Netherlands Enterprise Agency

Near term Photonic applications in Datacom and Telecom with impact on reducing energy demand

Commissioned by the ministry of Economic Affairs and Climate Policy

>> *Sustainable. Agricultural. Innovative. International.*

0. Foreword

The downside of a highly digitizing society is, among others, the high energy consumption of data centers which is increasingly becoming a bottleneck.

One of the solution perspectives mentioned for this is to increasingly use photonics for data transmission.

For network cables that serve to interconnect data centers, this is already commonplace. The next big step can be made by also using photonic integrated chips instead of electronic chips. Such a photonic integrated chip has no heat generation that causes energy loss as in an electronic chip. As a result, the air conditioning equipment in a data center would not have to cool as much as is does today either.

The Netherlands is showing great ambition in this field of integrated photonics by allocating a billion-dollar plan for this in its National Growth Fund. Recently, the PhotonDelta consortium received an award for this purpose to invest €1.2bn in a Dutch production line. This focuses on four application areas, of which data and telecommunications are undisputedly number one in terms of opportunities and challenges.

The question arises: are these applications already market-ready and can we pioneer them in the Netherlands?

The research presented here shows that the urgency is great, that fiber networks already reach deep into the data centers, but that the application of photonics in the servers themselves is still in its infancy. The photonic chips that should make a difference are still too expensive and unproven for large scale adoption by data centers. Waiting for the price curve to go down will lead to the loss of the initiative in the design and production of these photonic chips as well as a delay in making our strained data infrastructure more energy efficient.

One of the considerations LEAP would like to pass on to the government and other parties involved in the value chain is to put in place a highly focused "man-on-the-moon" program to design this set of photonic chips for data centers and take them into production in leading Dutch data centers now. Thus positioning the Netherlands as a pioneer in this field, allow it to be the first to reap the benefits of an energy-efficient data infrastructure and establish a new economic sector around photonics for a new generation in 2030.

1. Summary

Optical fiber networks are widely used in the telecom and datacom space due to their ability to efficiently transport large volumes of data over long distances. Photonics aim to do the same but now on a micro and nano scale. The key challenge is to be able to pack miniaturized optical and electronic systems so tightly together on one chip, that they truly become one integrated unit operating at very high speeds.

Theoretical scientists have proven it possible, engineers and entrepreneurs have made the first products. The race is now literally about the last inch. But the rewards are huge, every millimeter may be worth more than a billion euro!

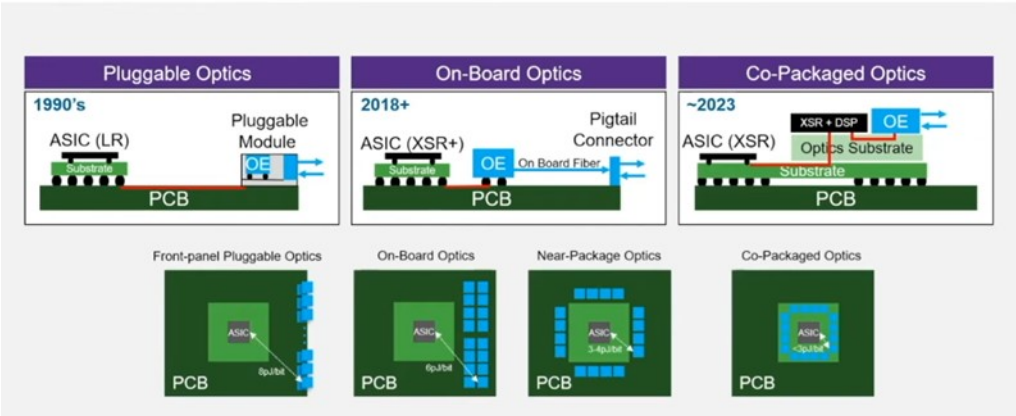

Figure 1. The last inch and its rewards : the impact of co-packaged optics on energy needs (source: Semiconductor Engineering)

In fact the expectation is that, over time, a tenfold reduction in power consumption could be achieved. A similar result has been achieved in the past decade in long-haul networks.

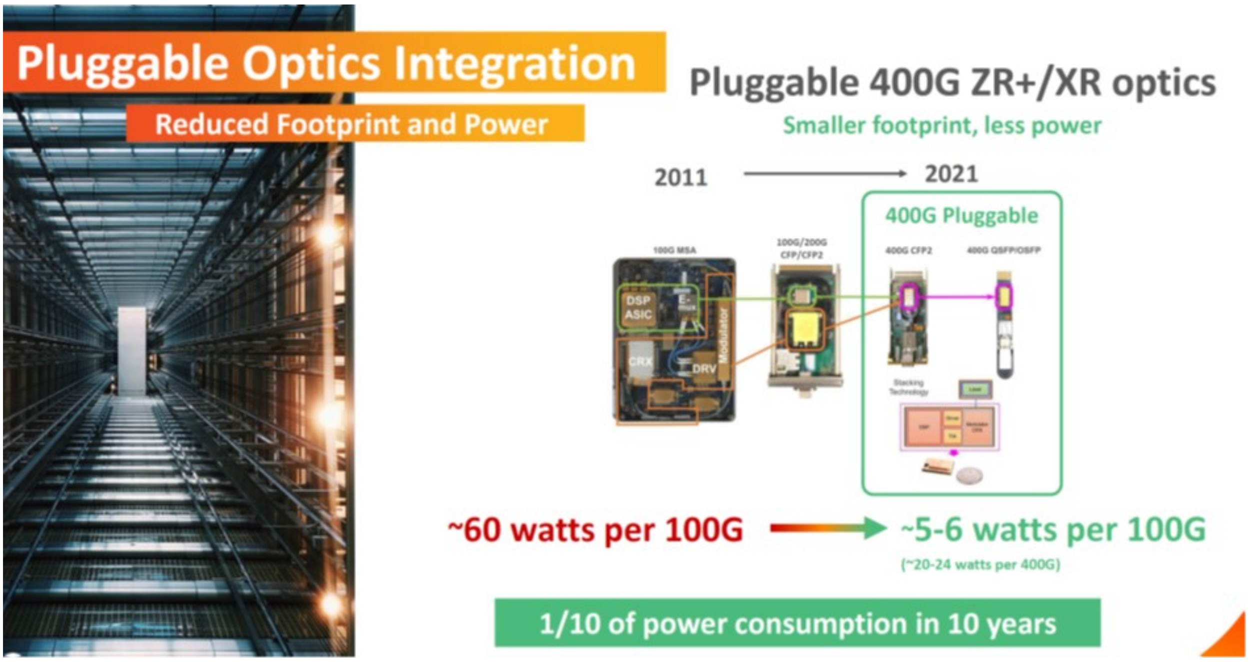

Figure 2. Historic power savings achieved in long-haul transport networks – Source: Infinera

In the next few years the following photonic products can be expected to become commercially available at scale:

Table 1: Near term photonic applications

What (Product type)	When (1st introduction)	Who (Vendor)	So what (Impact)
Photonic Transceivers	As of 2020	Various including EFFECT Photonics, Intel and Broadcom (see table 2 and 3)	Could save 50% of power used for data transport per year and prolong Moore's Law.
Photonic Switches	As of 2020	TU Eindhoven, Intel (through an acquisition of Barefoot Networks) and Broadcom	Enormous network capacities and power savings of 1 MWh per year per switch.
Photonic AI/ML accelerators	As of 2024 – 2025	Lightmatter https://lightmatter.co/	Could save 20% of power used for AI/ML applications per year
Photonic processors	As of 2024 – 2025	Lighttelligence https://www.lightelligence.ai/	Could save 20% of power used for computations per year

Photonic transceivers in particular are becoming available from a wide selection of vendors for telecom as well as for datacom use cases as illustrated by tables 2 and 3 and figure 2.

Table 2: Suppliers of Photonic Transceivers for Telecom Networks(Long-haul/DCI, Mobile Front/Backhaul & Metro), >10 km*

Company Name	Website	Speeds current products	Announced
1. EFFECT Photonics	https://effectphotonics.com/	10G	25G, 40G, 100G(2023)
2. Intel	https://www.intel.com/	100G	
3. Lumentum	https://www.lumentum.com/	10G,100G,200G, 400G	
4. Nokia	https://www.nokia.com/networks/	100G, 200G, 400G, 600G	
5. Cisco	https://www.cisco.com/	25G, 40G, 100G,400G	
6. Source Photonics	https://www.sourcephotonics.com/	10G,25G,40G,50G,100G	
7. Fiber Mall	https://www.fibermall.com/	400G,800G	
8. Molex	https://www.molex.com/	100G,400G	
9. II-VI	https://ii-vi.com/	40G,100G,400G	
10. Innolight	https://www.innolight.com/	10G,25G,40G,100G,200G, 40G,800G	
11. Infinera	https://www.infinera.com/	400G, 800G	1600G(2023)

* DCI = Data Center Interconnect

Table 3: Suppliers of Photonic Transceivers for Datacom (Intra-Datacenter connections & Campus Networks) 500m – 2km

Company Name	Website	Speeds current products	Announced
1. Intel	https://www.intel.com/	100G, 200G, 400G	800G(2023)
2. Cisco	https://www.cisco.com/	25G, 40G, 100G,400G	800G (2024)
3. Source Photonics	https://www.sourcephotonics.com/	40G,100G,400G,800G	
4. Fiber Mall	https://www.fibermall.com/	400G,800G	
5. Molex	https://www.molex.com/	100G,400G	
6. II-VI	https://ii-vi.com/	40G,100G,400G	
7. Broadcom	https://www.broadcom.com/	100G*,400G,800G*	1600G* (2024)
8. Innolight	https://www.innolight.com/	10G,25G,40G,100G,200G, 40G,800G	

*Chipset only

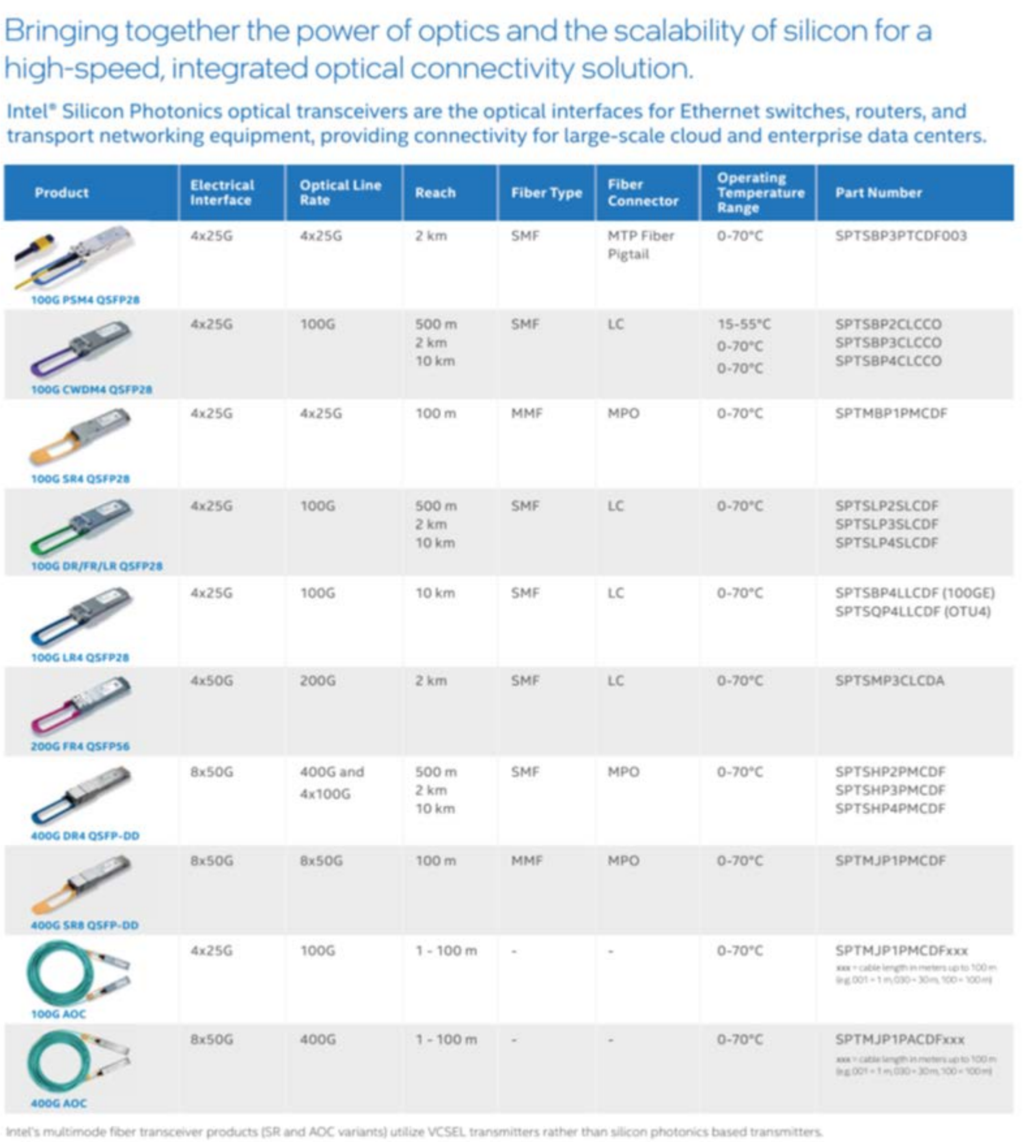
Bringing together the power of optics and the scalability of silicon for a high-speed, integrated optical connectivity solution.

