Automation of Virtual CCAP to Reduce OpEx and Enable New Revenue Streams in the Access Network

A Technical Paper prepared for SCTE•ISBE by

Michael O'Hanlon
Principal Engineer
Intel Corporation – Network Platforms Group
michael.a.ohanlon@intel.com

Eric Heaton
Platform Solutions Architect
Intel Corporation – Network Platforms Group
eric.d.heaton@intel.com

Table of Contents

List of Figures

List of Tables

Introduction

Network utilization is dependent on time, varying according to the hour, day, or week and reflecting activity spikes due to special events (e.g., sports games, news conferences). Figure 1 shows a typical profile of aggregate subscriber traffic throughout the day. This is just one example, but the key thing to note is that the demand fluctuates over some time period. In this case, one can see that there is approximately an eight-hour period of prime-time network activity with 16 hours of relatively low usage. In fact, primetime growth is outpacing average traffic growth, and thus driving CapEx spending for new access equipment to meet this demand.

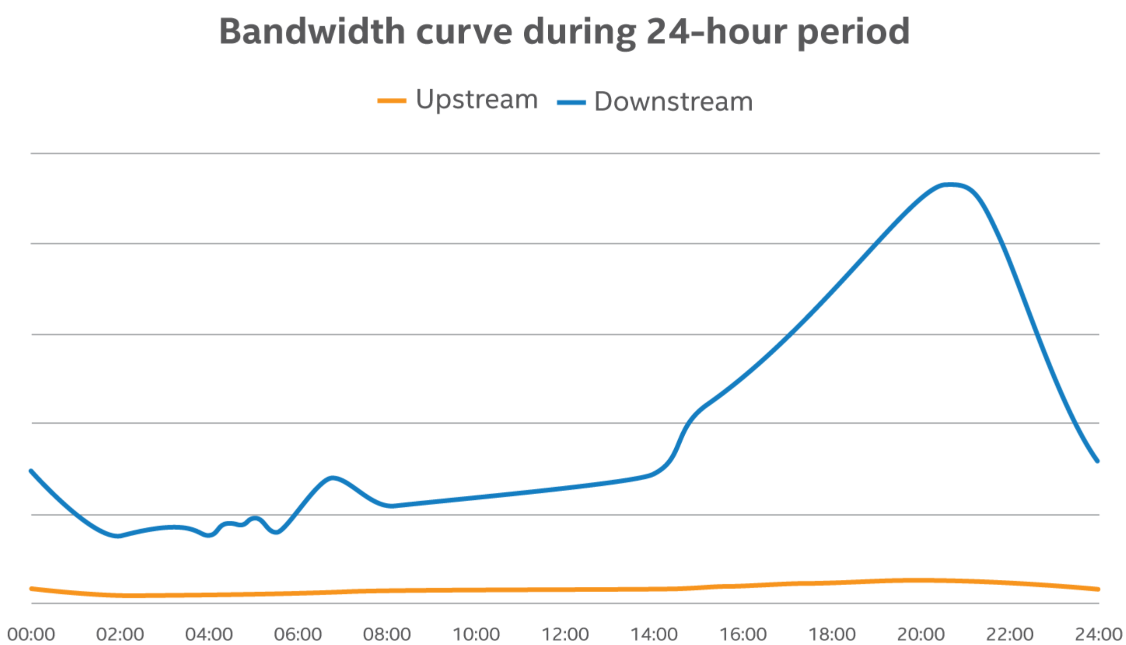

Figure 1 - Primetime driving peak network demands

It is important to consider whether current network infrastructure has the capability and flexibility to adapt to these changing needs and conditions. Can it provide the lowest cost-per-bit, while always meeting the real-time demand of users? This paper provides a road map of options to help with the planning and deployment of a next-generation access network that takes these requirements into account.

Fixed appliances used to deliver network functionality (switches, CCAPs, EPCs, firewalls, etc.) are hardcoded with certain features and capabilities despite a clear underutilization of their compute, network, and storage resources during large parts of the day. However, with the emergence and maturity of SDN and NFV, the network architect can take a more intelligent path and design a flexible system reactive to the needs of both users and operators.

In this case, network infrastructure can be seen as a flexible entity with behavior and parameters that can be optimized based on real-time technical and business needs. In other words, the infrastructure will have a “state” at any given time that you can control and manage with the right hardware and software, as discussed in the following sections.

This paper specifically focuses on optimizing power usage for a virtual Converged Cable Access Platform (vCCAP) data plane VNF running on standard COTS servers, but this research is applicable across any type of network function. Figure 2 shows the general network transformation from purpose-built network appliances to virtual software functions running on a common server-based infrastructure.

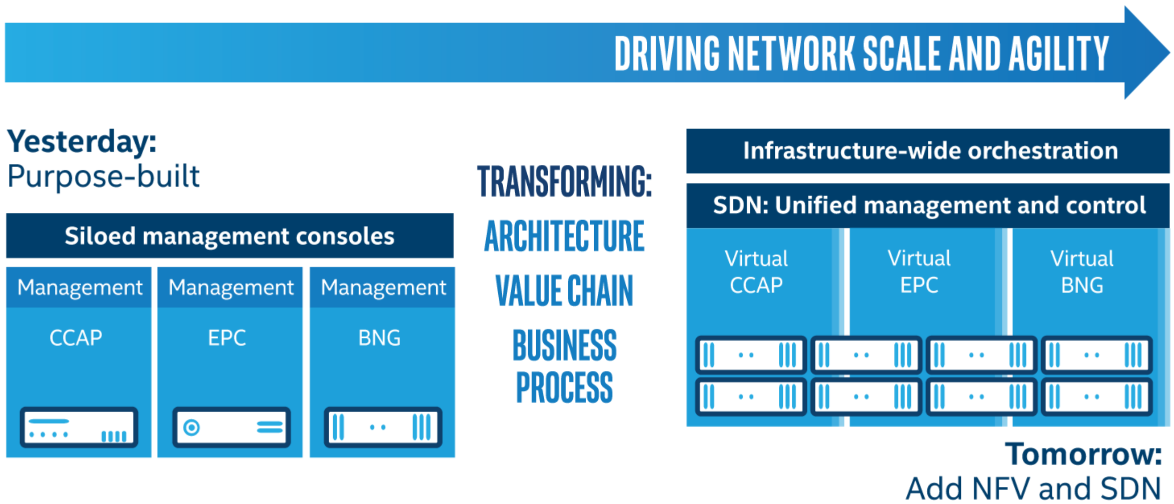

Figure 2 - Moving access functions to standard COTS servers

The power measurements discussed below can be understood in two ways: first, as literal savings in operational costs to pocket, and second, as a proxy for headroom in the infrastructure resources to perform other tasks. Further, the effort has pushed the state of the art for demonstrating best practices in hardware and software to realize the most efficient use of one's access infrastructure for the vCCAP or any other type of network function.

Our research provides a proof point of how a modern, container-based network functions virtualization infrastructure (NFVI) can be used to scale up or scale down various vCCAP operational parameters of the VNF itself or the platform it runs on through the collection of platform and network telemetry; a framework for decision-making through static or machine-taught policy engines; and, ultimately, the system automatically recognizing and reallocating resources to best meet its functional and operational requirements.

The learnings from this work will allow an operator to understand the potential of their SDN- and NFV-based network infrastructure to enable overall greater business agility and reduce network TCO. In fact, this paper presents a continuum of options that can be deployed in order to maximize the value of CapEx spent to upgrade from the legacy appliance model to one based on agile software.

