ENergy-efficient Light-wave Integrated Technology Enabling Networks that Enhance Datacenters (ENLITENED) Program Overview

B. PROGRAM OVERVIEW

1. Summary

The growing demand for datacenter services across a range of applications has resulted in significant and sustained growth in electrical energy consumption in the Information Communications Technology (ICT) sector. Currently, datacenters consume more than 2.5 % of US electricity and this percentage is projected to double in about 8 years [1-4, 6, 7, 9]. Efficiency improvements due to more efficient cooling, power delivery, and electronic processor chips via Moore's law improve overall efficiency, but do not significantly slow the current growth trend; to do so requires a transformative improvement. The overall objective of the ***ENLITENED*** (***EN***ergy-efficient ***L***ight-wave ***I***ntegrated ***T***echnology ***E***nabling ***N***etworks that ***E***nhance ***D*atacenter**s) program, therefore, is to provide a transformative change - to achieve an overall doubling in datacenter energy efficiency in 10 years through deployment of novel network topologies enabled by integrated photonics technologies. ARPA-E estimates that if the technical challenges posed by ***ENLITENED*** can be overcome, these alone would reduce projected US energy use by about 1% after 10 years and realize at least twice the number of datacenter transactions with the same amount of energy.

Industry projections show that in order to achieve future datacenter performance requirements, metal interconnects must be increasingly replaced by photonic technologies, yet costs for deployment are often prohibitive and energy efficiency is not necessarily the highest priority in the commercial sector [11, 15, 16]. Though a broad industrial consensus pushes toward further photonic integration in switches and other datacenter components [11,16,18], in some cases, large companies will build another 80 MW datacenter to meet demand, rather than increase efficiency with photonics, due to a combination of reliability risks, cost and limited component supply.

To overcome metal interconnect limitations on future datacenter energy-efficiency performance, ***ENLITENED*** will target the critical packaging and integration challenges needed to exploit the inherent performance advantages of dense photonic interconnects and switching technology at the chip-scale within datacenters. Specifically, ***ENLITENED*** will target packaging and integration of novel and efficient photonics-enabled hardware systems that can demonstrate at least a 2-fold increase in energy efficiency at the datacenter level. To validate hardware solutions, ***ENLITENED*** will also entail modeling and simulation of the new datacenter architectures and data traffic protocols under realistic workloads, to provide quantifiable measures to validate transformative design strategies for future datacenters and retrofits.

2. Background and Motivation

Impact on Datacenters: Energy Consumption and Industrial Trends

The energy consumed by datacenters has been the subject of much study [1-5], reaching 91 TWh in 2013, and increasing at a rate that doubles about every 8 years [1]. Currently, datacenters account for about 2.5% of electricity consumption in the US – while the total annual electricity consumption in the US has remained approximately constant at about 3800 TWh [6] from 2005 to 2012[1].

Datacenters in the US can be divided according to their electricity consumption by market segment as: Small and Medium-Sized (49%), Enterprise/Corporate (27%), Multi-Tenant (19%), Hyper-Scale Cloud Computing (4%), and High Performance Computing (1%) [4] – which shows not only a wide diversity in scale and function, but also highlights the importance of finding energy efficiency solutions that address all market segments. Trends indicate that Multi-Tenant and Cloud services

[1] US electrical power generation is projected (EIA reference case) to be 3.9 trillion kWh in 2016, 4.2 in 2026 and 4.7 in 2040. Annual Energy Outlook 2016, early release: http://www.eia.gov/forecasts/aeo/er/index.cfm

are growing, which may improve efficiencies due to the possibility of virtualizing smaller datacenter functions within larger server farms that allocate resources more efficiently. Nevertheless, the fundamental technologies used, and solutions envisioned are relevant to all datacenters. With growth in all sectors, the importance for reducing energy consumption across the datacenter application domain continues to grow.

Figure 1 depicts the current and projected trends in aggregate power consumption for datacenters in the US (red line). The demands for datacenter services have resulted in an exponential growth rate in the total power requirement. The exponential growth in service demand is actually outpacing the exponential growth in ICT equipment computing efficiency (due in part to Moore's Law scaling of processor performance). Even more concerning is the trend towards an increasing fraction of datacenter power consumed by the internal datacenter network [17]. This pending upswing in the projected power consumption is attributed to anticipated growth in interconnect energy demands with increasing ICT performance, but with existing network technologies it appears unlikely that network energy efficiency will increase proportionally to server performance. Ultimately, the electrical interconnects to the server chips themselves will not be able to keep up with the increasing computing power of those computing elements [14] – which could potentially add to the energy-efficiency burden of datacenters as depicted in Figure 1 for the years beyond approximately 2022.