Intel® Silicon Photonics optical transceivers are the optical interfaces for Ethernet switches, routers, and transport networking equipment, providing connectivity for large-scale cloud and enterprise data centers.

Product	Electrical Interface	Optical Line Rate	Reach	Fiber Type	Fiber Connector	Operating Temperature Range	Part Number
100G PSM4 QSFP28	4x25G	4x25G	2 km	SMF	MTP Fiber Pigtail	0-70°C	SPTSBP3PTCDF003
100G CWDM4 QSFP28	4x25G	100G	500 m 2 km 10 km	SMF	LC	15-55°C 0-70°C 0-70°C	SPTSBP2CLCCO SPTSBP3CLCCO SPTSBP4CLCCO
100G SR4 QSFP28	4x25G	4x25G	100 m	MMF	MPO	0-70°C	SPTMBP1PMCDF
100G DR/FR/LR QSFP28	4x25G	100G	500 m 2 km 10 km	SMF	LC	0-70°C	SPTSLP2SLCDF SPTSLP3SLCDF SPTSLP4SLCDF
100G LR4 QSFP28	4x25G	100G	10 km	SMF	LC	0-70°C	SPTSBP4LLCDF (100GE) SPTSQP4LLCDF (OTU4)
200G FR4 QSFP56	4x50G	200G	2 km	SMF	LC	0-70°C	SPTSMP3CLCDA
400G DR4 QSFP-DD	8x50G	400G and 4x100G	500 m 2 km 10 km	SMF	MPO	0-70°C	SPTSHP2PMCDF SPTSHP3PMCDF SPTSHP4PMCDF
400G SR8 QSFP-DD	8x50G	8x50G	100 m	MMF	MPO	0-70°C	SPTMJP1PMCDF
100G AOC	4x25G	100G	1 - 100 m	-	-	0-70°C	SPTMJP1PMCDFxxx xxx = cable length in meters up to 100 m (e.g. 001 = 1 m, 030 = 30m, 100 = 100 m)
400G AOC	8x50G	400G	1 - 100 m	-	-	0-70°C	SPTMJP1PACDFxxx xxx = cable length in meters up to 100 m (e.g. 001 = 1 m, 030 = 30m, 100 = 100 m)

Intel's multimode fiber transceiver products (SR and AOC variants) utilize VCSEL transmitters rather than silicon photonics based transmitters.

Figure 3. Examples of available photonic transceivers for Single Mode Fiber – Source: Intel

Deployment of photonic transceivers has been gaining traction in the past few years, particularly in the Telecom domain. In the Netherlands KPN, VodafoneZiggo and EuroFiber are examples of network operators that have deployed such products in their long-haul, Metro and Datacenter Interconnect networks. Hyperscale Datacenter operators in the USA such as Google and Amazon are also starting to deploy co-packaged optic transceivers in their campus networks. No Dutch examples were found.

The energy consumption related to running the IT workload in datacenters is typically referred to as the 'useful energy consumption' while the remaining portion is referred to as 'waste' or 'overhead' consumption. Because in datacenters the servers consume the bulk of the useful energy, the biggest energy savings in datacenters will be realized when the chips that power servers become photonics based, i.e. when photonic processors and AI accelerator chips become available.

It is important to note that due to the exponential growth in traffic, the savings are not likely to be in the form of an overall lower energy consumption in absolute terms but rather in relative terms, that is in the form of flat or flattish total energy consumption despite a steep growth in traffic volume.

Another point to note is that typically a significant portion of the energy used in datacenters can be lost on the physical equipment infrastructure including power equipment. This is due to built-in safety margins to ensure availability and reliability. Application of best practices in operations and planning can minimize that waste by up to 35%.

Table 4: Best practices for power savings

Best Practice	Description
1. Energy efficient Data Center Design	Make use of latest energy efficiency insights when building a new data center. Modular design allows for energy efficient growth.
2. Retiring systems	Conduct regular system reviews. Remove systems that are no longer in use. ('orphans')
3. Activating power management features	New servers are equipped with power management features. Should be activated.
4. Deploying blade servers	Replace old servers with new blade servers. These are more efficient and have power management and accurate power consumption reporting features
5. Virtualization	Virtualization reduces overall number of servers
6. Standardization	Standardize on 2 blade types
7. Use an effective server migration strategy	Limit the number of servers by using high density servers.
8. Right-sizing the physical infrastructure system to the load	Ensure capacities are in line with what is required for load. Don't heavily over dimension.
9. System design	Ensure best design for desired performance

2. Introduction

Data usage is increasing rapidly and with it the related energy consumption. To sustain the required growth rate in the long run a shift from electronic to photonic data transmission technology is required. The Dutch government sees photonics as one of the pillars on which the energy transition can rest as well as a strategic opportunity for sustainable economic growth.

In April 2022 the Dutch government through its National Growth Fund has announced the creation of a Public Private Fund totaling more than €1 billion for photonic research and the development of the photonics ecosystem in the Netherlands. The aim is to position the Netherlands as a frontrunner and world leader in photonics.

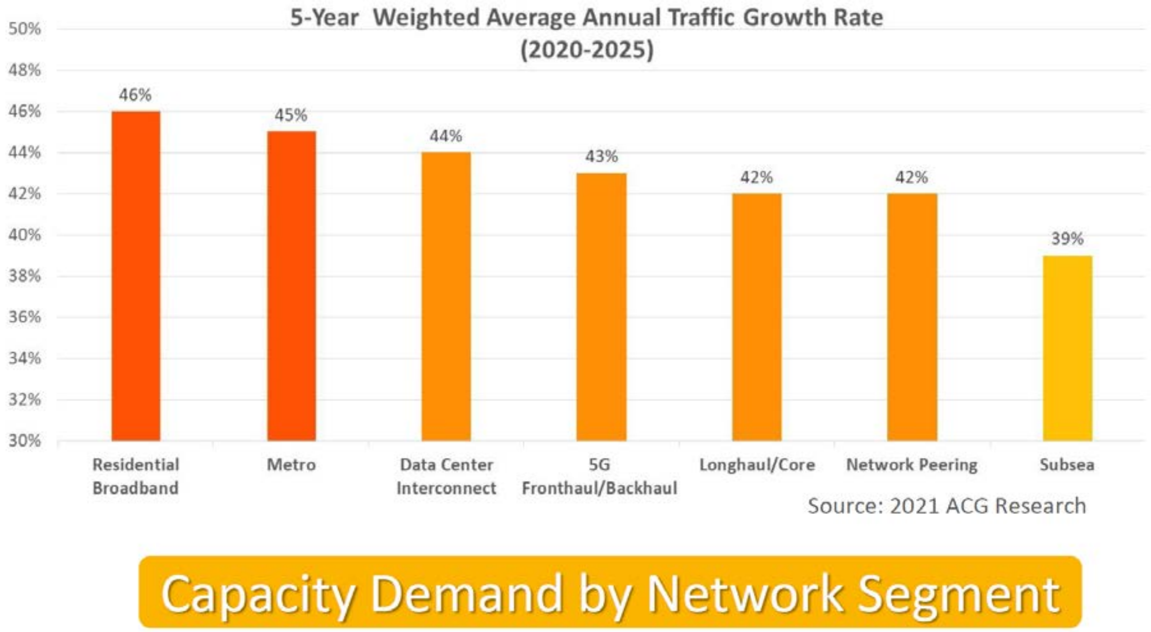

Figure 4. Data traffic growth projections – Source: Infinera

This document summarizes the findings of a research commissioned by platform LEAP of the Amsterdam Economic Board and RVO Nederland to investigate the possible applications of photonics in the datacom and telecom domain that are currently commercially available or expected to become available in the coming two years. The investigation was conducted by means of a desk research and interviews with representatives of various organizations in the datacom/telecom value chain.

3. Electronics vs Photonics

Data transmission is the transfer of data from one digital device to another. This transfer occurs via point-to-point data streams or channels. These channels can be in the form of copper wires, fiber optic cables or be part of a wireless network.

The transmission occurs by the encoding and decoding of digital signals using some kind of modulation technique. The effectiveness of data transmission relies heavily on the amplitude and transmission speed of the carrier channel. The amount of data transferred within a given time period is the data transfer rate, which specifies whether or not a network can be used for tasks that require complex, data-intensive applications.

Traditional digital signals are sent electronically, i.e. by way of transmission of electrons between high and a low voltage points. This requires energy. The required amount of energy grows exponentially with the increased volume and speed with which the data is transferred electronically.

Also, as the electrons travel they collide with the ions that form the solids in which they travel. This results in the electrons transferring the energy gained from the electric field to the positive charged ions with which they collide causing them to vibrate more vigorously. This vibration raises the kinetic energy of the ions and thus the temperature of the solid. Too much heat can cause the electronic devices to malfunction, burn or even explode. Hence the need of electronic devices for a cooling system, fans typically, which in turn requires energy.

A superior way for digital signal transfer is to do so optically instead of electronically, thus using light i.e. photons instead of electrons. The principles of digital data transmission remain in essence the same but now higher data rates are achieved by using lasers with different wavelengths (colors). Such technology is superior because it provides much more capacity, allows for much higher transmission rates and uses significantly less power.

Although the concept of optical transmission is not new, in fact it can be traced back to the centuries old practice of using sunlight and mirrors to signal messages, recent technological developments have made it possible to apply it on a microscale and to produce Photonic Integrated Circuits (PIC) lithographically on silicon wafers thus allowing the mass production of PICs at affordable costs.

4. The Dutch Photonics Ecosystem

In the Netherlands the photonics ecosystem is organized and coordinated through PhotonDelta a non-profit membership based organization representing a consortium of European industry partners and research centers with the aim of acting as an Industry accelerator through cooperation and co-creation. It offers a holistic eco-system that covers the photonics value chain from end-to-end.

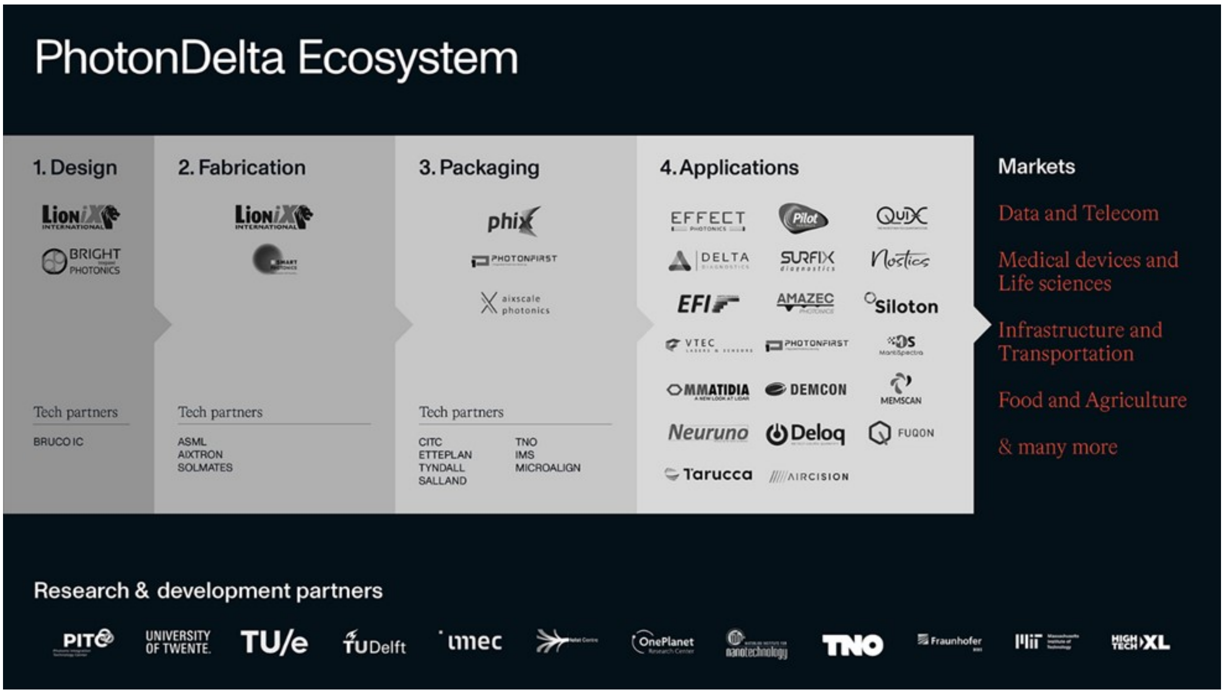

Figure 5. The PhotonDelta eco-system (source: PhotonDelta)

This consists of 4 main domains:

1. Design: The design and testing of PICs is tailored to meet the specific requirements of the application for which the PIC is intended to be used for. The specification is translated into a circuit design and then a mask layout is created for manufacturing. A design library containing highly functional predefined and tested building blocks is employed to reduce design costs.

