Data Center Effectiveness Metrics – 1H14

Client Challenges: Most data center executives are improving their data center effectiveness and efficiency within and across data centers as part of their cost cutting and optimization efforts. However, what still holds a number of executives back from making major strides in strategic optimization planning are the questions of where are the problems, where to begin and what the metrics, paybacks, and desired targets are. RFG finds IT executives want and need a baseline of best practices and best of breed targets against which to compare their current operations. RFG offers the following data points from the Greenway Collaborative Knowledge Base to use in a gap analysis and creation of a data center strategic optimization planning initiative.

Observations and Recommendations

RFG has written a number of research reports on the business case and rationale for pursuing data center best practices, resource optimization, and sustainability. However, the issue for a number of executives is not the generic rationale for optimizing IT operations. It is the understanding of what actions could be taken and the return on investment (ROI) on those actions. This calls for IT executives to know where to look, their current baseline, the extent of the gap, and the potential gains that can be achieved from improving data center operations. This research brief is the spring 2014 update and highlights the key areas and metrics associated with those areas that IT executives should examine.

*RFG studies find that most enterprises are not run at maximum effectiveness levels. As a result, a typical data center operation can **double capacity** while **reducing costs by up to 50 percent** and **cutting energy expenditures by up to 80 percent** within an existing data center. If one includes data center consolidation and use of cloud computing, the savings can be even greater.*

Key Data Center KPIs

According to APC/Schneider Electric, data center power and cooling infrastructure wastes more than 60,000,000 megawatt-hours per year worldwide. For example, in a typical 2N redundancy data center only 47 percent of the power actually is consumed by the IT load while the rest is used by the non-IT infrastructure. See Figures 1 and 2.

There are a few basic data center key performance indicators (KPIs) that IT executives need to follow to ensure overall data center energy efficiency and optimization of equipment resources. IT executives that meet or exceed these KPIs have been able to reduce their energy costs by up to 80 percent. For example, **Microsoft Corp.** built a third generation data center in Dublin that has a power usage effectiveness (PUE) of 1.25 using an N+1 redundancy model, an improvement of 50 percent over a traditional data center with the same reliability level. Moreover, the data center consumes less than one percent of the annual water consumption of a traditional data center. In Quincy, WA, Microsoft constructed a fourth generation data center with air cooled, pre-manufactured IT

containers, whose power usage effectiveness (PUE) is in the 1.15 to 1.2 range. The design cuts capital costs by 30 to 50 percent and reduces construction time in half. These data centers demonstrate that tier III data centers can do far better than they are today.
The most important factors to consider are:

- Further data center consolidation possible?
- PUE = 1.91 or better (top performing tier III centers have PUE of 1.40 or better; tier II centers can attain a PUE of less than 1.20.) Various studies place the PUE average somewhere between 1.9 and 2.9 today.
- Average cost per kWh < $0.099 (U.S.) (lowest average is $0.071 while highest average is $0.331)
- Lean provisioning / modular expansion / containerization of IT and non-IT equipment
- Rapid decommissioning of unused services and supporting hardware
- Use of the latest storage auto-tiering and management capabilities
- User chargebacks for products and services