As depicted with the dashed green line in Figure 1, the ***ENLITENED*** Program seeks to suppress this increasing power consumption trend by enabling energy efficient datacenters with packaged integrated photonic interconnects and switches and associated novel networking architectures. Specifically, ARPA-E sees an opportunity to exploit high-density integrated photonic chip-scale I/O at <1 pJ/bit and lower latency network fabrics enabled by integrated photonic links and switching technology. This combination will make possible significant increases in overall system efficiencies and thereby halt the increasing energy demand, and perhaps even lower the aggregate datacenter energy consumption to manageable levels as depicted in Figure 1.

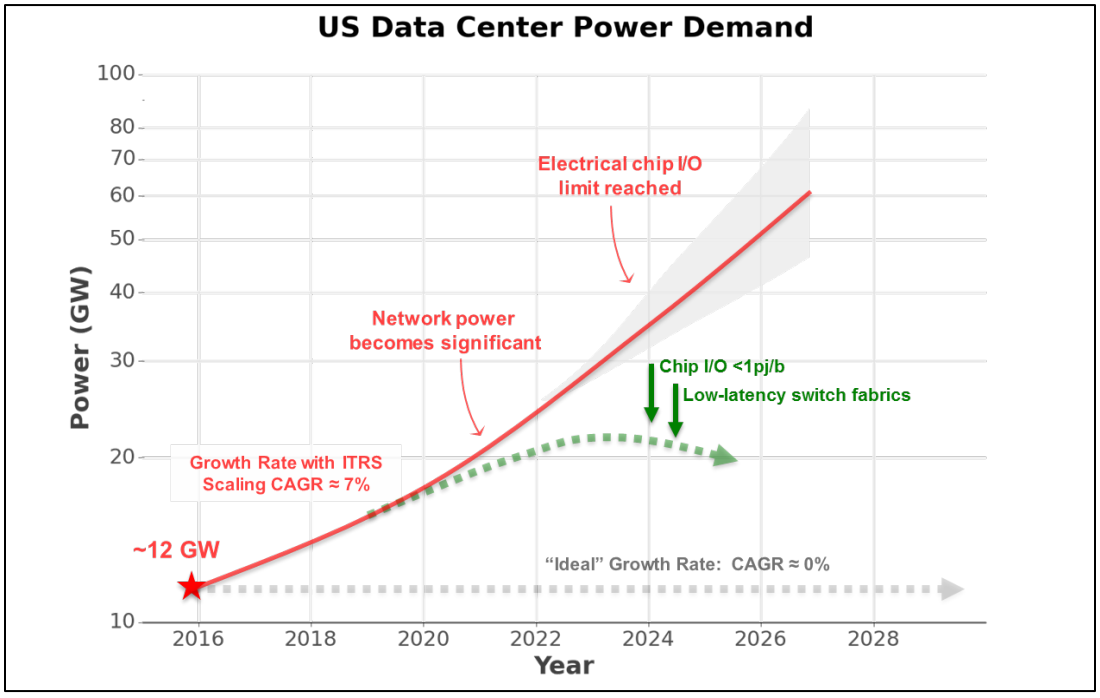

Figure 1. Aggregated power consumption of US datacenters as a function of time. The red line shows current trends and anticipates further electrical power growth due to the network becoming more significant and reaching the limits of metal interconnects. The green dotted line indicates the ***ENLITENED*** objective enabled by low-energy photonic interconnects and switching technology.[2]

[2] The estimated total electrical load of ~12 GW for 2016 is extrapolated from trends for data from 2000, 2005 and 2010 [1], and assumes a Cumulative Average Growth Rate (CAGR) of approximately 7%. The value obtained for 2016 is consistent with the prediction by the DatacenterDynamics (DCD) of 13.25 GW also for 2016. The extrapolation from 2016 to 2020 with the same CAGR yields 16.8 GW. The portion of the graph from 2020 to about 2023 assumes that network becomes a bigger percentage of energy utilization reaching as much as 50%, as predicted

Rationale for targeting IT Infrastructure vs. other sources of datacenter inefficiency

Opportunities for improving datacenter efficiency abound through potential improvements in Information Technology (IT) components, software, cooling techniques, electrical power delivery, and computing chip advances. As depicted in Figure 2, to analyze the energy consumed by a datacenter, it is useful to break up its energy efficiency into components, for example: electrical power delivery and conversion efficiency, cooling efficiency, and IT efficiency. Currently, power delivery, power down-conversion, and cooling are the subject of much interest and research. For example, companies operating very large datacenters are reporting cooling efficiencies approaching optimal values. Due, in part, to the existing strong level of investments in the infrastructure components, these elements are not considered under this FOA.

It is anticipated that future datacenter performance limits will stem from the inability of the datacenter network to keep pace with increased data loads to/from future server chip technology. Therefore, one opportunity for transformative improvement comes from enabling new network topologies and data management protocols through advanced integrated photonics systems. ***ENLITENED*** is targeting the efficiency of the IT elements of datacenters – and specifically the networking elements (interconnects and switching elements) of the datacenter.