2. Fabrication: PIC manufacturing facilities are called foundries. Depending on the required characteristics PICs can be produced based on 3 different platforms: Indium Phosphide (InP), Silicon Nitride (SiN) and Silicon Photonics (SiPh) PICs. InP PICs have active laser generation, amplification, control, and detection. This makes them an ideal component for communication and sensing applications. SiN PICs have a vast spectral range and ultra-low-loss waveguide. This makes them highly suited to detectors, spectrometers, biosensors, and quantum computers. SiPh PICs provide low losses for passive components like waveguides and can be used in minuscule photonic circuits. They are compatible with existing CMOS (electronic) fabrication.
3. Packaging: PICs are packaged for protection and mounting on a circuit board. Photonic chips and electronic chips can be combined (co-packaged) in a hybrid module, leveraging the benefits of both: cost reduction and best-in-class processes. It is also possible to configure different combinations of photonic chip types to address the demand for increasingly sophisticated solutions.
4. Applications: PICs can be used in a wide range of applications and industries. This includes applications for Food & Agriculture, Medical & Health, Engineering & Transport and the main focus area of this research: Datacom & Telecom.

Members of the PhotonDelta ecosystem that are active in the telecom and/or datacom industry include:

- **Bright Photonics:** https://brightphotonics.eu/; an independent design house for Photonic Integrated Circuits (PIC) in Silicon, III-V, SiN, Silica and Polymers. Bright Photonics supports companies and researchers in PIC development from application idea to prototypes and design for volume production.
- **Lionix International**: https://www.lionix-international.com/; a leading global provider of customized microsystem solutions. Specialized in photonic integrated circuits and MEMS since 2001. Lionix International is a vertically integrated company, working across all stages of the production process from design to delivery of a finished module.
- **Smart Photonics**: https://smartphotonics.nl/; a Pure-play InP foundry, producing only photonic components (discrete and ICs) for customers based on customer specified design.
- **Phix**: https://www.phix.com/; offers assembly services and contract manufacturing for photonic integrated circuits (PICs) and MEMS. Phix builds optoelectronic modules based on all major PIC technology platforms, such as Indium Phosphide, Silicon Photonics, Silicon Nitride, and Planar Lightwave Circuit. The company is specialized in chip-to-chip hybrid integration, coupling to fiber arrays, and interfacing of DC and RF electrical signals.
- **Quix Quantum**: https://www.quixquantum.com/; a start-up fabless, horizontally integrated company, focusing on quantum computing using integrated photonics. The company was founded in January 2019 by Dr. Hans van den Vlekkert, a veteran of the photonics industry and serial entrepreneur, Dr. Jelmer Renema, an expert in quantum photonics, and a team of professors from the University of Twente.
- **Aircision:** https://www.aircision.com/; a producer of telecom grade Free space Optics (FSO) solutions, providing point to point FSO links with speeds of 10Gbps – 100Gbps that can cover distances of 2.5 km – 5 km.
- **VTEC Lasers & Sensors**: https://www.vtec-ls.nl/; an innovation solutions company specializing in optics and data integration to create customized systems for a wide variety of industries, including telecom and datacom.
- **Effect Photonics**: https://effectphotonics.com/; delivers highly integrated optical communications products based on its Dense Wavelength Division Multiplexing (DWDM) optical System-on-Chip technology.

5. Datacom and Telecom network architecture

Datacom and telecom are interconnected networks and form an intricate global mesh of billions of connected nodes and devices. Though increasingly intertwined they are still distinct and separate networks. They are subject to different regulatory frameworks, are dominated by different players and have different business models and value chains.

The datacom network is popularly referred to as 'the cloud' and indeed cloud services are at the core of the datacom services.

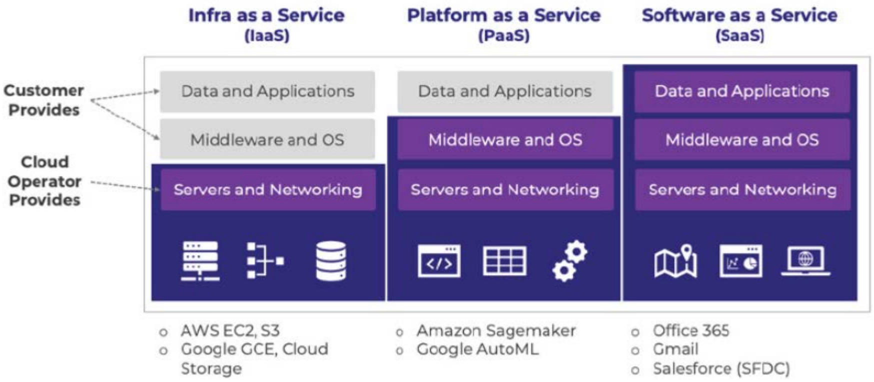

Figure 6. Cloud service models (source: Extreme Networks)

Cloud service providers vary from providers of cloud infrastructure only providers (Infra as a Service or IaaS model) to full stack cloud providers that provide Software as a Service (SaaS model).

IaaS providers offer compute, storage, networking, security and other infrastructure services in a hosted offering on a usage basis while their customers assume responsibility for building everything else on top of that themselves. SaaS providers provide complete applications as a ready to use product on a subscription basis. A model that is halfway these two extremes is called Platform as a Service, in this service model a certain platform for example Microsoft Azure is provided as a platform on which the customer can build and run his own applications.

Cloud services are physically located in and operated out of data centers and can be provided on a global, regional or local basis.

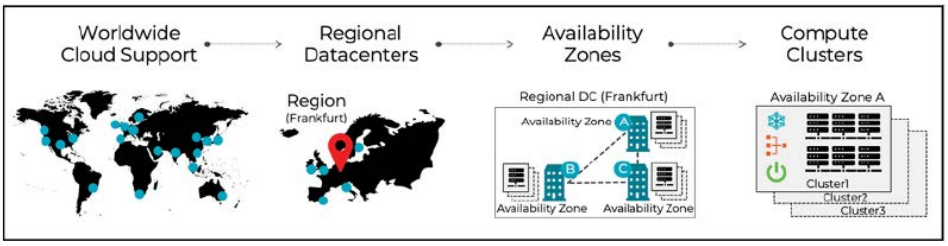

Figure 7. Geographical cloud architecture (source: Extreme Networks)

Data centers can serve one single or very many end-users. In their most basic form data centers simply offer facilities services i.e. rack space rental and utility services.

This is the most commonly used business model for co-location data centers. Hyperscale data centers are usually operated on a global basis and offer either PaaS or SaaS services.

	SINGLE TENANT	COLOCATION DATA CENTER			SINGLE TENANT
Type	PRIVATE/ENTERPRISE including server spaces	REGIONAL	NATIONAL	INTERNATIONAL	HYPERSCALE
Customers	SME Enterprise Public Semi-public	SME Public Semi-public	SME Enterprise Cloud Public Semi-public	SME Enterprise Cloud SaaS	Cloud SaaS
Space Impact	> 10 m2 Small	> 2000 m2 Small	> 2000 m2 Small	> 5000 m2 Medium	> 5 ha Medium/Large
Energy Impact	0,01 - 10 MW Small	0,5 - 10 MW Small	1 - 10 MW Medium	> 5 MW Medium	> 50 MW Large
Location (approximity)	In every province	In every province	In (regional) hubs	Amsterdam hub	Outskirts

Figure 8. Types of data centers (source: DDA)

The architecture of a typical data centers consists of a 3 tier structure. The servers that host the actual cloud services are organized in racks, the servers of each rack connect to one switch known as a Top of Rack (ToR) or Edge switch. The traffic from several Edge switches is aggregated in an interconnecting aggregation switch that in turn aggregate to core switches. The core switches form the interface to external networks such as the other data centers or an internet exchange.

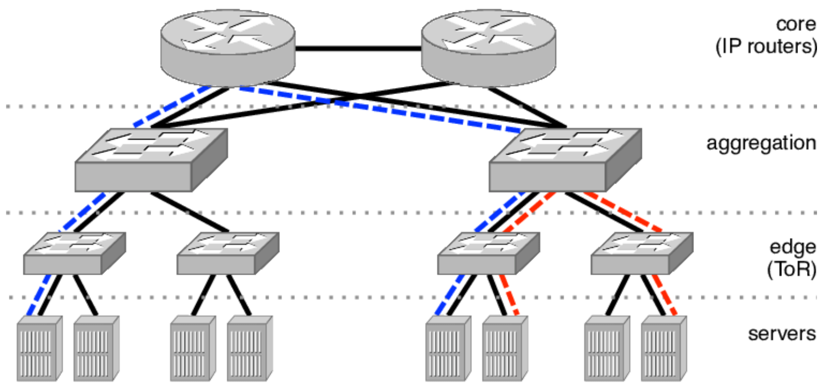

Figure 9. Generic data center typology (source: Research Gate)

This is where the datacom domain intersects with the telecom domain. Specialized telecom carriers typically provide the long-haul network that interconnects the global data centers between themselves as well as provide the connections to the various national telecom carriers.

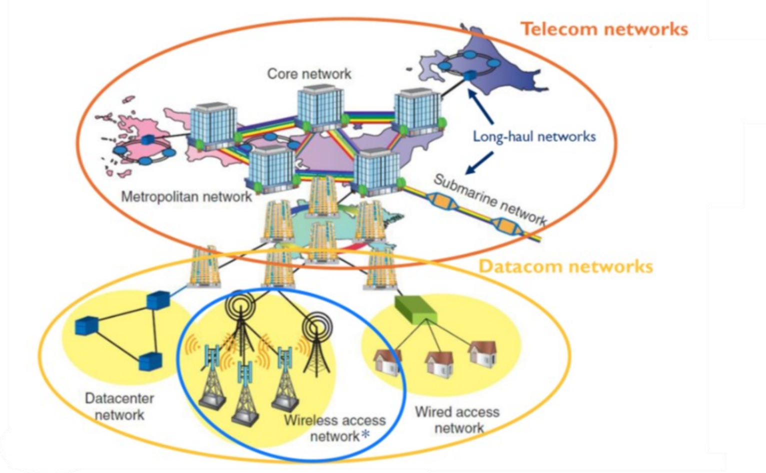

* Note: Wireless Access Networks are part of the Telecoms domain

Figure 10. Datacom and Telecom (source: Yole)

The Telecom network architecture distinguishes a hierarchy of three network segments:

1. an access network segment, sometimes referred to as Local Area Network (LAN) or last mile connection that connects premises to central offices with landlines or wirelessly connects mobile devices with cell towers;
2. a Metropolitan Area Network (MAN), sometimes referred to as city rings, that interconnects central offices in a given region or a mobile back-haul that connects a number of cell towers to a mobile network core node;
3. a long-haul network or Wide Area Network (WAN) that interconnects metro networks or mobile core nodes.

