A HEALTHY DATA CENTER BEGINS WITH A WELLNESS SELF-ASSESSMENT

By Mark Settle, Chief Information Officer, BMC Software

If you've got a headache, it's usually easy to locate the cause of your pain and take a pain reliever to ease your discomfort. If the problem persists, your doctor can give you a checkup to help identify the cause. You may even be able to take certain actions to prevent the headache from occurring.

In many ways, managing your health is like maintaining a data center. You need to identify the source of "pain" in the data center, and then take steps to alleviate that pain.

What does it mean for a data center to be healthy? You should be able to meet service expectations and commitments. You should be able to maintain control by defining and managing risk. The data center should help your company get the most out of every person, asset, project, and activity. And your data center should enable your IT organization to operate at the speed of the business.

Is your data center in good shape? This article outlines some questions to help you assess the health of your data center so that you can take the appropriate actions to achieving wellness.

Take Inventory

Gathering the vital statistics of your data center begins with finding the answer to this fundamental question:

» *Do I have an accurate asset inventory of my hardware?*

Imagine, for example, that you've purchased blade servers from multiple vendors. Your Vendor "X" representative makes a pitch to standardize all your blade technology. If you know that the value of your balance sheet investments of non-Vendor X blades is $1.4 million, you can begin negotiations. You might say, "If I were to purchase another million dollars of your company's blades, what could you do to help me offset the write-down of the value of the blades on the floor today?"

Having information about your current equipment gives you leverage with your vendors. Yet, in many companies, finding the current book value of the existing blades could tie up operations staff for days — manually checking what's out there, getting into the fixed asset records, researching other peoples' spreadsheets, and so on.

Rather than waiting until you need the information, ask yourself these questions today:

» *What do I own?*
» *How much of it do I have?*
» *Where is it located?*
» *What is the financial value of the technology I own?*
» *What's still on my balance sheet that I'm depreciating?*
» *What am I leasing?*
» *How many generations of technology have I fallen behind?*

Not knowing what technology you own can cost money, time, and a lot of frustration. For example, a large company was in the process of reducing their data centers from five to one. They had picked one of the data centers as the consolidation location. They purchased extra-large WAN circuits to migrate applications and data from the satellite centers back to what would become the new primary center. At the third data center, the technicians had removed all the equipment from the floor, and the room was empty. Then one of the network engineers said, "Don't disconnect this data center from the network. We're still getting signals that there's traffic on the network emanating from here, even though it now has no equipment."

Not knowing what technology you own can cost money, time, and a lot of frustration.

They tore up the raised floor in the data center and found servers underneath that were still operating. Yet there was no record of them! The network engineers had no knowledge of what the servers were doing. And they didn't know if it was safe to disconnect them or what might happen next if the servers were disconnected. This company had a lack of information about what they really owned. A complete, up-to-date inventory of all their equipment would have prevented this situation from occurring.

Address Capacity Management

Virtualization in the data center and the move to cloud computing has brought renewed interest in capacity management. In the past, as many IT shops moved from the mainframe to a distributed environment, they tended to buy heterogeneous equipment that was running at 20 percent or less capacity. Servers were so inexpensive that they tended to purchase more and more of them, even though the servers might be used only 15 percent of the time. However, the pendulum is swinging back again, and we're starting to realize that running below capacity is a frivolous waste of money. The servers not only take up space, but they also increase your energy costs.

Ask yourself these questions to help determine the cost of your investment in servers:
» *How many servers do I own?*
» *How much heat are they generating?*
» *How many kilowatts do they use?*
» *How much time is my operations staff spending on maintenance and wiring?*
» *What is the return on my investment in that hardware?*

Control Labor Costs

After you've purchased hardware, integrated it with the network, and paid your costs for running the network, you also have to understand the operational cost of labor. In fact, labor costs can represent one of the most significant aspects of your operational budget. Labor efficiency is a key measure of data center efficiency.

