Raritan.
Know more. Manage smarter.

Monitoring efficiently?

How to ensure your monitoring does what you need it to

DatacenterResearchGroup
A Division of DatacenterDynamics

Overview

How will proper data center monitoring assist owners and operators to manage through an upcoming period of enormous technological, corporate and legislative change?

While it is clear that most owners and operators have taken the mantra of 'you can't manage what you don't measure' to heart, it is also clear that the speed of change across data centers is fast exposing the inadequacies of some past practices. Accordingly, this White Paper has been written to indicate guidelines towards 'best practice' in overcoming these when monitoring in your data centers. It uses a survey of upper-level data center owners and operators across the United States to describe:

- Current practices in monitoring and reporting across this segment of US data centers
- Where improvements are looked for from current practices, particularly in meeting future trends that will impact on data centers, in the US and globally.

Key findings and implications:

- Monitoring in data centers responds to key industry concerns for energy consumption, availability and costs, maintaining an optimal environment for IT equipment and for coping with increasing total cost of ownership.
  - *These concerns are not going to go away and trend data indicates they will continue to impact adversely upon data center operation. Therefore monitoring needs to identify remedial actions swiftly and accurately as well as being able to cope with new and future requirements, most obviously new reporting and metrical standards*
- The vast majority of US owners and operators are monitoring energy consumption, temperature and humidity across all their facilities, usually on a 'continuous' basis.
- Energy efficiency, carbon output and power quality are also monitored but with greater tendency to do so irregularly and only within the primary/main facility.
  - *As data centers become more networked and delivery becomes focused more on the portfolio and less on the individual facility so monitoring systems will need to offer sufficient flexibility in monitoring new variables and/or extending into new facilities.*
- Reporting on data obtained through monitoring follows no set pattern between using the data purely to warn if a threshold is exceeded, reporting that is compiled manually, reports that are generated automatically and more sophisticated modelling and analysis applications.
  - *In some cases, reporting appears to be led by what the technology is able to deliver rather than what the facility may require. Reporting is critical in translating 'data' to actions and low confidence with this stage of the process emphasises the need for monitoring and reporting to be considered of strategic importance to the data center.*
- Levels of satisfaction with monitoring and reporting are at best 'subdued' although the linking of these processes to a formal process of continual improvement does much to improve satisfaction levels.
  - *Increasing satisfaction and confidence in data center monitoring has as much to do with organisational expectation and organisation of the systems in place as with technological excellence. 'Best practice' in relation to monitoring has enabled operators in this sample to integrate it more seamlessly into the requirements their data center is designed to fulfil and to the corporate processes that support and direct the data center.*
- There are concerns about the ease of integrating different monitoring processes to provide a 'whole of portfolio' picture, as well as the amount of data that can be generated through monitoring and the capability of existing systems to provide actionable information on cost savings as the pressures on budgets increases.
  - *To match the development of the data center towards a 'dynamic' facility where considerable flexibility is required in delivery, provisioning and planning, monitoring and its outcomes need to follow a parallel course. This means moving from simple to multivariate analyses, from simply issuing alerts to establishing control over a facility and moving from a reactive to a predictive role*

Why does best practice matter?

As the demand for data center services continues to increase so do the challenges facing data center owners and operators.

The pressures on the data center now come from many sources. They include energy consumption and availability, operating costs, data center capacity, downtime, cooling high density servers and the impact of legislation and regulation intended to reduce consumption. These issues (*Figure 1*) each represents significant concern among the US data center community.

The most critical aspect in dealing with these challenges is that all impact on each other. Deal with one and a new challenge emerges. The many layers of data center planning and operation will become more complex as the era of constant loads and availability is replaced by the requirement for 'dynamic' facilities to manage enterprise change, based on high data growth, scalable applications, blurred network lines, and variable loads.

For example: one of the key means of dealing with requirements for increased IT capacity is to deploy higher density servers. A snapshot of New York owners and operators in 2007 and then in 2011 (*Figure 2*) indicates that the proportion of racks operating at greater than 10kW/rack has increased from 4% to 11%. Over the same period the proportion operating at over 5kW/rack has moved from one in four to over one in two. As this process has unfolded, concerns about required environmental support for the IT (cooling, power protection) and the impact on the costs of operation have increased also.

This challenging environment and the acceptance that 'you cannot manage what you don't measure' has created strong growth in the technologies for monitoring and analysing what is going on in the data center. US market interest in technologies for monitoring energy consumption has increased by greater than 15% of sites year on year since 2007, and interest in temperature monitoring by a similar proportion.

So, given increasing pressures and the plethora of monitoring options available how can monitoring work most effectively in today's data centers?

