



OSIsoft®

*“Where Facilities and IT meet,
the Enterprise Data Center”*


Published: August 2007



The trend is growing that the CIO will have the responsibility to make the Enterprise Data Center the foundation for all strategic decisions within a competitive organization.

Increased Power Consumption increases cost and of course federal regulation heightens the risk factor. The US congress recently passed a bill through the Energy Star program that authorizes the Environmental Protection Agency to study Data Center power consumption. As risk and cost continue to become growing concerns the CIO must look at a converged strategy.

Even though we as a society use innovation to augment our brick and mortar commerce, with the advent of the internet, organizations are still limited by their facilities (Brick and Mortar) to power their ever-growing need for IP based technology. CIO's are faced with strategic IT decisions having to be made based on the facilities that can support the ever growing need for technology. Obstacles within the data center: *is there enough space for the hardware, is there enough cooling, and is there enough power to meet requirements.* The dependency on the facilities team will only continue to grow as any new technology change requires strategic planning. Power is the second total cost of ownership (TCO) today, and has doubled in price since 2002 as the hardware install base has grown exponentially.



Lawton, *IEEE Computer*, Feb 2007

That being said the ability to converge Operational information and IT information into a common real-time infrastructure is becoming more and more understood to be of value. The impact for companies to capitalize on an enterprise energy management program could cut 2% of power consumption within their organization. Customers that are planning on implementing this type of converged infrastructure predict a \$20 million return per year.



Optimizing the Data Center Common Sense Methods for Cost Reduction

Introduction:

In the past few years, there has been a tremendous focus on data center technology. A number of factors are driving this trend

- Expansion of the worldwide web and automation of on-line business transactions require increased processing capacity
- Increasing business need for 24x7x365 access to those resources
- Most of the data center capacity vacated by the Technology Bust in the late 1990s has been consumed, and demand is forecast to increase in the years ahead
- Increasing business needs for Rising cost of energy (electricity) that is expected to continue into the foreseeable future
- Push toward "GREEN" and reducing energy consumption and "carbon footprint"
- Sarbanes Oxley, HIPPA and other regulations require increasing levels of security and technical expertise to ensure compliance

Business Challenge:

Data centers are no longer considered just another cost center to most organizations. They are increasingly being viewed as a means of competitive advantage. Organizations are either investing in data center optimization, or outsourcing this business to those who can provide the high-availability, on-demand services in a secure environment at the most competitive cost. The largest driver of data center optimization are the cost savings associated with energy consumption. Clean, reliable electric power is required for operation of the server farms, IT infrastructure, lighting, cooling, and other environmental loads. Rack space and cooling capacity are also limiting factors in data center operation and optimization.

There are thousands – perhaps millions - of pieces of information that can be analyzed in the operation of a modern data center. Consider the following:

- When are the peak energy demands on the data center? Peak cooling demands?
- Do these correlate to outside weather conditions such as temperature and humidity?
- How do these demands relate to the actual cost of electric energy at that time of day?
- How are these demands changing over time? Are the peaks increasing or decreasing...? If they are increasing, at what point in the future will electric power or temperature limits be reached?
- What rack-level components are the largest drivers of energy consumption?
- What is the temperature profile at the rack or component level, and how does this vary?
- What is the normal load on an Exchange Server cluster on Monday morning at 9 AM? How about on Thursday afternoon? What is the associated power consumption and rack temperature?
- How does the UPS respond to these changes?
- How do these profile vary during the work week? On weekends...?



Solution:

Today's data centers require a comprehensive monitoring solution to ensure energy optimization, as well as performance, reliability, and security. The combination of real-time and historic information - plus the ability to analyze, filter, and visualize it all - is required for data center optimization. Historic data allows users to understand the baseline capacity loading, temperature, and power consumption profiles and their variability during normal operation. Rather than using a brief snapshot during a random survey or a 30-minute average from a limited number of assets, the ability to capture all of this data and store it for analysis provides the user with a detailed profile of asset utilization history and variability. Once the normal baseline capacity and performance information are understood, real-time data can be compared to the "normal" baseline using statistical analyses to determine trends and provide early warning whenever any of these items is out of its normal range. In addition, with data center staffs increasingly tasked to do more with less, automation and logic are needed to ensure the equipment runs reliably without searching for problems and the ability to predict a failure in advance. Any data center optimization solution should meet the following basic requirements:

1. The ability to monitor to every asset, regardless of the vendor or protocols. It should provide a combination of real-time and historic data to collect "all the data, all the time" from millions of data streams from all of the data center assets including:

- Hardware & devices (servers, routers, switches)
- Network connectivity (LAN, WAN, SAN)
- Security devices (VPN, Firewalls, IDS)
- Operating Systems
- Applications (SPAM, VOIP, AVVID, Help Desk, anti-virus, e-mail, databases)
- Building controls (HVAC, fire protection, alarm systems, etc)
- Power supplies and cooling systems

2. A time-series database capable of archiving large amounts of data in its original fidelity, and maintain it online for long periods of time. Such data is required for historic reference, and provides trends and baseline information from which normal network operations can be determined. It must be capable of handling data rates in excess of 75,000 events per second, and employ advanced compression algorithms to optimize data storage space and reduce the cost of digital storage.

