# Real-time, cloud-based RAN awareness

**A conversation with John St. Amand, CEO Movik Networks**

**By Frank Rayal**

**Frank Rayal.** Hello, and welcome to this conversation with John St. Amand from **Movik Networks**. I'm Frank Rayal. This conversation is part of a Senza Fili report on deployments of HetNets and the latest solutions from the wireless ecosystem that will make HetNets a reality.

Today we're speaking with John St. Amand, CEO of Movik Networks. Movik is a leader in intelligent RAN-awareness solutions.

John, I would like to start off by asking you to give us some background on Movik.

**John St. Amand.** Movik is the first company to look across all the domains of the mobile network – voice, data, SMS and MMS – and correlate, in real time, content-level awareness with radio access network awareness to drive a much higher level of quality of experience for every mobile user in the network at any single instance.

**Frank Rayal.** How do you do that? Can you go into more detail about the product and describe the features and the functionality of this solution?

**John St. Amand.** Movik's capabilities are derived from the way we dig deep into the radio access network user and control planes. Movik's software ingests the entire user and control planes of the mobile network, which can be 3G, 4G or Wi-Fi. This enables us to look at exactly what the users are doing and are experiencing. We can determine whether they have a high or a low quality of experience, and determine the interdependencies between a certain user and other users in a particular macro-cell, small-cell or Wi-Fi network. Based on that, we make real-time decisions to drive a higher level of quality of experience for a particular application. It's no longer good enough just to say a user may get a 1 mbps download speed; instead it's really how their video performs. Does his video player or the network stall? Does he drop calls because of heavy data usage on his device or in that sector?

What we do is in real-time. We have a lot of sophisticated algorithms that determine the quality of experience per user and what is driving that quality of experience. Is it network-level errors? Is it competition with other users? Is it a particular content type? We use that information to drive policy and traffic management functions into the network, whether it is into the radio access network elements for network selection across a heterogeneous network, whether it's policy information to drive rate limiting, or whether it's optimization decisions around content delivery networks or through core network elements to do video optimization, for example.

**Frank Rayal.** What do you mean by real-time?

**John St. Amand.** Real-time in the sense that we look at information in one-second bins. A lot of competitive solutions which are RAN-based generally operate in near-real-time or 15-minute bins. That's not good enough, because users are experiencing, in real-time, degradation in the

**Sponsored by**

movik

quality of experience. For example, when downloading a three-minute, high-definition video – which could be a very large file that would instantaneously exhaust a sector's capacity – a 15-minute bin is going to miss those perturbations in the network where the quality of experience drops very quickly. So it has to be in real-time, and the only way to do that is to ingest the entire user plane, which is all the bits and bytes going across the network from every single device, aggregated across the entire macro-cell, small-cell and Wi-Fi network, and the entire control plane, because we're looking at session establishments. So that's real-time.

**Frank Rayal.** You mentioned that the solution operates on both the control and user planes. That's a lot of data that will be flowing through the system. How do you scale to handle all that amount of data?

**John St. Amand.** Movik's solution, which is named cREACH, which stands for Cloud REACH, is a very scalable, cloud-architecture solution. It'll sit on any hardware platform or cloud infrastructure that an operator has deployed. Most operators are very hesitant to deploy a new hardware platform, which is why Movik, from the beginning, has been software-solution focused.

We can scale from single radio-network controller locations to ingest an entire highly distributed LTE network through our cloud deployment model, which is basically a multinode implementation of our software. This is very important, as you want to track users' quality of experience as they move across the network. We've got to keep track of things like mobility, because if a user is highly mobile, you do not want to put that user on a small cell, for example, because his data session may drop as he goes between cells. You want to keep him on a particular network so he has that high level of quality of experience and you don't make incorrect network selection decisions. So it's very scalable in that sense. If you deploy into a single-chassis or blade server architecture, we scale up to 40 Gbps per instance, but that instance then scales across an operator's cloud infrastructure or virtualized network of computing resources.

**Frank Rayal.** Does it present a single point of failure?

**John St. Amand.** No, it does not represent a single point of failure, because if we lose a particular second, three seconds or five seconds of data, the algorithms are tuned in real-time, where losing potentially a few seconds worth of data on a particular device is not going to affect how we score a sector or an eNodeB. The system is cloud-based, so if you lose a blade in a server architecture or a server in a cloud-based architecture, that traffic automatically gets load-balanced off to another set of servers.

