

Testimony of
Mark P. Mills, Executive Director of the National Center for Energy Analytics
Before
Subcommittee on Economic Growth, Energy Policy, and Regulatory Affairs
U.S. House Committee on Oversight and Government Reform
Hearing on:
“America’s AI Moonshot: The Economics of AI, Data Centers, and Power Consumption.”
April 1, 2025

Good afternoon. Thank you, Mr. Chairman and members of the Committee, for the opportunity to testify. I’m the Executive Director and founder of the National Center for Energy Analytics, a Distinguished Senior Fellow at the Texas Public Policy Institute. I am also a Faculty Fellow at the McCormick School of Engineering at Northwestern University, member of the Advisory Board of [Yotta](#), a consortium focused on digital infrastructure. And, for the record, I’m a non-operating partner in a venture fund focused on energy-related software, and I’m also a director of a public oil-field services company.

The words “revolution” and a “pivot in history” are frequently overused and mis-used. But, by now, it’s clear that such words *are* appropriate in describing the emergence of AI as a practical tool deployed in the infrastructure of the cloud. Debates about risks and benefits, over what’s hype and real, will continue. But the fact of an unfolding revolution is evident across nearly every domain.

Evoking the Moonshot *is* appropriate in terms of echoing the great power competition of the 20th century, and because of AI’s strategic implications. But the central promise of useable AI at scale is an economic one; and specifically in how AI will boost overall U.S. productivity

While we’re currently preoccupied with the politics of cutting government waste and the deficit, the Holy Grail is in finding a productivity boom. Consider: If, instead of the current Congressional Budget Office forecast for anemic annual productivity growth, AI raises that growth merely back to the late 20th century’s average -- that would, by 2030, add a cumulative extra \$10 trillion to the U.S. economy. And it could be more.

It’s that promise that underlies the private-sector’s drive to spend hundreds of billions, likely totaling a trillion dollars before 2030, to build more data centers at the epicenter of the cloud and where AI is democratized. The next *one* year’s spending on data centers will alone exceed the entire decade of government spending on the moonshot.

The unprecedented digital construction underway has led to the rediscovery of a basic truth: All software exists inside hardware that, in turn, uses energy, a lot of it. Each digital byte uses an infinitesimal amount of energy. But here we find salience for the euphemism that quantity has a quality all its own. Again, in moonshot terms: the amount of energy used to launch a rocket is consumed every day by just *one* AI-infused datacenter.

Or in monetary terms: \$10 billion worth of data centers will consume some \$10 to \$20 billion of electricity over a decade of operation. For context, \$10 billion of EVs will consume one-tenth as much electricity.

Such realities invariably invoke a response that we should focus on, even mandate, more energy efficiency. This is a rabbit-hole subject: suffice to say that efficiency is precisely why data power demands have risen. Consider: if one smartphone today operated at 1984 computer efficiency, that one phone would use more electricity

than this entire building. Instead, efficiency gains led to billions of phones and thousands of data centers. This is the oft-noted so-called Jevons Paradox which, I note he knew in the 19th century wasn't a paradox, but a reality.

Credible forecasts see, within a half-decade or so, new data centers needing somewhere between 70 GW and 130 GW *more* power capacity built than grid planners thought they'd need only a couple of years ago. This is the equivalent of *adding* generating capacity equal to the entire state of California, and in a half-dozen years.

Thus, a key question: Will there be enough power to fuel the velocity of the cloud's expansion to stay at the forefront of this revolution? Reality, and arithmetic, show that the needed power won't come from squeezing more out of existing assets, or stopping coal plant retirements, or utility-scale wind and solar, nor from building new nuclear plants. All that will help, but it won't be enough, soon enough.

Most of the new power will come from natural gas combustion turbines and engines. Those can be and *are* being built rapidly. That implies the need for a lot more natural gas too. One might be tempted here to invoke the politically charged, "drill-baby-drill." But politics aside, history shows that the wells and pipelines *can* be deployed fast enough at scales needed.