- The first section discusses power management features native within COTS servers and techniques that can be used to save power when demand is low.
- The next section evolves this basic approach, allowing intelligent orchestration layers to make use of workload consolidation across a pool of servers to bring average server power down to the absolute minimum.
- The last section look to using this downtime in one access application to do other revenue-generating work, in particular it focuses on the potential for looking across multiple fixed, mobile, and enterprise needs with the goal of maximizing total resource utilization, thus maximizing the value of the investments of the operator.

Equipping your SDN- and NFV-based access network properly can provide the capability to:

1. Pay as you grow where labor costs for upgrades are high—for example, installing based on future-forward requirements and waiting to activate until the right business case is defined

2. Reduce OpEx by pocketing power savings when demand is lower

3. Use these savings to support more system maintenance and/or security (e.g., equipment failure detection or prediction, optimized redundancy schemes, security scans)

4. Enable new commercialized and next-generation services, such as VR/AR, smart cities and homes, autonomous driving, and IoT, on the same server infrastructure

Active Power Management per Server

In order to be able to take advantage of opportunities in NFVI for cost savings and/or to have the flexibility to otherwise use it to deploy and run new workloads, key NFV features must be part of the solution. To illustrate these elements and show how they can be used to optimize the needs of a particular network or business, this paper looks at how to minimize the fixed and dynamic costs of running a vCCAP data plane on an individual server.

It is important to note that standard commercial off-the-shelf (COTS) servers running NFV software already include a suite of power management tools. These can be used to increase or decrease the clock frequency of many hardware elements in the system to put them in lower power modes or turn them off altogether. While these capabilities are generally available, they are not always fully utilized to reduce OpEx. This paper focuses on the server's general ability to change the core and uncore frequencies of the CPU as it will give the operator the greatest "bang for the buck;". The core frequency generally applies to the cores themselves[1] (ALU, FPU, etc.) and the L1 and L2 caches, and the uncore frequency applies to shared resources, such as the LLC, integrated memory controller interfaces, and a few other tightly integrated internal units.

Initial benchmarking takes a look at vCCAP data plane performance measured as throughput against the AC wall power consumed by the server. This enables the modelling of a system based on dynamic performance and, conversely, power demands. For example, as bandwidth needs go down, compute, memory, and network elements of the system can run slower and still keep up, and with those slower clock frequencies, one will see a proportionate reduction in power.

Figure 3 shows the main components of the servers used to run these tests. This is just one sample server configuration out of many possible in the market. The key items to note are that the server has two Intel® Xeon® SP processors, each with 20 cores with a default frequency of 2.4 GHz for both the cores and the uncore logic. Each core is running one instance of the vCCAP data plane VNF and handles one Service Group (SG) with its data traffic coming in on one of the twenty-four 10 GbE ports available in the system. The maximum throughput per core (i.e., per SG) when the system is running at the default

[1] Note that individual core frequencies can be varied independently of each other allowing different cores to run at different frequencies. For the purpose of this paper all cores are set to the same frequency.

frequencies[2] for the configuration is 6.4 Gbps. Note that more than one core could be used to saturate each 10 GbE port, but it was not necessary for this testing. Details of the vCCAP dataplane software implementation and test environment can be found in the published paper: "Maximizing the Performance of DOCSIS 3.0/3.1 Processing on Intel® Xeon® Processors."

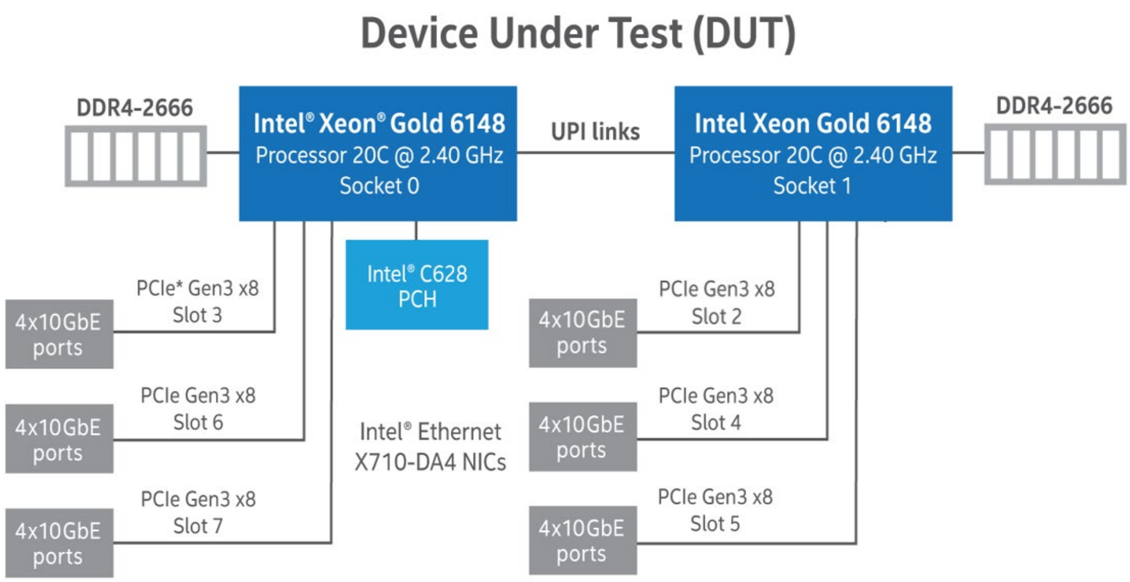

Figure 3 - Test server and benchmark components

Figure 4 shows server power measurements as one adjusts the core and uncore speeds when running 24 vCCAP instances. The core frequency is shown on the x-axis and the AC wall power measured for the system is the y-axis. Two different lines are used to represent the measurements, while the uncore frequency was held at either the default of 2.4 GHz (orange) or reduced to its minimum of 1.2 GHz (blue).[3] In other words, the chart shows how much active power is required for the given server to pass the maximum amount of traffic possible per core at the specified clock frequencies with zero packet loss for all 24 service groups. There is clearly a linear relationship between the clock frequencies and the power consumption of the server.

[2] The default frequency is also known as "Base" frequency. It is possible to achieve greater performance by increasing the frequency of the cores above the default frequency using turbo modes. This feature and its application to NFV is outside the scope of this paper.
[3] The uncore frequency can be adjusted anywhere in the range 1.2 Ghz to 2.4 Ghz.