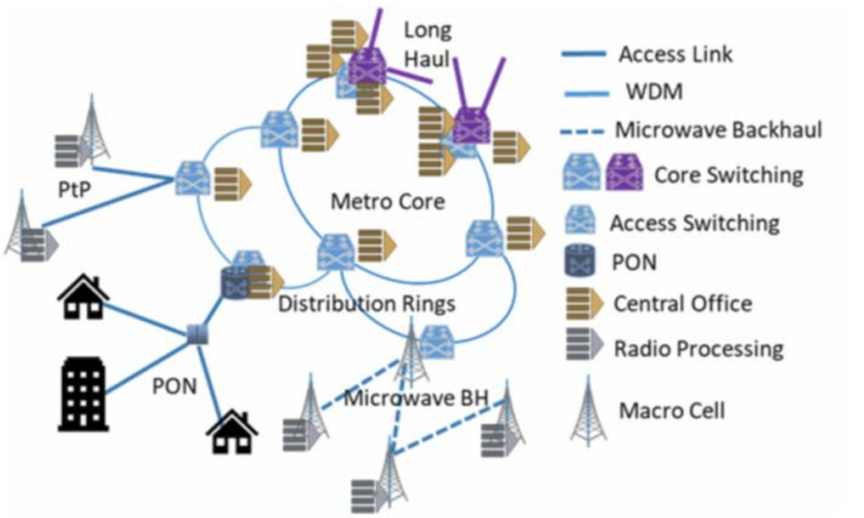

Figure 11. High level Telecom Network architecture (Source: IEEE)

6. Optical Technologies and Technological Developments

6.1 Optical networking in transport network

Over the last decades, many cable connections have become optical. Particularly long distance connections such as intercontinental submarine cables and connections between datacenters have long been making use of fiber optic technology.

Within data centers, cross connects are also almost all fiber based; in the Netherlands anyway.
Notable current exceptions are the cables connecting base stations or base band units to antennas in mobile networks and copper based coax cable in cable television networks.

An optical network is a communication system that uses light signals, instead of electronic ones, to send information between two or more points. Optical networks comprise optical transmitters and receivers, fiber optic cables, optical switches and other optical components.

Optical and electronic networks can take several different forms. Point-to-point networks make permanent connections among two or more points so any pair of nodes can communicate with each other; point to multipoint networks broadcast the same signals simultaneously to many different nodes; switched networks like the telephone system include switches that make temporary connections among pairs of nodes. The basic building blocks of these networks are fiberoptic cables-the so-called “pipes”-which carry signals from node to node, with switches directing them to their destination.

The Signal

An optical signal consists of a series of pulses produced by switching a laser beam off and on. Its speed depends on how fast the beam can be switched on and off, and how much the pulses spread in length during transmission, an effect called dispersion. The amount of dispersion depends on the type of fiber, the fiber length and the nature of the optical signal. The more dispersion, the more difficult it is to distinguish between adjacent pulses. Different types of fiber can be combined to reduce dispersion effects. A single fiber can transmit many separate signals simultaneously at different wavelengths of light, a technique called wavelength-division multiplexing.

The maximum number of optical channels is limited by the slice of spectrum used for each channel and the total amount of spectrum available. Devices called “demultiplexers” separate the optical channels and distribute them to separate optical receivers. Demultiplexers slice the spectrum into very narrow chunks, isolating each optical channel from adjacent ones. Multiplying the number of optical channels by the data rate on each optical channel gives total transmission capacity of a fiber.

Achieving high data rates and multiple channels requires sophisticated components. Semiconductor lasers (which generate the light pulses used in almost all fiber optic communications systems) must emit only a very narrow range of wavelengths to limit dispersion. Fibers are also designed to limit dispersion.

Optical Amplifiers

The optical fibers can transmit signals more than 100 kilometers without amplification. When the signal must span a longer distance, it is passed through an optical amplifier, which multiplies the strength of the optical signal. The most widely used optical amplifiers are fibers doped with atoms of erbium (Erbium Doped Fiber Amplifiers or EDFAs), a rare-earth element that absorbs light energy from an external pump laser. The erbium atoms then release that energy to amplify weak optical signals across the entire band of wavelengths that the laser transmits. With careful control, a string of dozens of optical fiber amplifiers can transmit signals thousands of kilometers across the ocean.

Another type is the Semiconductor Optical Amplifier (SOA). A semiconductor optical amplifier (SOA), also known as a laser amplifier, is an active medium of a semiconductor laser. An electric current is externally applied to the laser device that excites electrons in the active region. When photons travel through the active region it can cause these electrons to lose some of their extra energy in the form of more photons that match the wavelength of the initial ones. Therefore, an optical signal passing through the active region is amplified and is said to have experienced optical gain.

Optical Switches

One challenge to optical networking is how to switch light signals. When a signal arrives at its destination, it must be separated from the rest of the channels. To drop one signal at an intermediate point, an optical filter separates the proper wavelength from the rest. Equipment at that point may also add a new signal to the now unoccupied wavelength.

Optical switches may operate on a single wavelength, or on all the wavelengths transmitted through a fiber. A fixed filter corresponding with a single wavelength could be replaced by a switch that selects one of several filters to divert the desired wavelength to the intermediate point.
A third kind of switch separates the wavelengths into separate beams, and a moving mirror directs one or more of the wavelengths in a different direction.

Other optical switches simultaneously switch all wavelengths passing through a fiber; one example is a mirror at the fiber output that could tilt between two different positions to reroute all optical channels in case of a fiber break. This type of optical switch is called an optical circuit switch.

The preceding examples are called “all-optical” (OOO) switches because they operate on light signals only. A different class of switches convert optical signals into an electronic form which can be switched electronically; the resulting electronic signal then feeds into an optical transmitter to generate a new optical signal. These are called opto-electro-optical (OEO) switches. This conversion process results in a certain signal latency. Also, the required transponders are tied to specific traffic formats and data rates, they would thus need to be replaced if these transmission rates change over time.

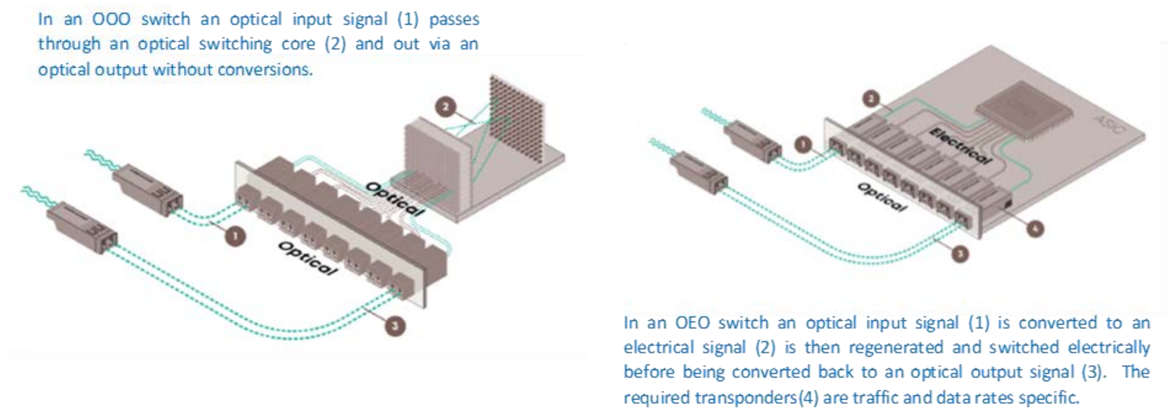

Figure 12. an OOO and an OEO switch. (Source: HUBER + SUHNER Polatis)

Wavelength Division Multiplexing

Wavelength Division Multiplexing (WDM) is a process in which a prism splits light into different wavelengths, which could travel through a fiber simultaneously. The peak wavelength of each beam is spaced far enough apart that the beams are distinguishable from each another, creating multiple channels within a single fiber strand.

Due to the physical properties of light, no channel can interfere with the next – they are completely separated from each other. Each channel is transparent to the speed and type of data, meaning that any mix of voice, data and video services can be transported simultaneously over a single fiber WDM system.

Optical Transceivers

Optical transceivers are devices capable of transmitting and receiving data used to transport high levels of data traffic over a network. Transceivers contain lasers that are wavelength specific and they convert electrical data signals into optical signals for transmission over optical fiber and vice versa. Once this process has been performed the data signals can then be transmitted over optical fiber. Each stream of data is transformed into a signal with a unique wavelength.

Transceivers are an element in a WDM system. In its most basic form a WDM system consists of 4 elements: transceivers, a multiplexer, patch cords connecting the transceivers to the multiplexer and an optical fiber cable.
There are three main categories of transceivers, grey (standard), single fiber (bi-directional) and

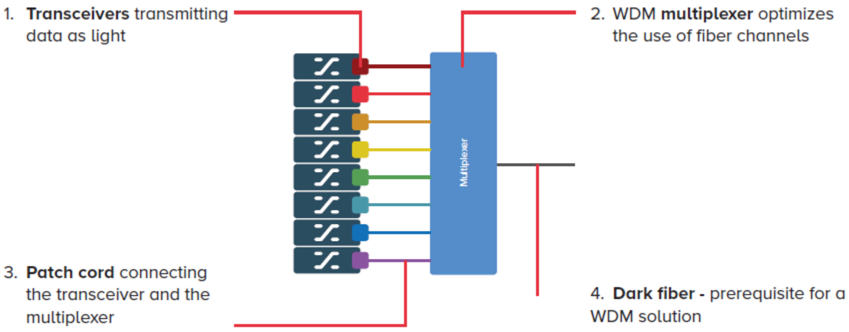

Figure 13. a WDM system (source: Smartoptics)

CWDM/DWDM. A standard transceiver, otherwise known as a grey transceiver, is a single channel device. It is known as grey because xWDM signals are colored wavelength channels, and any signal that is not xWDM is uncolored or 'grey'.

Grey transceivers have two main applications, firstly they can be connected directly to a single fiber channel or ethernet data switch to transport data over dark fiber, or they can act as an optical interface on the client side of a transponder based xWDM system.

There are four grey (standard) transceiver types, each with their own transmission channel and distance:

Table 5: Standard grey transceiver types

Type	Channel	Max Distance
Short Range (SR)	850 nm	300 m
Long Range (LR)	1310 nm	10 km
Extended Range (ER)	1550 nm	40 km
Further Extended Reach (ZR)	1550 nm	100 km

A single fiber bi-directional transceiver (BiDi) has two separate wavelength channels. One is used to transmit and one is used to receive traffic over a single fiber strand.

Coarse Wavelength Division Multiplexing (CWDM) and Dense Wave Division Multiplexing transceivers are collectively referred to as xWDM.

Again, they have two fundamental applications. Firstly, they are able to connect directly to a data switch to transport an xWDM wavelength over dark fiber. Secondly, they act as the output signal from a transponder based xWDM system.

CWDM supports up to 18 wavelength channels transmitted through a fiber at the same time. To achieve this, the different wavelengths of each channel are 20nm apart.

DWDM, supports up to 80 simultaneous wavelength channels, with each of the channels only 0.8nm apart. CWDM technology offers a convenient and cost-efficient solution for shorter distances of up to

Figure 14. Examples of transceiver form factors (source: Sopto)

70 kilometers. For distances between 40 and 70 kilometers, CWDM tends to be limited to supporting only 8 channels instead of 18.

DWDM connections can be amplified and can therefore be used for transmitting data over much longer distances. DWDM is able to handle high speed protocols up to 100Gbps per channel over distances of up to 80km.

Both CWDM and DWDM solutions are available as active or passive systems. In a passive, unpowered solution the XWDM transceiver resides directly in the data switch. The output from the XWDM transceiver connects to an unpowered multiplexer that combines and redistributes, multiplexes and demultiplexes, the various signals. As the XWDM transceiver resides in the data switch, it means that all XWDM functionality is embedded in the data switch.

Active XWDM solutions are stand-alone AC or DC powered systems separated from the switch. The task of the stand-alone system is to take the short-range optical output signal of the fiber or IP switch and convert it to a long-range XWDM signal. This OEO, (optical to electrical to optical), conversion is handled by a transponder. The converted XWDM signal is then transmitted with the help of transceivers and multiplexers. Due to the separation of the XWDM transport solution from the actual switch, active systems also tend to be more complex than passive, embedded solutions.
There are many different types of xWDM optical transceivers, which come in a variety of shapes and sizes. These variants are known as form factors.

The type of form factor required depends on the type of data, speed and distance required within the network. Different rules then determine how the data is transmitted. These rules are called protocols. The form factor of an optical transceiver specifies the dimensions of the transceiver, its shape and its size. The size differs according to speeds and protocols.

In general, manufacturers design according to the multisource agreement (MSA). This is a standard for ensuring that the same form-factor transceivers from different vendors are compatible in size and function, ensuring interoperability.

The most commonly used form factor in telecoms and datacoms is the Small Form factor Pluggable family of products (SFP/SFP+/QSFP).

6.2 Optical networking in access networks

In the last few years fiber roll-out in the access network layer of fixed line networks (Fiber to the Home or FTTH) has accelerated in the Netherlands. Currently about 4.6 million homes are connected to fiber (though not all are used). It is expected that by the end of the year 2026 all Dutch homes will have been connected with fiber optic cable. There are currently 8 million Dutch homes a number that is expected to grow to 9.8 million by 2026.