This is where the moonshot metaphor fails. The Apollo program had the government spending taxpayer's money to put a dozen men on the moon. Ensuring we achieve lift-off for an AI productivity boom will come from *private* capital be rapidly deployed to build the needed fuel systems for the digital infrastructures.

The government's role to ensure U.S. industries can win this great-powers race should start with asking private-sector players to identify the regulatory and policy-centric impediments to rapid deployment at the scales needed. And then removing those impediments. That could be the real moonshot challenge in all this.



Extended testimony for the record

[N.B. Reprinted from "Unplugging Growth? AI, the Cloud & Electricity Demands," By Mark P. Mills, *Dakota Digital Review*, October 21, 2024.]

The year 2024 will go down as the pivot when the growth in U.S. electricity demand reverted to normal. After a two-decade interregnum of no growth, many forecasters, including those in the electric utility industries, thought that was the new normal. Planning for a static future is quite different from meeting the needs of robust growth.

Now we find myriad power-demanding hot spots around the country, from Georgia and Virginia to Texas and California and a huge swath of the northeast, as well as states like North Dakota reporting radical increases in requests for power, and soon. In nearly every case, the demands that will materialize in the next one to three years vastly exceed current plans to build sufficient generating capacity, of any kind.

The reason for this surprise? If you believe the headlines and hype, it was because of the staggering electricity demands coming from the artificial intelligence (AI) boom. Thus we saw headlines echoing a new trope, such as "Can AI Derail the Energy Transition?"¹ With AI as the *piñata* for putting the transition at risk, Silicon Valley's potentates scrambled to explain what happened and to endorse—in some cases fund—new nuclear power

ventures, as it quickly became clear that favored wind/solar power can't come close to meeting the scale of demand coming soon enough, reliably enough or cheap enough.

It is true that AI is very energy intensive, as chronicled in one of my previous articles in Dakota Digital Review: "AI's Energy Appetite: Voracious & Efficient."ⁱⁱⁱ It's also true the emergence of more useful applications for AI tools is leading to a rapid increase in installations of energy-hungry data-center hardware. But, while AI is driving new demands for power, the reality is that conventional computing and communications the existing Cloud account for well over 90 percent of those forecasted rises in power demands.

It is also true that other sectors boost electric demand including, though far less impactful, electric vehicles. The more relevant and big-demand wildcard is whether and how soon goals to reshore manufacturing will be realized. Of course, if AI delivers on its promises, economic growth across the board will be greater than earlier forecast. Then there'll be the old-fashioned 'problem' of economic growth itself boosting electricity uses.

AI is accelerating a trend that was already underway. In the coming decade, even as AI takes up an increasingly larger share of total digital power appetites, the uses of conventional silicon will also expand and are forecast to still account for more than 70 percent of total digital electricity consumption by 2030.ⁱⁱⁱ And those expected digital demands are at scales shocking to local and state utilities and regulators everywhere. In North Dakota, The Bismarck Tribune reported near-term plans for a handful of new data centers that would alone consume as much power as all that state's residences.^{iv}

This all distills to a rediscovery of a simple reality: Planners and policymakers need to refresh their views about how to meet society's electricity and energy needs based on *growth*. One thing is clear, if planners fail to ensure adequate supply, and there isn't enough electricity, then the growth won't happen. Planners and forecasters might want to know whether this is a bubble or a trend and, in the latter, just how much more global and local digital demand is yet to come?

The idea of something as remarkable as a global, interconnected information and knowledge-supplying and knowledge-creating infrastructure is not new. Back in 1938, H.G. Wells's novel *World Brain* imagined such a thing during that era of revolutionary globe-spanning telephone and telegraph networks. A couple of decades later, in 1962, the Council on Library Resources launched a project to imagine the "Libraries of the Future" and asked MIT computer scientist J. C. Licklider to provide a technology how soon goals to reshore manufacturing will be realized. Of course, if AI delivers on its promises, economic roadmap, the same year that President John F. Kennedy made his famous "within this decade" moonshot speech.

