



The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search

<http://ageconsearch.umn.edu>

aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.

NCEE Working Paper

Data Center Energy Efficiency Investments: Qualitative Evidence from Focus Groups and Interviews

**Heather Klemick, Elizabeth Kopits, and Ann
Wolverton**

**Working Paper 17-06
November, 2017**

Data Center Energy Efficiency Investments: Qualitative Evidence from Focus Groups and Interviews

Heather Klemick, Elizabeth Kopits, and Ann Wolverton

Abstract:

The data center industry is one of the fastest growing energy users in the US. While the industry has improved its energy efficiency over the past decade, engineering analyses suggest that ample opportunities remain to reduce energy use that would save firms money. This study explores potential barriers to energy-efficiency investments in data centers. Given the scarcity of empirical data in this context, we conducted focus groups and interviews with data center managers to elicit information about potential barriers to investment and used content analysis to qualitatively evaluate the results. Split incentives between departments within companies and between colocation data centers and their tenants, uncertainty and imperfect information about the performance of new technologies, and tradeoffs with data center uptime were the most pervasive potential barriers discussed by participants. While these factors have moderately slowed investments in energy-saving technologies for many firms, only in the cases of uncertainty/imperfect information and split incentives are these barriers potentially indicative of market failures.

Keywords: energy efficiency paradox; market failures; data centers; technology investment barriers.

JEL classifications: Q52; Q48; Q58

DISCLAIMER

The views expressed in this paper are those of the author(s) and do not necessarily represent those of the U.S. Environmental Protection Agency. In addition, although the research described in this paper may have been funded entirely or in part by the U.S. Environmental Protection Agency, it has not been subjected to the Agency's required peer and policy review. No official Agency endorsement should be inferred.

1. Introduction¹

As the world has become increasingly reliant on electronic forms of communication and services—for instance, online banking and shopping, satellite-based navigation systems, smart phones, streaming of movies and television shows—demand for storing, managing, and distributing large amounts of data and information has grown. Data centers are the centralized repositories of computer servers that provide these services. US demand for these services is expected to continue to grow rapidly. A 2016 report projected that data center workload will more than double between 2015 and 2020, while the number of servers installed is expected to grow by 25 percent (Cisco 2016; Shehabi et al. 2016). This rapid growth has been accompanied by increasing energy consumption. In 2014, U.S. data centers consumed approximately 70 billion kilowatt hours, totaling about 1.8 percent of domestic electricity consumption, compared to roughly 30 billion kilowatt hours in 2000 (NRDC 2014; Shehabi et al. 2016). This electricity usage comes at a cost to the industry of about \$7 billion a year.

This study examines data center managers' investment decisions in energy-saving technologies and considers whether they invest less than what economic theory would predict due to a potential "energy efficiency paradox." The energy efficiency paradox is a term used to describe situations in which consumers forgo investments in energy efficiency that are costlier upfront but save money in the long term (Jaffe and Stavins 1994). The economics literature has devoted more attention to investigating the energy efficiency paradox in households than businesses. This study contributes to the emerging literature on firms' energy investment decisions by gathering qualitative data from a series of focus groups and interviews with data center managers to shed light on potential barriers to energy efficiency enhancing investments.

Data centers are comprised of information technology (IT) equipment—including servers, data storage, and networking devices—as well as the facility and infrastructure needed to house and maintain them, such as cooling and lighting. Adoption of energy-saving technologies and practices by data centers can

¹Focus groups and interviews were conducted with contractor support funded by the U.S. Environmental Agency. The authors thank Barbara Bauer, and David Cooley (Abt Associates); Linda Dethman and Jane Peters (Research Into Action); Beth Binns; Datacenter Dynamics; AFCOM; and Keith Sargent (EPA) for help coordinating and conducting focus groups and interviews. The authors also thank Cynthia Morgan (EPA) for her helpful input. All opinions expressed in the paper are those of the authors and do not necessarily reflect the views or policies of the U.S. Environmental Protection Agency.

weaken the link between the volume of computing services and the energy needed to power these services. The data center industry has already started to realize this potential. Electricity consumption has substantially slowed in the past decade even as demand for computing services has continued unabated. The 2008 recession temporarily dampened growth of the sector, but since then increases in electricity consumption have been modest despite a rebound in the demand for computing services (Shehabi et al. 2016). A 2007 EPA report on data center energy usage projected that electricity consumption would have topped 100 billion kilowatt hours by 2011 if trends in the early 2000s had continued (Brown, et al. 2007). By contrast, a 2016 update to the report estimated that actual electricity use in the sector has not approached previous projections and will likely remain under 75 billion kilowatt hours by 2020 (Shehabi et al. 2016). Data centers have realized energy savings through both IT and infrastructure related measures, such as improved air flow and temperature management and server virtualization, which allows a single physical server to run multiple operating systems simultaneously.

However, improvements in data center energy efficiency have not occurred evenly across the industry. Efficiency gains have been concentrated in hyperscale data centers, which are warehouse-sized facilities that provide computing services as their core business, including cloud computing and colocation services that provide space for other firms' servers.² While the hyperscale segment of the industry is growing rapidly, as of 2014 it only represented about 20 percent of installed servers (Shehabi et al. 2016). Enterprise data centers, which are medium- and large-scale facilities operated by firms or organizations outside of the IT industry for their own internal use (e.g., banking, health care), have also made efficiency improvements (Shehabi et al. 2016). A 2014 survey of data center managers in IT, finance, and other mid-level firms in North America, the majority of which owned their data centers, found that 76 percent ranked energy efficiency as important or very important when investing in new data center facilities (Forrester Research 2014). Smaller server closets and localized data centers that are typically located in commercial buildings remain the largest segment of the industry with more than 40 percent of installed servers, but they lag behind in terms of efficiency. Consolidation of computing services from smaller operations to hyperscale cloud service providers is expected to continue, which could lead to energy efficiency improvements in the industry as a whole by 2020 (Shehabi et al. 2016).

² Cloud computing allows companies to outsource computing infrastructure and applications to servers remotely through the internet. Colocation or multi-tenant data centers are shared spaces "where customers lease space and power to run their computing equipment rather than managing their own data center" (NRDC 2014).

Despite recent improvements, studies continue to point to potential opportunities for net cost savings from additional energy efficiency investments. The 2007 Report to Congress and more recent studies have projected that adoption of best practices—such as increased server utilization, realigning incentives between data centers and their clients, and decommissioning idle “zombie” servers that continue to draw power—could reduce electricity consumption below current levels by 25 to 45 percent by 2020, yielding billions of dollars a year in cost savings (Shehabi et al. 2016, NRDC 2014, Brown et al. 2007). Masanet et al. (2013) calculated a technical potential to reduce data center electricity demand of more than 80 percent by shifting all US business software to cloud computing.

A 2013 California regulation established the only requirements for data center efficiency in the United States, mandating the use of specific cooling, airflow, and fan technologies (California Building Standards Commission 2013). Energy consumption across the remainder of the industry remains unregulated, despite a few federal initiatives to promote best practices, such as a 2015 Executive Order setting targets for federal data centers and a variety of tools and resources developed by the Energy Star voluntary program.³ Analysts estimated average energy savings per data center of \$10.5 million annually as a result of energy efficiency improvements made in response to the California regulations, which would far outstrip the cost of adoption (California Statewide Utility Codes and Standards Team 2011).

The apparent large cost-saving opportunities in the data center sector could indicate an energy efficiency paradox. The literature on commercial buildings has investigated market failures such as imperfect and asymmetric information, split incentives (principal-agent problems), and credit constraints, as well as behavioral anomalies, which occur when investment decisions are inconsistent with profit maximization (e.g., inattention and loss aversion), as possible contributors to an energy efficiency paradox (Andrews and Krogmann 2009; Sallee 2014; Schleich and Gruber 2008).⁴ However, studies have also pointed to tradeoffs between energy efficiency improvements and other valued product attributes that, while not reflected in engineering analyses, make investment in energy-saving technologies genuinely less attractive (Gillingham et al. 2009, Klemick, et al. 2015; Klemick, et al 2017).

³ For example, see [http://www.energystar.gov/buildings/tools-and-resources?f\[0\]=field_select_sector%3A664](http://www.energystar.gov/buildings/tools-and-resources?f[0]=field_select_sector%3A664).

⁴ Unpriced externalities from GHGs and other air pollutants also likely contribute to suboptimal technology adoption. Because they do not enter into firms’ private profit-maximizing decisions, we do not consider them here.

For example, data center managers across all segments of the industry are reticent to adopt any new technologies that could compromise reliability or uptime, the amount of time the data center is operational (NRDC 2014).

In this paper, we analyze qualitative evidence from focus groups and interviews with data center managers to examine whether such barriers limit or slow adoption of energy-saving technologies and practices. We also consider whether these barriers represent market failures that might be mitigated through targeted policies versus real economic costs that could be incorporated into payback analyses to provide a more complete accounting of the net benefits of energy-saving strategies.

The paper is organized as follows. Section 2 describes the methodology and summarizes characteristics of participating data centers. Section 3 discusses adoption of specific energy-saving technologies among our sample and whether these adoption patterns show evidence of a “payback gap,” meaning that firms only adopt energy-saving technologies when the costs savings fully offset the higher upfront costs within the first few years of operation, neglecting those technologies that yield positive returns over a longer time horizon. Section 4 discusses the role of potential barriers limiting adoption of energy-saving practices discussed by participants, focusing on imperfect information and uncertainty, split incentives, financial barriers, and tradeoffs with valued attributes. The final section concludes.

2. Methodology

Researchers have used qualitative methods to collect information on building energy management and design (e.g., Gul and Menzies 2012, Pellegrini-Masini and Leishman 2011), energy efficient supermarket refrigeration investments (Klemick et al. 2017, Ochieng et al. 2014, and Sullivan and Gouldson 2013), and heavy-duty trucking fuel economy investments (Klemick et al. 2015). NRDC (2014) used interviews with data center stakeholders to inform a report on energy efficiency trends and barriers in the industry. Given limited data on investments to improve energy efficiency in data centers, we also used this approach, conducting focus groups and interviews to examine potential investment barriers.

Focus groups are useful when the interaction of participants can result in a more nuanced, richer conversation than speaking to each participant individually, highlighting possible areas of agreement and disagreement. In-depth interviews are viewed as useful for collecting detailed information on a complex or sensitive subject from busy, high-status participants. Frechtling and Sharp (1997) identify

several advantages of the focus group and interview approach compared to other data collection approaches: They directly capture participants' perspectives, permit in-depth exploration to gain better understanding of particular perspectives and experiences, and allow flexibility in how questions are asked to reflect the unique circumstances and background of participants. However, they also highlight several disadvantages (Frechtling and Sharp 1997). The quality of the information gathered depends on the facilitator's ability to moderate the discussion and is subject to biases common to research approaches that rely on statements from participants instead of directly observing behavior. Flexibility in how questions are asked also can result in inconsistencies across focus groups and interviews in what data are collected. Finally, recruitment of participants can be particularly challenging and is unlikely to result in a representative sample. We discuss this last issue in greater detail below.

2.1 Focus group and interview protocol

The moderator's guide followed a semi-structured funnel design to facilitate focus group and interview discussion (i.e., starting with a broad discussion about how companies make investment decisions within a data center and the main sources of information used to make those decisions before asking about the specific role energy efficiency plays in equipment and server upgrade decisions). Questions were open-ended so that participants could identify factors we may not have anticipated ahead of time (Table 1). The facilitator relied on additional prompts to elicit information related to specific hypotheses from the literature on an as-needed basis. Information collected via questionnaire on 22 specific facility or server technologies, as well as facility management or pricing practices that are often promoted as saving energy in data centers, was used to facilitate discussion of particularly appealing or unappealing features that influence adoption.⁵ The focus groups and interviews ended with a discussion of the extent to which these considerations are reflected in payback calculations, and reasons why firms may require that energy efficiency investments pay back over a relatively short time period.

We then performed content analysis to evaluate focus group and interview transcripts (Radcliff and Best 2005), relying on a directed approach to develop codes related to specific barriers highlighted in the literature (Hsieh and Shannon 2005). We tested the coding algorithm by having more than one person independently code the same transcript, identified any differences in coding across coders, and revised the coding algorithm or the way it was applied to achieve consistency across the coding of the remaining

⁵ Appendix A describes many of the facility management and server technologies included on the questionnaire. Appendices B and C contain the questionnaire and moderator's guide.

transcripts. While we offer descriptive evidence in support of or against a particular hypothesis, we do not conduct summative content analysis;⁶ because we use an open-ended design, basic statistical tests are not advisable.

⁶ Summative content analysis is defined as “identifying and quantifying certain words or content in text with the purpose of understanding the contextual use of the words or content” (Hsieh and Shannon 2005).