Figure 1. Power consumption in typical data center

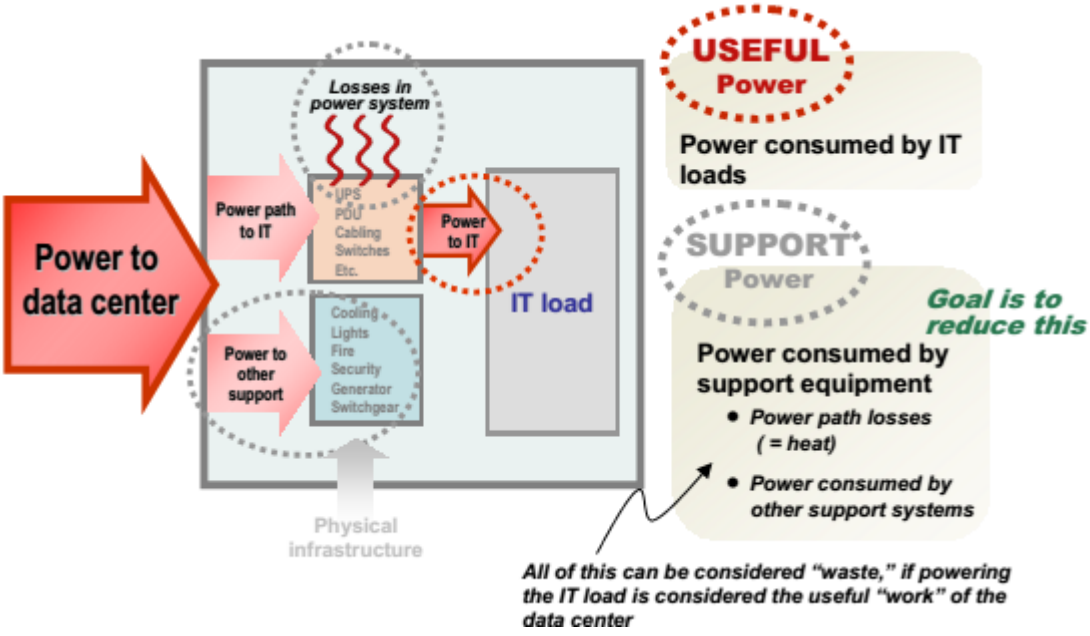

Figure 2. Power flow in typical data center

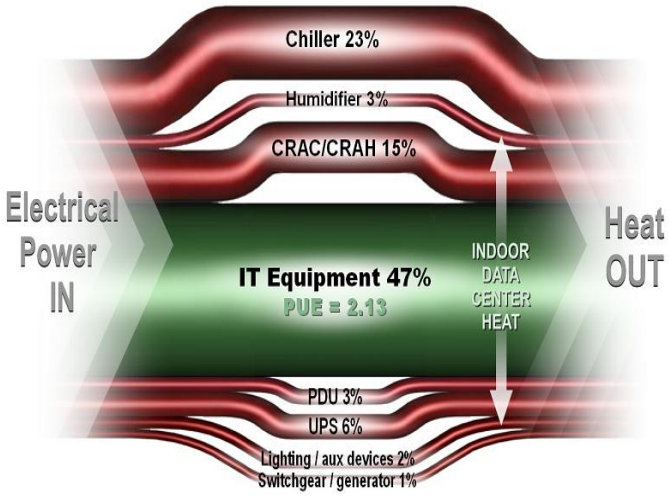

Source: APC/Schneider Electric

Other infrastructure energy conservation metrics are as follows:

- Air flow management – uniform static air flow in cubic feet/minute (CFM) ≥ 700 with static pressure of 0.1" of water
- UPS ≥ 85% efficient at IT loads ≥ 30%
- UPS redundancy = N+1 versus 2N
- High efficiency power distribution units (PDUs)
- Variable speed fans ≥ 50% of all fans
- Relative humidity between 40 – 60%
- Maximum dew point = 59 °F
- Temp > 72°F (22.2 °C) ±2°F and < 80°F ±1°F at inlet to IT equipment
- Rack cooling index (RCI) ≥ 96%

The rack cooling index (RCI) is a measure of how effectively air-cooled IT/network equipment is cooled and maintained within industry thermal guidelines and standards. (RCI is a trademark of **ANCIS Inc.**) Rack-based equipment that is improperly cooled will overheat – regardless of overall room temperature and potentially fail, thereby causing outages and/or loss of data.

Recent studies show that about one-third of organizations do not measure PUE. While 90 percent of all large enterprises measure their effectiveness, only about 50 percent of the small companies do. Service providers all recognize the value of PUE measurement; thus, virtually all track their PUE performance.

IT executives should note that some studies claim that enterprises can **save as much as four percent of their energy costs for every degree of upwards change in temperature**. Moreover, new technologies should allow hardware to run without problems with temperatures below 100°F. Thus, there is plenty of room for cutting cooling costs in all climates. IT executives should consider keeping data centers at room temperature.

The following actions can reduce costs as well.

- Rebalancing and reducing operating loads and overcapacity of power and cooling due to lighter IT loads
- Conversion to DC power or high efficiency 415/240 V AC power distribution
- Under floor cabling contains only current in-use cabling or use of overhead cabling. A switch to overhead cabling could cut cold leaks by up to 35 percent.
- Containment systems: Hot and cold aisles in proper use throughout data center; or ducted/plenum system in use
- Humidification and dehumidification in place with automatic cutoffs
- Energy Star equipment in use where possible
- Economizers in use
- Development of data center cooling zones
- High efficiency chilled water or direct expansion (DX) air cooling systems installed
- Modular cooling units installed
- IT equipment in rows according to industry guidelines and standards
- Cabling with patch panel connector
- Bundled cables – no airflow blockage

Key IT Operations KPIs

Best of breed data centers devote less than 50 percent of their IT budget on maintaining existing infrastructure whereas the typical enterprise spends 65 percent or more of its IT budget on existing infrastructure maintenance. IT executives that meet the following server metrics can reduce their overall operations costs by 25 percent or more. Depending on the current state of the equipment and software, IT executives can achieve significantly greater savings. For example, IT executives that have utilized the standard scale out approach to adding x86-architected servers can potentially reduce the number of servers in operation by 50 percent or more through a modernization program. Server virtualization can enhance the savings to a point greater than 75 percent.