Fig 1: Key Concerns of United States Data Center Owners & Operators
(Source: DatacenterDynamics 2010: N=554: Comments written in/grouped responses)

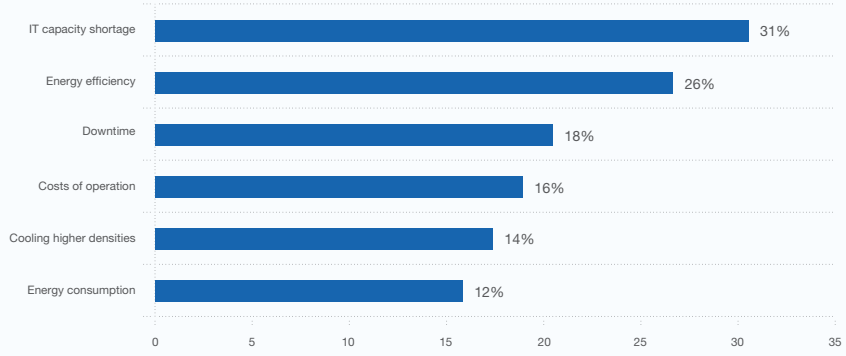

Fig 2: How Higher Density Racks Have Grown as % of All Racks
(Datacenterdynamics Samples New York Region: March 2008 (N=204) & March 2011 (N=175))

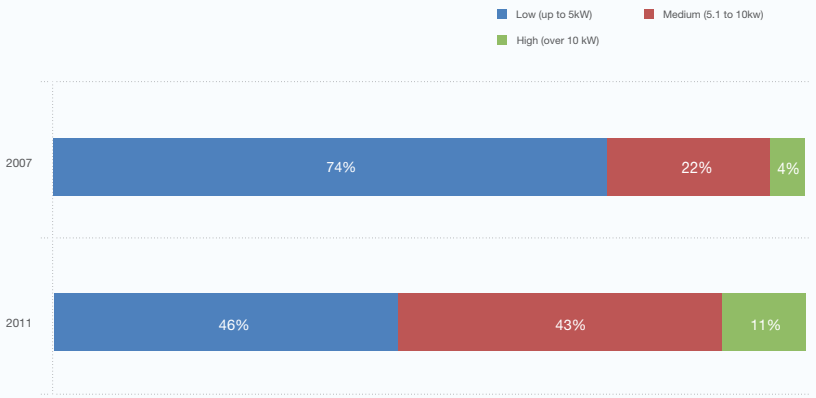

This White Paper presents the findings of research commissioned by Raritan and conducted by DatacenterDynamics Research Group in order to ascertain:

- The practices adopted in United States data centers for collecting information, analysing it and utilising it to improve the current and future operation.
- The improvements required of current monitoring, analytic, reporting and decision making practices.

Implications

The White Paper uses this information to propose a set of 'best practice' guidelines in relation to data center monitoring. 'Best practice' data center monitoring is important because:

- As the data center becomes more dynamic and sophisticated decisions will be based increasingly on data and less on 'hunch'
- You need to be able to trust the data you collect
- Ineffective monitoring deployment may represent missed oppotunities for greater facility efficiencies and cost savings
- It is important that you match monitoring and reporting to the requirements of your facility. In a number of cases observed in this research, less may be more!
- Increasingly data centers will be required to provide accurate reports to external agencies.

Why monitor?

What do data center operators hope to achieve from monitoring in their data centers?

Different reasons are given for different monitoring profiles, therefore, to answer this, the variables monitored in the data center have been divided on *Figure 3* into three distinct groups:

- Energy consumption
- Energy efficiency variables including the monitoring of carbon output, and
- 'Environmental' variables, including temperature, humidity, cooling air flow and pressure.

As the key concern among data center owners and operators, energy consumption is monitored for a variety of reasons: to maintain availability

Fig 3: Reasons for Monitoring Variables Within Data Center
(% weighted by number of racks: Multiple responses possible)

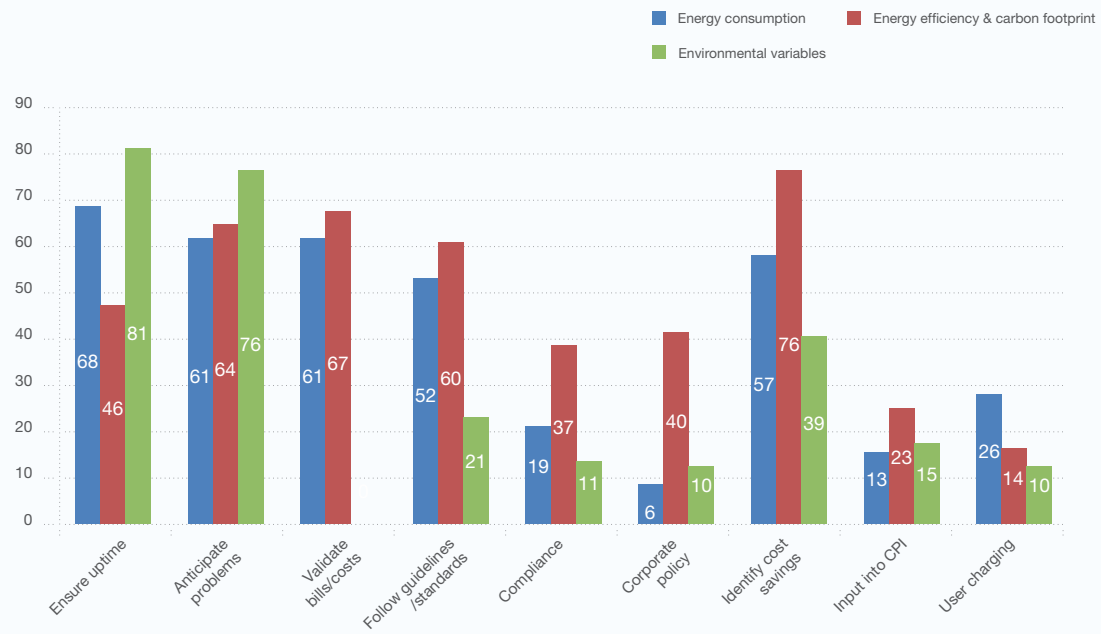

and to identify potential problems that may impact the availability, to validate energy bills, to follow guidelines and standards on energy consumption and to identify possible cost savings.