Passive Optical Network

The most dominant technology used for FTTH is that of the Passive Optical Network family of products or PON. In a PON access network there are two end-points with active (powered) electronic transmission equipment, connected by passive (non-powered) equipment known as outside fiber plant. At the subscriber premises, there is an Optical Network Termination (ONT) device that terminates fiber and connects home devices (TV, PCs etc). The second end point is the Optical Line Termination (OLT), typically located in a telecom central office. The role of OLT is to aggregate connections of multiple users and connect them to the core network.

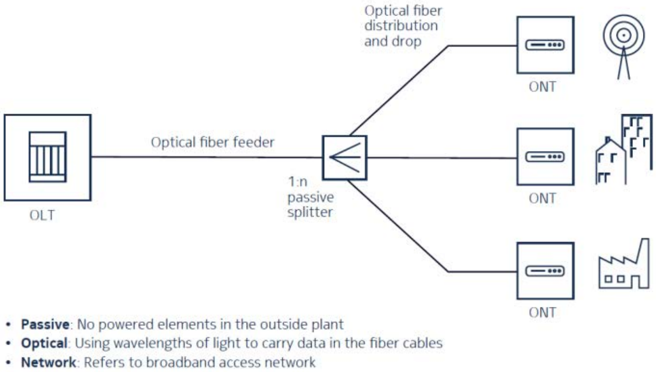

Figure 15. PON based access network (source: Nokia)

A PON network uses point-to-multi-point topology. This means that a single transmission point in the OLT connects multiple end-points. Optical splitters are used to split the signal into multiple branches. There could be several levels of splitters, which are separating the outside plant into different sections: fiber feeder, distribution, drop. The outside plant components (splitters, cables, distribution frames etc.) are completely passive and don't need power, making PON networks very power efficient.

The total available bandwidth on a PON is divided between all users that are connected to that PON, i.e. each user is allocated a portion of the total available bandwidth by a traffic manager. PON has advanced security mechanisms in place to ensure that each end point ONLY receives the signal that is intended for them.

There are two branches in the PON family tree: Gigabit PON (GPON) and Ethernet PON (EPON). And there have been many advances in each branch over the years, resulting in new flavors of PON with progressively better performance.

GPON and XGS-PON are by far the most widely deployed PON technologies in the world today.

Each flavor of PON uses a different wavelength pair (one in upstream, one in downstream) to transmit data. The wavelengths are specified by international standards and stretch from 1260 to 1600 nm. Upstream traffic mostly uses the lower bands, because lasers operating in these bands are more cost-efficient, which is important for ONTs that are deployed in big volumes. Different PON technologies that use different wavelengths are able to coexist on the same fiber optical cable. This makes it simple to migrate from one generation of PON technology to the next.

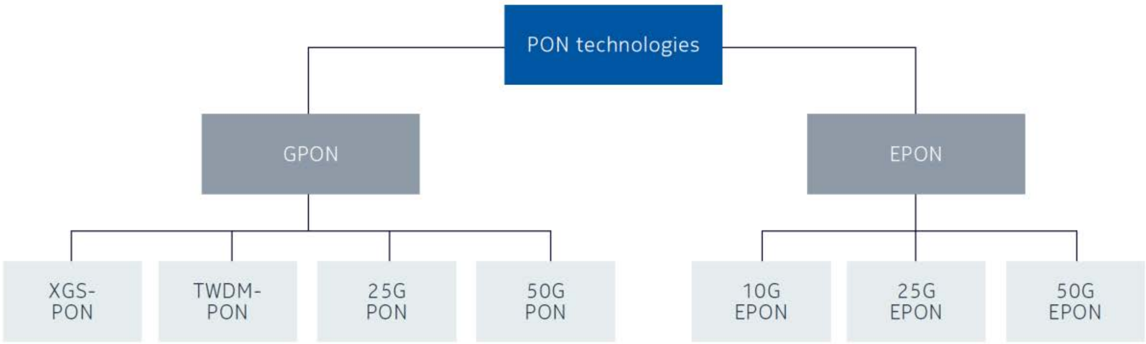

Figure 16. GPON and EPON (Source: Nokia)

Optical networking in 5G Mobile front-haul and mid-haul

With the advent of the roll out of 5G mobile networks the last copper based transmission links within mobile networks - the links between base band units and antennas - will now be replaced with a fiber optic alternative. This sometimes referred to as Fiber to the Antenna (FTTA).

The 5G Cloud Radio Access Network (C-RAN) architecture enables the disaggregation of the Radio Access Network. Baseband functionality is split into two functional units: a centralized unit (CU) and a distributed unit (DU).

Where these units reside depends on the specific architecture and geographical locations available. It can include putting them in sites other than the cell site where the radio unit (RU) is located.

In a centralized 5G C-RAN the DU is hosted in an edge cloud data center or central office, and the CU can be co-located with the DU or hosted in a regional cloud data center. The connection between the RU and DU is referred to as “fronthaul” while the connection between the DU and CU is often referred to as “mid-haul”.

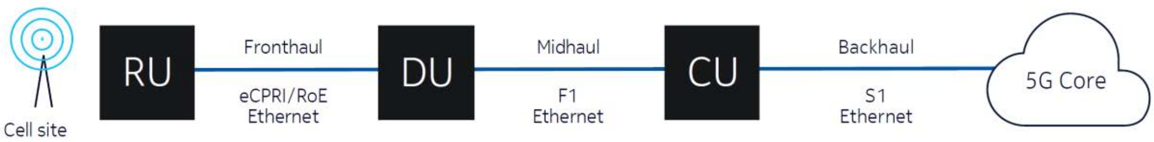

Figure 17. 5G disaggregated RAN

5G radio access networks place unique requirements (high capacity, ultra-low latency, high precision synchronization, high resiliency, etc.) on fronthaul connections. Thus, fiber based optical connections become paramount.

Standard colored or grey optical transceiver modules can be used depending on capacity and distance requirements for the particular configuration.

The dual small form-factor pluggable (DSFP) optical transceiver module is a good solution. As it supports both the Ethernet and wireless eCPRI protocols needed in the 5G front and mid haul. In Asia and the USA where 5G C-RAN has already been deployed the 25G SFP28 grey light transceiver modules are the most commonly used model.

6.3 Optical networking in indoor networks

A next wave of optic deployment is expected to take place within homes and offices. Two optics based inbuilding network technologies are emerging:

Plastic Optical Fiber

As fiber to the home becomes pervasive end-user internet connection speeds of 1Gb/s or even 10 Gb/s become possible. However, much of that speed is often lost because typically the main Optical Network Terminal (ONT) connects to a WiFi router that is located in a thick walled cabinet that is far away from the main dwelling areas.

Many users then resort to WiFi range extenders over powerline which improves the signal somewhat but still results in the loss of often half the speed or more.

Do it yourself home fiber installations are now starting to come to market based on Plastic Optical Fiber cables that are thicker and far less fragile than glass fiber cable making it much easier much easier to handle and to install yet still provide high performance for distances of up to 80 – 100 meters.

The cables can be installed in any existing conduits and used whether the ONT and Router are separate devices or integrated into one device. Ideally these are two sperate devices in which case the POF is used to connect the ONT typically located in a remote cabinet to the WiFi Router that should be placed in a central location like a living room (figure 15A). If the 2 units are integrated in one device then the POF is used to connect that device to centrally located WiFi Access point (figure 15B).

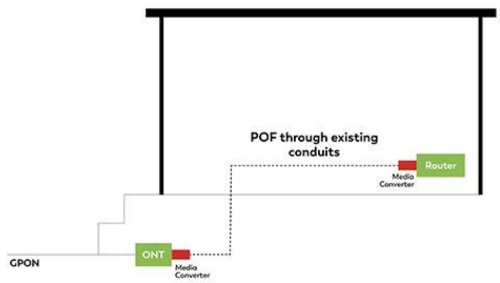

Figure 18A: POF usage for indoor cabling in case ONT and Router are sperate units (source: KDPOF)

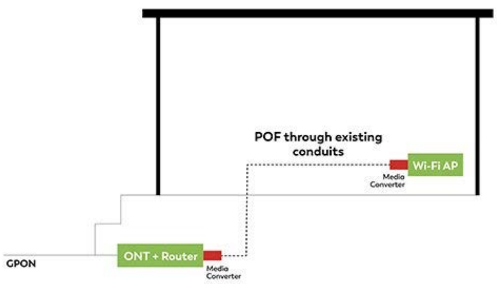

Figure 18B: POF usage for indoor cabling in case ONT and Router are one device (source: KDPOF)

In both cases the signal is processed and transmitted through affordable media converters at each end of the POF cable. Key elements in the media converters are low cost digital signal processors and low cost Fiber Optical Transceivers (FOTs).

Light Fidelity (LiFi)

LiFi, short for Light Fidelity, is a wireless communication technology that uses light waves instead of radio frequencies. The system is capable of transmitting data at high speeds through a range of light spectrum, such as visible light, ultraviolet and infrared light.

In terms of its end use the technology is similar to WiFi. The key difference being LiFi's use of light instead of WiFi's use of radio. The currently marketed system mainly targets office environments and consists of three main components: an Access Point, a transceiver and a USB connectable Access Key.

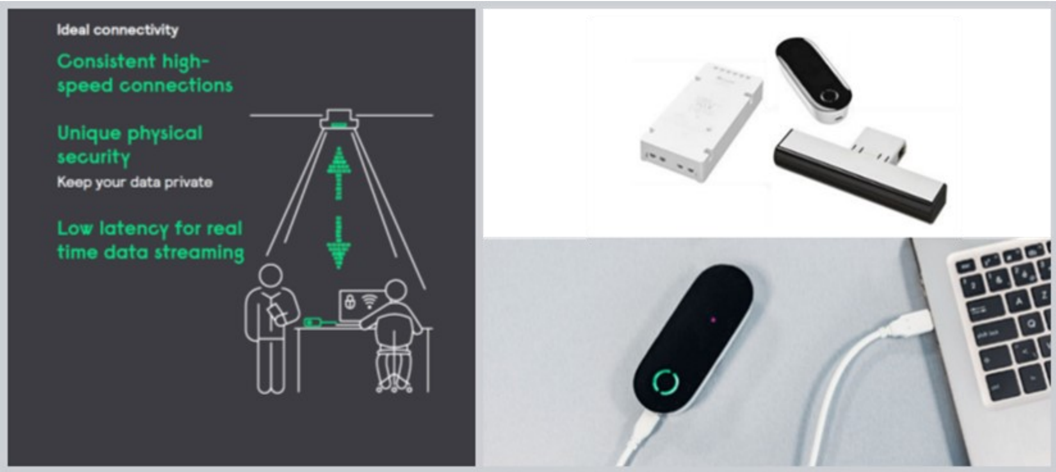

Figure 19. LiFi system main components (source: Signify)

The LiFi Access Point connects the data network (internet) to one or more ceiling mounted transceivers. The Access Key connects to an end device (laptop) via USB. The ceiling mounted transceiver can simultaneously connect multiple access keys within its footprint to the internet.

First generation ceiling mounted transceivers are special infrared transceivers. However, the product roadmap envisions the integration of the transceiver within the standard indoor LED based lighting fixtures thus integrating illumination and internet access in one system.

6.4 All optical data center interconnect

In data centers today, signals are typically transported optically but switched electrically, requiring numerous "optical to electrical to optical" (OEO) conversions. Transport is done with static point-to - point optical links, while switching is done electronically using OEO switches and routers. OEO technology is not considered future proof as the transponders are tied to specific traffic formats and data rates and need to be replaced when the line transmission rates increase or change over time. A transponder is a device that, upon receiving a signal, emits a different signal in response. The term is a blend of 'transmitter' and 'responder'.

The extreme growth of data traffic leading to an increased need of high speed low latency interconnects has led to the advent of an all optical circuit switch that routes traffic while maintaining the signals in the optical domain.

Such all-optical or OOO switches will not replace all OEO switches but serve to improve interconnect scenarios. OEO switches will remain useful for routing individual IP data packets from input ports to output ports or for when the optical signal needs to be regenerated due to impairments.