Table 1. Sample questions from the moderator’s guide

Where do you gather information on what new equipment, technologies and facility management strategies are available? Do these sources provide reliable estimates of likely energy use? Are there estimates out there that you don’t believe? Why?

Can you give me an example of a major investment that improved facility energy efficiency? Why these technologies or approaches? What attributes make them appealing? What factors did you consider when deciding whether to make this investment?

I have a list of factors here that you might consider when choosing technologies or strategies to reduce energy use [e.g., uptime/reliability, maintenance and repair, location-specific factors, financing]. I’m interested in whether there are tradeoffs or synergies between these factors and energy use; are these factors barriers or motivators for improved energy efficiency?

Thinking about all the factors you have mentioned so far that affect your investment decision, how do you weigh them against upfront cost? Do you incorporate any of them into an ROI or payback calculation?

2.2 Participant Summary

We conducted six professionally facilitated focus groups with a total of 27 managers involved in the purchase, servicing, and/or maintenance decisions for equipment in one or more U.S. data centers.⁷ Each focus group was approximately two hours in length. The smallest number of participants in any single focus group was three; the largest was six. In addition, we conducted seven phone interviews, which were mainly with managers of very large data centers in case competitiveness issues may have inhibited candidness of such large companies in a group setting. While our original intent was to host discussions with key segments of the industry separately (e.g. by sector or data center type) this proved infeasible. Recruiting data center managers to participate was unexpectedly difficult due, in part, to privacy concerns. Instead, focus groups occurred in three geographic areas with somewhat different sectoral emphases.⁸ In addition, New York City/Northern New Jersey and Dallas/Fort Worth are in the

⁷ Focus groups (FG) and interviews (INT) are labeled numerically throughout the document.

⁸ According to JLL’s 2016 Data Center Outlook, the largest sources of demand for data centers in Boston are the technology sector (35%), followed by life sciences (25%), financial services (25%) and institutions such as universities (15%). In Dallas/Fort Worth the leading industries with regard to data center demand are insurance (30%), technology (30%), and banking and financial services (25%). In New York City, 55% of data center demand derives from cloud computing and telecommunications, following by banking and financial services (30%) and manufacturing (15%). Finally, in northern New Jersey, the sectors that dominate data center demand are banking and financial services (45%), healthcare (20%), technology (15%), and retail/e-commerce (15%).

top ten metropolitan areas in terms of concentration of data centers (JLL 2016). Boston represents a smaller but rapidly growing data center market. Focus group and interviews occurred between October 2014 and June 2015. We excluded data centers providing services to state or federal government. We anticipated that these data centers may have different incentives regarding adoption of energy efficient equipment and practices compared to the private sector.

Participants were recruited through collaboration with national associations and industry experts.⁹ Company affiliation was used to avoid duplication when recruiting but was otherwise masked to protect participant confidentiality. To ensure we did not only include industry leaders in reducing energy use, we actively recruited firms that do not participate in the Energy STAR buildings program or belong to an organization that provides information to industry with the aim of increasing data center efficiency (e.g., Green Grid).¹⁰ In spite of these precautions, the vast majority of our sample participated in programs or belonged to organizations promoting energy efficiency. In addition, it was challenging to reach companies that operate small data centers. This may be because managers of large data centers are more likely to travel to association conferences, the main method that we used to recruit focus group and interview participants. As such, participants may be more interested in energy and technology issues and may be working for larger companies, on average.

In our sample, 41 percent of participants were employed by companies with more than 50,000 servers spread across their data centers, while about a quarter of participants worked for companies with less than 2,000 servers (Table 2). These servers are spread across companies' data centers in various ways. About 37 percent of the companies represented in our sample typically use 200 or fewer racks to house servers in their data centers, while 26 percent of the companies have medium-sized data centers (using 201 – 1,000 racks) and 41 percent have large data centers (using more than 1,000 racks), on average.^{11,12}

⁹ Four focus groups were organized in conjunction with Data Center Dynamics conferences in Dallas and New York, and two focus groups were organized in conjunction with an AFCOM conference in Boston.

¹⁰ EPA and DOE's Energy STAR buildings program recognizes top performers in energy efficiency through certification of individual facilities as well as portfolios of buildings or plants.

¹¹ A rack is a set of vertical mounting rails and a supporting metal framework in which servers, network switches, cables, and other computer hardware are housed. The number of servers in a single rack varies with rack height and depth, how much weight it can support, and the way in which it is configured. See <https://www.tripplite.com/shared/literature/White-Paper/Rack-Basics-White-Paper-EN.pdf>.

¹² Our rack categorization is informed by the way AFCOM defines data center size. See <http://www.datacenterknowledge.com/archives/2014/10/15/how-is-a-mega-data-center-different-from-a-massive-one/>

Table 2: Data Center Participant Summary

		Number of participants	% ^A
Total number of participants (focus groups plus interviews):		34	
Total number of servers in company:	< 2,000	9	26%
	2,000-50,000	6	18%
	50,000 +	14	41%
Average number of racks in a data center:	Small (<201)	10	37%
	Medium (201 -1,000)	7	26%
	Large (>1,000)	11	41%
Data center types (not mutually exclusive):	Cloud/hyperscale	12	35%
	High-performance/scientific	6	18%
	Colocation/multi-tenant	11	32%
Sectors supported (not mutually exclusive):	Information technology	15	44%
	Banking/financial services	14	41%
	Media/telecommunications	6	18%
	Health care	9	26%
	Education	6	8%
	Other/not specified	9	26%
Location (not mutually exclusive):	One state	7	21%
	Multiple states	23	68%
	Also international	5	15%
Number of data centers participant manages:	5 or less	16	50%
	6 - 20	13	41%
	21 - 80	3	9%
Average size of data center participant manages:	Server room/localized	7	21%
	Enterprise/corporate	8	24%
	Mega/utility-scale	19	56%
Server replacement frequency:	< 3 years	3	10%
	3-5 years	19	66%
	Varies	4	14%
Major infrastructure upgrade frequency:	< 2 years	10	33%
	5-10 years	6	20%
	10 or more years	8	27%
Third-party audit in last 3 years?	Yes	13	38%
E.E. program participant (e.g. Energy STAR, Green Grid)?	Yes	24	71%

^A Percentage based on total number of participants.

There is a fair amount of diversity in our sample with regard to the types of data centers operated by a given company. About one-third of participants hosted cloud or hyperscale computing; 18 percent hosted high-performance or scientific computing used by research facilities for data-intensive processes; and about one-third hosted colocation or multi-tenant data centers. The sectors supported by the data center companies that participated in our study also varied. Data centers whose main line of business is providing IT services were most common (44%). Participants also reported supporting a number of other industries, including banking or financial services (38%), health care (26%), media and communications (18%), and other industries ranging from energy to state and federal government to manufacturing (26%). The vast majority (about 70%) had data centers in more than one state; a few also had an international presence (15%).

We also gathered information about data centers directly managed by participants. All participants worked at companies with more than one data center (ranging from two to thousands, though the majority owned 10 or fewer data centers). Half of the participants had direct oversight over five or fewer data centers, while another 41 percent were involved in decision-making for six to 20 data centers. More than half of the participants also managed at least one mega- or utility-scale data center (larger than 10,000 square feet). This category includes most retail and wholesale colocation facilities. About one quarter were directly involved in the decision-making of an enterprise/corporate data center (greater than 5,000 square feet), which fills a support function rather than being the company's main line of business. About 20 percent of participants managed either localized data centers (less than 1,000 square feet) or server rooms or closets (less than or equal to 500 square feet) that typically serve only the needs of an individual office or call center.

Major server and equipment upgrades occurred frequently. One third of participants reported that they made some type of major infrastructure upgrade once every two years. More than half of participants also reported replacement of servers on a three- to five-year cycle, and less than ten percent reported doing so even more frequently. In addition, over 70 percent of participants took part in either a government or third-party energy efficiency voluntary program, while almost 40 percent had a third-party energy audit in the last three years.

3. Results: Evidence of a “Payback Gap”

We begin by discussing whether there is a “payback gap” in data center investment decisions – i.e., if firms require energy savings to offset adoption costs in substantially less time than the amount of time that the technology is expected to be in use. Statements made by all interviewees and by participants in every focus group indicate that their firms did indeed consider some type of calculation of how long it would take for an energy efficiency related investment to pay for itself when making purchase decisions, although the level of detail and sophistication of these calculations varied across firms and/or type of investment. As discussed below, in most cases the expected payback period for investments was likely shorter than equipment lifetime, suggestive of a “payback gap” in energy efficiency enhancing investment decisions in the data center industry. In addition, participants’ responses about their adoption of specific energy efficiency enhancing technologies and strategies were consistent with their general statements related to payback calculations. Adoption rates for the technologies expected to have relatively short payback periods (e.g., less than 2 years) were quite high, while adoption rates for other technologies were more mixed, especially among participants managing localized data centers.

3.1 Payback Calculations

Participants in all the focus groups and interviews stated that their companies calculate payback, total cost of ownership, or return on investment (ROI).¹³ Some firms emphasized the need to perform a more sophisticated analysis than a simple payback calculation in order to gain support within the company for a project, while three firms noted that they do not do a complete ROI calculation (FG#6). Several participants stated that the level of detail in their analyses varied by type of investment. For example, one interviewee mentioned that new payback analyses are not needed for well proven “no-brainer” technologies (INT#2) or for IT investments given the rapid replacement cycle of the equipment (INT#1).

At least four interviewees indicated they have fairly comprehensive total cost of ownership models, incorporating factors such as depreciation of the asset, water usage, labor costs, maintenance and repair costs, age of the facility, the time it takes to install the technology, discount rate, uptime impacts, and the opportunity cost of larger equipment in terms of lost space (INT#3, INT#4, INT#6, INT#7). Other

¹³ Total cost of ownership is a term used in financial analysis intended to capture both the direct and indirect costs of an investment, from the upfront acquisition cost and operating costs through the replacement or upgrade at the end of its life cycle. It provides a cost basis for determining the total economic value, or return on the investment (ROI). A high ROI indicates the investment’s gains compare favorably to its cost.

participants also mentioned accounting for some of these or other factors in their quantitative payback calculations, such as maintenance costs (FG#1, FG#2, FG#6), service contract expenses (FG#1), government or utility incentives (FG#3, FG#4), and installation costs (FG#2, FG#3), although not all factors could be included in a quantitative way (INT#1, INT#2). As one interviewee noted, *“An investment decision needs to make good sense...I purposely did not say economic sense because sometimes you make good sense for benefits to brand or benefits to acquiring a new customer or some other more external benefits...”* which would be complex to incorporate into a cost model (INT#2). Two interviewees highlighted the importance of discounting cash flows, because otherwise the payback calculations do not reflect the actual ROI (INT#2, INT#4). A couple of participants noted that staff retraining costs were considered if expected to be significant (FG#6).

Electricity usage and rates featured prominently in these calculations for most companies. One focus group participant noted that the data center is the largest energy user in the company; while energy use was considered an *“afterthought”* in the past, it is now considered in technology and siting decisions (FG#4). Participants in another focus group echoed this view, stating, *“ten to fifteen years ago [energy consumption] flew way under the radar, and it just didn’t cost that much to run and build a data center, and every year... our utility rates are going up. It’s become a huge expense, and it’s got the attention of all of our businesses”* (FG#1). These statements suggest that lack of attention to energy consumption was not a major factor driving new technology investment decisions. However, a participant in a different focus group suggested that energy consumption plays a less important role in the decision-making of data centers for companies whose main line of business is something other than IT (FG#3).

3.2 Required Payback

A few participants in every focus group and all but three interviewees required paybacks between one and five years. However, many participants’ explanatory statements suggested that these payback guidelines are not hard-and-fast rules. Some noted that there is less emphasis on achieving a short payback for equipment replaced on a regular cycle than for upgrades to existing facilities (FG#1, FG#6). For example, one participant stated, *“as a rough rule of thumb, there needs to be a really good reason to consider a retrofit with a really long payback”* (INT#6); another explained, *“There might be times when we could let it stretch out [the payback] longer than [two to three years] but if it is an upgrade of an existing facility, [the payback] would have to be fairly short to compensate for the risk of making changes”* (INT#3). A couple of participants said that regardless of the payback they may still look at an investment if there was a good marketing reason, such as to make the company look *“green”* (FG#4). As

one participant explained, *“If I am doing an efficiency improvement and I can’t make it pay for itself in three years, there is virtually no chance I am going to get the money. I will give you an exception to that. If there is what I will call a marketing reason...to help market the company as being a good company, then I have got a chance...”* (FG#2). See Section 4.4 for further discussion of how tradeoffs and synergies between customer concerns and energy efficiency factor into investment decisions.