- Server utilization ≥ 60% for Unix and mainframe
- Server utilization ≥ 40% for Microsoft Corp. Windows and Linux platforms
- For Intel servers ≥ 60% of servers should be virtualized
- Number of virtual servers to physical = 10:1 (minimum; best of breed ≥ 30:1; average today is 4-6:1)
- Number of servers more than 36 – 40 months old = 0 - 20%
- Performance/watt (peak business operations per second (BOPs)/watt) = > 2000 (average of all Intel servers)
- Number of operating system version levels of same OS in production = 1
- Unix/Wintel servers/administrator ≥ 60 - 100:1 (best of breed > 500:1)
- Power management software on > 10% of systems
- Servers powered off during downtime

As can be noted in Figure 3, servers consume a significant amount of power even when idle, often more than 50 percent of the peak power. Thus, if one can move the workload level from less than 10 percent to 50 percent through virtualization, one can obtain five times the performance for only 40 percent more power usage.

Figure 3. Performance vs Power

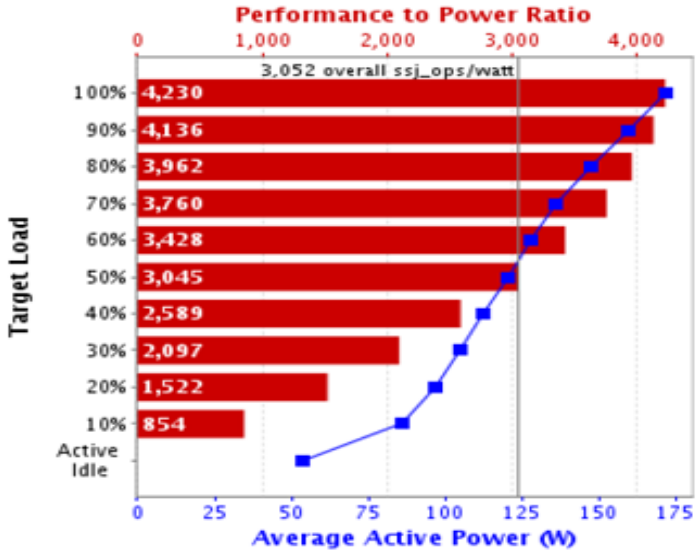

Source: Microsoft Corp.

IT executives should note that power management software can address power from two angles: power caps and power reductions. Most software addresses the sleep or quiescence modes. However, there is software that will look at the true power utilization of servers (versus ratings) and allow for controls on maximum usage across a set of servers as well as individual servers. The latter software can help contain power usage so that the requirement for an additional data center is deferred.

For those with or willing to consider mainframe environments, there are two additional savings options. Users can run highly virtualized Linux environments with either Linux on an IFL or an Enterprise Linux Server. The Enterprise Linux Server does not require companies to have or purchase a standard mainframe. It is a standalone IFL server that virtualizes and optimizes Linux environments. By taking advantage of the server consolidation, users can reduce overall Linux server costs by more than 60 percent.

The second option is to use the mainframe as the data server for Linux and **Microsoft Corp.** Windows applications. RFG studies find that data storage capacity can be reduced by more than 80 percent and overall system costs by more than 50 percent, when the shared storage functionality of mainframe storage is employed.

There have been tremendous advances in storage technologies as relates to energy conservation as well. In one case study that RFG looked at a company was able to expand its total storage from less than one petabyte to almost six petabytes over four years while

reducing its overall costs by more than $12 million. In other studies RFG found the use of flash technology can reduce storage costs by half. Since enterprises constantly upgrade storage and in many cases turn it over every three years, this is another excellent place to look for savings.

Use of any of the items below can result in significant savings.

- Policy based archiving process
- Flash storage or SSDs (solid state drives) ≥ 4%
- Deduplication ratio ≥ 20:1
- Compression after deduplication
- Use of golden records (copies not archived or backed up) = 50% of all key files/databases
- Tiering and auto-tiering -- tier 1 for high performance only < 35% of total storage
- Thin provisioning and storage virtualization
- Percentage of storage as external storage (versus direct attached) ≥ 80%
- Actual usage GB / raw capacity GB as a percentage = 60-70% (file or objects)
- Block allocated GB / raw capacity GB as a percentage = 60-70%
- TBs/storage administrator ≥ 10 TB (maximum achieved is > 100 TB)
- Percentage of storage devices older than 36 – 40 months ≤ 20%
- Backup success rate ≥ 90%
- Snapshots for testing

IT executives can also reduce costs by better management of applications. Some key examples are as follows.

