NOKIA

Getting more with public safety SDN

How SDN optimizes microwave network use while maintaining stringent SLAs

Strategic white paper

NOKIA

Contents

The advent of software-defined networking technology

Software-defined networking (SDN) is a new way of operating networks that enables better end-to-end control, automation and, service agility.[1] With SDN, the network's control and forwarding planes are separated physically (Figure 1). The control plane function resides in a centralized platform called the SDN controller. This is equipped with an open, northbound application programming interface (API). The network is programmable, allowing other applications to programmatically control and monitor the network via the controller.

Figure 1. An SDN network architecture

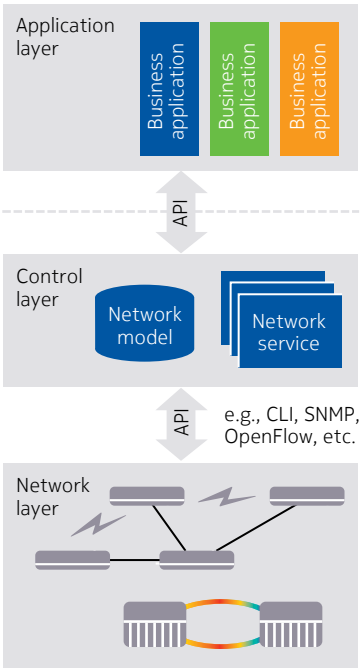

SDN enables service providers to be agile and meet ever-changing, on-demand dynamic service requirements. However, in a public safety network where the use of microwave transmission is dominant, connectivity required by major applications such as land mobile radio backhaul is usually static. Consequently, at first sight, SDN does not seem particularly applicable to public safety. However, the public safety networking landscape is shifting. This paper describes how public safety network operators can exploit SDN in this new environment.

1 For a detailed discussion of what SDN is, please read the Metro Ethernet Forum SDN whitepaper. (https://www.mef.net/Assets/White_Papers/Carrier_Ethernet_and_SDN_Part_1_-_An_Industry_Perspective_08-14-14.pdf)

The shifting landscape of the public safety network

The emergence of a shared public safety network architecture

Beyond on-line shopping or entertainment, instant connectedness extends to all forms of public service—e-government services included. From IT applications used by government staff to fast Internet access from public libraries and schools, the demand for network connectivity back to data centers, ministry headquarters, or the Internet has risen inexorably. To satisfy these demands, different government departments have resorted to building their own network or leased-line services. However, because governments at all levels are faced with increasing budget constraints, they are being required to do more with less. The current network operating paradigm of disparate service networks is proving too costly and inefficient for today's growing network demands (Figure 2).

Figure 2. Disparate service networks

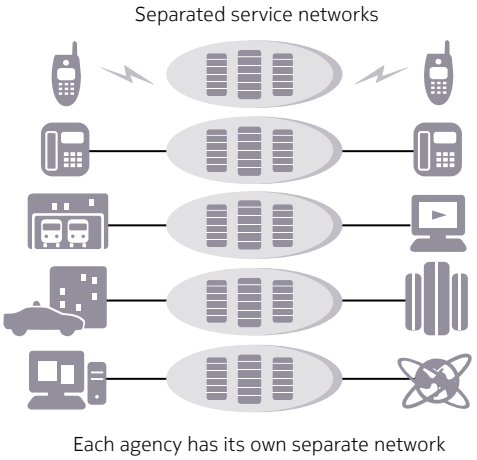

As a result, many governments are looking for an innovative network architecture capable of leveraging the full capabilities of the latest networking technologies in order to achieve a more efficient, cost-effective operation while, at the same time, delivering improved quality of public services for citizens. The same goal is at work in the public safety domain. As public safety agencies' operations become ever more dependent on broadband data to increase situational awareness to provide faster and more targeted responses to incidents, these same agencies are swiftly adopting the latest packet microwave technology when deploying new public safety backhaul networks. This rapidly growing trend provides an ideal opportunity to explore this new shared network architecture.[2]

2 For a more detailed discussion of a shared architecture, please read, "Delivering a Public Sector Shared Architecture." (http://resources.Nokia.com/asset/183929)