A commonly-used measure of datacenter efficiency is PUE (Power Usage Effectiveness). PUE is defined as the ratio of total energy consumed by the datacenter (E_{total}) to that consumed by the IT equipment (E_{IT}). The total energy used by a datacenter can therefore be expressed as:

$$E_{total} = PUE \cdot E_{IT}$$

Often, the "PUE" factor lumps in all other factors such as cooling, power conditioning and delivery. Other smaller facilities-related consumption, such as lighting, is also added in certain estimations. Many companies report decreasing PUE values, yet this number does not reveal the amount of energy consumed. In most cases, the majority of the energy is consumed by IT equipment and the proportion is growing, due in part, to lower PUE achievements. It is also true that if IT energy consumption is reduced, this results in less energy consumed by cooling and power delivery in roughly the same proportion, because both are directly driven by IT energy use. It is this IT energy that is the target of the ***ENLITENED*** program.

A number of studies suggest that the IT network is an increasing fraction of the total IT power, in some case over 50% [12, 15, 18, 19]. Innovation in networking components and topologies can reduce the power consumed by the IT network, and further increase system efficiency in other ways, such as enabling a reduced percent of idle time of processor chips that consume significant energy while not being utilized [8,13] via higher overall network bandwidth and lower network latency. For example, for the datacenter schematic depicted in the bottom portion of Figure 2, higher bandwidth components could enable fewer switch nodes with higher port count per switch, in turn reducing the number of "hops" and idle time between servers, resulting in less energy consumed for the same transaction.

in [8], leading to an approximately 50% additional growth to 30 GW for 2023. After 2022, the effects of reaching the limits of metal interconnects for delivering the required throughput to the anticipated many-core chips is added. As an illustration of this last point, in [14] we observe that by 2022, assuming a constant IO B/FLOP, the metal interconnects fail to deliver the required I/O throughput to the future server chip – resulting in a further limitation of energy efficiency growth in datacenters. The shaded region indicates uncertainty in predicting future demand.

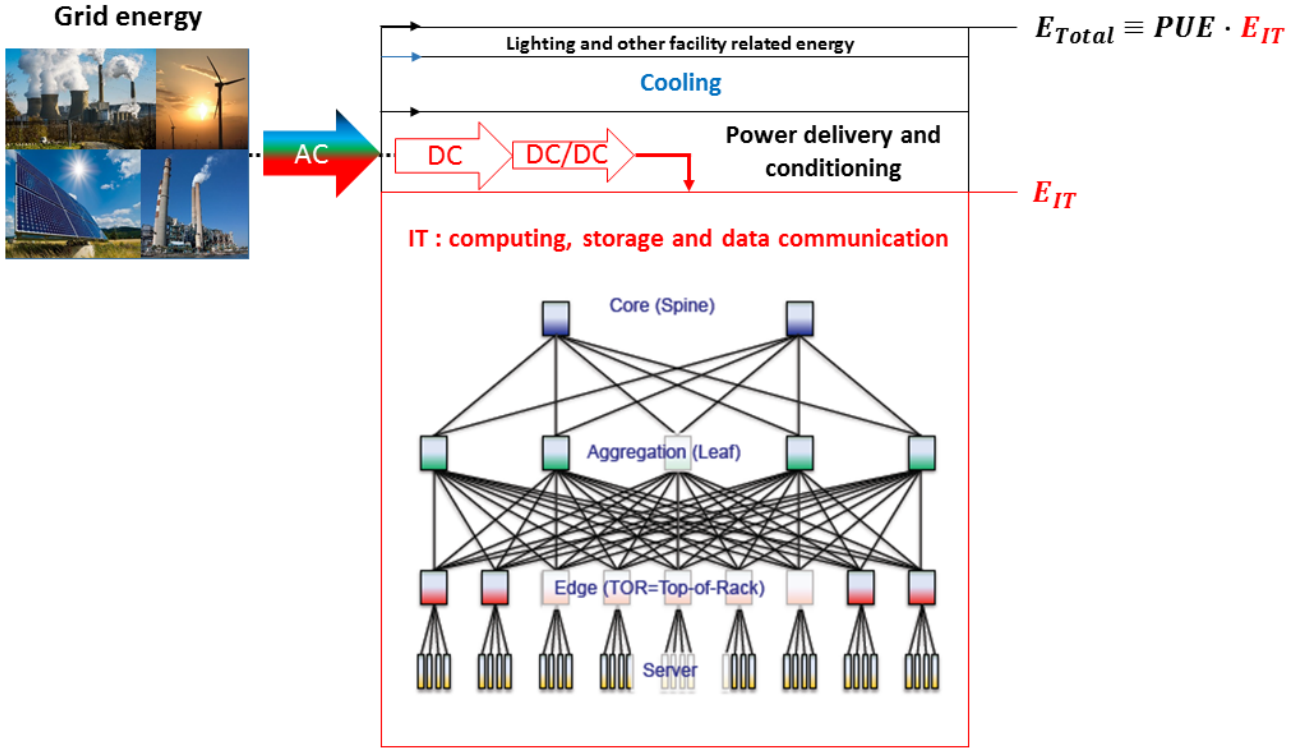

Figure 2. Schematic depiction of a datacenter and sources of energy consumption, the primary being IT, followed by cooling and power delivery and conditioning. The schematic at the bottom represents a datacenter network with servers as the bottom thin yellow boxes and switches as red, green and blue boxes.