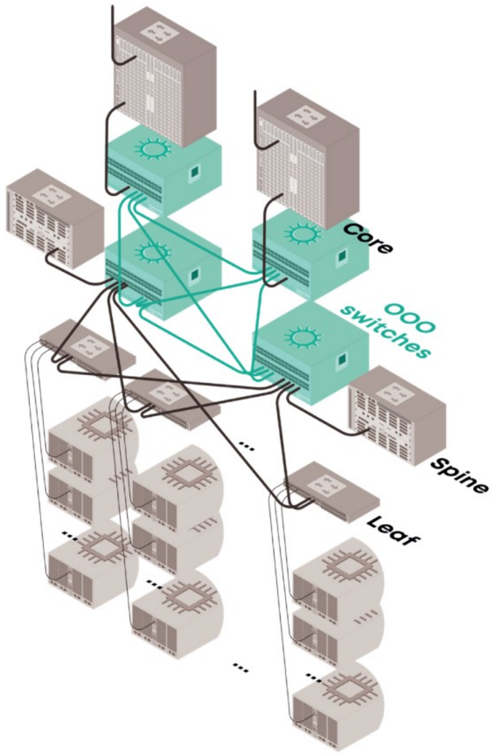

Figure 20. Hybrid OOO and OEO network (Source: Polatis)

6.5 Radio over Fiber

Radio over fiber (RoF) or RF over fiber (RFoF) refers to a technology whereby light is modulated by a radio frequency signal and transmitted over an optical fiber link. Main technical advantages of using fiber optical links are lower transmission losses and reduced sensitivity to noise and electromagnetic interference compared to all-electrical signal transmission.

One main application is to facilitate wireless access, such as 5G and WiFi simultaneously from the same antenna. In RoF systems, wireless signals are transported in optical form between a central station and a set of base stations before being radiated through the air. The advantage is that the equipment for WiFi, 5G and other protocols can be centralized in one place, with remote antennas attached via fiber optics serving all protocols. It greatly reduces the equipment and maintenance cost of the network.

RoF transmission systems are usually classified into two main categories (RF-over-fiber and IF-over-fiber) depending on the frequency range of the radio signal to be transported. In RF-over-fiber architecture, a data-carrying radio frequency signal in the same radio frequency that needs to be transmitted by the antennas is imposed on a light wave signal before being transported over the optical link. Therefore, wireless signals are optically distributed to base stations directly at the desired frequency and converted from the optical to electrical domain at the base stations before being amplified and radiated by an antenna. As a result, no frequency up–down conversion is required at the various base stations, thereby resulting in simple and rather cost-effective implementation is enabled at the base stations.

In IF-over-fiber architecture, an intermediate frequency radio signal with a lower frequency is used for modulating light before being transported over the optical link. Therefore, before radiation through the air, the signal must be up-converted to RF at the base station. This method is useful if the required RF frequency is very high like in mmWave.

An important application of RoF is its use to provide wireless coverage in the area where a wireless backhaul link is not possible. These zones can be areas inside a structure such as a tunnel, areas behind buildings, mountainous places or secluded areas such as jungles.

6.6 Power over Fiber

5G cellular networks are expected to increase the system rate around 1000 times higher than 4G systems. The reduction of the cell sizes for providing this high bandwidth requirement foresees a huge increment in the power consumption demand by the massive installation of remote antenna units (RAUs). Thus centralized/cloud radio access network (C-RAN) approaches are being considered as a 5G fronthauling solution to minimize this impact, to allow a sustainable scalability and to reduce the geographic distribution of maintenance sites.

Moreover, power consumption demands of future 5G-based RAUs can be dramatically reduced, especially if Radio over Fiber (RoF) mobile fronthaul is considered, where processing is done at the Central Office (CO).

This upcoming 5G technology opens up new application niches for the power-over-fiber (PoF) technology, where low power simplified antenna units (around 80 mW) are required as part of future 5G-based Remote Antenna Units (RAUs) together with some strategy for energy saving to reduce the impact of the power consumption of the remote antenna units, including the capability of turning into sleep mode some RAUs with total energy savings estimations of more than 30%.
Additionally, remote local battery charging with energy harvesting devices can also be used to backup operations and to provide feedback to the central office (CO) in case of failure.

The use of multicore fibers (MCF) in the C-RAN scenario provides compact designs to develop optical fronthaul technologies targeting spatial division multiplexing (SDM) with increased aggregated capacity. MCF can also contribute to downsizing the footprint by using optical fiber composite low-voltage cables.

Radio-over-Fiber (RoF) mixed with Power over Fiber through multicore fibers is suitable for small cell operations in advanced radio access networks (RAN).

6.7 Free Space Optics

Free-space optical communication (FSO) is an optical communications technology that uses light propagating in free space to wirelessly transmit data for telecommunications networking. "Free space" means air, outer space, vacuum, or something similar. This contrasts with using optical fiber cable.

Though the technology was known for quite some time, recent technical breakthroughs have made it a viable optical alternative for microwave point-to-point connections for mobile backhaul where terrain or other conditions make fiber connections too difficult or too expensive to provide.

6.8 Optical (GPS independent) Network Timing

Cellular network require frequency synchronization between the various components. Timing signals are used for generating the radio frequencies for transmission and therefore the synchronization must be very accurate and very stable. The advent of 5G based Ultra Reliable Low Latency services and the use of Time Division Duplexing makes that even more important.

Currently, time synchronization in mobile networks is done based on GPS signals from Global Navigation Satellite Systems (GNSS), however this has 4 drawbacks:

1. Politics – GNSS systems are run by national governments (and frequently, the military) and there are only a few of these systems. So, a country has to be willing to base its cellular network on the goodwill (and competence) of another country.
2. Cost – GNSS receivers and the electronics necessary to generate high quality synchronization from the GNSS are expensive.
3. Visibility – To use a satellite-based solution, the receiver must be able to see multiple satellites.
4. Security – GPS signals can be jammed and spoofed, making networks vulnerable for attacks.

As 5G rollout expands, with the aim of providing very high data rates to a large number of connections including Internet of Things (IoT) devices, the use of small cells becomes increasingly important. Per unit, a small cell is much less expensive than a macro cell. But they come in larger numbers, including GPS components on every one of them wrecks the economics. Furthermore, they are often inside homes, businesses, stadiums, etc. and therefore won't always have a clear view to multiple GPS satellites. Also, outdoor small cells are typically used in cities where large buildings also block the view.

An emerging alternative is an Optical Positioning, Navigation and Timing solution developed by a Dutch company in Amsterdam. The company, OPNT, is a spin-off of the VU Amsterdam university. The company offers a terrestrial network based system using existing fiber optic cables that is capable of providing sub-nanosecond accuracy using atomic clocks from the various European Meteorological Institutes as reference clocks.

6.9 Co-packaged optics

Co-Packaged Optics (CPO) is an advanced heterogeneous integration of photonic and electronic Application Specific Integrated Circuits (ASICs) on a single layer that is combined with highly advanced Digital Signal Processing (DSP) aimed at addressing next generation bandwidth and power challenges.

Co-packaged optics for Small Form Pluggable (SFP) coherent Dense Wavelength division Multiplexing (DWDM) transceivers

DWDM allows multiple wavelengths to be used to modulate data streams on, stack up and transport through a single strand of fiber. This is then separated again on the other side of the fiber thus simultaneously transporting a huge amount of data on a single fiber.

Typically there is are trade-offs between module size versus delivered performance; module price versus delivered performance, and energy efficiency versus delivered performance. The use of co-packaged optics changes these equations and allow for low power, miniaturized transceivers at affordable prices without performance degradation. Effect Photonics has achieved this by combining a Photonic Integrated Circuit (PIC) with a Semiconductor Optical Amplifier (SoA) on a single chip.

Co-packaged optics for Photonic Switching

Pluggable optics installed in the switch faceplate are connected to switch serializer/deserializer (SerDes) ports using an electrical trace. But as the data center switch bandwidth grows, connecting the SerDes to pluggable optics electrically will be more complex and require more power. With co-packaged optics, the optical chip is placed near the switch silicon within the same package, thus reducing power and enabling continued switch bandwidth scalability. This photonic switching concept has be achieved by the Technical University in Eindhoven. Intel claims a similar achievement.

Co-packaged optics for general purpose Machine Learning accelerator

This co-packaged chip combines photonics and transistor-based systems in a single, compact module. A silicon photonics processing core is used for most computational tasks, providing offload acceleration for high performance AI inference workloads with extreme performance and efficiency.

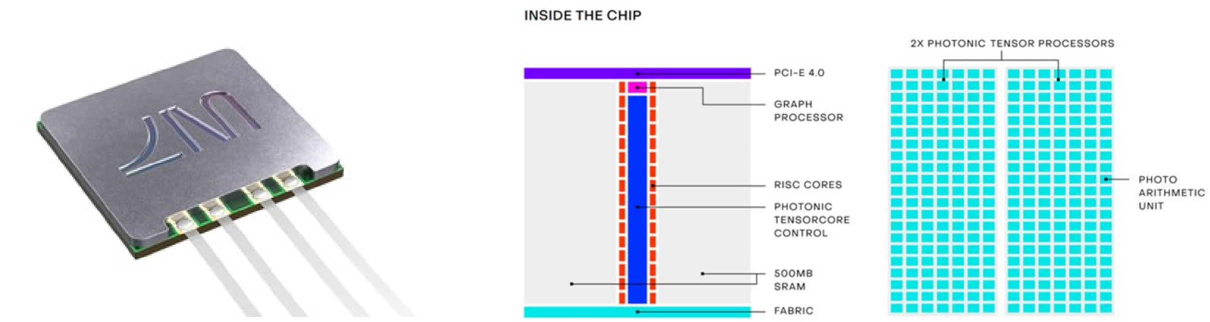

Figure 21. Photonic Processor (source: Lightmatter)

Co-packaged optics based photonic computing

Combining large scale photonics chips and a CMOS microelectronic chips in advanced packaging methods an optical computer is designed by integrating roughly 10,000 photonic devices on a single chip running a system clock of 1GHz. The aim is to solve complex mathematical problems.

Figure 22. Photonic Computing (source: Lighttelligence)

7. Power consumption and power savings

7.1 Power consumption

According to the Dutch National Statistics Agency (CBS) The electricity consumption of data centers in the Netherlands was 3.2 TWh in the year 2020, a figure equal to 2.8% of total electricity consumption in the Netherlands.

An earlier study conducted in 2016 by the Delft University of Technology that aimed to estimate electricity consumption in the period 2020 – 2030 analyzed the electricity consumption of both the supply and the demand side of the ICT sector of the year 2013. It concluded that the total adds up to

8%. In the year 2013, the share of the ICT supply side (= data centers and telecom providers) was 2.5% of the electricity consumption in the Netherlands and 5.5% was that of the ICT demand side (= businesses and consumers). These figures are of course no longer valid and unfortunately a more recent similar study could not be found but they do serve as a reminder that there is also an energy consumption component on the demand side that may still be larger than the supply side.

Useful consumption and waste

When considering the power consumption of data centers it should be noted that not all consumed power goes towards powering the actual IT equipment. Significant amounts are used for overhead items such as lighting, security, cooling etc. This can add-up to 50% of total power consumption.

Because the aim of data centers is to process IT workloads, the portion of power used for the IT equipment is labeled 'useful consumption', while the remainder is considered overhead or 'waste'.

There are three types of losses:

1. **No-Load Loss,** Energy losses regardless of how much of the available server capacity is actually utilized.
2. **Proportional Loss,** as equipment is used losses occur in proportion to the load.
3. **Square-Law Loss,** as equipment capacity utilization increases beyond a certain threshold losses occur that are proportional to the square of the load.

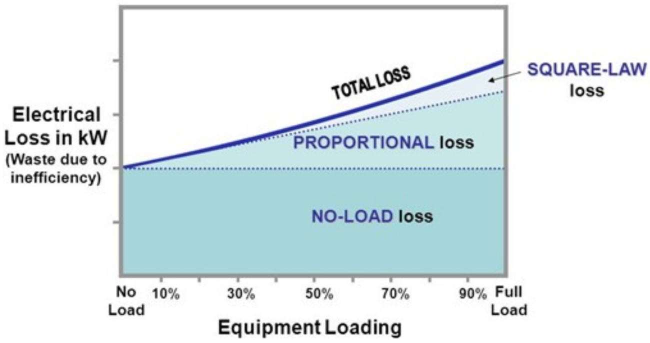

Figure 23. Types of power loss

PUE and DCiE

The Power Usage Effectiveness (PUE), defined as total power consumption divided by useful power consumption is a widely used ratio to rate the energy efficiency of a data center. If all consumption would be useful the value of that metric would be 1; if 50% is waste the value would be 2.

A related equally commonly used metric is that of the Data center infrastructure Efficiency (DCiE), which is the PUE inverse value (useful consumption divided by total consumption).

