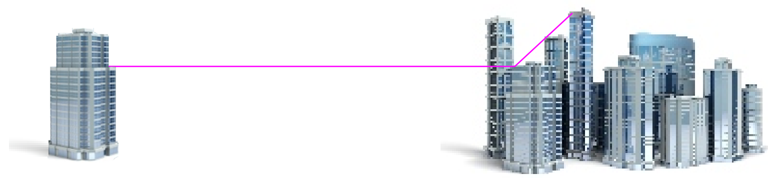

Ultra Low Latency Wireless Point-to-Point Networks

White Paper

Dr. Heinz Willebrand

LIGHTPOINTE™ WIRELESS

Table of Contents

Executive Summary

For more than a decade, increasing the speed of a network to fulfill the exponential rise in demand for bandwidth has been one of the major goals of infrastructure providers and carriers worldwide. Significant progress has been made by transitioning to faster packet based protocols, deploying optical fiber infrastructure, and utilizing associated technologies. Advances have included technologies such as wavelength division multiplexing (WDM), and an overall increase in the speed of electronic and electro-optic components—microprocessors, FPGAs, and optical transceivers just to mention a few.

Yet, with more and more real-time applications coming online, not only in the wireline but in particular in the Mobile Wireless sector, network latency (a synonym for 'delay') has become increasingly important. The new paradigm in designing network infrastructure is now driven by careful consideration of achieving both **high network capacity *and* very low/ultra-low latency performance**.

Ultra-low latency networks can be beneficial in numerous industries and applications, such as emergency response, medical networks, military applications, and other time-sensitive environments.

One particular high profile application, where ultra-low latency is extremely important, is in High Frequency Trading (HFT) networks. Since reducing the time it takes for electronic trading platforms to make a trade can result in millions of dollars of gain or loss, this industry sector has deployed significant numbers of ultra-low latency wireless networks. A 1999 decision by the Securities & Exchange Commission (SEC) made it possible for electronic exchanges to compete with the NYSE and other trading organizations. This laid the foundation for HFT and, according to estimates from Rosenblatt Securities, as much as two-thirds of all stock trades in the U.S. from 2008 to 2011 were executed by high-frequency firms. In 2012, HFT accounted for 1.6 billion trades a day. The HFT organizations are usually not interested in holding onto a stock for more than a few seconds, and they virtually start from scratch each day without any holdings.

Since financial markets are driven by high-speed, high-return activities, fast execution of a transaction (before a competitor finds out) can determine who gets the profit from a trade. The old days of ticker tape and phone calls are long gone, and response times have gone from milliseconds to microseconds. With billions of

"There's a huge technology arms race to drive out latency, because when you're talking about latencies in nanoseconds and microseconds, a millisecond is an eternity." [1]
Dave Lauer, High-frequency Trading Professional

trades a day and billions of dollars traded, markets and traders are heavily investing in the latest state-of-the-art in networking and computing technologies to access and analyze data, and to perform a trade virtually in real-time.

To optimize High Frequency Trading networks, there are two key elements: (1) the speed of computing hardware used, and (2) the speed and latency of the exchange interconnections. Regardless of how fast the computing hardware is, or how excellent the trading software utilized is, the trading system is highly dependent on how fast the information required for a trade is made available via the network. An HFT operation might have the fastest super computer crunching the numbers, but if the numbers are not fed to it quickly enough it is of little benefit.

Optical fiber lines have traditionally been the preferred infrastructure medium to interconnect exchanges. Light travels in optical fiber at 2/3 of the speed of light in a vacuum, which is 186,282 miles per second. In other words, at a speed of 124,194 miles per second it takes information about eight microseconds to travel one mile in a fiber optic cable. While this has, in the past, been considered phenomenal speed, today it is not necessarily fast enough for High Speed Trading platforms and other time-sensitive applications, where saving microseconds can provide tremendous advantages.

As a logical consequence, and with the goal to minimize the distance and network latency between computing platforms connected to the exchanges, trading firms have placed their computing infrastructure at buildings as close as possible to the exchanges.

In New York City, many firms decided to move their data centers out of Manhattan to reduce cost and reduce risks, after the 9/11 terrorist attack. Many trading companies now operate out of ordinary and unmarked buildings located in New Jersey, or on Long Island. Similar relocations of data centers and trading companies have also taken place in other financial hubs around the world.