3. Condense, filter, and summarize this information so that humans can quickly understand what is happening. Powerful analysis tools should enable allow advanced calculations, event framing, and correlation to synthesize all that data into useful information. Each user should have unique, role-based dashboard views that provide the "right data, right now" in order to speed their decisions. These dashboards provide context based on the user's job functions, and should include more than just tables. People can be overloaded by information, so they need trends, graphs, and animation, plus an alerting function that allows management by exception. This information should then be distributed to users in real-time, anywhere in the world. Users should be able to view their responsibilities at a high level, with the ability to immediately drill to down through a series of customized displays to get at the details immediately.

In addition, with new assets and systems constantly being added to the data center, any solution must be flexible enough to accommodate the addition of these new assets from different vendors within a period of 90 days.

4. Advanced intelligence capabilities are required to analyze trends and predict failures and events before they occur. The analytics engine should be able to baseline historic data and determine what is "normal"



based on statistical analysis and patterns. The engine should also have ability to sweep through oceans of this data quickly to identify patterns, trends, or abnormalities for further investigation and action once a pattern is detected.

5. For organizations with multiple data centers, the system should scale to allow analysis, load balancing, and optimization across multiple locations. This includes receiving millions of data streams distributed throughout a global geography, while serving many users simultaneously. The information could be rolled up from a single data center all the way to a global view.

6. Agentless technology should be utilized, so as not to artificially increase the load on data centers or networks that are already busy.

7. The solution should provide connectivity to other databases and systems, to access data elsewhere without the need to replicate the master data again. Standard interfaces to other data sources provide the capability to bring in data from other sources & monitoring tools to perform KPI summaries and financial calculations, enabling users to understand not only what is happening, but how much it is costing them.

8. The solution would support an N-tier architecture with the appropriate bandwidth and hardware, plus the ability to accommodate millions of data streams on a single server. In addition, the solution must allow accessing this information via the web using hand-held devices as well as laptops and desktop PCs.

9. The solution should utilize COTS technology throughout in order to benefit from widely accepted industry standards. This will ensure timely migration to the latest technology, while avoiding large expenses in cost of custom coding and maintenance.

Benefits (single data center location):

Energy savings from reduced server and cooling

As energy costs continue to rise, energy cost reduction will continue to escalate in importance. For a single data center, load reduction, loading efficiency, and load shifting to off-peak hours will reduce the costs associated with server operation and the associated cooling loads. For organizations with multiple data centers, the ability to shift processing loads to geographies where power is less expensive (“follow-the-moon”) will provide additional benefits.

Identify and reduce unnecessary capacity

Many new IT projects establish requirements for separate server hardware during installation and configuration. However, after several months in operation the applications may consume only 5-10% of the assets' capacity. Combining a number underutilized assets onto a single server can significantly reduce server power and cooling consumption, plus free up rack space for additional capacity needs. The ability to determine this from historic data provides the data center manager with the information to make that decision – and the data to justify it to the appropriate stake holders. SQC methods, including Six Sigma, can then be applied to further optimize the use of these assets by load shifting or load sharing through the use of VM environments.

Capacity Planning:

The ability to monitor & track the utilization of all assets for the data center, and to predict when an asset's capacity will be exceeded. Again, historic data provides the data center manager with the information to make that decision – and the data to justify it to the appropriate stake holders.



Increased MTBF /reduced maintenance costs:

In general, it is much less disruptive and less expensive to plan for the maintenance of a particular component, such as a server or HVAC compressor, than it is to run that component to failure. The ability to use trending information to predict in advance the loss of a critical service or the failure of a component allows the data center management to have the asset serviced during a scheduled outage – rather than on an emergency basis – avoiding service interruption and costs associated with emergency repairs.

Public perception and corporate image: Carbon Footprint Reduction:

Consumers are increasingly using their wallets to reward companies that are “green” or environmentally friendly. Reducing energy consumption and improving efficiency can help in promoting the image of a good corporate citizen, and can be very good for business

SLA Compliance Tracking:

A secondary, but very important benefit of historic data retention is the ability to monitor and track SLA Compliance of hardware or service vendors. The ability to retain detailed historic data provides the data center manager with the information needed to detect non-compliance from an underperforming OEM or Service Provider, and to provide documentation when seeking performance improvement and/or restitution.

Conclusion:

Many IT budgets are constrained by enormous percentages allocated for maintenance and personnel. Using a converged architecture to lower operational expenditures for IT spending may free up budgets for innovation in technology for companies to become competitive through technology advances. Standardization of software platforms allow for the increasing cost of trained professionals. Clearly the ability to converge Operational information and IT information allows for better decision making and better use of resources and spending.