Participants’ statements largely indicated that the required payback periods were shorter than the amount of time that they would hold on to equipment, which was often noted to be 10 or more years for many facility-related investments. For example, one participant calculated returns over a 15-year lifecycle of the equipment but needed a three to five-year payback to justify the investment to other departments (INT#7). One participant stated that his company does not have a required payback period and will generally consider anything that has a payback period less than the lifetime of the equipment. In the case of IT equipment or software, this means more stringent guidelines may be relevant, *“...But for the data center itself, because again, the lifetime is 20 to 30 years....you are really looking for any payback period. If it was 30 years you probably wouldn’t do it, but typically no payback period is really longer than five or six”* (INT#5). However, even this participant went on to amend his statement to say that he still needs a *“reasonable payback period,”* citing onsite electricity generation as an example of technology with a payback period of over 10 years that his company declined to adopt.

3.3 Adoption of specific energy efficiency enhancing technologies

To help facilitate the focus group and interview discussions of potential payback gaps and their causes, we asked each participant to complete a questionnaire about their use of 22 specific energy efficiency enhancing technologies and strategies. We asked about management strategies related to the facility infrastructure and opportunities related to the IT equipment, hardware and software that can be employed across a variety of facility types. In addition, participants managing multi-tenant or colocation data centers were asked about the use of pricing strategies to incentivize energy efficiency. Figure 1 summarizes the questionnaire responses by data center size.

Facility Management Strategies. There are many ways to reduce data center energy use. Strategies that focus on management of air flow within the building are often cited as very cost effective investments. For example, investments in barriers that prevent the mixing of incoming cold air and hot exhaust air

(cold- or hot-aisle containment), tiles and panels that guide cool air directly to servers (e.g., vented or perforated tiles in the floors, blanking panels to cover unused rack spaces), devices on existing air conditioning units that allow variation in air flow as cooling demand fluctuates (variable fan or speed drives), and use of higher temperature set points are frequently described as paying back in less than two years.¹⁴ Investments in the method of cooling utilized often have significantly higher upfront costs but can yield large energy savings in specific circumstances. For example, in some locations it is possible to use outside air or water that is already sufficiently cool instead of running equipment to reduce the temperature of hot air or water for recirculation or use, technologies called airside or waterside economizers. Companies' stated experience and research on airside and waterside economizers suggests multi-year payback periods for these technologies, especially for retrofits.¹⁵


As Figure 1 indicates, in several cases adoption patterns appear to correlate to some extent with expected payback estimates, with air flow management strategies generally being used at higher rates than some HVAC system adjustments. Consistent with recent industry trends showing higher energy efficiency among larger data centers (Shehabi et al. 2016), adoption of many technologies and practices was noticeably higher among utility-scale data centers, followed by the enterprise category, compared to localized data centers. Air flow management strategies – hot/cold aisle configurations, blanking panels, variable speed fans, and energy efficient air handling units – as well as energy efficient lighting have been used by all but one (95%) of the utility-scale data center companies. Managers of enterprise data centers reported slightly lower use of variable speed fan drives, energy-efficient air handling units, and efficient lighting (75%). Adoption by localized data centers ranged from 57-71 percent across all of these technologies. Use of strip curtains (i.e., a curtain of overlapping strips of plastic that hang in a doorway or partition a space) followed this same pattern, though adoption rates are lower than for many other technologies across all three data center sizes. In fact, three enterprise data center participants stated they would never consider using them (see Section 4.4 for more discussion).

¹⁴ See Appendix A for a more detailed description of these and other technologies. The Green Grid (2011) and Energy Star provide an overview of common ways save energy in data centers:

https://www.energystar.gov/products/low_carbon_it_campaign/12_ways_save_energy_data_center/ .

¹⁵ See Wathaifi (2009), Pacific West Air Conditioning (2014), Van Greet (2017), and Energy Star's website at: https://www.energystar.gov/products/low_carbon_it_campaign/12_ways_save_energy_data_center/ .

Figure 1. Participant Use of Energy Efficiency Enhancing Technologies, by Data Center Size



Adoptions of HVAC system adjustments followed the same pattern but were a bit lower across the board. Approximately 80 to 90 percent of participating utility-scale data centers and over half of enterprise data centers had used air- and/or water-side economizers, while only one localized company had adopted these technologies. Nearly 80 percent of all respondents (25 participants) reported raising temperature set points as a way to reduce energy use. Of these, all but three managed enterprise or mega/utility-scale data centers. The reported typical temperature across these facilities was higher than those who have not yet tried this strategy (averaging 76 and 71 degrees, respectively).¹⁶ The remaining facility management strategies on the questionnaire – infrastructure management software, modular facility design, and predictive modeling of future IT deployment - were also widely adopted by companies managing utility-scale data centers but less so by enterprise or localized data centers.

IT Management Strategies. Participants also reported widespread use of several IT opportunities with relatively short expected payback periods. Over 80 percent of respondents¹⁷ have adopted energy-efficient server hardware, virtual servers, and decommissioned unused servers, and 74 percent reported consolidating lightly used servers. These are mostly measures that reduce energy consumption by reducing the number of servers needed to manage a data load at any given time. For example, instead of using a separate physical server for each application, which results in only a small fraction of computing resources being actively utilized, the creation of virtual servers allows multiple applications to run on a single physical server simultaneously. Interestingly, even though past surveys have found virtualization to be less common in smaller data centers (NRDC 2012), all of the localized data centers in our sample reported using server virtualization, which may suggest that this technology is becoming a more universal practice throughout the industry.

Over half of all respondents had adopted measures to improve management of data storage and/or IT power management software, although adoption rates were somewhat lower among localized data centers. The practice of under- or over-clocking—dynamically adjusting the clock speed of a computer’s processor or memory to match IT workloads to save power or reduce heat—was less widely adopted. Overall, 26 percent of participants reported they had tried this strategy as a way to reduce energy use,

¹⁶ Most of the respondents who employed hot/cold aisle containment strategy did not specify whether they were reporting the typical temperature in the cold or hot areas of their data centers.

¹⁷ The response rate on this section of the questionnaire was less than 100 percent. Three mega/utility scale companies did not respond, and another utility-scale company reported “n/a” for all of these IT opportunities.

and a greater share (38%) reported not knowing whether their company had tried it; the remainder said they would consider trying it in the future (15%), would never consider it (9%), or did not respond (12%).

Pricing Structures. The last section of the questionnaire asked about the use of pricing structures to incentivize energy efficiency. Participants were asked whether pricing was based on the amount of space rented in a data center and/or the amount of power consumed (either within set block amounts or according to actual electricity uses). Since these questions were aimed at multi-tenant data centers, response rates are lower for these questions. Among the respondents who manage mega/utility-scale facilities, pricing models that charge for both space and actual power consumed (37%) were somewhat more popular than pricing based on space alone (32%), and 16% said that they would not use space-only pricing again or consider it in the future.

4. Results: Barriers to Technology Adoption

In this section, we discuss the evidence from the focus groups and interviews about potential barriers to the adoption of energy-saving technologies and practices in data centers. We also discuss the extent to which any identified barriers may result from market failures or instead be characterized as costs that make technologies genuinely less attractive to firms but that are not accounted for in typical engineering-based NPV analyses. When possible, we identify evidence consistent with behavioral anomalies, though it is often not possible to separate out this effect from other explanations.

In particular, we organize the discussion around four general categories of potential barriers to technology adoption: uncertainty and imperfect information, split incentives, financial constraints, and tradeoffs with other valued attributes. For each of the four categories, Table 3 summarizes our results based on how often a particular barrier was discussed as a factor of potential concern, the degree to which this barrier appears to be limiting adoption of energy efficiency related investments, evidence that the barrier may result from a market failure, and potential improvements that could be made to engineering-based NPV calculations to address analytic gaps.

While these measures are inherently subjective—they are based on our interpretation of participant statements—in general, we find that some aspect of each potential barrier was a considerable concern among participants. As discussed in detail below, there is also evidence that many of these factors moderately limit data center investment in energy efficiency. However, only split incentives and imperfect information/uncertainty appear to be barriers that may result from a market failure. In those

cases, we characterize the evidence as mixed because participants often developed strategies that at least partly overcome the primary investment barrier. Finally, in several areas, we suggest potential improvements to engineering-based NPV calculations that would better align these calculations with the way data center managers actually make investment decisions.

Table 3. Evidence of barriers to investment and potential market failures

<u>Potential barrier to investment</u>	<u>Discussed as Factor of Concern</u>		<u>Effect on limiting adoption</u>	<u>Evidence of market failure</u>	<u>Potential improvement to NPV calculations</u>
	Number of Focus groups	Number of Interviews			
Uncertainty and imperfect information	6/6	4/7	Moderate	Mixed	Include info. on real-world performance and cost to gather it
Split incentives					
<i>Within company</i>	6/6	7/7	Moderate	Mixed	Better account for costs of adoption (e.g., coordination, training)
<i>Colocation facilities</i>	6/6	5/7	Moderate	Mixed	
<i>Firm-contractor</i>	3/6	0/7	Weak	Weak	
Financial constraints					
<i>Limited access to outside financing</i>	0/6	0/7	Weak	Weak	Use private instead of social discount rate
<i>Opportunity cost of capital</i>	4/6	4/7	Moderate	No*	
Tradeoffs with other attributes	6/6	7/7	Strong	No *	Develop methods to explicitly value these tradeoffs

* No is used when a barrier is not classified as a market failure, while weak is used to signal that while it is possible a barrier could be due to a market failure we find little evidence that this is the case.

4.1 Uncertainty and Imperfect Information

This section investigates whether imperfect information and uncertainty pose a barrier to adoption of new energy-saving technologies in data centers. If firms that adopt new technologies generate information about their performance that is valuable to other data centers (sometimes termed “learning-by-doing” or “learning-by-using” externalities), and if this type of externality causes firms to wait for others to test out new technologies before making an investment, then the pace of technology adoption will occur slower than is socially optimal.

However, if at least some firms in the industry can capture enough of the returns to investing in new technologies themselves, then they may still proceed with adoption. In this situation, firms can still benefit from information spillovers from a small group of firms—often the largest in the industry—that undertake most of the innovation (Olson 1965). A study of supermarket investment decisions found preliminary qualitative evidence suggesting that the largest supermarket chains often serve as early adopters of new energy-efficient technology, providing beneficial information spillovers to their competitors and smaller chains (Klemick et al. 2017). However, the supermarket industry is highly concentrated, with the top five chains comprising 80 percent of the market, whereas small and medium sized companies make up 40 percent of the data center industry (NRDC 2014). Thus, there may be less capacity among the largest data center firms to serve as “guinea pigs” for the entire industry.

Participants in all six focus groups, as well as four out of the seven interviewees, discussed imperfect information and uncertainty as factors that could potentially limit or slow their adoption of energy-saving technologies (Table 3). Their comments suggest that imperfect information and uncertainty have a moderate impact on adoption of energy-saving technologies in data centers. While many participants avoid being on the “*leading edge*” (INT#1) and prefer to wait for others to test out new technologies, most report making investments to improve the efficiency of data center facility infrastructure and have confidence that these investments will yield a positive return on investment due to information available from their own testing or from other firms. Information most valuable to firms is typically company- or facility-specific, and it is also costly to gather, suggesting that information barriers represent a real transaction cost to firms rather than a market failure in most cases. However, there are some information spillovers among firms, so we cannot completely rule out a small role for information market failures in the context of data center investment decisions.

Information about Current Energy Use

Data centers typically track and communicate information about their energy efficiency in terms of power usage effectiveness (PUE), a metric that is defined as the amount of power used by the data center as a whole divided by the amount of power used by the IT equipment (Belady et al. 2008). A PUE of one indicates that the data center uses no additional power for lighting, cooling, power distribution or facility operation beyond what is drawn by the IT equipment, whereas a PUE of two indicates that for every unit of power consumed by the IT equipment, another unit is used for facility operation. Therefore, PUE measures infrastructure efficiency rather than data center efficiency as a whole. The industry has not yet coalesced around a more comprehensive energy efficiency metric due to the difficulty in quantifying server utilization and computing services output (Horner and Azevedo 2016).

Use of the PUE metric was widespread among companies participating in our study but varied by size of the data center. More than 80 percent of mega-scale and enterprise facilities tracked this information, but only 29 percent of localized data centers and server closets measured PUE. PUE among firms that tracked and reported this information also decreased with facility size (indicating an improvement in efficiency), with the median dropping from 1.7 for localized facilities to 1.55 for enterprise to 1.4 for mega/utility-scale data centers. These figures compare favorably to industry-wide trends: a 2014 survey of more than 1,000 data centers managers found that 72 percent of respondents' data centers measured PUE and reported an average PUE of 1.7 (Stansberry 2014). Shehabi et al. (2016) assumed a PUE of 2 for server closets and localized facilities, which may suggest that participants in our study are somewhat more interested in energy efficiency than the industry at large. However, less than half of participants in our study track the server utilization rate at their facilities.

