SCTE CABLE-TEC EXPO'15
OCTOBER 13-16 NEW ORLEANS

SCTE Society of Cable Telecommunications Engineers

Energy Efficient Cable Plant Facilities: Strategies to Increase Density through Capacity Reclamation, Site Configuration and Subscriber Based Financial Modeling

A Technical Paper Prepared for the Society of Cable Telecommunications Engineers
By Daniel Marut

Daniel Marut
Principal Engineer
Comcast Cable
1701 John F. Kennedy Boulevard
215.286.7319
Daniel_marut@cable.comcast.com

Thomas M. Murphy, Ph.D., University of Pennsylvania

Table of Contents

Title **Page Number**

Abstract 3
Introduction 4
Energy Intensity 5
1. Background 5
2. The MSO Energy Picture 6
3. Density as it pertains to National Data Centers 7
4. Density as it pertains to hubs and headends 12
    4.1. Introduction 12
    4.2. Measurement Methodology 13
    4.3. Measurement Implementation 14
    4.4. Foxboro's new services 16
        4.4.1. Site Characteristics 16
        4.4.2. Launch of cloud DVR services impact on density 17
    4.5. A Representative Hub Consolidation Effort 19
        4.5.1. Site Characteristics 19
        4.5.2. Site Density 19
    4.6. Lyndon regional consolidation to virtual-hubs 21
        4.6.1. Site Characteristics 21
        4.6.2. Regional Density 21
    4.7. Large scale regional CCAP deployments 22
        4.7.1. CCAP site characteristics 22
        4.7.2. CCAP impact on density 23
    4.8. Berlin policy driven density 24
5. Improving density using regression analysis 27
Conclusions and Recommendations 29
1. Conclusions 29
2. Areas for Further Investigation 30
Abbreviations 31
Appendix A 32

Abstract

The Energy 2020 program set energy use targets for the year 2020. These goals are achievable by continuing to implement solutions that focus on increased network data-throughput capacity. However, much of the energy in a cable plant is consumed by the outside portion of the cable plant, largely not impacted by current solutions, therefore the goals at inside plant facilities must be relatively more aggressive. Additionally, some low cost improvements to site infrastructure that improve energy use intensity at inside plant facilities are often left untapped. This paper clarifies the definition of energy intensity as applied to cable sites and uses various metrics to help drive decision-making for cable site improvement. It also describes a preliminary statistical model used to quantify the relationship between cable infrastructure and impacts on energy use, in order to improve planning of cable site upgrades.

Introduction

The Society of Cable Telecommunications Engineers (SCTE) Energy 2020 program demonstrates that the energy used by Multiple System Operator's (MSO) cable network plant, compared to data centers and office infrastructure, can be an order of magnitude higher.[1] The industry must make these network facilities more energy efficient in order to empower continued capacity and service growth, in a manner that prevents energy cost and availability from becoming barriers to that growth. In addition, there is growing policy pressure to address energy consumption from a sustainability perspective, e.g. lowering total greenhouse gas emissions.[2] The Energy 2020 program is exploring energy efficient technology, adoption of existing data center and new edge facility practices for energy management, and reduction of grid dependency through power cogeneration from alternate energy sources such as natural gas and solar.[3] Also, the program seeks to identify and promulgate practices and standards to improve the overall density of networking facilities through the implementation of power reducing strategies.

In this operational practices paper, the authors will explore a range of operational strategies focused on historical and real-time measurement used to demonstrate density changes of sites undergoing capacity reclamation and site re-configuration. This paper explores the use of subscriber-based financial, statistical modeling and real-time continuous measurement. This will be done by investigating cable infrastructure sites that are undergoing or underwent site upgrades. Given the variation in power measurement and benchmarking currently done in many of these facilities, effective strategies can be devised to benchmark density in a wide range of facilities, using both the utility supply and the equipment supply consumption meters. Several ongoing projects will be discussed in terms of density to help identify areas with high return on investment. Metrics such as kilowatt-hour (kWh) per subscriber are shown as useful tools for benchmarking site power density in terms of subscribers served, allowing establishment of financial modeling tied to subscriber lines of business for better long term planning for greater energy efficiency in cable critical facilities.

Energy Intensity

1. Background

The authors are members of the multidisciplinary Energy Density and Consolidation working group within SCTE Energy 2020 program.[4] This working group focuses on the multi-faceted area of energy density and site consolidation. Multiple issues cloud the area of energy density and site consolidation of cable plant facilities, which systematically reduce the importance of operating energy use.[5] The concept of energy density was unclear at the outset, primarily because of a lack of an industry-wide consensus about how density should be both measured and used to drive financial decisions. In addition, reliance on finance and capacity-based site management, varied disciplinary customs, technological convergence of data center and cable site infrastructure, management of legacy cable plant sites, changing transmission technology and a growing demand for high transmission density further clouds the understanding of energy density.[6,7] Often, data center and plant engineers naturally rely on capacity measurement to drive decision making, yet financial managers tend to focus on investment returns for site expenditures. However, neither of these measurements sufficiently incorporates other important factors, like energy expenditures, improvements to operating energy use during periods of growth, climate, co-location, and electrical system power-factor.

Because of the rise in large, centralized data centers, many operational practices are emerging that address a site's energy use. The top five examples include i) improved airflow management, ii) power distribution optimization, iii) increased site operating temperature, iv) real-time monitoring of site cooling efficacy and v) optimization of site cooling systems to use natural sources of negative thermal energy. However, even these practices are not all well understood, and their effects can be counterintuitive at varying economies of scale.[8] Measuring a site's density and energy intensity can be an important tool to fully understand a site's energy use, and to help improve density, reduce a site's overall energy expenditure and provide for more sophisticated analysis.[9]

In addition to measuring site density, a useful approach to infrastructure assessment can be borrowed from the building science industry. The work is variously called integrative design, or design and decision making based on a multidisciplinary collaboration.[10] *Building science engineers* typically measure and benchmark sites based on Energy Intensity;[11] however, this metric is often and mistakenly dismissed as not applicable to the benchmarking of cable sites.[12] On the contrary, this type of measurement can be used to help build basic measurement capability. It also bridges the gap between non-plant facilities management and data center management, which often share the same financial decision mechanisms used by large companies. Although a site's energy intensity based on square footage alone is not accurate enough to make engineering decisions for cable plant sites, it nonetheless forms the basis of more relevant metrics that engineers and financial managers can use for decision making. This paper will take an integrative approach to examine several cable plant sites, drawing comparisons to data centers and showing that the energy intensity of a site is an essential metric for cable plant operators. Furthermore, energy intensity based on subscribers and data throughput are also powerful decision making tools that cable plant and financial managers can use to make useful engineering decisions. Statistical modeling is described to identify relationships between independent site variables and energy use, as well as predict site energy use. Finally, the chief issue that this paper begins to address is identification of the most cost effective method to meet the Energy 2020 goal of a 20% increase in density, in terms of what must be measured and how to meet this goal.

2. The MSO Energy Picture

As shown in the Energy 2020 Energy Pyramid, most of a cable operator's power consumption is spent on cable plant infrastructure.[13] In practice, cable companies have to manage tens of thousands, if not hundreds of thousands of utility accounts in order to accurately account for its energy use. Utility bill management and energy procurement have become essential parts of cable operations. In 2014, Comcast spent over $320M on utilities, through approximately 250,000 different utility accounts.[14] Of the 250,000 accounts, approximately 90% were comprised of outside plant power supply accounts, which consume up to 70% of the network's total electricity expenditures.[15] Many of these power supplies are not metered and present a significant opportunity for savings, by either closing unmetered/unused accounts or negotiating new or more accurate billing rates.[16] However, as analog signals transmitted over a an HFC network sunset, and new network technologies like the Converged Cable Access Platform (CCAP) emerge, investment in density solutions that include outside plant power supplies may not be the best strategy --even though they represent the vast majority of utility accounts.[17] For example, a large scale network design upgrade or improvement including state of the art technology such as remote media access control (MAC) via CCAP, in markets with many distributed hubs (like the Bay Area of San Francisco, or Detroit, as two of many such examples) may prove to have the largest impact on energy use. That's because sites can be consolidated and or collapsed, thus eliminating the cooling aspect of energy consumption, while also providing higher bandwidth capacity.[18] In addition, historical energy measurements taken at multiple hub sites indicate that several relatively low cost site improvements can significantly and permanently reduce cooling costs, through the implementation of variable speed compressor drives, variable speed supply air fans, and free cooling economizers. Lastly, it is well known in the data center field that mechanically air cooled equipment is sub-optimal from an energy use perspective, and that significant gains in energy density can be achieved by implementing liquid cooled equipment.[19]

3. Density as it pertains to National Data Centers

Because large cable hub sites increasingly resemble data centers, the energy density in a data center is an important topic.[20] Most hub sites in a cable operator's infrastructure can be considered legacy buildings. By contrast, many data centers are newly built and managed using state of the art power and cooling equipment, often employing best practices that mitigate expensive cooling problems. For instance, Comcast owns, operates and leases a diverse portfolio of data centers across the country. Specifically, its data center footprint is comprised of more than 64,000 operating system (OS) instances running in 150 locations.[21] Over 80% of these instances are located in national data centers (NDCs) managed by the Deployment Engineering Division. As Figure 1 indicates, there is tremendous growth in the number of operating systems located in NDCs -- which grew from 4,255 in 2010 to 55,366 in 2015. It follows that energy use is also on the rise, growing from 8.5 million kWh/month in 2011, to 11.8 million kWh/month in March 2015. One of the chief operating principles of data center engineering, beside capacity management, is the systematic management of under-utilized device capacity, which is the same thing as increased energy density.

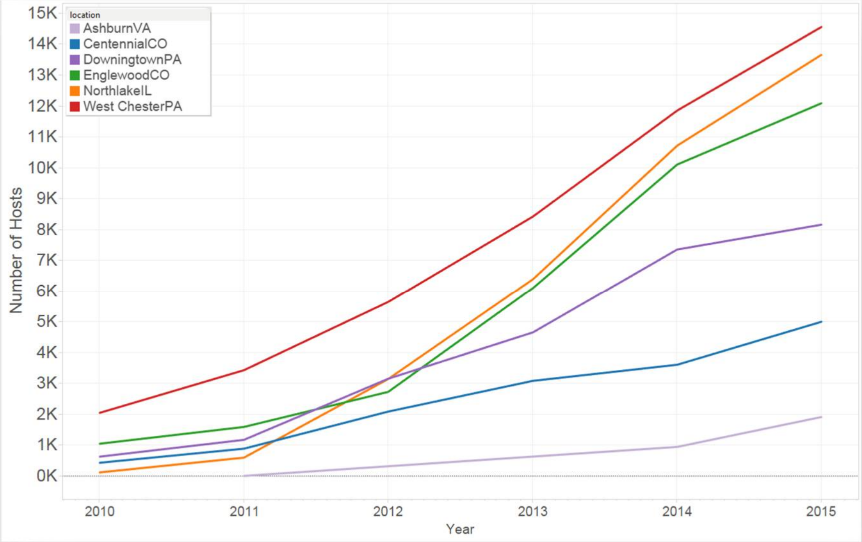

Figure 1 – In 2015 over 80% of data center devices reside in these 6 data centers. Source: The ITRC-database, an in-house catalog system.