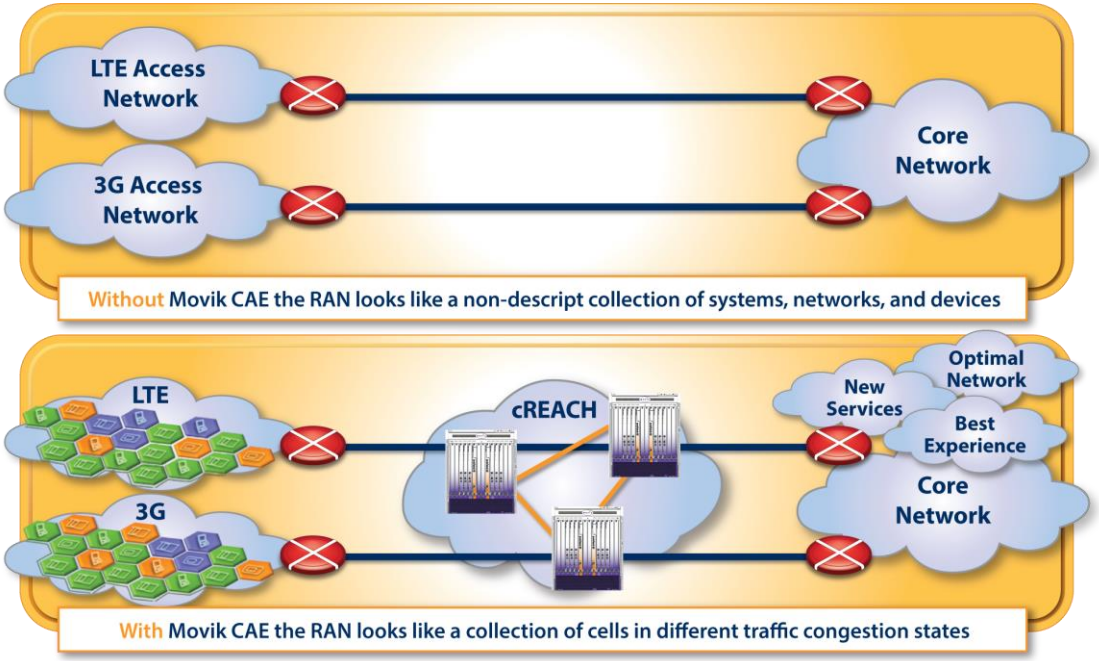


**Figure 1. Movik's Content Awareness Engine imparts intelligence to the RAN. Source: Movik Networks**

**Frank Rayal.** How does your solution compare with other solutions in this space? For example, how does it differentiate from DPI engines, which can also look into the traffic, make decisions, and interface with other network elements?

**John St. Amand.** DPI engines do have some of this information, but because of where they sit in the core network, they do not have the radio access network awareness, so they do not see sector-level or mobility-level information. They will see the full data plane, but they will not necessarily see it tied to a device and to a particular sector, and then be able to take that information and aggregate and correlate it across all the other users in a particular sector or eNodeB. That information is lost. They do not sit on the control plane, as well, because control plane interfaces are distributed much closer to the access network. Without the control plane and without sector- and mobility-level information, they only see people's behavior – for example, what they are browsing and how quickly they are getting their data, which is interesting and important, but it's not the full picture.

**Frank Rayal.** Can we zoom in on some of the benefits that your solution would provide in a HetNet context? How does it complement or enhance the HetNet?

**John St. Amand.** It's very complementary to heterogeneous network solutions that are out there today. So SON, in the traditional sense, was architected for autonomous radio access network configuration, interference mitigation, neighbor relationships, and things like that, which are important, and Movik does not necessarily do. SON allows for the configuration of the network, but as the network adapts to the demand being placed upon it by devices, the demand may outpace how the network is configured.

Movik complements SON to add content-level awareness and real-time quality of experience to dynamically decide how the demand is placed across the various networks: small cell, macro cell, and Wi-Fi. We do that based on which content should be where. We dynamically account in real-time for the load that's being placed upon the radio access network elements, such that the user's quality of experience is preserved. The traffic load is going to contract and expand as people come into and out of sectors; we tie that to content-level awareness.

**Frank Rayal.** Is this a complementary function to the radio resource management that's available at the base station?

**John St. Amand.** Extremely complementary, and it's something that resource management solutions do not do right now. They are historically focused on signal-to-noise ratio, interference management, neighbor configuration, and things like that. They have not focused on the content-level awareness to see the type of carried traffic.

Content-level awareness allows you to look into data sessions to see how the network should behave to continue to deliver the quality of experience the user would like. It's no longer "Can you hear me now?" which was very popular for a long period of time. Everybody's got good coverage, everybody has the same devices, so now it's really differentiating on "How well do my applications perform on that 4G network, or that iPhone or Android phone?" and that's how the operators are going to differentiate themselves. It's going to be the carriers who have the most intelligent radio access network to deliver what people are demanding, most efficiently, that are going to be the ones rated the highest.