To begin with, the new network architecture enables shared use of network resources among different government departments, beyond just public safety agencies. This capability enables government to conserve financial resources, while also maintaining the performance integrity of mission-critical public safety applications without compromising the privacy and integrity of all participants. The network can carry a multitude of applications—from first responders' mission-critical applications to government departmental IT applications, to public amenities' best-effort Internet traffic (Figure 3). Each can be given a corresponding priority in order to attain the necessary quality of service (QoS),[3] thus meeting its service level agreement (SLA).[4]

Figure 3. Shared public safety network

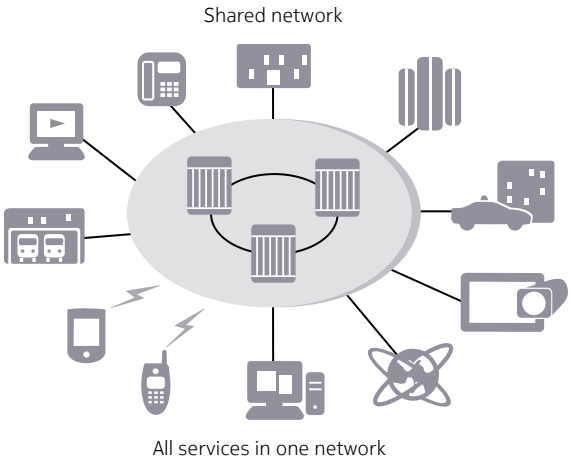

In the shared network architecture, mission-critical traffic (hereafter called "critical traffic") is forwarded with high priority and assurance to attain deterministic QoS while non-mission-critical traffic (hereafter called "non-critical traffic") is forwarded with low priority. This allows it to be discarded when network congestion arises. Because packet discard of non-critical traffic neither threatens nor results in loss of life, it is generally considered acceptable network operation practice. Nonetheless, because packet loss causes lower data throughput resulting in slower application response times, packet loss impacts productivity and Internet surfing experiences from public amenities.

Extracting more from the shared network

Not surprisingly, the shared network architecture has been attracting considerable attention from network and IT managers due largely to its potential to deliver economic benefits to governments.[5] Governments can benefit from defrayed network operating costs while departments and agencies can gain access to advanced, reliable, resilient and secure services delivery. The more non-critical IT and Internet applications that ride on the shared network, the more savings government can achieve. As noted earlier,

3 For a more detailed discussion of a converged public safety backhaul network, please read "Mission-Critical Communications Networks for Public Safety." (http://resources.Nokia.com/asset/170422)

4 An SLA is usually a set of committed communication performance metrics, including metrics such as availability.

5 For the economic benefits of a shared network architecture, please read "The economic benefits of sharing government communications networks." (http://resources.Nokia.com/asset/187929)

however, network congestion causes productivity loss and deters such movement. As a result, if congestion can be avoided so that applications can consistently attain the pre-agreed SLAs, more government departments and agencies will look to use the network, sharing the network cost.

Quantum jump in microwave link capacity

One obvious way to increase microwave link capacity is to increase its channel bandwidth. But this is not always feasible. Suitable microwave spectrum is usually already crowded. However, with the advent of packet microwave technology, new techniques such as packet compression (also known as packet throughput booster) and cross-polar interference cancellation (XPIC)[6] have become available to increase capacity without increasing the channel size.

Another promising technique is high-order adaptive modulation. Adaptive modulation is a technique to fully use microwave spectrum by changing the modulation scheme, according to propagation availability. By adjusting the modulation scheme as atmospheric conditions change, more radio throughput is attained. With more processing power, the modulation level can reach beyond the traditional ceiling of 256 QAM. High-order adaptive modulation at 2048 QAM can increase the highest link bandwidth by as much as 50 percent (Figure 4). However, it should also be noted that operating at such a high modulation level requires a high signal receive level, which can be dampened by atmospheric events such as rain (Figure 5).

