Why C-RAN deployments will give you more headaches than D-RAN rollouts

and how to avoid them

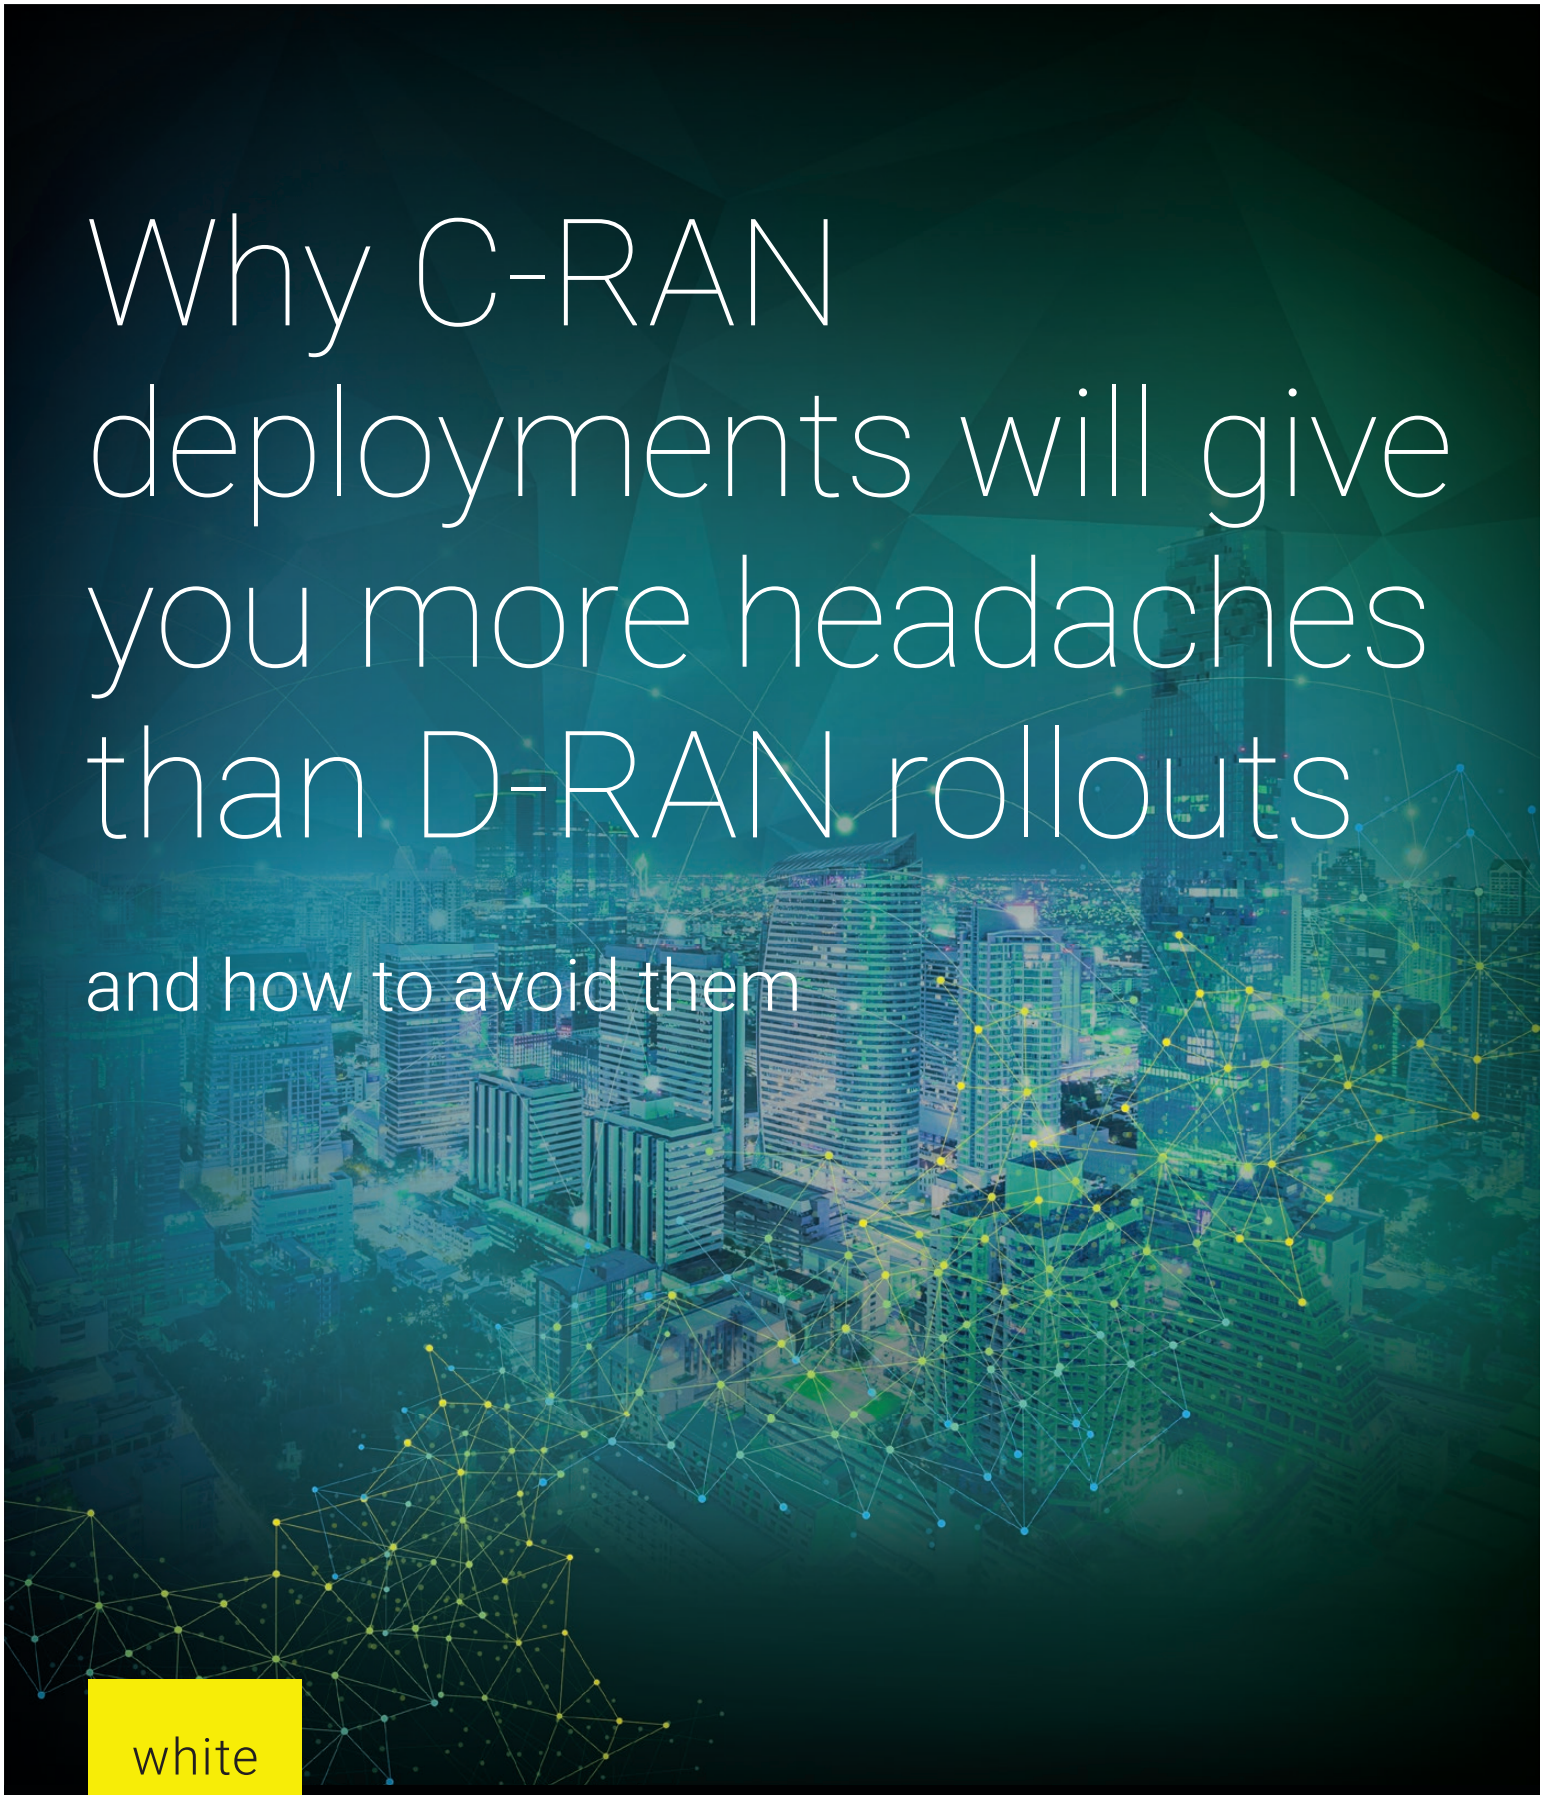

white paper

EXFO

Why C-RAN deployments will give you more headaches than D-RAN rollouts

and how to avoid them

white paper

By Jean-Sébastien Tassé
Product Manager

Each step in the evolution of radio access networks (RANs) brings unique opportunities and challenges for mobile network operators. Today, fiber has mostly replaced coaxial cables in new mobile fronthaul builds, otherwise known as fiber-to-the-antenna (FTTA). The next technological innovation in mobile fronthaul is centralized radio access network (C-RAN), a trend already present in some parts of the world. The advantages of C-RAN have been widely discussed (please refer to EXFO's white paper entitled The path to 5G requires a strong optical network—from C-RAN to Cloud-RAN). For the purpose of this white paper, we