

Real-time Operational Intelligence (RtOI) for the Data Center

Practical Guidelines for Managing the “Big Data” in Your Data Center

Modius Turns Big Data into Operational Intelligence for Solving the Real-World Problems of Today’s Modern Data Centers



Table of Contents

- Introduction 2
- A Little History 2
- A Case for RtOI in the Data Center..... 4
- The Benefits of RtOI 4
- The Modius approach to RtOI 5
- Big Data, Big Problem 6
- Solving the Data Problem 7
- Summary of Modius’ Recommendations..... 10
- About Modius 11





The modern data center represents one of the most challenging management problems facing organizations as we head into an uneasy future of rising operational costs and increasing demand on IT infrastructure. Not only is there a diverse mix of heterogeneous equipment that must be managed as a single entity, the data center typically has two groups (Facilities and IT) that need to work in concert to efficiently manage this dynamic operational environment.

The answer to this management problem in the Data Center was software, specifically Data Center Infrastructure Management (DCIM) software designed to help Facilities and IT personnel better manage this disparate mix of equipment and systems. In 2009, Gartner Research and Forrester Research agreed on Data Center Infrastructure Management (DCIM) as the official name of this emerging software category and all vendors with solutions addressing any or all of these data center problems were grouped together under the DCIM umbrella. But not all DCIM solutions solve this management problem in the same way. Some vendors build models to explore the impact of changes, and others, like Modius, believe that the only way to truly manage a data center is to continuously measure its performance over time.

A Little History

The data center is a heterogeneous mix of equipment designed to house and support IT infrastructure in an environment that is primarily focused on reliability and risk management. In years past, insuring that this environment would provide adequate risk management and support future IT growth, led to over-provisioning the support systems for power, cooling, and physical space. As capital expansion and electricity costs rose in recent years, this over-provisioning significantly increased data center total cost of ownership. Organizations realized that without better management tools, these costs would adversely impact organizational goals.

The over-provisioning problem wasn't the only driving factor for creating the software category we now call Data Center Infrastructure Management (DCIM). Other problems in the data center that were solved with a component of DCIM include the following:

- Poor Documentation - (Asset Management, CMDB)
- Long MTTR and Unplanned Outages - (Alarm & Alert Monitoring)
- Inconsistent Process & Procedures - (ITSM, Workflow)
- Poor Capacity Planning - (CFD, Simulation Software, Advanced Analytics)
- Inefficient Operation - (Power & Cooling Management, VM, Load Balancing)

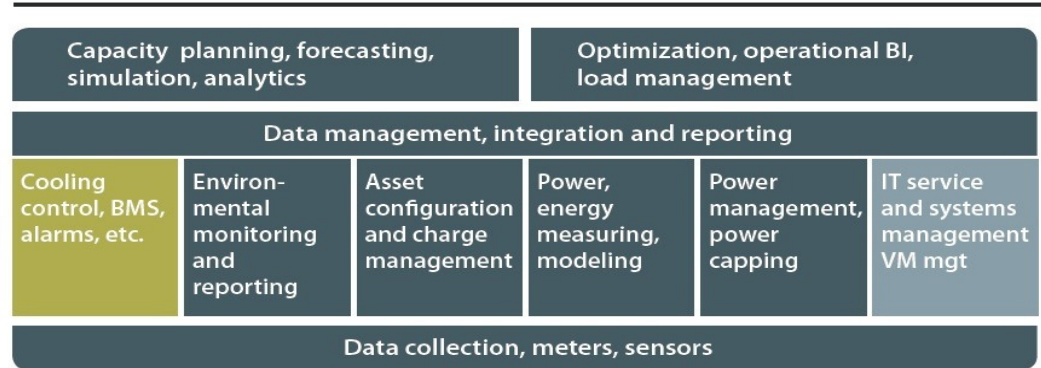
This list of data center problems spawned a variety of vendors offering solutions to one or more of these issues. Initially, some vendors attempted to provide an all-inclusive suite to address all of these data center management challenges, and other vendors elected to deliver a more focused solution to address a subset of the total data center



management problem.