Monitoring of energy efficiency and carbon footprint is conducted for similar reasons with the key difference that it ties in also with corporate energy policy and compliance requirements.

The profile of environmental monitoring is focused more firmly on maintaining availability and warning of threats to that and, to a lesser extent, in identifying cost savings.

It is noticeable that the core monitoring of energy consumption, temperature and humidity are rarely linked to compliance policy, or continuous performance improvement - this creates the problem that there is nothing 'strategic' underpinning monitoring activity and this will reduce the scope for such monitoring to contribute to decision making and planning.

The requirements as identified above will not remain static and therefore again indicate the requirement for a monitoring and analysis system that can meet the changing needs of the facility and of the corporation. The recent history of monitoring in the data center indicates very clearly that as advances are made in the technological capabilities of systems so further advances are required as new questions emerge. Trends in power metrics follow a similar path, from consumption data to ensure uptime through efficiency metrics to reduce energy costs towards performance-based metrics that match energy input to IT output.

Reasons for monitoring reflect closely the key concerns of and challenges facing the data center industry in the United States, in particular those related to energy consumption, operating temperatures and threats to uptime, total cost of ownership and, to a lesser extent, to meet compliance and policy requirements.

Who monitors and what do they monitor?

Monitoring of any kind is virtually universal in this sample. 94% monitor energy variables (principally energy consumption, energy efficiency and carbon footprint) and 91% environmental variables (principally, temperature, humidity and cooling air flow).

This high level of monitoring is not surprising given the 'Fortune 500' profile of respondents – six in ten primary facilities are Tier III or Tier IV and 77% of facilities here are defined as mission critical to their organization. Possibly the only surprise in these results is the handful of facilities where monitoring is not done – these tend to be older, lower demand and lower redundancy facilities.

Fig 4: Proportions Monitoring Energy & Environmental Variables in Data Center
(% weighted by number of racks)

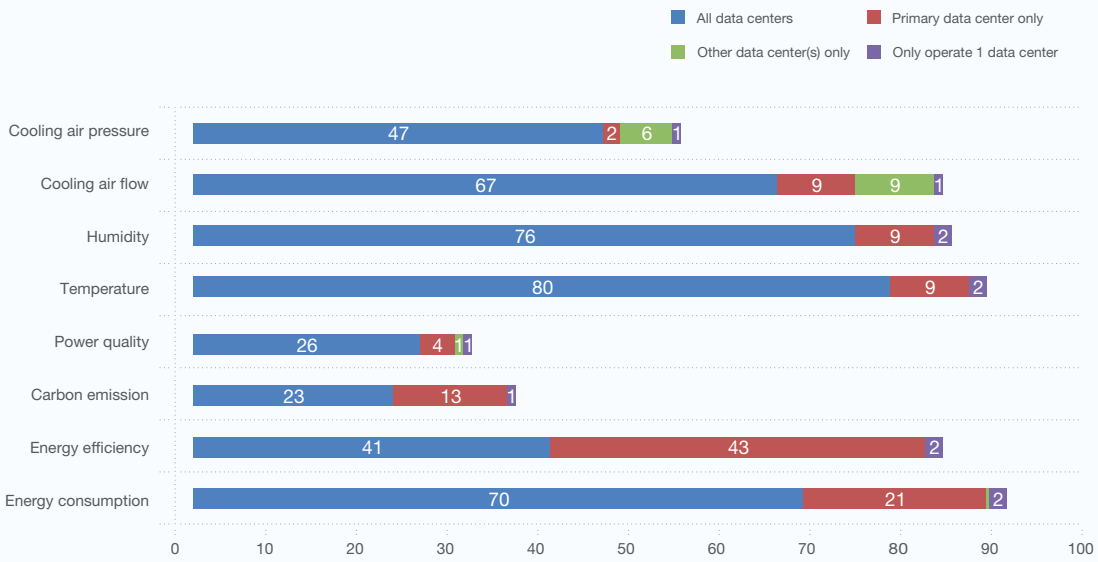

So which data center variables are being monitored across this sample of 'upper-quartile' American data centers owners and operators (*Figure 4*)?

- Monitoring of energy consumption is virtually universal while more than 80% also monitor energy efficiency. Fewer than 40% monitor carbon emissions and 30% measure other energy variables, principally related to power quality.
- Four out of five organisations monitor temperature, humidity and cooling air flow. Fewer monitor cooling air pressure.

In the context of a fast-growing and fast-changing data center workload two key missed opportunities emerge from this data.

The first 'missed opportunity' concerns *what* is monitored.

The key requirements of energy consumption, temperature and humidity are well catered for and a further set related specifically to monitoring 'efficiency' and carbon emission are growing in use as data center monitoring and reporting requirements develop.

"Power meters are changing the way in which data centers operate – they have allowed temperature set points to be increased safely, additional capacity to be added without risk, and ghost servers to be turned off or powered down.
They basically allow data centers to be more efficient …"

Herman Chan, Raritan

Current levels of monitoring of energy efficiency and of carbon are unlikely to be sufficient for the coming tide of regulation and of increasing energy costs, and of continuing efforts to eradicate inefficient use of energy in the data center. While legislators in the United Kingdom have taken the carbon tax principle furthest to date, other legislators in developed data center markets will not be far behind.