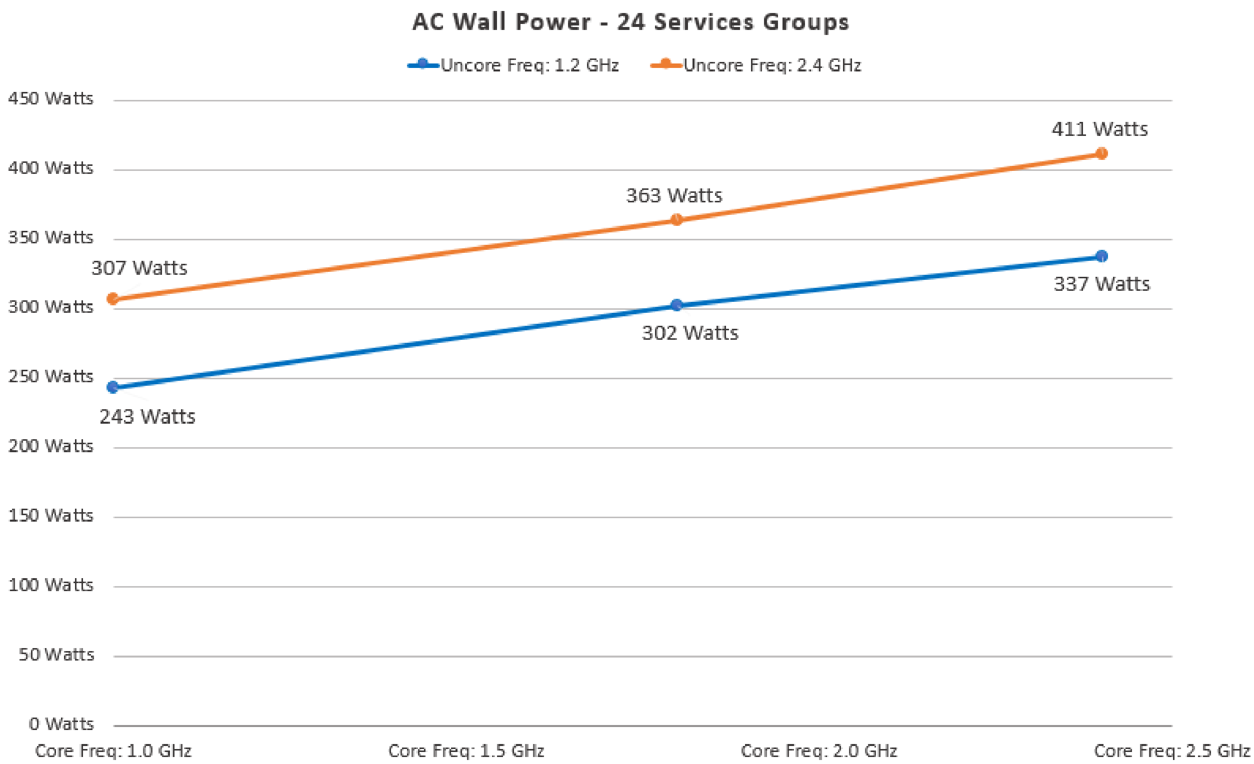

Figure 4 - Wall power measurements for vCCAP system as clock frequencies are varied

Figure 5 summarizes the power measurements made for the 2RU server described above while varying the clock frequency of three different entities: (1) the cores running the vCCAP data plane workload; (2) the cores running other applications; and (3) the uncore.

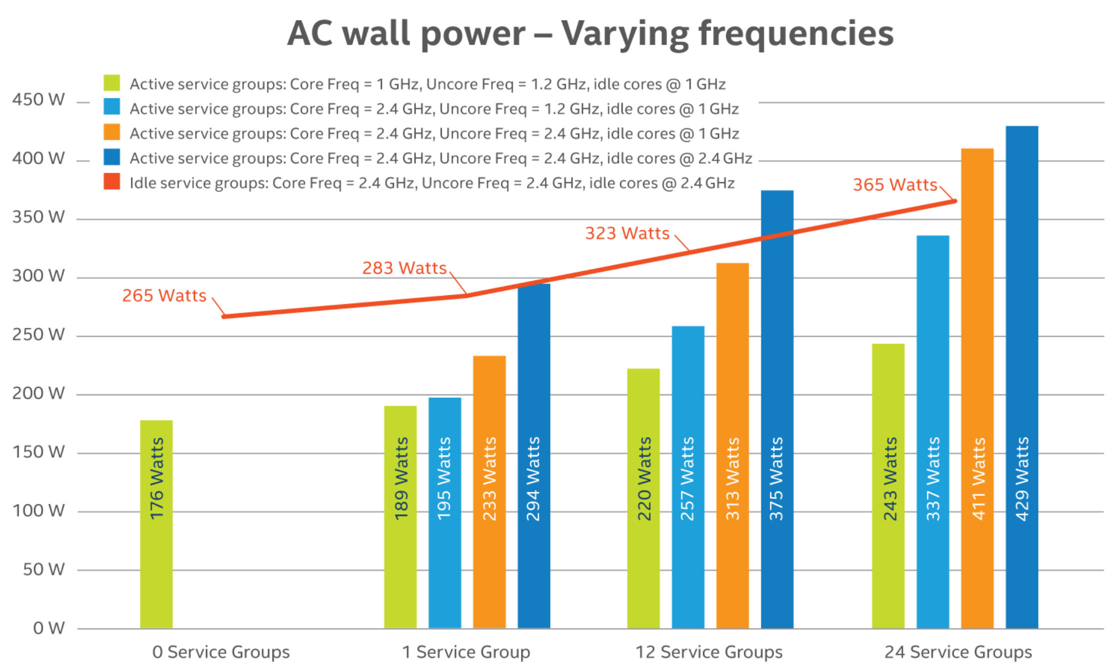

Figure 5 - Wall power measurements based on varying clock frequencies

The left side of Figure 5 shows the idle power for a given server when not running any applications and taking advantage of minimum core and uncore frequency settings. More advanced settings that could increase power savings but also affect the responsiveness of the system (e.g., to start up and instantly run a new workload) were not used. For our purposes, this measurement represents the baseline or "static" power per server. Note that while this research focuses on power savings available through frequency scaling of the CPUs, there are other components in the test system (memory, NICs, etc.) that contribute secondary levels of power. A future analysis could add these into the optimization model.

Moving right from the optimal idle power measurement, the other power measurements are bunched based on how many service groups are being handled by the server for that test. For example, the next four bars show the power measured in different frequency permutations when only one service group is being handled, and then the next four bars show the measurements for 12 service groups, and so on. The right-side bar in each of these bunches is the default power of the system (when the core and uncore frequencies are at their defaults). Conversely, the left-side bar shows the lower power possible when these frequencies are dialed down to their minimums.

As vCCAP data plane instances are loaded and running, there is a linear increase in the "dynamic" power of each system to account for packets being received by the network interface controller (NIC), sent to the cores and/or memory, processed, and then sent out to another NIC port. In this way, Figure 5 shows how the wattage demands increase across all clocking permutations.

This data establishes a couple of things. First, even the most basic power management features, like adjusting clock frequencies, do have a tangible effect on the power used by a server and therefore the OpEx of the system. Second, there is a maximum and minimum amount of power each server will consume for a given workload, depending on how fast various elements in the system are clocked. Of course, if the clock frequency of the cores running your vCCAP data plane is reduced, it will handle less throughput. But what is the derating factor and how do we map this all back to meeting the real-time demands of the network?