To get control over labor costs, ask yourself the following:
» *What repetitive, manual tasks can be eliminated, streamlined, or automated?*
» *Which tasks routinely expose us to human error?*

Maintain Security

The physical security of data centers is generally well managed. It's more likely that security could be jeopardized by an electronic invasion. As hackers develop new technologies and look for different kinds of targets, it's important to develop ways to protect your infrastructure.

To verify that you have the right security safeguards in place, ask yourself these questions:
» *Do I have a third party performing unannounced penetration testing? (The CIO and the VP infrastructure would likely know about these tests.)*
» *Do they test my internal defense systems so that problems can get caught and/or reported?*
» *Am I able to test various applications randomly at different points in time?*
» *Is IT able to intercept viruses?*
» *Is my company protected against a hacker trying to get into the systems?*

Provide Effective Patch Management

Microsoft puts out patches frequently. It's a good idea to trickle out the patches so that they are not as network-consumptive as releasing them all at once. If you are not patching effectively, there will be a lot of calls to the service desk relating to the consequences of patch inconsistencies.

Ask yourself these questions about your internal ticket flow:
» *Am I doing more reactive work or proactive work?*
» *Which incidents have disrupted user productivity?*
» *Which incidents are simply user requests for increased IT capabilities and are not related to patch management?*

Track the Right Metrics

According to Gartner, "Through 2015, 80% of mission-critical outages will be caused by people and process issues, and more than 50% of those outages are caused by change/configuration/release integration and hand-off issues."[1] That's why it's so important to measure the availability of your tier-one systems, including the outage times. Begin by tracking these metrics, and have reporting information available as well. Answer the following questions:

» *Am I tracking my tier-one systems?*
» *Am I getting better at tracking?*
» *What are my resolution times around sub-one problems?*

A sub-one problem would be an outage of a component or service that wouldn't necessarily derail the entire service. A classic example is a high-availability cluster of servers where you lose just one server, yet the application stays up and running. The resulting problem is more related to a degradation of that service than of failure. The users might not see a performance issue, but if you track this information, you would know that you're running a riskier configuration because you lost one of the four servers. Address the following questions:

» *What is the outage time around my tier-one applications?*
» *How am I monitoring that outage?*
» *How robust is the monitoring?*
» *Am I using external, synthetic, transaction-generator-type tools?*
» *Are my monitoring tools hiding data from other devices?*

This kind of monitoring is similar to wearing a wrist device that measures your blood pressure, pulse rate, and temperature and routinely sends that information off to your health care provider. The program receiving this data looks at static and dynamic thresholds. It recognizes that a rise in your heart rate from soccer practice every afternoon at four o'clock or from time on the treadmill every morning is to be expected. A rise in your heart rate at an unexpected time, however, would signal an alarm. With that in mind, be sure to answer the following question:

» *Do your tools function as an external interrogator?*

Having an external interrogator is similar to having your doctor do more than merely collect your vital statistics. When the doctor asks you probing questions, your condition becomes clearer. Similarly, when your tools function as an external interrogator, you learn much more about the health of your servers.

Look at the Ultimate Metric

Another principal, and often overlooked, measurement is to keep a tally of all sub-one outages that affect significant numbers of people (five to ten people or more). These should be automatically reported as part of your monitoring strategy. For example, you shouldn't have to wait for a call at the service desk to say the Amsterdam switch to the network has gone down.

The problem with many monitoring tools is that they generate a tremendous number of false signals. You don't know which signal is valid and which one is not. Address these questions:

» *What is the monitoring strategy, particularly for sub-one issues related to the tier-one systems?*
» *How effectively do I use that data stream to identify problems in advance of my users?*
» *Am I keeping statistics around my ability to resolve those problems — the true mean time to repair (MTTR)?*

Become Proactive Instead of Reactive

So many IT shops are overwhelmed because they don't have the bandwidth, skill sets, tools, or cultural model to do things proactively. And they're sometimes rewarded for tactically solving problems in a crisis.

There's almost an expectation that they'll lose one of their tier-one applications each quarter, because every quarter some new disaster happens. This quarter, the problem might be related to the ERP system. The next quarter, it might be related to the supply chain or customer support, and so on.

