Managing Application Upkeep and Performance

Executive Summary

Application management has become complex and time-consuming. Many application types – from customer relationship management (CRM) and enterprise resource planning (ERP) to voice, video, data and even telepresence applications – are converging onto a common infrastructure. This paper explains what's involved in managing applications in this new networking environment, whether done in-house or using a managed service provider.

at&t

Introduction: Today's Management Landscape

A number of factors are making the application management landscape more complex. First, global mergers and acquisitions are driving the need for interoperability between and access to disparate networks. Data centers are becoming virtualized as businesses decouple application software from physical hardware resources to improve efficiencies. Meanwhile, many application types with varied behaviors have converged or are being introduced onto a common infrastructure, resulting in complex environments that are difficult to mange and demand broader levels of network and application expertise. The expertise needs to continually evolve as new technologies emerge and are integrated into the new network.

Adding to this complexity are cost pressures of a down economy and the need to comply with the latest corporate governance mandates, such as Sarbanes-Oxley (SOX), Gramm-Leach-Bliley (GLB) and Healthcare Information Portability and Accountability Act (HIPAA). To contain costs and comply with governance procedures, many organizations are centralizing their application management and reporting. At this juncture, businesses must make an important decision: Should their own IT staff take on installing, maintaining, managing, patching and supporting all the business's applications? Or is it beneficial to contract for third-party assistance with some or all aspects of application management?

As a general rule, the larger and more complex its environment, the more likely the business will be to benefit from a managed service partner's economies of scale and rich pool of resources. This paper examines a few steps that should be taken when rolling out a new application, as well as some basic decisions associated with considering whether to manage applications in-house or use a managed service provider.

Step One: Choosing Apps and Evaluating Existing Resources

First, the enterprise must select the appropriate application software that will meet their needs. For example, a business unit will decide upon a particular inventory management, human resources, payroll or customer relationship management package.

From there, the organization should examine the capacity and utilization levels of its current computing infrastructure (servers, storage, desktops) and networking resources (data center switches, LAN switches, WAN routers, WAN services). This step will help determine whether any additions or upgrades are required for the new software module(s) to run well and whether the new application(s) will impact the performance of existing applications. To determine the physical infrastructure capacity, it is important to have operational processes and procedures in place to ensure that this is done accurately and in a timely manner. Each of these is discussed in this section.

Server and Storage Capacity

From a hardware perspective, a technical architect should be used to determine the existing utilization levels of the computing infrastructure in the data center and on user desktops. That information will then enable the architect to determine whether there is enough CPU, memory, disk and storage space to accommodate the new software, while allowing some wiggle room for growth. This determination should take into account such variables as the size of the user population and number of items in the application database.

Disaster Recovery

Despite the trend toward data center consolidation, using at least two redundant data centers to host core business applications are recommended for disaster recovery. If designing their own network, the enterprise should set up mirrored server farms in at least two alternate data centers with requisite real estate, power and cooling. Data replication to keep them in sync with one another would be required.

This is generally achieved by installing storage-area network (SAN) devices that support mirroring in each site and writing custom scripting so that the proper data is replicated consistently between or among data centers on the appropriate timetable. The enterprise then sets up all the processes and procedures for data center failover, and maintains a network operations center (NOC) for monitoring all activity.

If the enterprise is using a provider to help design the solution, hosting the applications in the provider's redundant data centers is an option. In this case, the enterprise provisions connectivity to the provider's data centers, and the provider leverages its setup for redundant site connectivity, replication, storage and mirroring on behalf of the enterprise. One connectivity option in this type of setup is for the enterprise to turn the provider's data center and NOC into network connectivity spokes of its WAN service. This can be a smooth implementation if the enterprise has a Multiprotocol Label Switching (MPLS) WAN service, which inherently provides any-to-any connectivity among sites without having to set up individual virtual circuits.

Here again, having a single entity in the roles of both the hosting supplier and the network service provider bodes well for end-to-end infrastructure and network management flow and ultimate performance results.

