

Taxonomy of Data Center Instrumentation

The Definitive Guide to the Instrumentation Landscape within the Modern Data Center

By Mark Harris,
Vice President, Product Management, Modius, Inc.

<http://www.modius.com>

January 2010

Preface

This paper is intended to provide a comprehensive background on the various types of instrumentation available within the modern data center, with a clear comparison of each of these types including their suitability towards the ITIL¹-inspired goal of “continuous optimization” or more commonly “operational intelligence”.

It will be shown that modern datacenters are complex systems with a tremendous quantity of physical infrastructure devices already in place, some of these components with monitoring capabilities built-in, some with optional monitoring features optionally available, and finally others without any monitoring capabilities whatsoever. This paper will highlight where visibility gaps exist within this measurement/monitoring landscape and provide guidance on various technologies that can be deployed to fill these gaps.

Furthermore, this paper will show that more real-time monitoring information is always better to help make better informed decisions. Granular concise, real-time information will allow trends to be seen, thresholds to be set, plans to be made. We will discuss many of the various approaches to device instrumentation over the years while at the same time emphasize the reasons why most IT situations will actually require a combination of several instrumentation technologies to work together, allowing a complete picture of status, availability, capacity and efficiency. We will consider what is required to normalize or transform otherwise raw instrumentation data points into this actionable information and metrics and also consider the business aspects and opportunities available when a well conceived instrumentation plan is executed.

We will hypothesize that the current less-than-optimal state of active monitoring through instrumentation within the modern data center is a direct result of the historical lack of comprehensive distributed enterprise-class solutions to gather, analyze and make informed management decisions across the enormous source of raw measurement points available across any IT structure. Lastly, the technology to provide this continuous monitoring of vast discreet sources of data points has only become available in the last few years to be sourced, deployed and consumed at will.

Executive Summary

Since the Industrial Age began, nearly every piece of commercial equipment invented has had some form of associated telemetry (also known as instrumentation) to go along with it. In most cases this telemetry has been put in place to confirm, modify and

¹ http://en.wikipedia.org/wiki/Information_Technology_Infrastructure_Library

diagnose its operational status. Analog or Digital alike, nearly every device can benefit from instrumentation.

While very early systems were mechanical in nature (such as the simple water capacity indicators in a steam locomotive), with the advent of electrically powered equipment came the ability to use much more advanced technology to look deeper, in new and exciting ways and extend the range of this telemetry. In addition, the analysis of raw data was no longer limited to a single point in time by a single user, but could be correlated, viewed and studied by any number of users at any point in time.

A wide range of Instrumentation within the design of any type of complex system has become critical. Consider the complexity of the cockpit within a modern jet airliner. There are literally thousands and thousands of points of data that are being gathered in real-time sourced by a wide array of individual manufacturer's components, in a variety of measurement units. Some key discrete points are presented without interpretation to the crew as they are gathered, while other points may be grouped and considered relative to other discrete points before the crew is made aware of this condition.

A modern datacenter is also a complex system. For a data center to be operated efficiently, it's critical to gather as much information as possible from a vast array of dissimilar vendor equipment, normalize the measurement units, and then provide the analysis of this instrumentation to allow highly-confident business intelligence. It will be seen that a modern datacenter must be designed in such a manner to be able to take advantage of all of the discrete instrumentation points, across IT systems, Power Chain components, HVAC systems and building security systems. With all of this data available, the need to aggregate and normalize these data-points (in real time) is critical to leverage the opportunity to meet efficiency level previously unattainable. Ultimately the opportunity to significantly increase the efficiency of all data center operations is possible only when carefully choosing and combining multiple instrumentation systems and the analytics that is possible with that information, and doing so in a manner that incorporates the processes and discipline to enable change in the data center.

Backgrounder and Historical Approaches

The deployment of instrumentation for complex systems has always had its Achilles' heel in connectivity. Typically, while some level of sensing technology itself has been available in the devices themselves, the ability to transport the resulting data created by these distributed sensors has been difficult. In early mechanical systems, many sensors had their presentation physically limited to the same location as the sensors (such as a railroad track switch position mechanical flag), or this connectivity was addressed with

simple technologies such as pressure transfer lines and round pressure gauges to indicate status some distance away from the sensor point of origin.

Until very recently, modern electrical systems have emulated this design center, either visually representing metrics or status very close to the sense point, or with longer runs of individual discrete wire or cabling from the source of the sensor device to the desired location of the presentation device.