will provide a field operations perspective to C-RAN deployments by reviewing the additional challenges they bring in comparison to "traditional" D-RAN rollouts. Specifically, the paper will examine in detail how C-RAN can introduce pressures on optical loss budget and interconnection complexity—and what you can do about it.

The next technological innovation in mobile fronthaul is centralized radio access network (C-RAN).

About C-RAN

When discussing C-RAN, it is important to first describe two essential pieces of equipment used in typical fronthaul architecture: the base station (also called BBU or baseband unit) and the remote radio head (RRH):

1. As the name implies, the base station used to be located at the base of the cellular tower, and its main function is to convert the IP/Ethernet data coming from the mobile backhaul into digitized radio frequency signal (I-Q data) that is transmitted in the fronthaul network using CPRI (Common Public Radio Interface) or OBSAI (Open Base Station Architecture Initiative) protocols. The BBU demarcates the boundary between the backhaul (towards the core network) and the fronthaul (towards the radio antenna).

2. The RRH converts the digitized RF data into analog RF signals to be sent over the air in the downlink direction, and vice-versa in the uplink direction.

In the common distributed RAN topology (D-RAN), the BBU and the RRH are collocated, which means that the BBU is located at the bottom of the tower (Figure 1). The main innovation in C-RAN topology is to centralize all the BBUs in a BBU hotel (Figure 2) that can be positioned up to 20 km away from the RRH, allowing for energy and space reduction, better synchronization, etc.

The main innovation in C-RAN topology is to centralize all BBUs in a BBU hotel that can be positioned up to 20 km away from the RRH.