The 451 Group's definition of the DCIM Stack recognizes that data collection is the foundation of a successful DCIM implementation. ¹ Real-time Operational Intelligence (RtOI) for the Data Center is the result of the Modius approach for DCIM -- continuously measuring data center system performance and transforming this data into knowledge for making better management decisions.

The 451 Group DCIM Model



So where does Modius fit in the DCIM model? Modius OpenData was designed to provide an open platform for collecting measured data from all the infrastructure in the data center across a network of data centers. Modius differentiates itself from the rest of the DCIM vendors in the following ways:

- Support for both new and legacy Infrastructure
- Designed Specifically for Real-time Data Collection
- Highly Scalable System Architecture
- Self-service Implementation, no Vendor Services Required
- Comprehensive Data Management
- Integrated Real-time Operational Intelligence Tools

The original focus on scalable, real-time data collection has paid significant dividends for Modius' evolution as a DCIM vendor. Because Modius OpenData was built specifically for real-time data management, it doesn't have the issues facing other DCIM vendors that don't have data collection as a foundation, or that simply skip data collection altogether and just use manufacturer's plate values for estimating power and cooling capacity. Because Modius can get this data in real-time, correlate it to produce useful information, integrate data mining tools, and then deliver the results to a mobile workforce, we are uniquely qualified to provide enterprise-grade Real-time Operational Intelligence (RtOI) for the data center.

¹ Andrew Lawrence, Research Director, The 451 Group, Uptime Institute Symposium, Santa Clara, CA May 2011



A Case for RtOI in the Data Center

By definition, “Operational Intelligence is a category of real-time dynamic, business analytics that delivers visibility and insight into data, streaming events and business operations. The purpose of OI is to monitor business activities and identify and detect situations relating to inefficiencies, opportunities, and threats and provide operational solutions. Some definitions define operational intelligence an event-centric approach to delivering information that empowers people to make better decisions. In addition, these metrics act as the starting point for further analysis (drilling down into details, performing root cause analysis — tying anomalies to specific transactions and of the business activity).”²

According to the Aberdeen Group, organizations with Real-time Operational Intelligence are 29% more likely than those without to get pertinent information they need in time to apply it to deadline driven decisions.³

Clearly, data center operators can benefit from getting the information they need to make better operational decisions -- especially when dealing with decisions that might have catastrophic consequences. For example, if they were using a DCIM solution that was only collecting power data once per day (or never) and they were using those figures to determine if there was sufficient capacity to add another server they could overlook meaningful changes in conditions. One team member could identify a location for adding servers, not knowing that another staff member on the night shift added two more blade servers to that rack. The first staff person wouldn't know that plugging in the servers would blow the circuit breaker and shut down the e-commerce application because the measured power information was stale or not available.

There are other examples that make a case for Real-time Operational Intelligence in the data center. In that same Aberdeen Group report, they state that RtOI adopters are more likely to achieve operational success because 70% of these organizations have an open exchange of operational data across business functions. In the data center, having a central repository of real-time data with mobile access by all personnel, facilities and IT, can significantly improve service delivery, collaboration, and boost key performance indicators like MTTD and MTTR.

When it is all about making the right decisions in a critical environment like the data center, you can't rely on stale data or theoretical models, you need RtOI.

The Benefits of RtOI

By capturing data in real-time, Modius offers a breakthrough solution that allows data center operators to apply Real-time Operational Intelligence to achieve energy cost savings and maximize the efficient use of data center capacity in several important

² Wikipedia, Operational Intelligence, Last Edited - July 2013

³ Aberdeen Group, Real-time Operational Intelligence: There's No Time Like The Present, June 2013



ways:

- First, energy costs can be dramatically reduced. OpenData enables optimization adjustments to key data center equipment – CRAC units, HVAC systems, chillers and generators – that can translate into substantial cost savings and rebates.
- Second, data center capacity that appears to be maximized is often simply poorly configured. OpenData provides critical information on the myriad microclimates within the data center, helping facilities professionals to reorganize equipment to harmonize the IT load. This prolongs the life of existing data centers and allows companies to defer or eliminate CapEx investments in new data center facilities.
- Third, the holistic view on data center health and reliability that OpenData provides helps impart peace of mind that facilities management and IT executives alike can appreciate. Modius allows data center teams to take a proactive approach to managing the data center by identifying conditions in real-time that have drifted from their original design, instead of perpetually reacting to emergency situations after they occur.