Consider a very complex system, the soon-to-be-retired space shuttle orbiter built by Boeing Corporation in the late 1970s and early 1980s. Within its 186 feet of fuselage, are more than 1.2 Million feet of wire² used to connect discrete systems and their instrumentation. Each sensor is individually wired to its intended display. This can be clearly seen as an arduous task to install and deploy a complex system in this manner, not to mention the maintenance and change management processes that are required to work with systems built in this manner.

Now consider the data center. For the same fundamental complexity reasons, the vast majority of datacenters have had limited or no instrumentation deployed over the past 10 years, and the strategic value of the available metric information has been minimized due to this limited coverage. In nearly all cases, monitoring for IT devices has been satisfied with the inclusion of status LED or LCD indicators on the hardware devices themselves. In some cases more complex devices have had some level of remote monitoring installed, but the capture and presentation of this instrumentation data has remained unavailable to all but a few individuals. In many cases, we see large portions of critical infrastructure being left unmonitored.

Sadly, while the management of logical production IT systems has become a huge business today, the real-time management of the critical infrastructure underneath these systems, and the interaction of these logical and physical infrastructure management schemes have been left alone.

The State of Instrumentation Today

What is needed in a modern data center is the ability to gather and normalized performance data using a combination of in-band and out-of-band instrumentation technologies. (In-Band refers to the usage of common general-purpose networking transports such as the existing LAN to carry information, while Out-of-Band refers to the usage of discrete application-specific transport, most commonly discrete cabling).

² http://www.boeing.com/defense-space/space/hsfe_shuttle/docs/shuttle_overview.pdf

It can be argued that since modern data centers are comprised of a wealth of discrete subsystems from a variety of different manufacturers, the instrumentation of the data center must (by design) allow for multiple technologies. The requirement to use multiple sensing technologies stems from three factors; 1) corporate structures which change over time, adding or divesting companies or divisions with a wide array of existing technology already in place. 2) The fiscal requirement to leverage the almost ubiquitous inclusion of performance instrumentation already provided within many active devices by major suppliers within the data center, and 3) the need to augment all available sources of instrumented data with additional and specific data types such as branch circuit consumption, under floor pressures or outside temperature.

Device Included Instrumentation

Active devices today (such as servers and networking equipment) routinely include some form of platform management interface via in-band technologies, most commonly a LAN connection. Once connected, these devices use various protocols built upon the LAN transport to allow access to a wealth of instrumentation data. Some of the most common management protocols include SNMP³, WMI⁴, IPMI⁵ and ILO2⁶ and the choice about which protocol to use is commonly made by each manufacturer. In some cases, the choice changes over time, within a vendor's own product catalog or across various series of product types.

These in-band management interfaces provide a tremendous number of management capabilities, ranging from power-cycling a device or reconfiguring its hardware options to the ability to read the various sensors contained within the unit itself. Most commonly, these sensors may allow operational status, temperature and/or power consumption values.

To augment the device-included instrumentation discussed above, two additional types of environmental and power instrumentation have arisen over the past decade, wired and wireless.

Wired Instrumentation

Wired instrumentation became popular nearly a decade ago when a limited set of additional sense points were desired and the technology for low-power networking was not yet a viable option. Total costs for wired instrumentation (product and services)

³ <http://www.snmp.com/>

⁴ http://en.wikipedia.org/wiki/Windows_Management_Instrumentation

⁵ http://en.wikipedia.org/wiki/Intelligent_Platform_Management_Interface

⁶ <http://h18013.www1.hp.com/products/servers/management/iloadv2/index.html>

were fairly steep, and the physical requirements to install such systems were in many cases prohibitive. As such, utilization of wired instrumentation was relegated or limited to only the most critical items or areas within an infrastructure. Accordingly, the vast majority of industrial control, HVAC and building and facilities systems manufacturers created vendor-specific options for their own products to allow wired access to their device's instrumentation for monitoring and control purposes.

Again, referring back to the Shuttle Orbiter program, data centers could be instrumented with miles of cables if so desired, but rarely did this happen due to the intrinsic complexity of doing so, and the limited return on investment due to the very nature of different systems, each with different value types and limited integrations with any other system.

As such, the industry has largely left the technology choice for remote management and monitoring up to each individual vendor, although various organizations have put forward proposals for standards to help rectify such widespread diversity. SMASH CLP⁷ is a good example of one such standard proposed by the DMTF⁸ and currently being considered.