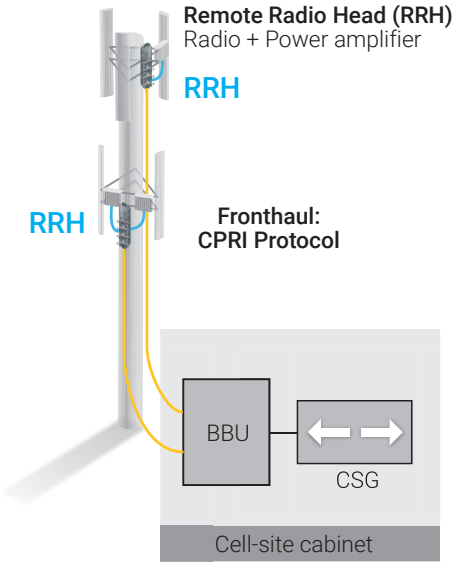

Figure 1. D-RAN architecture

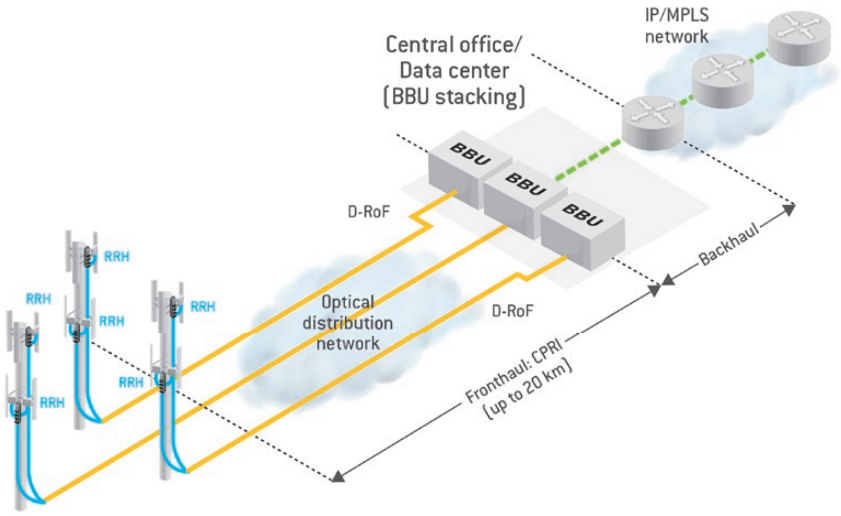

Figure 2. C-RAN topology

C-RAN headache #1: attenuation

Attenuation, which is the power reduction experienced by the optical signal as it propagates in the fiber, is a critical consideration for successful fiber optic communications because receivers might fail to detect low-power signals. Attenuation originates from several factors: the fiber itself through phenomena such as Rayleigh backscattering, connectors (dirty, damaged, or clean), macrobends (loss due to tight fiber bends), among others.

If we compare D-RAN and C-RAN deployments, we see that both can suffer from dirty or damaged connectors, and macrobends. **However, in terms of attenuation, C-RAN poses a much greater challenge than D-RAN because of two main reasons: the longer fiber spans and the number of connections.** First, fiber attenuation is usually around 0.2 dB/km at 1550 nm and 0.35 dB/km at 1310 nm. This implies that fiber attenuation is small for an D-RAN deployment that consists of a few tens of meters of fiber running from the bottom to the top of a tower.

However, fiber attenuation can be several dBs in C-RAN rollouts of 10 to 15 km of fiber. In addition, given the sheer distance of C-RAN rollouts and the more complex topology involved—sometimes a ring architecture—these rollouts generally have more connectors between the BBU and the RRH than is the case with D-RAN deployments. This increases the likelihood of having dirty connectors in the light path. **In summary, optical attenuation poses a bigger threat to network stability in C-RAN rollouts than in D-RAN deployments.**

To avoid attenuation issues, it is necessary to first determine the optical loss budget, which is the maximum attenuation that the optical signal can undergo and still be detected at the receive side.

Loss budget determination

To avoid attenuation issues, it is necessary to first determine the optical loss budget, which is the maximum attenuation that the optical signal can undergo and still be detected at the receive side. Most C-RAN deployments use small form-factor pluggable transceivers (SFPs) so the loss budget can be determined from the SFP spec sheet, which usually gives minimum and maximum launch powers.