While these changes have made the data centers more secure and drastically reduced costs for building leases, this has also resulted in the challenge to overcome system latency, due to the longer fiber runs required to connect to the exchanges. These longer runs can add several hundreds of microseconds of latency or more, simply because the fiber routes follow natural and man-made obstacles like rivers, streets, buildings, subways routes, etc. In High Frequency Trading environments, and when the fiber distances are several times the direct line-of-sight distance between buildings to be interconnected, even fiber optic connections can be too slow.

[1] http://www.pcmag.com/article2/0,2817,2424495,00.asp

Wireless Ultra-Low Latency Technologies

The latency penalties, induced on High Speed Trading networks by longer fiber routes, caused the designers and operators to look at alternative technologies, such as wireless point-to-point bridges. The core advantages of the latest ultra-low latency (ULL) wireless bridges over the use of fiber are two-fold:

- ULL Radio Frequency-based wireless bridges are faster. An electro-magnetic wave travels at a speed of roughly 189,000 miles per second through the air vs. light traveling at only about 124,000 miles per seconds through an optical fiber cable.

- A direct line-of-sight wireless point-to-point connection is the shortest distance between two remote locations. Furthermore, point-to-point connections do not incur the distance/latency penalty caused by more complex fiber routes and rights-of-way issues.

In High Frequency Trading exchanges, these two distinct latency advantages over fiber optic transmission can result in an organization achieving additional profits, and in reducing risk caused by delays.

MICROWAVE RADIO TECHNOLOGY BELOW 40 GHZ

Point-to-point microwave systems operating over narrow frequency blocks of less than roughly 50 MHz, and frequencies below 40 GHz, have been used for many decades to establish a point-to-point networking connection between remote locations. Frequencies below 10 GHz are very resilient to weather conditions and offer very long distances. Moderate capacity links of 100 Mbps speed or less can be established even under heavy rain and snow conditions. When using more advanced coding technologies such as higher order QAM modulation to improve spectral efficiencies (a.k.a. bit/Hz performance), and error correction methods, much higher link capacities can be reached. When these techniques are combined with dual-polarization antennas, full duplex connections with a bandwidth of several hundred Mbps can be achieved over narrow frequency channels.

Unfortunately, signal processing, coding and error correction schemes massively impact the latency of these higher capacity radio links. And they can be very expensive to deploy. Although still very useful in standard networking applications requiring point-to-point connectivity, they are for the most part no longer useful for high speed trading and other time-sensitive applications where ultra-low latency is required.

Widely used Ethernet based point-to-point microwave radios, that also process Ethernet frames/traffic, can have radio equipment latency of 100 microseconds or higher. Such equipment latency alone compares to a latency induced by roughly 12.5 miles of fiber. And for many applications, especially in urban based/metro networks, the latency advantage of using such sub-40 GHz radio solutions over a fiber connection is dramatically diminished.

V/E BAND MILLIMETERWAVE RADIO TECHNOLOGY

Millimeterwave radio solutions operating in the 60 GHz (V- Band) and 70/80 GHz (E-Band) frequency spectrum are excellent candidates for high capacity and ultra-low latency wireless point-to-point connections. In the US and the UK, two of the largest financial hubs in the world, a total of two 5 GHz frequency blocks in the 71…76/81…86 GHz band, and a 7 GHz frequency block in the 57…64 GHz band, were released for lite-licensed and unlicensed use.[2]

Because of the large spectrum blocks available, no complex and higher order signal modulation is needed to transmit signals at Gbps speeds and higher across a V/E Band radio link. Simple ASK/OOK modulation (as used in the LightPointe physical layer 1 and protocol transparent Aire X-Stream products) allow for a total equipment latency of less than 10 ns per terminal while transporting full duplex GbE traffic. When using dual-polarization antennas, total link capacity can be doubled and provide the same ultra-low latency levels. ***In terms of optical fiber latency, which is 20 cm/ns, the radio terminal latency of 10 ns is equivalent to only 2 meters or about 6 feet of optical fiber.***

Figure 1:
LightPointe Ultra-Low Latency Aire X-Steam MMW radios (up to 2.5 Gbps)

This same technology approach is used by LightPointe to build ultra-low latency 60 GHz radio systems at GbE speed and equipment latency of less than 10 ns/terminal.