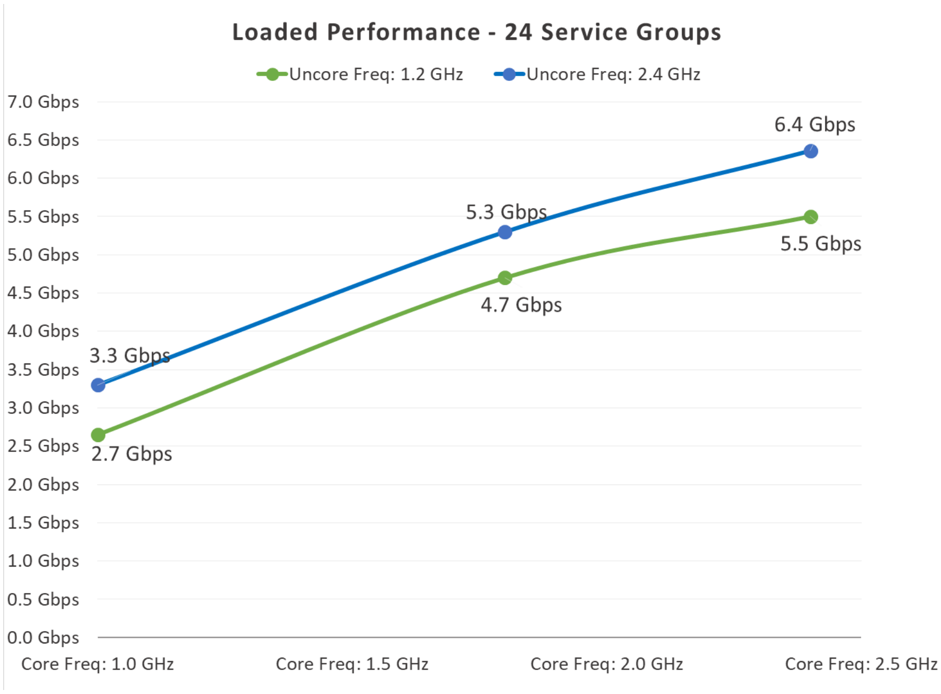

Figure 6 - Mapping vCCAP performance to CPU Frequencies

Figure 6 shows how throughput per SG per core was affected as the different clock frequencies varied. Table 1 below the graph takes the calculations further and summarizes the power per SG and throughput per SG across all core and uncore frequency permutations that were part of the testing.

Table 1 - Calculating expected bandwidth per SG as frequencies vary

Core Frequency	Uncore Frequency	Power per Service Group	Service Group Bandwidth
1.0 GHz	1.2 GHz	10 W/SG	2.7 Gbps/SG
1.8 GHz	1.2 GHz	13 W/SG	4.7 Gbps/SG
2.4 GHz	1.2 GHz	14 W/SG	5.5 Gbps/SG
1.0 GHz	2.4 GHz	13 W/SG	3.3 Gbps/SG
1.8 GHz	2.4 GHz	15 W/SG	5.3 Gbps/SG
2.4 GHz	2.4 GHz	17 W/SG	6.4 Gbps/SG

To select a specific example from the data above: if you reduce the uncore frequency from the default of 2.4 GHz down to 1.2 GHz (the green line in the chart) and also reduce the core frequency from 2.4 GHz to 1.0 GHz, then you can expect the vCCAP data plane VNF running on that core to handle about 2.7 Gbps of traffic. To put it another way, if the throughput demand of a given SG is only 2.7 Gbps, you can reduce your uncore frequency from 2.4 GHz down to 1.2 GHz and the core frequency from 2.4 GHz to 1.0 GHz. This reduces the server power needs by 40 percent—from 411 watts to 243 watts.

Consider that 40 percent power reduction across 10, 100, or 1,000 such servers at a given location making up the access infrastructure and it adds up to considerable cost savings!

Software Infrastructure for Active Power Management

In order to realize these savings, there needs to be software in the system that can automatically adjust the aforementioned frequencies in response to real-time system behavior. This will require the solution architect to define a set of Key Performance Indicators (KPIs) for the network; choose the system telemetry that best represents these KPIs; define policy and the associated actions engine to maintain the KPIs within a desired range; choose the right tools for the job; and figure out how to automate the process going forward.

Most Linux* distributions will include many of the tools you need to scale core and uncore frequencies at runtime, but you may have to implement new logic within your applications or create your own "glue" software at a higher layer to take advantage of them for maximum effect. For example, Figure 7 shows how a data plane application running in a Kubernetes-based NFVI can use the Data Plane Development Kit's power management library to automatically detect opportunities to save power, while still meeting the required latency and throughput demands of the network operator.

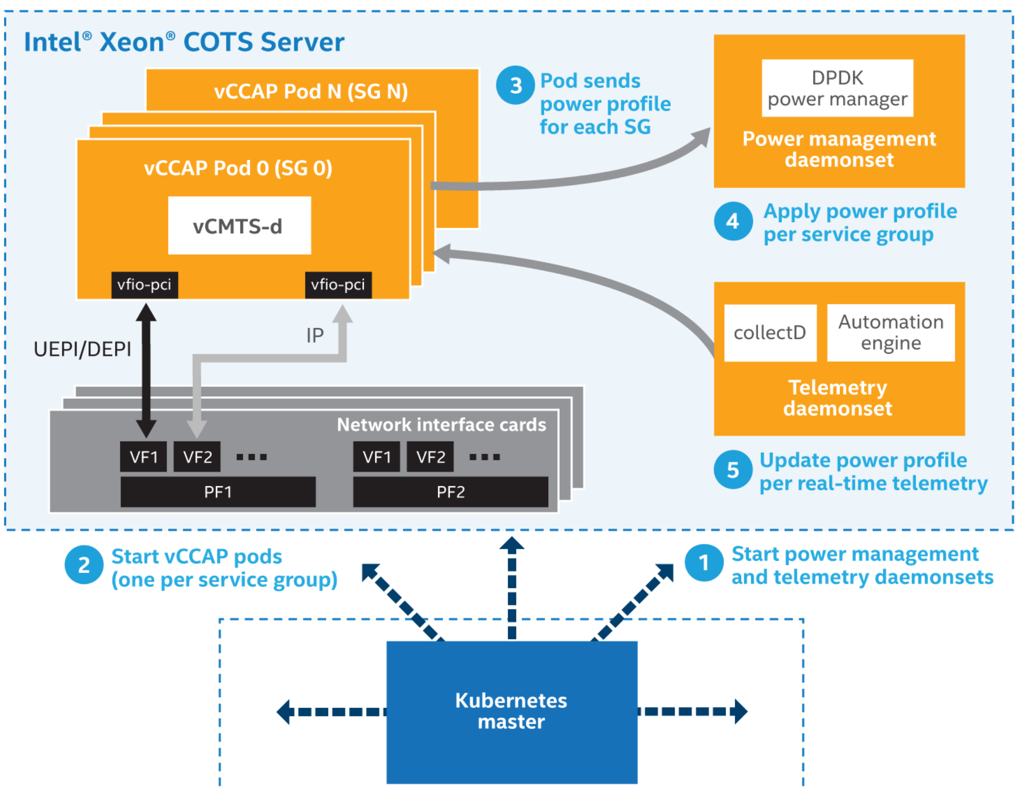

Figure 7 - Example of automated frequency scaling for vCCAP data plane

This example shows all the elements discussed earlier to implement a network infrastructure that can respond to real-time demands: system telemetry collection; an engine to make decisions based on that data; and then a harness to be able to execute those decisions with minimal to no operator input. Here, Kubernetes-based container orchestration and management infrastructure is used to deploy all of these elements onto a COTS server. Once they are in place, these elements are able to understand what is happening in the system and adjust the core and uncore frequency of the platform according to a power management ruleset.

Looking back to our original demand curve, by deploying the frequency scaling techniques discussed above, a power curve similar to the one shown in Figure 8 can be achieved. The dotted red line at the top is the power consumed by the server running fully loaded at the default clock frequencies, and the green line is the power measured using optimal frequency scaling to accommodate the demand. There is a large amount of savings possible in this particular example, evident in the gap between the dotted red and solid green lines highlighted with the large arrow. Again, in those off-peak times, the server power usage is about 40 percent lower than the maximum. And over the full 24-hour period, the total savings accrue to approximately 33 percent relative to the default settings.

