HFCL

The Role of Optical Connectivity in AI Infrastructure

This is what AI runs on

HFCL

Contents

HFCL

The Day the Internet Nearly Broke and What It Taught Us About AI Infrastructure

ChatGPT crossed one million users in just five days after launching in November 2022, at the time, the fastest consumer product adoption in digital history, beating Instagram's 2.5 months and Netflix's 3.5 years.

The servers behind that product, massive GPU clusters in data centers, were running so fast and so hard that the team had to physically ration access. Users were put on waitlists. The system creaked. And this was just one AI product, in its early days, before the real scale arrived.

That moment was a preview of a much larger reality now unfolding: the world is building an entirely new layer of infrastructure to support artificial intelligence. And at the center of that infrastructure, connecting every chip, every server, every data center to every user, is fiber.

HFCL

So, What Exactly Is AI Infrastructure?

Think of AI as a city. A very demanding, very power-hungry, very data-hungry city. Like any city, it needs roads, electricity, water, and buildings to function. AI infrastructure is all of that, the physical and digital foundations that allow AI systems to be trained, run, and delivered to users.

It has several interconnected layers:

Layer	What It Is	Why It Matters for AI
Compute (Silicon)	The GPUs and specialized chips that do the actual AI processing	AI models require massive parallel computation, thousands of chips working simultaneously
Optical Fiber and Networking	The cables and connections that move data between chips, servers, data centers, and users	Without high-speed, high-bandwidth fiber, AI data cannot move fast enough to be useful.
Data Centers	The physical buildings that house the servers and cooling systems	AI workloads generate enormous heat and require purpose-built facilities
Power Infrastructure	The electricity grid, backup systems, and increasingly on-site generation	A single AI data center can consume as much electricity as a small town
Storage	The systems that hold training data, model weights, and inference outputs	AI models train on trillions of data points; storage must be fast and vast
Software and Orchestration	The platforms that manage, monitor, and optimize AI workloads	Hardware alone cannot run AI at scale; intelligent software is required to coordinate it all

Notice that optical fiber appears in the second row, and it is not an optional component. Every layer of AI infrastructure depends on high-speed, high-bandwidth data movement. Without fiber, the GPUs sit idle. Without fiber, the data centers cannot talk to each other. Without fiber, the AI model on your phone cannot receive the processed response in time for it to be useful.

Fiber is not a supporting actor in AI infrastructure. It is load-bearing.

HFCL

Why Optical Fiber Is at the Heart of AI Infrastructure

Here is a story most people have never heard.

In 2016, a single undersea fiber optic cable called MAREA was laid across the floor of the Atlantic Ocean, stretching 6,600 kilometers from Virginia to Spain. This cable is designed with a capacity of up to 160 terabits of data per second, enough bandwidth, in theory, to stream millions of HD films simultaneously.

That cable, and thousands like it, is what makes the modern internet possible. And it is what will make AI at a global scale possible.

Here is why fiber specifically matters for AI infrastructure, not just any connectivity:

- **AI Data Is Enormous and Moves Constantly**

When an AI model is being trained, it processes billions of data points. Those data points must move from storage to compute chips, back and forth, thousands of times. When you use an AI application, your query travels from your device to a data center, gets processed across multiple servers, and the result is sent back, all in under a second. Copper cables and wireless connections simply cannot carry the volume and speed required. Fiber can. A single strand of optical fiber can carry data at the speed of light. Multiple strands bundled together create bandwidth that scales with demand. This is exactly what AI infrastructure requires.

- **Inside the Data Center: Intra-Datacenter fiber**

Most people think of fiber as the cable in the ground connecting cities. But an equally important application is inside the data center itself. Modern AI data centers connect thousands of GPUs to each other and to storage systems using high-density fiber optic interconnects[M1], short cables running between server racks, routing tables, and storage bays. As GPU racks have become more powerful, the data they generate and consume has grown exponentially. The networking inside a modern AI data center is, in many ways, as sophisticated as the networking between cities. High-density, bend-insensitive fiber cables, products that HFCL specializes in, are central to making this work. This is precisely the challenge HFCL addresses through OptiQ AI™, its unified portfolio of optical connectivity solutions engineered specifically for AI, cloud, and hyperscale data centre environments.
[M1]Get link to dci solutions LP

- **The Last Mile: Connecting AI to Users**

The most brilliant AI system in the world is useless if the user cannot access it quickly. The “last mile”, the final fiber connection from a local exchange to a home, office, or cell tower, determines whether AI applications feel instant or sluggish. FTTH (Fiber to the Home) and FTTB (Fiber to the Building) deployments, which HFCL has been actively building across India and globally, are the delivery mechanisms for AI at the user end.

