VERTIV.

IMPROVING VISIBILITY INTO DATA CENTER OPERATIONS

Table of Contents

Executive Summary

Most IT organizations are subject to an 80-20 rule where on the first day of the fiscal year 80% of the budget is spoken for and only 20% is left for supporting new initiatives. Over time that number gradually increases until less and less funding is available to fund new initiatives (or increases in salaries). Most of the budget has to be spent on items of intangible benefit to the business such as software licenses, maintenance contracts, and energy. Recent advances in data center operations such as virtualization and software defined networking have made it possible to begin reducing that ratio, but they have also made some operational issues more complex to visualize and address. In general, managing data center operations has gone from being a simple to a complicated problem, with the potential to turn into a complex dilemma. As an analogy, driving the only car on the road on a straight, level grade is simple. Negotiating busy highway traffic at speed can be complicated. But operating an automobile truly becomes complex when driving blindfolded. Simple data center problems need to be automated so that more complicated problems such as planning can be addressed and complex dilemmas, such as retrofitting, avoided.

Three Types of Problems

Two academics from Canada, Sholom Glouberman and Brenda Zimmerman, have categorized problems into three groups: simple problems, complicated problems, and complex dilemmas [i]. Installing a new server in an existing rack is an example of a simple problem in the data center. Complicated problems are different. Determining which rack in the data center has adequate thermal and power capacity to accommodate another server, or planning a new data center are examples of complicated problems. High levels of expertise in a variety of fields are necessary for success, but new data centers have the advantage of at least being similar to each other. This is the difference between complicated problems and complex dilemmas. For instance, making sense of the data obtained from operating the infrastructure to improve processes and developing policies that align with business requirements can be a complicated problem as it requires special domain expertise. Retrofitting an existing data center while keeping it running is a complex dilemma because each data center has a unique mix of technologies, configurations, and tolerance for outages. Having experience retrofitting one data center helps, but does not necessarily guarantee, a successful retrofit because each different data center faces different issues.

Where do data center managers start to make improvements? They have a facility to maintain and infrastructure to keep running – they cannot afford to work on every opportunity. Even still, there is a cost to doing nothing. Every day, cash literally flies out of the building in the form of heat and other wasted resources that could be used to improve services and have the data center perceived as a strategic asset. One place to start is to better align the type of problem with the solution so that the data center can concentrate on the tasks that are most important to the business, such as offering timely services. That begins by categorizing, analyzing, and prioritizing problems, and determining the best approach to solving them.

SIMPLE PROBLEMS	COMPLICATED PROBLEMS	COMPLEX DILEMMAS
A detailed formula is all that is required to solve the problem	Approaching the problem scientifically is critical and weighing different solutions is a requirement for success	Detailed processes and the scientific method have limited application
Processes are tested to assure consistent results	Building one data center increases assurance that the next data center project will be successful	Retrofitting an existing data center offers experience, but will not be an assurance of success with other projects
No particular expertise is required, but an IT background increases the likelihood of success	High levels of expertise in a variety of fields are necessary for success.	Expertise can contribute, but is neither necessary nor sufficient to assure success
Detailed written processes produce standard results	New data centers are similar in critical ways	Existing data centers are unique and no two have the same equipment or configurations
The best processes give consistent results every time	There is a high degree of certainty in a successful outcome	Uncertainty of outcome remains; process shifts have to be managed and controlled
A single person can usually solve the problem	The complicated nature of the problem is beyond the ability of any one person and requires a team to solve	It takes a village to solve the problem and requires the involvement of multiple groups, often with different agendas and competing goals
Usually there is an optimal solution to a problem	Usually there are multiple solutions to a problem	No single solution to the problem can be found

[i] Glouberman, Sholom, and Brenda Zimmerman (2002). "Complicated and Complex Systems: What Would Successful Reform of Medicare Look Like?" Discussion Paper 8, Ottawa: Commission on the Future of Health Care in Canada.

Solutions to Simple Problems Need to Be Documented as Formulas

Routine tasks in the data center need to be on the path to automation so that IT can concentrate on activities that are more important to the business like time-to-market. Simple problems can be completed by following a recipe or a protocol. Recipes are specialized kinds of formula that can be followed again and again without requiring any specialized expertise. They can be tested often, refined over time, and can usually be completed by one person. Data center operators should strive to document the processes of solving simple problems with an eye towards improving them over time. They should aspire to make the documentation as easy to follow and consistent in results as recipes. At an organizational level, this requires an understanding of how frequently simple problems occur as well as the best process for solving them and an organizational commitment to refining the process over time by making it a part of everyone's job. Elegantly solving simple problems also requires data and analytics to gain actionable insights. For example, identifying the people best at solving certain problems can be achieved through measuring their progress and achievement over time, giving you the opportunity to reward them. However, effectively solving simple problems alone is not going to result in the kind of improvements needed to reach a nirvana-like state of optimization. IT organizations need to be more efficient in addressing routine problems, such as replacing a dead server, so that they can spend more time focused on complicated problems such as improving business time-to-market and even tackling complex problems, such as security.