Figure 4. Increased link capacity with high order modulation

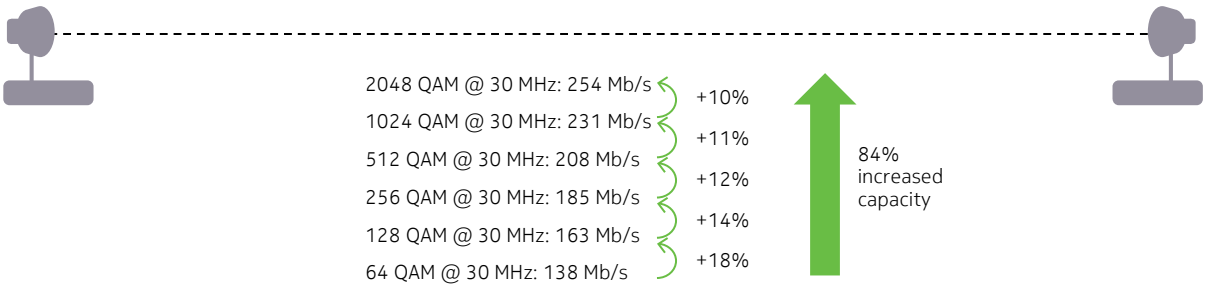

Figure 5. A study of availability of high order adaptive modulation link

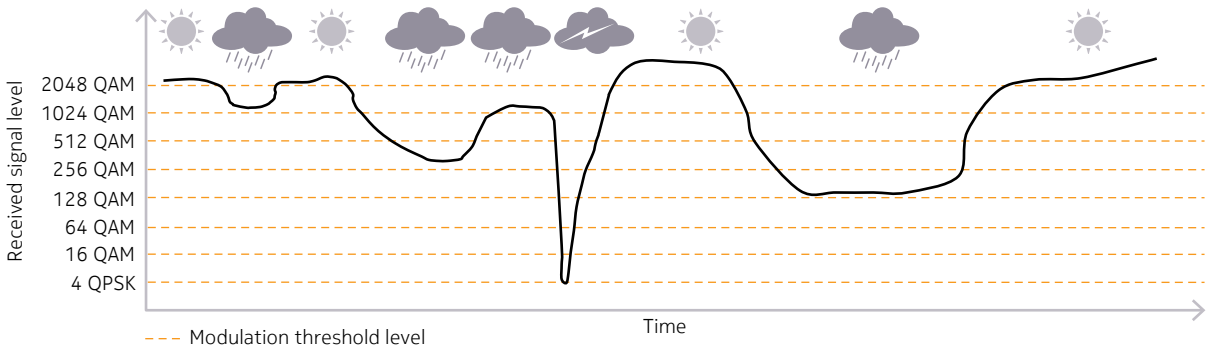

6 For more information on multichannel technology, please read "Creating Faster, More Efficient and More Reliable Microwave Links" (http://resources.Nokia.com/asset/168173)

When the modulation level drops, so does the available bandwidth. With a proper QoS policy and network engineering, high priority critical service traffic can be delivered with no compromise. However, low priority non-critical traffic is discarded due to link congestion. As explained earlier, this does not lead to life-and-death situations. Even so, government workers' IT application query responses and citizens' Internet surfing experiences at public amenities could be significantly slowed down. This results not only in lower office productivity but also degrades citizens' Internet surfing experiences during periods of low modulation. More importantly, it could also violate the pre-agreed SLAs and discourage government network operators from migrating more non-critical traffic to the new network. Consequently, the full potential of the public safety network remains unleashed.

Boosting non-critical applications performance

Despite these challenges, there is now an opportunity to remedy the situation—to unleash the network infrastructure's full potential. With the advent of software defined network (SDN) technology, network operators can improve non-critical application performance during inclement weather, raise employee productivity, and elevate the quality of citizens' surfing experience; in sum, network operators can meet SLAs of non-critical applications more consistently, giving them the confidence to move more non-critical traffic, thus unleashing the full potential afforded by the architecture.

This paper now shifts its focus to look more closely at the keys to consistent SLA fulfillment in a public safety microwave network. The two most important are as follows:

1. Build the microwave network with rich path diversity.
2. Deploy SDN to continuously monitor network key performance indicators (KPI),[7] and optimally re-route non-critical traffic to ensure SLAs during performance degradation.

