based network

2. The maturity of the IP packet based network.

For a number of years at AEP, packet network deployment for substation backhaul focused on SCADA connectivity. However, over time a number of other substation applications became packet network-aware and gradually moved into the substation network technology space. For example, the introduction of microprocessor relays with Ethernet and TCP/IP capability not only allowed for end-to-end packet connectivity for SCADA, but also enabled the implementation of Distributed Digital Fault Recording (DDFR). (With a DDFR implementation, protective relays create disturbance records and a variety of other files useful for post-event analysis, rather than relying on dedicated disturbance-monitoring equipment – these files must be managed and moved across a communications network.) Likewise, in recent years an increasing amount of data has been collected from Intelligent Electronic Devices (IEDs) in the substation for long-term trending and analysis to supplement the real-time situational awareness information provided by SCADA systems. Furthermore, in addition to substation-centric technologies, utilities have followed technology trends and moved enterprise applications to packet networks. Examples of this include Voice over IP (VoIP) for telephony and IP-enabled video cameras for security. Mobile work forces have also become more reliant on enterprise information systems, driving more demand for corporate network connectivity at substations, especially those located in areas with little or no mobile wireless coverage. Metcalfe's Law, which states "the value of a network goes up as the as the square of the number of users", starts to become realized [3]. As represented in Figure 2, an increase in packet network-enabled end points drives up the network usage. The utility industry is taking advantage of the new technology, resulting in the communications network handling more traffic than it was previously.

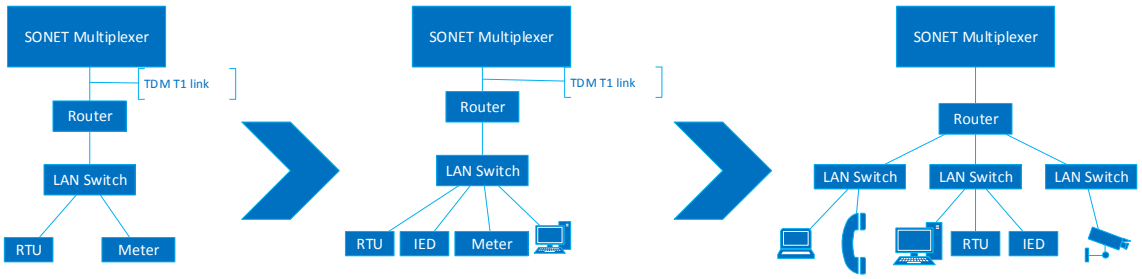

Figure 2: Maturity of the packet based network

Initially, there were very few problems with the communications network and management of the network was simple. IP-enabled SCADA is not particularly bandwidth-intensive, and most backhaul circuits were more than adequately sized to handle this traffic. However with the addition of the new applications on the network, the bandwidth requirements for each application had the potential to overload the network. This could and has resulted in degraded performance for the network devices or in extreme cases, even an application outage. One of the earliest examples of this issue occurred with AEP's Station Data Repository (SDR) platform, a Microsoft Windows-based substation computer responsible for gathering and transmitting DDFR files from protective relays to a central repository. The DDFR application itself did not have high bandwidth requirements. However, the maintenance of the computer itself through remote application of operating system patches and anti-virus updates would often tax substation network connections beyond their capability. Most concerning, the file transfers required to perform these updates would often cause temporary SCADA outages. This situation was a real-life example of the classic "Mice and Elephant flow" problem where a single large-volume data flow (the file transfer, which represents the "Elephant flow") can dominate a constrained network resource to the detriment of lower-volume flows (represented by the SCADA connection, a "mouse flow"). [4]

The problem can be attributed to the fact that at the time there was not a strong focus on managing the bandwidth at the station. Additional applications would be added without analyzing the requirements of all the existing as well as the new applications. While a packet networking approach is more efficient than a TDM approach because network capacity can be shared between applications dynamically, it is not immune to problems, especially in cases where oversubscription can occur. In an ideal situation, sufficient network capacity would be provided to meet all application needs at peak utilization. Often, though, compromises must be made when capacity is constrained. With TDM networks, these compromises were managed with static capacity allocations of 64 kbps timeslots. With packet networks, careful attention must be paid to network sizing to make sure the network can handle nominal traffic patterns. In addition, Quality of Service mechanisms must be in place to meet minimal application requirements during time of heavy network loading.