The real-time information OpenData provides is essential in managing power capacity in today's constantly changing data center environment. Simply put, Modius helps to reduce the cost, and stress, of managing the enterprise's most volatile environments.

The Modius approach to RtOI

In general, Modius has built its approach to RtOI from the ground up by providing a highly scalable architecture for real-time data collection across all data center infrastructure. Depending on the issue you are trying to address, Modius recommends that the data center operators start collecting data continuously throughout the day (i.e. approximately once a minute), and then increase or decrease the frequency, device coverage, and comprehensiveness (i.e. the granularity of the instrumentation points) over time.

This recommendation is based on the fact that a continuous sample rate is absolutely necessary to understand important fluctuations during the day. Even if only relatively few instrumentation points are being used to record data center performance, continuous collection allows more insightful correlation with other factors such as outside air temperature. Once performance adjustments are made, additional instrumentation points can then be added to provide new insights.

In general, the more granular the data, the easier it is to perform trend analysis and detect correlations. But more importantly, to increase efficiency, the data center operator also needs to identify transient phenomena – such as hot spots or cold spots, or temporary compute loads – that impact efficiency during the course of the day. By capturing data continuously throughout the day, operators increase their ability to make incremental adjustments to their infrastructure that drive increased efficiency.



Once you establish your baseline data for the areas and systems you are monitoring, you can set your high and low thresholds for normal operation based on real data, not on theoretical models, or manufacturer's plate values for normal operation. Modius OpenData will continuously poll these devices, check for threshold violations, monitor KPI's, and intelligently alert data center personnel of any possible issues. Because OpenData is web-based, any user with access to an Internet Browser can receive the threshold violation and access the information in OpenData to fix the problem quickly.

For example, let's say that monitoring your power consumption was showing that Fridays were especially heavy due to additional CPU Cycles for payroll. Based on this knowledge, we set our power monitoring thresholds higher and lower the room temperature to compensate for the additional server heat output every Friday. But this Friday, payroll doesn't run! OpenData can recognize that power consumption hasn't increased as predicted and can alert in real-time that the air conditioning temperature can be increased accordingly, thereby saving the cost associated with over-cooling the room. The data collected by OpenData also provides the Operational Intelligence needed to set the optimum temperature based on correlating historical data like outside temperature, humidity, server load, and power consumption.

Big Data, Big Problem

In a recent article published in Mission Critical magazine, the author claims that, "to date, vendors have approached DCIM as a hardware problem, offering a variety of specialized devices and appliances as solutions. But DCIM is not a hardware problem, it is a data problem."⁴ At Modius, we couldn't agree more. Your data center produces a mountain of data, every day, from a wide variety of sources:

- Facilities Infrastructure - Cooling & Power Infrastructure
- IT Infrastructure - Servers, Telco, Storage Systems
- Software Management Solutions - BMS, ITSM, EPMS, CMDB
- Security and Access Control Systems
- Other Data - Weather forecasts, financial systems, CMS

The volume, variety, and velocity of data today are rapidly growing, and many data center managers are finding that their confidence in data – and in data-driven decision-making – is dwindling. There are a few discrete functions related to the "data problem" facing today's data centers: Collection, Normalization, Correlation, Management, and Visualization. Each of these functions has numerous strategies for implementation, but some methods are better than others.

⁴ David Schirmacher, The Problem of Stranded Data and DCIM, Mission Critical Magazine, February 2014



Solving the Data Problem

It is important to understand how your DCIM vendor has addressed each function for resolving the data problem in your data center.