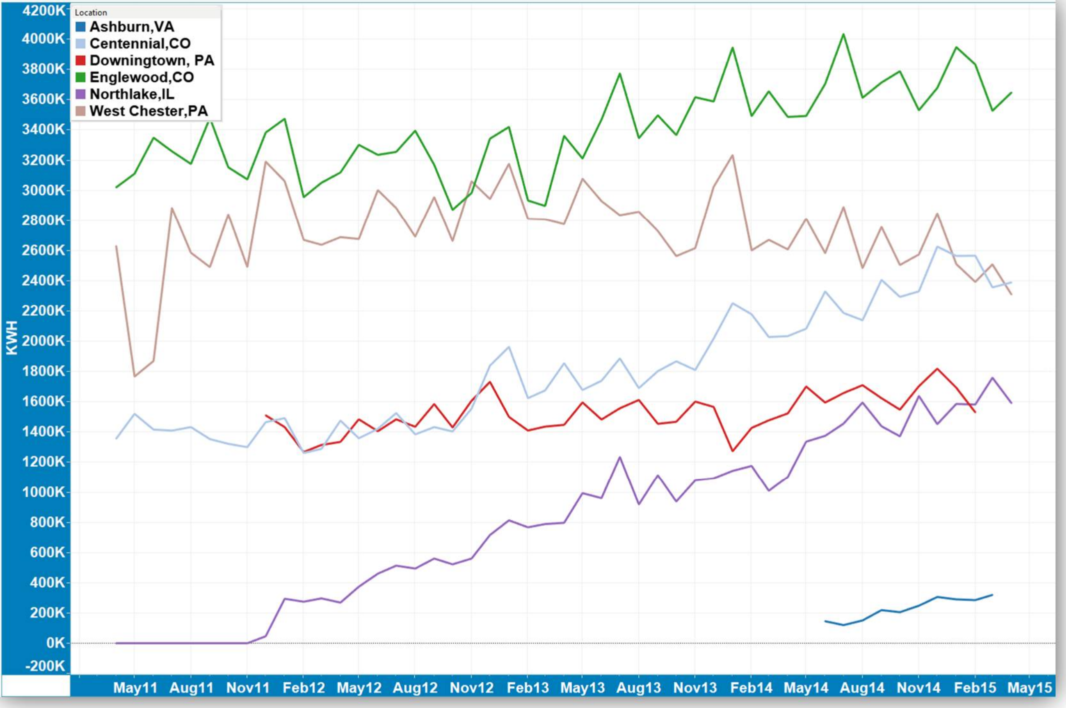

Figure 2 - Data Center Energy Consumption Increased 38% over 4 years. Source: Electric utility electronic data interchange feeds, paper bills.

Neither of these figures fully encapsulates an operational benchmark. To obtain a complete snapshot, it is important to account for a myriad of factors including density of devices, energy use per relevant unit and location together. For example, Figures 3, and 4 illustrate a recent method relating to location that Comcast employed to move equipment and start new data centers in more northern climates (zones 5 -7), to take advantage of free-air cooling and milder summer climates.[22]

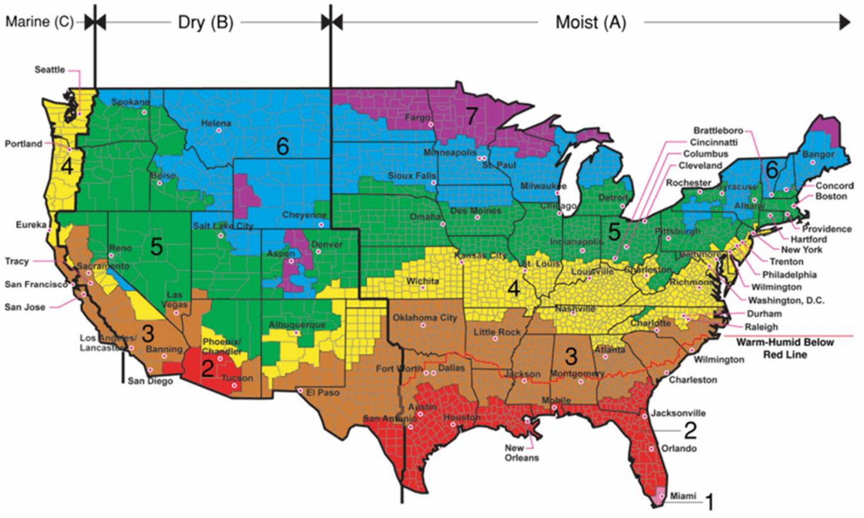

All of Alaska in Zone 7 except for the following Boroughs in Zone 8: Bethel, Dellingham, Fairbanks, N. Star, Nome North Slope, Northwest Arctic, Southeast Fairbanks, Wade Hampton, and Yukon-Koyukuk

Zone 1 includes: Hawaii, Guam, Puerto Rico, and the Virgin Islands

Figure 3 – United State Climate Zone Map, International Energy Conservation Code (IECC) climate regions.[23]

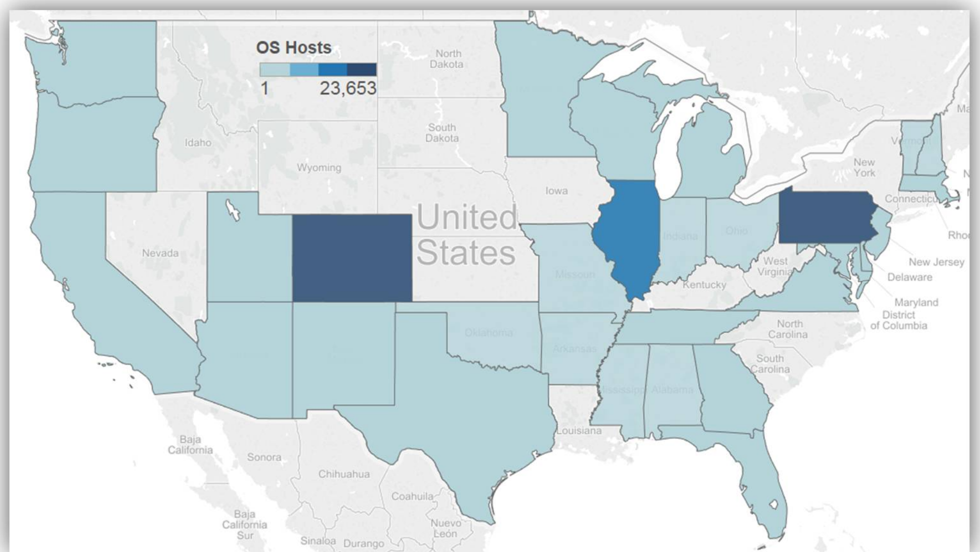

Figure 4 - 2015 Geographic Distribution of devices. Source: The ITRC-database, an in-house catalog system.

In addition to locating data centers in cooler climates, another longer term technique virtualizes legacy systems, in order to reduce energy consumption on a per unit basis by utilizing unused capacity. Figure 5 shows the density of energy use per operating system.[24] This figure along with Table 1 demonstrate the significant impact that device-level consolidation can have on energy use intensity. Table 1 shows the percentage of virtualized hosts at each site. Note that the data centers with the highest amount of virtualized hosts typically are the most energy dense and have achieved the highest per unit reduction in energy use. That's because a virtualized operating system host takes advantage of under-utilized computing resources.

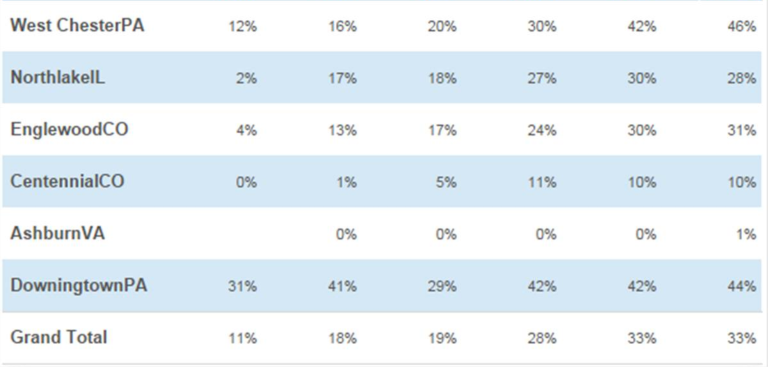

	2010	2011	2012	2013	2014	2015
West ChesterPA	12%	16%	20%	30%	42%	46%
NorthlakeIL	2%	17%	18%	27%	30%	28%
EnglewoodCO	4%	13%	17%	24%	30%	31%
CentennialCO	0%	1%	5%	11%	10%	10%
AshburnVA		0%	0%	0%	0%	1%
DowningtownPA	31%	41%	29%	42%	42%	44%
Grand Total	11%	18%	19%	28%	33%	33%

Table 1 – Percentage of the number of operating systems hosted at each data center site that are virtual. Source: The ITRC-database, an in-house catalog system.

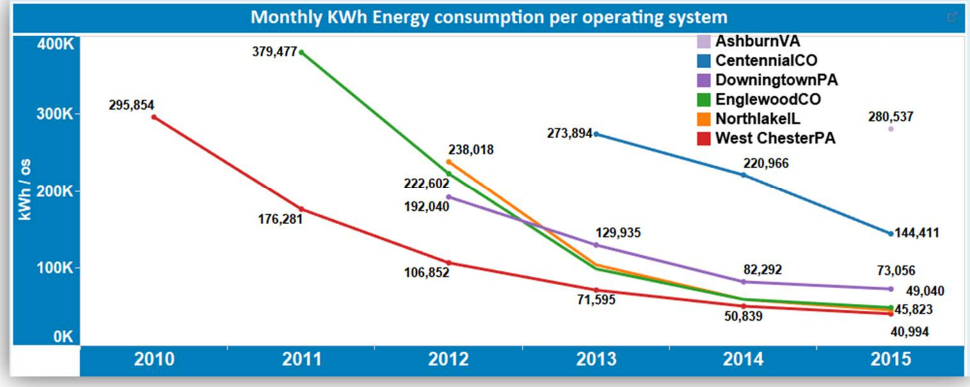

Figure 5 - Historical Energy Density of National Data Centers.