When compared to standard microwave radios, a shortcoming of V/E Band radios is the shorter reach caused by higher levels of signal attenuation, in particular due to rain and the additional high level of oxygen absorption in the case of 60 GHz radios. While E-band radios with higher power transmission amplifiers can still reach distances of more than 8 miles at GbE speed under clear weather conditions, the distance of a 60 GHz radio operating at GbE speed under the same weather conditions is limited to about 1 mile. When higher overall system availability is required, the distances should be kept shorter. This shorter reach generally restricts the use of V/E Band systems to urban/metro environments. However, when systems incorporate appropriate clock recovery circuitry, a cascading multiple radio link is possible and total reach can be increased.

[2] The regulatory landscape can differ from country to country. In Europe e.g. the 70/80 GHz spectrum has been released by the ETSI for commercial use, but individual national regulators can impose own specific rulings. The same is true for operation in the 60 GHz band.

Figure 2:
LightPointe Aire X-Stream FSO system with four autotracking, overlapping laser beams

INFRARED LASER BASED FREE SPACE OPTICS (FSO) TECHNOLOGY

Even long before the widespread use of millimeter wave radios, high capacity and ultra-low latency Free Space Optics communication links were used to connect remote networking locations. The infrared spectrum is license-free worldwide for use in communication systems and this makes the spectrum appealing. The narrow divergence angle of the optical beam, which is typically only a fraction of a degree, ensures that FSO systems can be deployed virtually right next to each other. And data is well protected from physical interception/eavesdropping. Modern high performance FSO systems (Figure 2) employ an active tracking system to keep the narrow beam(s) on target. Available spectrum is virtually unlimited (several hundreds of GHz). One of the highest capacity FSO systems built was able to deliver a total capacity of 1.2 Tbps using 30 wavelengths each modulated at 40 Gbps (lab environment). In 2012, American scientists from the University of Southern California in Los Angeles, NASA's Jet Propulsion Laboratory, the University of Tel Aviv, and China's Huazhong University of Science and Technology demonstrated a 2.5 Tbps Free Space Optical transmission using an orbital angular momentum (OAM) technique.[3] As for distance records, in September of 2013 NASA deployed a lunar satellite which will utilize a ½ Watt laser to transmit to/from the Earth.

FSO bridges are utilized in ULL networks, especially where there is a high degree of RF congestion. In some cases, MMW radios are utilized in conjunction with an FSO bridge (Figure 3 below).

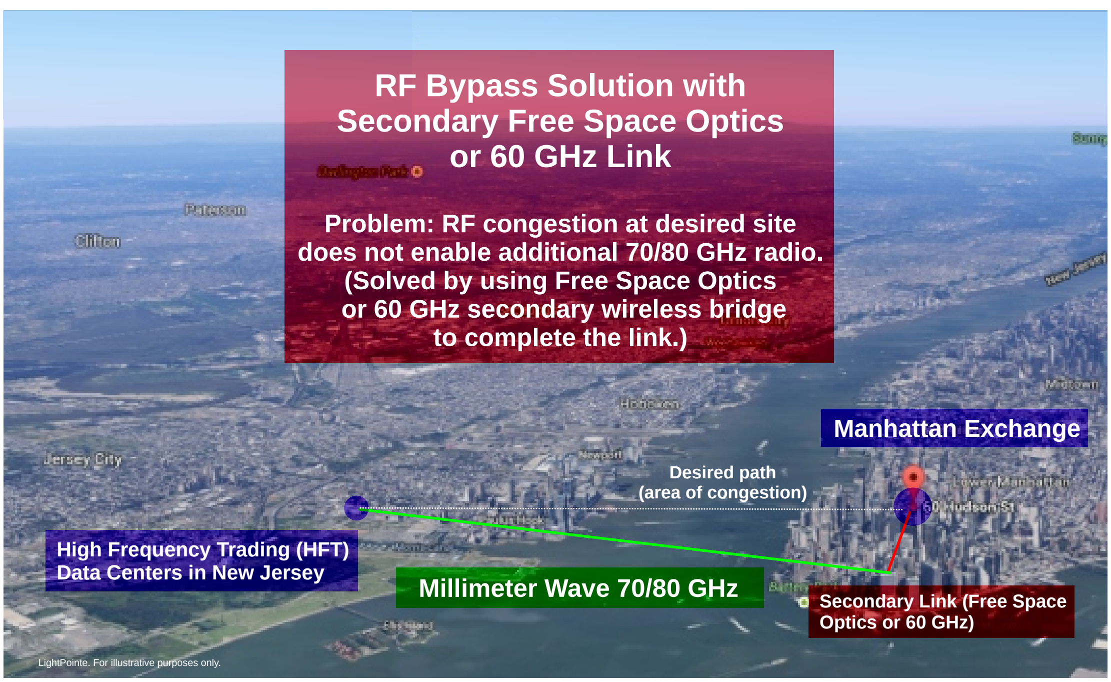

Figure 3: Example of ultra-low latency connection which bypasses an area of radio frequency congestion by utilizing FSO or 60 GHz.

