

Optimizing backhaul capacity

A conversation with
Richard Lord, CTO &

Mats Vilander, VP EMEA Sales
& Head of Global DatE,
AltoBridge



By Frank Rayal
Senza Fili Consulting

Frank Rayal. Today we are joined by [AltoBridge](#), and we have on the call with us Richard Lord, CTO of AltoBridge, and Mats Vilander, VP of EMEA Sales and Head of Global DatE (Data-at-the-Edge) product. This conversation is part of Senza Fili's report on small-cell backhaul that gives an update on small-cell backhaul solutions and on the evolution of the mobile operator requirements for small-cell backhaul.

AltoBridge is a leading developer of wireless network solutions that reduce network operating costs in delivering mobile voice, mobile broadband and big data services. The company developed a series of patented and patent-pending intelligent mobile broadband data optimization and delivery technologies, across wireless networks, that reduce backhaul costs. AltoBridge is a privately held company headquartered in Ireland, with offices in the US, China, Malaysia and Indonesia.

Welcome, Richard and Mats.

Let me start by asking you to give us an introduction about your solution – and that would be the Data-at-the-Edge product – and how it works in the small-cell environment.

Mats Vilander. Thank you very much for that. Data-at-the-Edge, first of all, is a backhaul optimization product. It operates by saving backhaul cost in one environment, where it is driven out of satellite for rural backhaul, and secondly, in 3G and LTE congested networks, where you have a massive amount of content that is being distributed to the end users, and the down time is one of the most driving factors that prevents the end user from getting his content in time. So it is driven out of quality experience. We are also

the only company that supports the local call interception, the LCI, which is already an existing 3GPP standard allowing the operators to overview the traffic over the luh and lub interface. We are the only backhaul optimization company for byte caching that is supporting the lub interface.

Frank Rayal. I take it that your product is a software solution – do you port it onto different vendors' hardware? Can you elaborate on exactly where the product is located in the mobile network?

Richard Lord. The solution is based on byte caching, which is a symmetrical caching technology – in other words, the caching engine required at both sides of the link to be optimized. In our solution we are optimizing the link between the remote base station and the core network.

As Mats indicated, we support all the major 3GPP cellular interfaces, so we are one of the few companies that have ported this solution onto an lub interface. We also support the luh interface, which is for small cells, and the S1 interface, which is for LTE small cells and macro cells.

Sponsored by

ALTOBRIDGE™

The solution is composed of two modules: one is called the central gateway and the other is called the remote gateway. The central gateway sits near the core network, and the remote gateway sits next to the base station. Basically, it is what people might refer to as a bookend solution, so you have a monitor on either side of the backhaul link that you are trying to optimize.

Frank Rayal. Backhaul is a major issue that is gating small-cell deployment. How does your product address the backhaul challenge, and how does it enable operators to deploy small cells?

Richard Lord. Backhaul is always, or very often, an issue with small cells. Small cells can be backhauled over a number of different media. In an urban environment, you will typically find DSL lines and you will find PTP or PTMP wireless backhaul solutions, microwave solutions, licensed and unlicensed spectrum. In remote areas, you will find small cells that are backhauled over DSL, over leased lines and also over satellite.

We came from a satellite background, therefore we have very restricted and a very expensive bandwidth, and that is where we learned our trade.

We have taken the same technology and adapted it to address the other types of backhaul – the microwave and DSL that you find in the urban environment. It allows the operator to effectively reduce their backhaul bandwidth requirements by approximately 50%.

You can look at that in a number of ways. You can say, “Well, I can get double the capacity for the same price.” Or “I can pay less and cut my cost, reduce my cost for the same capacity, while at the same time reducing the congestion or the amount of time that my backhaul is congested, therefore giving the users a much better experience when they are watching videos or downloading internet content.”

Mats Vilander. I would like to add is that congestion due to increase in data traffic is driven by companies such as Google, YouTube, Akamai, ChinaCache and

Netflix, just to mention a few. They provide mobile applications over different wireless networks so we have a convergence scenario, where you have 3G, LTE and Wi-Fi all combined in this so-called HetNet environment in urban areas. The most important part is that we are the only one that can supply shared caching.

Frank Rayal. So you are saying that it applies to 3G and LTE, and it is the same solution that will be deployed and will be used by both networks at the same time.

