

Automating the Data Center

The First Steps Make All the Difference

By: Harris Kern's Enterprise Computing Institute

The past decade had brought tremendous technological advances to the data center: open systems, client/server technology, distributed platforms and web-based computing. While these changes have allowed organizations to take advantage of unprecedented new business opportunities, they have also presented IT management responsible for running the data center environments with a host of new challenges. There are more systems and platforms to support, a steady flow of emerging applications to incorporate, new processes to manage and fewer qualified technical professionals available to meet these needs. In addition, IT executives are facing escalating pressure to facilitate the sharing of back-office information across all lines of business, divisions and departments, putting a heretofore unseen emphasis on enterprise infrastructure. In short, competitive mandates are forcing businesses to stay online 24 hours a day, seven days a week. To complicate this even further, enterprises are expanding portions of their operations into Internet-based co-location facilities, presenting the data center with the added challenge of managing high-availability production operations in physical computing environments that do not belong to them.

In response to these forces, increasing numbers of organizations are exploring automation techniques so their data centers can provide more-sophisticated levels of computing at lower costs. Performance monitoring, storage and network management are just some of the data center components that are becoming targets for automation. When properly defined, planned and implemented, automation of these components offers savings in staff time, an increase in system reliability, a faster resolution of problems and less downtime.

But notice the caveat: when properly defined, planned for and implemented. Many automation efforts stall because the scope of the project is unrealistic and the components targeted are not strategic or, often, because the specific terms and benefits of automation have not been clearly delineated. Additionally, inherent complexities, such as managing new data center processes and staffing issues, are not addressed during the planning and implementation of automation initiatives. However, when data center automation projects have well-defined goals and parameters, and when the implementation methodology is systematically laid out to meet those goals, companies of all sizes can achieve cost-effective, predictable IT operations and make significant strides toward the ultimate aim of a streamlined and cost-effective data center.

Define For Success

So where to begin? The first step for any organization considering automation for their data center model is to define and set their parameters. What does this mean? It means understanding what *automation* really means and how it applies to your specific situation. Defining and setting automation and benefit parameters will help companies get expectations in line with what's achievable, as well as help them identify the technical and production components they should focus on to best ensure success for their individual business imperatives.

In the broadest sense, automation can be defined as those processes or components that can be monitored and provisioned without human intervention. If you really think about it, in the context of the data center operations, automation encompasses little more than usage and capacity monitoring and provisioning to support the availability of your system, its resources, and the performance of your product or application.

The summary technical components that must be accounted for when considering usage and capacity monitoring and provisioning, include: hardware, operating system, network components, I/O, the database, supporting storage and batched processes (i.e., database loads, etc.). Non-technical components include those processes associated with change management and disaster recovery.

When analyzing automation opportunities, it is very important to consider the production processes – those procedures that support the technical processes you are targeting – that ultimately control the technology and the dynamic changes that automation will bring. For example, if you are implementing technology that will ultimately effect automated provisioning of storage, but bypass altogether the change management policies and protocols that the data center has in place, you will have a major problem. This lack of planning for change management is one of the biggest stumbling blocks encountered by IT management and one that is critical to consider before implementing new technologies.

Once automation has been defined, companies investigating data center automation must identify the strategic components most suitable for automation, critically analyze the complexity of those automation processes, and then assess—as realistically as possible—the potential results. Doing this successfully is no easy proposition, as many data center directors can attest to. But by understanding the true definition of automation, identifying the components most appropriate for automation and assessing the complexity involved in automating those components, businesses can realize the returns they expect and implement a successful automation initiative.

Identifying Automation Components

Identifying the number and types of operating environments under consideration is also an important part of preliminary planning. For example, does your business have data centers spread throughout the world or a variety of co-location installations? Once you've identified all the different geographies under consideration, you can start breaking those environments down by the following components:

- Physically Distributed Environments (Geo)
- Hardware
- Operating systems (platforms)
- Network platforms (backbone, wide area network, internet, etc.)
- Interfaces (I/O paths)
- Databases
- Storage (subsystems)

Automating each of these technical components requires unique considerations and technologies, but the same evaluation process holds for all: Develop a clear-cut and realistic list of goals for each component, analyze the complexity of the automation required to meet

those goals and understand the gap between your business' current level of automation and the targeted level. In addition to ensuring that your automation project will meet your budgetary constraints and be in line with your business strategies, this honest evaluation will also help you with another critical part of the automation initiative: making sure you address your requirements as to how you will interface to existing "production" requirements, such as change management issues.