In an Uptime Institute survey of 500 data center operators (reported in May 2011) 45% indicate that reporting carbon emissions is 'important' to their organization.

The requirement for integration of energy and environmental monitoring is best illustrated by the steady rise of inlet air temperatures in the samples of American data centers drawn between 2007 and 2010 (*Figure 5*). The Raritan sample broadly follow ASHRAE recommendations in tracking temperature. However, the success of this process and the risks associated

Fig 5: Rise in Temperature of Air to IT equipment (Mean sample temperatures 2007 - 2010)

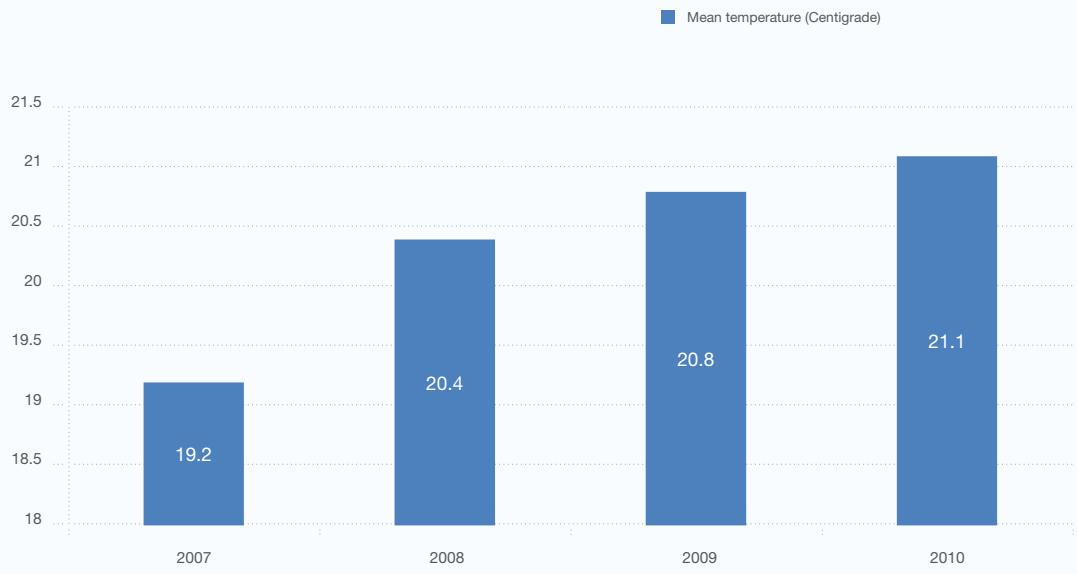

with it cannot be established without the integration, or at least the parallel analysis, of energy consumption and temperature monitoring information. Once the scope of monitoring is expanded into more complex measurement and modelling tasks such as identifying cost-savings, the requirement for analysis across different data center variables becomes more acute.

The second 'missed opportunity' concerns *where* monitoring is undertaken and again there are different practices according to what is being monitored.

The core environmental variables (temperature and humidity) tend to be monitored across all data centers in a portfolio – only a very small minority distinguish between their primary data centers and others that they operate.

- This widespread deployment may be accounted for by a greater familiarity with these monitoring systems.
- The fact that concerns about downtime due to thermal events have been longer established in data center thinking than concerns about efficiency.

The monitoring of energy efficiency and, to a lesser extent, of energy consumption as well tends to occur exclusively within the primary data center. Primary/main data centers tend to be the largest in the portfolio, to use the most energy, to produce the highest amount of carbon and therefore to be the first deployment in the rolling out of the relevant monitoring technologies. This replicates (and ran in parallel to) the roll-out situation in the early stages of the PUE metric.

As energy becomes scarcer, more expensive, and the need to report on consumption becomes greater so the systems deployed to monitor its consumption in terms of quantity and efficiency will be required in all data centers rather than just primary data centers. In this sample primary data centers account for around half of total portfolio space and IT capacity so extension to current monitoring systems will be fairly extensive (and possibly expensive) particularly if it involves replacing the whole system.

On the evidence of the research, the principle of being able to extend monitoring into new 'variables' within the data center (carbon tracking, for example) is inhibited by problems of deployment.

Only 15% of the sample are totally satisfied with the ease of deploying the required technologies to initiate or improve their data center monitoring. A higher proportion – 28% - are not satisfied. The level of satisfaction appears low for a system designed to work seamlessly in support of a mission critical IT system. Profile information indicates that the handful of operators who are 'entirely' satisfied are:

- Responsible for smaller portfolios (fewer than 3 data centers)

Fig 6: Satisfaction with Ease of Deploying Required Technologies (% organizations)