The datacenter depicted in Figure 2 is one with a "Fat-Tree" network topology, with lines representing data interconnects between switches and servers – the lines and nodes in the diagram are "the network." Typically, switches receive, store and forward data, and play a role in data routing decisions. Within each server are central processor and memory chips mounted on printed circuit boards that perform the majority of calculations and data storage. The first level of hierarchy for data communications outside of the server is called the "Top-of-Rack" (TOR) switch, which usually has the largest number of connections compared to all other network switches. Data communication at the level of the TOR and above is commonly referred to as "east-west" and from the TOR to the processor and memory chips within server boards as "north-south." Although the "network" is often considered to be all data transactions outside of a server, in fact a large amount of data communication occurs "on-board", right up to the edge of the memory and processor chips inside server boards. ***ENLITENED*** seeks to invest in technologies that can be integrated right up to the edge of the chip (i.e., integrated directly with the chip carrier) to enable both increased performance as well as transformative gains in energy efficiency.

Meeting Future Datacenter Needs: Overcoming the Bandwidth Bottleneck

There are two areas of concern that will limit future datacenter system performance. Chip-scale metal interconnects limit the scalability of silicon-chip-based electrical switches and therefore, are a fundamental source of future datacenter inefficiency [15]. Similarly, the metal interconnects connecting server chips and memory units to each other and the network interface will not be able to keep pace in an energy-efficient way with the bandwidth demands [14]. As bandwidth and capacity requirements of datacenters grow, the limitations of electrical interconnects will become further exaggerated when compared to integrated photonics technologies that are inherently more energy-efficient and capable of having far greater bandwidth-per-channel and bandwidth density at the chip scale [14]. There is an ever-increasing fraction of datacenter power requirements and costs dedicated to the intra-datacenter network, with projections suggesting further increases [12, 13].

Although integrated photonics technologies have been shown to achieve much lower link and switch energies than metal and electronic technologies, wide-spread deployment (such as in future datacenters) must overcome significant fabrication cost barriers to replace metal interconnects on server boards and to switching elements. Consequently, ***ENLITENED***

specifically targets the development of fabrication/manufacturing techniques that address the integration and packaging of photonic data-com technologies within server environments and the development of novel networking architectures that take advantage of these photonic technologies.

Figure 3 is a comparative schematic depiction of chip-scale integrated electronic and photonic interconnects which highlights the key issues for metallic interconnects and the key challenges that must be overcome in next-generation integrated photonics technologies. State-of-the-art server or switch chips (depicted as orange squares) are mounted on chip-carriers (depicted as blue squares with black dots indicating metal connection points to the underlying board) that provide the metal trace interface to the underlying circuit board. Two key attributes determine the I/O capability of the system: the "escape bandwidth density" for the chip package and the energy-per-bit (E/b) consumed for each link. The escape bandwidth density is defined as the linear bandwidth density around the periphery of the chip or chip package and is limited by the geometric and physical constraints of the chip-carrier on pin density, number of wiring layers, cross-talk, line termination requirements, etc. Traditional multi-layer chip-carrier techniques approaches are limited to approximately 1 Tb/s/cm in escape bandwidth density [20]. Furthermore, due to the fundamental loss mechanisms in high-bandwidth electrical transport, energy requirements for high-density integrated metallic links are approximately proportional to the length of the link. State-of-the-art metal links can achieve approximately 0.5 pJ/b/cm and therefore the length of required inter-chip links is a critical concern. As server and switch chip technology continues to scale, the ability of metal interconnects to keep pace is projected to be limited [14].

In contrast to integrated metallic interconnects, integrated photonic interconnects offer the potential to provide much higher escape bandwidth and E/b performance that is essentially independent of the chip-scale interconnect distance, as depicted in the lower half of Figure 3. At the inter-chip distances considered here, waveguide and fiber transport is essentially lossless. E/b performance for optical links have been demonstrated at approximately 1 pJ/b [14] and projected performance for integrated photonic links may achieve even much lower levels of E/b [14]. What's more, due to the ability of waveguides to carry highly multiplexed signals, and be packed in high density with little crosstalk, the potential escape bandwidth of integrated photonic interconnects far exceeds that of metal interconnects.