4. Density as it pertains to hubs and headends

4.1. Introduction

In total, Comcast operates approximately 2.6 million square feet of critical facilities in about 1,800 locations. These locations contain over 363,000 devices. Each device has a manufacturer specified number of watts that amount to about 42.8 "nameplate" watts per square foot.[25] About half of the critical facility footprint is comprised of sites that occupy less than 4,000 square feet. These facilities possess different economies of scale than data centers and overall are more diverse in terms of power system configuration, co-location, square footage, cooling system and types of information technology (IT) equipment. Table 2 shows that the majority of critical infrastructure is located in hub, headend, master headend and optical transport sites, which tend to house similar IT equipment. However, several aspects vary between sites. For example, hub sites are routinely standalone facilities or are co-located with technical operation facilities, network operation centers or divisional headquarters, in addition to a myriad of other combinations. Nonetheless, as with data centers, equipment density at these sites is an important aspect of energy density because the equipment landscape is rapidly changing. Historical measurements indicate the that typical site upgrades at hub sites improve density year over year-- by single digit single multipliers as well as orders of magnitude gains for deployments that include new technology like CCAP.[26]

Critical Facilities	Sqft.	% Total
Other	49,315	1.8%
Data Center	140,915	5.3%
Disaster Recovery Site	26,620	1.0%
HeadEnd	982,408	36.6%
Hub	919,581	34.3%
Master Head End	256,449	9.6%
Optical Transport	211,858	7.9%
Satellite Farm	23,453	0.9%
Tower	72,332	2.7%
Total Stand Alone	947,849	35.3%
Total Critical Owned	518,812	19.3%
Total Critical	2,682,930	11.0%
Total All Facilities	24,440,326	

Table 2 - Square Footage Breakdown of Critical Facilities

4.2. Measurement Methodology

In data centers, it is essential to measure density based on operating system hosts or a similar formulation of size of the data center; however, hub site configuration is more complex and therefore should include other site aspects including square footage, number of connected subscribers, and aggregate data throughput.[27] In the above data center density example, operating system instances coupled with the creation date for each instance proved to be an essential measurement. However, other measurements exist that may be just as demonstrative.[28] Ultimately, many different metrics are useful provided that they are measured continuously and are available for turnkey reporting. For hubs and headends, the simplest measurement to use is square footage as a proxy for work, because generally speaking larger hubs and headends support larger numbers of subscribers. However, more accurate energy measurement at these sites should directly account for historical changes in subscribers -- that way, if multiple hub sites are collapsed into one, or vice versa, data collection about energy use remains relevant. Additionally, if subscriber count is used as a metric, the energy density can also be used to demonstrate the greenhouse gas emissions per subscriber, and directly link investments in a location to subscriber lines of business. Lastly, square footage and subscriber density are very good measurements for the typically flat energy growth that represents the vast majority of cable sites. Using only these metrics, typical IT equipment deployments and capital improvements to building envelopes and to cooling systems can show a significant improvement to energy density.[29] However, as new (e.g. cloud-based) services emerge that demand higher energy use, and the number of subscribers does not change significantly, measurement of *data throughput* becomes essential to show improvements in site energy density.

4.3. Measurement Implementation

For Comcast, relevant information related to energy is contained in many disparate data sources, including: Electric company bills and databases, electronic data interchange (EDI) feeds, accounts payable databases, financial data warehouses, smart meters, energy probe systems, building management systems, automatic transfer switches, HVAC controls, facility databases, real estate databases, network monitoring software, cable plant monitoring databases, and critical facilities monitoring software. The key is identifying the systems that are easily integrated using application programming interface (API) components that enable extraction of useful data. This traditionally existed as an obstacle, because many software as a service (SaaS) platforms cannot easily be plumbed for API access because of corporate security and firewall constraints and SaaS vendor lack of motivation to interoperate. In many cases, utilizing file transfer protocol (FTP) "get" commands in a periodic shell script can alleviate this issue. Nonetheless, identification of a corporate internal repository or a closely integrated, centralized repository for this data is an important requirement for ease of reporting and analysis. A sample architectural diagram is shown in Figure 6, where a cross-divisional capacity database (CDCD) exists as a centralized repository that enables both operational and reporting capability. Operational components include a site scorecard for each energy-consuming location that is integral to maintaining facility power systems. Reporting capability includes structured query language (SQL)-based database reporting, as well as web service based integration with a diverse set of energy use data tools.

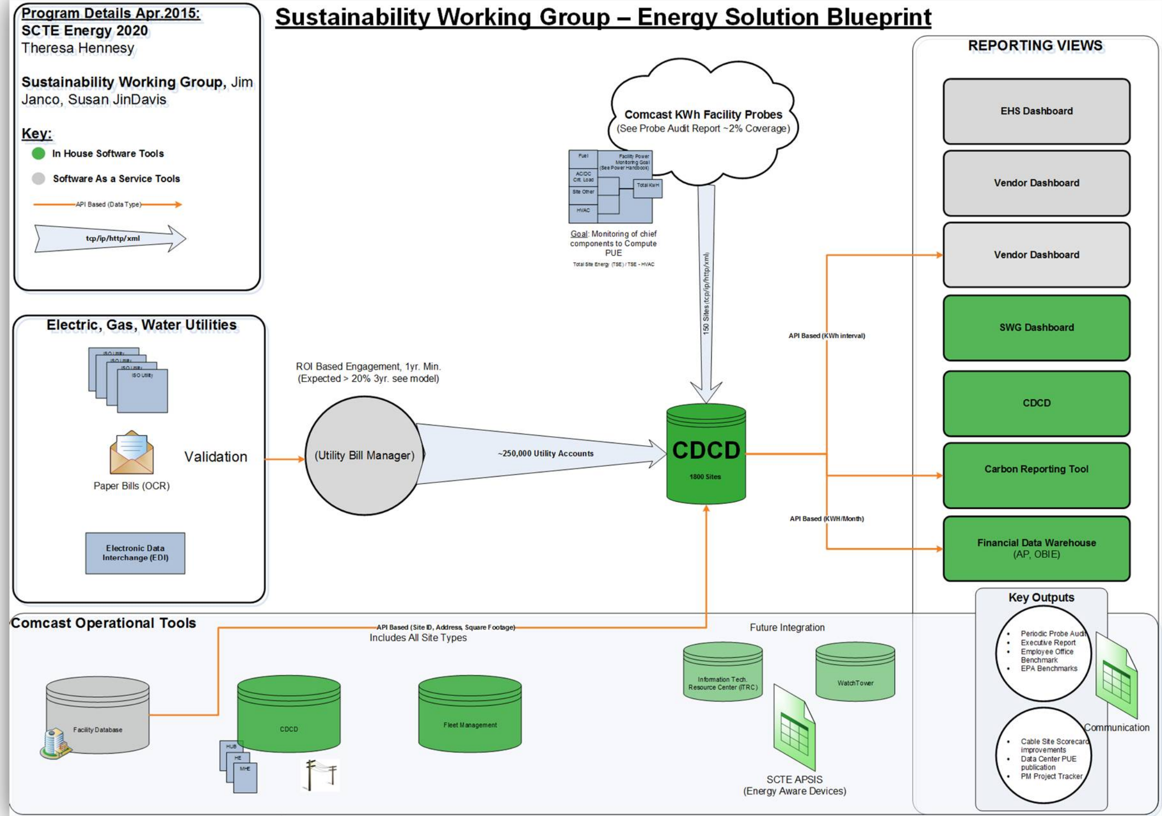

Figure 6 – Software Architectural Roadmap for Comcast's Energy Management System

4.4. Foxboro's new services

4.4.1. Site Characteristics

Foxboro hub supports about 20,000 subscribers and exemplifies a significant shift in technology deployment, such that typical density measurements will not suffice.[30] Also, and as shown in Table 3, the HVAC plant at Foxboro apparently grew incrementally, although the site does contain some HVAC economization. This site is also notable because of its partnership with the local electric utility provider. For example, the electric provider at this site contributed over 60% of the HVAC equipment costs. Electric companies collect fees from customers to make reductions in electrical grid demand and many of them have programs and partnerships to spend these fees to make these improvements in partnership with large consumers of electricity. This partnership impacted the cost of facility upgrades by reducing Comcast's capital expenditure by about $4.50 per subscriber. Lastly, shown in Figure 7, Foxboro also took partial advantage of a white roof surface, although it is as yet unclear how this impacts energy use at this site.

	Foxboro	**Noblesville**	**Lyndon**	**San Jose**	**Sacramento**	**Berlin**
Square Footage	11,462	1,880	2,832	13,620	12,300	8,315
Service Amps	1,600	800	1,200	2,000	2,000	800
Cooling Tons	4 x 30 9 x 20 1 x 20 5 x 20 2 x 10	3 x 20	2 x 5 1 x 3 1 x 5 3 x 15 crac 1 x 20	1 x 10 5 x 20 5 x 22	6 x 20 1 x 30	3 x 30 2 x 3 1 x 20
Economization	Partial	No	No	Unknown	Unknown	No
Racks	198	35	103	273	346	142
Subscribers	21,858	52,164	81,214	231,924	300,788	32,147
Climate Zone	5	5	5	3	3	5

Table 3 – Critical site characteristics. Source: In-house site catalog (CDCD)

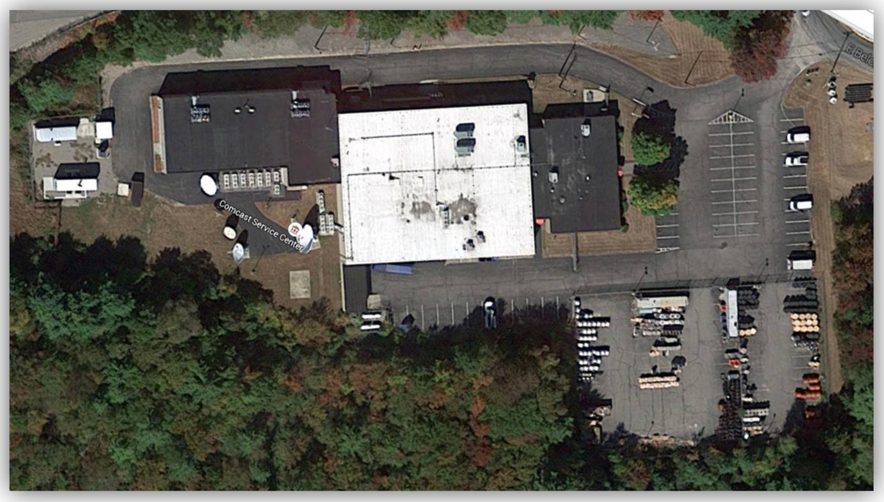

Figure 7 – Foxboro MA, Aerial photo showing white roof material.