Several participants noted that real-time monitoring of PUE and energy consumption at the level of the individual facility—or even the individual rack or server—is an important prerequisite to improve efficiency (FG#1, FG#4, FG#6, INT#4). One interviewee explained, *"We have a real time PUE monitor for each of our data center sites that everyone can go to at any moment and find out what the PUE... is. The site techs are constantly looking at that to try and figure out how they can be more efficient, and they're very competitive"* (INT#4). The same interviewee called monitoring and control systems the *"best tool you can put in the data center"* to motivate energy efficiency investments (INT#4). Another interviewee

noted that, *“The level of measurement and the specificity have gotten better over the past. I’d say it’s been... a three-year journey to make that better and better over time”* (INT#2).

However, not all participants have advanced monitoring capabilities. One participant said his data center does not meter energy use at all (FG#5), and another cited a need for more research to understand the relationship between new technologies and overall energy use (FG#6). Granular energy monitoring systems are particularly useful in colocation data centers, making it possible to give tenants real-time feedback on PUE and charge them for actual electricity consumption, an approach that can strengthen tenants’ incentive to conserve energy (INT#2, FG#4). The next section discusses split incentives between colocation providers and clients and between departments within the same company.

PUE is the most widespread energy efficiency metric in the industry, but participants in four focus groups and one interview noted it is an incomplete representation of a new technology’s energy efficiency. Because PUE does not include a measure of the effectiveness of IT equipment—or more generally, the amount of computing services provided per unity of energy—it is not a holistic measure of energy efficiency (FG#1). A focus group participant discussed this dilemma: *“There’s an innate problem out of the gate in terms of even engineers telling you what the energy consumption of a data center is going to be.... It’s a performance payoff, and there’s no method for calculating performance”* (FG#5). Furthermore, some participants complained that PUE is not always comparable across data centers because of a lack of standardization in the way it is measured and calculated (FG#2, FG#3, FG#4).

Information about New Technologies

While most participants are able to assess the energy consumption of their current technologies through monitoring, they evaluate new energy-saving technologies using a variety of information sources. As shown in Table 4, the companies in our sample rely on multiple sources of information to learn about new technologies and practices. The most common information sources that data centers consult are manufacturers, in-house testing, conferences, and peers. External consultants and trade publications were less commonly mentioned, particularly among the focus group participants, which included most of the localized and enterprise data centers in our sample. Several focus group participants mentioned subscribing to trade journals but suggested that these publications are not influential sources for information. One participant commented that, *“I have a bookcase of magazines I don’t have time to read”* (FG#1). Information provided by government or regulators was not mentioned

as a primary information source by any of the participants, with the exception of one interviewee who consults information provided by Energy Star.

Table 4: Sources data centers use and trust for information on new technologies and energy efficiency

Information sources	Focus groups	Interviews
Manufacturers/suppliers	6/6	7/7
In-house testing	6/6	6/7
Conferences /trade shows	6/6	6/7
Peers	5/6	6/7
Outside contractors & consultants	6/6	3/7
Trade publications & internet	5/6	2/7
Regulators & government	0/6	1/7

Participants in every focus group and interview mentioned using information from manufacturers and vendors as a starting point for research about new technologies. A focus group participant explained, *“An important stream of information about what is new and upcoming in the industry technology comes from our vendor partners and OEMs [original equipment manufacturers]... those folks are always bringing to us what’s new, and what’s cutting edge”* (FG#1). However, other participants stressed the limitations of energy efficiency data provided by manufacturers. A focus group participant noted that, *“The technology is growing and changing so fast that [manufacturers and vendors] don’t have the proper metrics, or at least they’re not staying far enough ahead of the technology to get hard data back to the customers they’re trying to help”* (FG#4). Because data centers are individualized facilities that vary in size, climate, load and other dimensions, average estimates are of limited usefulness (INT#2, FG#5, FG#6). Another focus group participant summed up these limitations by saying, *“Have you bought a car that has a sticker that says that they usually get 25 miles to the gallon, but when you take it out on the road, you drive it, and you only get 18 to 20? The sticker value isn’t how it actually runs”* (FG#3). Participants also cautioned that while data about facility energy efficiency is usually reliable, information about the performance and energy consumption of IT equipment is *“kind of worthless”* (FG#4).

Because of the need for facility-specific data about the performance of new technologies, participants in all focus groups and interviews mentioned additional ways they collect information to supplement what

they learn from manufacturers and vendors. Many firms conducted their own testing before widespread adoption. For instance, a focus group participant thought of directional floor tiles as *“a gimmick at first, but we actually tested them and they worked exactly as designed... it was the biggest no brainer. We saw rack temperatures go down by two degrees, just by swapping out the floor tiles”* (FG#6).¹⁸ Some of the larger firms had their own in-house R&D divisions to pilot and develop new technologies (INT#2, INT#3, INT#4, INT#5, INT#6). Manufacturers sometimes facilitated in-house testing by allowing data centers to pilot new technologies for free to demonstrate effectiveness (INT#1, FG#2). Others—particularly focus group participants from small- and medium-sized firms—cited a lack of resources and expertise and concerns about losing uptime limiting their ability to conduct testing (FG#4).

Learning from peers in the industry was also repeatedly highlighted as an important source of information. One interviewee discussed the value of sharing information through industry initiatives like the Open Compute Project, describing it as *“some of the most open sharing that’s going on. Where the participants developed technology, they open source it and make all the details available and talk a lot very frankly about the positive and negative experiences, so we can all learn from each other”* (INT#7). Participants reported that their willingness to adopt a new technology is much higher once other companies report a positive return on investment, and few want to be on the leading edge of technology adoption (FG#2, INT#1, INT#7). A focus group participant noted that he was motivated to reconsider adopting blanking panels and cold aisle containment, *“now that I know that [a fellow focus group participant] saves 16 percent”* on energy costs (FG#5). A few technologies are already widely proven and not in need of further analysis. One interviewee noted, *“Nobody needs to write another white paper on the benefits of airflow and the payback. It has been proven to be an extremely economically rational decision”* (INT#2).

The most valuable information spillovers do not always come from the largest data centers. One interviewee noted that, *“Often the stuff that’s more innovative is not coming from the big guys. It’s coming from some of the smaller shops that are really doing things differently”* (INT#4). Some also pay particular attention to European data centers for the latest trends in energy efficiency (FG#5), while others mentioned looking at energy efficiency practices in other industries, such as manufacturing (FG#4, INT#4). The reliance on peers and hesitation to be the first mover means that the pace of

¹⁸ A directional floor tile “directs” the air flow towards the rack instead of exhausting it straight out. Appendix A also contains definitions for various types of energy efficient facility and server technologies and practices.

adoption can be slow for the most nascent technologies. One interviewee discussed a decision to pass on geothermal cooling, saying, *“it’s got to have spent some time in the field. We would not purchase the latest and greatest.... I don’t know that many people [who] have it”* (INT#1). A focus group participant summarized that, *“In the IT world and the data center world, everyone is still trying to come up with what is the best solution, so everything is a novelty, everything is expensive”* (FG#4).

Uncertainty about New Technology Performance

Even data centers that take advantage of in-house testing or knowledge spillovers from peers still face some degree of uncertainty about future energy savings and the price of energy. Decision-makers’ attitudes towards risk and uncertainty can influence their technology adoption decisions. The hesitancy to adopt a new technology in the face of uncertainty may in part be due to loss aversion. Loss aversion is a term from the behavioral economics literature used to describe situations in which people put greater emphasis on avoiding a loss than on achieving a gain of a similar magnitude when making decisions under risk (Tversky and Kahneman 1992). Loss aversion can hamper adoption of energy-saving technologies if people view the increased upfront cost of the technology as a certain loss that is weighed against an uncertain gain of reduced energy costs (Greene 2011). In the context of data centers, managers are especially concerned about loss of uptime as a potential risk when considering new technologies. While participants acknowledged that the financial gains from energy-saving technology can be substantial, the threat of loss of uptime looms large in the industry. Focus group participants described this risk by saying, *“If I screw up, I could take the company down,”* (FG#2), and *“the risk [of shut down] may kill [the proposed upgrade] dead in the water. These are the most risk averse human beings you’ll ever meet, in the data center industry”* (FG#3). In contrast, an interviewee from a large data center noted, *“We are a little bit more tolerant to risk. We are a little bit more experimental”* (INT#3).

Temperature set points provide one example of the risk of system failures discussed by focus group participants. While new equipment can withstand temperatures up to 110 degrees, yielding substantial savings on cooling costs, most data center managers do not let facility temperatures rise above the 90s. They worry that if the system goes down, it will *“cost the firm a billion dollars,”* and *“they’re going to be on the hook...for any errors”* (FG#4). Section 4.4 provides more discussion about the potential for tradeoffs between data center energy efficiency and reliability.

Another type of uncertainty that data centers struggle with is predicting how intensively the facility will be used. This uncertainty affects investment decisions because some energy efficiency investments, like cold aisle containment, only achieve a positive return on investment in high density operations (INT#7). An interviewee from a large data center explained, *“the most difficult things to model are customer behavior and how much they’re going to use a piece of infrastructure; if they don’t use it much, the efficiency is bad,”* and, *“just like a bus driving around with only one passenger in it... it’s just far less efficient than if you had a hundred people on there”* (INT#2). A focus group participant complained that because manufacturer energy efficiency estimates are contingent upon the assumptions about utilization, it is difficult to develop firm payback estimates (FG#6). Some companies address this uncertainty by building more capacity—and hence using more energy—than they need initially to allow for growth over time, even though they acknowledge that energy efficiency performance is poor while the facility is under-utilized.

4.2 Split Incentives

Split incentives can occur when the person that makes capital investment, use, or maintenance decisions is not the same as the person who pays for electricity. More generally, individuals that do not bear the electricity costs may behave in ways that undermine the effectiveness of energy-saving technologies (Jaffe and Stavins 1994; Gillingham, et al. 2009). Questions posed to focus group participants and interviewees aimed to assess the potential for possible split incentive problems for data centers. In this section, we discuss split incentives between the facility and IT departments within a data center; between data center owners and their lessors; and between the company purchasing equipment and the contractor repairing or maintaining it.

We find moderate evidence of split incentives between departments within a company, and between lessors and lessees, but only weak evidence of split incentives between outside contractors and data centers (see Table 3). In many cases, companies noted they have employed strategies to reduce cross-department split incentives such as creating integrated oversight teams, training employees, and charging individual departments for electricity use. Realigning incentives between lessors and lessees, however, appears more difficult unless the data center charges the lessee for actual electricity use. While still relatively rare, such pricing strategies appear to be growing in popularity, which suggests that the adoption of monitoring technology that would remove one barrier to understanding and then passing along the costs of power use to its customers.

Company

A data center typically has both IT and facilities departments. The IT department is responsible for ensuring the service, safety, and security of data and makes purchase and upgrade decisions for servers and software with those goals in mind. The facilities department is responsible for maintaining the buildings and infrastructure, and placing equipment to ensure air flow and heating and cooling needs are met. The facilities department also typically pays for the electricity for the entire data center (NRDC 2014). Without coordination between them, it is possible that the IT department will make purchasing decisions that undercut energy efficiency improvements to air flow or other aspects of heating and cooling on the facilities side. In other words, it is possible to have misaligned or split incentives between the two departments. A recent case study found that organizational changes to help align incentives were critical to recent energy efficiency improvements in eBay's data centers (Schuetz et al. 2013).

The potential for split incentives within their own data centers due to disconnects across departments was discussed by participants in every focus group and interview. Among focus group participants, siloed decision-making was fairly common. As one participant described, *"IT and facilities are separate parts [of the company] that don't meet until you get very, very high up"* (FG#2). In addition, several participants noted that the IT department is often more concerned with maintaining service than with saving energy. For instance, one participant stated *"IT people, for the most part, don't care about power"* (FG#4). Another agreed, noting that *"my IT department, frankly, [is] not very concerned about the efficiencies, and the technologies of the data center. They want a service, and they want it when they want it"* (FG#1). Several participants noted that a breakdown in decision-making across departments for energy efficiency investments is a problem for the industry as a whole. One manager stated that *"based on the data centers that I've personally audited, in 95 percent they [the facilities and IT departments] are separate. The IT organization is not involved with the energy cost, while the facility engineering department provides and therefore pays for the power"* (FG#3).