As depicted in Figure 3, the critical implementation challenge for integrated photonic links is the transduction from the electronic to optical domains (E→O) at the transmitter and the inverse (O→E) at the receiver. Integrated photonic technologies may achieve this through direct modulation of a laser diode (LD) and then detection of the photonic signal with a photodiode (PD) at the receiving end, as depicted, or other types of photonic modulation and detection of laser signals may be employed. In any integrated photonic scheme, the electronic/photonic transductions present the critical integration and packaging challenges for efficiently interfacing to the server or switch chip and the waveguide structure that comprises to actual interconnect fabric.

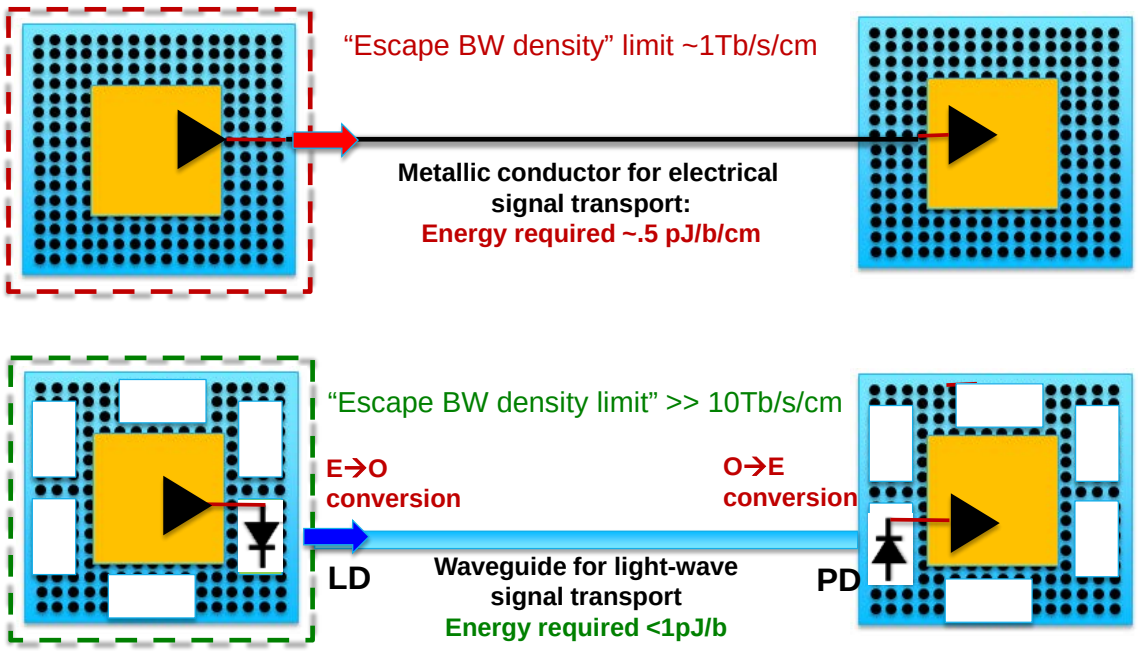

Figure 3. Schematic demonstrating an "interconnect" between processor chips and limitations using electrical metal links (above) and optical waveguide (below) interconnection technologies. Due to practical and dimensional design limitations of about ~200 W/processor chip and other physical constraints of metal wires [14] , data transfer rates into and out of processor chips are limited and distances over which electrical signals can travel are limited due to signal loss (~0.5 pJ/b/cm). Optical interconnects are virtually lossless as a function of distance and enable a much greater "escape bandwidth" to increase the chip input/output data rates, however, E/O and O/E (O= optical, E = electrical) conversion steps must be efficient to minimize heating and achieve higher overall efficiency.

C. PROGRAM OBJECTIVES

The overall objective of the ***ENLITENED*** Program is to create new technology platforms, components, and evaluation methods to enable a > 2-fold improvement in the energy efficiency of datacenter ICT infrastructure.

Although the energy needs and functions of datacenters vary, ARPA-E posits that efficiency can be improved across all datacenter applications via the incorporation of advanced integrated chip-scale photonic interconnects and switching technologies. Although no one universal metric for datacenter efficiency may be applied to all cases, ARPA-E considers a doubling of the average "Transactions/Joule" (interpreted as a "wall-plug" efficiency for the entire ICT system of the datacenter) to be a plausible general target. One objective of ***ENLITENED*** is to tie different use-cases and data loads to a standard system-level metric: "Transactions" in the numerator and energy consumed by the entire datacenter (E_{IT} in Figure 2) during those transactions in the denominator. This would enable greater clarity for innovation to optimize for energy at the system level and evaluate other performance trade-offs. Examples of intermediate system-level metrics from which "Transaction" units can be derived could be measures such as FLOPS/Joule and bytes/FLOP for a series of queries, the number of intermediate switch hops and overall time for a given byte or data packet to travel between servers, and so forth. Such analysis will enable objective measures of datacenter energy-efficiency and easier comparison across computational models and hardware sub-systems.