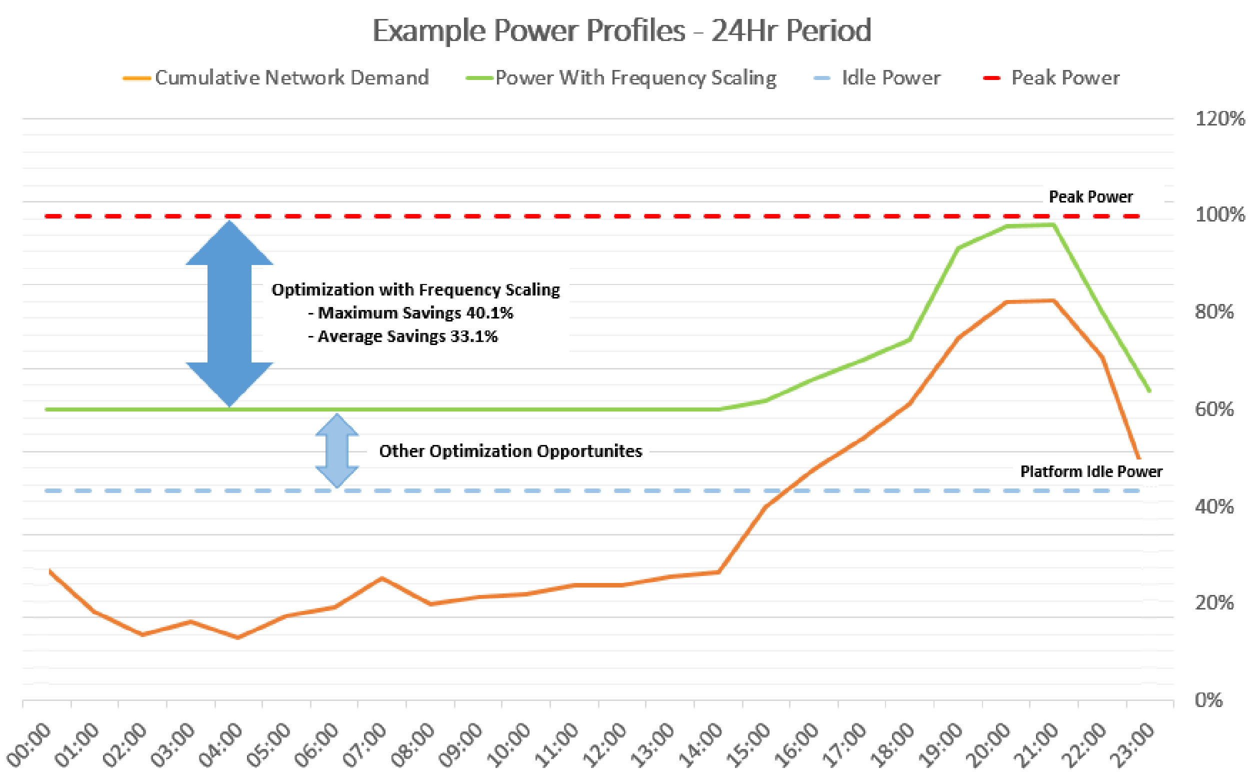

Figure 8 - Power measurements mapped to vCCAP throughout demand

The flat part of the green line representing the optimized power using frequency scaling techniques (approximately from the 0:00 hour to the 14:00 hour), shows a minimum server power (i.e., about 250W when all 24 vCCAP data plane instances are active) coming into play for these lower-demand parts of the day. It is beyond the scope of this paper, but knowing that the true minimum server power as shown in Figure 5 is about 176W[4], there are other power management opportunities in the platform. For example, the vCCAP data plane software used in the benchmarking assumes that maximum performance is expected all the time and thus generates a lot of work polling for new packets on the network interface, whether they are actually there or not. Of course, in lower-demand parts of the day, the network driver

[4] 176W is idle power of a system when no services are running the system, but cores and uncore are at base frequency. As stated previously, further power reduction is possible using other features outside the scope of this paper.

could be configured to reduce the amount of polling it does or moved to an interrupt-driven mechanism to further reduce system power.

These additional efforts would be rewarded with a possible further 30 percent reduction in the power usage (i.e., from 250W to the ideal of 176W). Alternatively, with the view that these power measurements are a proxy for excess resources in the system, the operator could decide to take advantage of this gap by running other applications on these servers "for free."

In short, the data above definitively shows that there are real operational savings to be had if the SDN- and NFV-based solution for vCCAP or any other VNF has the hooks in place to frequency scale different parts of your system in response to demand. In fact, with an understanding of the particular demands of your network, along with the particular performance curves for the desired VNFs, one can calculate the OpEx savings of the system and drive some of that investment back into more powerful servers up front.

Further, Moore's Law continues to bring down the fixed power costs of Intel® architecture-based servers, generation over generation, allowing more complex and intelligent power management features to become standard in the resulting hardware. In other words, as the performance per watt of the server CPUs and associated chipsets increases and new power management features within the silicon are developed, the platform idle power for newer equipment will be naturally lower.

The next part of this paper shows that there are even more ways to save power and/or use idle compute for running complementary applications by looking beyond the capabilities of an individual server and taking a pooled approach to network infrastructure through the use of smart orchestration tools.

Savings Through Orchestration

The next strategic approach starts by thinking of the network infrastructure as a pool of resources that can be managed in real time and not just stand-alone appliances to be individually controlled. To this end, one employs a full suite of software infrastructure and tools that allow the complete orchestration and management of all network functions and applications in such a way that in periods of low demand, workloads can be consolidated onto fewer and fewer running servers. This allows the infrastructure to be optimized for the lowest possible operational power and/or creates the space to run complementary applications. This approach can be seen as an evolution of the one discussed in the previous section or pursued independently, as one starts moving from network appliances to SDN and NFV.

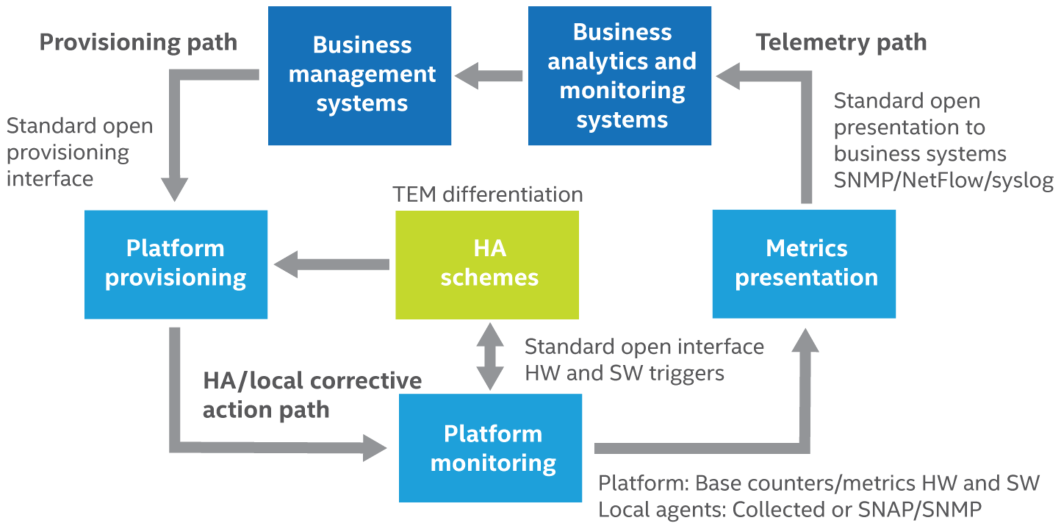

Figure 9 - Software elements and logical control over the hardware

Figure 9 shows the software elements required to create this type of environment—real-time telemetry collection, monitoring and analytics engines, business management and policy engines, and various action engines—holistically across all servers being used for access and edge service. With these elements in place, a fuller version of SDN and NFV is realized, where servers and attendant hardware are seen not as individual entities, but as a truly homogeneous pool of compute, network, and storage resources.