[3] http://www.extremetech.com/extreme/131640-infinite-capacity-wireless-vortex-beams-carry-2-5-terabits-per-second

Thousands of full duplex Gigabit capacity FSO systems have been field-deployed all over the world. A Layer 1 transparent FSO system, such as the LightPointe Aire X-Stream FSO system, operates at full duplex GbE speed and a latency figure below 10 ns/terminal. FSO systems offer a unique combination of ultra-low latency, high capacity, and unlicensed and interference-free operation. Unlike many RF solutions, thousands of FSO links can be deployed in a downtown environment.

Although the reach of FSO systems under clear weather conditions, or even snow, can be long, the signal attenuation in some regions prone to foggy or, to a lesser extent, heavy rain conditions can decrease availability. For this reason, most FSO installations are limited to distances below one mile, which is about the same as 60 GHz wireless bridges.

Ultra-Low Latency System Design Criteria

When considering wireless connectivity solutions as a means of reducing a network's latency, there are many factors to consider in order to achieve optimum performance.

- *Shortest Physical Path Objective:* The selection of the shortest possible route is the key in minimizing the overall "distance induced latency" between locations, which is one of the greatest contributors to the end-to-end latency. The most straight forward approach is to use the direct line-of-sight path, as shown in Figure 4 below. Compared to a high capacity fiber connection, the direct wireless line-of-sight path is always the shortest and consequently lowest latency path between two remote locations. Add the fact that an optical signal over fiber only travels 2/3 of the speed of an electromagnetic signal's propagation through the air, the latency-decrease benefit of using a wireless connection is significant. Nevertheless, the selection of the optimal wireless route can be complex and will depend on factors such as roof-right/tower availability, the actual distance between locations, and the number of repeater locations. Figure 4 illustrates how fiber can be bypassed.