In short: the data center needs fiber. The network backbone needs fiber. The last mile needs fiber. AI infrastructure, from end to end, runs on optical fiber.

HFCL

The Scale of What Is
Being Built Right Now

Numbers, when they get large enough, stop feeling real. But the numbers being invested in AI infrastructure right now deserve to be understood, because they represent the largest coordinated infrastructure buildout in human history.

What	The Number	Source
Total AI infrastructure spending projected by 2029	USD 1 trillion by 2029	IDC, 2025
Annual AI capex estimated for 2026 alone	USD 765 billion	Goldman Sachs Global Institute, 2026
Cumulative AI infrastructure investment, 2026–2031	USD 7.6 trillion	Goldman Sachs Global Institute, 2026
Global data center electricity consumption by end-2026	~1,000 TWh (equal to Japan's annual usage)	International Energy Agency, 2026
Number of national governments with AI infrastructure mandates	47 countries	American Action Forum, 2026

Let that sink in for a moment. USD 7.6 trillion over five years. That is more than the GDP of most countries. It is being spent on data centers, chips, power systems, and the networks, including fiber, that connect them all.

HFCL

The Stargate Story: The Largest AI Infrastructure Commitment in History

In January 2025, the United States announced a project called Stargate: a USD 500 billion commitment to build AI infrastructure on American soil. The project involves technology leaders, SoftBank, Oracle, and the backing of the White House. The flagship site is an 875-acre facility in Abilene, Texas, that will ultimately house enough computing power to run some of the most advanced AI systems in the world.

When complete, this single facility will draw 1.2 gigawatts of electricity, enough to power a city of nearly one million people. And it will need kilometers of high-performance fiber optic cable to connect its thousands of servers to each other and to the outside world.

Stargate is not an isolated story. It is the American chapter of a global story that is playing out simultaneously in the UAE, Japan, India, Saudi Arabia, and across Europe. Every country in the world that wants to participate in the AI economy is building infrastructure. And the backbone of that infrastructure is fiber.

The Data Center Is the New Factory

There is a useful mental model for understanding what is happening. In the Industrial Revolution, the factory was the fundamental unit of economic production. Whoever had the most advanced factories had economic power. Today, the data center, specifically the AI-optimized data center, is the new factory. Whoever has the most advanced data centers, connected by the fastest networks, will hold structural advantages in the AI economy.

This is why governments are treating data center infrastructure the same way they once treated steel mills and shipyards: as strategic national assets. And it is why the fiber that connects these facilities, within them, between them, and to end users, has become a matter of national interest.

The Energy Problem Nobody Warned Us About

Here is a fact that tends to stop people mid-sentence: a single rack of the latest generation of AI servers requires 120 kilowatts of electricity. Traditional data center racks were designed for 8–17 kilowatts. The new AI racks are roughly ten times more power-hungry.

The International Energy Agency reported in 2026 that global data center electricity consumption grew 17% in 2025 alone, during a year when global electricity demand overall grew just 3%. AI is not just demanding more computing power. It is demanding more energy at a scale that is straining power grids.

In the United States, regions like Northern Virginia, the self-declared "data center capital of the world", are facing grid stress that engineers have not seen before. The queue for large power connections in prime U.S. markets now runs three to seven years.

What This Means for Fiber

The energy crisis in AI infrastructure has an interesting and underappreciated connection to fiber. Traditional copper-based data center networking consumes significantly more power than fiber optic networking at comparable speeds. As data center operators look for every watt of savings they can find, and they are looking very hard, the shift from copper to fiber inside data centers accelerates.

Additionally, the push to build data centers in locations with cheaper and cleaner energy, Scandinavia for its hydroelectric power, Chile for its wind energy, and India for its combination of land, labor, and growing renewable capacity, means that long-haul and subsea fiber connections become even more critical. The data center may be built in a remote location for energy reasons, but it must be connected to users in dense urban centers. Fiber carries that distance.

HFCL

Sovereign AI: Why Countries Are Racing to Build Their Own Compute

In 2019, the idea of a government building its own AI data center the way it once built a dam or a highway would have seemed strange. In 2026, it is entirely normal. As of today, 47 national governments have active AI infrastructure investment mandates. Something fundamental has shifted.