This cycle of disasters becomes expected, but it shouldn't be. You should be building quality into the front end.

This cycle of disasters becomes expected, but it shouldn't be. You should be building quality into the front end. This seems counterintuitive when bonuses are sometimes doled out for solving problems after hours. But that kind of thinking is akin to dramatically saving the patient when he has a heart attack, as opposed to helping him stay healthy by eating well and exercising regularly.

Automation Is the Path to Data Center Efficiency

Automation can keep your data center running smoothly, improve efficiency, and help increase your data center's overall health. People can only scale so much, and automation helps your organization reach a level of maturity in processes that you can't easily achieve on our own. Repeatable processes bring an immense amount of value, ultimately affecting how well you satisfy the needs of the business. By addressing the questions discussed in this article, you can help ensure that your data center is healthy and well equipped to manage physical and virtual environments.

For more information about data center automation, visit www.bmc.com/solutions/esm-initiative/Data-Center-Automation.html.

END NOTES

1. "Best Practice: Executing Release Management with Change and Configuration Management," Gartner, Inc., May 10, 2010.

Steps to Data Center Maturity and "Wellness"

Here are some of the key areas in which you can improve data center efficiency and reduce costs through policy-driven and operator-assisted automation of processes and activities. This approach is based on Business Service Management (BSM), a comprehensive approach and unified platform for running IT.

CONFIGURATION AUTOMATION: SUPPORTS RAPID CHANGE

Configuration automation automates client (PCs), network, server, and application deployments and changes across physical, virtual, and cloud environments. Leading industry analysts have estimated that 80% of incidents are related to configuration changes either being made improperly or out of process. Because changes made through configuration automation are made in a predictable, repeatable manner and according to policy and process, they're done consistently and correctly, without incident. This dramatically increases the percentage of successful changes made in the environment.

COMPLIANCE AUTOMATION: ENABLES CONTINUOUS COMPLIANCE

Compliance automation assures that PCs, networks, servers, and applications remain compliant with policies and guidelines across physical, virtual, and cloud environments. This ensures that security patches are current and that configuration settings do not allow unauthorized access or use. Because changes are made in a consistent manner, compliance becomes a byproduct of the process. This approach increases the number of devices in compliance over time.

PROACTIVE AVAILABILITY AND PERFORMANCE MANAGEMENT: AVOIDS OUTAGES

Industry studies have shown that most IT organizations are reactive and often fail to meet Service Level Agreements (SLAs). With a BSM approach to availability and performance management, you can avoid outages and extend mean time between failures for business services. This approach includes event automation, which automates the detection and processing of events from infrastructure and application resources across the enterprise, including both mainframe and distributed, as well as virtual and physical, systems. Events are analyzed based on known normal behaviors. Abnormalities, those events that are outside of the normal operating range, are then correlated and further analyzed against rules and policies so that the root cause can be automatically identified. This improves availability and efficiency while reducing costs.

ABOUT THE AUTHOR

Mark Settle, chief information officer for BMC Software, joined the company in 2008. He has served as the CIO of four Fortune 300 companies: Corporate Express, Arrow Electronics, Visa International, and Occidental Petroleum. Settle has worked in a variety of industries, including consumer products, high-tech distribution, financial services, and oil and gas. He received his bachelor's and master's degrees from MIT and a PhD from Brown University. He is also a former Air Force officer and NASA program scientist.

**BUSINESS RUNS ON I.T.
I.T. RUNS ON BMC SOFTWARE.**

Business thrives when IT runs smarter, faster, and stronger. That's why the most demanding IT organizations in the world rely on BMC Software across distributed, mainframe, virtual, and cloud environments. Recognized as the leader in Business Service Management, BMC offers a comprehensive approach and unified platform that helps IT organizations cut cost, reduce risk, and drive business profit. For the four fiscal quarters ended June 30, 2010, BMC revenue was approximately $1.92 billion.

bmcsoftware

175720