Since the loss budget is determined considering the worst-case scenario (i.e., minimum launch power), it can be calculated with the following formula:

Loss budget = minimum launch power minus receiver sensitivity

In this formula, sensitivity stands for the minimum power that the receiver can detect. It should be noted that the loss budget varies quite a lot between models and vendors, even between SFPs that are all meant to cover the same distance (e.g., "10 km SFP"), as shown in Figure 3.

SFP vendor	SFP model	Reach	Loss budget
A	#1	10 km	4.4 dB
A	#2	10 km	8 dB
B	#3	10 km	6.4 dB
B	#4	10 km	9 dB
C	#5	20 km	18 dB
D	#6	20 km	14 dB
B	#7	20 km	10.4 dB
A	#8	40 km	15 dB
C	#9	40 km	22 dB
D	#10	40 km	25 dB

Figure 3. Examples of SFP+ loss budget vs. stated reach

We see, for instance, that 10 km SFPs show loss budgets from 4.4 to 9 dB—a huge variability—while a 40 km SFP loss budget ranges from 15 to 25 dB. There are even some incongruences, such as a 40 km SFP (model #8) with a loss budget of 15 dB, which is lower than a 20 km SFP (model #5, with 18 dB loss budget). **It is therefore clear that system performance is dictated by the SFP loss budget—not the rated SFP reach.**

Depending on the fronthaul architecture (e.g., passive, semi-passive, active), there could be one or several loss budgets between the BBU and RRH to determine. In a passive architecture, the wavelengths are distributed to the various RRHs using purely passive optical add/drop multiplexers (OADMs) or MUX (filters), such that the signal remains in the optical domain from BBU to RRH. Accordingly, a single loss budget must be determined from the BBU hotel to the RRH, as shown in Figure 4. In contrast, active fronthaul and semi-passive architectures feature optical-to-electrical-to-optical conversion at OADM locations, thanks to the SFPs located at each OADM site. Therefore, two or more optical budgets must be determined.

Depending on the fronthaul architecture (e.g., passive, semi-passive, active), there could be one or several loss budgets between the BBU and RRH to determine.