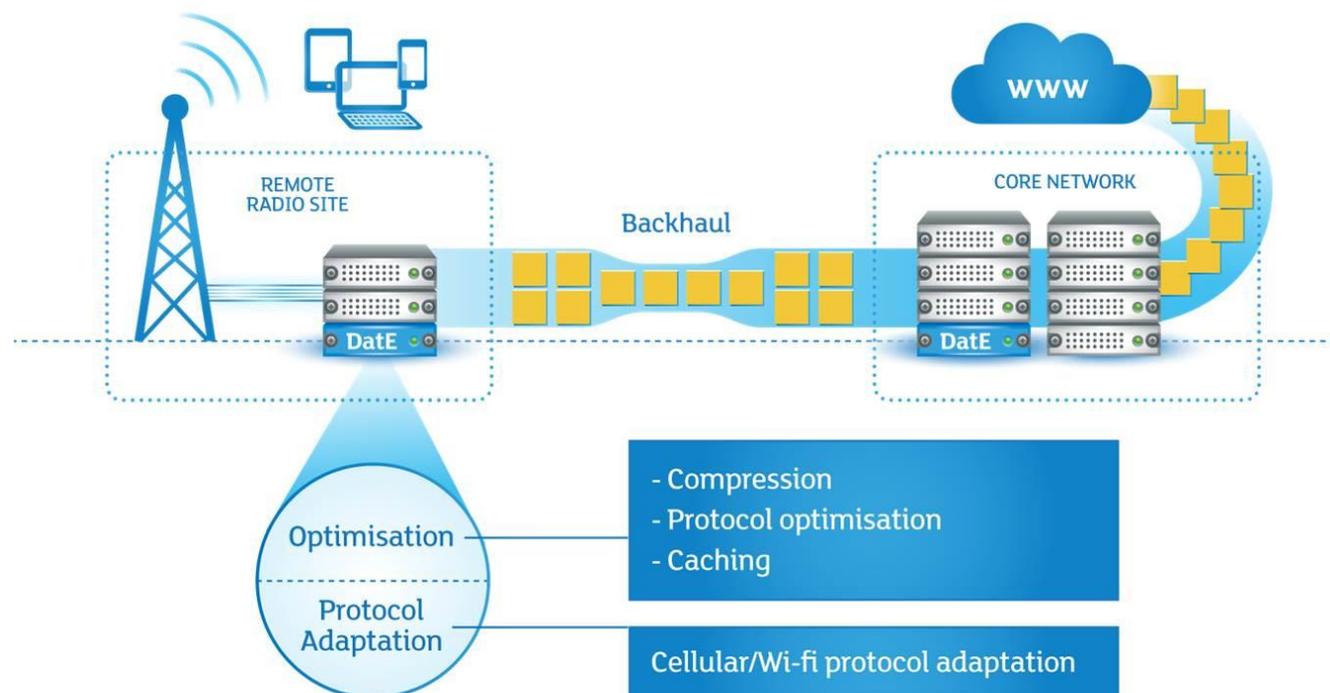


Figure 1 Data-at-the-Edge (DatE) backhaul optimization solution. Source: Altbriidge

Mats Vilander. We can currently do shared caching or mesh caching 3G and LTE and Wi-Fi, separately.

Frank Rayal. Richard and Mats, I have a question about the advantages of this technology regarding the business case of the operator and some of its features. You have already mentioned some of those, but can you expand on that in the small-cell context?

Mats Vilander. Sure, let's just start by adding on that in the business case for satellite backhaul, typically the current savings are 50% of the backhaul costs – that is one business case which is very specific. By saving 50% of the backhaul cost, from our experience, our solution is paid back within twelve months.

In a small cell urban environment, the cost savings are driven very much by two parameters: one is the cost for the backhaul, i.e., the leased lines, and the second is the entire capacity in traffic. By increasing traffic capacity, the backhaul cost savings will of course be realized.

The other very important aspect, moving in from the satellite transmission market to the urban market, is very much driven, as well, by the quality of experience. For operators, backhaul congestion is a big problem currently. And that is reflected in what was in the recent study of one of the transmission companies globally, announcing that over \$9 billion is needed to cope with increased traffic of 3G and LTE.

Many of the operators are in situations where they have to upgrade their network and double its capacity to cope with increased traffic. The business case again

lies with byte caching – transparent byte caching – that can save up to 70% in backhaul costs.

What we see in urban scenarios is that byte caching through pre-positioning of content optimizes bandwidth. At the data edge, there is an algorithm that can recognize traffic patterns by end user, so it can store the data of the most popular content close to the user's node. Thus the content will not be pushed back to its core network, where it would, of course, jam up the transmission.

Frank Rayal. I would like to ask you if there is intelligence in terms of distinguishing the different types of traffic.

Richard Lord. Absolutely, yes. We use deep packet inspection effectively to identify the type of traffic and decide, firstly, whether or not this traffic should be cached. We try and cache as much as we can.

Voice traffic, compressed traffic, or encrypted traffic are the types of traffic that do not lend themselves to caching– and therefore we don't cache it. Caching is most effective with streaming video, for example, or any form of large pictures, like JPEG, GIF or PNG. This is where caching is at its most effective, and so we focus on those protocols and those data types for caching.