The model builds upon the calculations covered in the previous section, but focuses on reducing fixed power costs per system by consolidating work onto fewer servers. Each server introduces a fixed minimum power cost when running workloads. The idea is to reduce the cumulative fixed cost by powering only enough servers to meet the demand at any given time of the day. Of course, this only works in deployments where more than one server can be dedicated to the applications of interest.

First, total server needs are identified based on peak network throughput requirements and thus create a "pool." Next, the capability is enabled through software infrastructure to be able to fully move applications to any available server in the pool. Finally, a set of functions must be added to be able to detect when certain KPI thresholds are reached and then react per operator policy. When demand is low, this type of system allows application consolidation onto the minimum number of servers to still meet demand and fully shuts down any that are not used (saving 100 percent of the power they would use just to be "on").

To illustrate this point, Figure 10 shows that when demand is high for a particular virtualized application, like a vCCAP, workloads may need to run exclusively on three different servers to achieve the necessary peak performance. However, as real-time demand for that application drops, it may be possible to consolidate application instances onto one server. This degree of control allows the overall system to operate at power levels that almost exactly mirror the demand curve.

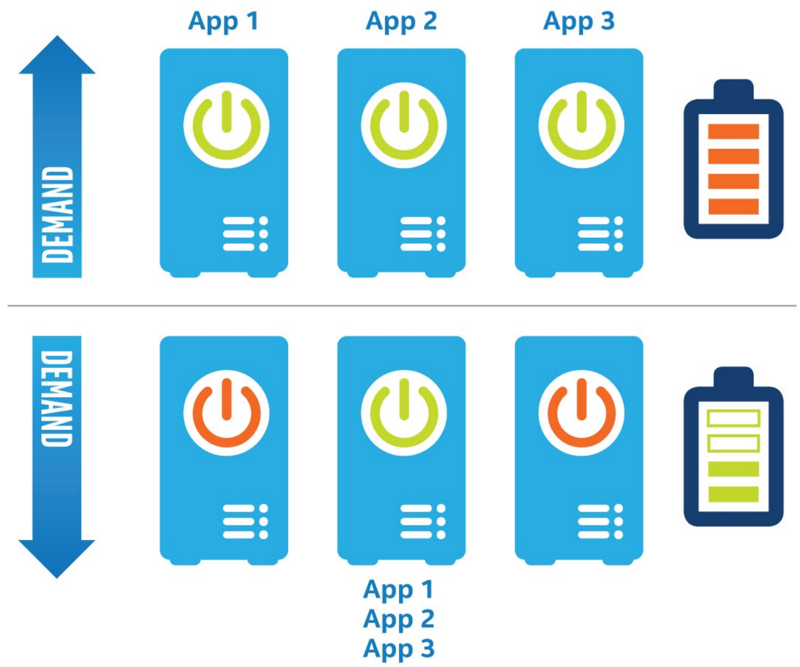

Figure 10 - Consolidation of vCCAP workloads

Calculating Potential Efficiency using Consolidation

Expanding the model starts with an assumption that all SGs will be serving the exact same throughput requirement for a given time of day. While this is unlikely, it makes the model calculations much simpler and illustrates the potential gains of workload consolidation. Taking a time-of-day example at 5:00 a.m., we see from the traffic capture data that each SG needs to support an aggregate bandwidth of 1.31 Gbps. This 1.31 Gbps represents 20.5 percent of the maximum possible per SG per core.

In theory, a core should be able to handle four service groups at 20.5 percent of peak demand with ease given the cumulative demand is 82 percent of peak. However, consolidating more than one application (i.e., vCCAP data plane instance) onto a core adds overhead due to context switching and for low-level resource sharing (e.g., cache). Consequently, there is a reduction in the per-core throughput. Table 2 shows the performance degradation measured as up to five vCCAP data plane instances are deployed to a single core.

Table 2 - Throughput for consolidated vCCAP data plane downstream instances

vCCAP Instances Per Core	Throughput Per Instance	Total Throughput Per Core	Performance Degradation per core
1	6.40 Gbps	6.40 Gbps	N/A
2	2.72 Gbps	5.44 Gbps	15.0%
3	1.56 Gbps	4.68 Gbps	26.9%
4	1.04 Gbps	4.16 Gbps	35.0%
5	0.72 Gbps	3.70 Gbps	42.2%

After taking the overhead into account for the original 5 a.m. example above, a single core can support up to three of those 1.33 Gbps SGs at the same time. Figure 11 repeats the calculation and maps these "consolidation factors" to the whole 24-hour network demand curve. The larger the consolidation factor, the fewer servers are required to meet the network demand and, thus, the greater amount of power savings possible.

In reality, each SG may have different throughput needs at any given time. In this case, the ideal solution would have the telemetry layer monitor how the servers are tracking demand over some timescale, and then have a machine learning-driven algorithm solving for the problem of packing SGs into the minimum amount of servers necessary. To return to our early morning example, at 5:00 a.m., if two SGs require the aforementioned 1.33 Gbps but two others require 2.74 Gbps, then at least two cores will be needed. In this particular case, each core will be allocated one 1.33 Gbps SG and one 2.74 Gbps SG, such that the maximum throughput required for all SGs per core does not go above the 6.4 Gbps maximum. All of these calculations can be handled in real time by a utility or an automated decision-making engine built into the orchestration software layer.

So while frequency scaling provides some very compelling power savings on the individual server level (as described in the previous section), each server will retain a minimum power requirement on the order of 176 watts[5,6] simply to be "on." By having a pooled view of resources being controlled under the same management domain, the number of servers active to deliver a particular service, like vCCAP, can be scaled down to allow for great savings in average power usage per SG. In this way, the system effectively breaks the minimum power barrier at the individual server level by amortizing the fixed costs of power supplies, memory, NICs, etc., by packing more SGs per core (and hence needing less servers to do the job).

---

[5] Minimum power varies from system to system and depends on several factors including but not limited to selected CPU, memory quantity and size, plugin cards, and storage.
[6] Lower idle power is possible where advanced power features are available and used.