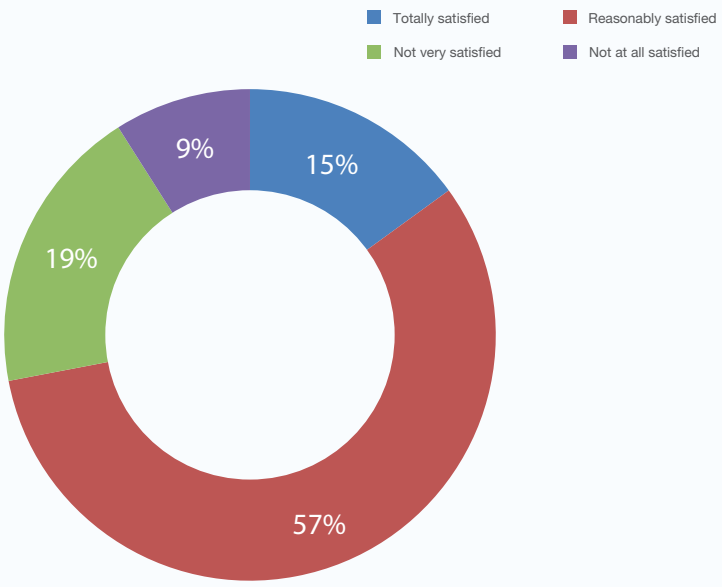

- Working from a set of defined policies and procedures in relation to monitoring and reporting within the data center

- Without exception linking their monitoring to a process of continuous improvement in their data center(s).

As the requirement grows for the 'dynamic' data center to be operated as a single 'live' environment and as data centers become increasingly networked and virtualised so the need to coordinate monitoring across variables will grow also. This again indicates the importance of a monitoring system which is flexible enough not just to extend across facilities but to coordinate across variables as well.

Implications

The first element of data center monitoring 'best practice' is that it is part of a strategic process which will ensure that monitoring at whatever level of technological sophistication meets the requirement of the facility. This will involve answering the questions of:

- What will be done with the information obtained – who will it go to? Who will act on it? How will they coordinate?

- What action standards will be set on the basis of the information, whether when setting thresholds for alert-based monitoring to setting inputs for planning scenarios?

- What levels of monitoring is required (in terms of frequency, placement within the data center, validation) and, since we are dealing with statistics here, what level of accuracy/ risk is considered acceptable?

- How can monitoring and reporting systems be 'future proofed' to ensure that they remain sensitive to the changing needs of the portfolio?

These key questions asked at the establishment or review of a monitoring system will avoid the major causes of dissatisfaction – of incompatible monitoring systems collecting data that is difficult to coordinate, or the failure to obtain a 'whole of data center' perspective as different sets of data find their way to different parts of the organisation, or the discipline of systems which generate far more information than can be analysed properly or usefully.

How are they monitoring?

It is self-evidently a critical element of best practice in data center monitoring that the technologies and systems deployed are able to deliver the data required for current and future data center management, and to do so with a level of reliability necessary to support the decisions made from it.

At the facility level, there are a number of decisions to be made as to the practical implementation of monitoring. These can be summarised as:

- Where in the data center do I monitor the variable? For 'environmental' variables – temperature, humidity, cooling air flow and pressure – there are logical points in the data center for measurement dependent upon how cooling is delivered to the IT equipment. There is greater variation in where power and energy-based measurement are taken as it courses through the data center – from the power feed or sub-meter, at the UPS, at the rack PDU, at the individual outlets of the rack PDU or at the CPU.

- How often do I take the measurement? Failure to match the frequency of measurement to what the data is actually required to do is the major cause of data 'overload' (where continuous readings are taken with no evident need for that frequency) and, conversely, for problems of confidence in the data obtained (usually where data is taken too irregularly to confirm the decisions that need to be made)

- When do I take the measurement? Given the fast-changing environment within most data centers, different readings will be obtained dependent upon load, local environmental conditions etc.

- To what extent do I need my readings validated? At the most basic level energy consumption readings are used to validate the facility energy bill or to determine how costs need to be shared between customers. Readings at multiple levels within a facility (for example, energy at the (sub-) feed, PDU and device level) permit the identification of unexpected variation in the monitoring process and improvement in the accuracy of the process. It is also a proven method of identifying sub-optimal equipment performance.

The frequency of monitoring and the extent to which the readings are validated at different levels distinguishes between three sets of variables (*Figure 7*):

- Energy consumption which is almost always monitored continuously and validated by at least two sets of readings.

- The set of environmental variables which tend to be monitored continuously at a single point in the data center.

- Monitoring of energy efficiency, of carbon footprint and power quality which are intermittent, monitored on an 'as needed' basis and, with the exception of energy efficiency, largely reliant on a single set of readings.

In terms of satisfaction with the accuracy of results generated by monitoring (*Figure 8*), mostly operators are reasonably satisfied (76%). The minority who are entirely satisfied conduct validation through multiple monitoring methodologies and monitor on a continuous basis, and are able to link their readings into pre-set criteria for action.

Implications

Whether deploying a new monitoring system or refreshing an existing system, there are a number of steps to be considered:

- Government entities and industry advisory groups have focused increasingly on monitoring within the data center as part of the push towards greater energy efficiency. In terms of established standards there are a number of good starting points including:

  - *Energy Star monitoring and reporting requirements for participants in their energy efficiency programs*

  - *The US Department of Energy Practitioner Program*

  - *ASHRAE guidelines as to operational standards*

  - *The Green Grid.*

Fig 7: Monitoring Diagnostic Summary
(Averages of frequency of monitoring & validation of results)

- *The European Union Code of Conduct for Data Centers*

- It is important to establish benchmarking data as the starting point for measuring improvement or to identify accurately the impact of changes to the data center.