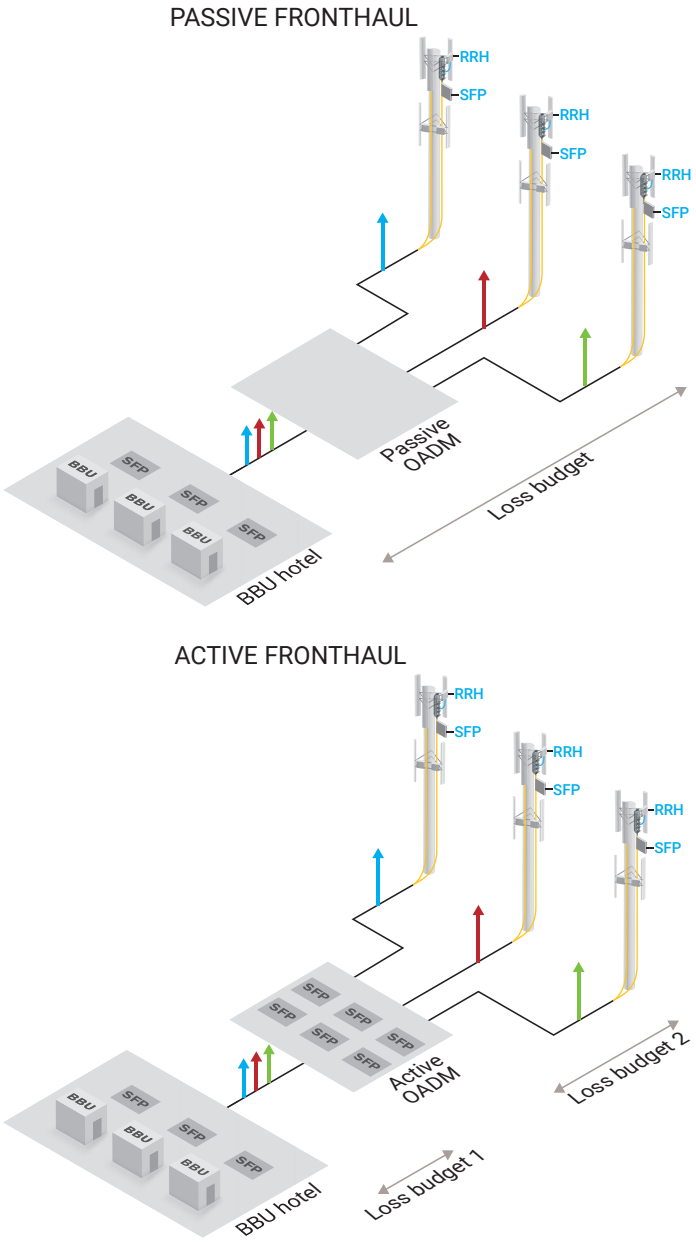

Figure 4. Loss budget determination in active vs. passive fronthaul

a. Universal Virtual Sync is available on the following EXFO Active Verifier models: BV-110/1100/1500/3000/3100

An optical time domain reflectometer (OTDR)—or its more advanced version, EXFO'S intelligent optical link mapper (iOLM)—allows for precise measurements of the attenuation in the C-RAN network.

An optical time domain reflectometer (OTDR)—or its more advanced version, EXFO's intelligent optical link mapper (iOLM)—allows for precise measurements of the attenuation in the C-RAN network, and identification of any elements that might cause excessive loss (dirty or damaged connectors, macrobends, bad splices, etc.), as well as their precise location. **It is therefore critical to perform an OTDR measurement on any new C-RAN fiber installation.** More information on the highly automated iOLM, an efficient tool to reduce operational expenses is available here: www.exfo.com/iolm.

The C-RAN trials in which EXFO participated revealed that 20% to 50% of connectors typically fail industry standards for cleanliness—a huge percentage! Should any dirty connectors be found using the OTDR/iOLM, a fiber inspection probe is essential to determine whether the cleaning was properly done; it does so by automatically comparing the captured image against industry standards such as the IEC or IPC.

Why using long-reach SFPs carry huge hidden costs

Some service providers have told EXFO that they prefer to employ long-reach SFPs (such as 40 km or 80 km SFPs) instead of short-reach SFPs (such as 10 to 20 km) in C-RAN deployments to compensate for dirty connectors and avoid testing their networks with OTDRs and inspection probes. As the following example will reveal, this approach entails significant extra costs.

We will consider a typical C-RAN network, with 7-8 km point-to-point connections between the BBU and the RRH, and calculate the extra cost of using 40 km SFPs to compensate for dirty connectors, instead of relying on 10 km SFPs if the connectors were clean. Figure 5 discloses the SFP prices of some well-known vendors. By computing the average for the two vendors, we find that a 40 km SFP costs approximately $300 more than a 10 km SFP.

SFP vendor	SFP model	Reach	Cost
A	#1	10 km	$94
B	#2	10 km	$238
A	#3	20 km	$178
B	#4	30 km	$258
A	#5	40 km	$369
B	#6	40 km	$600
B	#7	50 km	$700
A	#8	60 km	$641
B	#9	80 km	$960

Figure 5. SFP cost as a function of reach

Next, we have to find the number of SFPs used in each cell site. Typically, a cell site will feature many RRH types for the different types of services offered (PCS, AWS, LTE, etc.). In addition, each sector will require an RRH for the different service types to ensure proper cellular coverage, each sector typically covering 120º (Figure 6). In our example, we will use the realistic assumptions that there are 3 sectors in the cell site and that 3 types of services are available, which gives a total of 9 SFPs located at the tower top (one per RRH). These SFPs communicate with an equal number of SFPs at the BBU hotel, so a total of 18 SFPs are required for this typical cell site. **With an extra cost of $300 per SFP, we conclude that employing 40 km SFPs rather than 10 km SFPs carries an additional cost of $5,400 per cell site!** This figure will vary depending on the specific type of SFP, the SFP vendor and the number of services offered, but the bottom line is that this represents a huge hidden cost. We invite the reader to calculate the actual cost for his or her own situation. **Given that many mobile providers have thousands of cell sites, we see that a proper optical testing approach (OTDR, connector inspection) combined with the use of shorter-reach SFPs can save millions of dollars in rollout costs.**