Foxboro, MA	
HVAC Capital Investment	$2.29 / sub
Utility Incentive	$4.57 / sub
Subscribers	21,858:21,858
kWh/sub month before:after	11.4:18.3

Table 4 – Energy density per subscriber-month before and after service upgrades

4.4.2. Launch of cloud DVR services impact on density

Figure 8 is the most significant historical energy use curve discovered in this analysis and a chief reason that this study began with a closer look at data center density. Because Foxboro is now operating a new, cloud-based digital video recorder (DVR) infrastructure, it is stepping towards an all Internet Protocol (IP)-based service mix that's housed in a data center located closer to subscribers. It also demonstrates that the density metrics of kWh per Subscriber and kWh per square foot are insufficient measurements for this type of infrastructure change. For example, both square footage and number of subscribers remained relatively unchanged in the denominator, but the numerator nearly doubled. Thus, measurement using those metrics illuminated a decrease in energy density -- despite doing more work.

This is where the metric that is data throughput -- and specifically, kWh per data throughput -- comes into play. To fully show improvements in density, kWh per site throughput capacity or actual cumulative throughput is required-- because as the energy use doubled, both the site capacity measured in Megabits per second (Mbps) and site cumulative throughput measured in Terabytes increased by an order of magnitude.[31] This resulted in a significant improvement to density -- well beyond the energy 2020 density goal at this site.

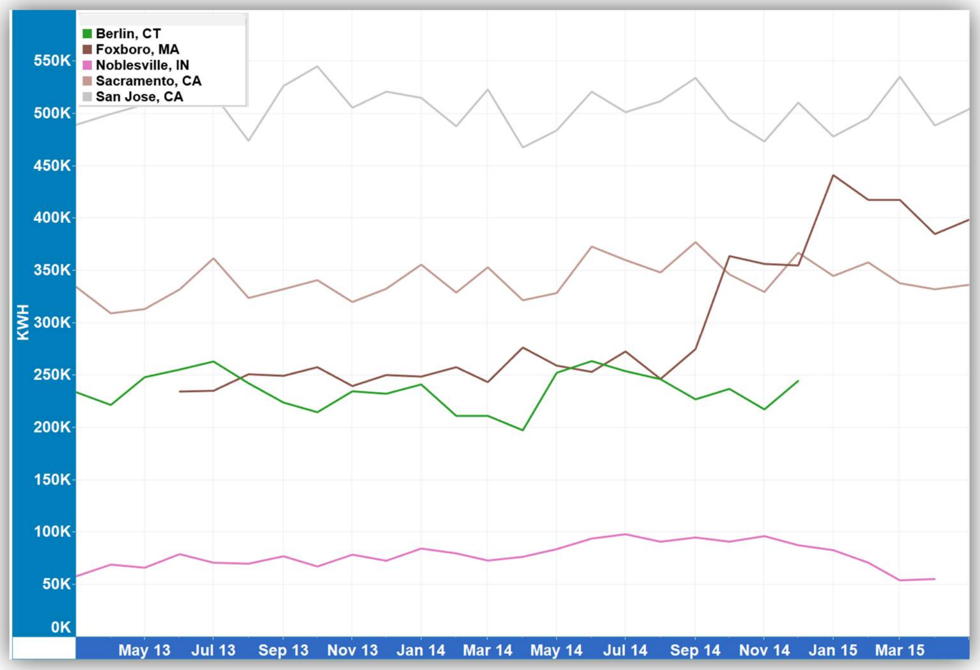

Figure 8 – Foxboro MA, kWh per month shown increasing usage, shown in comparison to other medium and large sized hubs and head ends. Source: Electric utility EDI feeds, paper bills.

4.5. A Representative Hub Consolidation Effort

4.5.1. Site Characteristics

The Noblesville, Indiana hub serves approximately 50,000 subscribers and is classified as an optical transfer network.[32,33] Because of a combination of site power constraints, increased cooling demands and service capacity issues, the site was expanded from 20 to 35 racks, which required the new, adjacent building shown in Figure 9, as well as upgraded power and cooling to support the new racks.[34] The new site characteristics are shown in Table 3. In addition, a small hub supporting approximately 5,000 subscribers was collapsed into the new Noblesville site.

Figure 9 – Aerial photo showing physical site configuration.

4.5.2. Site Density

The Noblesville site collapse exemplifies the use of subscriber-based density measurement to sufficiently show that typical improvements to critical infrastructure can impact site density. Specifically, improvements including heating ventilation and air conditioning (HVAC) equipment and new construction methods, coupled with a modest ~10% increase in subscribers, garnered a 31% density improvement, shown in Table 5. For this density improvement it was not necessary to measure data throughput, although the implementation of two CCAP-capable devices in Noblesville also increases the theoretical data throughput and its capability to provide higher levels of service to subscribers.[35] As expected and demonstrated in Figure 10, the site density improvements were not fully realized until twelve months after the initial turn up of the new facility, beginning February 2014. Additional improvements to density are also expected when the CCAP capable devices begin hosting video services. When these service are moved to CCAP the legacy video equipment can be removed from the site, resulting in lower electricity usage and increased rack space. Here, it was essential to use historical power usage measurements, in order to show the breadth of the improvements.

Site	Springlake, IN	Noblesville, IN
Cost to Collapse	$95.85 / sub	
Subscribers before:after	5000:0	47,164:52,164
kWh/sub month	1.25:0	1.52 : 1.05 (31%)

Table 5 – Energy density per subscriber month before and after consolidation

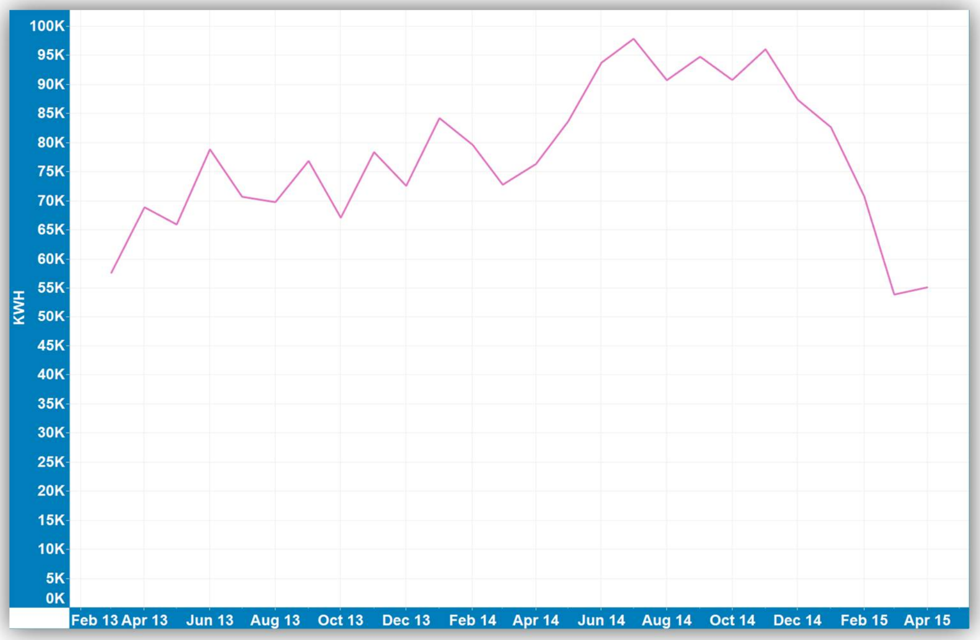

Figure 10 – Noblesville, kilowatt-hour (kWh) per month. Source: Utility feeds, paper bills and site power probes. In February of 2014, new equiment was installed in the new building and in February of 2015 the old euquipment was either moved to the new site or de-provisioned.

4.6. Lyndon regional consolidation to virtual-hubs

4.6.1. Site Characteristics

Comcast’s Lyndon, Michigan hub supports about 80,000 subscribers is another example of medium sized hub that’s classified as an optical transfer network.[36] Due to antiquated system components and the operational costs of supporting facilities across a dispersed geographic region, the Engineering Team decided on a virtual hub strategy. Specifically, virtual hubs were deployed in the infrastructure, and cable television signal equipment was consolidated into a single master site in Lyndon.

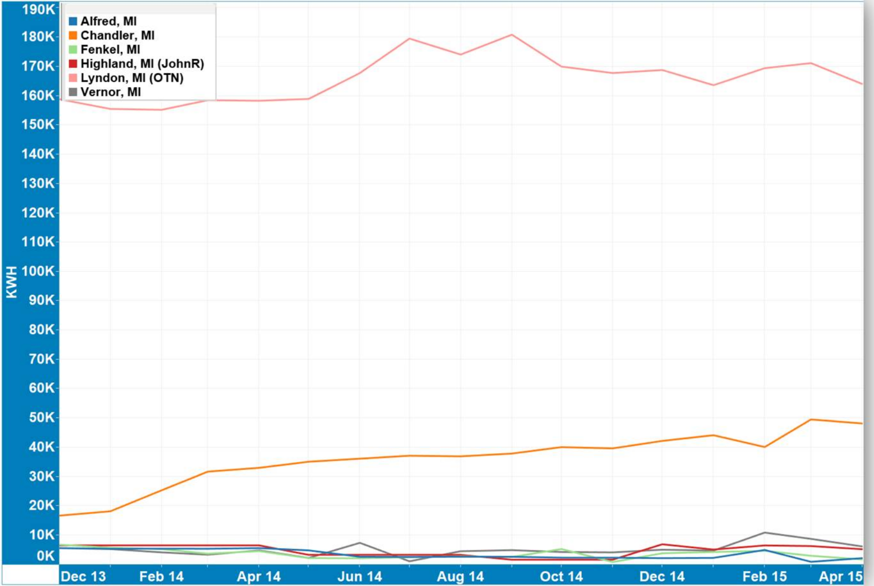

Figure 11 – Detroit area, kWh per month, note the low consumers at the bottom modestly reduced energy consumption and the Master site increased beginning in May 2014. Source: EDI feeds, paper bills.

4.6.2. Regional Density

The site characteristics are shown in Table 3. Notably, the cooling systems at this site appear to have grown incrementally as well, which may represent an opportunity to install HVAC equipment containing economization methods, for additional improvements to density. Even without this improvement, as shown in Table 6, subscriber energy density improved from approximately 8 to 2 kWh per sub-month at the Lyndon site. However, the energy use at the consolidated sites was only reduced nominally, (see Figure 11.) in contrast with the Indiana site collapse, and as a result these network improvements do not fully take advantage of potential regional energy density. That's because the virtual hubs, along with commercial account equipment, were not eliminated from the consolidated sites. Furthermore, compared to the Indiana hub consolidation effort, this site still uses almost double the amount of energy per subscriber.[37]

Lyndon, MI	
Cost to Consolidate	$40.48 / sub
Subscribers before:after	20,145:81,214
kWh/sub month	7.75:2.02

Table 6 – Energy density per subscriber-month before and after consolidation

4.7. Large scale regional CCAP deployments

4.7.1. CCAP site characteristics

As of June of 2015, Comcast deployed approximately 1,300 CCAP-capable devices. Thirty-one CCAP devices are located in the Sacramento and San Jose hub facilities.[38] Sacramento and San Jose are classified as a master headend and headend, respectively, and as shown in Table 3 and Figure 12 are relatively large cable hub sites serving over 500,000 subscribers and each use over 300kWh per month. An important feature of these CCAP deployments is revealed in Figure 12. Where, despite the increased wattage of the CCAP racks, the site power consumption increased by 2 and 4% respectively.