Making sure that there is enough computing power and storage to serve both the existing application environment and the new software is a critical pre-deployment step. If the computing and storage infrastructures prove to be inadequate, the environment will have to be upgraded before the application(s) is deployed. That upgrade can be handled by investing directly in additional capital or by tapping into the resources of a managed services partner. These options and their capital and operational expense (capex/opex) tradeoffs are discussed in more detail in the section, "Ownership and Responsibility Considerations."

The same utilization calculation exercise that has been performed for the computing environment should be performed with respect to network capacity and equipment ports.

Network Capacity

Just as there must be sufficient computing resources to take on the addition of a new application module(s), there must also be enough network resources to handle the added load. Enterprises should understand what applications are on the network, the current network load, application behavior and daily traffic patterns in order to accurately plan for the impact of the new software on the existing environment. The network resources to evaluate include LAN switch ports (both the number and speed required), wireless LAN infrastructure capacity, WAN router ports and sizes and WAN service bandwidth. Enterprises should also understand their disaster recovery requirements and design the WAN accordingly (see disaster recovery excerpt).

Many enterprises are unaware of all the traffic that is actually traveling on their network. They might assume a certain utilization level based on the application traffic they know about, which might not be a complete picture. Enterprises are often surprised by some of the applications they find on the network as well as the bandwidth the applications are consuming. Once recognized, the organization can decide whether to block, "shape" (limit how much bandwidth it can consume) or otherwise prioritize traffic to manage capacity and accommodate new application traffic. For this reason, enterprises should include some upfront 'application aware' network monitoring that identifies specific applications that are traversing the network. This requires the ability to examine packets at layer 4 or above using either a specialized appliance or a Netflow-enabled router on the edge of the WAN. This information is reported to a centralized reporting tool that analyzes and summarizes the data so the organization can see which applications are driving network utilization at various points in the network, and understand the relative performance of key applications across the WAN.

The network capacity calculation is particularly important in large distributed organizations that rely on WAN communications for much of their user access to applications. Independent software vendors (ISVs) tend to build and optimize their application suites for use on high-speed LANs, while the WAN is often more bandwidth-constrained and incurs distance-induced delays that impact application performance.

New networking trends, such as virtualization, are pushing more applications across the WAN, making it critical to identify appropriate network capacity needs. In fact, special tools and techniques are sometimes required to make the application perform as well across the WAN as in the LAN environment. The function of this software and equipment is often referred to as "WAN optimization" or "application acceleration," and it can be acquired and managed by the enterprise in the form of customer premise equipment (CPE) or procured as a managed network service from a network service provider.

An enterprise may already have processes and procedures in place internally to measure capacity levels and help ensure that the appropriate hardware and network resources are available. If not, they will need to develop them. Alternatively, enterprises can use the buy-it approach. Using a managed service, they can leverage the real-world metrics that the service partner has accumulated over many years from working with multiple clients of similar sizes and with comparable requirements.

Scenario: Accelerated Deployment Needed

An enterprise decided to spin off a division of 10,000 employees. It wanted to retain its own IT staff and Oracle application expertise within the parent company, yet wanted the new division to be open for business and publicly traded right away. The new division engaged AT&T to migrate the spinoff from an IBM to Linux hosted computing environment running redundant Oracle Real Applications Clusters (RAC). Because AT&T had the expertise and resource scale required, it pre-staged the application on hardware already on hand and had the new division up, running and publicly traded in 30 days. In this case, buying expertise in the form of a provider helped enable a quick application deployment to the new division.

Scenario: Better Performance Needed

A company with a large warehouse had a poor-performing inventory management application. As workers picked products from the warehouse and put them on pallets for shipping, they used bar-code scanners to deduct the products from the company's inventory count. Each scan took 15 to 20 seconds, in part because the IT team had fallen behind in the thousands of software patches required to fix bugs.