ARPA-E anticipates that this level of datacenter system efficiency improvement will stem in part from the fundamentally lower energy-per-bit achievable with integrated photonic links. The remaining significant efficiency improvements will stem from changes in network architecture enabled by the ultra-high bandwidth density of photonics and low latency photonic switch fabrics. The closer integration of switching and computational elements that will emerge from ***ENLITENED*** will potentially blur the distinctions now made in techno-economic analyses of server, switch, and transceiver costs for datacenters. Nevertheless, ARPA-E estimates that an effective cost structure is needed to bring ***ENLITENED*** technologies to market by putting packaged technologies on a path to < 10 cents/Gb/s for integrated transceivers and aggregate switch bandwidths > 50 Tb/s, capable of radix > 128 [18]. Early market insertion opportunities may be possible in other areas such as high-performance computing (HPC), where performance and energy efficiency are prioritized over market cost metrics.

At the end of the program, ARPA-E expects datacenter sub-systems and architectures, supported by detailed modeling and simulation that achieve a 2-fold or greater average improvement in energy efficiency as well as a path to commercialization.

D. PROGRAM STRUCTURE AND TECHNICAL CATEGORIES OF INTEREST

ARPA-E anticipates a 2-phase, 4-year program for ***ENLITENED***, however ***the initial period of performance for awards resulting from this FOA will be for Phase I and will not exceed two years. Phase 2, also a maximum of 2 years, may only be awarded contingent upon successful completion of Phase 1 and will be subject to the availability of appropriated funds. All Applicants must provide budgets and task descriptions that cover both Phase 1 and Phase 2 (a maximum of 4 years total).*** Details are described below. The program objective is a packaged hardware solution consistent with a computational datacenter model and/or calculations that demonstrate the potential to achieve an average 2-fold improvement in energy efficiency for a broad range of use cases.

Applicants are encouraged to have vertically-integrated teams that bring together technology, packaging and integration, and network architecture expertise. ARPA-E may also consider **partial solutions**, for example, of **a component that has**

a critical path to being packaged and integrated into datacenter sub-systems, a new integration and packaging approach that could have broad impact, or a computational modeling approach that can be used, together with realistic hardware specifications and datacenter constraints, to design solutions that meet the program metrics.

Full Solutions

Vertically-integrated teams are expected to develop the integration and packaging concepts needed to enable chip-scale ultra-high-density photonic link and switching technologies within future datacenters. It is anticipated that, depending on the integrated photonic technologies proposed, potential solutions will have differing levels of hybridization/integration between electronic and photonic elements. Applicants must propose, and articulate a rationale for, an integrated technology platform (that includes links, transceivers, switching elements, control strategies, and all interfacing technologies and techniques) needed to couple the photonic elements to the anticipated electronic computing elements at the chip, chip package, board and intra-rack-levels, as appropriate in future datacenters. Applicants must present a network structure enabled by the integrated technologies that would lead to transformative energy-efficiency improvements.

Implementation of Phase 2 for the ***ENLITENED*** program will depend on satisfactory progress across the program during Phase 1, subject to the availability of appropriated funds and other factors. Project teams will be expected to deliver analyses, perform sub-system experiments (as detailed below), and refine a Phase 2 project plan that will enable ARPA-E to make informed selections for the second phase. If ARPA-E selects a project to continue to Phase 2, the team will develop more significant experimental validations of their concept that culminates with critical analyses, simulation, and hardware-in-the-loop experiments that will enable valid assessment of the technologies and progress towards commercialization.

An expectation of Phase 2 success is demonstration of the concept at a minimal system level. For example, a new transceiver or switch architecture must take into account all system elements that consume energy such as clock and data recovery, re-timing circuitry, and other system level inefficiencies. Applicants pursuing full solutions should propose goals that are ambitious and exceed the current goals of other commercial efforts in terms of energy efficiency, such as those of the Consortium of On-Board Optics (COBO) [11].

Partial Solutions

ARPA-E anticipates that some teams may not be fully “vertically integrated” or “full solution” as described above, but may provide a valuable partial solution that could contribute significantly to achieving the program objectives. Such projects may have a hardware solution that is significantly less than a full system, for example a laser or modulator array that is to be packaged into a transceiver or switch. Others may have novel datacenter network architectural notions that promise to take full advantage of emerging photonics technologies in future datacenters using realistic components and show a clear commercial path to achieving the program metrics, yet not currently be part of a team that has the matching hardware solutions.

ARPA-E anticipates that partial solutions may be more suited to Phase 1, but a Phase 2 path must at least be addressed, albeit in less detail than vertically integrated teams. As listed in Section IV.C “Content and Form of Concept Papers,” and in the Concept Paper template document, teams pursuing partial solutions are expected to submit Concept Papers of no more than 4 pages in length. For certain common components that have well-demonstrated performance, for example, high efficiency/bandwidth lasers, modulators, couplers, etc.; a two-year window may be critical for integration and packaging into existing product lines.

