SME Technology Development Program

Anchor Firm	ciena
Challenge Statement	**STRAT – STructured Re-Arranged Topology for Enterprise Data Centers**

Challenge Launch Date	November 6, 2019
Challenge Deadline	November 28, 2019
Challenge Statement	Efficient data centers are critical to the evolving 5G infrastructure, which is expected to provide both connectivity and low-latency compute and processing to a number of applications spanning a range of industrial, automotive, entertainment, civic infrastructure, etc. Data centers become a more critical component throughout the network, starting with micro data centers positioned at antenna sites, to data centers serving metropolitan areas, all the way to the mega data centers providing global content processing and distribution. Ciena is interested in developing and demonstrating a new Data Center network, i.e. a collection of switches that provide data interconnection amongst servers. More specifically, Ciena are interested in working with an Ontario based SME to prototype and demonstrate three aspects: a) flat-mesh network switch interconnection topology b) resilient and efficient routing and forwarding algorithms c) possible inclusion of optical switching to improve cost and performance
Project Partner	Ciena Canada
Timeline	2 years
Available funding	Up to $500,000 CDN
Applicant Type	Ontario based SME Scale company
Location	Ontario
Project Details	Current data center networks are built with rigidly structured architectures. Some examples known in the data center industry are Fat Tree (Clos), Dragonfly, Slim Fly, and B-Cube. Specifically, a Fat Tree or Clos network is frequently used in modern data centers. The network is easy to visualize, can be built from smaller building blocks and provides high bisectional bandwidth. Further, the number of hops is at most 4, traffic is easily routed with equal cost multi-path (ECMP) and is resilient to failure, and the network can be scaled further at the cost of additional layers. At the same time, structured networks suffer from some well-known problems, enumerated as follows:

	1. Increased latency due to many hops, especially as the number of layers grows. Low throughput also results in filled switch buffers, increasing latency. 2. Discrete implementation sizes, and higher layer ports are unused. 3. Difficulty in horizontal scaling—requires multiple layers. 4. Structured (Clos and other) networks are susceptible to cluster-packing problems, i.e. CPU, Storage, etc. resources in the cluster must then be sized to anticipate largest load, and may be underutilized otherwise Our target network architecture STructured Re-Arranged Topology (STRAT) changes the current industry approach (multi-tier Clos) such that ALL switches are actually Top of Rack switches (TORs) and have servers attached. Unlike existing products, there are NO switches which serve a purely network interconnect function, i.e. there is no Aggregation or Spine layer. Each TOR has a number of ports which face towards other TORs in the network. The optical interconnection between TORs is established in a uniform random, but non-repeating way. These connections are implemented via passive patch panels, are completely static, and do NOT need any exotic optical switching. Such networks are known to have a number of excellent properties: 1. Uses only commercial off-shelf components (Switch ASIC, Optical pluggables, etc.), fully packetized 2. Zero new hardware technologies, either electrical or optical (though some could be introduced for further benefit) 3. Use of relatively low port count switch ASIC is sufficient – do NOT need huge switches 4. Completely de-layers network 5. Higher throughput and very few hops (5 or lower) with high diverse path count (simulation details in a later section) 6. Alternately, can achieve fewer optical links and switches, i.e reduced hardware at fixed throughput (simulation details in a later section) 7. Smooth, continuous trade-off of hardware cost and throughput (no large discrete increments) 8. Easily Scalable: New server and switch insertion is simple, even with different port counts and technologies 9. Day 1: pay only for network servers use; no underutilized centralized fabric 10. Ports can be aggregated into a few large bundles, improving optical Transceiver cost and front-panel density – typically 16 distinct port bundles are sufficient. (In contrast, Clos scalability requires very large port fan-out to realize large network E-W scalability) 11. Substantially increase resilience to failure through richer interconnect and through workload spread across the full data center

Project Goals/ Outcomes	The project's primary goal is to develop the required architectural refinements and hardware and software infrastructure appropriate for successful development of a higher performance data center topology. The project will build, test and evaluate a prototype STRAT network using white-box off-the-shelf switch hardware (i.e. switches from Broadcom, Marvell, Barefoot, etc.), including sufficient routing/forwarding protocol stack modifications for basic functionality. In order for STRAT to be considered a credible approach for the "Inside" DC space, Ciena needs to develop a much deeper understanding of the performance and interactions of data center architectures, protocols, traffic patterns, etc. These technical capabilities will also allow STRAT to gain some mindshare with the DC industry, which is becoming much more important to overall communications market. By building a small-scale replica white-box switch-based enterprise network that can be programmed to support realistic CLOS and STRAT topologies, running on state-of-the-art merchant switch/router silicon and an industry standard network stack. This approach provides an ideal vehicle for demonstrating moderate sized enterprise data center implementations that are feasible, flexible and beneficial. Initial project proof point may be shown in a packet simulator (OMNET++ or similar) or in a network emulator (MININET) to demonstrate routing and forwarding protocol operation and network performance benefits.
Applicant Capabilities	• Ontario based SME with the capabilities to deliver the project outcomes noted above SME staff should have working knowledge of internal operation of networking protocol stacks such as BGP, EIGRP, RIP. • SME staff should have working experience with white box switch hardware (from Accton/EdgeCore, Alpha Networks, etc.) • SME staff should have working experience with Linux (such as ONL), ONIE, and Linux-based host software tools, SDKs and APIs for developing packet switching and routing functionality on modern merchant silicon-based router/switch ASICs (e.g. Tomahawk 3) • SME staff should have working experience with a set of programming languages that may be used in the project (C/C++, Python, etc.) • Familiarity with PISA switch concepts and P4 programming may be beneficial
Additional Information	• Ciena may provide white box switching hardware for the project

Launched in 2018, the ENCQOR 5G SME Technology Development Program partners Ontario based SMEs with ENCQOR 5G Anchor Firms on 5G technology development projects. Areas of research interest are defined by Challenge Statements submitted to OCE by the ENCQOR 5G Anchor Firms and posted to the OCE website on a rolling basis.

If you are interested in developing an expression of interest, please visit the program guidelines for information on next steps.

For any questions about new Challenge Statements or the ENCQOR 5G SME Technology Development Program please contact Jennifer Moles at Jennifer.Moles@oce-ontario.org.