To do this, let's go back to the storage example we used earlier. You understand the definition of automation, you've identified the technical components suitable for automation and you're now examining the technologies that will effect automatic provisioning of storage—one of the key technical components under consideration. Because storage management involves the managing of disk and tape archival requirements through backup/restore, catalogue and storage functions, the ultimate goal of automation would likely include the following:

- automatically provisioning new disk space, within stripes, volumes and hardware units available.
- automatically notifying staff when hardware is to be dealt with.
- automatically provisioning for backups to start at scheduled times on system-labeled tapes and for data files to be automatically restored for users;
- automatically provisioning for robotic tape loading and archiving with built in cataloguing and organizing tapes for third-party vendor pick-up and delivery;
- automatically provisioning for off-site storage of archives;
- and automatically provisioning for third-party vendor contacts for return of tapes.
- throughout all, automatically communicating with the appropriate members of the IT enterprise that changes are taking place, triggering associated processes (i.e., billing or charge-back, etc.).

These goals, you believe, will make the data center run more efficiently as well as offer benefits throughout the organization.

Aligning Goals with Reality

Time for a reality check: while these goals are admirable—and would certainly go a long way toward making the storage component of the data center more efficient—they may be far fetched, depending on your current situation. Understanding this before deploying automation technologies is critically important to the project's ultimate success and acceptance. In this example, some of the goals listed above are realistic, but others may not be. For example, while the first one is achievable with the right technologies and processes, the second might prove to be tricky.

Once the goals have been listed—and given a good reality check—then it's time to assess the complexity involved in automating those technical components. Going back to our storage management example, we can determine that the complexity level of automating those functions will be in the medium-to-high category; medium because many tools exist today to address these automation requirements, but high because the complexity increases with the number of platforms under consideration. These days, most data centers these days operate off a number of platforms so expect the level of complexity to be quite high. Knowing in advance what level of complexity you're facing can not only help you plan for the attendant process changes that will be involved, but will also help you to prepare senior management for the road ahead and effectively sell your automation project to the finance department.

For budget and time considerations, you may very well end up choosing some platforms and subsequent configurations over others in your first phases.

Now that you have developed a working list of the components of your data center that seem suitable for automation, reality-checked your goals for these components, and identified the complexities involved, the next step is to analyze the current level of automation versus the realistic automation goals. As with complexity assessment, this “gap analysis” will provide valuable assistance in financial planning and selling the initiative to management.

Planning Buy, Build and Deployment

It is now time to move on to planning the design and deployment of the solution. While this phase obviously requires more tutelage than can be effectively covered in the scope of one article, a broad-stroke outline of some key points for consideration can help ensure the success of an automation project.

The design work for the automation of the technology architecture and the administrative processes should occur both separately and simultaneously. Doing this planning in tandem is important because effecting the automated provisioning of storage without considering change management policies and protocols, for example, is a recipe for disaster.

Creating a design plan for the technical architecture basically comprises evaluating and choosing the system management software to support the technical components identified for automation. This phase of the design planning should be completed with the participation of a project manager and will result in system management software solutions that have implications for system, storage, and peripheral hardware; network hardware and software; and database applications.

The design process is fairly straightforward, but that doesn’t mean that it is easy. As a starting point, however, it helps to use the automation criteria outlined in this piece to develop requests for proposals to be distributed to external vendors and staff. The responses will help you determine what is required to meet your automation goals and assess whether the solution should be purchased from an outside vendor, developed internally or created by a combination of the two.

In this planning phase, give special consideration to the platform architecture requirements because features such as supportability and cost will have substantial impact on system management software and the automation of operations. Additionally, multiple platform requirements will be the most expensive to fulfill.

As you plan for technical architecture design and deployment, consider your budgetary constraints and decide whether outsourcing the product(s), implementation, and/or ongoing service is the most cost-effective way to proceed. Like most other aspects of data center automation, the outsourcing question is a complex one the answer to which could fill (and has) an entire book and a careful assessment of the technologies, processes and financial considerations of your outsourcing project is in order when weighing this question.

Finally, designing the automation to interface with administrative processes is often overlooked. It involves more than the purchase of new tools. It requires taking an inventory of existing processes and refining them to support the new levels of automation. Without this work, the entire initiative can fail. In fact, its omission is the most common cause of failure for data center automation projects. So, at this phase, IT managers must be developing strategies to enact the following production process components:

- Production assessment—process that identifies the operational requirements to implement and manage new and changing applications
- Problem management—a centralized process to manage and resolve user network, application, and systems problems
- Change management—a process that coordinates all changes that affect the production environment
- Asset management—process to query, discover, track, and store enterprise-computing resources, including hardware, operating systems, and applications
- Disaster recovery—process to enable recovery in the event that a disaster should render mission-critical systems inoperable.

Again, I cannot emphasize the importance of planning to interface heavily with these administrative processes enough. If you do not provide resources for process redesign and training, your initiative will surely fail.

As the speed and reach of business grows, increasing numbers of companies—both large and small—will turn to data center automation to leverage new information opportunities, streamline processes and cut costs. By understanding what automation means and entails, and planning appropriately for both the technical and administrative processes, IT managers can do much to ensure the validity and long term success of their initiatives.