**Frank Rayal.** John, what are some of the gains achieved with your solution? Can you provide some quantitative numbers on KPIs?

**John St. Amand.** Our applications and use cases are all centered on quality of experience improvements and network utilization increase. Among the user experience improvements we have made are 35% to 40% faster page download time, and 35% to 40% faster video start time.

That's very important, because the user's abandonment threshold, which is how long do you wait before your Amazon, eBay, or CNN page displays, is three to four seconds before people get bored and move to a different site. This is very important to ecommerce sites and to content delivery networks. A 35% improvement in download time is a big number.

Another big number is a 40% reduction in overall bandwidth consumption. Operators design their network for peak usage, and to design for peak continues to be very expensive. As you exhaust

spectrum, you tend to increase the number of base stations to increase the capacity of the access network. It's very difficult to balance cost and revenue.

Movik smooths over the peaks in demand placed on the network, such that you don't have to design a network for high peaks. Once you've smoothed out those spikes and you've delivered very low page-download times and video start times, you've now freed up resources on base stations so they can serve more subscribers.

For example, if you get a heavy user on the edge of a particular macro cell, it's going to use a lot of that base station's resources. If we reduce the amount of resources that he's using, because we've modified his consumption to some extent, we can have more subscribers in that particular geography, on the order of 10% more.

That basically frees up capacity in terms of subscriber count, which translates to lower capex and spectrum cost, which is very important for operators. Those are the big numbers that we hit on: 10% increase in subscribers, 35% improvement in page download times, video start times, time to first byte, and 40% reduction in overall bandwidth consumption.

**Frank Rayal.** Everybody talks about the user traffic, but there's also a huge amount of signaling traffic that goes along with that. Does your solution do anything on that side?

**John St. Amand.** Absolutely. There are different levels of signaling – for example, network-level signaling such as TCP, and call-control or data-session signaling. There's also application-level signaling. Take, for example, fast dormancy: when you put your smartphone down, it goes to sleep and releases its assigned radio access bearers, and then as you pick it up, it reestablishes everything. But when background applications wake up and want to pick up their locations – so they can find out whether they should tell you what the weather is or whatever happens to be going on – that's imposing a much higher signaling load on a network than people would have ever predicted.

LTE in general has a high level of signaling itself. We've known some operators who have had some challenges with signaling storms due to LTE.

What Movik does is, we see what type of application usage is going on and we flag high-signaling resource utilization. We also consolidate that in real-time (ignoring some of the signaling that really doesn't matter to the functionality of applications) so there is less signal flowing through the core of the network. In effect, we do a lot of proxying – for example, we'll proxy TCP, radio access–bearer requests, and HTTP requests. Overall, that quiets the network down considerably.

**Frank Rayal.** You mentioned your solution is multimode: it operates with LTE, 3G or Wi-Fi. Does it provide more benefit in one of those technologies over another, or is it basically the same benefit across all technologies? How does it compare where it works with these different technologies?

**John St. Amand.** Really, all that changes is the signaling interfaces across HSPA, CDMA, LTE or Wi-Fi. The bearer is very similar, but how we get the bearer is different per network. You have a different interface depending whether the network is LTE or HSPA, but the bearer itself is TCP, HTTP or HTTPS.

The signaling piece is very different in LTE versus HSPA versus CDMA. That's where a lot of the heavy workload is.

In LTE, for example, you have to sit on many interfaces to get the full picture of what is actually going on in the operator's network. For example, you could be in the same city as I and we both have LTE service, but we could be connected to different MMEs. So you have to sit on multiple interfaces in the operator's LTE network to put together that full picture of where the data and voice paths are going and the associated control and user planes.

So it's a very highly distributed network that dictates cloud architecture. If you do not have a cloud solution in LTE, you're going to have a very difficult time stitching together the exact story that's going on in real-time for an operator's network. It's a little bit easier in 3G, but like I said, the user plane is very similar across the different networks.

**Frank Rayal.** Are the interfaces you mentioned standardized or proprietary?

**John St. Amand.** Not proprietary. That also differentiates us substantially from the traditional vendors or the purest SON vendors. Their solutions will generally pick up per-call measurement data and different management data that exists in the network that are specific to a particular vendor's radio network. We sit on 3GPP standard interfaces, which in HSPA is the Iu interface, and in LTE are standard interfaces like the S1U, and the S11 MME interface.