  - *As with any statistically-based program, a key task will be to identify sources of error in the collection of data through data validation and through the full analysis of variation in data that cannot be explained through existing application of the data.*

- When choosing the monitoring system to be adopted (or refitted) be aware of some of the factors that are causing concern in this sample:

  - *Accuracy claimed of the readings. For carbon reporting requirements and any situation in which the data obtained feeds into direct financial dealings, for example charging, incentive programs accuracy levels of +/- 5% may be required while to meet ISO/IEC specifications a +/- 1% billing grade accuracy is likely to be a requirement. The level of statistical accuracy will depend substantially on the levels of risk defined as acceptable for the readings obtained.*

  - *A system that is open, flexible and able to work within wider IT management systems, to incorporate new modules as new requirements emerge - for example, moving from reading of efficiency to the identification of cost savings, or moving from basic 'warning' systems to more control based systems (to predict when problems will occur).*

  - *Of particular importance is the evolution of key data center efficiency and performance metrics which as they evolve will place new demands on what, where and how often to monitor in the data center. For example, one newly proposed Green Grid metric – ERE (Energy Reuse Effectiveness) – will entail the measurement of energy use at the 'control volume boundary' to calculate a rating of how much heat generated within the data center is re-used outside it.*

  - *Work back from the outputs/decisions that will be made on the basis of the information, through the reporting that will be required to reach those decisions back to what the technology is able to deliver. This will avoid generating too much information, or too little, a danger if you work forward from what the technology can offer and allow that to shape your requirements.*

"We do lots of monitoring but my problem with this is that there is little to drive information collected towards knowledge"

Financial institution, Texas

Fig 8: Satisfaction with Accuracy of Measurements Provided (% organizations)

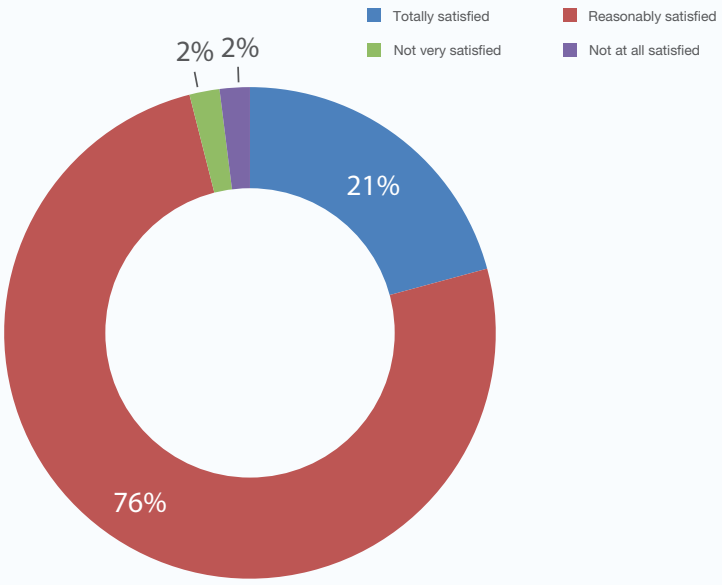

What do they do with the information?

While it is possible to detect a relationship between the monitoring profile described in the previous sections and facility/corporate needs, this relationship breaks down when the information gained is analysed. This is important since analysis and reporting are the critical stage between measurement and management. The level of reporting across this sample appears to be related more to technological capability than to requirement, and subdued levels of satisfaction suggest that not everyone has followed the principle of using requirements to define technology rather than vice versa.

The level of reporting on key data center energy variables (*Figure 9*) shows that energy consumption is reported through alerts, manually and as an input into performance evaluation. It should be noted that almost one in five organizations that monitor carbon emission appear to have no reporting structure for it although in this sample this is partly a reflection of the recency of deploying these systems.

Less analysis is done on environmental monitoring data – reporting on temperature is almost entirely restricted to alerts, and reporting on the other major variables of humidity and cooling air flow/pressure reported through a mix of manual and automated reporting.

One in four organizations are totally confident in the decisions made on the basis of information provided by monitoring (*Figure 10*). Levels of satisfaction are lower for the clarity of output and for the speed of identifying problems and this will obviously undermine confidence in decision-making. **Once again, the highest satisfaction band are those with greater integration of monitoring systems, using data to build into a continuous improvement program and which have adopted a monitoring and reporting policy that is flexible enough to meet their changing facility requirements.**

Implications

Again, the inputs into 'best practice' can be traced from the reasons for dissatisfaction in reporting in the sample. The end-game of monitoring is best facility decision making, either in speed of response to emerging localized problems or in fine-tuning the operation and planning for the whole facility. Any break in the whole chain which goes before it will be shown in sharp relief at the reporting stage, therefore 'best practice' here might include:

- Ensuring the accountability of monitoring system deployment

Fig 9: Reporting & Analysis Conducted on Data Collected by Monitoring within Data Center
(% organizations monitoring variable: Multiple responses possible)