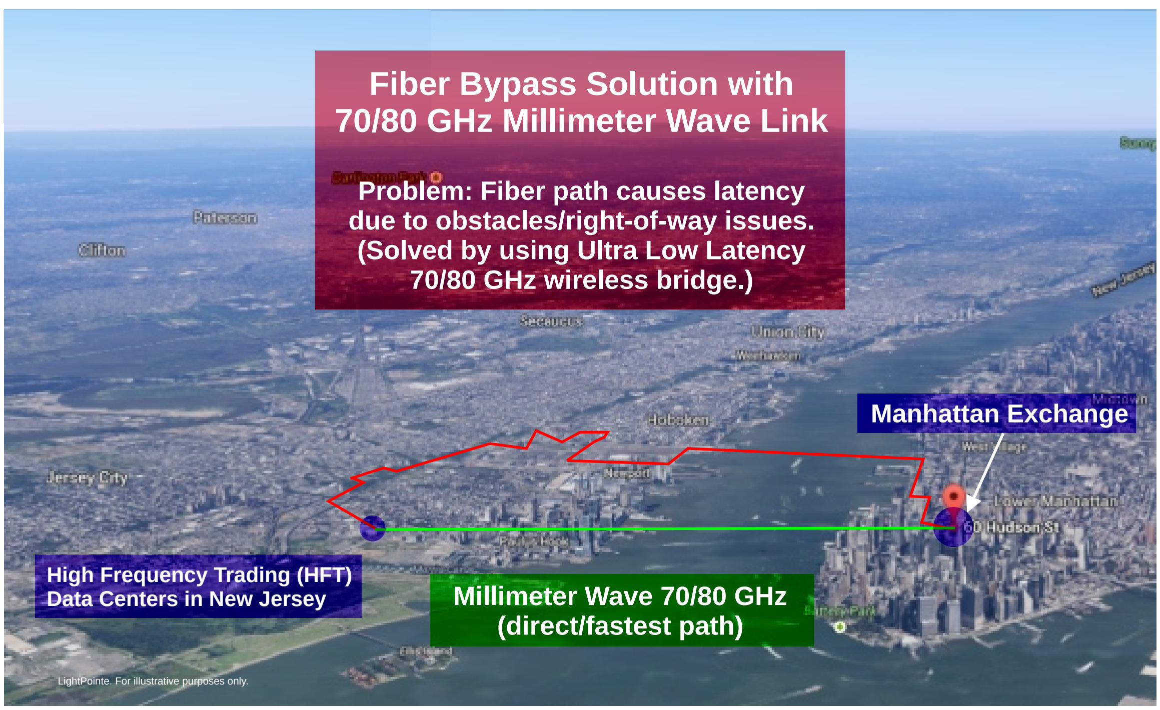

Figure 4: Example of ultra-low latency connection, fiber vs. wireless.

- *Lowest # of Hops Objective:* In case there is no direct line-of-sight between locations and/or when the distance is simply too long to be covered by a single point-to-point connection, reducing the number of hops needed helps ensure the lowest latency possible. Each microwave-to-microwave (or FSO) link is an additional hop, and more hops consequently introduce more processing latency. The number of hops needed will depend directly on the maximum link distance required in the geographic area of deployment. In general, determining the right 'number of hops' requires finding a balance between the 'total number of hops' needed to meet the overall availability objective, the financial model, and the latency penalty induced by adding additional hops. The latency penalty is of minor concern for layer 1 systems (e.g., LightPointe Aire X-Stream systems), because each hop simply acts as a layer 1 signal repeater with full clock and data recovery (CDR) capability, and only adds about 20 ns of additional latency. Some wireless bridges, however, can add much more latency due to their performing additional signal processing, which severely impacts the end-to-end low latency design objective of the network.

Addressing network topology issues and minimizing end-to-end latency by optimizing overall distance and the number of hops is extremely important. However, and particularly in the case of using licensed or lite-licensed wireless microwave or millimeter wave radios, the network designer additionally needs to take into account other factors such as national frequency spectrum plans, frequency availability and coordination in the region, as well as zoning restrictions.

Using different ultra-low latency wireless systems based on licensed/lite-licensed 70/80 GHz as well as unlicensed 60 GHz[4] and FSO technologies provides additional flexibility when designing the network. For example, in scenarios where frequency availability or congestion of the longer reach 70/80 GHz spectrum is problematic, the creation of a hub location in close vicinity to the final target location could solve the problem.

On the following page, Figure 5 illustrates a scenario where a High Frequency Trading data facility (labeled location **A**) first connects to a hub location **B** using a 70/80 GHz link because there isn't spectrum available to directly connect to the Exchange location **C**, or there is RF congestion/interference potential at location **C**. Due to the narrow antenna beam profile of the 70/80 GHz radio, just a few degrees of angular change opens up additional 70/80 GHz spectrum and consequently the hub location **B** can be located close to the target Exchange location **C**. Due to the very short distance between the hub location and the Exchange, the final link can be accomplished with a short distance and ultra-low latency FSO system or 60 GHz radio link. The "latency penalty" resulting from the additional second connection can be expressed in terms of length of the individual connections "**b**" and "**c**" as well as the angle "**a**" between these two segments, and use of the cosine law (next page).

[4] V Band operation is not unlicensed in all parts of the world. E.g. the US, Canada, and the UK allow outdoor unlicensed operation in the 57...64 GHz band. Other countries limit the amount of spectrum available for outdoor use.