Once an enterprise falls behind in patching, trying to apply selected new patches can cause errors because certain prerequisites aren't in place from the patches that have been skipped. The company turned to AT&T, which used proprietary tools it has developed over the years to inventory the application itself. AT&T was able to quickly update the application and improve scanning performance to 1 to 2 seconds. In patching the application, AT&T also fixed a problem in the same application suite's accounting package, which would disconnect after a user would enter 50 lines of data, requiring the user to start over. The culprit, again, was out-of-date software patching. In this situation, experience in inventorying applications enabled AT&T to help the manufacturer reduce scanning times and improve productivity in both the warehouse and accounting.

Scenario: Better Performance Needed

A global holding company centralized the hosting and management of its human resources and payroll applications with AT&T. After tracking its savings over a five-year contract, the enterprise determined that it had saved $4.2 million over the contract term. Based on its experience and this cost measurement, the company handed over the hosting and management of all back-end applications to the service provider.

Organizations should also frequently revisit their overall operational processes and procedures. In today's dynamic computing environment, processes must be flexible enough to adapt to change. For example, old ways of administering an application physically tied to a single server, may not be adequate in a virtualized server environment.

This is why the next step is to determine the expertise and support requirements of the new application environment, as discussed in the next section.

Step Two: Understanding Expertise and Support Requirements

Determining whether an organization has – or can cost-effectively procure – the relevant experience for managing the varied mix of application types sharing a computing and network infrastructure is the next important step. Most organizations select their core business applications based on what best fits their business needs, rather than on the existing skill set in the IT department. Once the application module(s) has been selected, evaluating whether the appropriate expertise exists in-house is imperative.

Note that some enterprise-wide application suites might have thousands of online screens with 15 to 20 configuration options each and will likely require some degree of customizing. So the enterprise is likely to require staff highly seasoned in the intricacies of the application module or suite.

Generally, installing and managing a business application suite will require a group of specialists in the following areas:

- Web server technology being used
- Database type
- Programming language
- Operating systems (OSs) in use
- Actual application module(s)

If these skills aren't part of the IT department's expertise, the organization must determine if they are available for hire, at what cost, and what the learning curve is likely to be. The recruitment process and training for application-specific expertise will add time to the test and rollout process, so those factors should be built into the project timetable. Perhaps most significantly, it can take an individual several years to become fully educated about the intricacies and best practices associated with a given ISV's application software suite. The learning curve, then, should account not only for basic training but also for initial missteps and reconfigurations while the specialist(s) gets fully up to speed.

One way to procure these resources is through an outsourcing provider that hires, trains and hones senior professionals steeped in the expertise needed for a given application. Because these individuals often have many years of experience in the application at hand, they can be counted on to provide the requisite amount of experience from day one, which is a significant consideration. How this expertise is reflected in accelerated deployment times, reduced costs and application-performance improvements is demonstrated through scenarios described at the end of this paper.

Step Three: Choosing which Parts to Own and Manage

Managing the Application: Servers, Storage and the Network

Based on internal expertise, timeline to implement and other IT priorities of the company, enterprises need to determine whether they have the trained in-house personnel needed to successfully deploy a new application. A provider has years of experience with applications and understands computing and storage requirements. They can rely on best practices established through years of experience to guide them in understanding application needs.

Let's consider server procurement and management. An enterprise using the build-it approach purchases its own servers and operates them in its own redundant data centers. The enterprise then procures network connections to both data centers from its various business sites and also installs a heavy-duty communications link between data centers for backup protection.

In this scenario, the enterprise can opt to manage the applications running on the server farm on its own or it can use remote management services from a provider. In the latter scenario, a provider uses network links to remotely monitor, manage and troubleshoot applications, as well as to automatically push regular patches and security updates to the devices. The gain for the enterprise is in alleviating the opex associated with the human resources needed for the ongoing management tasks. With this solution, the enterprise does not realize any savings in real estate, power and air conditioning associated with housing the hardware itself.

To improve capex, the enterprise can also opt to outsource the hosting, as well as the management, of its application suite. In this instance, the business would turn to a provider who would use its own data center resources for hosting applications, enabling the enterprise to deflect costs in real estate, power and cooling.

Who Owns the Hardware?