What changed? Governments realized, slowly and then all at once, that AI is not just a technology product. It is a strategic capability. The country that controls advanced AI infrastructure controls something close to economic and, in some cases, military advantage. Depending on a foreign country's cloud servers to run your government's AI systems started to seem unacceptable, the same way as depending on a foreign country for your defense hardware seemed unacceptable.

This is the idea behind "sovereign AI": the principle that a country should have AI infrastructure on its own soil, under its own laws, within its own power grid.

- **What Sovereign AI Looks Like in Practice**

The UAE built a 1-gigawatt AI compute cluster in Abu Dhabi, partnering with leading global technology firms. Japan committed the equivalent of approximately USD 13 billion to semiconductor fabrication and 2nm chip R&D. Saudi Arabia's Public Investment Fund has committed USD 40 billion to AI infrastructure. India has multiple national-scale AI initiatives underway, with a proposed USD 100 billion renewable-powered AI data center platform.

Each of these sovereign AI projects requires the same fundamental ingredient as commercial AI infrastructure: high-performance optical fiber, running from the power source to the data center, from server rack to server rack, and from the data center to the citizens who use its services.

- **The Data Sovereignty Dimension**

Sovereign AI is not only about who owns the chips. It is also about where data lives. Regulations across the European Union, India's data protection frameworks, and emerging laws in the Middle East all restrict where citizen and enterprise data can be stored and processed. This is driving the construction of local, regional AI infrastructure rather than a single global system, and each regional hub requires its own fiber backbone. For a company like HFCL, which already exports to more than 60 countries, the sovereign AI trend is a direct commercial opportunity. The government's building sovereign AI infrastructure needs fiber. They need it manufactured locally or from trusted partners. And they need it fast.

HFCL

India's Moment: Why the World Is Looking East

India is simultaneously the world's most populous country, one of its fastest-growing economies, and one of the most underserved in terms of digital infrastructure relative to its size. That combination, huge demand, growing investment, and still large gaps to fill, creates an extraordinary opportunity window.

- **The BharatNet Foundation**

India's BharatNet program, the world's largest rural broadband project, has been laying optical fiber to connect India's 600,000 gram panchayats (village councils). This is the physical substrate that AI can eventually run on. But BharatNet is a foundation layer, not a finished product. The AI infrastructure buildout requires a second layer: high-capacity fiber connecting cities, data centers, 5G towers, and enterprise campuses at speeds and densities far beyond what earlier generations of telecom infrastructure required.

- **The 5G Multiplier**

India's rapid 5G rollout is the single most important near-term driver of fiber demand in the country. Here's the counterintuitive truth: 5G is wireless, but it runs on fiber. Every 5G tower needs a fiber backhaul connection, a high-capacity optical cable running from the tower back to the network core. India needs tens of thousands of these connections. HFCL, as the country's leading optical fiber cable manufacturer and a provider of 5G transport products, is directly in the path of this demand.

- **India as an AI Data Center Hub**

Global technology companies are increasingly looking at India not just as a market for AI products but as a location for AI infrastructure itself. The reasons are compelling: India has abundant land, a rapidly growing renewable energy sector, a large pool of engineering talent, and improving connectivity. Several hyperscale data center campuses are under development or planned in Mumbai, Chennai, Hyderabad, and Pune. Each of these facilities needs fiber: inside the campus connecting racks, within the city connecting the campus to the internet exchange, and across the country connecting the campus to users in tier-2 and tier-3 cities. The AI infrastructure buildout in India is a fiber buildout. They are the same story.

HFCL

What AI Infrastructure Means for Businesses Today

If you are running a business, whether you are a telecom operator, an enterprise CIO, a data center developer, or a government planner, the AI infrastructure buildout has specific, practical implications for you right now.

Your Network Is Probably Not Ready

The biggest finding from enterprise AI deployments in 2025 and 2026 is consistent: organizations that tried to run AI workloads on existing network infrastructure were disappointed. Latency was too high. Bandwidth was insufficient. The response times that feel fine for a web application feel terrible for a real-time AI assistant.

Multiple industry surveys have found that organizations consistently underestimate AI infrastructure costs, with actual deployment spending frequently exceeding initial budgets by a significant margin.

The lesson is clear: AI infrastructure must be designed for AI from the start, not adapted from what existed before. That means high-bandwidth fiber connections at every layer, from the server rack to the building, from the building to the data center, from the data center to the internet.