Rich path diversity

A public safety network infrastructure blueprint typically has a hierarchical multi-ring topology (Figure 6):

1. Fiber core rings (R1 and R2 in Figure 6)
2. Microwave aggregation rings (R3, R4 and R5 in Figure 6) and,
3. Microwave spurs for remote sites where forming a ring is challenging or cost prohibitive

One key advantage of this topology is multi-fault tolerant network resiliency.[8] With rich path diversity, when multiple failures occur simultaneously, there is a better chance to restore connectivity for mission-critical applications whose accessibility could be a matter of life and death to first responders and citizens in need.

7 Key performance indicators (KPIs) are a set of network parameters performance objectives. These can include parameters such as bandwidth use, one-way delay, packet loss, and availability.

8 For more details on a multi-fault tolerant network, please read "Building a multi-fault tolerant microwave backhaul network." http://resources.Nokia.com/asset/175593

Figure 6. Reference public safety network infrastructure topology

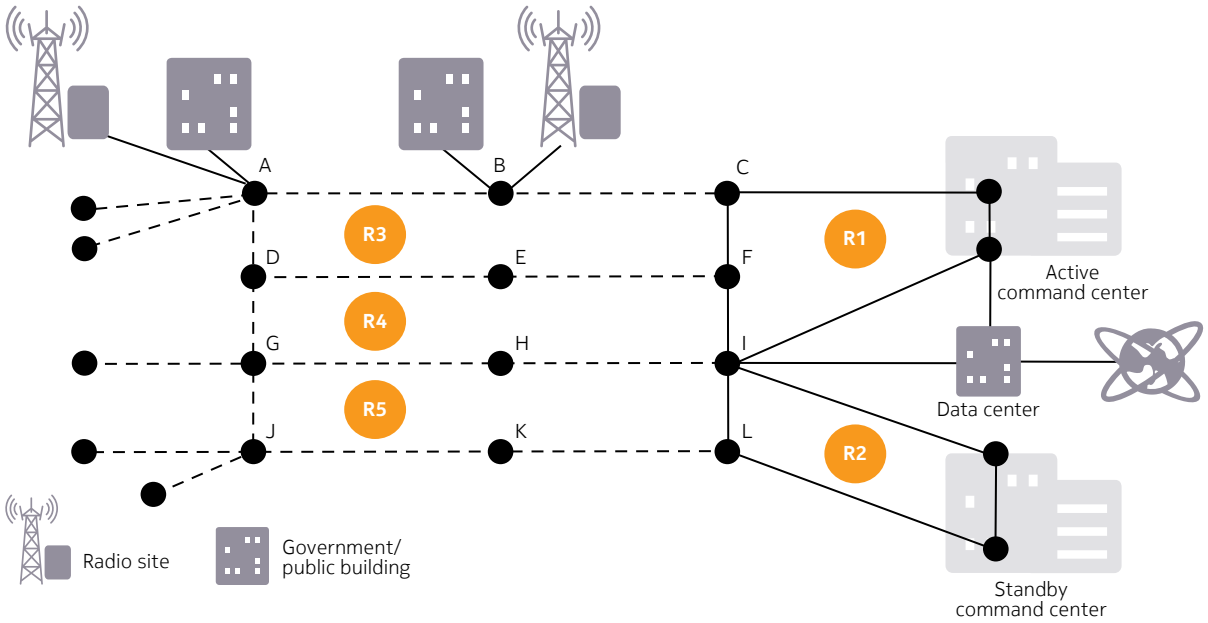

This hierarchical multi-ring topology provides rich path diversity with multiple options for traffic to reach its destination. For example, site A and B are connected to locations supporting both critical and non-critical traffic: critical traffic from a land mobile radio (LMR) base station, which is destined for an active command center, as well as non-critical traffic from government offices or public amenities for a data center and the Internet. For site A, there is a rich selection of paths (A-B-C, A-D-E-F, A-D-G-H-I and A-D-G-J-K-L-I) to reach the fiber ring (R1) before critical and non-critical data part ways to their respective destinations (Figure 7).