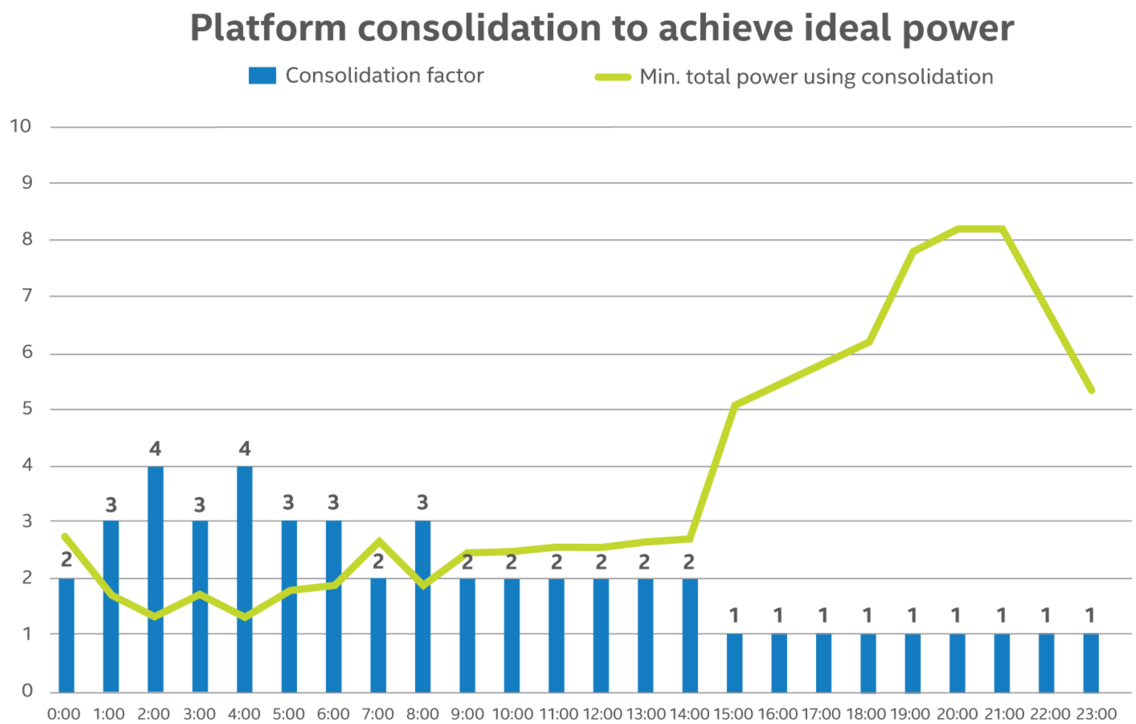

Figure 11 - Consolidation factor calculations mapped to 24-hour demand curve

For instance, if you were to consolidate four servers down to one server running at maximum frequency, one server might be maxing out its power profile, but there are also three other servers now turned off and contributing zero watts to the overall power draw of the system. The most savvy software infrastructure should actually implement algorithms to save power at both the individual server level and pool level in order to deliver maximum value for the infrastructure.

Getting back to the original demand curve and using the simplified assumption of homogenous SG needs and the calculations above, average server power requirements for a vCCAP data plane were plotted over a full day of traffic (see Figure 12). Again, the dotted red line is the power of the servers running the full vCCAP load with no power management enabled, the green line is the power measurements for the frequency scaled case described in the previous section, and the dotted blue line is the measurements when the fully orchestrated consolidation scheme described in this section is used. This last approach allows the network operator to truly tune the operational expenses represented by power to a minimum, while still meeting the needs of the users.

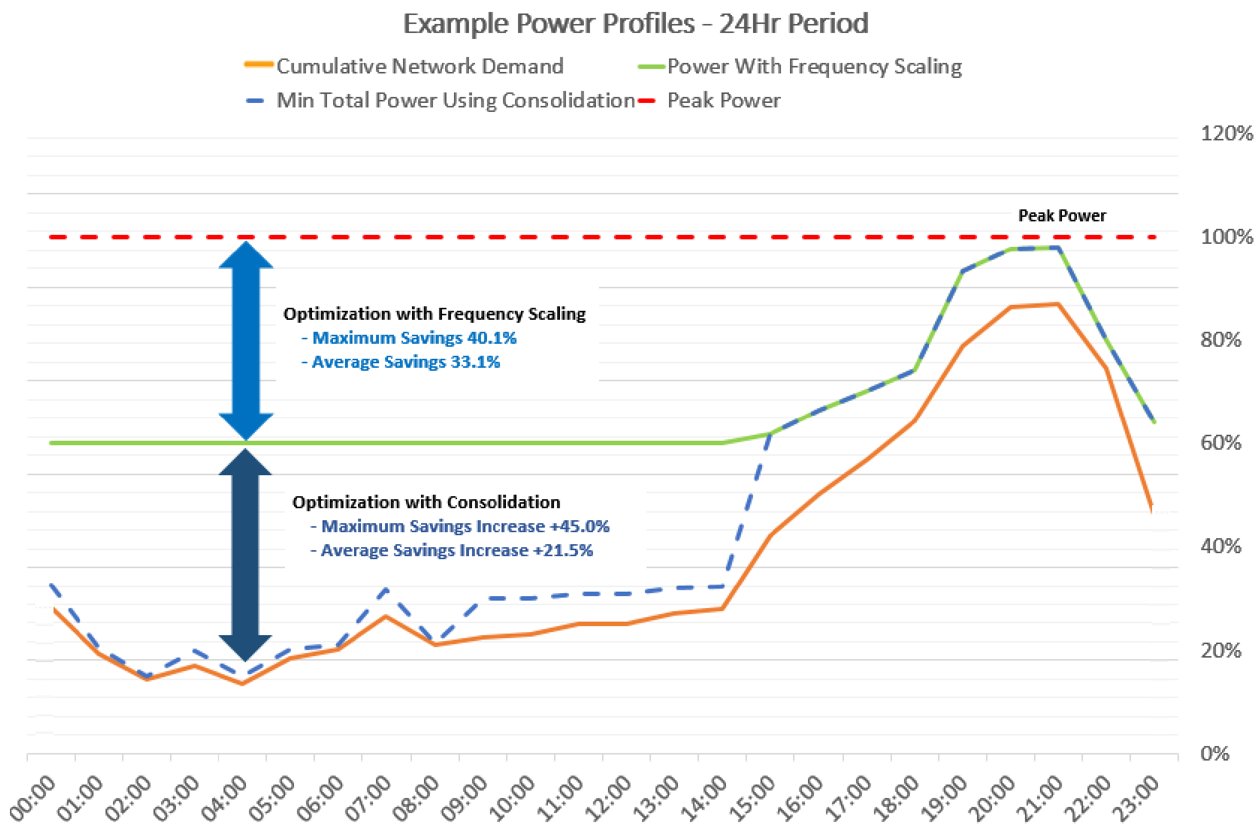

Figure 12 - Power measurements mapped to vCCAP through demand

The previous section introduced the frequency scaling opportunity for power savings at the individual server level; this is highlighted with the top blue arrow in Figure 12. This section took a "pool of servers" view toward the goal of saving power (or reusing the compute represented by that power), and broke the per-server minimum power barrier to achieve up to 85 percent savings in the lowest demand periods of the day. This is highlighted with the bottom dark blue arrow in Figure 12.

This work represents the start of what is possible and it is expected that the industry will continue to bring down the TCO of NFV-based infrastructure as more telemetry-gathering, decision-making, and automation layers are refined and added to deployed solutions. There will be up-front costs to develop or buy these new capabilities, but the data above shows that it will be made up many-fold over the lifetime of the equipment.

Consolidating Fixed, Mobile, and Enterprise onto the Same Infrastructure

As discussed, the power savings outlined in this paper can be seen as literal OpEx savings for providing the electricity to run the virtualized access infrastructure. The other view of the power savings metric is that it is also possible to reuse the spare compute, networking, and storage resources to run other access or enterprise applications on the same servers. In other words, by transforming the headend or central office to a distributed data center, a network operator can realize the full vision of network functions virtualization shown in Figure 2, where a common infrastructure can be used to support whatever functions are demanded by the operator and users in real time. There are several different approaches for this.