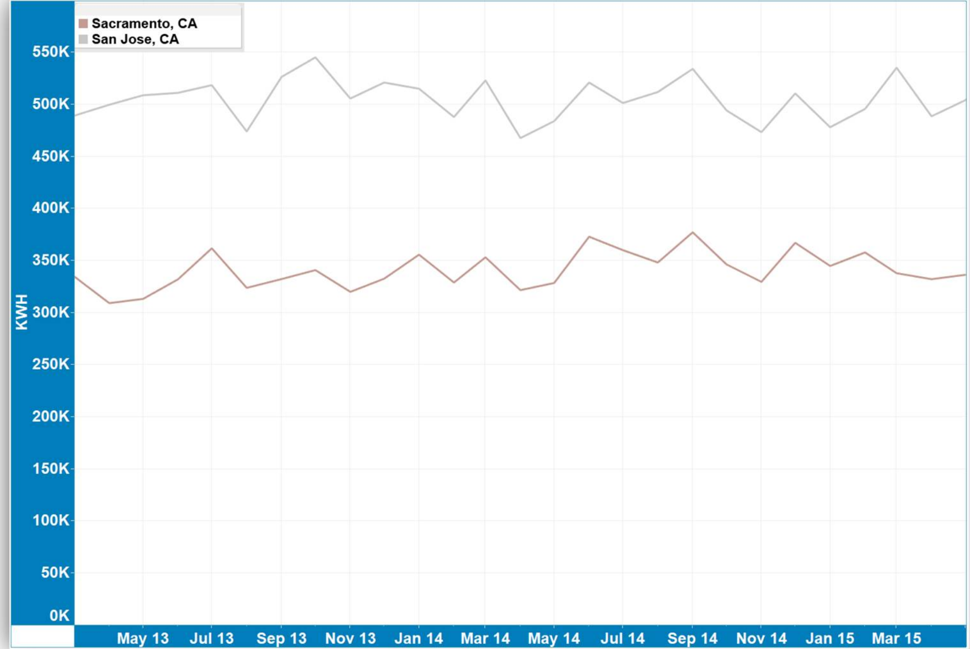

Figure 12 – Sacramento, San Jose, CA, kWh per month, during this period 31 CCAP upgrades were installed, note a 4% and 2% increases respectively. Source: Utility bills.

4.7.2. CCAP impact on density

Deployment of CCAP-capable devices carries numerous benefits that include energy density. Figure 13 shows a tremendous 85% increase in energy density, based on theoretical limits of CCAP devices. Ideally, this should be measured using actual data throughput, however, this at least demonstrates the theoretical limits of the CCAP deployments as it pertains to the SCTE Energy 2020 goals. Notably, in large hub sites like Sacramento and San Jose, economies of scale begin to reach similar proportions to a data center. In terms of energy use and the cost per subscriber for state-of-the-art IT equipment, the equipment only cost is less than one-third that of some smaller site upgrades. (see Table 7.)

Figure 13 – Sacramento, San Jose kWh per throughput capacity (Mbps) Source: Utility bills and internal equipment site catalog (extrapolations based on planned capacity work)

Northern California	
Capital Investment	$30.04 / sub
Subscribers before:after	532,712: 532,712
kWh/Mbps-month	24.8:3.6 (85.5%)

Table 7 – Energy density per subscriber-month before and after service upgrades

4.8. Berlin policy driven density

Berlin is a relatively small site serving about 32,000 subscribers, classified as a master headend, and located in Connecticut, with typical cable infrastructure.[39] The Comcast Engineering Team chose this site to host Comcast's first off-grid Natural Gas Fuel Cell for three reasons: To save on electricity costs, reduce greenhouse gas (GHG) emissions, and trial an off-grid technique. Ultimately, the natural gas fuel cell system saves approximately $25,000 per month in electricity costs (see Figure 15). However, even with utility incentives, this improvement had significant per subscriber costs (see Table 8.). On the other hand, because of the new system configuration, site power is now n+2 redundant and more reliable than the typical n+1 systems, using the electric grid only as a secondary power source. Although costs per subscriber were 50% higher than Noblesville and similar hub consolidation efforts, the successful implementation of this system demonstrates the Comcast Corporate Social Responsibility and Sustainability Policy by reducing greenhouse gas emissions by 6.4%. Additional site improvements, like HVAC economization and white roofing material (see Figure 14.), are still available at this site to further its energy optimization and GHG reductions.

Berlin, CT	
Capital Investment	$124.43 / sub
Utility Incentive	$1.18 / sub
Subscribers before:after	32,147: 32,147
CO2e / sub-year	2.8e-2:2.62e-2 (6.4%)

Table 8 – Green House Gas per subscriber-year before and after site improvements measured in metric tons[40]

Figure 14 – Berlin, CT, Aerial photo. The black rectangle indicates the Fuel Cell Location.

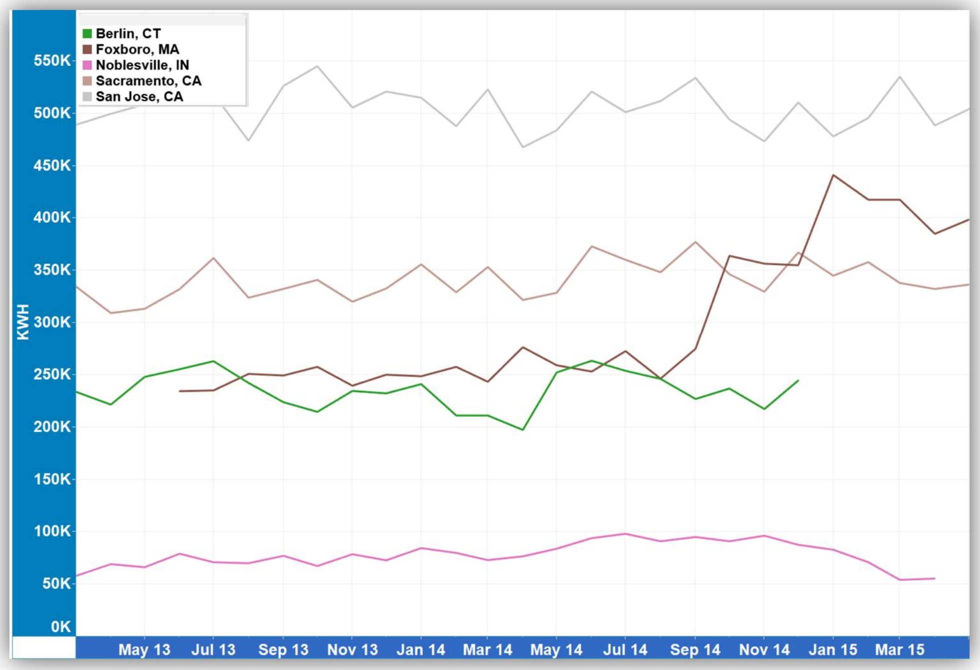

Figure 15 – Berlin, CT, kWh per month obtained through paper bills shown in comparison to other medium and large sized hubs amd head ends. Note that this site went off-grid in 2015.

5. Improving density using regression analysis

We use simple statistical linear regression to probabilistically model some characteristic data. As a preliminary step, we perform some rudimentary data preparation, such as validating the facility site IDs, selecting only standalone cable sites, ensuring that critical areas and 12-month energy usages are valid, confirming that critical areas and total areas have a proper relationship, and computing kilowatt hour averages.

Restricting ourselves to only one predictor variable, we use least squares / maximum likelihood estimation to fit a basic regression model with additive Gaussian noise. In this simple model,

$$Y = \beta_0 + \beta_1 * X + \varepsilon,$$

the intercept term, β_0, represents the expected value of the response Y when the predictor X is 0. The slope term, β_1, represents the expected change in the response Y for a unit change in the predictor X. We model such a relationship in Figure 16. Here, the response Y is the energy usage, summed over the last 12 months and measured in kilowatt hours. The predictor X is the critical area of the facility, measured in square feet.

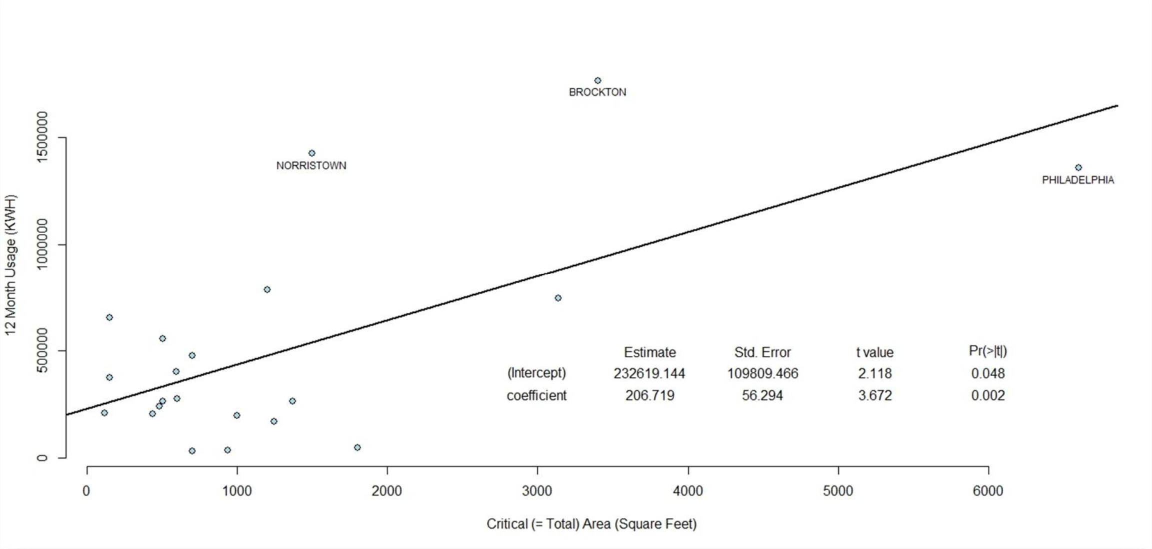

	Estimate	Std. Error	t value	Pr(>|t|)
(Intercept)	232619.144	109809.466	2.118	0.048
coefficient	206.719	56.294	3.672	0.002

Figure 16 - Simple preliminary model predicting energy use at stand alone facilities.