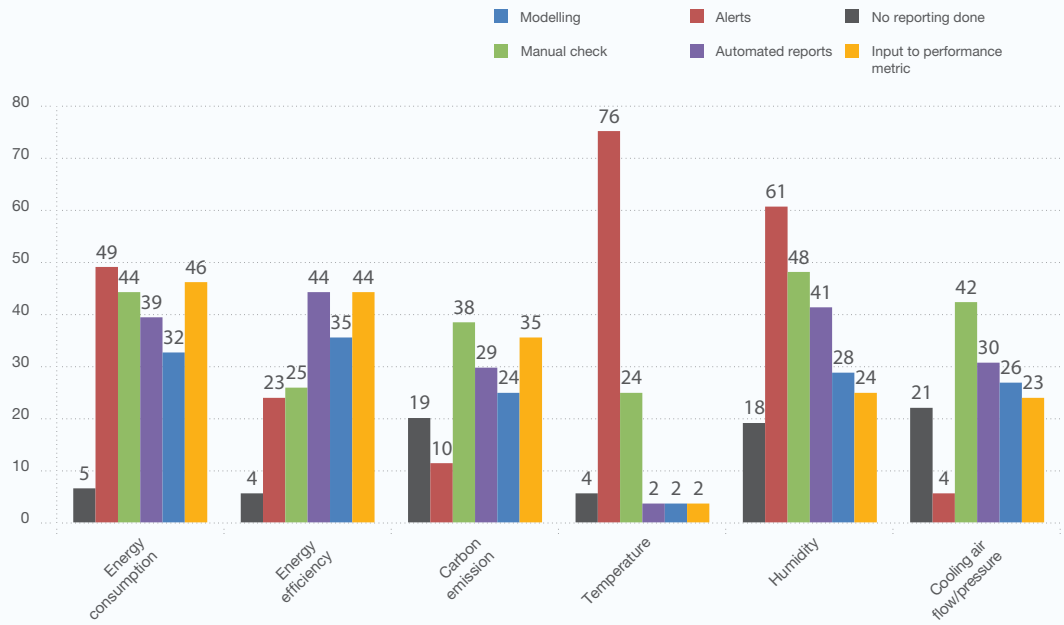

"Storage - we are monitoring a lot at this point, and we plan on monitoring more; however, a lot of the programs we review use databases and storage that are not capable of quickly processing the information to a state where we can use it"

Telecoms, Mid West

to the operational standards and expectations set for the facility

- The translation of these standards and expectations into the practical 'when, where, how often, how much validation?' of monitoring deployment (to avoid too much or too little data)

- Basing reporting practices on reporting requirements, whether these are internally-driven (for example, tracking change management within the facility) or externally-driven (through requirements of regulation and compliance) and avoiding the common happening whereby an extra stage between data collection and decision making is created as the data is re-shaped into a usable format

- Recognising the links that enable basic standards of analysis and reporting to evolve into more proactive and valuable measures that allow data center personnel to exercise greater control over the facility, rather than simply reacting to problems. The emerging requirements of flexibility of load and support within the 'dynamic' facility require the integration of all elements of management within the facility across resource availability, capacity management and load shifting, IT environmental control, and power quality.

Fit for Purpose?

It is evident from the research that different organisations and different data centers have different needs of monitoring, analysis and knowledge: this is based on their mission critical status and their profile characteristics. Given these differences within the sample there are some clear indications of where current monitoring practices might act as a pointer to best practices in the future.

Satisfaction with current monitoring procedures (*Figure 11*) can be described at best as 'subdued'. The research indicates that following suggested best practice guidelines will increase the likelihood of higher satisfaction and that the 20% of organisations which have already linked their monitoring – reporting – decision making process to some form of continuous improvement have

Fig 10: Satisfaction with Monitoring Outputs (% organizations)

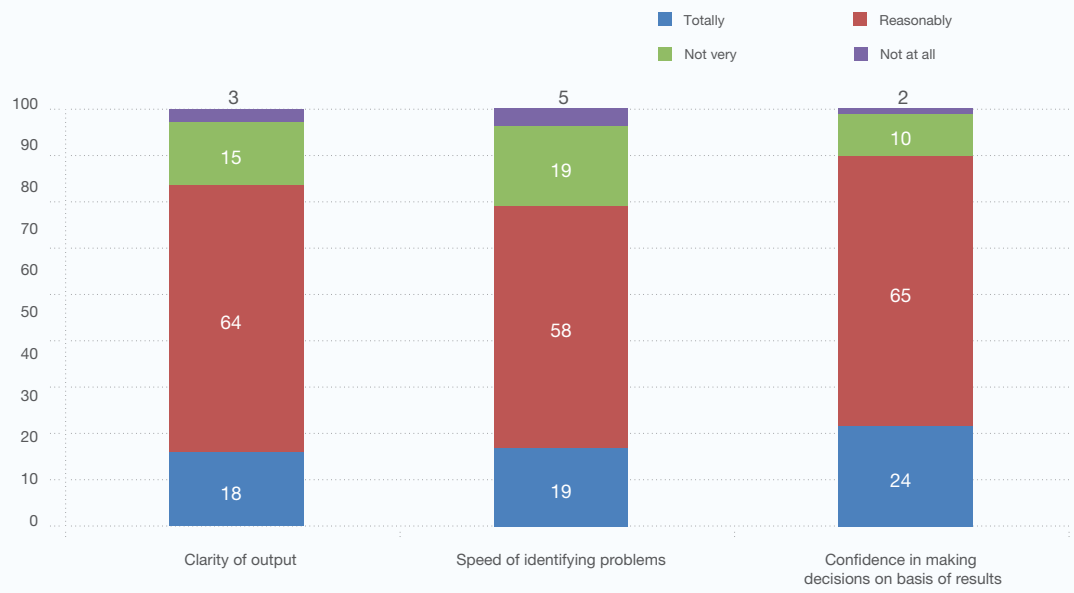

increased their satisfaction levels by doing this (*Figure 12*).