Futurists of the 1960s were not only inspired by rocket ships but also by the arrival of the mass production of transistors, the proliferation of mainframe computers and the first communications satellites. In 1962, there were about 2,000 computers in the world, then a seemingly remarkable number given that, just two decades earlier, there were only two, the ones built secretly here and in England during World War II.^v

The possibilities of a computerized "world brain" were clear to the prescient Professor Licklider, and he knew that it would be much more than a mere electronic "library," that it would be an information system. He also wrote that it wouldn't be possible without radical technological progress, including new inventions.^{vi} He was right. It turned out to be far easier and cheaper to put a dozen men on the moon, than to build an information infrastructure

connecting, in real time, billions of people on earth. But now, global Cloud-centric information services businesses are approaching \$1 trillion a year,^{vii} built on tens of trillions of dollars of hardware, providing an entirely new classes of commerce, much more than libraries of cat videos or AI-created “deep fakes” for fun (and malicious meddling). At the core of that infrastructure are the massive data centers—the thousands of “warehouse-scale” computers, *each* consuming more power than a steel mill. For students of the history of technology forecasting, none of the futurists of only a few decades ago anticipated the nature and scale of today’s Cloud infrastructure.

One thing that can be forecast today is that we’re at the end of the beginning, not the beginning of the end, of building it all out. Since the Cloud is an energy-using, not energy-producing infrastructure, the implications are consequential. The future will see far more, and far bigger, data centers—the digital cathedrals of our time.

The world’s first modern data center was built in 1996 by Exodus Communications in Santa Clara, California, a half-century after the first computer centers. The Exodus facility was a 15,000-square-foot building dedicated to housing the silicon hardware for myriad Internet Service Providers (ISPs).^{viii} By coincidence, that year also witnessed a massive power blackout over seven western states, shutting down everything including already internet-centric businesses and services. That is, except for those ISPs that had equipment housed in the Exodus data center, which remained online with its backup power system, enshrining one of the key benefits and design features of data centers.^{ix}

The term “data center” is an anemic one, failing to telegraph the nature and scale of how those mammoth facilities have evolved. It’s a little like calling a supertanker a boat. The data center term evolved naturally from the early days of computing, when buildings had rooms called computer centers in which banks, universities and some companies placed mainframes. But when it comes to the physical realities of scale, including energy use, there is a world of difference between a shopping center and a skyscraper. So too data centers, more aptly skyscraper-class energy demands are, in commerce terms, the digital cathedrals of our time.

In 1913, the world witnessed the completion of the world’s first skyscraper, the 792-foot-tall Woolworth building, then the world’s tallest habitable structure. It had taken 600 years to best the 524-foot tower of England’s Lincoln Cathedral (completed in 1311). Public awe wasn’t only inspired by the height but also about the technologies—electricity, elevators, engines—that made it possible in the first place, and in particular the commercial implications of such structures. Hence, in 1913, The New York Times enthusiastically declared that building a “cathedral of commerce.”^x

In the three decades since the Exodus data center, the expansion and networking of digital cathedrals has led to the creation of an entirely new, and now essential infrastructure, the Cloud. It is an infrastructure that, even more so than in 1996, must be insulated from the vicissitudes of the (un)reliability of public power grids. But unlike skyscrapers, digital cathedrals are essentially invisible in daily life. That may help explain how easy it is to believe a popular trope that the digital revolution somehow promises a kind of dematerialization of our economy—that the magic of cyberspace, or instantaneous banking and advice (from mapping to shopping), has de-linked economic growth from hardware and energy use. It has not.