In this Figure, we also provide an “Estimate” for the Intercept (β_0) as well as the coefficient (β_1, the slope), the “Std. Error” (a measure of the statistical variability) for both, the “t value” (Estimate / Std. Error, used for hypothesis tests) for both, and finally “Pr(>|t|)” (a measure of statistical significance) for both. Using the given values, we can calculate confidence intervals for each of the β terms. For example, using

$$\text{Intercept}_{\text{Estimate}} \pm t_{0.975,\,df} * \text{Intercept}_{\text{Std. Error}}$$

with n–2 (=19) degrees of freedom, we can estimate with 95% confidence that a 1 square foot increase in critical area in a facility results in a 88.896 to 324.543 increase in kilowatt hours used over a 12 month period.

Typically, a linear regression that is as simple as this does not yield the best prediction results. However, the models that result are parsimonious and easily interpretable. Also, quantitative estimates of the predictions’ uncertainties are easily developed. These factors sometimes outweigh any interest in incremental increases in accuracy. We identify several data points in Figure 16 representing sites (Norristown, Brockton and Philadelphia) with large residuals, i.e., significant differences between the observed and the predicted outcomes.

We can also utilize this regression model for prediction. For example, using

$$Y = 232619.144 + 206.719 * X,$$

we would expect 1,059,496 kilowatt hours of energy usage over a 12 month period (= Y) in a facility having a critical area of 4,000 square feet (= X). Of course, predicted responses have standard errors and predicted and/or expected response intervals can be calculated.

Our future work will include a multivariable regression analysis, in which we will generalize the simple linear regression to incorporate more than one regressor for prediction. Linear regression modeling is a methodical, nontrivial and dynamic process – iteratively deciding what the appropriate variables for inclusion in the model are. The models resulting from these techniques are arguably among the most important in applied statistics and machine learning.

Conclusions and Recommendations

1. Conclusions

The Energy 2020 program set energy use density targets for the year 2020. These goals are achievable by continuing to implement solutions that focus on increased network capacity for data throughput. However, much of the energy in a cable plant is consumed by outside plant power supplies and are largely not impacted by current solutions. For the most part, the outside plant will remain unchanged until IP-based technologies are implemented, like remote CCAP. That means that the density targets at inside plant facilities must meet higher thresholds in order to meet those overarching goals. The good news is that significant increases in energy density can be made by deploying CCAP devices. Additional increases can be realized when the CCAP devices are fully utilized and legacy infrastructure is removed. Furthermore, untapped opportunities, like building improvements, will also increase energy density. Although this is changing, most cable plant facilities do not follow the top five emerging data center best practices, especially: Improved airflow management, real-time monitoring of site cooling efficacy, and optimization of site cooling systems to use natural sources of negative thermal energy. In addition, even when presented with the opportunity to implement the most efficient cooling solutions, the default position is usually to “value engineer” the capital expenditure and eliminate free cooling as an option. This must change! Engineers should be empowered to recommend changes to critical infrastructure that incorporates free cooling. They should also be empowered to deploy energy monitoring at automatic transfer switches and with HVAC equipment, so as to continuously monitor and measure cooling efficacy. Especially since electric companies are positioned to contribute to these infrastructure improvements. To achieve the Energy 2020 goals and corporate sustainability goals forevermore, cable operators must put in place centralized and continuous monitoring of all of its facilities that track energy use and costs as well as historical device counts, historical sub counts and cumulative historical network throughput.

2. Areas for Further Investigation

This paper would be more complete with a better understanding of the impacts of remote CCAP on the outside plant infrastructure energy use. For example, when the choice is made to use remote CCAP, will the energy-impacting infrastructure be placed in the outside plant, or will it remain inside traditional hubs?[41] Additionally, the apparent benefits of virtualized devices at data centers indicate further research, so as to estimate their impacts on energy use at hubs. That includes network function virtualization as well as virtualization of customer premises equipment and new cloud DVR infrastructure. Further investigation is still needed for all sites that incrementally implement energy savings techniques, to ascertain the best returns on investment. Lastly, the impacts of liquid cooled equipment under in-situ conditions at distributed cable hub infrastructure sites should be investigated.

Abbreviations

API	Application Programming Interface
CCAP	Converged Cable Access Platform
CDCD	Cross Divisional Capacity Database
CT	Connecticut
DVR	Digital Video Recorder
EDI	Electronic Data Interchange
FTP	File Transfer Protocol
GHG	Greenhouse Gas
HVAC	Heating Ventilation Air Conditioning
ID	Identification
IECC	International Energy Conservation Code
IP	Internet Protocol
ITRC	Interstate Technology and Regulatory Council
kWh	Kilowatt Hour
MA	Massachusetts
MAC	Media Access Control
Mbps	Megabits per second
MSO	Multiple System Operator
NDC	National Data Center
OS	Operating System
SaaS	Software as a Service
SCTE	Society of Cable Telecommunications Engineers
SQL	Structured Query Language

Appendix A

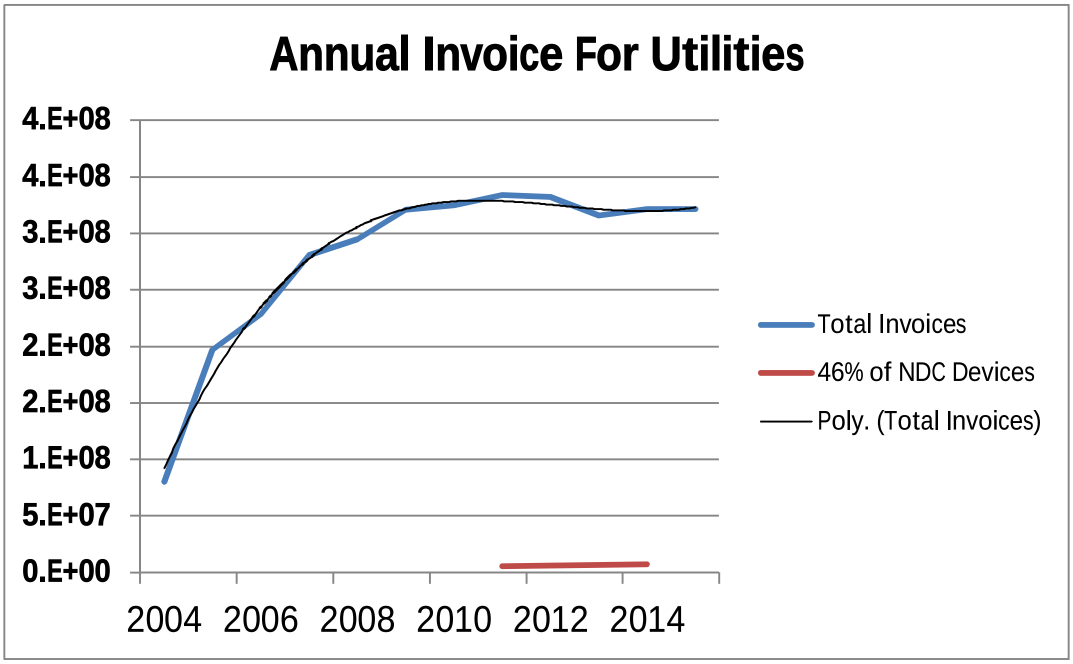

Figure A1 - Total invoices for utilities in dollars, measured from accounts payable, compared to overall electricty expenditure at national data centers.

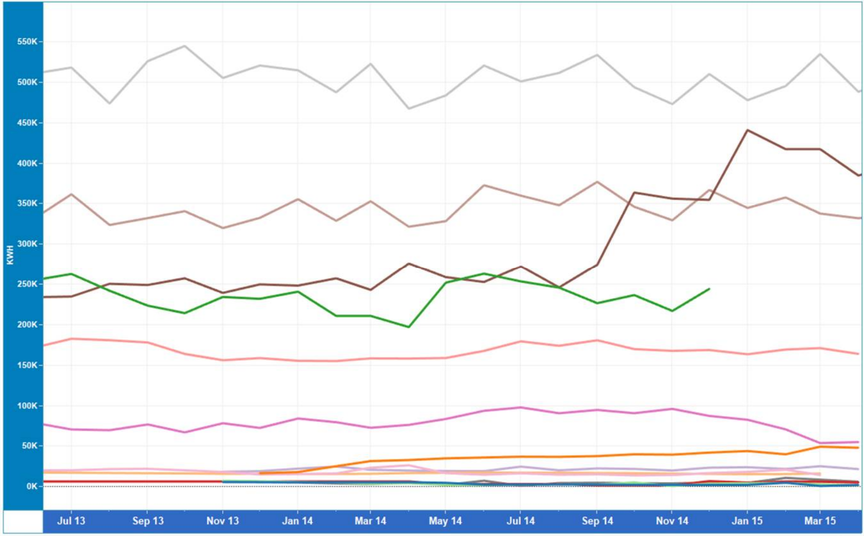

Figure A2 - Showing representative cable site energy usage. Source: EDI and paper bills.

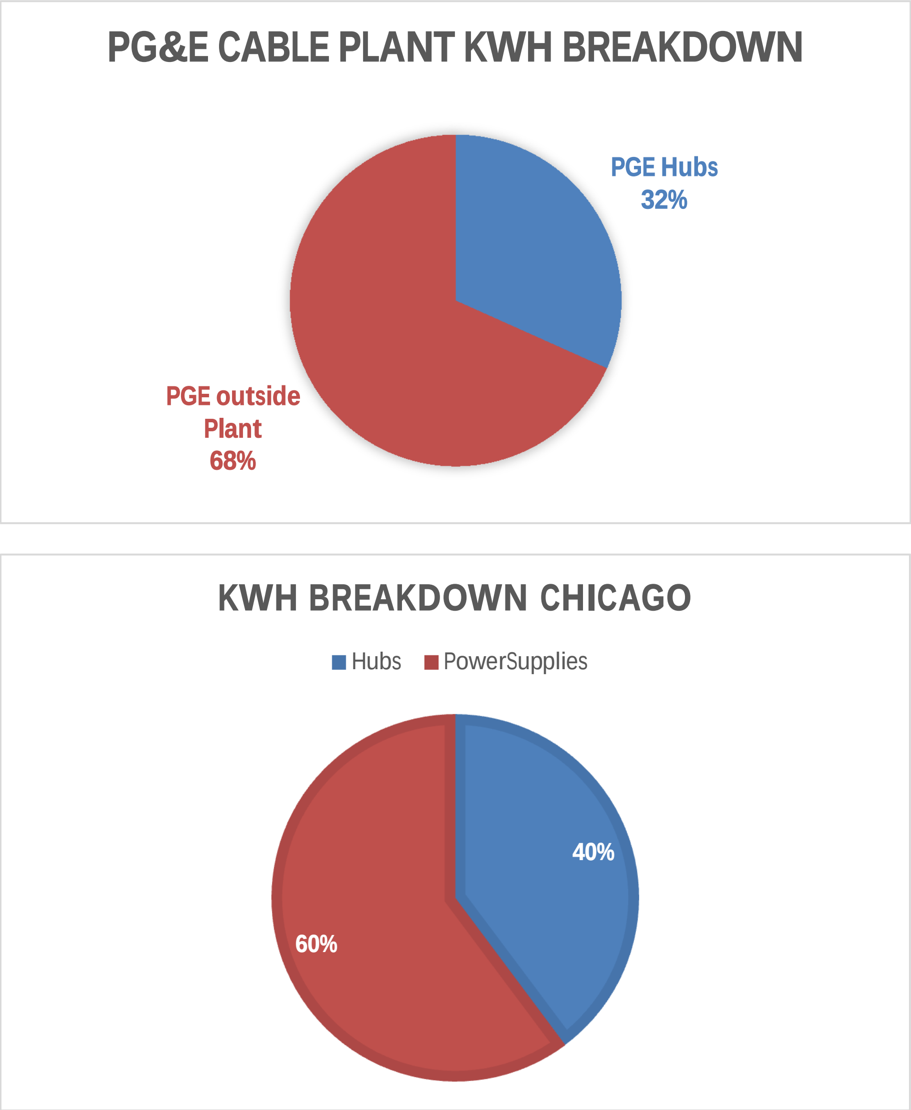

Figure A3 – Showing actual kWh breakdown for two large energy providers percent allocation of outside and inside plant. Source: EDI feeds.