For teams pursuing partial solutions to design and simulate novel network architectures only, Submissions must explain how their proposed network design will leverage projected performance advantages of photonics-enabled sub-systems. Teams must use realistic [10] component technology specifications and techno-economic analyses and address all technical performance targets, even though they may not have experimental measurements as part of their Phase 1 project plan.

E. TECHNICAL PERFORMANCE TARGETS AND DELIVERABLES

The ***ENLITENED*** technical performance targets are designed to demonstrate the ability of any proposed solution to meet the transformative energy efficiency objective described in this FOA when compared to *anticipated* state-of-the-art electrical

interconnect and networking technologies. In other words, the benefits of the proposed solution must be in addition to any future improvements in CPU, memory, or other datacenter electrical components or architectural efficiency enhancements (e.g., anticipated silicon CMOS scaling, or energy-proportional computing strategies).[3]

Applicants must address the overall 2x average system efficiency improvement target in the context of datacenters at all scales (large, small, multi-tenant etc.), and diverse distributions of data loads. The justification of the efficiency improvement should include, in addition to datacenter diversity and data loads, the effects of format conversion between single-mode and multi-mode for North-South and East-West communications, number of E-O-E (or O-E-O) conversions, etc., if appropriate. Table 1 provides targets for Program Metrics. **The target in Metric 1.1 is the overall program objective and must be addressed in all submissions. If an Applicant proposes a concept to meet the target in Metric 1.1 without meeting the targets in Metrics 1.2-1.5, the Applicant must provide alternative targets for Metrics 1.2-1.5 and still provide a rationale showing how Metric 1.1 will be met.** Comprehensive simulation results are not expected in concept papers, but a quantifiable argument for how the 2x efficiency improvement (excluding cooling and power down conversion effects) will be reached must be presented. Metric 1.6 refers to the technology economic assessment. **For Full Applications**, a more comprehensive justification for achieving Metric 1.1 is expected. **Applicants should provide a detailed energy and bandwidth budget table(s) in the Technical Volume that include all energies for each subsystem required to evaluate the transaction per Joule figure of merit**. Example computational parameters and types of analyses that may be appropriate to include as part of energy budget and impact projection calculations detailed in Full Applications include:

- Definition of a transaction or types of transactions(e.g., search, image analysis);
- Size of, and/or range of sizes and size distributions of, transactions – in, e.g., bytes or other appropriate measures;
- Number of operations per transaction and transaction time;
- Energy consumed during a transaction (referenced to "wall plug" energy, with assumptions provided);
- A simulation and/or modeling analysis, based on the above parameters, or other appropriate measures, that compares transaction processing performance in current data center networking architectures with the proposed future datacenter network architecture enabled by the proposed integrated photonic platform, with reasonable assumptions about other future datacenter components, such as Moore's law-type CPU scaling. This analysis could specify and include, as appropriate, key datacenter network properties, such as: topology and number of nodes; number of servers; types of CPU chips in servers; memory; number of switches and their aggregate bandwidth and port-count (radix); description of photonic links and their length and bandwidth; description of any required copper links and their lengths and bandwidths; total hops per bit and latency across the datacenter; etc.
- Comparative analyses that estimate/project the performance enhancements provided by integrated photonics across a range of important data center service types and data center sizes.
- Comparative analyses that estimate/project energy consumption as a function of work load – to highlight, e.g., the energy-consumption benefits of the proposed architecture under quiescent and reduced traffic conditions.
- Comparative analyses, that, to the extent possible, project impact at a national level, under the assumption of wide-scale deployment of the proposed technology

In general, any provided energy and bandwidth budget tables should inform and illustrate how Metric 1.1 could be met with the proposed concept, and how well the projections can be generalized to other datacenter types to understand the breadth of the impact. There may be components of full solution concepts that are required for operation, but are not specifically addressed such as energy required for clock and data recovery (CDR), or memory refresh. Although the task descriptions may not address elements such as these explicitly, the energy budget estimates should take into account how the overall datacenter architecture will be affected by other critical elements, or how these may be bypassed.

Applicants should identify the specific datacenter structures targeted for innovation. Examples of this would be creating a switch architecture that combines many top of rack and core switches into one and meeting the FOA metrics by reducing the number of "hops" for the majority of data packets, while increasing bandwidth and processor utilization. This example may include both "short" and "long reach" hardware elements that could be addressed by both VCSELs over multi-mode fiber as well as Silicon photonics, for example. Another example would be targeting the infrastructure "below the Top of Rack", which would likely be dealing with mostly "short reach" communications that take place between CPUs and memory elements inside the "server unit" of current datacenter designs. Of interest are approaches that aim to disaggregate

[3] For example the system efficiency improvement could be shown by exhibiting simulations with and without the photonic technologies but both including other future improvements

datacenter elements and blur boundaries between servers, as such architectures are predicted to have advantages in performance and energy efficiency and potentially maintenance and reliability.