**Frank Rayal.** What is your vision and roadmap for your intelligent RAN awareness platform?

**John St. Amand.** The future, really, for us is to start driving a lot of the network-selection activities. That involves us being complementary to the SON network elements and the access-network selection-function elements. That's where we've started to migrate now. As you migrate into the monetization strategies with content-delivery networks and origin servers, you're driving how an operator monetizes user behavior by matching user's dynamic needs and interests with that of content publishers like eBay, Amazon, CNN and ESPN. Putting that together in real-time is a big project and is often encompassed under the big-data umbrella. Movik positioned itself as a big-data vendor with uniqueness in the aspect of its real-time RAN awareness and content level–awareness correlated data.

**Frank Rayal.** You define your architecture as being cREACH: Cloud, Report, Export, Act, Control, HetNet. Explain what these elements mean.

**John St. Amand.** cREACH is the way that we approach our operator customers to help them first understand the dynamic challenges that are going on inside the network, which includes user behavior and QoE provided by the network. We then take them through a progression of steps of how Movik is not just a reporting or an analytics engine. We predict in real-time that a user will have challenges, or that he'll want a certain type of information, whether it's in the form of promotion, advertising or content. We'll drive the actions that actually allow improvement in QoE.

So the R stands for "reporting and analytics." We generally start our engagements by offering network operators a real-time, all-encompassing view of what is going on in the network, down to the sector level. We answer questions like: How are my pico cells performing? Have they taken load off of my macro network? How's the user experience? Is it improving or declining? Where and why is it declining? What are people doing on my network? So we'll give operators that fully correlated view in real-time.

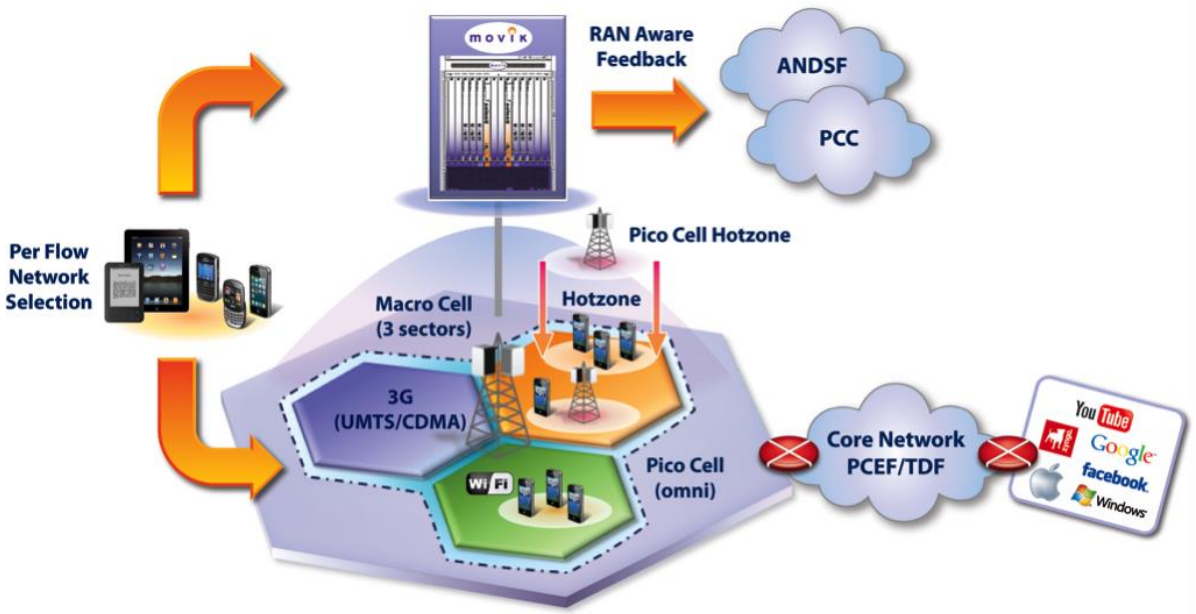


**Figure 2. Movik enables unified, real-time control over any RAN while providing subscribers optimal connection anytime, anywhere. Source: Movik Networks**

From that we move to the next step: What can you do to help me solve this problem? That's when Movik gets into the E of the cREACH strategy, which is "export."

We export our intelligence in real-time to the policy and optimization domains, as well as, perhaps, to their marketing domain, where they may want to make marketing decisions as far as partnering with their content-delivery networks. All that happens in real-time. We have projects in deployment now where we export intelligence to the operator's policy domain or optimization engines, where they may want to do video optimization, rate limiting or admission control – they "act," which is the A in cREACH.