Figure 7. Rich selection of paths available

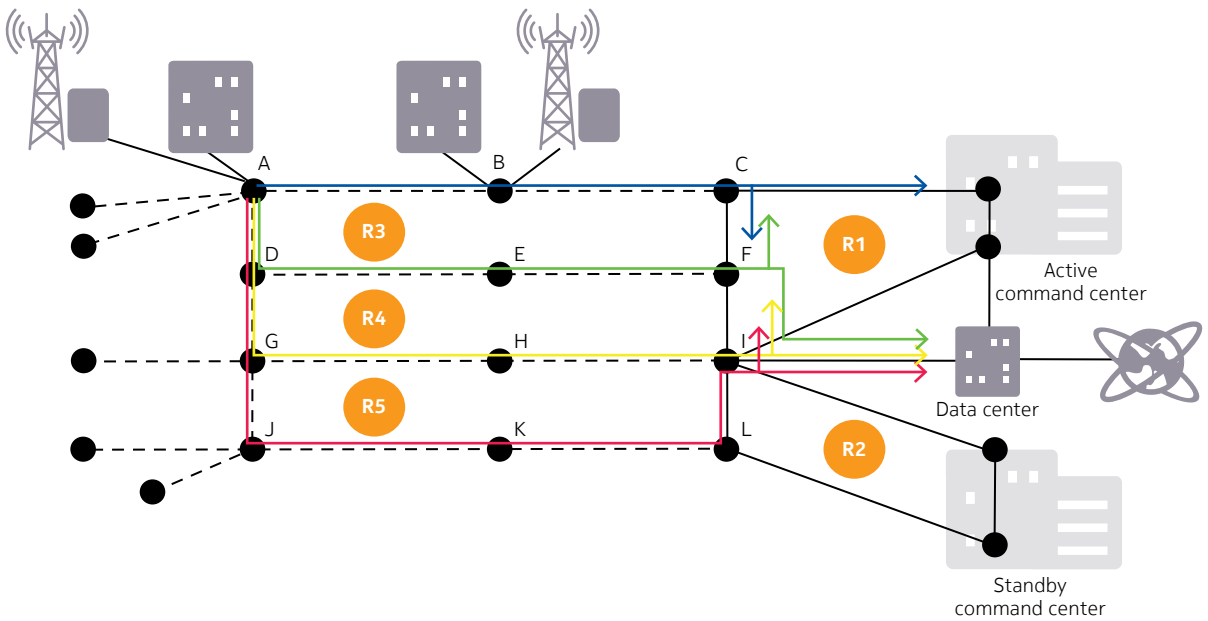

Under normal weather conditions, microwave links operate at 2048 QAM and can offer enough bandwidth for all traffic (critical and non-critical). If a severe rain storm hits the area's central north portion, links BC and EF are buffeted (Figure 9). Adapting to the worsened atmospheric conditions, the links can quickly lower the modulation scheme from 2048 QAM to a lower level such as 64 QAM, thereby reducing the link capacity by as much as half. As explained earlier, while the delivery of critical service traffic from the radio site is never compromised, the same cannot be said of non-critical applications from government offices and public amenities connected to sites A and B, as well as sites E and F (traffic from E and F are not included in this discussion). Non-critical traffic from A and B, being lower-prioritized, is partially discarded, due to insufficient bandwidth in the microwave links. This, in turn, causes a drop in traffic throughput, affecting the performance of IT applications, as well as the Internet surfing experience. The net result is that SLAs are violated.

Normally, with shortest-path based route computation, services concentrate at the shortest path. In this example, non-critical traffic from site A and B follows path B-C-F-I to reach the data center (Figure 8).