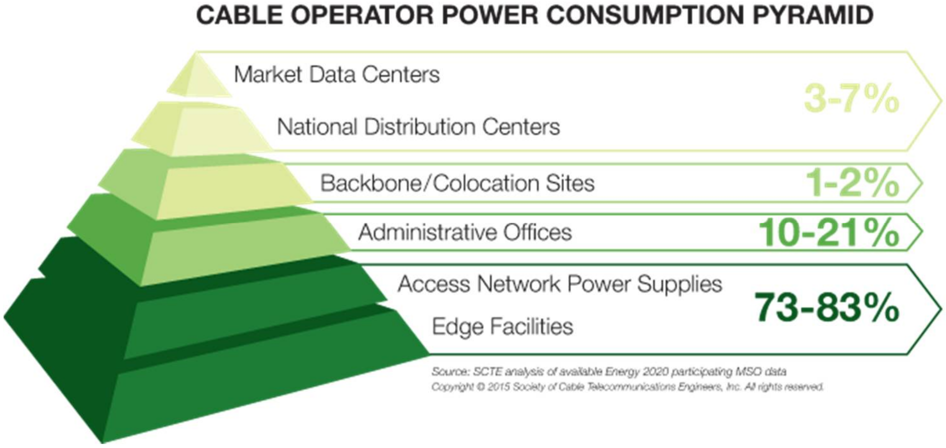

Source: SCTE analysis of available Energy 2020 participating MSO data
Copyright © 2015 Society of Cable Telecommunications Engineers, Inc. All rights reserved.

Figure A4 – Showing the breakdown of typical energy use at a cable operator, provided by SCTE

Quantity	Model	Watts
7	C4 CMTS	19600
28	Laser Link 3 Chassis	18200
2	E6000 CMTS 8000W	16000
2	ASR 9010	12560
2	7609 Router 6000W	12000
11	Proliant DL360p Gen 8	8250
13	NSG 9000 553W	7189
3	HCU-1500	5400
4	ONS M12	5376
4	720XD Server	4400
16	SEM v8	3680
16	CH3000N Chassis	3456
3	FlashWave 7500	2790
11	DM-6400 Cherrypicker	2200
2	Laser link 2 Chassis	1920
1	MX-240	1680
3	Prisma II Chassis	1449
16	RPD 2000	1280

Table A1 – Equipment Manifest for Noblesville, IN, >1000 watt total, measured using internal site catalog database

Quantity	Model	Watts
36	NSG 9000 665W	23940
8	C4 CMTS	22400
73	CH3000N Chassis	15768
159	RPD 2000	12720
2	7609 Router - 6k PS	12000
1	MX-960 XD	9348
4	HCU-1500	7200
5	ONS M12	6720
8	NSG 9000 40G 780W	6240
1	MX-960	5100
5	EX 4550	3250
3	Optera Metro 5200	2880
6	GX2 Chassis	2310
1	4507 R {208v}	1444
9	ARPD 1000	1296
1	ASX 4000	1200
1	ASX-1000	1200
12	OM2000	1008

Table A2 – Equipment Manifest for Lyndon, MI, >1000 watt total, measured using internal site catalog database

Quantity	Model	Watts
365	S410	368650
52	UCS-5108	161408
68	Recorder Cluster Manager	88400
40	WOS-CSN(Web Cloud Storgae)	48000
25	WOS 7000	46625
4	7609 Router	24000
2	cBR8	23940
50	UCS C220 M3	22000
6	Nexus 6004	19800
16	M3-RCM-HA-AA-4H 146GB-DC-A-SRVR-SET	19200
23	DL 360 Gen8	17250
1	CRS 1	13900
2	ASR9000	12600
2	ASR 9010	12560
2	7609 Router - 6k PS	12000
6	S400 (M3-S400-72H1TB-6X)	11700
10	720XD Server	11000
12	CDE 250 Content Delivery Engine	10800
10	UCS 6296 UP	9500
8	NSS-APP-4H 146GB-DC-A-SRVR (APP Server)	6720
7	ONS M6	5880
5	NSG 9000 QAM Modulator Turbo	5880
11	DS2246	5676
2	OME6500 (60Amp)	5040
10	CDE 220 Content Delivery Engine	5000
2	WOS7000	4992
2	1830	4800
3	TRX-24000	4800
3	1696	4800
2	UBR10000	4800
6	UCS C240 M3 LFF	3900

8	Medius	3840
5	Proliant DL360p Gen 8	3750
8	NSG 9000 QAM Modulator	3600
3	PowerEdge 720XD	3300
4	UCS-6248 Fabric Switch	3000
1	7504E	2900
13	XG1 Cable Box	2652
1	6500 Transport Shelf	2520
12	ORX/OTX Housing HX1281EC	2400
8	4948e	2400
3	CDE250 Pump	2400
2	XMS FLEX	2400
1	OME 6500 14 Slot NTK503ADE5	2399
10	SEM v8	2300
88	VFA 750	2200
6	CAP-1000	2100
1	(JDSU) HCU-1500	1800
6	Apex 1000	1542
4	CHP Max 5000	1516
2	Flash 600(2400) ADX	1500
7	Proliant DL360-G7	1470
2	NMX Server	1464
2	Nexus 5548P and 5548UP	1460
1	4507 R-E Switch	1444
2	PTS22600	1440
4	2800 Series	1344
6	7050	1320
6	Server Model#SCSU11022T	1260
15	ARPD 1000	1102.5
3	2960G Switch	1080
18	TM402P	1080
9	Omnistar Chassis	1080
2	Arris CHP Max 5000 Shelf	1080
5	ASR9000V	1050
3	Cherrypicker CAP1000	1050

Table A3 – Equipment Manifest for Foxboro, MA, >1000 watt total, measured using internal site catalog database

Quantity	Model	Watts
11	E6000	88000
2	ASR 9922	69000
55	CDE460	55550
26	WOS 7000	48490
22	PowerEdge R720xd	48400
6	E6000 CMTS 8000W	48000
36	CDE 460	46080
6	ASR 9000	36520
162	CH3000N Chassis	34992
10	Nexus 6004	33000
10	UCS-5108	31040
9	1830 PSS-32 70 Amp	30240
40	Power Edge R610	28800
25	720XD Server	27500
53	CDE 220 Content Delivery Engine	26500
49	NSG 9000 Octal QAM Modulator	22050
5	7609 Router - 4k PS	20000
8	WOS7000	19968
4	Netra 1290	15752
15	UCS C220 M3 SFF	13950
18	PowerEdge 2950	13500
2	ASR 9010	12560
2	7609 Router - 6k PS	12000
26	LaserLink III	11125
12	CDE 250 Content Delivery Engine	10800
7	NET-NET 9200	10640
14	PowerEdge 1950	9380
31	T2000	9300
96	ARPD 1000	7056
16	Laser Link 3 Chassis	6275
27	NC-1500	6210
2	1830 PSS-32	5760
71	AT 1601 M	5680
26	ASR 9000v	5460
7	FAS2240 220 vac	5250
10	Cable Box-Dolby PX001ANM	5160
2	MGX 8800	5040
6	CDE 420 Content Delivery Engine	4980

4	WOS-CSN(Web Cloud Storage)	4800
8	xSeries 336	4680
9	HCU-1500	4320
13	DL360	4225
12	NET NET 4250	4200
7	Proliant DL365	4200
4	SBC 5200	4000
4	FlashWave 7500	3720
6	Proliant DL360-G5	3600
2	5500 NGX	3360
16	EX 4200	3040
7	Ex 4550	2891
2	4500 E-Series Switch	2800
6	Proliant DL380-G6	2760
14	EX4200	2660
4	Laser Link III Power Supply 254047	2600
13	NET NET 4500	2600
3	PowerEdge R710	2520
10	TIGPT1U	2500
9	Apex 1000	2313
2	FLASHWAVE 7500 SHU3	2304
2	Flashwave 7500S OLC	2280
2	Flashwave 7500S OADM	2280
18	Server Model#SCSU11021T	2268
3	DL 360 Gen8	2250
28	OM 2000 Out of Band Modulator	2240
6	Cherrypicker CAP1000	2100
9	SEM v8	2070
6	CableVista Chassis	2040
5	Sentry	2000
2	BMR 1200A	1850
4	Maxnet Chassis (active)	1728
5	2800 Series	1680
1	7510	1680
1	NET NET 9200	1536
2	UCS-6248 Fabric Switch	1500
2	Nexus 5548P and 5548UP	1460
5	WS-C4948E-AC	1375
2	Sunfire V440	1300
2	Nextra X4250	1300
2	UCS C240 M3 LFF	1300
2	DSNSA7-GE500SX	1200
2	Proliant DL385	1200
5	Maxnet 11 Chassis	1175
2	3845 Router	1110
5	Pro Stream 1000	1100
5	Proliant DL360-G7	1050
8	HTR2000	1024
1	Flashwave 7500S	1008
2	Sun Fire X2250	1000

Table A4 – Equipment Manifest for San Jose, CA >1000 watt total, measured using internal site catalog database