Real-Time AI Changes the Requirements

Early AI applications were mostly batch processes: you sent a query, waited a few seconds, and received a result. That was tolerable. But the AI applications gaining traction in 2026 are real-time: AI customer service that responds mid-conversation, AI-assisted medical diagnosis that works during a consultation, and AI-powered fraud detection that acts in milliseconds.

Real-time AI has a specific infrastructure requirement: ultra-low latency. And the only way to achieve ultra-low latency at scale is through fiber, specifically, fiber that connects computing resources as close to the user as possible (what engineers call edge computing). Deploying fiber to edge nodes, which is an active growth area in India and globally, is the infrastructure answer to real-time AI.

The Questions Every Infrastructure Decision-Maker Should Ask

- s our fiber backbone built for AI-level bandwidth, or was it designed for the internet of ten years ago?
- Do our data center facilities support the high-density fiber interconnects that modern AI server racks require?
- Have we planned for fiber to our 5G towers as part of our AI edge strategy?
- Are our last-mile connections fast enough to deliver AI applications to users in real time?
- Do we have a partner who can supply high-performance, tested, and certified fiber quickly as demand grows?

HFCL

OptiQ AI™: HFCL's Answer to the AI Connectivity Challenge

Everything discussed so far in this whitepaper, the 100,000-GPU clusters, the 120-kilowatt racks, the race toward 800G and 1.6T networks, the three-to-seven-year infrastructure queues, converges on a single practical question: who can supply the optical connectivity that AI data centres actually need, at the speed they need it?

In July 2026, HFCL answered that question by launching OptiQ AI™, a unified brand identity for its end-to-end optical connectivity portfolio purpose-built for AI, cloud, and hyperscale data centre environments.

Why a Unified Framework Matters

Hyperscalers building AI facilities today face a fragmented supply landscape. Cables come from one vendor, assemblies from another, enclosures from a third, each with its own specifications, lead times, and interoperability risks. When GPU clusters are scaling toward 100,000 units and every week of deployment delay carries an enormous opportunity cost, fragmentation is a liability.

OptiQ AI™ eliminates that fragmentation by bringing HFCL's proven data centre connectivity products under one cohesive, interoperable framework:

Product Category	Role in AI Infrastructure
Intermittently Bonded Ribbon (IBR) Cables	Ultra-high fibre density in constrained pathways, the backbone of GPU-to-GPU and rack-to-rack connectivity
Fibre Cable Assemblies	Pre-terminated, factory-tested links that cut installation time and eliminate field-splicing errors
Patch Cords and Pigtails	The high-performance final connections between equipment ports and the structured cabling plant
High-Performance Trunks	High-count backbone links between rows, halls, and meet-me rooms
High-Density Cassettes and Enclosure Panels	Organised, scalable fibre management that maximises connections per rack unit

The Five Pillars of OptiQ AI™

The portfolio is built around five attributes, each mapping directly to a pain point identified earlier in this paper:

1. High Quality.

AI clusters run continuously at full utilisation. A single degraded link can idle GPUs worth millions of dollars. Every OptiQ AI™ product is manufactured and tested to hyperscale-grade standards at HFCL's facilities.

2. Quantum Bandwidth

As AI networks migrate from 400G to 800G and onward to 1.6T, the passive fibre layer must not become the bottleneck. OptiQ AI™ solutions are engineered for seamless migration to 800G and 1.6T network architectures, meaning the fibre installed today will carry the workloads of 2030.

3. Densely Quantified

The 10X jump in rack power density described earlier has a fibre corollary: fibre counts per rack, per tray, and per duct are rising exponentially. IBR cables and high-density cassettes pack more fibres into less space, which is exactly what space-constrained, high-power AI halls demand.

4. Quick Rollout

With grid connections backlogged and AI demand compounding, deployment speed is a competitive weapon. Pre-terminated assemblies, plug-and-play cassettes, and an integrated single-vendor portfolio compress deployment cycles and reduce interoperability risk.

5. Q-Class Uptime

Real-time and agentic AI applications cannot tolerate link failures. OptiQ AI™ is designed to meet the reliability profile required by mission-critical AI infrastructure.