Managing Data Center Operations has Become a Complicated Problem

In the past few years, technology evolution, converged infrastructure, and higher utilization though virtualization have increased the density in the rack. Products like NX-OS, Nexus switching, Unified Computing Systems, vBlock, and blade servers are more configurable than their predecessors. Rack density in most enterprise and hyper scale data centers has reached the point where significant opportunities exist to optimize these data centers' composition with an immediate benefit in efficiency and savings in cost. However, optimization is not a simple problem. According to our definition, it is a complicated problem because it has multiple possible solutions and requires a team of people to implement.

Using the Scientific Approach to Solve Complicated Problems

Complicated problems can be solved, but how? Fortunately there is a popular start-up methodology, based on the scientific method, which can be applied to complicated problems. It was developed in a workshop called Lean Startup Machine to help startups align their actions with customer needs. Although it was developed for another part of the technology business, the process can be adapted to data center operations. It is known as the validation board. The validation board is based on the understanding that you can't solve a scientific problem with false assumptions.

Define your hypothesis

In science, a hypothesis is a conjecture made on the basis of limited evidence but with insight into the problem, as a starting point for further investigation. Complicated problems tend to have multiple solutions. This is one of the reasons they are complicated. When you build your hypothesis it is better not to infer a solution. Starting with a solution might cause you to miss a better one. This is the most common mistake people make in solving complicated problems. There is even a folk aphorism for it: "fire, aim, ready".

So at this point it is important to make an assumption, but not to be wedded to any one solution. It is instead more productive to test your assumptions about the best way to improve the processes by modeling the data center and then conducting what-if scenarios. What if I were to change the ambient temperature of the data center by 5 degrees, 15 degrees, 25 degrees? What effect would this have on my costs and my reliability? Like any kind of experimentation, you need data to test your assumptions. This requires a greater insight into data center operations.

When you start this process you might be surprised to find how many assumptions you have made. However there is one that is paramount. This is the riskiest assumption. The riskiest assumption is the one that, if invalidated, will cause the entire project to fail. For example, the riskiest assumption in a cost containment project might be that the savings matter to the CFO. Business may be so good that time-to-market matters more than cost to the business.

Test and experiment

Once you have identified the weakest link among your assumptions you need to test it. The point of the experimental phase is not to try to test every assumption. The point is to test the riskiest assumption. If that assumption is invalid then none of the others matter and testing them is a waste of time. So the key question to ask is: what kind of data do I need to do to get this insight? You have to be disciplined about your thinking and not fall in love with your assumptions. Once your riskiest assumption is validated, it is time to test your next riskiest assumption and then the next until it is clear that your project is going to be an unqualified success. Only then is it time to execute the solution to a complicated problem.

Solving Complex Dilemmas

Complex dilemmas are problems in a class by themselves. There is no formula to follow and they are too multifaceted to yield gracefully to the scientific approach. There is really only one way to solve complex dilemmas and that is to see them as a series of simple and complicated problems. To use an old saying, when it comes to eating an elephant, the only way to do it is bite by bite. Breaking large problems into smaller pieces is a problem solving method common to many different disciplines. In engineering, the technique is known as modularity, where larger systems are built by combining smaller subsystems. The concept of modularity can be applied to solving complex dilemmas such as retrofitting a data center while maintaining service. Individual data center modules such as compute, storage, network, cooling, or power can be optimized separately to increase overall efficiency.

DCIM Capability Maturity Curve

Almost every data center is looking for ways to become more efficient—reducing power consumption, reducing or eliminating underutilized resources, deferring equipment purchases, and optimizing how people get their jobs done while maintaining reliability. Following these steps will help cut operating expenses and improve the consistent delivery of services. Insight into data center status is obtained when the input data to every step of the process is captured and visually displayed. When it falls out of bounds (beyond certain thresholds) then alerts and alarms are provided to signal the shift. Once the process is codified and kept from shifting, we can articulate policies that are tied to business processes. When policies are violated, they are tackled by addressing the problems created by the shift or by managing it as an exception. When all these work in harmony, we can look to automating by integrating to systems management software. The initial steps of aggregating data are quite simple, but analyzing the data to derive information and define or refine processes is a complicated task. The next steps of articulating policies around the operational processes and linking them to how business gets done require special expertise. The final step is quite complex. However, none of these would be possible without visibility into infrastructure performance within the data center. This is a necessary step to start the journey towards an optimized data center that is self-diagnostic and self-healing.