Based on these metrics Data centers are classified as following:

PUE	Level of Efficiency	DCiE
3.0	Very inefficient	33 percent
2.5	Inefficient	40 percent
2.0	Average	50 percent
1.5	Efficient	67 percent
1.2	Very efficient	83 percent

Figure 24. Data center efficiency ratings (source: Nlyte Software)

The PUE of a data center is not a static value, it varies with work load and changes in configuration if servers are added or removed. The PUE is therefore a metric intended to be measured regularly so as to learn the impact of changes made and managed so as to ensure optimal efficiency.

Design PUE

The design PUE is the calculated value for a modern data center as its construction is being planned. Typically the design PUE is part of the information that needs to be submitted when applying for a data center construction license and that needs to comply with government standards.

In the Metropolitan Area of Amsterdam, a design PUE of 1.2 or less is required for a newly built data center or for adding an expansion to an existing one.

The design PUE is a theoretical value based on peak utilization of the IT equipment. However, in practice most data centers are rarely operated at peak but typically on a utilization rate of 60% – 70% of capacity. The actual PUE is therefore often higher. A more accurate figure is the annualized PUE which is the ratio of the data center's annual energy consumption to the annual energy consumption of its IT equipment.

7.2. IT workloads

In computing, a workload is any program or application that runs on any computer. In the context of data centers it is typically a complex software application hosted on one or more servers with large amounts of user systems connected and interacting with the application servers across a vast network.

Workload can also refer to the amount of work (or load) that software imposes on the underlying computing resources. Broadly stated, an application's workload is related to the amount of time and computing resources required to perform a specific task or produce an output from inputs provided. A light workload accomplishes its intended tasks or performance goals using relatively little computing resources, such as processors, CPU (central processing unit) clock cycles, storage I/O (input/output) and so on. A heavy workload demands significant amounts of computing resources.

Different workloads consume different amounts of energy. The total amount of useful power consumption in a data center will thus vary depending on the type of tasks performed.

workloads may be classified as static or dynamic. A static workload is always on and running, such as an operating system (OS), email system, enterprise resource planning (ERP), customer relationship management (CRM) and many other applications central to a business's operations.

A dynamic workload is temporary; loads and runs only when needed. Examples include temporary instances spun up to test software or applications that perform end-of-month billing.

real-time software, emphasizes high-throughput and low-latency performance to operate in sensitive real-world environments such as medical and industrial systems.

Analytical workloads are a growing type of workload that are driving the increase in data center energy consumption. Such software applications analyze enormous amounts of data, sometimes from varied and disassociated sources. Designed to find trends, make predictions and drive adjustments to business operations and relationships. This is the underlying notion behind more advanced programming such as big data and machine learning software technologies.

High-performance computing (HPC) workloads, frequently related to analytical workloads, perform significant computational work and, typically, demand a large amount of processor (CPU) and storage resources to accomplish demanding computational tasks within a limited timeframe, even in real time.

Artificial Intelligence

Artificial Intelligence (AI) is a specific type of analytical workload designed to mimic the human thinking capabilities. AI's primary goal is to make decisions based on large data sets and provide information to maximize high accuracy and likelihood of achievement.

There are many industries that are advancing and progressing because of AI. Some of these industries include: healthcare, finance, cybersecurity, manufacturing, daily transportation, weather forecasting, advertising/marketing etc. It is a rising, impactful, and in-demand technology in our modern society, which has unlimited potential in any field of interest.

There are 6 basic AI workloads that are at the core of AI functions. These workloads are:

1. Predicting & Forecasting

This is a concept that involves predicting future events based on statistic and data supported reasoning. Forecasting is based on large amounts of historic information to make predictions and "educated" guesses. An example of this workload is the use of AI weather forecasting systems. The AI model is provided with an abundant amount of data of previous weather conditions to predict weather changes, and long-term climate impact.

2. Anomaly Detection

Anomaly detection is the use of machines to observe and find unexpected values, within an already established range of values. It can help with identifying previously unnoticed, out-of-place events in any given situation. The way it works is that the machine learning model has been fed with large amounts of data and based on that information it developed a fitting range for all the info. If it then detects something out of this existing range, it triggers a call to action. An example of such an application can be smart watches. They track your heart beat rate and based on the vast amount of data it has already collected can define usual and detect unusual heart rhythms.

3. Computer Vision

Computer vision analyzes videos and images to extract specific information. This information could include, key words in an image, the mood/context of a video or the subject matter in an image. Computer vision involves a model being trained with thousands of pictures, labels, videos and key content to be able to recognize and classify images. Examples of computer vision in action could be every time Safe Search is run when surfing on Google. This feature's targeted goal is to classify explicit and non-explicit videos and images based on a user's search to filter out inappropriate images.

4. Natural Language Processing

Natural Language Processing (NLP) is the main tool for machines to be able to communicate in human languages as well as convert speech to text, and vice versa. It can allow computers to understand questions, requests and commands as well as respond to them. It can perform translations, and analyze text within larger documents for key content or specific targets.

Within NLP, the whole concept moves around the idea of training a model with trillions of phrases, sentences and key words. Once again, it's related to datasets and responding/making decisions based on the data the model is given and trained with. It's also one of the key factors for ordinary human beings to be able to communicate with computers and machines. Examples are Siri, Alexa and Google translate.

5. Knowledge Mining

There are many instances where it is difficult for computers and machines to extract specific, valuable data from given unstructured files or content. This includes images, videos, pdf files, text etc. This is where knowledge mining comes into play.

There are three basic steps when it comes to knowledge mining: ingest, enrich, and explore. The first step is to feed the program with content from a range of sources, using connectors to first and third-party data stores (data can be structured or unstructured). The second step is to enrich the data with AI capabilities (language, speech, vision etc.) to extract value from the gathered data. The final step to knowledge mining is analysis of data through exploration. An example is for an airline engineering department to use knowledge mining to quickly identify the cause of a defect based on previous reports on similar defects collected in various formats.

6. Conversational AI

This feature allows simple AI solutions to be able to converse and talk with humans. The biggest use of conversational AI is through chatbots. Chatbots can be used for things like customer support, making appointments, medical consultations, personalized assistants etc.

In a basic sense, this model works by trained data sets as well. The data sets will contain detailed information about specific things to look for as well as collect when interacting with users.

7.3. Possible energy savings

A 2021 IEEE Access publication aiming to model electricity consumption in data centers resulted in the following breakdown of electricity consumption within data centers:

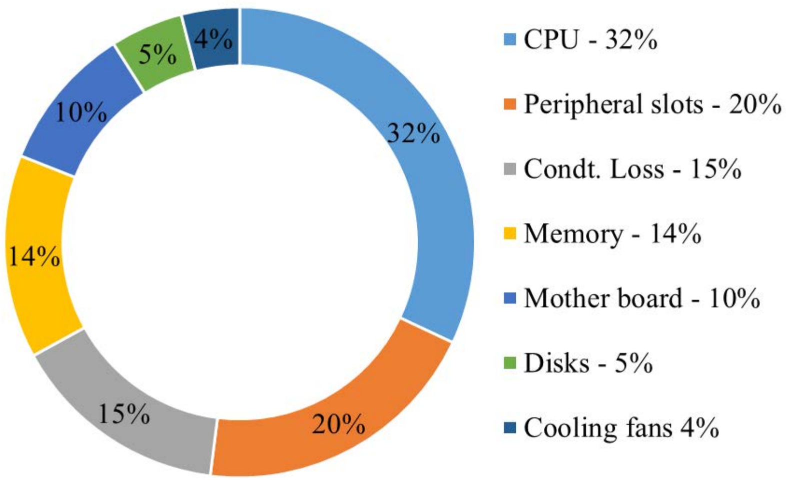

Figure 25. Breakdown of electricity consumption in an average US data center (source: IEEE Access)

The analysis in the article is based on US averages. They show about a third of total electricity consumption to be used by the CPU. Obviously these figures may not be valid for the Netherlands but can still give us an indication of the possible energy savings that could be expected if co-packaged photonic processors such as the ones developed by the likes of Lightmatter and Lightelligence become available at mass manufacturing volumes in the coming few years.

Energy savings from Photonic Processors

According to Semiconductor Engineering estimates, with (near) co-packaged photonic processors the required energy per bit would go down from 8pJ/bit to 3pJ/bit. As a rough indicative figure for the order of magnitude savings that can be achieved in a data center where about a third of the consumed power is used for data processing is then: 32% * -5/8 = saving of 20% of annual electricity use.
Thus if we assume that on average the power in all Dutch data centers is used for data processing the potential energy savings, if all servers would be fitted with photonic processors would amount to around 20% x 3.2TWh = 0.64TWh per year.

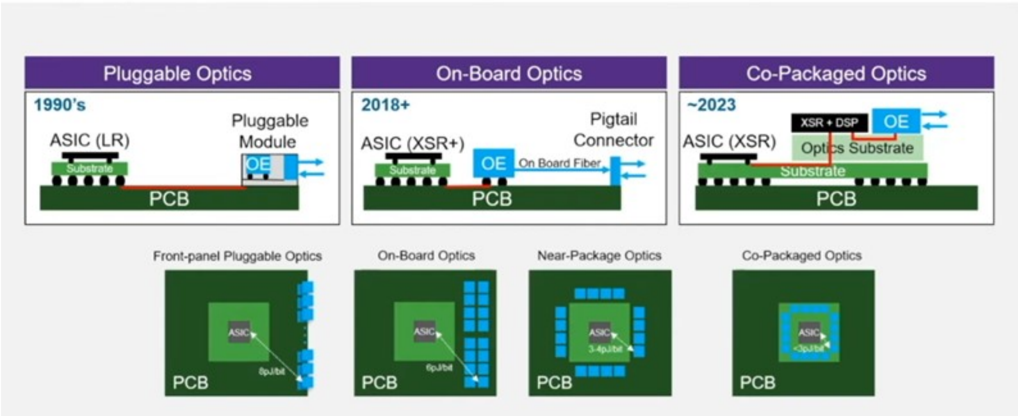

Figure 26. impact of co-packaged optics on energy needs (source: Semiconductor Engineering)

Energy savings from coherent photonic DWDM transceivers

Another possible saving should be achieved through the co-packing of photonic integrated circuits with a semiconductor optical amplifier. This achievement of companies like Effect Photonics should serve to at least extend Moore's Law allowing capacity of the transceivers to double every 18 months.

If we again assume the breakdown of the power consumption in the Dutch data centers is roughly equal to that shwon in figure 25 thus 20% is used for data transport; then this would potentially result in a savings of: 50%*20%* 3.2TWh per year = 0.32TWh per year.

Energy savings from photonic Switching

Ethernet switches equipped with 32 ports of 400G per port are now on the market. Optics co-packaged with a switch silicon on a chip reduces the power consumption from 8 to 3 pJ/bit. An energy saving of 5pJx32x400G= 0,064 J/s per switch.

Assuming the switch is on 50% of the time this equates to 0.064 J/s * 3600 * 24 * 50% * 365 = 1009 kWh per year per switch or roughly 1 MWh per year per switch.

Developments are going fast, on 16 August 2022, Broadcom has announced that it is shipping its StrataXGS® Tomahawk® 5 switch series. A co-packaged photonic switch module capable of delivering switching capacity of 51.2 Tb/s in a single monolithic chip. This allows for increased rack density thus further optimizing power consumption.

Energy savings from Fiber to the Home

There is a clear acceleration in the rollout of fiber in the Netherlands. Besides KPN, who is currently responsible for around 50% of fiber roll-out in The Netherlands other players are also active. These players are primarily funded by large Venture Capital companies such as EQT (Delta), KKR and Warburg Pincus (Open Dutch Fiber) and represent the bulk of the other 50% of roll-out.

The year 2021 ended with a net addition of 885.000 FttH connections. The ambition expressed by the leading three fiber operators – KPN, Delta and Open Dutch Fiber – translates to the addition of more than 1 million fiber to the home connections each year for the coming 5 years.

KPN announced that as of 2023 the copper lines based xDSL service will be discontinued. The KPN customers would thus need to switch over to fiber .

KPN estimates that the energy reduction resulting from the transition from copper cable to Fiber will be around 11 kWh per year per line. This would thus result in savings of roughly 11 GWh/year if all current xDSL users switch over to fiber (and not to Ziggo's coax TV cable!).

Other energy efficiency gains from deployment of photonic/optical technologies

The growing utilization of photonic/optic technology will facilitate further growth of network function virtualization, disaggregation, process automation and software defined networking. This in turn will create much more flexibility in the network architecture thus allowing greater freedom in work load balancing and power management optimization.