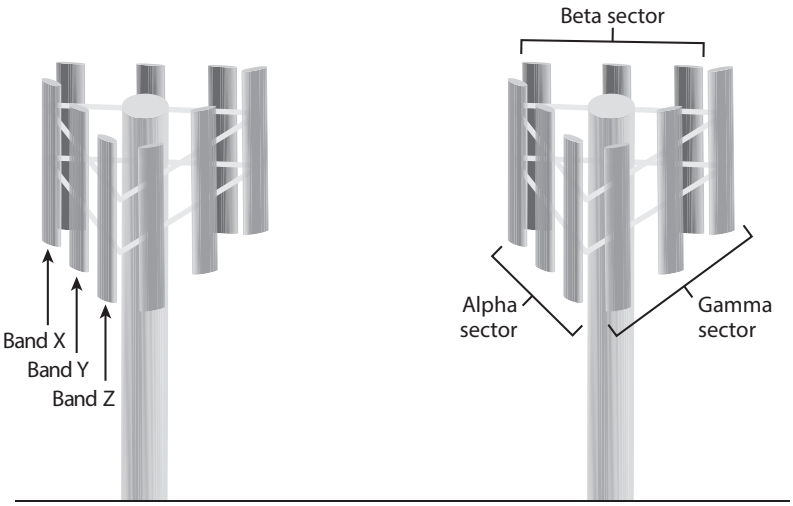

Figure 6. Concept of sectors and bands

We anticipate that interconnection issues will be a major problem in C-RAN networks, a situation already experienced by early C-RAN adopters in countries like South Korea or the United States.

C-RAN headache #2: interconnection issues and wavelength addressing

In D-RAN deployments, the number of connections between the BBU and the RRH is rather small because of the short distances involved. C-RAN is an entirely different scenario because of the long distances, in the kilometer range. Based on EXFO's experience with fiber-to-the-home (FTTH) deployments and broadband access networks, which are similar in many ways to C-RAN, we anticipate that interconnection issues will be a major problem in C-RAN networks, a situation already experienced by early C-RAN adopters in countries like South Korea or the United States.

Interconnection issues occur when a connector is plugged into the wrong port, or when a fiber is spliced with the wrong fiber. Although interconnection issues might seem trivial, they are actually very common and not easy to troubleshoot without the right tool. Figure 7 illustrates a simplistic C-RAN ring network where the BBU hotel emits three wavelengths in the ring, and one wavelength goes to each tower (blue to Tower A, red to Tower B and green to Tower C). The OADM, which is referred to by many different names including *node*, selects a specific wavelength and drops it to the right tower. This function is also referred to as wavelength addressing. If the interconnections at the OADM level or anywhere else in the network are not properly made, then the wavelength might end up in the wrong location (e.g., blue wavelength ends up in Tower B).