1. Collection: Data collection from networked devices in the data center isn't particularly hard to do. The problem is that tons of legacy equipment on the facilities side aren't network enabled, and you need data from these devices to get a complete picture of how your data center is functioning as a whole. Important questions to ask your DCIM vendor about their data collection capabilities include the following:

- How do you support data collection from non-networked devices?
- What device protocols do you support natively in software, (e.g SNMP, BACnet, Modbus), ect.)?
- What is your fastest polling rate, (seconds, minutes, hours, daily)?
- Is the polling rate variable by device or type of device?
- Is there any limitations to the number of data points I can collect and store?
- How do I setup a new device for collection?
- Can I setup a new device myself, without Vendor assistance?
- Can you collect data from other applications?

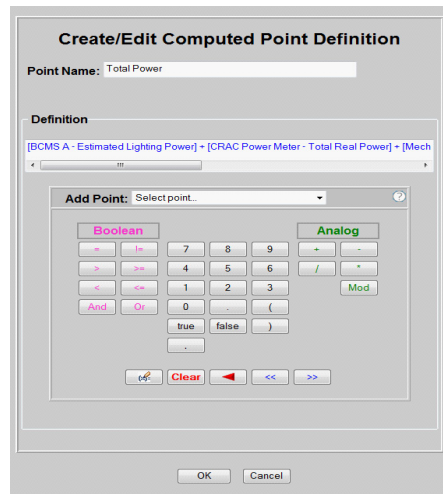
In a recent speech at the Open Compute Project Summit in San Jose, California, Facebook's VP of Infrastructure, Tom Furlong, reviewed Facebook's on-going DCIM project. One of the most important lessons learned while implementing DCIM at Facebook was to pay close attention to discrepancies in the latency of data collection between systems. If your DCIM solution only collects power data every hour, it is impossible to correlate it accurately with IT load data that is collected every second. You need to make sure that your DCIM solution has the capability to gather data points from the facilities infrastructure with the same speed that things are changing on the IT side of the house, so these data points can be correlated to make accurate comparisons.⁵

Modius, by design, has solved the issues with data collection by building a highly scalable platform for collecting data points from both networked and non-networked equipment and other management applications found in the data center. Using a library of standard device templates, you can quickly and easily set up a new device, define the data points you want to collect, set management parameters like polling rate and storage rules, all without any vendor assistance. This self-service capability eliminates the need for costly professional services and can significantly improve your return on investment.

⁵ Tom Furlong, VP of Infrastructure, Facebook, OCP Summit V, San Jose Convention Center, January 2014



2. Normalization: Another important lesson learned at Facebook was that data quality plays a critical role in determining a successful outcome. Mr. Furlong used an example of trying to compare total power between facilities when one UPS was providing data points in Watts and the other UPS was reporting in Kilowatts. The need to normalize this unstructured data requires the ability to apply math to a data point on a conditional basis. Without the ability to normalize the data to a defined data structure, these comparisons would not be possible.



Modius provides the ability to normalize each data point or a collection of data points by applying math equations and boolean logic to actual or newly created “synthetic” data points using a simple to use, calculator-style interface. In addition, you also have the ability to use this calculator to create Key Performance Indicators (KPI’s) from this unstructured data.

3. Correlation: In order to compare data points from different devices and systems, the data must be correlated so these associations can be viewed in context. One basic correlation is time and date stamping, but others include understanding the relationship hierarchy of these devices and how they are connected as part of a system (i.e., Power Chain).

Modius provides both time and date stamping of the data points collected as well as the relationship hierarchy of the device producing the data point. For example, we collect the metadata that defines where a device is located, so data points for PDU’s in a row of racks can be compared with another row in the same location.

4. Management: It’s hard enough collecting data, but once it is collected it must be managed properly to insure that the right metrics are available to the right people at the right time to provide the Operational Intelligence needed to make sound decisions.

Modius provides a complete set of data management tools to control access to the information collected, triggers for alerting key personnel that KPI thresholds are about to be reached, and automatic archiving or purging of older data that is no longer relevant.