Wireless Instrumentation

“Surely there must be a better way to communicate sensor information” was the mantra in the early 2000s. The growing importance for instrumentation in the data center was beginning to be seen, but there simply was no easy way to get this important operational information to the users that needed it. In response, a handful of wireless technologies rose to the surface.

The first mature technology was based upon PARC's⁹ years-earlier “Smart Matter”¹⁰ research work which provided the basis for a new DARPA funded project called Smartdust¹¹. Smartdust and related projects were ultimately conducted within several of the University of California (UC) campuses (Berkeley, Davis and others) to study low power, miniaturized instrumentation.

All of the subsequent commercial environmental instrumentation offerings based upon this UC research work can be described or approximated by the current standards

⁷ http://en.wikipedia.org/wiki/Systems_Management_Architecture_for_Server_Hardware

⁸ http://en.wikipedia.org/wiki/Distributed_Management_Task_Force

⁹ [http://en.wikipedia.org/wiki/PARC_\(company\)](http://en.wikipedia.org/wiki/PARC_(company))

¹⁰ <http://www2.parc.com/spl/projects/smart-matter/>

¹¹ <http://en.wikipedia.org/wiki/Smartdust>

outlined within the “Zigbee” (or officially known as 802.15.4-2003¹²) specification. The created products each offer very low power consumption, high data reliability and due to their limited distance capability, the ability to send its data upstream via a “hop” to its closest neighbor thereby creating a segmented upstream path.

While successful in realizing the dream to transport instrumentation metric data wirelessly in a small form-factor, all of these UC-inspired low-power product offerings available today tend to be susceptible to severe distance limitations due to interference with the structures they are deployed within and signal type itself. Technically, this can be attributed to their common reliance on higher-frequency signals (2.4Ghz or higher) that are shared with many other wireless technologies, most notably the many flavors of 802.11 or “Wi-Fi”¹³. (The ability of any signal to penetrate dense construction materials, such as concrete and steel is inversely proportional to the frequency itself, i.e. lower frequency waves penetrate dense materials more easily). As such, these systems must typically deploy a larger number of instrumented sensor and receiver combinations, each with smaller ranges, and they themselves intercommunicate to create a mesh topology to hop from receiver to receiver before getting data delivered to the intended point of data collection.

More recently, a second mature wireless technology to make its appearance in support of data center instrumentation was a novel usage of a mature technology called, “RFID”. While various engineers had been working with various frequency versions of RF based ID technologies¹⁴ for more than 50 years, it was in the early 1970s with proof-of-concept projects at Los Alamos Labs and the New York Port Authority that their true potential became obvious. Later in the mid 1990s, general-purpose UHF RFID products were finally introduced to the masses due to a wealth of applications stemming in the retail inventory management space.

Recently, commercial-grade inexpensive RFID solutions have been introduced which have the required level of performance and distance to satisfy the commercial IT segment. These new RFIC-based products have been introduced using an “active” version of RFID running at a very building-friendly 433MHz frequency. The active version of RFID takes the original technology one step farther by allowing each individual sense point to periodically report its metric data without the need to manually interrogate the unit for information. This combination of battery powered RFID and the frequency

¹² <http://en.wikipedia.org/wiki/ZigBee>

¹³ http://www.wi-fi.org/discover_and_learn.php

¹⁴ http://en.wikipedia.org/wiki/Radio-frequency_identification

choice of 433MHz provides for years of self-contained battery life across an impressive distance, regardless of the data center construction materials, rack configurations and enclosure steel specifications. Active RFID technology receivers can be placed in the middle of a circular cell of more than 1000 devices, each cell with a radius of up to 100 meters. These cells can overlap as needed to provide coverage, and the receivers themselves simply report instrumentation data upstream, where other intelligence eliminates duplications or errant payloads. By using mainstream lower frequency RFID technology components, these commercial Active RFID offerings can be extremely cost-effective and thus allow for more extensive deployments, in more places, with fewer wires of any type, across any type of data center configuration.

Putting it all together

It can be seen that the choice for instrumentation within an IT structure is really not about choosing ONE specific technology that seems to be best suited for any given company. Strategically, it's about which combinations of solution(s) that when put together, form a COMPLETE view of actionable management information. Each given portion of the instrumentation goal should be carefully considered, based upon availability, cost and deployment requirements. When properly put together in a cohesive fashion, this results in a highly granular and actionable consolidated view of the complete infrastructure.