Frank Rayal. I would like to also ask you about the competitive environment, in that you have a product that is fairly unique, but are there other ways of approaching this problem? Can you highlight some of the different approaches and how your solution would compare to them?

Richard Lord. I think that, from sort of a competitive point of view, there are two main areas, really. One is what we have always done in the past, is that we have always respected the philosophy that says the more intelligence you push out to the edge of the network, the more efficient it is from a backhaul perspective. I think that we were the first guys to put that sort of optimization intelligence and caching right out at the base station, and we have done that for all types of base stations: classic 3G, small-cell architecture, LTE architecture. We are the only company who has implemented successfully the lube caching, which is something that is very complex from the technology perspective.

Then the second area, from a competitive point of view, is that rather than implement object cache or URL cache, we implement transparent byte caching, as Mats alluded to earlier, and byte caching has several unique features or aspects. One is that it is more effective as a caching technology than traditional URL caching or object caching.

That is because it works at a lower level: it is able to ignore a lot of the do-not-cache tags and indicators that you find on many web pages with dynamic content. So it gets around that dynamic content aspect by downloading the adapter – only the parts of the web page that have changed since the previous download.

It is also more effective because it is very good at caching partial content. This is a particular problem within applications like YouTube, for example, because in many cases people only download partial YouTube

video. For example, if you download the first 30 seconds of a video, and I come along afterwards and download the first minute of that video, then the first 30 seconds comes out of the cache and the second 30-seconds interval comes across the network – so it is a very effective caching solution.

Because the byte caching is a symmetrical solution that is implemented on both sides of the backhaul and it is only focusing on optimizing that backhaul, it is completely transparent of some of the key features in the core network. Those key features in the core network include billing, legal intercept, content filtering, and service management. Our solution is completely transparent to those key features for the operator, which I think causes a lot of our competitors headaches, trying to meet requirements such as legal intercept and content filtering.

Frank Rayal. Latency is an important parameter on the backhaul side. How does caching impact backhaul latency?

Richard Lord. Clearly as soon as you start manipulating traffic on the backhaul link, you are going to introduce latency; even if you are doing something as simple as DPI, you introduce latency. Caching for sure introduces some latency.

Now, on the other hand, as soon as your link becomes congested – and this is what is happening all over the world, in small cell and in the macro cells – then, caching helps reduce that latency because you are not waiting for your content. Your content is effectively being delivered directly from the cache at the remote

site to the base station, and then onto the user’s device. Many parameters can affect the latency of the solution, but it is really a tradeoff. At the end of the day, if your backhaul is congested, then for sure the caching is going to reduce the latency on average, particularly during the busy times. And if your backhaul isn’t congested, then you don’t really need much caching, so you know the latency introduced isn’t really relevant.

Frank Rayal. That is a very good point, Richard. Another question that I have is about the future roadmap, specifically talking about the DatE product (Data-at-the-Edge). Can you elaborate on where this product is now? And what are the unique features and capabilities that you plan to implement, and their timeline?

Richard Lord. Well, I think that the key thing, we have mentioned some numbers during the call: we have seen averages of 50% savings on the backhaul in certain peak times; as Mats said, we have seen up to 70% savings. Most of that is based on caching, and so there are two factors, I guess, which affect the effectiveness of your caching solution.

The first is how good are you at recognizing duplicate content, and we have chosen this byte caching technology, which is very good at recognizing duplicate content.

The second part of the question is: how much commonality is there across the content that users are looking at on a particular cell site? That obviously depends on the community, and then, depends on the

size of the community, and the type of the content that they look at.

In order to improve the hit rate on the cache, which is what we are always striving to do as we move forward, we introduce more and more intelligence into the system, to allow the system to learn what the usage patterns are, what the popular data is and what type of content people are looking at, and what time of the day they are looking at it. We use that analytic-based information to pre-position the data and the remote caches.

You can take some very simple examples. You might discover in the morning that some particular news story is popular in New York. Well, you can be pretty sure that three hours later in California that news story is going to be popular. So you can take the initiative to push that content out to the remote caches out in California early in the morning while the system is still relatively lightly loaded.

There are many other examples of that, right down to the point where you can almost have a personal cache that follows a user around the network as he moves from one base station to the other.

There is really no limit to how intelligent you can get in terms of determining what content people are going to be looking at. All the more data that you can pre-position at the caches, the better the hit rate and the more effective the solution is in reducing the bandwidth required and delivering better-quality experience.

That is the main thrust of our development over the next couple of years. We already have a basic implementation of this pre-positioning, and it is something that we will develop over the next couple of years.

We are also fortunate to manage a number of the networks that we actually sell. We don't just sell boxes, but we are in a very good relationship with operators, and we help them manage their networks. This allows us to get to the firsthand feedback of what the problems they are facing are, and what the usage patterns are like. That drives our roadmap going forward.