Figure 8. Shortest path-based computation concentrates links on common path

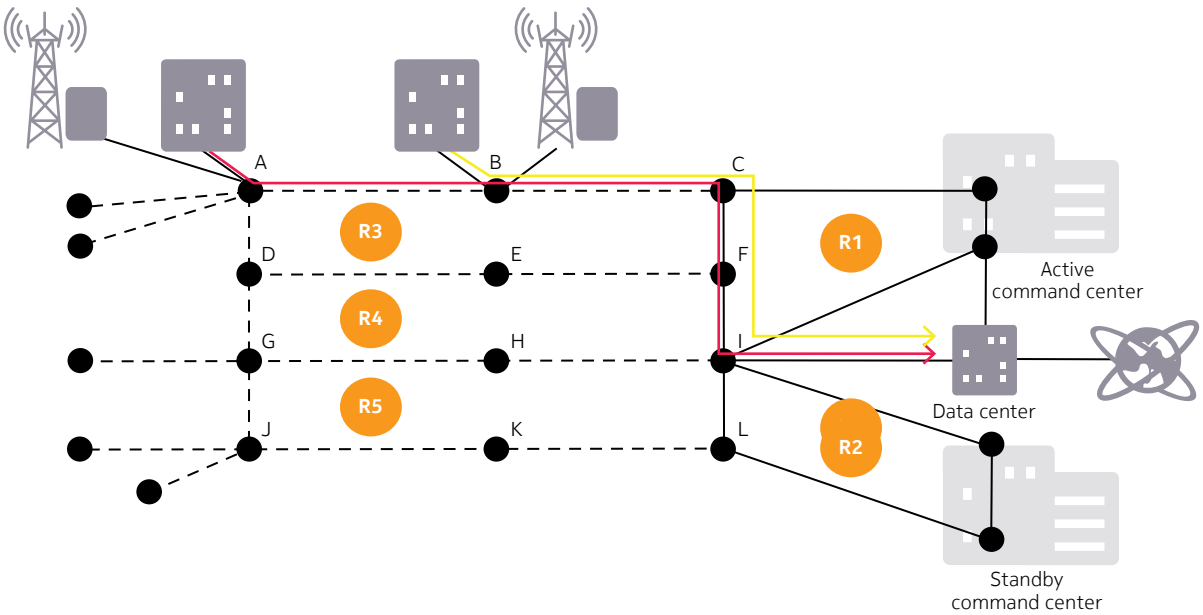

Meanwhile, because areas in the southern segments of R4 and R5 are not hit as hard by the storm, it is probable that there is unused bandwidth in those microwave links (in other words, paths A-D-G-H-I and A-D-G-J-K-L-I). Even though these paths have more hops, thus incurring a larger latency, they are still acceptable to non-delay-sensitive IT applications. However, because these paths remain idle, network resources are not being used maximally while some non-critical applications performance is penalized.

Optimal re-route of non-critical traffic

To restore the performance of affected non-critical applications originating from public offices and amenities, the traffic can be re-routed. It can be rerouted from the current congested paths that traverse the buffeted microwave links to the unaffected links, which likely have bandwidth available. In this way, the use of deployed network assets is maximized to attain optimal applications performance.

To accomplish this goal, there are two general approaches:

1. Distributed node-based optimization
2. SDN-based optimization

Distributed on-node optimization

As microwave links BC, CD, FG and GD step down the modulation level in response to the storm, the link capacities joining these nodes are reduced. Upon detecting the associated service impact, each node (A, B, E, F and G) that employs an intelligent path computation algorithm will look for new paths that have higher available bandwidth over which to route affected non-critical traffic. This distributed approach, though, has two shortcomings:

1. Hot spot and congestion point re-occurrence after re-routing

   Information on link utilization, bandwidth availability, and services is locally known. Therefore, each node calculates alternative routes based on the limited information pertaining to its local view instead of a network-wide view of services and link restrictions. Furthermore, due to the distributed nature of this computation, each affected node, running the same path Dijkstra-based algorithm, selects the same new path (in this example, A-E-H-I-J-K) although alternative paths (for example, A-E-H-L-M-N-O) are available. This results in all traffic being re-routed to the same new path creating new congestion points. As a result, this distributed approach fails to consider network-wide bandwidth and resources availability to re-route affected applications in a balanced and optimal way.

2. No support of policy- and constraint- and metric-based optimization[9]

   The on-node approach lacks the capability to apply more sophisticated business policies (e.g., re-route government IT traffic before public library Internet traffic) and to apply constraints (e.g., end-to-end path must not exceed 10 hops).

9 Constraint- and metric-based optimization goes beyond the typical cost-based computation. It includes the number of hops, spans, latency, distance or any combination thereof.

NOKIA

SDN-based optimization

A SDN-based optimization approach is based on the SDN controller embedded with an adaptive and powerful path computation engine (PCE).[10] This controller, equipped with full service information (endpoints, priority, and QoS etc.), as well as network-wide information (topology and link use etc.), can make use of advanced analytics algorithms and techniques, such as linear programming and graph theory in order to compute new paths for affected service based on network KPI information. Its analytics engine is also able to derive the necessary information to make automated business policy decisions.