The full range of instrumentation data sources to be considered will ultimately include many of the following technologies in wide-scale usage today:

1. IT Appliances, the "1U" or "2U" device

Most modern IT equipment has some form of instrumentation built-in. Each of these appliances uses various forms of communications protocols to present this hardware-based instrumentation. For networking devices, it is very simple and SNMP has arisen as the defacto standard to communicate nearly all aspects of a given platform in operation. For server devices, most IT vendors have focused on using IPMI or a flavor of it (DELL and IBM for instance use protocols called "RSA"¹⁵ and "DRAC"¹⁶ which are based upon IPMI and Sun has used a number of protocols including IPMI, ALOM and SSP¹⁷). HP has remained committed to their own proprietary protocol called "ILO" with some SMASH-CLP¹⁸ extensions, but essentially provides similar capabilities as the IPMI standard for servers of any type.

¹⁵ http://en.wikipedia.org/wiki/IBM_Remote_Supervisor_Adapter

¹⁶ http://en.wikipedia.org/wiki/Dell_DRAC

¹⁷ http://en.wikipedia.org/wiki/System_Service_Processor

¹⁸ <http://www.dmtf.org/standards/mgmt/smash/>

2. High Density Chassis

The largest IT providers (such as CISCO, HP, DELL and IBM) have created complex chassis based offerings over the years to combine IT functions. These chassis include high density 'Blade Servers', Storage fabrics and Network Switching devices. Each of these chassis has a combination of function cards, management cards, power cards and cooling trays. Subsequently, each chassis product has a wealth of information which can be communicated about the operational status of each system and sub-component. Everything from the percentage of CPU utilization to the speed of any one of the system fans can be interrogated via SNMP or XML¹⁹.

3. Discrete Individual Networked sensors

A very traditional approach to monitoring a few environmental points has created the market for single-purpose, wired instrumentation products. These can be characterized with products from companies such as Axis and IT Watchdogs. These vendors create fairly low-cost small form-factor sensors which most typically have a single LAN port and an external power supply. Since they do not maintain connectivity to any larger system, they typically provide the current values for the sensor itself. (I.e. the instantaneous temperature at the device, or a snapshot image of a walkway or aisle in a datacenter).

4. Wired Instrumentation Systems

Building upon the discrete concept of instrumentation, a number of vendors have gone to the next level and built 'Platform' systems based upon the concept of aggregating many discrete wired instrumentation products. Vendors such as Sensaphone, Avtech, Netbotz and Network Technologies have create fairly complete catalogs of sensor devices and aggregation hubs which can be connected to any given LAN and provide the current and historical values of any of the sensor points deployed within their own system.

5. Wireless Sensor Systems (802.15.4)

The wireless category of instrumentation originated with the early commercialization of the 802.15.4 or Zigbee standard, and has been applied as a transport by sensor monitoring companies such as Synapsense, ArchRock, Dust Networks and PacketPower. These wireless providers offer a growing lists of instrumentation sensors and can all be thought of as low-power, short-distance hotspot technology. Many hotspots are created in overlapping structures to create a topology loosely described as a 'mesh'. All of these mesh systems are deployed in the traditional 2.4ghz frequency range, so each vendor has take various approaches

¹⁹ <http://en.wikipedia.org/wiki/Xml>

to minimizing the interference and corruption of packets due to other co-resident wireless technologies, most notably the 802.11A/b/g/N “WiFi” networks.

6. Wireless Sensor Systems (Active-RFID)

Perhaps one of the newest categories of instrumentation is based upon the RFID technology and commonly referred to as “Active RFID”. The technology has been introduced at the time of this writing by a company named Rfcode. Products developed using Active-RFID exploit the lower frequency, building-friendly 433Mhz frequency-range which provides for highly reliable, long distance communications in a crowded data center.

7. Building Management Systems

Historically, nearly all environmental information was instrumented, gathered and presented by the chosen Building Management System. The chosen BMS did a adequate job at gather the more traditional HVAC style componentry, in some cases perimeter/security information, but did very little outside of their realm. In additional, most of these systems locked-away the resulting information for all but the select few individuals that went to the effort of interacting with these limited systems. The traditional BMS is an island of information.

Normalizing the data

It is important to note that each type of instrumentation discussed in the preceding pages (regardless of type) uses any one of a number of general purpose transport layers and some type of vendor chosen protocol to communicate and present its values to the user, and each device type does so differently. The usage of a common network protocol like SNMP for example still leaves a great deal of vendor-specific representation uniquely defined, making it fairly difficult to interpret these values when interacting with this device. Most common issues include the varying type(s) of data returned and the scale of units of data which vary greatly from device to device. (I.e. some devices may report temperature in Celsius rounded to a whole number, while others report in Fahrenheit with a tenth of a degree precision, with or without decimal point and possible the letter “F” included).