Frank Rayal. And I take it that this roadmap would apply to both the 3G and the LTE technology, and even to Wi-Fi, is that correct?

Mats Vilander. All three of them.

Richard Lord. Absolutely. There is minimal difference to us between 3G and LTE. Obviously, the data volume in LTE is potentially much larger, but from a technology perspective, it is largely transparent to us. We have implemented a layered solution that can be shared across various access technologies

Frank Rayal. Well, Mats and Richard, is there anything else that you would like to add, on the company or on the solution?

Mats Vilander. I think we all – including Richard and myself – have been in this telecom industry for more than 20 years, in various roles and various RAN

suppliers. I think that the biggest challenge right now for mobile operators— because in the end, it is an operator business case – is how to fund the additional investments needed to cope with the increased traffic.

We see an enormous number of new players that come in with entirely new content, but they are all struggling with the fact that many of the 3G networks are congested. And the downtime, if we compare with the fixed broadband services, is the biggest challenge.

Caching by itself, as we foresee it now, will be in every network globally: how do you optimize for more traffic on every node in the coming future, and how do you pre-store popular content close by its end users? Because we are talking about the fixed-mobile convergence that we have been talking about for fifteen years.

This will definitely happen in the 4G world. But it requires the caching functionality.

Frank Rayal. Now, you did mention 50% cost savings and twelve-month payback period. I am curious to understand: how much on average does your product save in terms of backhaul capacity?

Richard Lord. That is a key parameter, because that affects the whole business case to introducing a solution like this. And when we started out this whole sort of program a couple of years ago, we gave ourselves a target of saving 40% – in other words, taking 40% of the bandwidth requirement away from the backhaul – and what we are seeing with real deployments in the field is that we are doing better

than that. We are taking 50% of the backhaul bandwidth out, and we are seeing during certain periods, particularly the busier periods, that number go up to 70%, and that is when it is important.

What we are working on is basically technology that will improve the hit rate of the cache and that will allow us to go beyond 70%. In other words, we would be able to take out more than 70% of the bandwidth requirement from the operator's perspective.

Frank Rayal. Richard, thanks a lot. It is very significant to be able to provide 70% or even over that in the future in capacity savings over the backhaul.

Well, Mats and Richard, thank you very much for the conversation today.

Richard Lord. Thank you.

Mats Vilander. Thank you.

Frank Rayal. This conversation is part of a Senza Fili report on small-cell backhaul that provides an overview of small-cell backhaul solutions along with an in-depth conversation like this from leading vendors who opted to participate in the report. The report can be downloaded from the Senza Fili website at www.senza-fili.com. Thank you for listening in, and, Richard and Mats, thanks again for this conversation.

Acronyms

3G	Third generation
4G	Fourth generation
3GPP	Third Generation Partnership Project
CTO	Chief technology officer
DatE	Data-at-the-Edge
DPI	Deep packet inspection
DSL	Digital subscriber line
HetNet	Heterogeneous network
Iub	The interface between NodeB (3G base station) and Radio Network Controller
Iuh	The interface between Home NodeB and the Home NodeB gateway
LCI	Local call interception
LTE	Long term evolution
PTMP	Point to multipoint
PTP	Point to point
RAN	Radio access network
S1	LTE interface between an eNodeB (LTE base station) and a mobility management entity or a serving gateway
URL	Uniform resource locator
JPEG	Joint photographic experts group
GIF	Graphics interchange format
PNG	Portable network graphics

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 author



Frank Rayal is a telecom industry professional with over 20 years of experience working with network operators and system vendors to develop and deploy innovative wireless solutions. He's is a founding partner at Xona Partners a boutique management and technology advisory firm specialized in TMT and a founding member of small cell backhaul pioneer BLINQ Networks. Frank 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, OH, and a MSc in Electrical Engineering and an MBA from the University of Toronto, Canada. Frank is a Senior Member of the IEEE, and a member of Professional Engineers Ontario.

© 2013 Senza Fili Consulting, LLC. All rights reserved. This paper was prepared on behalf of Altobridge Limited. The views and statements expressed in this document are those of Senza Fili Consulting LLC, and they should not be inferred to reflect the position of Altobridge Limited. The transcription of the conversation has been edited for consistency and readability. The document can be distributed only in its integral form and acknowledging the source. No selection of this material may be copied, photocopied, or duplicated in any form or by any means, or redistributed without express written permission from Senza Fili Consulting. While the document is based upon information that we consider accurate and reliable, Senza Fili Consulting makes no warranty, express or implied, as to the accuracy of the information in this document. Senza Fili Consulting assumes no liability for any damage or loss arising from reliance on this information. Trademarks mentioned in this document are property of their respective owners. Cover-page and last-page graphics from Milos Dizajn/Shutterstock.