Getting more from a public safety SDN

For these reasons, a public safety SDN is well-suited to optimally balance traffic during re-routing to restore affected non-critical applications based on network KPIs. Based on a global view of network use in addition to each application's business policy constraints and SLA requirements, the SDN controller's PCE can intelligently compute optimal new paths. These re-direct traffic from affected services along possible paths without creating new congestion points.

Based on the Figure 9 rainstorm example, the following step-by-step guide shows how SDN can self-adapt to shrinking microwave link bandwidth and optimally re-balance non-critical traffic in order to attain higher throughout performance:

1. The SDN controller monitors bandwidth use of all network links and network performance for both critical and non-critical service. Network performance monitoring can be done by continuously running OAM sessions (for example, ITU-T Y.1731 Continuity Check) or by a dedicated performance monitoring station to issue alerts when traffic loss occurs or by simply monitoring the link use percentage.
2. The rainstorm causes microwave links BC and EF to operate in lower modulation levels. If the level is lowered to 64 QAM from 2048 QAM, about half of the bandwidth is lost. With proper QoS, critical traffic is always preserved. However, the network link's KPI indicates congestion, thus potentially violating non-critical applications' SLAs.
3. The SDN controller learns about the service degradation for affected non-critical applications. If degradation continues after expiration of a pre-configured time, the SDN controller starts the PCE. The PCE algorithm performs analytics on the use of network-wide bandwidth to compute a set of the most optimally balanced paths.
4. The SDN controller commands corresponding nodes to re-route affected traffic along the new computed paths. If bandwidth in the network permits, it can re-route all non-critical traffic from nodes A and B to paths A-D-G-J-K-L-I and B-A-D-G-J-K-L-I, respectively (Figure 9).

10 1A PCE, as defined in IETF RFC4655 (https://tools.ietf.org/html/rfc4655), is a network entity or application that is capable of computing a network path or route based on a network graph with computational constraints applied. For a technical discussion of PCE and its benefits, please read "Computing a path to more profits: Centralized PCE using Bell Labs Self-Tuned Adaptive Routing." http://resources.Nokia.com/asset/186905

If available bandwidth is insufficient to accommodate all re-routed traffic, the SDN controller can apply a business policy in order to re-route the higher prioritized traffic flow (for example, government office applications over the public Internet). This kind of business intelligence gives operators full flexibility when applying SDN's network optimization capability.

As a result, throughput performance is optimally improved for non-critical applications. When the storm ends, the microwave link modulation returns to 2048 QAM. Similarly, all affected non-critical applications are re-routed back to their original paths.

Figure 9. SDN controller re-routes affected non-critical traffic in public safety SDN