The disconnect between departments can also affect investment decisions for a new data center. One participant noted that project designers have a disincentive to invest in energy efficient equipment that has a higher upfront cost, even if it saves money for the company over time, because they receive a bonus if they come under budget: *"A lot of people are still first cost. 'This is the cheapest and best one. That's what we're buying.' You say, 'Well, hang on a sec. I ran the numbers and that unit's going to cost \$500,000 over the next five years. This one's going to cost \$350,000. So, in five years, you're going to pay \$150,000 [more].' [They respond], 'Yeah, well, we don't have it in the budget to buy the more expensive*

one.’ What’s [their] incentive? They’re going to get a performance percentage [if they] come in under a certain amount on a project” (INT#1).

In many cases, efforts are underway to better integrate decision-making across departments. A subset of participants in three focus groups and every interview worked in companies that were implementing strategies to better coordinate across groups or realign incentives. For instance, in some companies a team or person was assigned an explicit oversight role in the company to help insure more integrated decision-making. One interviewee observed, *“It’s really one team. It’s managed in a very integrated way. Obviously, it’s a big team so there are subgroups...Each of those [subgroups] has a budget, but that’s all managed and coordinated by one central [team] within the whole organization” (INT#6).*

Likewise, a participant noted the advantage of having the IT and facilities departments report to a single person in the company: *“I run the data center. I run the boilers and the generators and when the toilet backs up, they come to my office. But I’d rather deal with that so that I can stipulate exactly what kind of [attention each] will get and how it will get done” (FG#3).* For other companies, coordination was encouraged between the departments on a more routine basis. For instance, one interviewee stated, *“We have a bidirectional relationship with our operations staff and our engineering staff” (INT#3).*

Training was also utilized in some companies as a way to encourage IT staff to consider energy efficiency in server purchase decisions. One interviewee noted, *“We have a training program that everyone who work[s] in the data centers takes. It highlights, these are the things that are important to us. One of them is energy efficiency and finding opportunities for energy efficiency” (INT#4).* Similarly, a facilities manager in another focus group stated, *“I’m the one that has to go educate them on why this server or this system is more appealing than that system, because it has an efficiency, and it fits in our data center, of which they have no concern unless I tell them” (FG#1).*

A few participants addressed within-company split incentives by charging electricity use back to individual departments. One manager noted that *“the electricity costs come out of the data center budget, and [then] we provide a consumption-based charge back to our product groups or product units internally” (INT#3).* Another participant observed that, while it is not easy to charge the electricity back to the users, it is important to *“make our user realize what the real cost of their equipment is” (FG#1).*

Colocation Facilities

A “colo” or colocation facility is one that leases space in a data center to another party. The data center itself typically provides and manages the cooling, heating, and other support infrastructure—including

paying for the electricity—while the lessee or tenant often determines the IT equipment utilized, such as the servers, racks, and management systems used to store data and supply computing services (Bullock 2009).¹⁹ The potential for split incentives occurs when the cost of electricity is not passed through to the lessee, resulting in equipment choices that may undercut the overall energy efficiency of the facility. NRDC (2014) points to a number of barriers to aligning incentives between service provider and tenant in colocation facilities, including the desire to keep costs low, competing priorities (e.g. reliability, security), adequate monitoring to reflect the actual cost of power and cooling in pricing, and utility incentive programs aimed at the service provider instead of the tenant.

Within our sample, the possibility for misaligned incentives between lessees and lessors was discussed in every focus group and five interviews. As highlighted in the summary statistics in Table 1, about a third of the participants managed data centers that leased space to other companies. The degree to which a data center operated as a colocation facility varied widely among participants, from less than 10 percent of available racks to over 90 percent. In addition, more than half of the sample leased space from others, including some overlap with companies that *provide* colocation services, though a quarter of these data centers leased a relatively small portion (i.e., 20 percent or less) of their space. In addition, the majority of the data centers outsourced facility management, at least on occasion.

Participants varied in the degree to which they attempted to influence tenant equipment choices (short of directly pricing electricity use). One interviewee noted, *“I don’t get involved in anything inside the racks”* (INT#7). Others reported that they offer guidelines and investment advice to tenants based on observed energy use (FG#3, FG#2). The degree to which tenants were receptive to this type of information ranged from the opinion that tenants are attracted to data centers that are more energy conscious or *“green”* (FG#1, FG#2) to one in which *“the primary expectation of the customer is that the data center provide reliability (100 percent uptime). They don’t want to bear the risk of energy costs and don’t want to have to think about energy efficiency”* (FG#4).

¹⁹ Other arrangements are also possible. An enterprise service agreement stipulates that the lessor fully manages and make all necessary upgrades to both infrastructure and IT equipment. In these cases, we would not expect split incentives to occur. Only one participant noted that they regularly lease or outsource both facility management and servers from another company. However, another participant who provides colocation services noted, *“We have two different operations. For web hosting, we have client managed services... if [the client] wants to manage his own equipment one hundred percent and upgrade it, that’s up to him, that’s how we write his agreement. Or we’ll pay [for] an enterprise service, which is fully managed by the corporation and everything is taken care of from soup to nuts, they don’t have to come near the place”* (FG#6).

We asked the participants that routinely lease space in colocation facilities who pays for their electricity. Several participants stated that they are charged for the electricity they use. As one participant observed, *“We know what the power use is. We pay the bill”* (FG#4). Others noted that the owner paid for the electricity and did not directly pass these costs along to the tenant. In addition, a participant who leased space in his data center to other companies noted that pricing varied with the size of the tenant: *“We have customers that [are charged for] actual usage of power consumption; for smaller [customers], we’ll just give them a flat rate”* (FG#6). Questionnaire responses indicate that the majority of colocation facilities in the sample use a combination of pricing strategies: space-based, space and power block pricing based on expected use, and space and actual power use. See section 3 for more information on pricing strategies.

According to participants in several focus groups, a typical contract in the colocation industry is based on square footage or the number of racks utilized, adjusted to account for expected (but not actual) energy use. Some contracts also account for power usage effectiveness or PUE (INT#6, INT#7). While the status quo has been to only account for expected energy usage when leasing space, a number of participants observed that it is becoming more common in the industry to charge for actual power usage. As one interviewee stated, *“the idea is to align incentives between us and the customer so that both parties win when the customer is more efficient”* (INT#2). Discussions in several focus groups identified two main reasons why charging for actual usage is growing in popularity (FG#1, FG#2, FG#3). First, square footage is a poor proxy of energy use for some types of users (e.g., scientific computing allows one to use a given physical space in a data center more intensely than other types of computing). When a data center fails to account for these differences, it charges the same amount to a tenant that uses 5 KW per rack as one that uses 50 KW per rack. Second, monitoring technology has become more widely available, which allows the data center to track how energy use varies at the rack or server level. One interviewee noted that, once you have this type of monitoring data, pricing energy use to more directly affect tenant investment choices is possible: *“As a colocation provider, we will always be siloed from the IT decisions of the customer. So in some ways there’s no way to overcome that barrier. The way we try to do that is with information and incentives, providing real-time data and aligning billing incentives to enable us and customers to make optimal decisions”* (INT#2).

Firm-Contractor Relationships

About one third of focus group and interview participants indicated that they outsource maintenance and repairs in their data centers, while approximately 40 percent indicated that they sometimes outsource this service. In spite of fairly widespread use of contractors to service equipment, potential

split incentives between contractors who maintain and repair the equipment and data center managers were only discussed in three focus groups (FG#4, FG#5, FG#6).

In two of these focus groups, no one highlighted split incentives as a major issue. One participant noted that, while he outsourced maintenance and repairs because he lacked in-house staff, he didn't think this had any bearing on the data center's overall energy efficiency. Another participant stated that he monitored equipment closely such that if something was not working as it should, they could raise it with their contractors immediately (FG#5). In the remaining focus group, participants discussed the possibility that service technicians may not be as aware of energy efficient technologies as the data centers themselves: *"Field technicians aren't necessarily up to speed with their own products. That's a big problem for us as the end users in a mission critical environment"* (FG#6). In this instance, the data center relied on in-house staff to correct or circumvent problems.

4.3 Financial constraints

In reviewing participants' statements on whether firms faced financial constraints when considering energy-saving investments, we distinguish between two types of constraints: limited access to outside financing, also referred to as liquidity constraints, and competition from within the firm for other uses of existing funds, also referred to as the opportunity cost of capital. While liquidity constraints—when they exist—may result from a market failure, having a high opportunity cost of capital does not. We do not find much descriptive evidence to suggest that firms are not making energy efficient investments in data centers due to limited access to outside financing. On the other hand, several participants highlighted that subsidies and incentive programs often help tip the scale towards some investments with longer expected payback periods. In addition, competition among projects within the firm for available funds appears to be a substantial barrier to investment (see Table 3). To more adequately account for the role these factors play, we suggest that analysts consider using private discount rates—which are typically higher than those used to inform government policy—when evaluating the implications of private investment decisions.

Access to Outside Financing

Liquidity constraints may contribute to limited adoption of technologies if firms cannot obtain sufficient funds to cover upfront costs (Gillingham et al. 2009). While participants varied in their use of cash or external financing for major investments, none of the participants mentioned the inability to obtain a

loan as a reason for limited or delayed energy efficiency investments in their data centers. As one participant explained, *“Financing for me has no input whatsoever. It is not a money availability issue...What’s the total cost of ownership, what’s the ROI?”* (FG#2). Another noted, *“...[t]here is a lot of financing available from manufacturers these days”* (FG#3).

While financial constraints were not a limiting factor in investment decisions, participants in every focus group and interview mentioned using rebates from utilities or government programs as a way to make energy-saving investments more attractive. Many participants characterized them as a major factor in tipping the scale toward adopting an energy-saving technology: *“I would say I have spent many millions on things that I wouldn’t have spent if the tax incentive hadn’t been in place”* (FG#2). Examples of specific investments where utility or government (tax) incentives were a deciding factor in adoption include: investing in automated controls, airside and waterside economizers and cold aisle containment, LED lighting, and even some shifting to wind or solar power (FG#2, FG#4, INT#7). One participant explained that various government incentives were *“the only thing that ...took it [the LED lighting retrofit] from being completely unrealistic from a price and return on investment [perspective] to something we couldn’t turn down”* (FG#2). Analogously, another mentioned that they had not invested in energy efficient lighting in part due to the lack of utility incentives in their area (FG#4). Another participant identified incentives from power providers as the only reason for adopting hot aisle containment in one data center even though the company’s standard elsewhere is cold aisle (INT#5). For some, tax incentives also played a key role in decisions about where to locate a new data center (INT#3, INT#6), though one participant noted that siting incentives are not necessarily related to energy efficiency (FG#2). Several emphasized that working closely with the utility is essential, and some can even negotiate electricity rates directly when making upgrade decisions.

Opportunity Cost of Capital

At least some participants in most of the focus groups and interviewees raised competing uses of capital as one reason for limited investment in energy-saving technologies that appear financially attractive based on engineering-based NPV calculations. For example, one data center manager explained, *“I have to show paybacks, short-term payback because I’m competing with other people within the company. You know, somebody wants to build a building or buy a piece of medical equipment or do something, I’m competing with them”* (FG#4). Another stated, *“If I am doing an efficiency improvement and I can’t make it pay for itself in three years, there is virtually no chance I am going to get the money...I am not going to*

get it funded if it is not paying for itself that quickly because we can take that same money, invest it in something else that is going to grow the business, and return profit” (FG#2). Similarly, others indicated that all capital investments come out of one budget, so energy efficiency-enhancing projects need to make the same return as other capital investments (e.g., FG#5, FG#6). One interviewee noted that at his company, *“we have a responsibility to our shareholders to get the most return for the amount of money that we have” (INT#3).*

Participants’ statements also highlighted the importance of the discount rate assumed in payback calculations: *“[The existence of a payback gap] speaks to alignment of incentives and also includes ideas about time value of money and needing to discount future cash flows... payback period does not usually include discounted cash flow and so that’s one of the major reasons why I think it can appear like there’s this discrepancy between the payback period that I want and the useful life of the equipment” (INT#2).* If a firm’s internal rate of return is higher than the discount rate assumed in engineering-based NPV calculations, then lack of adoption may indicate a rational decision to pursue more attractive opportunities available to the firm. Use of a consistent rate of return for investments within a company does not suggest the presence of a market failure. Analyses of new technologies better reflect the likely private benefits to firms when calculated using private rather than social discount rates, which are often used in engineering-based analyses of regulations or other energy efficiency policies.

4.4 Tradeoffs between Energy Efficiency and Other Factors

This section summarizes participant statements about the potential for tradeoffs between improved energy efficiency and other important attributes that could present real costs to firms when adopting a new technology or strategy. While sometimes challenging to quantify for inclusion in a NPV calculation, these factors may make a new technology genuinely less attractive or costlier. In general, a key concern of participants was how energy efficient investments affect data center reliability. Heterogeneity in climate and other geographic- or data center-specific factors may hinder adoption of some types of energy efficient technologies. While not discussed in every focus group and interview, many participants also considered how easy a new technology would be to maintain and how it would be viewed by customers. While examples of tradeoffs dominated in these discussions, participants also acknowledged potential synergies in some instances.