Just as skyscrapers grew in scale and proliferated in number, so too have data centers, but only far more so. There are more than 5,000 enterprise-class data centers in the world now, compared to 1,500 “enterprise-class” skyscrapers that is, Woolworth-sized and bigger (in square feet).^{xi} Smaller data centers number some 10 million.^{xii} In the U.S., the square footage of data centers—under construction or planned for the next few years—is greater than the entire existing inventory. Today’s typical digital cathedrals is bigger by some tenfold, or more (in square feet) than those of three decades back. And each square foot of data center inhales 100 times more electricity than a square foot of a skyscraper.^{xiii} The latter reality is intimately tied to the rise in computing horsepower, a trend that AI accelerated.

Because the Cloud is an information infrastructure, instead of square feet, one could count its growth and scale in terms of the metric of data traffic, the bytes created, moved and processed. Today, a few days of digital traffic is greater than the annual traffic of just 15 years ago. And the growth hasn’t slowed; indeed, with AI, the appetite for and use of bytes has accelerated data traffic in a way that’s impossible for the physical traffic of humans and automobiles. That unprecedented reality has implications when it comes to forecasting energy consumption.

The energy cost to move one byte is miniscule and still shrinking, but there is an astronomical quantity of bytes. And the creation, transport and storage of bytes—that is, their physical and energy costs—rests with the power-using transistors that create, store and move bytes. Civilization now fabricates more than 10,000 times more transistors annually than the combined number of *grains* of wheat and rice grown on the planet.^{xiv}

The nature of the Cloud’s energy appetite is far different from that of many other infrastructures, especially transportation. For the latter, consumers literally see where 90 percent of energy is spent when they fill up their gas tank. (The other 10 percent or so is consumed in producing the hardware in the first place.) When it comes to smartphones or desktops, more than 90 percent of energy is spent remotely hidden, so to speak, in the electron-inhaling, digital cathedrals and in the sprawling information superhighways.

The physics of transporting bits leads to a surprising fact: An hour of video using the Cloud infrastructure uses more energy than a single person’s share of fuel consumed on a 10-mile bus ride.^{xv} That leads, at best, to a tiny net energy reduction if someone Zooms instead of commutes in a bus, and a net increase if a student, say, Zooms instead of walking or bicycling to class. But the fact is there are exponentially more uses for Zoom and all forms of software services, than for replacing older energy-using alternatives.

Apps—application-specific programs—provide a window on the appetite for software services. The power of apps resides in the fact that an easy-to-use interface (on a smartphone) taps into, seamlessly and invisibly, the Promethean compute power of remote data centers to provide all manner of services and advice that have become a staple of everyday life for billions of people. There are already millions of apps competing to meet consumer and business needs and desires. Those services, in every sector of the economy, are what collectively drive the scale of data centers that aren’t just measured in bytes or square feet, but now more commonly in megawatts and even gigawatts. Hidden from sight within *each* of the thousands of nondescript digital cathedrals, there are thousands of refrigerator-sized racks of silicon machines, the physical core of the Cloud. *Each* such rack burns more electricity annually than do 50 Teslas.

The global information infrastructure has grown from non-existent several decades ago to one now using twice as much electricity as the entire country of Japan. And that's a stale estimate based on the state of hardware and traffic of several years ago. Some analysts claim that as digital traffic has soared in recent years, efficiency gains have muted or even flattened growth in data center energy use.^{xvi} But such claims face countervailing factual trends.

Over the past decade, there's been a dramatic acceleration in data-center spending on hardware and a huge jump in the power density of that hardware, not least now with the advent of widely useful but energy-hungry AI. How much more power the Cloud infrastructure will demand depends on just how fast innovators innovate new uses and services that consumers and business want. The odds are the pace of that trend will exceed the pace of efficiency gains in the underlying silicon hardware. The history of the entire century of computing and communications shows that demand for bytes has grown far faster than engineers can improve efficiency. There's no evidence to suggest this will change, especially now with the proliferation of AI, the most data-hungry and power intensive-use of silicon yet invented.^{xvii}