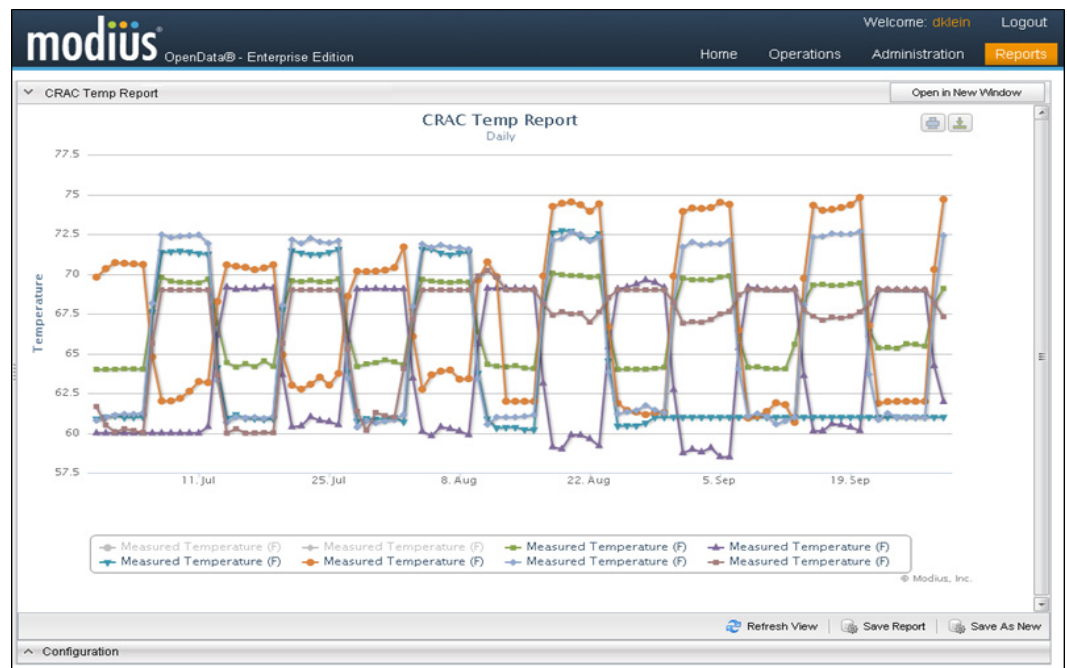
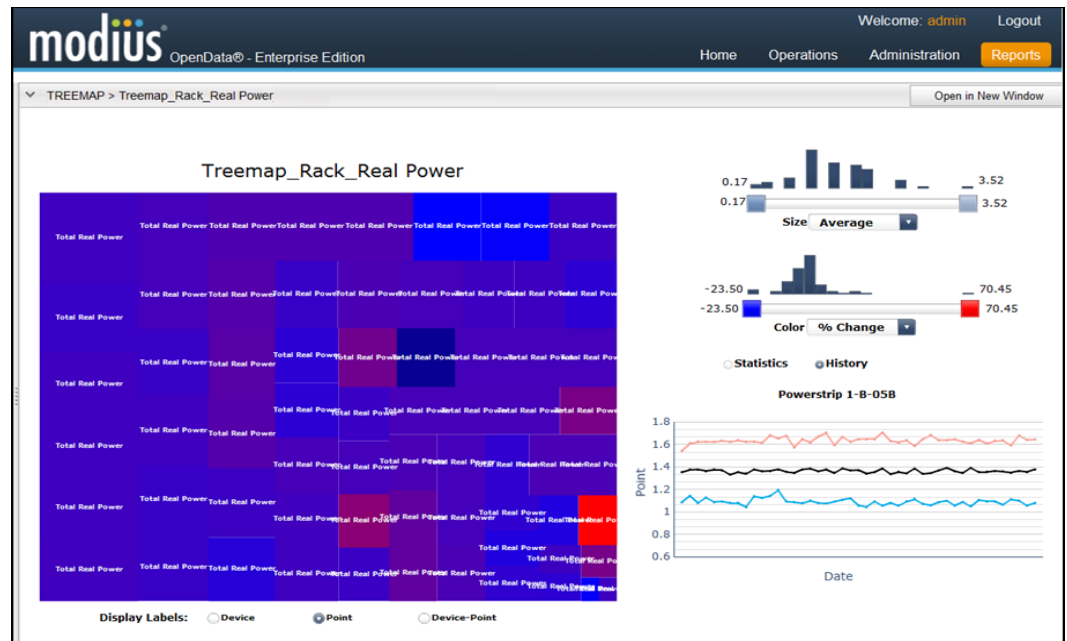
5. Visualization: It doesn’t do you a lot of good to collect data if you don’t have a robust set of tools to analyze the data in ways that provide meaningful insight into how your data center is operating and how changes can potentially improve its efficiency. Data visualization capabilities range from simple reports to complex interactive dashboards with data feeds from multiple sources. These tools need to be flexible, easy to use, and provide a rich set of user-configurable visualization options for presenting the data in



the most meaningful way.