The C in the equation is "control." Movik actually does the control, which, for example, can be moving a stationary user from a macro cell to a small cell or Wi-Fi because he's downloading heavy video traffic. We'll drive that control to the network-selection function and proxy this information based on our knowledge of the user's activities and QoE.

Then the H in the strategy is how we bring this across all the components in the "heterogeneous network." And that is really what makes up the cREACH strategy for Movik.

**Frank Rayal.** Your solution is about intelligence. This is something that is driven by the networks becoming all-IP. What's your take on the market dynamics? How quickly are operators adopting these kinds of solutions to reduce traffic congestion on the network?

**John St. Amand.** They're moving very quickly because it's software based. We're not introducing another network element in terms of a box on the network. The deployment is quick, and we make the elements that are already deployed work better, which is important. We do not necessarily do the optimization ourselves by sitting inline, which is a bit challenging.

We will make their core network elements do their job better, because they now have access to information that allows them to understand what is going on. We enable bridging the RAN domain and the core domain, which never talked to each other before.

The RAN domain is very focused on coverage, interference and how you get access to the radio network; the core domain is focused on IP and on delivering packets. Neither really has any regard for the other. Movik is now in the middle to drive a certain behavior in the RAN. This allows us to deploy an operator's network quickly, to give them a real-time correlated view within days.

**Frank Rayal.** John, I would like to thank you for joining us today in this candid and informative conversation.

I would also like to thank the viewers for watching this conversation with John St. Amand, CEO of Movik Networks. This conversation is part of the Senza Fili report on HetNet deployments and the latest solutions from the wireless ecosystem that will make HetNets a reality.

The report can be downloaded from the Senza Fili website at **www.senzafiliconsulting.com**.

## Acronyms

| | |
|---|---|
| **3G** | Third generation |
| **3GPP** | Third Generation Partnership Project |
| **4G** | Fourth generation |
| **ANDSF** | Access network discovery and selection function |
| **CAE** | Content Awareness Engine |
| **CDMA** | Code Division Multiple Access |
| **cREACH** | Cloud, Report, Export, Act, Control, HetNet |
| **DPI** | Deep packet inspection |
| **eNodeB** | Evolved NodeB |
| **HetNet** | Heterogeneous network |
| **HSPA** | High Speed Packet Access |
| **HTTP** | Hypertext Transfer Protocol |
| **HTTPS** | HTTP Secure |
| **IP** | Internet Protocol |
| **Iu** | The interface between radio network controllers and core network nodes in 3G networks |
| **KPI** | Key performance indicator |
| **LTE** | Long term evolution |
| **MME** | Mobility management entity |
| **MMS** | Multimedia messaging service |
| **PCC** | Policy and charging control |
| **PCEF** | Policy charging enforcement function |
| **QoE** | Quality of experience |
| **RAN** | Radio access network |
| **RNC** | Radio network controller |
| **S11** | The interface between the MME and serving gateway in LTE networks |
| **S1U** | The interface between the base station and serving gateway in LTE networks |
| **SMS** | Short message service |
| **SON** | Self-organizing network |
| **TCP** | Transmission Control Protocol |
| **TDF** | Traffic detection function |
| **UMTS** | Universal Mobile Telecommunication System |

# About Senza Fili

Senza Fili provides advisory support on wireless data technologies and services. At Senza Fili we have in-depth expertise in financial modelling, market forecasts and research, white paper preparation, business plan support, RFP preparation and management, due diligence, and training. Our client base is international and spans the entire value chain: clients include wireline, fixed wireless and mobile operators, enterprises and other vertical players, vendors, system integrators, investors, regulators, and industry associations.

We provide a bridge between technologies and services, helping our clients assess established and emerging technologies, leverage these technologies to support new or existing services, and build solid, profitable business models. Independent advice, a strong quantitative orientation, and an international perspective are the hallmarks of our work. For additional information, visit www.senzafiliconsulting.com or contact us at info@senzafiliconsulting.com or +1 425 657 4991.

# About the interviewer

Frank Rayal is founding partner at Xona Partners, a boutique management and technology advisory firm specializing in telecom, media and technology. He is a telecom industry professional with more than 20 years of experience working with network operators and system vendors to develop and deploy innovative wireless solutions. Frank co-founded small-cell backhaul pioneer BLiNQ Networks. He held senior product management, marketing and business development positions at Ericsson, Redline, and Metawave. He holds a BS in electrical engineering from Case Western Reserve University, Cleveland, Ohio, and an MASc in electrical engineering and an MBA from the University of Toronto, Canada. He is a senior member of IEEE, and a member of Professional Engineers Ontario.