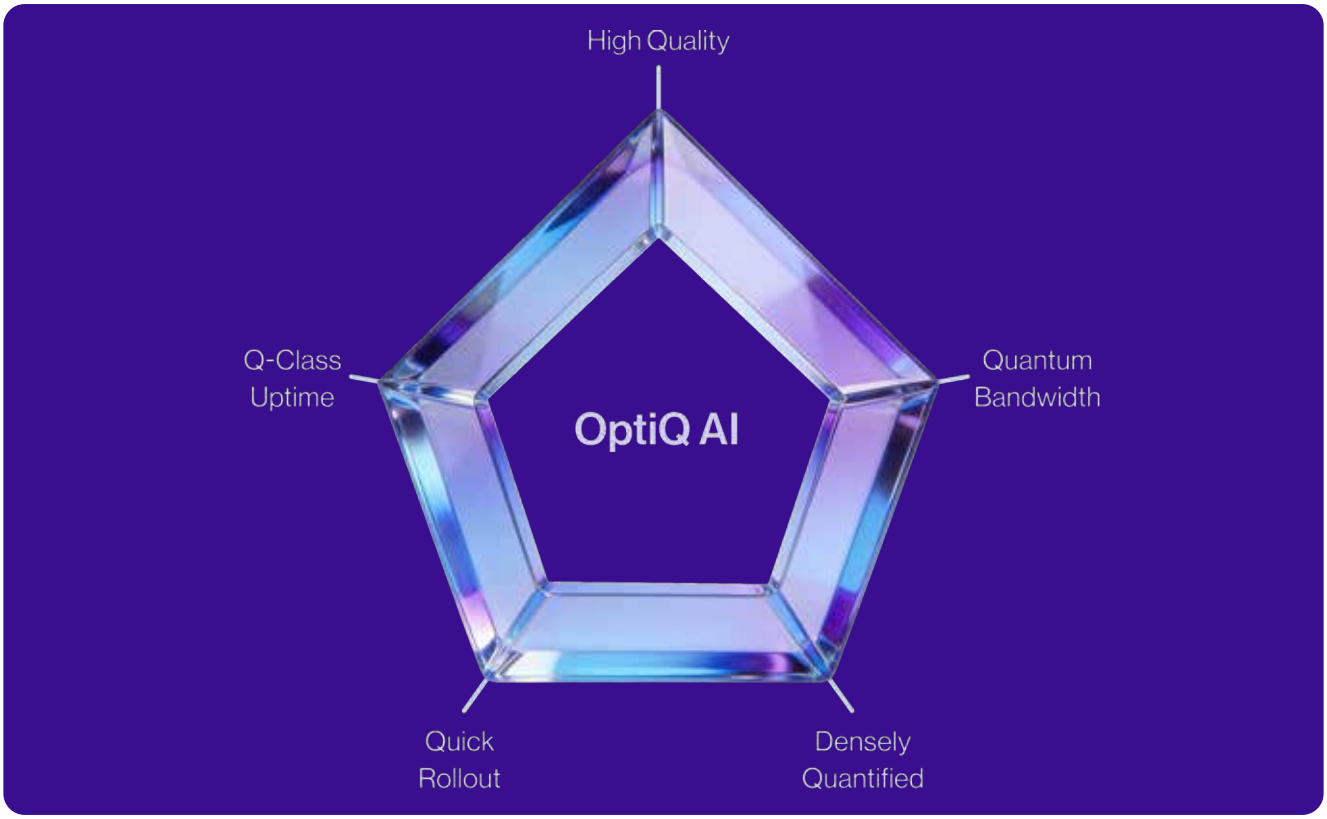

A Market Growing Fivefold

The timing is not incidental. Industry projections cited by HFCL's leadership indicate the AI optical market growing from roughly USD 14 billion today to USD 73 billion by 2030 (Bank of America Global Research, 2025), driven by investment in AI data centres, high-performance computing, and hyperscale cloud networks.

Critically, the global supply chain for these products remains concentrated in a small number of geographies. For hyperscalers and governments pursuing supply-chain resilience, and for the sovereign AI programmes described earlier in this paper, a competitively manufactured, high-quality alternative from India's resilient and agile supply chain is strategically valuable. OptiQ AI™ positions HFCL and India squarely within that opening.

What OptiQ AI™ Means for Different Buyers

- Hyperscalers and AI data centre developers get faster deployment cycles, reduced interoperability risk, and a fibre plant ready for 800G/1.6T migration.
- Telecom operators get a partner whose portfolio spans from the data centre core to the 5G tower to the last mile, one supply chain, one quality standard.
- Enterprise CIOs get hyperscale-grade connectivity practices scaled to campus and edge deployments.
- Government and sovereign AI planners get a trusted, domestically manufactured supply source aligned with data sovereignty and Make in India priorities.