Modius provides a complete set of data visualization tools, report generators, advanced charting and graphing capabilities, and BI dashboards to provide Real-time Operational Intelligence for solving real issues in the data center. Because Modius OpenData was designed as a self-service application, it comes with a comprehensive library of standard reports, charts, graphs, and dashboards that can be easily configured to address your specific needs.

Below are some examples of interactive charts, and dashboards that have been built using OpenData's BI Analytics package.





Summary of Modius' recommendations

Modius offers three guidelines for implementing RtOI for your Data Center, and in turn has designed its solution to meet these requirements:

1) Focus on the data

Modius recommends that data center operators start by implementing a solution that was designed to capture and store data in real-time. A significant number of DCIM solutions haven't solved this problem or do a poor job collecting data from all the infrastructure in your data center.

Modius solution: Modius provides a robust, scalable platform for collecting hundreds of thousands of data points from networked and non-networked equipment. Modius provides all the tools needed to efficiently collect, normalize, correlate, manage, and visualize all the data in your data center, regardless of the source, to support Real-time Operational Intelligence for decision making. The Modius OpenData platform can extend your legacy applications (No Rip and Replace), seamlessly collecting data from your infrastructure and federating it across your organization to the people that need it most.

2) Automate the computation of metrics

Effective Real-time Operational Intelligence involves many calculations and adjustment factors. Because enterprise standardization is desired and because every data center is different, these transformations will vary from one facility to the next. As such, the data capture solution needs to have a flexible computational capability built into it.

Modius solution: Through its innovative in-database analytics technology, Modius provides an extremely easy and flexible way to create a standardized set of metrics, as well as make adjustments to the calculations to account for the differences in facilities and infrastructure. Modius enables users to change these underlying computations without programming, complicated SQL queries or other procedures seen in other systems. Instead, Modius provides a simple and intuitive interface.

3) Capture data continuously in Real-time

Finally, and most importantly, Modius believes that operators should continuously capture all the necessary data throughout the day. Rather than collect data only intermittently, Modius recommends that the system be able to record measurements as often as once a minute.

Modius solution: To achieve this requirement, Modius has developed a monitoring and measurement technology that actively polls each device multiple times a minute to collect the most up-to-date performance information. The solution then records the data in a centralized database for easy trending and analysis over time.



About Modius

Modius is an independent software vendor based in San Francisco, California. Founded in 2004, Modius develops and delivers intelligent measurement systems for mission critical facilities that improve business continuity, energy performance, and carbon management. Modius solves the challenge of integrating both IT and Facilities Management information into a single, comprehensive measurement system. Modius empowers Real-time Operational Intelligence for improved data center management through:

- Widespread, practical, low-cost collection of all physical-layer performance data
- Trustworthy and reliable analysis tools based on comprehensive data and rich-analytic capabilities
- Useful and actionable Real-time Operational Intelligence through highly-configurable business logic
- Customized workflows, delivering the right intelligence to the right people at the right time.

For more information on Modius, please visit www.modius.com.

[Modius, Inc.](#)

71 Stevenson Street, Suite #400
San Francisco, California 94105
USA

888.323.0066 (Toll-free)
415.655.6700 (Main)
415.655.6601 (Fax)

© Modius Inc., 2014, Unauthorized use and/or duplication of this material without express and written permission of Modius Inc. is strictly prohibited. Excerpts and links may be used, provided that full and clear credit is given to Modius with appropriate and specific direction to the original content.

All rights reserved. Modius and OpenData are registered trademarks of Modius, Incorporated. All other trademarks are the property of their respective owners.