At this point, it is important to bring up the fact that increased backhaul capacity alone is not the sole solution to the challenges of operating packet networks. AEP, like many other utilities, has been able to increase the number of company-owned fiber optic backhaul connections to substations. However, in AEP's experience, the core packet switching infrastructure often struggles to keep pace with the growth of the backhaul connectivity. This means attention to network utilization and architecture needs to be placed across the entire network, from the individual substations throughout the core of the network. This is not unlike the traditional TDM approach; however, TDM networks forced core network upgrades to be made before substation bandwidth upgrades could be enabled. With packet network

connectivity, it is entirely possible to upgrade individual elements of the network, leaving others untouched and potentially oversubscribed. Because of this, AEP network engineers have often found it necessary to artificially constrain traffic on some links until core infrastructure can catch up with fiber optic transport being deployed at the edge of the system. Measures must be taken to help mitigate these challenges associated with the growth of utility communications systems.

3. Solution for Dealing with a Modern Power System Communications network.

While packet networks allow for more efficient use of network capacity by eliminating static allocation of link resources, there are limits to this flexibility. If links are oversubscribed by streams of relentless application traffic, it is impossible for a packet network to compensate, and network performance will be severely degraded if not completely compromised. The seemingly obvious way to approach this issue is to carefully plan for the applications to be supported, tallying up the capacity required for each of them. This was the required approach for TDM-based networks since if an application was not assigned at least one 64 kbps timeslot, it would not be able to communicate. However, this approach is what led to inefficiencies and inflexibility in the first place. Specifically, by planning for worst-case (maximum) throughput in increments of 64 kbps, as was necessary in a TDM network, the intermittent nature of many types of traffic cannot be exploited. Unfortunately, many modern network applications are not entirely predictable, so developing sophisticated, highly accurate predictive models to maximize efficiency while completely preventing link oversubscription is extremely difficult if not impossible. AEP's solution to this issue is to build a simplified network traffic model based on calculations and empirical network traffic measurements, using it to size the WAN connection to the station. In addition, combine the model with standard Quality of Service (QoS) mechanisms to protect the most sensitive traffic types from degradation of performance should a link be driven to oversubscription.

The intent of the simplified traffic model is to prevent gross oversubscription by addressing fairly consistent network capacity needs with a fairly static allocation and then adding additional "headroom" for the less-consistent, less-predictable traffic. Some utility applications, such as SCADA and synchrophasor measurements, have rather constant traffic profiles with rigorous capacity needs in order to meet real-time requirements. Other applications, such as file transfer and remote engineering access, have unpredictable traffic patterns yet are more flexible in the amount of network capacity that they require. For example, file transfers occur only occasionally and may operate just fine on reduced network capacity although they may take longer to complete – a perfectly acceptable level of service for non-real-time functions. Still other classes of traffic represent a combination of these two types; their operation is unpredictable but requires a minimum amount of bandwidth while in operation. As seen in Figure 3, the AEP traffic model accounts for all applications in the "Technology" column and attempts to break the application traffic into two components – a "constant" bandwidth portion and an "on-demand" bandwidth portion. The constant bandwidth components represent the application traffic that is always occurring. These portions are simply summed to obtain a "constant" total. The on-demand components represent parts of the application that tend to be interactive or periodically send information. To represent this intermittent use, each application is scaled by a demand factor before summing with the constant bandwidth components. (In certain instances, the demand factor may be overwritten in order to guarantee sufficient capacity for those intermittent services that have a minimum required throughput.) This approach ensures a minimum level of

capacity for on-going application needs while attempting to reduce the over-provisioning of bandwidth for applications that require service less often. Once the anticipated traffic is modelled, a backhaul with sufficient capacity is selected. This approach is not perfect; occasional oversubscription is possible, but the effects of this oversubscription can be mitigated by utilizing QoS mechanisms.