Quantity	Model	watts
13	E6000 CMTS 8000W	104000
2	MX 2020	75600
80	Omnistar GX2 Chassis	42400
73	NSG 9000-6G	40369
44	NSG 9000 40G 780W	34320
7	7609 Router - 4k PS	28000
3	ASR 9000	22680
15	NSN HiT7300	21300
3	7609 Router 6000W	18000
7	OME 6500-14	17640
2	T1600 Router	16700
2	E6000	16000
3	MX-960	15300
15	CDE 250 Content Delivery Engine	13500
163	ARPD 1000	11980.5
26	Switching Power Supply - SYS-AC-Q36191E	11856
61	EX 4200	11590
23	CHP Max 5000	10925
1	TX Matrix Plus	9600
39	SEM v8	8970
4	Netra 1280	8400
13	Servers 2	8112
11	Weather Channel / Weatherscan Unit	7920
39	DM 6400 Cherry Picker	7800

9	Servers 1	7560
20	CAP-1000	7000
82	RPD 2000	6560
26	TIGPT1U	6500
1	7750 Service Router	6480
24	Apex 1000	6168
8	N5 Video Server	5840
32	CH3000N Shelf	5616
25	Pro Stream 1000	5500
9	Intellistar 2 HD	5400
2	Edge 6500	5040
2	MGX 8880	5040
2	OME 6500 Switch	4798
6	Proliant DL380 G7	4500
34	DM 6400 Cherry Picker - DC	4406.4
4	R720XD	4400
6	PowerEdge 1950	4020
1	7609 Router 4K	4000
10	Continuum Modulator Chassis	4000
4	BMR 1200A	3700
17	CH3000N Chassis	3672
1	MEG-400	3600
13	Divicom Ion	3380
7	HCU-1500	3360
4	CDE 420 Content Delevery Engine	3320
8	SEM v12	3072
6	DiviCom ION Multichannel Encoder	2880
1	1830 PSS-32	2880
8	Cherrypicker CAP1110	2810
1	Catalyst 4507R	2800
8	Cherrypicker CAP1000	2800
1	Cadant C4	2800
12	NC-1500	2760
2	System X3550 M2	2700
4	Serial 8 channel mux/demux (SCMD8-0)	2600
17	MV-100 Encoder	2550
1	OME 6500	2520
3	S/DMS Transport Node	2520
3	PowerEdge R710	2520
3	System x3650	2505
1	1830	2400
5	Sentry	2400
4	TV Guide Unit w/ Tray	2400
2	720XD Server	2200
55	DFR-8110A	2200
3	xSeries 345	1980
7	Catalyst 4948e	1925
6	ProStream 9000	1830
1	Master Switch Network Power Controller	1800
5	CableVista Chassis	1700
15	DVM 150E Receiver	1620
4	Sunfire V120	1600
1	DV6408ES	1512
3	Proliant DL320 G6	1500
2	DL 360 Gen8	1500
2	Proliant DL380	1470
1	4507 R	1444
1	DACS 5500 NGX	1440
5	Observer Scout	1400
9	9952 power amp	1357.2
1	ONS - M6 - 30 Amp	1260
1	VS-128 Power Supply (Video Commander)	1200
2	MN20-4	1200
5	DSR 6000	1150
1	PowerEdge R720xd	1100
1	Catalyst 4500+E	1092
9	HSM 1000	1080
21	D9858 Receiver	1050
16	DSR-4400MD	1040

Table A5 – Equipment Manifest for Sacramento, CA >1000 watt total, measured using internal site catalog database

Quantity	Model	total_watts
4	NSG Pro (CCAP)	17600
4	6506-E	16000
1	CRS 1	13900
16	Proliant DL360p Gen 8	12000
9	UCSC-C220-M3L	11700

9	NSG 9000 QAM Modulator Turbo	10584
5	6500 Transport Shelf	10080
4	UBR10000	9600
1	MX-960 XD	9348
15	TV Guide Unit w/ Tray	9000
5	TRX-24000	8000
22	CAP-1000	7700
1	ASR 9000 - 7.6 KW	7600
9	DV6016ES	7500
8	BMR 1200A	7400
6	PowerEdge 720XD	6600
1	7609 Router 6000W	6000
8	Weather Channel / Weatherscan Unit	5760
2	OME6500 (60Amp)	5040
12	CHP Max 5000	4548
1	7609 Router 4000W	4000
26	MV-100 Encoder	3900
6	RFGW-1-D	3744
3	VS-128 Power Supply (Video Commander)	3600
2	(JDSU) HCU-1500	3600
7	Intellistar 2 HD	3500
9	Intellistar 2 JR	3240
4	nCube n4x	2920
5	NSG 9000	2765
17	D9479-2 Modulator	2720
36	D9032 Encoder	2700
1	7600 Series	2700
11	NC 1500 Platform 6	2640
4	EX 4550	2600
20	DM 6400 Cherry Picker - DC	2592
7	Cricket Housing (IQK-CRKTRCK-001)	2520
7	Cherrypicker CAP1000	2450
2	Catalyst 6500 Series	2400
4	VIPr	2016
3	CMD44 44 Channel Mux/Demux	1950
5	Vista CV1107	1750
6	ES-247A Quad 1x6 Video Amp	1738.416
3	CR200 Card Cage	1620
6	DL 320 G6	1557.594
9	D9479 - 2 GQAM modulator	1440
1	4507 R-E Switch	1400
4	Intellistar	1400
5	Observer Scout	1400
6	SEM v8	1380
2	PowerEdge 1950	1340
2	EX-4550	1300
7	DVM-150E	1260
2	7200 VXR	1246
2	5618 16x16 Switcher	1200
1	XMS FLEX	1200
4	Star XL	1200
3	Atlanta LaserLink III	1125
7	D9479 - 2 GQAM modulator	1120
2	PowerEdge 1850	1100
12	6380A BTSC Encoder	1080
4	Apex 1000	1028

Table A6 – Equipment Manifest for Berlin, CT >1000 watt total, measured using internal site catalog database

---

[1] See Appendix A Figures A1-2, 4 showing that data center electricity expenditures are at least one order of magnitude lower than the national foot print of cable infrastructure utilities; also showing that representative cable sites average about 75,000 kWh monthly. For example, Comcast has approximately 1800 cable sites, excluding outside plant power supplies.

[2] Oró, Eduard, et al. "Energy efficiency and renewable energy integration in data centres. Strategies and modeling review." *Renewable and Sustainable Energy Reviews* 42 (2015): 429-445.

[3] See SCTE's Energy management program and goals http://www.scte.org/energy/, accessed on Thursday May 14 2015.

[4] Using an integrative design approach to achieve broader appeal and acceptance, the team includes Data Center Managers, Facilities Managers, Plant Managers, Vendor Professionals, Software Engineering, Network Engineering, Facility Engineering, Data Center Engineering, Capacity Management and Finance.

[5] Site engineers often focus on site capacity maintenance and upgrades, because the financial justification to improve energy performance can be difficult to justify, absent compelling financial data.

[6] E.g. new technologies that have unknown impacts to energy density include CCAP, SDN, and all IP-based network similar to data center infrastructure.

[7] Many cable plant sites originated prior to the ballooning of large centralized data centers.

[8] See, Patterson, Michael K. "The effect of data center temperature on energy efficiency." *Thermal and Thermomechanical Phenomena in Electronic Systems, 2008. ITHERM 2008. 11th Intersociety Conference on*. IEEE, 2008, demonstrating that thermal efficiency does not necessarily increase when cooling temperatures are reduced.

[9] See, Qian, Xiaodong, Zhen Li, and Zhixin Li. "Entransy and exergy analyses of airflow organization in data centers." *International Journal of Heat and Mass Transfer* 81 (2015): 252-259.

[10] See, Lovins, A. (2010, March). Integrative Design: A Disruptive Source of Expanding Returns to Investments in Energy Efficiency. Retrieved May 1, 2015, from http://www.rmi.org/keysolutionsindustry2, for several related case studies.

[11] Building industry measurement of Energy Intensity uses Kbtu / square-foot-year, 1kWh = 3.412Kbtu

[12] Including data centers

[13] See Appendix Figure A4

[14] Roughly $13.00 per year per subscriber

[15] See Appendix Figure A3, for a sample breakdowns of two large utility providers.

[16] Several case studies reveal anywhere from 5-10% savings off electricity expenditure.

[17] Chiefly because the impacts on energy use and power supply deployment of a remote CCAP implementation remain unknown.

[18] Toy, Mehmet. *Cable Networks, Services, and Management*. Vol. 13. John Wiley & Sons, 2015.

[19] Gao, Tianyi, et al. "Experimental and numerical dynamic investigation of an energy efficient liquid cooled chiller-less data center test facility." *Energy and Buildings* 91 (2015): 83-96.

[20] E.g. see inf. "Foxboro's new services"

[21] Typically, data centers contain equipment that run an operating system, like Windows Server, Linux or UNIX. This can be used to quantify the size of a data center.

[22] Although, not all new data centers are located in climate zones 5-7, indicating a need to highlight climate as an important design consideration.

[23] See, Baechler, Michael C., et al. *Building America Best Practices Series: Volume 7.1: Guide to Determining Climate Regions by County*. No. PNNL-17211 Rev. 1. Pacific Northwest National Laboratory (PNNL), Richland, WA (US), 2010, for a listing of coded climate zones by U.S State and County name.

[24] Note, not square footage density.

[25] Incidentally, with over 1,300 CCAP capable devices.

[26] See e.g. inf., "Large scale regional CCAP deployments"

[27] See inf., section 5, improving density using site characteristics and energy use.

[28] E.g. often data centers use the number of transactions to charge back to business units.

[29] A building envelope is a term used to describe the elements that comprise the enclosure of a building, i.e. walls, roof, insulation, windows.

[30] See Appendix Table A3 for a catalog of equipment.

---

[31] Based on estimated network throughput, see also “Large scale regional CCAP deployments” for similar estimated improvements to throughput.

[32] See Appendix Table A1 for a catalog of equipment.

[33] Note that SCTE Energy 2020 Facility Classification is currently not referenced.

[34] Note also that the new building takes advantage of a white roofing material that can significantly offset summer heat.

[35] See also infra, “Large scale regional CCAP deployments” for improvement to throughput.

[36] See Appendix Table A2 for a catalog of equipment.

[37] Here it will be instrumental to measure the site efficacy of its HVAC system to determine the size of the opportunity present in replacing or tuning the cooling at this site.

[38] Similar deployments of about 50 CCAP devices were described at the SCTE Energy 2020 May plenary meeting in the New York City market.

[39] See Appendix Table A6 for a catalog of equipment.

[40] Greenhouse Gas (GHG) Emissions are the carbon dioxide (CO2), methane (CH4), and nitrous oxide (N2O) gases released into the atmosphere as a result of energy consumption at the property. GHG emissions are expressed in carbon dioxide equivalent (CO2e), a universal unit of measure that combines the quantity and global warming potential of each greenhouse gas.

[41] See Lyndon case study where implementation of virtual hubs did little to change the regional energy density because commercial data infrastructure remained at consolidated sites and remote nodes remained at the inside plant site due to safety concerns, whereas it may have been more efficient to move nodes to outside buildings eliminating cooling and lease expenses.