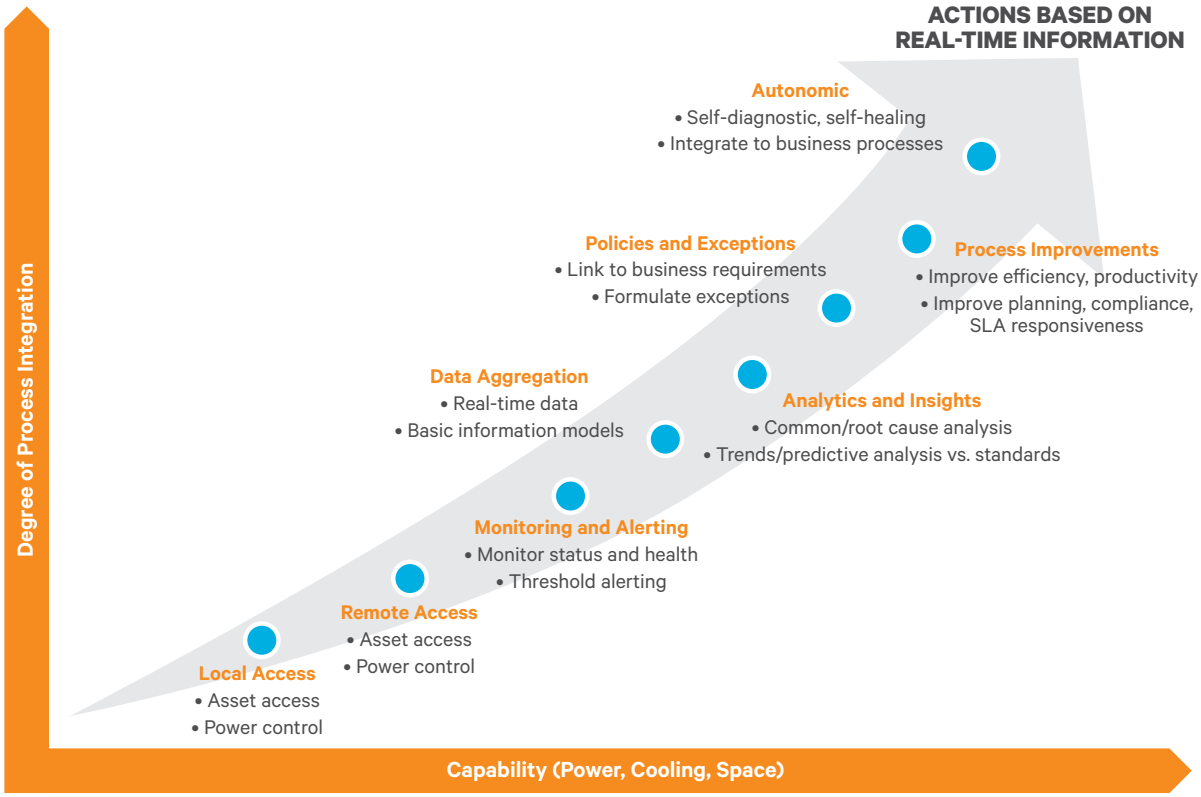

Visibility leads to agility

Being able to see ahead promotes IT agility. One example of a company that uses visibility to create agility is a financial service firm in London. It has a program it calls time-to-market-servers. At first, it buys equipment in bulk. It may acquire 500 servers and switches at the same time, hiring all the support staff and setting it all up in a data center before services are requested. If anybody in the organization needs computing that can fit onto those servers, the IT department simply relocates, renames, and sets it up as quickly as it possibly can. In this way, it avoids the delays traditionally associated with installing and procuring equipment. Everything is powered up and ready to go to market before the formal request. Insight into the future needs of the business allows the IT department to simply relocate, name and set it up as quickly as it possibly can. Having an accurate handle on data center resources allows it to respond in a timely manner to new requests.

Conclusion

Simple problems can be treated as formulas and refined over time. Advancing technology, changing business needs, and new business models, such as infrastructure-as-a-service, have all reached the point where high levels of expertise in a variety of fields is required to solve complicated problems. Taking a scientific approach of developing a hypothesis, building a test, and analyzing the results can allow data center management to apply a popular start-up technique to this problem, improving operations in ways that will most benefit the business. Complex dilemmas can only be solved by breaking them down into a series of simple and complicated problems. How closely a data center comes to being a strategic asset depends on where it sits in the DCIM Capability Maturity Curve. Driving the only car on the road on a straight, level grade is simple. Negotiating busy highway traffic at speed can be complicated. Operating an automobile truly becomes complex when driving blindfolded. In the data center, as in driving, negotiating the curve comes down to a question of visibility.

VertivCo.com | **Vertiv Headquarters,** 1050 Dearborn Dr. Columbus, OH 43085, USA

AV-29105 (R07/12)