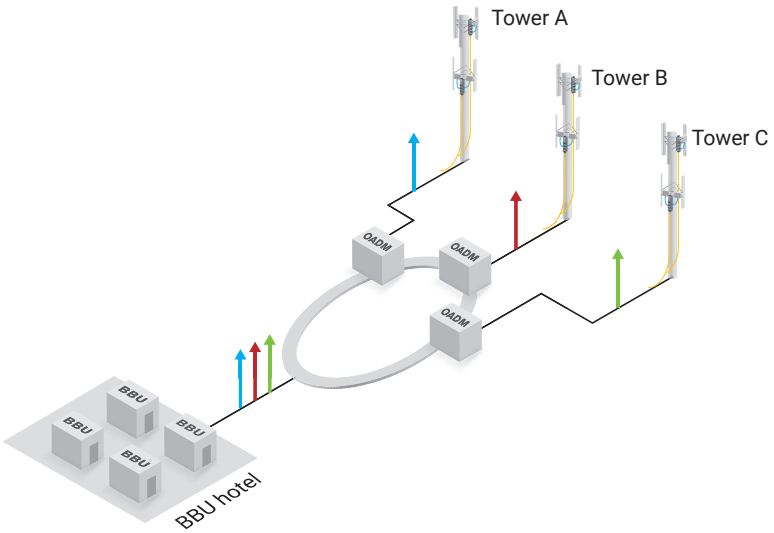

Figure 7: Wavelength addressing in C-RAN

Although interconnection issues might seem trivial, they are actually very common and not easy to troubleshoot without the right tool.

The best way to confirm that all the interconnections are correct is to test them with a tunable OTDR, which is an OTDR that allows the user to select the wavelength. For instance, in the simplified example above, the user could choose the “green” wavelength in the OTDR and run a measurement. If the interconnections are good, the green OTDR pulse will travel around the ring, get dropped by the OADM towards Tower C, be reflected at Tower C, and then travel back in the same fashion towards the BBU hotel (see Figure 8). The tunable OTDR would then give the distance traveled by the pulse, which should be compared against the network plan and thereby enable the user to validate the interconnections if the distances match. The same process can be repeated with all the wavelengths. Even better: all measurements can be performed from a single location, the BBU hotel in this example. There are two types of tunable OTDRs: CWDM OTDR and DWDM OTDRs, depending on the WDM technology used in the C-RAN network.

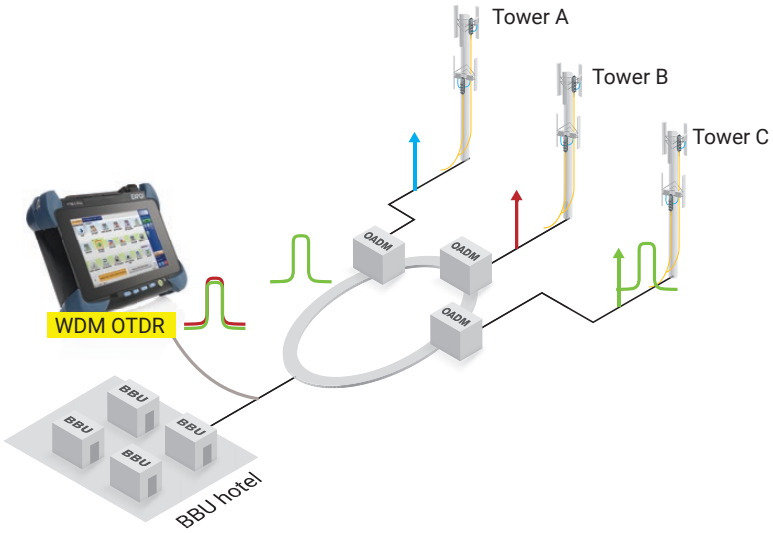

Figure 8: C-RAN testing using a tunable OTDR to validate interconnections

Conclusion

This white paper discussed the migration of fronthaul networks from D-RAN to C-RAN in the coming years. This transition increases the likelihood of some optical impairments, such as loss of signal due to excessive attenuation or wavelength addressing issues. Excessive attenuation is due to the longer span distances in C-RAN as compared to D-RAN, as well as the higher number of connections involved. Wavelength addressing issues are most common in C-RAN architectures, and result from connecting the wrong ports together due to the increased number of interconnections. The best way to avoid both issues is to conduct proper testing using OTDRs/iOLM and fiber inspection probes during construction. Finally, we have demonstrated that using long-reach SFPs of 40 or 80 km to avoid cleaning of connectors carries huge hidden costs, in the thousands of dollars range per cell site. A much better—and more cost-effective—approach is to inspect and clean connectors especially during construction and turn-up, and use short-reach 10 or 20 km SFPs.