As with all energy-intensive industries of every kind, ultimately any such business will seek to locate to paraphrase the great Walter Wriston's aphorism about capital—"where it's welcome and stays where it's well-treated." For data centers, that translates into sufficient power, and soon, that's both reliable and cheap enough. Thus, we should expect to see the rising attractiveness of those great shale basins, from Texas to North Dakota, where the on-site surplus of natural gas can be rapidly translated into a surplus of digital power. ☉

ⁱ <https://oilprice.com/Energy/Energy-General/Can-AI-Derail-the-Energy-Transition.html>

ⁱⁱ <https://dda.ndus.edu/ddreview/ais-energy-appetite-voracious-efficient/>

ⁱⁱⁱ <https://www.goldmansachs.com/insights/articles/AI-poised-to-drive-160-increase-in-power-demand>

^{iv} <https://www.govtech.com/computing/north-dakota-prepares-for-data-centers-to-come-online>

^v "Computer History for 1960." Computer Hope, November 30, 2020. <https://www.computerhope.com/history/1960>.

^{vi} Licklider, J.C. *Libraries of the Future*. Cambridge, Mass.: M.I.T. Press, 1965.

^{vii} <https://www.idc.com/getdoc.jsp?containerId=prUS52343224>

^{viii} Mitra, Sramana. "Anatomy of Innovation: Exodus Founder B.V. Jagadeesh (Part 1)." October 18, 2008.

<https://www.sramanamitra.com/2008/10/13/entrepreneurship-and-leadership-through-innovation-3leaf-ceo-bv-jagadeesh-part-1/>.

^{ix} Mitra, Sramana. "Anatomy of Innovation: Exodus Founder B.V. Jagadeesh (Part 1)." October 18, 2008.

<https://www.sramanamitra.com/2008/10/13/entrepreneurship-and-leadership-through-innovation-3leaf-ceo-bv-jagadeesh-part-1/>.

^x Sutton, Philip. "The Woolworth Building: The Cathedral of Commerce." The New York Public Library, April 23, 2013. <https://www.nypl.org/blog/2013/04/22/woolworth-building-cathedral-commerce>.

^{xi} <https://data-economy.com/data-centers-going-green-to-reduce-a-carbon-footprint-larger-than-the-airline-industry/> <https://www.skyscrapercenter.com/countries>

^{xii} <https://www.statista.com/statistics/500458/worldwide-data-center-and-it-sites/>

^{xiii} "Data Center Power Series 4—Watts per Square Foot." Silverback Data Center Solutions, November 15, 2020.

<https://teamsilverback.com/knowledge-base/data-center-power-series-4-watts-per-square-foot/>. Gould, Scott. Plug and Process Loads Capacity and Power Requirements Analysis. NREL, September 2014. <http://www.nrel.gov/docs/fy14osti/60266.pdf>.

^{xiv} Author calculation: credit for the idea of comparing transistors produced to grains grown belongs to: Hayes, Brian, "The Memristor," American Scientist, March-April 2011.

Annual transistor production from: Hutcheson, "Graphic: Transistor Production Has Reached Astronomical Scales," IEEE Spectrum, April 2, 2015.

^{xv} "Does online video streaming harm the environment?" Accessed April 7, 2021. <https://www.saveonenergy.com/uk/does-online-video-streaming-harm-the-environment/>.

^{xvi} Jones, Nicola. "How to Stop Data Centers from Gobbling up the World's Electricity." Nature. September 12, 2018.

<https://www.nature.com/articles/d41586-018-06610-y%20>.

^{xvii} Hao, Karen. "The Computing Power Needed to Train AI Is Now Rising Seven Times Faster than Ever Before." MIT Technology Review, April 2, 2020. <https://www.technologyreview.com/2019/11/11/132004/the-computing-power-needed-to-train-ai-is-now-rising-seven-times-faster-than-ever-before/>.