And it is clear also that this trend cuts across advances in technology and applies equally to relatively simple 'alert-based' systems, through more sophisticated power and energy monitoring systems as far as 'whole of portfolio' DCIM and DCEM systems since it is as much about usage and expectations as technology.

A major cause for dissatisfaction is that deployment of technology is not moving as fast as corporate and facility requirements are moving. (*Figure 13*) In particular current monitoring and analysis systems appear stretched when answering the critical questions of:

- How and where can I save energy?
- How can I accurately track costs and identify cost savings?
- How can I track power availability (so I can ensure that power within the facility goes where it is needed while reducing wastage where it is not needed)?

What is perhaps most striking about the improvements required to these instrumentation functions is that they come almost entirely from organisations that are attempting already to monitor and analyse on the basis of these outcomes. This indicates that the speed with which the requirements for actionable information and analysis has grown and changed has not been satisfied by the original technologies deployed. The monitoring and analysis process has not been 'future proofed'.

Fig 11: Aggregated Satisfaction with Monitoring Process
(% organizations)

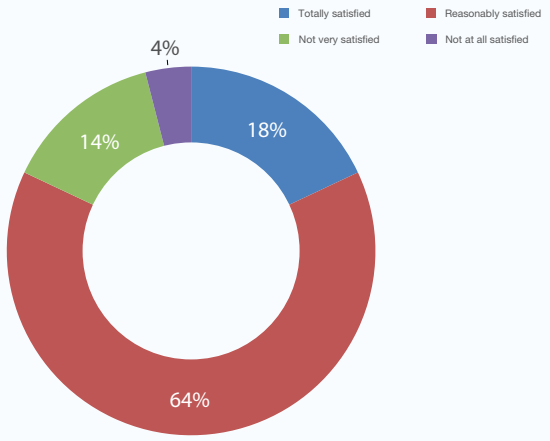

Fig 12: Linkage of Monitoring to Continuous Improvement Program
(% organizations)

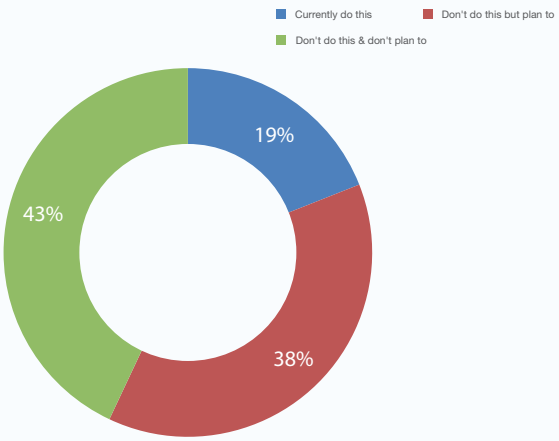

Fig 13: Instrumentation Functions Requiring Improvement
(% organizations)

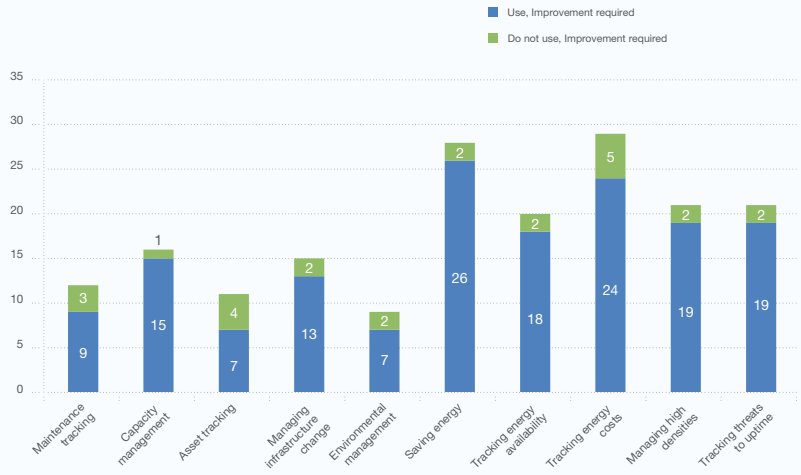

About the Research

The research findings presented in this White Paper are based on interviews with 75 data center owners and operators collectively responsible for almost 750 data centers (of 20 racks or larger). The sample includes representation from all relevant private and public industry sectors and of personnel from facilities management, IT/networks/applications management and engineering/support services within those organizations. The sample is comprised very largely of representation from America's 'Top 250' data center owners and operators (as defined by DatacenterDynamics).`

[Note that due to rounding percentages on some charts may not add exactly to 100%]

About Raritan

Raritan is a proven innovator of power management, infrastructure management, KVM and serial solutions for data centers of all sizes. In more than 50,000 locations worldwide, our hardware and software solutions, including intelligent PDUs, energy management software, award-winning KVM-over-IP and Serial-over-IP access products, provide IT and facility directors, managers and administrators with the control they need to increase power management efficiency, improve datacenter productivity and enhance branch office operations. Raritan's power management solutions won five key industry awards in 2010 and were recognized by the EPA for contributions to the agency's data center initiative.

Raritan. Know more. Manage smarter.