Reliability Tradeoffs

The literature points to the possibility that investments in energy efficiency may affect the reliability of the service a data center provides to its customers. The terms reliability and uptime are used in the industry to refer to the amount of time a data center is operational. As one participant stated, *“We are shooting for 100 percent uptime”* (INT#4; INT#1). To ensure reliability of the data center under a variety of conditions (e.g., unanticipated power outages or problems with servers), data centers often build redundancy or backup into their systems (Brown, et al. 2007).

Potential tradeoffs between energy efficiency and reliability were discussed in every focus group and all but one interview. There was broad consensus among participants that they would not invest in a more energy efficient technology or practice that compromises the reliability of a data center’s system. As one participant described, *“everybody wants you to save money, but doing certain things where you might shave some kilowatts is not worth the risk of any impact [it] would have on reliability”* (INT#1). Another participant observed, *“if [an energy efficient option] in any way compromises uptime, it is a non-starter”* (FG#2). Examples of energy efficient technologies that participants said they did not use due to their impact on reliability included power management systems for servers and flywheels as backup power. One manager noted, *“in a large scale data center environment, if the servers decided to put themselves into a reduced power state, performance could go significantly where we do not want it to”* (INT#3). With regard to rotary flywheels, a focus group participant stated, *“I like being old school. I like to know there is 15 minutes left on the battery just in case I have to do a shutdown for whatever reason. While [rotary flywheels] are greener, batteries give you 15 minutes to a half hour”* versus less than a minute (FG#3).²⁰ Some participants also mentioned that concerns about how higher ambient temperatures may affect system reliability have slowed the adoption of adjusting set points to allow for a higher ambient temperature in their data centers. Recent studies show that this tradeoff is relatively small and manufacturers have begun to account for higher temperatures when designing new servers (FG#4). However, others mentioned that higher set points resulted in the server fans running so hard that they undercut the expected energy savings, and the fans started to fail at a higher rate as a result (INT#1).

Participants disagreed on the extent to which energy efficiency should even be considered when making decisions about system redundancy. Generally speaking, participants were less open to considering energy efficient investments that could impact critical operations within the data center (e.g., tracking

²⁰ A rotary flywheel mechanically spins a large, heavy disk at high speed to extract and store energy.

missiles or making drugs versus opening a Word document). Some participants looked for the most energy efficient option for a given level of redundancy, while others did not consider it at all. As one participant stated, *“There is no tradeoff for reliability. It must be reliable”* (FG#3).

That said, several participants recognized that many data centers set higher redundancy requirements than they actually need operationally, which implies that there may be opportunities to lower energy costs without compromising the reliability of the system. One participant noted, *“we will spend the money to draw more power because we want to protect ourselves from failure. We’ve done that and we’ll continue to do that for those applications that we deem critical. But [for] those that aren’t, we won’t spend that kind of money on them”* (FG#4). Another participant agreed, noting that it is important for the industry to rethink its approach to redundancy for non-critical uses: *“There’s this trend among university and noncritical IT shops to not have as much redundancy as we used to have...We [used to have] dual power feeds from the utilities [but] we’ve never had both of them go down at the same time in the last 15 years, so [we] start wondering whether all that redundancy is necessary...There are energy savings to be [gained by] reexamining the whole resiliency thing”* in some industries (FG#5).

Finally, a few participants also noted that the tradeoff between reliability and energy efficiency will likely become less important in the future and could even be viewed synergistically in some cases. One participant noted that greater reliability allows for more stable temperatures, which then allows them to increase the density of their racks and to cool them more effectively (FG#4). Increased virtualization—meaning that demand for data services can be met with fewer physical servers—also improves energy efficiency. As one participant noted, this allows for greater control over data center operations, making it easier to manage and keep them running (FG#6). Another participant stated that as old equipment reaches the end of its life, there are opportunities to invest in more reliable equipment that also improves energy efficiency (FG#5).

Data Center Heterogeneity

Another theme discussed by several participants is that a technology that works for one data center may not be appropriate for another due to differences in location and existing system constraints.

Climate was acknowledged in all focus groups and interviews as a factor that can affect how well certain energy efficient technologies function in a data center. For instance, managers paid attention to differences in ambient temperature and humidity when assessing whether to use airside or waterside economizers or outside air to cool a facility because they do not work well in hot climates. Likewise,

hotter climate makes it more difficult to run servers in a power-saving mode and shortens the longevity of equipment. Local safety and fire codes were also mentioned in several focus groups as a potential barrier to investment. An example that illustrates this point is adding containment to racks: *“There are hidden costs that you have to be aware of like the sprinkler head cost and then the smoke detectors, another cost”* (FG#1). Two other participants said their companies were reluctant to consider some containment measures like strip curtains because the *“fire marshal has been adamantly against it”* (FG#5). In addition, a participant mentioned that they considered ambient noise levels when purchasing a more energy efficient generator because of their urban location (FG#3).

With regard to siting a new data center, temperate climates were viewed as more attractive from an energy standpoint but unlikely to outweigh factors such as distance to clients, customer demand, zoning, and real estate prices. Several participants noted that the largest data centers in the country are located in very hot climates. One manager we interviewed observed, *“cooler climates lead to greater efficiency of the data centers, so that’s a consideration always. But in the scheme of things, the heat generated within the data center far outstrips the heat that might come in through the [building] envelope”* (INT#2). That said, another participant stated that, in the past, siting a facility in a good climate for energy efficiency *“was an afterthought. Now, it’s more part of the discussion”* (FG#4).

Participants in five focus groups and four interviews discussed space constraints and the ease with which new energy efficient technologies could be integrated into an existing system as factors when evaluating whether to make potential energy efficiency upgrades. One participant stated, *“you can’t really start from scratch”* (INT#4). If a new technology cannot be supported by the current platform or requires substantial adjustments to other parts of the data center system then it is less likely to be adopted. In the words of one participant, *“[you evaluate] how you would incorporate that [new equipment] into what you already have, rather than having to go in a whole new direction. If somebody had something that did the equivalent for cheaper, that wouldn’t necessarily influence the purchasing decision”* (INT#1). Two managers we interviewed also pointed out that since they already are substantially more energy efficient than the average, many of the technologies available are not expected to deliver additional savings for them. On the other hand, as one manager observed, *“new facilities are much easier to put new technology into”* (INT#3).

Maintenance Tradeoffs

The relationship between energy efficiency and maintenance of new equipment was discussed in four focus groups and all but one interview. Several participants noted that it is important to evaluate how easy or difficult a new energy efficient technology will be to maintain prior to investment. As one participant described it, *“there’s equipment out there that might be really efficient but is designed in a way where replacing anything is a royal pain. We work really hard to make sure that all those kinds of jobs can be done easily and efficiently”* (INT#6). Another participant opined, *“I don’t think across the board you could say that [more] efficient products require higher maintenance. You could make an argument both ways. You could say that the more efficient it is, the less infrastructure and upkeep you need because it’s just more elegantly designed and requires less cooling and less power. You could also make the argument that efficient technologies especially at the forefront are newer and emerging. LEDs probably have more issues than incandescent light bulbs because one has been around for a hundred years”* (INT#2). Table 5 summarizes examples of maintenance tradeoffs or synergies discussed by participants for several energy efficient technologies.

Table 5: Examples of Maintenance Tradeoffs and Synergies for Energy Efficient Technologies

Energy Efficient Technology	Tradeoff or Synergy?	Participant description
Flywheel UPS	Tradeoff	Cost of flywheels with maintenance factored in was triple the cost of batteries
Transformless UPS	Tradeoff	More technology inside the UPS that requires more maintenance and upkeep
New computer room air conditioning (CRAC) units	Tradeoff	Old CRAC units had a longer useful lifespan with minimal breakdowns; only had to change the belts
Liquid cooling of servers	Tradeoff	Concern that if a connector for the tubes failed, servers would be damaged by liquid
Variable fan drives	Synergy	Reduced maintenance on CRAC units by 70 percent
Virtualization	Synergy	Reduced maintenance and support requirements; easier to monitor what is happening

Several participants also voiced concerns about the degree of familiarity the engineering staff has with new more energy efficient equipment. One participant stated, *“I’d rather have an OE [operations engineer] who’s more comfortable with the equipment that he has, that is less efficient, than have a brand new piece of equipment come in and have to re-learn it”* (FG#6). A related issue voiced by a few

focus group participants was concern that fewer companies provide maintenance services or parts for newer technologies than for existing but less efficient technologies.

Customer Concerns

Customer views about energy efficient investments were discussed in five focus groups and three interviews. Customers' desire for sustainability seemed to encourage more investment in energy efficiency at many data centers. One participant noted that investments in energy efficiency helped the company "*tell a good story*" to customers and nudged them toward greener technologies, even though the primary motivation for investment remained cost savings (FG#1). Energy efficient investments that improved public relations were described as a "*bonus*" by participants in another focus group (FG#3). Still another focus group pointed out that if there was a marketing or public relations opportunity associated with energy efficiency, it gave them greater flexibility with regard to the expected return on investment (FG#2). While many participants expressed the desire to invest in energy efficiency where it made the most economic sense, the possibility of greenwashing was raised (i.e., doing something that looks good but doesn't actually reduce energy use). That said, a few managers stated that customer perception was rarely a factor in their energy efficiency investment decisions (FG#4, INT#5).

While there appears to be some degree of synergy between energy efficiency and customer preferences, some participants also gave examples when customer perceptions may slow investment in energy efficiency or preclude certain technology options. For instance, several participants mentioned that they would not use strip curtains "*[because of] the aesthetic factor. Regardless of how much energy savings there [are] you cannot bring people into the data center and see meat locker curtains*" (FG#6). Some participants also mentioned that they adjusted temperature set points in their data centers but stopped short of industry recommendations due to customer concerns (FG#4).

5. Conclusion

Participants in our study discussed a variety of technologies and strategies to improve energy efficiency at the data centers they manage. As expected, technologies with longer payback periods were not as widely adopted as those that yielded net financial gains within a few years.

Many participants discussed potential barriers to adopting energy-saving technologies, including split incentives between IT and facilities departments within the company and between colocation data centers and their tenants, uncertainty and imperfect information about the performance of new

technologies, and potential tradeoffs with reliability and other factors. While evidence was mixed, split incentives and imperfect information may represent market failures – which offer the clearest justification for government intervention. However, participants also described actions they had taken to at least partially mitigate these problems, such as improved monitoring of equipment energy use that allows colocation facilities to charge tenants for actual electricity used and extensive in-house testing of new technologies.

We also find that engineering-based NPV calculations used to evaluate new energy-saving technologies may be inconsistent with data center decision-making in a few respects. Data centers used a private discount rate consistent with the rate of return required on all investments in the company, which is typically higher than the social discount rate. In addition, they placed a premium on avoiding risks to uptime that is not typically quantified in NPV calculations. Finally, many data center companies conducted their own testing to evaluate the performance of new technologies, which represents a real transaction cost. To minimize such inconsistencies, analysts could use private discount rates that reflect firms' opportunity cost of capital, incorporate more refined real-world estimates of technology performance, and include the cost to firms of conducting testing needed to gather such information, as well as the value of lost uptime in NPV calculations.

We remind the reader of a few caveats when interpreting the results of this study. By its very nature, information from focus groups and interviews is only suggestive. It is not drawn from a representative sample of the industry. In particular, participants may be more interested in technology and sustainability issues than nonparticipating data centers. Despite these caveats, the results may help inform future analyses, the development of surveys, and policy discussions on barriers to energy-efficient investments in the data center industry.

References

Andrews, C. J. & Krogmann, U. 2009. "Explaining the Adoption of Energy-Efficient Technologies in U.S. Commercial Buildings." *Energy and Buildings*, 41, 287–294.

Belady, C., A. Rawson, J. Pflueger, and T. Cader. 2008. "Green Grid Data Center Power Efficiency Metrics: PUE and DCIE." Green Grid White Paper No. 6.

Brown, R., Masanet, E., Nordman, B., Tschudi, W., Shehabi, A., Stanley, J., Koomey, J., Sartor, D., Chan, P., Loper, J., Capana, S., Hedman, B., Duff, R., Haines, E., Sass, D., and A. Fanara. 2007. *Report to Congress on Server and Data Center Energy Efficiency: Public Law 109-431*. Lawrence Berkeley National Laboratory, Berkeley, California. LBNL-363E.