- Live application mobility capability = 100% on mainframe and Unix servers; ≥ 60% of Intel servers
- Application redundancy < 20%
- Limits on size of email inbox = yes
- Limit on age of e-mails in inbox = yes
- Single instance ERP in place enterprise-wide
- DC operations provides input into application hardware requests
- Selective use of cloud computing

On the network front there are multiple items that data center executives should also examine.

- Number of virtual ports to physical ports = 1.4:1 (minimum)
- For side vented Cisco units, use of baffling to create front/back airflow vs side to side
- Network utilization in the 30 to 40 percent range or better
- Unified computing
- Converged networks or software defined networks (FCoE, 10 Gbps, etc.)
- Age of routers and switches
- Ratio of network administrators to routers

Summary

IT executives should approach data center optimization as an opportunity to have each manager look closely at his/her area of focus and find innovative ways to reduce the cost of operations through new processes and technologies. This should not be a one shot deal but an ongoing annual effort that should be done as part of process improvement efforts or spring/fall planning. See Figure 4.

Figure 4. IT Transformational Approach

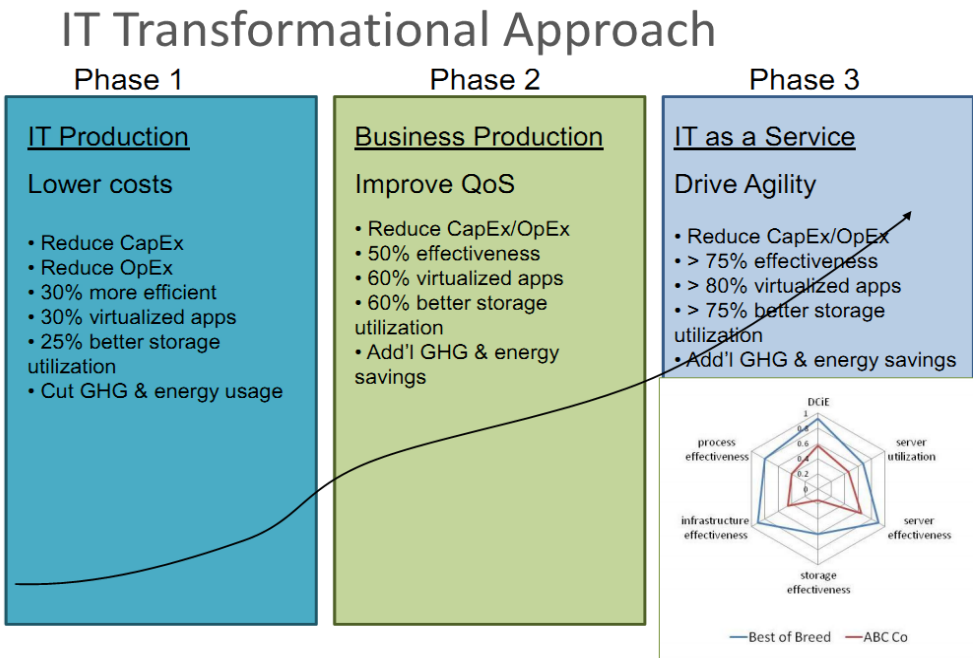

Moreover, if at all possible, IT executives should work with their CFOs and financial team to develop and implement a chargeback or showback system. Once business executives realize IT is not free, they will more readily work to trim costs.

There is every reason to believe that through detailed analysis **IT executives should be able to cut operations costs annually by a minimum of 10 percent per year over the next decade while still expanding capacity to meet the business workloads.** IT executives should also consider working with outside firms, like RFG or the GreenWay Collaborative, to assist with the baseline and workload analyses as well as "what if" scenarios.

RFG POV: Data center architecture and architects are needed to drive maximum efficiency and effectiveness within and across data centers. By using the KPIs, metrics and suggested practices outlined in this research report, IT executives should be able to benchmark their environments, perform a gap analysis, and then drive optimization and process improvement initiatives that can yield incremental benefits year after year. IT executives, working with facilities and finance staffs, should bend the IT cost curves so that each year they are taking at least 10 percent out of the total Opex costs and 15 percent out of the energy expenses. Moreover, IT executives should be able to find some of these projects that are self-funding (i.e., payback in less than 12 months), whose added savings can be used to fund the initiatives that have a longer payback period.

Additional relevant research is available at www.rfgonline.com. Interested readers should contact RFG Client Services to arrange further discussion or interview with Mr. Cal Braunstein.

46 Kent Hills Lane, Wilton, CT. 06897; (203) 429 8951;
http://www.rfgonline.com/; *Contact: inquiry@rfgonline.com*