Table 1: Program Metrics

ID	Metric	Target	Description
1.1	System Energy Efficiency	>2x "Transactions /Joule"	• Applicants must show how their proposed integrated photonics and new switch network provides a credible path to >2x efficiency.
1.2	Link Demo	< 2 pJ/bit	• Must show a path to <1 pJ/bit for inter-chip board-level links of length 1 to 100 cm.
1.3	BW Density Demo: Chip or Chip carrier I/O	> 1Tb/s/cm	• The proposed packaging approach enables aggregate bandwidths of 10 Tb/s or greater for future server or switch chips.
1.4	Board Level I/O	Applicant-provided	• Board-level I/O will be determined by specifics of proposed network architecture.
1.5	Switch Concept	Applicant-provided	• Must show how integrated photonics enables an efficient architecture (> 2x efficiency from item 1) • Specify required switch metrics (e.g., radix > 128, end-to-end Energy/bit, etc.) • Define how store and forward steps, latency affect overall system performance and energy • Switch concepts must show path to > 50 Tb/s aggregate bandwidth per switch
1.6	Technology Economic Assessment	Applicant-provided	• Show cost effective packaging compatible with server/communication equipment • IP enables path to < $0.10/Gb/s at inter board level • Network cost: Applicants must provide TEA to eventual wide scale deployment of the proposed IP –architecture

Deliverables

"Full solutions" applicants must experimentally validate the listed program metrics in Table 1. These deliverables must make a convincing case for commercial transition. Applicants must provide a technical approach and milestone plan for achieving the deliverables. Full solution Applicants must propose the following five deliverables for Phase 1:

1. Laboratory demonstration of a chip-level integrated photonic link with a clear path to < 1 pJ/b for 1-100 cm link lengths and an integrated interconnect demonstration that shows a path to 10 Tb/s for future server and switch chips.
2. Laboratory demonstration of an integrated switching concept that exploits the proposed photonic link concept and enables the proposed networking architecture.
3. Network architecture and datacenter simulations that are consistent with the experimentally-measured results from 1 and 2 and show a path to at least a 2-fold improvement in average energy efficiency.
4. A technical plan for the packaging and integration of all system components considered and methods of manufacturing identified with accompanying techno-economic analysis.
5. Submission, at the end of Phase 1, of an updated, detailed Phase 2 plan that builds upon the Phase I deliverables and refined simulations and proposes prototype integrated packaged sub-systems with sufficient numbers of links and switching elements to validate the proposed concept.

"Partial solutions" applicants must specify their own quantifiable milestones and deliverables as appropriate, but consistent with showing well-defined paths to the overall program objective via detailed simulations.

Additional Clarifications for Partial Solutions

Applicants developing partial hardware solutions must address all of the above metrics during the ***ENLITENED*** program, even though the metrics may not be directly measured as part of proposed project plans. Applicants must provide intermediate metrics and relate those to system level performance. For example, if a partial solution that involves the packaging and integration of a laser array into an existing product line were proposed, other metrics would be expected such as modulation rate, total wall-plug power per port, coupling efficiency, error-free bit-rate per port, inside of a prototype packaged system.

Similarly, for partial solutions consisting of modeling and simulation (not experimental work) , teams must demonstrate the ability to communicate and understand what industry customers require, as well as the inner workings of both enterprise scale and hyper-scale cloud datacenters. Teams must be technically competent to perform state-of-the-art simulations to evaluate experimental results from hardware developers and simulate datacenter operation by varying the size of the datacenter (number of servers and switches), the topology of the datacenter (Clos-type, HyperX, All-to-All, Fat-Tree, etc.) and software-defined datacenter algorithms.

Additional Clarifications for Computational Modeling

All Submissions that include network architecture/system modeling, whether vertically integrated teams, or partial solutions (without experiments), must:

- Develop models that show innovative ways to reach a global energy efficiency improvement > 2 x Transactions/Joule. This universal metric must be related across datacenter sizes, with different data load size distributions, activity patterns and computational metrics. Other measures such as FLOPS/Watt, FLOPS/Joule, bytes/FLOP must be related to expected global gains.
- Specify topologies and compare to the state-of-the-art.
- Specify data loads, where the data will be obtained and how relevant these data sets and traffic patterns are to evolving datacenter trends.
- Have models that use realistic system components, including the most advanced and anticipated integrated photonics capabilities being developed by other teams in the ***ENLITENED*** program and enable modularity to evaluate how the global metric is affected by changes in components and topologies. Physics-based constraints must be taken into account as well, such as the trade-offs in bandwidth density per area, heat generated per area, etc.
- Provide a techno-economic justification for the computational model to evaluate the feasibility of adoption. This must take into account realistic models of CAPEX and OPEX and creative ways of how new datacenters could be manufactured or retrofit with emerging technologies in a cost effective way. Thought leadership for creative ways that industry can adapt to limitations of power, real-estate and other constraints are welcome.