Technology	Quantity	Constant Bandwidth Per Device (Kbps)	Constant Bandwidth Required (Kbps)	On Demand Bandwidth Per Device (Kbps)	On Demand Multiplier	On Demand Bandwidth Per Technology (Kbps)	Demand Factor	On Demand Bandwidth Required (Kbps)
VOIP	1			33	1	33	0.5	16.5
METERS	2			10	2	20	0.05	1
CORP LAN	1			100	1	100	0.2	20
CAMERAS	3			512	3	1536	0.2	1536
SDR	1			100	1	100	0.4	40
AH DR-E3	3	2	6	100	1	100	0.2	20
AH OCU	3	10	30	100	1	100	0.2	20
RTU	1	16	16					
PMU	3	58	174					
TOTAL CONSTANT BANDWIDTH			226		TOTAL ON DEMAND BANDWIDTH			1653.5
					TOTAL STATION BANDWIDTH			1879.5

Figure 3: Network Throughput Model Calculation

Quality of Service (QoS) mechanisms provide the ability to manipulate how packets are processed through the network. The initial step in implementing a QoS architecture is to categorize traffic flows associated with various applications into separate classes and assign each class the associated priority based on the importance of the application. The priority will decide which applications continue uninterrupted in the face of congestion. To accomplish this, Ethernet switches identify these traffic flows through a variety criteria including Transmission Control protocol (TCP) ports or IP address. This information is found in each IP packet header and is kept the same as the packet moves throughout the network. Every packet for each type of traffic flow is marked using a specific value in the Differentiated Services Code Point (DSCP) field in the IP packet header. Packet flows are now marked and can be grouped into particular classes based on the DSCP bits that were set. Each class is associated with a queue and class-based weighted fair queuing is utilized to assign a minimum guaranteed throughput for the class. At this point, the traffic can fall into two profiles; the first is "under the set throughput" and second is "over the set throughput". As the queues are managed, traffic above the set throughput will not benefit as much from the queuing as traffic below the set throughput. Allocated bandwidth that goes unused can in turn be used to service other classes that are in the "over the set throughput" profile. This traffic will be prioritized and shaped based on the class that the traffic falls into. Implementing QoS in this way ensures the benefits of dynamic capacity allocation associated with packet networks can be retained when a link is not running at full utilization. As the throughput increases, traffic will be shaped to ensure each class is able to use at least the amount of bandwidth allocated as well as giving preferential treatment to higher priority traffic. [2]

When the traffic is modelled and the substation WAN connection is sized properly, QoS is applied to the network to minimize any unforeseen traffic performance issues. An example of

this can be seen in Synchrophasor Data collection. A Synchrophasor is a time synchronized measurement of the phase angle and magnitude of the sine waves in the power grid. These phase angle measurements are taken with high-speed Phasor Measurement Units (PMU). The data then is continuously streamed, requiring low latency communication to be effectively used in real time. If the PMU traffic is placed in the same class as other high throughput applications, such as video surveillance equipment, it will not be given a higher priority. Since synchrophasor data is used for real-time continuous calculations, significant latency (currently understood to be that greater than 500 ms) will limit the abilities of Wide Area Monitoring Systems (WAMS) and degrade the solution of Linear State Estimators (LSEs). Those systems will consider the late data as missing and perform their functions with the low-latency synchrophasor data they have received on a timely basis. To solve this problem, the PMU traffic can be placed in a class that has a higher priority, allowing the time-sensitive traffic to be processed first. The figures 4, 5, and 6 below demonstrate this situation. In this case, the WAN connection was sized properly but there were still traffic performance issues. Once the QoS at the substation was updated, the performance issues were eliminated. Figure 4 shows two plots, the first represents traffic from video surveillance equipment measured with a packet capture software; the second is a plot of the measured latency from a PMU that shared the same backhaul connection. Figure 5 illustrates the impact of video traffic on the PMU latency. When video throughput increases, the PMU shows a corresponding increase in latency. Figure 6 is a plot of the same PMU latency while using the video equipment after classification of PMU traffic priority to place it in a higher priority queue