All of this instrumentation diversity creates the absolute requirement to understand what specific vendor-provided data is being retrieved and then to normalize all of this data into some agreed common unit of measure. This ultimately allows the analysis of data for business intelligence and can provide the consolidated visibility into the data center. The consequences of not providing or incorrectly providing this normalization of data units can be catastrophic as was seen with the 1999 crash of the \$327M Mars

Climate Orbiter²⁰ when various subcontractors intermixed metric and imperial data points.

Conclusion

A well-conceived IT instrumentation plan calls for the ability in a single monitoring structure to allow for ALL of these aforementioned instrumentation technologies. These include native device embedded, hardwired, wireless and building management systems. For a wide range of technical and organizational reasons, most companies will find themselves having to gain access to various combinations of these instrumentation technologies in order to gain a comprehensive picture of their complete infrastructure.

Secondly, all data collected will be required to be normalized. For example, the conversion of all temperature metrics from all sources, to always be reflected in Fahrenheit with precision to one-tenth of a degree.

Lastly, once instrumentation data has been made accessible, and has been normalized, it is critical that this data can be aggregated over vast geographies as most corporations find themselves with more than a single location. Large and small, all companies deal with multiple sites, and various quantities of IT infrastructure at each location. While knowing the temperature of a user's office on the other side of the country may not be as relevant to this discussion, the ability to see each data center's capacities in terms of power and cooling may be instrumental in determining how to distribute processing load over the course of any given day, across seasons, etc. Aggregation of normalized data allows business decisions to be made upon the complete picture

Recommendations

Once the reader has an understanding of the wealth of instrumentation choices available to them, a simple set of steps can be executed to begin taking advantage of this previously untapped source of knowledge:

- 1) Take Inventory of the IT structures contained within. A fairly comprehensive baseline of all devices should be established which helps to identify the instrumentation already available as well as areas that need to be enhanced. It is important to include devices that have instrumentation already contained within but are left unconnected today.
- 2) Identify key subsystems which would affect business practices or corporate production. Consider core IT devices as well as the underlying power distribution

²⁰ http://en.wikipedia.org/wiki/Mars_Climate_Orbiter

components and associated cooling structures as each of these are valid sources of resource information.

- 3) Consider the most relevant key performance indicators relevant to running your core IT operations efficiently and set as a near-term higher priority goal the creation and visibility of these KPI metrics in a real-time. Look for ways to leverage this analysis, using common ITIL²¹ 'continuous optimization' principles; optimizing structures, balancing loads, etc.
- 4) Determine the instrumentation systems that are currently available as well as where additional system(s) will be required to fill in critical gaps. Create a strategic plan that identifies the risks associated with NOT monitoring and actively managing each component, and look for the best instrumentation solution(s) for each major area. Most likely a selection of wired, wireless as well as the inclusion of a legacy BMS system will be required.
- 5) Look for an overall higher level vendor-neutral monitoring solution that is not tied to any single hardware instrumentation platform, and can be rapidly deployed. This system should be able to gather in real-time metric and key performance information from any type of instrumentation. Deploy sooner rather than later.

²¹ http://en.wikipedia.org/wiki/Information_Technology_Infrastructure_Library

References:

Discussion about STS Shuttle Program,

http://www.boeing.com/defense-space/space/hsfe_shuttle/docs/shuttle_overview.pdf

Discussion about SNMP by SNMP Research,

<http://www.snmp.com/>

Encyclopedia entry regarding IPMI,

http://en.wikipedia.org/wiki/Intelligent_Platform_Management_Interface

HP's product feature-set for Integrated Light-Out 2,

<http://h18013.www1.hp.com/products/servers/management/iloadv2/index.html>

Discussion about the Smartdust project under DARPA,

<http://en.wikipedia.org/wiki/Smartdust>

Facts about RFID,

http://en.wikipedia.org/wiki/Radio-frequency_identification

Discussion about the failed Mars Climate Orbiter due to unit mismatching,

http://en.wikipedia.org/wiki/Mars_Climate_Orbiter

Discussion about Service Processors within IT equipment,

http://en.wikipedia.org/wiki/System_Service_Processor

Discussion about the DMTF's SMASH-CLP program,

<http://www.dmtf.org/standards/mgmt/smash/>