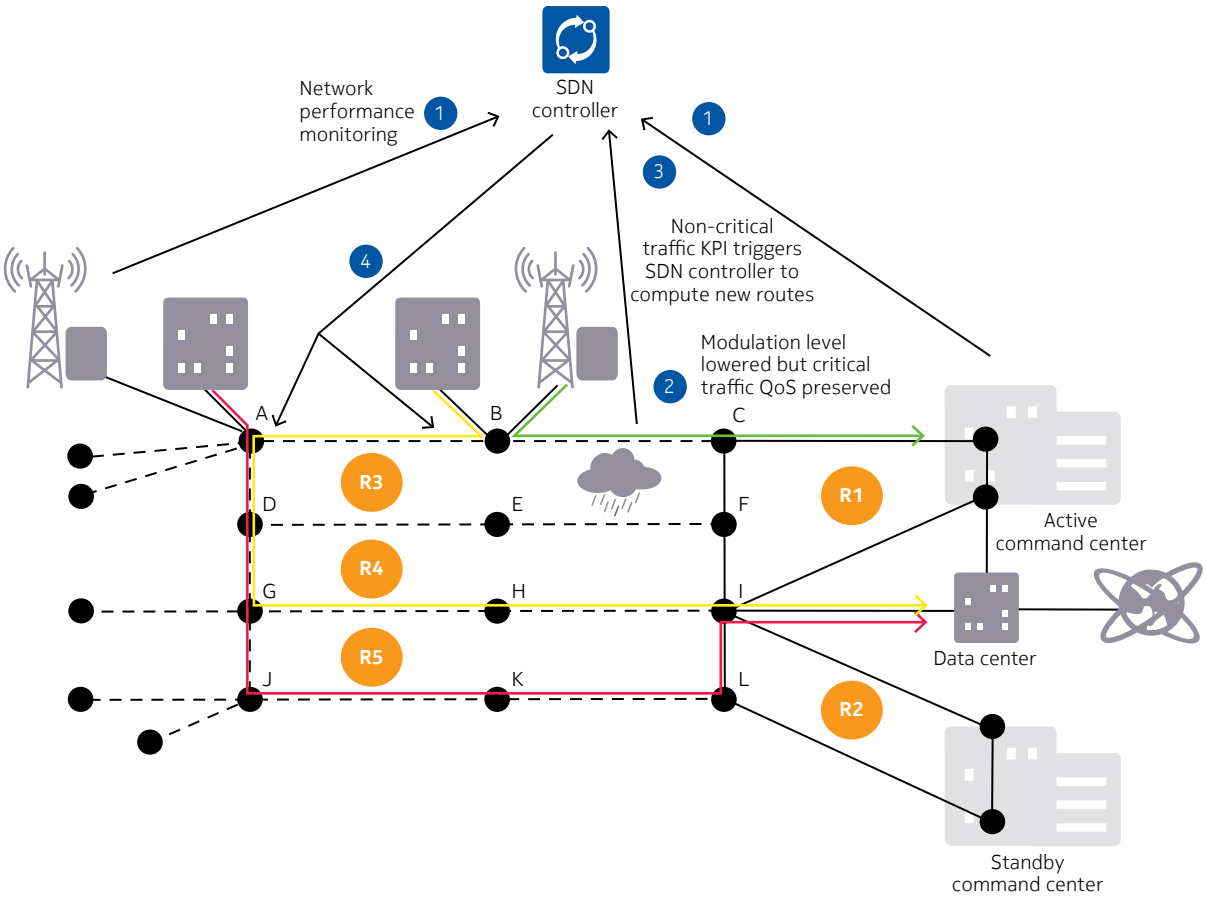

NOKIA

Conclusion

Governments today are at a critical juncture. With rapidly growing urban populations, stringent green environmental regulations, and escalating demands for an efficient, interactive government, government IT needs to adopt new technologies that empower first responders to provide a sharper response, as well as increase government office productivity—all while capping network costs.

Sharing a public network infrastructure among government agencies can effectively conserve financial resources to meet budget constraints. Without compromising critical applications, a public safety SDN lets operators transport more non-critical applications. This enables operators to maximize economic benefits while dynamically monitoring network KPIs to proactively maintain non-critical applications' SLAs with available network resources. Outcomes include increased government office productivity, a significantly improved e-government experience along with maximized use of public safety network assets.

A successful execution of government network transformation rests on more than technology. Professional services such as network design and consulting are equally crucial. Complemented by the comprehensive and innovative product portfolio that spans microwave and optics transport to IP/MPLS and SDN, Nokia can help governments build a network that serves citizens better, faster, and more efficiently.

Acronyms

CLI	Command-line interface
IT	Information technology
ITU-T	International Telecommunication Union - Telecommunication
IWF	InterWorking Function
KPI	Key Performance Indicator
LMR	Land Mobile Radio
MEF	Metro Ethernet Forum
MPLS	Multi-protocol label switching
PCEP	Path computation element protocol
QAM	Quadrature Amplitude Modulation
QoS	Quality of Service
SDN	Software Defined Network
SLA	Service Level Agreement
SNMPv3	Simple Network Management Protocol version 3
WDM	Wavelength Division Multiplexing
XPIC	Cross Polar Interference Cancellation

References

1. Nokia Network Services Platform
   https://networks.nokia.com/products/network-services-platform
2. Nokia 7705 Service Aggregation Router
   https://networks.nokia.com/products/7705-service-aggregation-router
3. Nokia 9500 Microwave Packet Radio
   https://networks.nokia.com/products/9500-microwave-packet-radio
4. Metro Ethernet Forum
   http://www.mef.net

Nokia Oyj
Karaportti 3
FI-02610 Espoo
Finland
Tel. +358 (0) 10 44 88 000

Product code: PR1605020304EN (July)