Some applications may be run opportunistically at any time if they do not have particularly high technical or business demands. For example, back-end machine learning-assisted analysis of network or user telemetry data. Others may have particular and unwavering demands of their own (similar to vCCAP), and thus can use the same type of demand curves and power/performance calculations described in this paper to harmonize how they share compute and other resources.

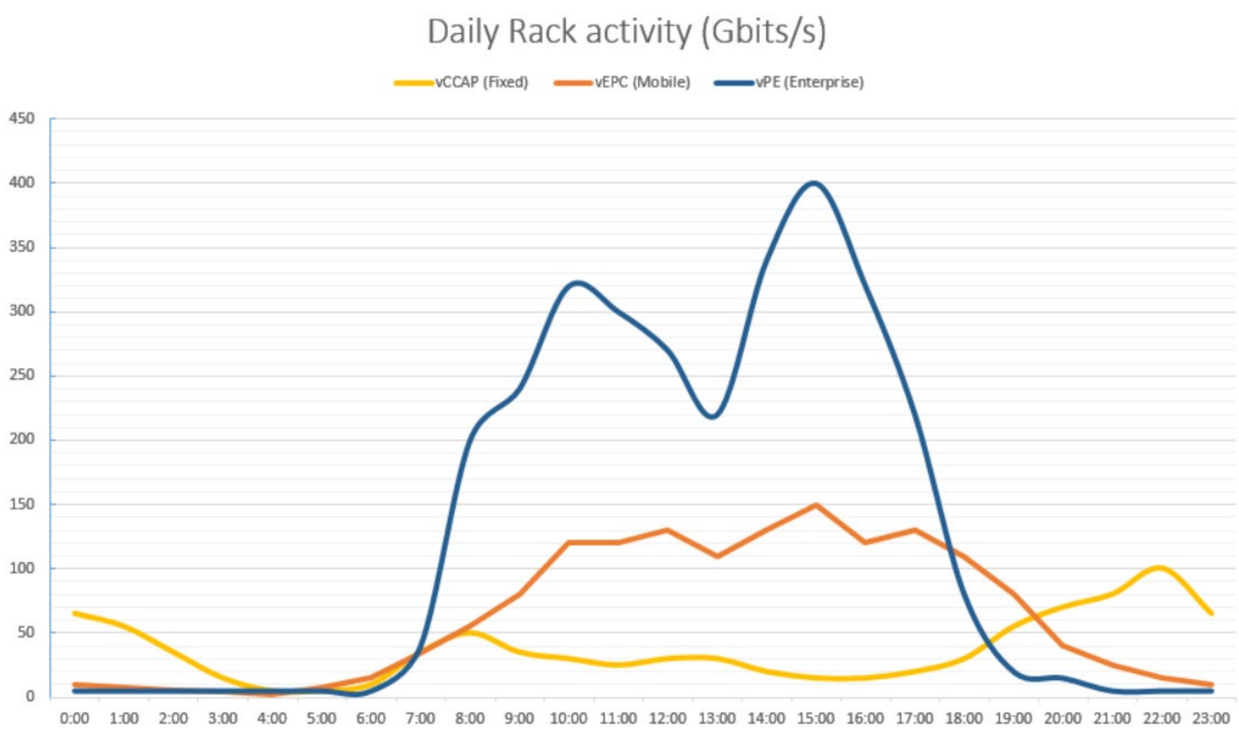

Figure 13 - Demand curves for fixed, mobile, and enterprise traffic over a 24-hour period

For example, Figure 13 shows that generally fixed, mobile, and enterprise workloads have different time-of-day demands. These usage measurements came from internal research for a Next-Generation Central Office (NGCO) that aims to support a mix of all of these services on the same COTS server infrastructure as an alternative to deploying parallel fixed function appliances with uncoordinated management facilities. In this way, expanding the intelligent orchestration and management concepts discussed in the previous section to comprehend more than one application at a time will allow the same equipment to be dynamically reused across all of these workloads, providing a cost-effective, flexible, and future-looking approach to network facility architecture.

Of course, Figure 13shows that while there are times where the demand is complementary (i.e., toward the start and end of the day), there is a bulk of time in the middle where at least two of the applications have high requirements at the same time.

This presents an area for further study, as supporting multiple demanding applications may require breaking up the notion of “one big pool” and instead utilize sub-pools dedicated to running only one type of mission-critical application. Or a sub-pool may be designed to run any number of instances of application one and two, but not application three. These decisions will be based on the characteristics of each application, such as whether all applications require access to common hardware elements, have demand curves that coincide (or are mirror images), etc.

Regardless of the particular implementation, it is clear that if the servers can be shared effectively using intelligent orchestration solutions, then inefficiencies can be driven out of individual server use and thus allow for fewer capital expenditures up front to support this wide range of services.

The details of implementing these strategies, and therefore the server needs, can be determined beforehand using theoretical or empirical data and then codified via a hardcoded policy engine. Or, an automated decision-making and policy-modifying engine can be deployed that uses machine learning algorithms to optimize the behavior of the system in real time. Strategies may also be modified, based on the availability of certain hardware features. For example, workloads with real-time requirements may initially be segregated onto a special sub-pool of servers, but it may be possible to use a single pool if the CPU features can provide resource determinism for workloads. The value of static versus automated decision-making is left for further study.

The beauty of SDN and NFV is that this is all defined in software, so that new sources of telemetry can be enabled and new decision models deployed. Essentially, all system operation can be managed in a flexible, agile manner, if the right amount of intelligence is employed in the solution.

Conclusion

SDN- and NFV-based solutions in the access network promise benefits over legacy hardware appliances in the realms of flexibility, manageability, and scalability. However, this paper highlights that while these systems might have the potential to deliver on these promises, not all solutions are created equal.

Using the vCCAP data plane VNF as a representative workload for other types of access technologies that could run in COTS servers, this paper outlined a continuum of options to reduce power usage as demand rises and falls over a given 24-hour period. At a minimum, the solution should take advantage of the frequency scaling of the CPU cores and uncore logic, as they are the largest contributors to both power and performance of the system for this type of workload. Savings per server can be on the order of 33 percent overall, with a peak savings of ~40 percent relative to the default server configuration!

In addition, if the solution adds intelligent orchestration and automation frameworks to the NFVI that can autonomously determine opportunities to consolidate the vCCAP (or other VNF workloads) onto the fewest servers possible to meet real-time demand, then an additional average of ~21.5 percent savings can be unlocked, with a peak of 45 percent for a total of ~57 percent lower power usage on average over 24 hours and upwards of ~85 percent in times of minimal demand. The power per SG can track demand very closely and thus provides the lowest TCO for the equipment.

The power it takes to run the access equipment can also be seen as a proxy measurement for the ability to run other workloads on the server. This allows an operator to take better advantage of the fixed costs of the equipment and/or create a foundation for new services on the same infrastructure. To this end, an evolution of the consolidation technique would account for not only a single workload (e.g., the vCCAP), but instead look across all the access and enterprise application needs of the network operator in order to take full advantage of the equipment at all times. In this way, a truly intelligent solution could find the maximum flexibility and savings across all infrastructure requirements.

By understanding the power and performance impact of specific features that can be made part of the NFV-based access network, operators can make better decisions for designing and deploying next-generation infrastructure to support the ever-increasing data throughput needs of their users over time and be able to nimbly respond to competitive pressures in a cost-effective manner. The ultimate aim is to unlock the potential of SDN and NFV to improve the user and operator experience, save costs, and create a foundation for new workloads and services in a world that requires constant evolution.

Bibliography & References

1. https://www.theatlas.com/charts/H1tALGE4