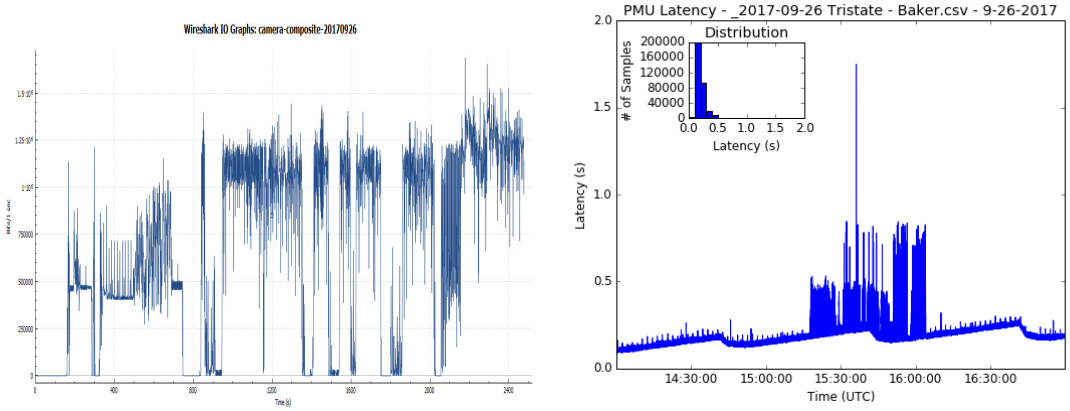

Figure 4: Plot of Bandwidth usage for video equipment (Left), Plot of PMU latency (Right)