Bullock, M. 2009. "Data Center Definitions and Solutions." *CIO Digital Magazine*. August 14. Available at: <https://www.cio.com/article/2425545/data-center/data-center-definition-and-solutions.html>

California Building Standards Commission. 2013. California Energy Code 2013. California Code of Regulations Title 24, Part 6. <https://archive.org/stream/gov.ca.bsc.2013.06/#page/n0/mode/2up>

California Statewide Utility Codes and Standards Team. 2011. Draft Measure Information Template – Data Centers. 2013 California Building Energy Efficiency Standards. http://www.energy.ca.gov/title24/2013standards/prerulemaking/documents/2011-04-11_workshop/review/Data_Center_33011.pdf#page=3

Cisco. 2016. Global Cloud Index: Forecast and Methodology, 2015-2020. Cisco Public White Paper. <https://www.cisco.com/c/dam/en/us/solutions/collateral/service-provider/global-cloud-index-gci/white-paper-c11-738085.pdf>

Forrester Research, Inc. 2014. "Data Center Demand Survey Results." Study commissioned by Digital Realty. Spring.

Gillingham, K., Newell, R. G. & Palmer, K. 2009. "Energy Efficiency Economics and Policy." *Annual Review of Resource Economics*, 1, 597–620.

Gul, M., and G. Menzies. 2012. "Designing domestic buildings for future summers: Attitudes and opinions of building professionals." *Energy Policy* 45: 752–761.

Horner, N. and I. Azevedo. 2016. "Power usage effectiveness in data centers: Overloaded and underachieving." *The Electricity Journal* 29: 61-69.

Hsieh, H-F., and S. Shannon. 2005. "Three Approaches to Qualitative Content Analysis." *Qualitative Health Research* 15, 9: 1277-1288.

Jaffe, A. B. & Stavins, R. N. 1994. "The Energy-Efficiency Gap: What Does it Mean?" *Energy Policy*, 22, 804–810.

JLL. 2016. Data Center Outlook. North America. <http://www.us.jll.com/united-states/en-us/Research/US-North-America-Data-Center-Outlook-2016-JLL.PDF> .

Klemick, H., E. Kopits, A. Wolverton, and K. Sargent. 2015. "Heavy-Duty Trucking and the Energy Efficiency Paradox: Evidence from Focus Groups and Interviews." *Transportation Research Part A: Policy and Practice*, 77, 154–166.

Klemick, H., E. Kopits, and A. Wolverton. 2017. "Potential Barriers to Improving Energy Efficiency in Commercial Buildings: The Case of Supermarket Refrigeration." *Journal of Benefit Cost Analysis* 2017: 1–31.

Koomey, J. and J. Taylor. 2017. "Zombie/Comatose Servers Redux." Stanford University and Anthesis. <http://anthesisgroup.com/zombie-servers-redux/>

Masanet, E., A. Shehabi, L. Ramakrishnan, J. Liang, X. Ma, B. Walker, V. Hendrix, and P. Mantha. 2013. "The Energy Efficiency Potential of Cloud Based Software: A U.S. Case Study." Lawrence Berkeley National Laboratory, Berkeley, California.

- Miller, R. 2010. "The Next Efficiency Frontier: Underclocking." *Data Center Knowledge*. October 27. Available at: <http://www.datacenterknowledge.com/archives/2010/10/27/the-next-efficiency-frontier-underclocking> .
- National Resources Defense Council (NRDC). 2014. "Scaling Up Energy Efficiency Across the Data Center Industry: Evaluating Key Drivers and Barriers." Issue Paper IP:14-08-a. August.
- National Resources Defense Council (NRDC). 2012. Small Server Rooms, Big Energy Savings. February. www.nrdc.org/energy/files/Saving-Energy-Server-Rooms-IssuePaper.pdf.
- Ochieng, E.G., N. Jones, A.D.F. Price, X. Ruan, C.O. Egbu, and T. Zuofa. 2014. "Integration of energy efficient technologies in UK supermarkets." *Energy Policy* 67: 388-393.
- Olson, M. 1965. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Cambridge, MA: Harvard University Press.
- Pacific West Air Conditioning. 2014. "Water-Side Economizers." November. Available at: <http://www.pacwestac.com/energy-efficiency-water-side-economizer.html>
- Pellegrini-Masini, G., and C. Leishman. 2011. "The role of corporate reputation and employees' values in the uptake of energy efficiency in office buildings." *Energy Policy* 39: 5409-5419
- Radcliff and Best. 2005. Eds. *Polling America: An Encyclopedia of Public Opinion Research*. Westport, CT: Greenwood Press.
- Sallee, J. M. 2014. "Rational Inattention and Energy Efficiency." *Journal of Law and Economics*, 57, 781–820.
- Schleich, J. & Gruber, E. 2008. "Beyond Case Studies: Barriers to Energy Efficiency in Commerce and the Services Sector." *Energy Economics*, 30, 449–464.

Schuetz, N., A. Kovaleva, and J. Koomey. 2013. "eBay Inc.: A Case Study of Organizational Change Underlying Technical Infrastructure Optimization." Stanford University.

Shehabi, A., S. Smith, D. Sartor, R. Brown, M. Herrlin, J. Koomey, E. Masanet, N. Horner, I. Azevedo, and W. Lintner. 2016 "United States Data Center Energy Usage Report." Lawrence Berkeley National Laboratory (LBNL), Berkeley, California. LBNL-1005775.

Stasberry, M. 2014. "2014 Data Center Industry Survey Results." Uptime Institute Symposium 2014, Santa Clara.

Sullivan, R., and A. Gouldson. 2013. "Ten years of corporate action on climate change: What do we have to show for it?" *Energy Policy* 60:733-740.

The Green Grid (2011). "Case Study: The ROI of Cooling System Energy Efficiency Upgrades." White paper #39. Available at:
https://www.energystar.gov/sites/default/files/asset/document/CaseStudy_TheROIofCoolingSystemEnergyEfficiencyUpgrades_0.pdf .

Van Greet, O. 2017. "A Method for Estimating Potential Energy and Cost Savings for Cooling Existing Data Centers." National Renewable Energy Laboratory (NREL). Available at:
<https://www.nrel.gov/docs/fy17osti/68218.pdf> .

Wathaifi, M. H., 2009. "Free Cooling Using Water-Side Economizer: Process Cooling Application at a Plastic Extrusion Plant." *Energy Engineering* 106(6): 45-78.

Appendix A: Technology Descriptions

Facility Management

Technology/Practice	Description
Isolated hot/cold aisles	When rows of server racks face the same direction, hot exhaust air from the back of the first row gets sucked into the cool air intakes of the second row. Alternating hot and cold aisles orient the rows of server racks so that the fronts of servers face each other and the backs face each other
Hot-/cold-aisle containment	Physical barriers (e.g. strip curtains) between hot and cold aisles further prevent mixing of incoming cold air and hot exhaust air
Vented, directional, or perforated tiles	Tiles that guide cool air from ducts in the subfloor directly to servers at a specific flow rate
Blanking panels	Panels that cover unused spaces in racks so that cool air passes through the existing servers instead of around them
Variable fan or speed drives (VFDs or VSDs)	Allow fans on air handling units to vary air flow with changes in cooling demand.
Energy efficient lighting	In addition to updating to LED lighting, sensor-driven, motion-activated lighting allows lights to be on in specific areas of the data center as needed.
Energy efficient air handling units	Upgrading computer room air conditioning or air handling (CRAC or CRAH) units. ²¹
Airside economizer	Use cool filtered outside air when available instead of running compressors to reduce temperature of hot exhaust air for recirculation
Waterside economizer	For centers with chilled water plants, use the evaporative cooling capacity of the cooling tower to produce water that is sufficiently cool instead of running the chiller to reduce the temperature of hot water for use
Raising temperature set points	Industry standards recommended a server air intake temperature range of 65°F to 80°F. However, many data centers set their temperatures as low as 55°F. Likewise, chilled water temperatures are often set too low. Energy can be saving by raising the temperature consistent with recommended levels. ²²
Data center infrastructure management (DCIM) software	A management tool that may/may not emphasize energy use monitoring and management. Using sensors, it is a centralized data repository that offers an integrated view of facility infrastructure (space, power, cooling) and IT systems (servers, storage, networking, applications)
Modular facility design	Separate “Lego-like” discrete, prefabricated data center facility components that can be assembled quickly to expand a data center facility instead of initial over-building in anticipation of long term growth in demand
Predictive modeling of future deployment	Computer modeling to simulate impacts of proposed changes/deployment of new equipment in a data center; uses DCIM data as an input

Sources: “12 Ways to Save Energy in a Data Center,” Energy Star, www.energystar.gov; www.greengrid.org; www.datacenterknowledge.com .

²¹ Cooling for a CRAC or CRAH unit occurs by blowing air over a cooling coil filled with refrigerant or chilled water, respectively. For a CRAH unit, chilled water typically is supplied by a chilled water plant (i.e. chiller). While CRAC units only turn off or on, it is possible to add VFDs to CRAH units to control the supply of chilled water.

²² It is also possible to modify humidity set points consistent with industry standards to save energy. When set too stringently, given the proximity of CRAC units' exhausts and intakes, one CRAC unit may work to cool or humidify the air while another works to dehumidify and/or reheat the air.

Servers

Technology/Practice	Description
Virtualization software	Instead of using a separate physical server for each application, which results in only a small fraction of computing resources being actively utilized, it is possible to create virtual servers that allow multiple applications on a single physical server simultaneously
Under/overclocking	Overclocking causes a computer's processor or memory to go faster than its factory-rated speed. The extra speed boosts the work being done by the processor. Underclocking slows clock speed and is used to save power or reduce heat from the processor. This technique is designed to capture potential gains from machines where the workload varies over time. As the workload diminishes, the CPU slows. As it rises, the CPU speeds up.
Energy efficient server systems (e.g., for uninterruptible power supply (UPS), fans)	UPS systems provide backup power to data centers to reduce the risk of disruptions; can be based on battery banks, rotary flywheels, or fuel cells. ²³ Software that manages its use can increase efficiency of the UPS or CRAC/CRAH fans (see VFDs or VSDs)
Decommissioning of unused servers	Retiring servers that are not being used because they are too old or obsolete but are still running
Consolidation of lightly used servers	Moving away from one server per application and one backup server per active server via virtualization, clustering servers to reduce the number of backup servers, consolidating redundant applications.
Tiered data storage	Store data that is used less frequently on servers with lower-speed drives, which use less energy, while reserving higher-speed drives for data that is more frequently accessed
IT power management software (e.g. power capping)	Set a limit on the amount of electricity consumed by a server based on an assessment of real-time use instead of under an assumption of maximum utilization

Sources: www.datacenterknowledge.com; "12 Ways to Save Energy in a Data Center," Energy Star, www.energystar.gov; www.datacenterdynamics.com/power-cooling; www.infoworld.com; Miller (2010).

²³ A rotary UPS uses rotating mechanical motion to generate kinetic energy and bridge the gap in power when an outage occurs. It can use batteries or a flywheel as an energy storage source. A flywheel usually can only fill a gap of less than a minute, while a battery has a reserve of 5 – 15 minutes.

Appendix B: Participant Questionnaire

1. Which of the following do you do in your position?

a. Participate in decisions about purchasing equipment for your data center (e.g., servers, software, racks, lighting, cooling equipment)? ☐ Yes ☐ No

b. Participate in decisions about servicing or maintaining equipment for your data center?

☐ Yes ☐ No

2. How many data centers, server rooms, and server closets does your company own or operate?

Enter number here _____

3. [If more than one data center] For how many data centers are you involved in major equipment purchase and maintenance decisions?

Enter number here _____

4. The majority of the data centers you help manage are:

- a. Server closets or rooms (less than 500 square feet)
- b. Localized data centers (500 – 5,000 square feet)
- c. Enterprise data centers (> 5,000 square feet)
- d. Mega or utility-scale data centers (>10,000 square feet)

5. What is the average number of racks in your typical data center?

Enter number here _____ or ☐ if > 2000

6. How many servers are there in your organization as a whole?

- ☐ < 2000
- ☐ 2000 – 10,000
- ☐ 10,000 – 50,000
- ☐ > 50,000

7. What type of data centers does your company operate?

- ☐ Server rooms or closets
- ☐ Enterprise/corporate data centers

- ☐ Cloud/hyperscale computing
- ☐ High-performance/scientific computing
- ☐ Multi-tenant/colocation data centers

8. Which sectors does your data center support? (Name top three)

- a. IT services
- b. Energy and utilities
- c. Media and telecommunications
- d. Banking and financial services
- e. Health care
- f. Education
- g. Co-location/shared
- h. Other: _____

9. What state(s) are your data centers located in?

Enter states here: _____

10. When did you last pursue a major equipment purchase/upgrade for your data center?

- a. Within the last 3 years
- b. 3-5 years ago
- c. More than 5 years ago

11. How often do you replace servers in your data centers?

12. How often do you make major infrastructure upgrades to your data centers?

13. Do you actively participate in the EnergySTAR Buildings, Green Grid, or another third-party energy-efficiency program?