Regardless of where the hardware resides, the enterprise can opt to buy it and own it at the end of the contract. In the case of using hosted services, the enterprise then effectively leases data center space for the duration of the service contract.

Alternatively, and often most cost-effectively, the partner can provide the server hardware as part of its hosting service. In this case, the services provider would purchase the servers and own them at the end of the contract. In such a scenario, usually the enterprise is charged a one-time upfront fee for an application-hosting rollout and then a fixed monthly recurring fee covering ongoing deployment, provisioning, maintenance and support services.

Enterprises can take this model one step further and outsource the server hardware in a utility computing model. In this case, the provider generally runs a virtualized server environment. Enterprises rent the use of the provider's capital resources in a time-sharing manner with other customers, much as they would do with metered electricity, gas and water services, to leverage economies of scale and reduce hardware expenses.

Application Performance on the Network

A trusted third-party provider has likely turned years of experience and metrics into a set of best practices, such as how much network connectivity to apportion to each user for a given application. In that respect, the provider is likely to be accurate and quick in determining if, where and to what degree, the infrastructure should be expanded.

Many of these providers also have the capability of conducting deep packet inspection to see what applications are traveling on the network. Through traffic prioritization and WAN optimization, enterprises can enable optimum performance of both the new and the existing network applications.

About AT&T's Application Management Offerings

AT&T's Application Management Services help enterprises realize the full value of business applications, such as customer relationship management (CRM), enterprise resource planning (ERP), and supply-chain management (SCM). We take over the day-to-day tasks of managing your application including patches, fixes and updates so you can move past "keeping the lights on" to focus on what's really important – the growth and profitability of your business. In addition, AT&T offers remote management for clients that have already invested in the physical infrastructure for their application environment but need support for the management of their applications. AT&T's broad application management portfolio include service offerings from Ariba®, IBM WebSphere Commerce®, Microsoft®, Oracle®, PeopleSoft®, SAP® and Siebel®, as well as Software-as-a-Service (SaaS) enablement solutions.

Network Management and SLAs

When using hosting services, it can be beneficial to closely consider an outsourcing partner that is also a global WAN service provider. The reason is that all the piece parts of the outsourced solution can be provisioned from one entity that has visibility into and control over all those components, resulting in an optimized end-to-end experience.

For example, the application management group at AT&T works closely with the network services group to determine bandwidth needs and establish service-level agreements (SLAs) for the network. Network management and troubleshooting, of course, are significant contributors to the overall availability of an application. This pairing, then, combines software, hardware and network SLAs to deliver an overall application performance SLA.

This SLA is likely to be important to those corporate IT departments that are accountable to their internal users for establishing and upholding certain levels of application and network service quality and availability. Doing so becomes more challenging in highly distributed organizations, which often can't afford to place IT staff in each site. Guaranteeing service levels in such situations, where service depends in large part on the WAN, is elusive when the IT department can control only so much of it.

Application performance SLAs require that the services partner can see and manage computing software and hardware, either on the customer's site or hosted on the provider's site, as well as the network. The combined hardware, software, application and network uptime SLAs essentially reflect a composite guarantee of the minimum end-user experience.

Conclusion

IT departments are challenged to get their arms around large and growing complex application deployments while scaling their environments with the requisite hardware, software, network and human resources. As such, they are increasingly considering service offerings from large, trusted partner organizations that have turned years of experience into best practices for application-specific deployment and management. Such partner organizations employ human resources steeped in experience with the particular application(s) to be rolled out and leverage their own specialized tools and metrics databases to improve an enterprise's capex, time-to-deploy, application availability, ongoing patching and support capabilities.

When selecting an application management services partner, distributed enterprises should remember the role that network management plays in overall SLAs and end-user experiences. Choosing a global WAN services provider that also offers computing infrastructure services, application hosting and management can help raise SLAs, because the provider has visibility and management control over all the components that contribute to end users' application experiences.

For more information contact an AT&T Representative or visit www.att.com/business.

06/24/10 AB-1592-01

at&t