Such as:

- Dynamically shifting non time critical workload for processing in locations where renewable power is available
- 5G Mobile networks with centralized cloud-RAN and remote active antennas thus sharing base station equipment capacity between sites and minimizing power use of network
- Automatic power down/hibernation mode for remote radio heads while not in use
- Use of all-optical circuit switches for network interconnect and/or traffic rerouting will allow for remote operation thus minimizing the need to enter climate controlled areas.

While advancements in photonics and other technological areas bring the promise of substantial increases in power efficiency, other trends project increased usage of power hungry ICT applications such as Machine Learning and Big Data Analytics; Internet of Things and pervasive computing.

How that will play out remains to be seen. However, most industry experts expect the ultimately the ICT energy consumption will in the next few years stay flat or flattish while at the same time the data usage will at least double in size. The overall increased energy efficiency will thus offset the increased usage.

7.2. Main causes of energy in-efficiency in data centers

Conventions

Historically, power usage was not considered a key design criterion for data centers. The design of in particularly older data centers may therefore not be optimal for energy efficiency. For example the energy cost of Cooling and Air Conditioning Units may be high yet the data center design may make it difficult to introduce new cooling methods and more energy efficient server and rack configurations for cloud computing environments.

Tools for modeling the electrical costs of data centers are not widely available and are not commonly used during the design phase of a data center.

Also, electrical bills are typically sent out long after charges are incurred making it difficult to clearly link an increased cost of electricity to particular decisions like the installation of a new zone of equipment in the data center or operational practices.

Billed electrical costs are sometimes not within the responsibility or budget of the data center operating group or the electrical bill for the data center may be included within a larger electrical bill and may not be available separately.

Decision makers may not have sufficient information during planning and purchasing decisions regarding the electrical cost consequences of these decisions.

Operational conditions

The cloud computing environment comprises tens of thousands of server machines. Studies show the energy consumption of these machines to be far from uniform. Servers can consume 80% of the peak power even at 20% utilization. This energy non-uniformity of servers is a key source of energy inefficiency in the cloud computing environment. Servers are often utilized with between 10% to 50% of their peak load and servers experience frequent idle times. This means that servers are not working at their optimal power-performance tradeoff points mostly, and idle mode of servers consumes a substantial portion of overall power.

Yet another reason for energy inefficiency in cloud data centers is the need for multiple power conversions in the power distribution system. The main Alternate Current (AC) power supply from the grid is first connected to the Direct Current converter so that it can be used to charge the battery backup system. The output of this electrical energy backup system then goes through an inverter to produce the power which is then distributed throughout the cloud environment. These conversions are necessary due to the oversized and highly redundant uninterrupted power supply (UPS) modules, which are deployed for voltage regulation and power backup in cloud computing environment. However, most UPS modules in cloud data centers operate at 10%-40% of their full capacity.

7.3. Best practices for energy efficiency in data centers

Energy efficient Data Center Design

When building a new data center it is important to incorporate energy efficiency in the design criteria, make use of advanced energy consumption modeling tools and the latest power efficient technologies.

Operators of hyperscale data centers are best positioned to adopt this best practice as they can afford to build or rebuild data centers based on the latest design insights. Microsoft for example builds a new generation Data Center every 7 to 8 years that incorporate the latest and most advanced cooling techniques. They also ensure an optimal design of their data center equipment by procuring and replacing equipment every 3 years based on very detailed hardware specifications optimized for the specific workloads that they will be used for.

Retiring systems

Most data centers have old technology platforms that remain operational for archival or research purposes. In fact most data centers actually have application servers which are operating but have no users. It is useful to inventory these systems and create a retirement plan, in many cases systems can be taken offline and powered down even if they are not physically retired.

A related opportunity exists where multiple old technology platforms can have their applications consolidated onto new servers essentially reducing the total server count and thus power consumption. A reduction of up to 20% is possible in typical cases even if the floor space is not recovered due to the deployment of higher density IT equipment.

Activating power management features

Most new servers are equipped with power management features that make it possible to reduce power consumption at times of reduced computational load. This was not true a few years ago when the power consumption of virtually all IT equipment was constant and independent of computational load. Users should be aware of this change in Information Technology and be aware of the status of the power management features on their IT-systems. Where possible, power management should be enabled on all devices with such capabilities.

Deploying blade servers

Many data centers have so-called low density servers that are three to five years old. Most of these servers draw more or less the same power per server than today's blade servers and are physically much larger per server.

When a new server deployment is planned the use of blade servers as opposed to alternative server form factors will generally result in a 20 percent reduction in power consumption. This is because blade servers generally have higher efficiency power supplies and share some overhead functions such as fans.

Today's major blade server OEMs provide user configuration tools that accurately report actual power consumption for various blade server configurations. This gives much more accurate information and is more convenient than the conventional method of determining power consumption values of the legacy servers by measuring a representative server with a watt meter.

Virtualization

The deployment of virtualization in the data center can work to reduce the total server count. This in turn reduces the overall IT requirement. Virtualization almost always reduces the number of installed servers dramatically. The elimination of servers is a structural reduction op power consumption of 200 – 400 W per server depending on technology.

Standardization

For users who standardize on a blade server system and dedicate servers by application the option exists to standardize on two types of blades: a high performance blade requiring higher power and a low performance blade requiring less power. A logical strategy would be to deploy applications on the lower performance blade by default and only move to a higher performance blade if a certain application demonstrates the need for extra processing power

Use an effective server migration strategy

Some of the most effective server migration strategies include the following strategies:

- Using a two-way server or a single processor dual core server to replace two or more old servers.
- Replacing an old server with a blade based server on a low voltage or mid voltage processor.
- Replace a dual processor server with a single dual core processor
- Using a two-way dual-core server in place of a four-way server

Right-sizing the physical infrastructure system to the load

The data provided by manufacturers during a purchasing process is typically not sufficient to determine actual energy consumption differences. Furthermore, right sizing and system design each have a much higher impact on the electrical consumption than the selection of the physical infrastructure devices. Of all the techniques available to users, right-sizing the physical infrastructure system to the load has the most impact on electrical consumption. This is because there are fixed losses in the power and cooling systems that are present whether the IT load is present or not and these losses are proportional to the overall power rating of the system. For installations that have light IT loads the fixed losses of the physical infrastructure equipment commonly exceed the IT load power consumption. Whenever the physical infrastructure system is oversized the fixed losses become a larger percentage of the total electrical bill. The compelling economic advantage of rightsizing is a key reason why the industry is moving towards modular scalable physical infrastructure solutions.

System design

In data centers system design has a much greater effect on the electrical consumption then does the efficiency of individual devices. In fact two data centers comprised of the same devices may have considerably different electrical bills. For this reason the system design is even more important than the selection of power and cooling devices in determining the efficiency of a data center system design.

Issues that commonly reduce the efficiency of data centers include the following:

- Power distribution units and or transformers operating well below their full load capacities
- Air conditioners forced to consume excessive power to drive air against high pressures over long distances
- Cooling pumps which have their flow rate adjusted by throttling valves which dramatically reduce the pump efficiency

These kind of problems routinely cause data centers to draw much more physical infrastructure power as is necessary. Simple design decisions can help to avoid these problems at little or no expense for example by ensuring the design has been fully engineered and has undergone testing such as computational fluid dynamics modeling and comprehensive commissioning. This helps to position the data center to run much more efficiently.

Obtaining an integrated physical infrastructure system based on a standardized design comprised of modules that have been pre-engineered and pre-tested can also contribute significantly to the overall efficiency levels.

Implementing these best practices typically renders savings of 15% – 35%.

8. Other Findings

Overall, sustainability is top of mind with most members of the eco-system. Though not necessarily energy efficiency nor photonics. In datacenters most attention is given to the use of green energy and heat recovery.

Performance, operational and financial considerations are the main drivers for decision making. This tends to make many players in the eco-system risk averse. Only tried and tested technology is used. Co-location operators seem to be hands-off where it comes to technology selection. That is seen as predominantly a choice to be made by the tenants.

Hyperscale operators seem to be more at the forefront of technology adoption but due to their scale and need for supply chain security will only adopt technologies that can procured in large volumes from multiple suppliers. This builds in a certain adoption delay.

It is worth remembering here that the datacom and telecom eco-system is based on a series of interworking components, equipment and links in a long interoperable chain. The interworking between each part is standardized and just as dimensions and form factors, ‘power envelopes’ are often also standardized for a given component or interface such as for example a port on a server. The various equipment vendors then use that power envelope as a design parameter for their product, component, equipment or interface on which to optimize the performance of that given item.

Higher power efficiency is then typically implemented in the form of a higher performance within the same power envelope rather than by delivering the same performance at a lower power consumption. This maximizes the marketability of a given item for it thus becomes compatible with different generations of adjacent links in the overall chain.

It is also worth noting that within the ecosystem there are different business models applied with various degrees of sensitivity to energy efficiency.

One business model used by operators of smaller colocation data centers is very much akin to that of a real estate developer. A data center is built, first tenants are attracted as quickly as possible till a certain threshold occupancy rate is reached and then the data center is sold to a larger data center operator at a much higher price because now it is with tenants thus with regular cashflows yet with room to still for growth.

Sometimes and not always in support of the above mentioned business model, data center operators would seek to draw high profile tenants to their data center by offering them very attractive multi-year contracts (usually a low all-in price for rent/power/security etc.) in anticipation of these tenants attracting many other tenants that want to be co-located with these high profile tenants to facilitate direct interconnect. Lumpsum, all-in monthly fees that do not specify electricity fees are in any case not uncommon.

Another business model relies heavily on heat recovery. An entrepreneur buys second hand servers tailored for compute applications and makes an agreement with the owner of a hotel, swimming pool or apartment building to install a heat exchange in the basement of that building and recycle the data center waste heat into the hot water system of the building of the host in exchange for rent free floor space and electricity.

Several people interviewed emphasized the need for an integral holistic approach from the government. There was a general sentiment expressed that the current policies were too scattered, not in tune with each other, with an insufficient attention given to the Dutch national interest. A specific concern expressed was that if a freeze on data center construction is imposed that it would lead for data centers to be built outside the Netherlands subsequentially eroding the Netherlands' position as a major digital hub.

There was little awareness of the expressed Dutch government ambition to become a leader in the area of photonics. Though most were familiar with PhotonDelta, associated entities such as EFFECT Photonics were totally unknown.

Once informed of the Dutch ambitions and the availability of a significant amount of funding allocated in support of it , most were genuinely impressed and applauded the initiative. A few expressed the fear that once the first results become visible the successful companies and technologies would be lost to the Netherlands through a purchase by American or Chinese investors.

9. Conclusions & Recommendations

Technological development in co-packaged and integrated photonics is progressing well. Perhaps faster than initially anticipated. This will definitely have a great impact on datacom and telecom energy efficiency.

Most progress is made in the area of optical transceivers, a wide selection of products from a large number of suppliers is readily available.

Progress in the area of photonic switching and photonic processors is also clearly on the horizon, though a certain delay in adoption is to be expected due to the industry's need for proven technology, availability of high volumes and focus on high system reliability.

Photonics will help ensure the required capacity growth while increasing energy efficiency. It should be noted though that this will not lead to a lower energy consumption in absolute terms but rather in relative terms. A larger traffic volume will be processed with the same amount of energy.

Less active equipment and novel architectures become possible thus further reducing overall network complexity and further enhancing performance and energy efficiency. However, the use of photonics based systems will initially impact efficiency in the useful power consumption only.

To ensure minimal waste it is important that energy efficiency is embedded in all aspects of the business. Starting with the data center business model and design.

Energy consumption and costs should be made identifiable as a separate item and actively monitored & managed so as to be able to quickly see impact of configuration changes on energy consumption.

Energy efficiency best practices should be observed, particularly those ensuring the rightsizing of capacities. Equipment should be upgraded every three years.

Taking a holistic view of the value chain, power station capacity and locations should be harmonized with data center locations and capacities while taking into consideration the long lead times involved. Instead of a construction freeze it might be wiser to consider promoting data center reconstructions so as to reduce the number of old data centers with outdated energy inefficient designs.

This is a publication of:

Netherlands Enterprise Agency
Prinses Beatrixlaan 2
PO Box 93144 | 2509 AC The Hague
T +31 (0) 88 042 42 42
Contact us
www.rvo.nl

This publication was commissioned by the ministry of Economic Affairs and Climate Policy

Publication number: RVO-212-2022/RP-DUZA

NL Enterprise Agency is a department of the Dutch ministry of Economic Affairs and Climate Policy that implements government policy for agricultural, sustainability, innovation, and international business and cooperation. NL Enterprise Agency is the contact point for businesses, educational institutions and government bodies for information and advice, financing, networking and regulatory matters.

Netherlands Enterprise Agency is part of the ministry of Economic Affairs and Climate Policy.