HFCL

The Road Ahead:
Future-Proofing Your AI Infrastructure Strategy

Global Institute estimates USD 7.6 trillion in cumulative AI infrastructure investment between 2026 and 2031. Here are the specific trends within that story that will shape demand for connectivity and fiber over the coming years.

- **Agentic AI: The Next Wave of Compute Demand**

The AI that most people have experienced so far is conversational: you ask a question and get an answer. The next wave is called "agentic AI", AI that does things autonomously on your behalf. An AI agent might manage your calendar, coordinate a supply chain, or monitor a power grid, making decisions and taking actions without human intervention.

Industry analysts expect agentic AI to become a significant share of enterprise software spending by 2027 and 2028, with Gartner projecting that 33% of enterprise software applications will include some form of AI agents by 2028.

- **The Subsea Cable Surge**

A quieter but deeply significant story in AI infrastructure is the surge in subsea fiber cable investment. The connections between continents, the undersea highways that carry data between data centers in different countries, are being upgraded and expanded specifically to handle AI workloads. New subsea cable projects announced in 2025 and 2026 represent billions of dollars in investment across the Indian Ocean, the Pacific, and the Atlantic.

For India, this is particularly relevant. Submarine cable landings in Mumbai and Chennai are critical connection points for AI data flows between India, the Middle East, and Southeast Asia. Domestic fiber networks that connect these landing stations to data centers and enterprise users are the final link in a chain that stretches across oceans.

- **AI in Defense and Critical Infrastructure**

AI infrastructure is not only commercial. Defense organizations globally are building AI-enabled command and control systems, surveillance networks, and logistics platforms. These require secure, hardened fiber networks, often physically separated from commercial networks. HFCL's active presence in defense communications, including its recently approved INR 230 crore defense manufacturing facility in Andhra Pradesh, positions it to serve this growing segment.

- **The Efficiency Curve**

One important counterpoint to the "unlimited growth" narrative of AI infrastructure: AI systems are becoming more efficient over time. Models that once required enormous compute clusters can now run on smaller ones. This is a healthy dynamic; it makes AI more accessible, not less demanding. The Jevons Paradox applies here: as AI becomes cheaper to run, more applications get built, and total demand goes up, not down. Efficiency gains drive adoption, and adoption drives infrastructure demand.

Building AI Infrastructure Starts with the Right Connectivity Partner

For over 35 years, HFCL has been connecting India and the world, from the first generation of telecom networks to the fiber backbones that carry India's internet today. As the AI infrastructure buildout accelerates, HFCL is positioned at the center of the story with the introduction of OptiQ AI™, the unified identity for HFCL's data centre optical connectivity portfolio, which brings together IBR cables, fibre cable assemblies, patch cords, pigtails, trunks, high-density cassettes, and enclosure panels, engineered for the bandwidth, density, and deployment demands of the AI era. Alongside OptiQ AI™, HFCL's broader portfolio spans FTTH and FTTB solutions for last-mile connectivity, underground and aerial cables for backbone networks, and 5G transport products. Whether you are building a hyperscale AI data centre, deploying 5G towers, or upgrading your enterprise campus to handle AI workloads, HFCL has the answer.

To learn more about OptiQ AI™ and HFCL's optical fibre and networking solutions for AI infrastructure, visit hfcl.com or speak with our team.

HFCL

About HFCL

HFCL is a leading technology company specialising in creating digital networks for telcos, enterprises, and governments. Over the years, HFCL has emerged as a trusted partner offering sustainable high-tech solutions with a commitment to providing the latest technology products to its customers. Our strong R&D expertise, coupled with our global system integration services and decades of experience in fibre optics, enable us to deliver innovative digital network solutions required for the most advanced networks.

The Company's in-house R&D Centres located at Gurgaon & Bengaluru, along with invested R&D houses and other R&D collaborators at different locations in India and abroad, innovate a futuristic range of technology products and solutions. HFCL has developed capabilities to provide premium quality Optical Fibre and Optical Fibre Cables, state-of-the-art telecom products including 5G Radio Access Network (RAN) products, 5G Transport Products, Wi-Fi Systems (Wi-Fi 6, Wi-Fi 7), Unlicensed Band Radios, Switches, Routers, and Software Defined Radios.

The Company has state-of-the-art Optical Fibre and Optical Fibre Cable manufacturing plants at Hyderabad, an Optical Fibre Cable manufacturing plant in Goa, and its subsidiary, HTL Limited, at Chennai.