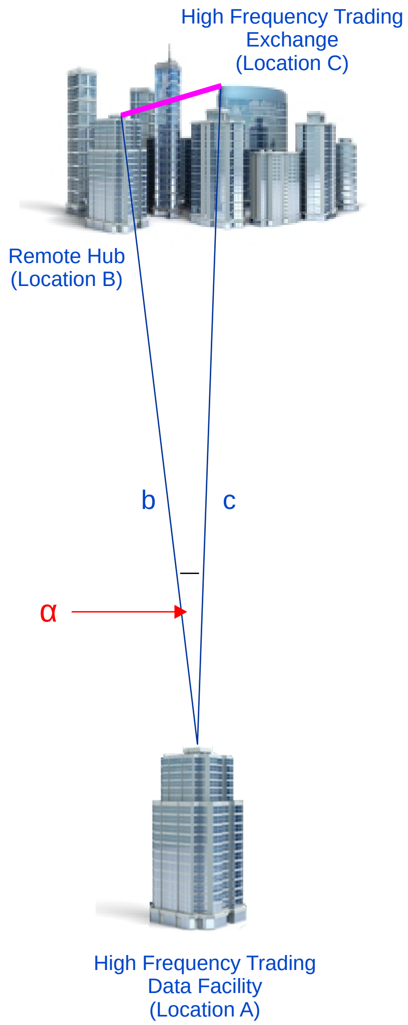

Figure 5:
Remote Hub Bypass Application
(Used when line-of-sight to target building is not available, and/or installation of additional radios is not possible on target building due to RF congestion or rights issues.)

(1) $a^2 = b^2 + c^2 - 2bc*\cos\alpha$

The latency penalty t_p = using two instead of one link can be expressed as:

(2) t_p [ns]= radio link latency [ns] + (a+b-c) / 0.3;   a, b in meters

Formula (2) can be used to calculate the latency penalty of a two segment versus a direct connection.

EXAMPLE:
c = 10,000 m; b = 10,000 m, a = 2 degrees yields a = 350 meters when using formula (1).

Using this value in formula (2) yields a total additional latency of 1.116 µs.

Latency of direct line-of-sight connection (A→C): **33.333 µs**

Latency of hub connection (A→ B→ C): 33.333 µs + 1.116 µs = **34.449 µs**

Although the latency number for the remote bypass application is about 1 µs higher compared to the direct line-of-sight connection, a comparison with the latency induced by a hypothetical "direct" line-of-sight distance fiber optic connection still shows the benefit of using this approach. Realistically, the fiber distance in a more densely populated urban environment is around 2 – 2.5 times the line-of-sight distance and this further dramatically increases the total end-to-end connectivity latency.

Latency of "direct" 10,000 meters line-of-sight distance fiber connection (A→C): **49.999 µs**

Latency of distance weighted 10,000 m urban fiber connection (2.5 distance factor): **124.998 µs**

Summary

High capacity and ultra-low latency network connections are the lifeline of many real-time networking applications. By using a combination of high capacity and ultra-low latency wireless bridges (e.g., E-Band radios operating in the 70/80 GHz frequency spectrum; V-Band 60 GHz radios; laser-based FSO solutions), an infrastructure designer can engineer a network that vastly out performs fiber connections, from a latency perspective.

The best ULL wireless bridges available offer virtual zero latency (VZL) designs, and have active clock and data recovery (CDR) circuitry. This enables the bridges to be cascaded—"daisy chained"—to increase distances without having to consider the increase in latency at the hub location(s). In addition, in the case of ULL radio deployment, the use of a dual polarization adapter doubles capacity, yet only requires a single antenna installation.

Whether deploying ULL radio bridges, or ULL FSO bridges—or a combination of radios and FSO units—organizations can now transmit data far faster than fiber networks.

For more information, please visit www.lightpointe.com.

About the Author

Dr. Heinz A. Willebrand is widely regarded as being among key driving visionary forces behind the Free-Space Optics (FSO) and millimeterwave wireless bridge movement, and is an inductee into EE Times list of *"Forty Innovators Building the Foundation of the Next-Gen Electronics Industry."* Well known and recognized for his contributions in the advancement of FSO, Heinz regularly speaks at national and international conferences on technical and business related aspects of FSO and Millimeter-Wave (MMW) technology. Heinz obtained his Ph.D in Applied Physics at the University of Muenster/Germany and has held research and teaching positions with the University of Muenster and the University of Colorado in Boulder. He holds more than 10 patents on fiber-optics and wireless technologies.

About LightPointe

LightPointe, an ISO 9001:2008 certified company, is the number one manufacturer of Gigabit capacity all-weather hybrid wireless bridges, 4th generation Free Space Optics links, and the world's fastest Ultra Low Latency (ULL) 60 and 70/80 GHz point-to-point radios. Since 1998, LightPointe has deployed over 14,000 systems with Enterprises, 4G/LTE Carriers, High Frequency Trading (HFT) networks, Government Agencies, Defense/Military Departments, and Security organizations.

LIGHTPOINTE™ WIRELESS

LightPointe Communications, Inc.
11696 Sorrento Valley Road, #101
San Diego, CA 92121
+1 858-834-4083