- ☐ Yes ☐ No

14. Have you had an energy audit performed by a third party within the last three years?

- ☐ Yes ☐ No

15. Do you track energy use on a regular basis? _____

a. What is the typical annual electricity consumption (kwh) of one of your data centers? _____

16. Do you track power usage effectiveness (PUE) or any other energy efficiency measures?

a. If so, what is the average PUE of your data centers?

17. Do you track server utilization?

a. If so, what is the average server utilization rate of your data centers? _____

18. Do you rent data center space *from another company*? ☐ Yes ☐ No

If yes, what percentage of your data center needs are met by renting space from others?
_____ %

19. Do you lease or outsource *from another company* any of the following elements of your data centers?

	Yes	No	Sometimes
b. Servers	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
c. Facility management	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
d. Maintenance & repairs	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
e. Other, specify: _____	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

20. Do you lease space in your data center(s) *to other companies*? ☐ Yes ☐ No

If yes, what percentage of your racks are leased to other companies?
_____ %

21. If you lease DC space to or from another company, does the tenant or facility owner typically pay for electricity costs?

22. Has your company considered or used any of the following technologies and energy saving strategies: (check all that apply)

	Has used		Never used		Don't know
	would use again	would not use again	would consider in the future	would not consider	
<u>Facility Management</u>					
Isolating hot/cold aisles	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Blanking panels	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Enclose server racks (e.g., strip curtains)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Data center infrastructure management software	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Raising temperature set points	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
List typical facility temperature: _____					
Variable speed fan drives	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Energy-efficient air handling units	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Modular facility design for data center expansion	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Airside economizer	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Waterside economizer	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Energy-efficiency lighting	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Predictive modeling of future IT deployments	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Other, specify: _____	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<u>Servers</u>					
Energy-efficient server hardware (e.g., power supply, fans, etc.)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Underclocking	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Virtualization	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Decommission idle servers	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Consolidate lightly used servers	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Improve data storage efficiency	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
IT power management software	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Other, specify: _____	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<u>[Multi-tenant DCs] Adjust pricing models to incentivize energy efficiency</u>					
Space-based pricing	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Space and power block pricing (differentiate power and energy)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Space and actual power pricing (differentiate power and energy)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Other, specify: _____	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Appendix C: Moderator's Guide

Appendix C: Moderator's Guide

I. WELCOME AND INTRODUCTION (5 minutes)

A. Introduction of the moderator

Good morning/afternoon. My name is <insert name>. I have been hired by an international research and consulting firm to moderate this focus group. Our client is interested in hearing your opinions about how data centers make investment decisions related to the adoption of new and existing data center technologies and what factors influence the decision process.

B. What is a moderator?

Before we begin, I want to let you know that I'm not an expert on data centers, IT, or the investment decisions you and your companies make. My job as a moderator is to:

- Help guide the flow of conversation
- Make sure everyone's comments are heard
- Ensure that the questions our client has are covered

You will see me referring to this outline during our session. The outline includes all the issues I need to raise with the group, and helps me keep the discussion on track. Since we have a lot to cover, I may have to break off the conversation to move on to another area. The questions at the end of the guide are just as important as the ones at the beginning.

C. Informed Consent

[Moderator: All participants will have reviewed and signed the informed consent form for their focus group participation prior to coming into the focus group room. This will be handled by designated staff at the focus group facility. Have a copy of the form that you can hold up.]

Before coming into the room, you were asked to review and sign an informed consent form. I want to highlight key points on the form to make sure we are in agreement.

- We are audio taping our discussion today. The tapes will be transcribed and our client will use the transcription to prepare a report.
- Your full names and the companies you represent will never be made known to the client and will not appear in the transcripts or any report that is written.
- For this reason, please refer to each other by your first name only. If you slip up, that's okay. We'll delete identifying information from the transcript.

D. Ground Rules

Before we begin, I'd like to review some ground rules for today's discussion.

- Because we are taping the discussion, I ask that you speak loudly and clearly and one at a time. If I think you are speaking too softly to be heard, I will ask you to speak up.
- Sometimes I'll go around the table and ask several of you for your input. At other times, I will just throw a topic open for general discussion. We want everyone to participate, but you don't have to answer every question.
- There are no right and wrong answers. Please feel free to disagree or question each other. We expect differences in how people see things. We want to know about these differences. It's important to tell us YOUR thoughts, not what you think others think or want to hear.
- Some observers are listening in by phone. They want to hear what you have to say about the topics we'll discuss, but we don't want you to feel constrained by their presence. They have each signed an observer confidentiality form.
- If anyone needs to use the restroom, they are located (specify). We are also providing water and snacks, which are available at (specify location) of the room. There is no need to stop the discussion to avail yourself of either.
- If you have a cell phone, please turn it off or set to vibrate.
- The session will last about 1 1/2 hours. I will do my best to get you out on time.

Does anyone have any questions?

E. Introduction of Participants

Let's start with introductions. Please introduce yourself—first name only—and tell us:

- What type of data centers does your company operate and what industries do you serve?

[Moderator: If they need prompts on “type”, offer these options: cloud/hyperscale, traditional enterprise, colocation, server closets, high performance/scientific computing.]

- What's your role in the company with regard to data center management?
- For how many data centers do you make purchase and management decisions?

NOTE TO MODERATOR: Questions 1 is a **warm-up question** to get the discussion going but should not last more than five minutes.

Probe **why** the factors mentioned are important, but no need to understand which rank first, second, or third or get consensus on what ranks in the top three. If you hear similar factors from several people, then a way to shorten the conversation is to ask the remaining participants if they agree with this list or would include different factors than those already mentioned.

II. OVERVIEW OF DATA CENTER INVESTMENT DECISIONS (20-25 minutes)

1. To get us going, what are the top 3 factors you consider when setting up a new data center or doing a major upgrade in an existing data center (e.g., involves replacing a major IT or building system)? *[Moderator: write factors on board.]*

To group (not one by one) once you have a list: Why these factors?

Prompt: [ONLY IF ENERGY USE IS NOT MENTIONED] No one mentioned energy efficiency as a top factor. Is it important, somewhat important, not important?

2. Starting from the idea that you want to set up or upgrade a data center ...
 - a. **[To the group]:** Where do you gather information on what new equipment, technologies and facility management strategies are available?

[Moderator: if group has trouble answering this question, mention possible sources such as: manufacturers/vendors, colleagues or larger (esp. cloud computing) companies, conferences/industry shows, in-house testing, outside organizations, general contractors]

- What is the role of the manufacturer in the types of technologies you consider?
 - Corporate headquarters (when relevant)?
 - Engineers that maintain the system?
 - Customers?
- b. Do these sources provide *reliable* estimates of likely energy use? Are there estimates out there that you don't believe or bother looking at? Why?

NOTE TO MODERATOR: Questions 3 - 6 relate to **split incentives** between those maintaining the system and those purchasing and/or using the equipment.

3. From the online questionnaire you filled out prior to this focus group, I see that [insert number] of you lease data center equipment or facility space *from* other companies.
- Do you pay based on space alone (racks or square footage), or are there charges based on other factors such as power use?
 - Who determines when major power and cooling infrastructure systems need to be fixed or upgraded? What is your involvement in this process?
4. On the flip side, from the questionnaire I see that [insert number] of you lease data center equipment or space *to* other companies. How is this relationship structured? Who determines when something needs to be fixed or upgraded?
5. For those who lease space or equipment *to or from* other companies, are bills based on *actual* measured or *estimated* electricity use? Do they include a single energy charge or some combination of energy and demand charges? *[Moderator: These are often called "service level agreements", esp. by colos and some enterprise DCs.]*
6. Is your data center run as a single organization with one budget, one team, one boss? Or is it broken into silos with separate budgets, teams, bosses?

- For companies that pay electricity bills in-house, how involved is the department responsible for paying the bill in technology purchase and maintenance decisions?

III. THE ROLE OF ENERGY EFFICIENCY (35-40 minutes)

NOTE TO MODERATOR: In this section, we want to know how energy efficiency fits into purchase decisions: are there **barriers to energy saving investments**? How is energy use **weighed against other factors** (e.g. performance or reliability)?

We are interested in why they do or do not adopt particular technologies. What factors make them more or less appealing? Please drill down a bit, including technologies already mentioned briefly in earlier discussion. Please ask for specific examples when none are provided.

At the very beginning of this focus group, you all mentioned energy efficiency as *[important, somewhat important, not that important]* in data center management.

7. How many of you have recently made major investments that improved *facility* energy efficiency? **[Show of hands – please note number of hands for tape.]**
 - a. **[To the group]** Can you give me a specific example? *[Moderator: If they have a hard time, suggest one or two areas: How about a technology related to cooling, lighting, air flow, power distribution equipment?]*
 - b. **Why** these technologies or approaches? What attributes make them appealing? What factors did you consider when deciding whether to make this investment?
 - c. Based on the questionnaires, some of you also considered ***[choose in advance several specific technologies from questionnaire; ask about ONE or TWO that have not already been discussed above]***. What are the main reasons you considered this technology/strategy?
8. How many of you have recently made major investments that improved *server utilization*? **[Show of hands – please note number of hands for tape.]**
 - a. **[To the group]** Can you give me an example? [e.g., virtualization]
 - b. **Why** these technologies or approaches? What attributes make them appealing? What factors did you consider when deciding whether to make this investment?

- c. Based on the questionnaires, some of you also considered ***[choose in advance several specific technologies from questionnaire; ask about ONE or TWO that have not already been discussed above]***. What are the main reasons you considered this technology/strategy?
 - d. Do you monitor equipment utilization? If so, for which equipment (servers, storage, cooling, power systems)? If not, why not?
 - e. Do you currently use all the energy efficiency technologies that are built into your servers? ***[If they say they turn these features off, ask why?]***
9. Are there energy or utilization efficiency technologies you have chosen not to pursue? Can you give an example?
- **Why** did you decide against purchasing these features? What factors did you consider when deciding whether to use them?
 - ***[Skip if consider all technologies on questionnaire or already brought up those not considered above]*** Some of you indicated on the questionnaire that you did not consider *[insert specific technology]*. What are the main reasons you don't consider them?
10. **[To the group]** How do you approach investing in new, cutting edge technologies that could potentially improve data center energy efficiency? ***[Moderator: ask for a specific example of something they are exploring now that was not listed on our questionnaire]***
- Do you try them out in a few facilities first? What generally tips the scale toward wide-spread investment?
11. We've been talking about major investments to reduce energy use. Do you also use maintenance or training (i.e., non-technology based) approaches to reduce energy use? **Why or why not?** Can you give an example?
12. I have a list of factors here that you might consider when choosing technologies or strategies to reduce energy use that I want to ask you about.

I'm interested in whether there are **tradeoffs** or **synergies** between these factors and **energy use**; are these factors **barriers or motivators for improved energy efficiency**?
[Ask about factors not already discussed in this context]:

- Performance
 - o Uptime/reliability (e.g., redundancy)
- Maintenance and repair issues
 - o Ease of installation, other maintenance and repair concerns
- Climate, other location-specific factors (e.g., cooling choices)
- Financing
 - o Borrowing constraints,
 - o Utility or state incentives
- Energy policy
 - o Company-wide energy efficiency policies or approaches

IV. OPINIONS ON PAYBACK (15-20 minutes)

NOTE TO MODERATOR: The next two questions are related to how they incorporate these factors into decision-making – do they calculate a **return on investment**? If yes, how are these aspects of the decision factored in?

13. Thinking about all the factors you've mentioned that affect your investment decisions, do you weigh them against the upfront cost and incorporate any of them into an ROI or payback calculation?

- a. If so, how?
- b. If not, how do you weigh those costs against upfront cost and energy savings?

[**Prompt:** Examples: uptime, maintenance costs, increased revenues from more compute cycles]

14. Some companies in other industries have told us that the payback period they need to justify investing in an energy saving technology/feature is less than the amount of time before they perform a major upgrade (for instance, a technology pays back in 1 -3 years but lasts 10 years).

- i. Is this true for you?
- ii. Why do you need an investment to pay back sooner?

[**Prompt:** If they don't say much or say it's a company rule of thumb, ask: Do you know why it's a company rule of thumb? Is it to hedge against uncertainties - e.g., future energy prices, true technology effectiveness?]

V. WRAP-UP

That is all the questions I have for today. I want to thank you for your input and sharing your expertise. Your time and ideas are incredibly valuable in helping us understand your decision making process.

15. Does anyone have any last questions or comments?