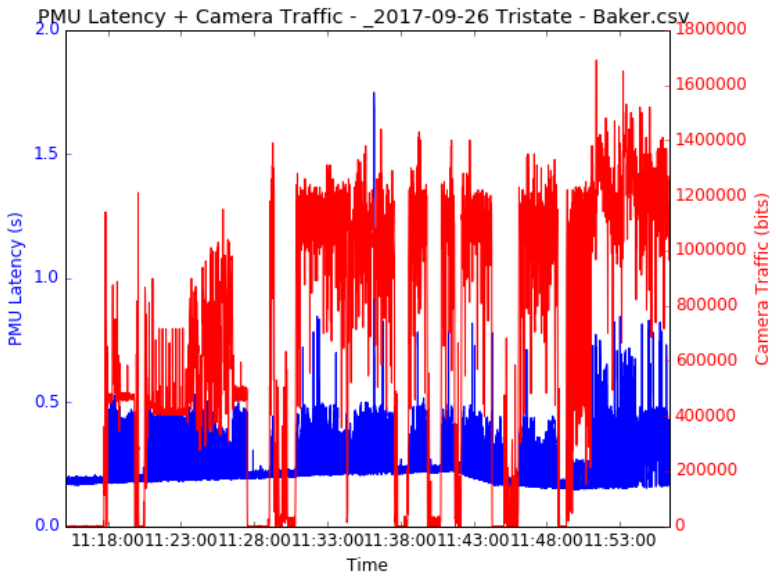

Figure 5: Video Equipment and PMU latency Plot Overlay

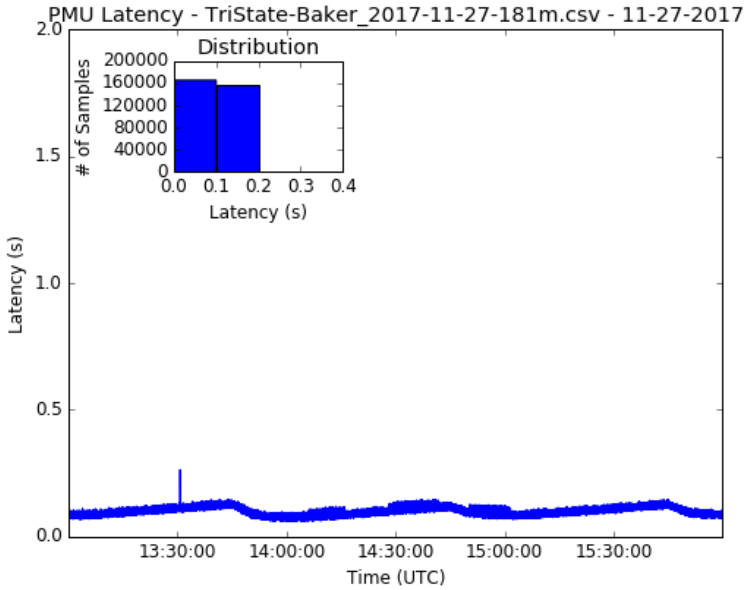

Figure 6: Plot of PMU latency during Video equipment usage after the QoS update.

Looking ahead, AEP hopes to enhance active management of the substation network by building the capability to continuously measure utilization of backhaul connectivity for every substation. This will allow the verification of the modelled performance of the network and help ensure substation applications have sufficient network capacity to function properly. IP Flow Information Export (IPFIX) protocol implementations, often referred to as "netflow", can provide switches, routers, and other network devices the ability to collect and summarize information about network flows then forward that information to a central network management system. [5] (A "flow" is a group of packets that make up a distinct application conversation across the network and is characterized by the so-called "5-tuple": IP source and

destination addresses, the IP protocol type, and the source and destination ports.) By collecting flow information from across the network and storing the information in a historian for long-term trending and analysis, network performance can be compared against a stored model to identify unanticipated traffic growth and aid capacity planning for the telecommunications network. Since the data can also be organized on a per-DSCP class basis, performance of the QoS policy can likewise be monitored and analysed on a per-substation basis. The ability to collect and analyse flow data will essentially close the loop of the plan-construct-verify engineering cycle, giving AEP a great deal of insight in the planning and performance of the substation WAN.

CONCLUSION

The substation communication network has evolved as new telecommunication technology trends led to the replacement of voice-centric TDM approaches in favor of a data-friendly packet-based approach. This new packet-based system was simple in the beginning, requiring little maintenance and low upfront planning. As both utility-based applications and the communications network were able to benefit from technology advances, the growth of applications outpaced the growth of the communications network, creating a more complex network. This resulted in a more complex network that can place competing demands on the WAN packet network capacity. Even in cases where there is enough bandwidth at the station, applications can have negative effects on each other, as seen with latency-sensitive applications being delayed because of other higher-bandwidth applications. The increasing number of applications in the utility industry provides us with great advantages for the power system. All of these advantages come at the small cost of using tools such as IP QoS, network throughput modeling, and long-term performance monitoring for the active management of our communications resources. To achieve this, a firm knowledge of both the power system applications and how the telecommunication networks operate is required.

ACKNOWLEDGEMENT

The authors would like to thank Brad Holt for technical expertise and assistance with the content. In addition, thank you to all colleges for providing comments to improve the paper.

BIBLIOGRAPHY

[1] Isenberg, David. *The Rise of the Stupid Network.* (**Computer Telephony***,* August 1997, pp. 16-26)
[2] Paquet, Chrystian & LaRoy, Stephen. *Quality of Service in power utility telecommunications networks*. (Utilities Technology council of Canada, September 2017)
[3] Shapiro, Carl and Varian, Hal R. *Information Rules.* (Harvard Business Press.1999)
[4] Guo, Liang and Matta, Ibrahim. *The War between Mice and Elephants*. (Computer Science Department, Boston University 2001)
[5] B. Claise, B. Trammell, and P. Aitken. "*Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information*"(IETF, RFC 7011, September 2013)