

# Linux Scalability Testing: Part II — Economics and the Linux Solution

**L**inux for S/390 and the virtual server concept can be used to dramatically reduce the cost of ownership for a large server farm.

**A**s introduced in last month's article on the Test Plan Charlie environment, Linux on System/390 is providing a new and innovative model for designing and deploying large-scale Internet service provider (ISP) and Internet data center (IDC) infrastructure solutions. The testing done in David Boyes' 41,000+ Linux virtual server farm not only demonstrates that open source tools and the IBM VM/ESA operating system are a natural match for a solution requiring large horizontal scalability, but also promises a new financial model for deploying hosting services.

This concluding article discusses the business case for a Linux on S/390 implementation. Historically, the justification for workstation and mini-computer class systems has been based solely on the discussion of acquisition price for the base hardware; only a few years ago a simple upgrade, such as a 64MB memory upgrade for a medium-range S/390 processor, could easily have cost more than \$50,000 USD. Today, given IBM's efforts at reengineering the S/390 to use commodity parts without sacrificing the system's legendary reliability and self-management features, it is more difficult to justify the financial case for workstation-based computing vs. centralized large system computing. A more sophisticated analysis is necessary; one needs to consider not just the hardware, but also the cost of operating and maintaining the facilities and communications infrastructure required, the staffing required

to deploy, maintain, and operate the environment, and the time required to deliver a solution to the marketplace.

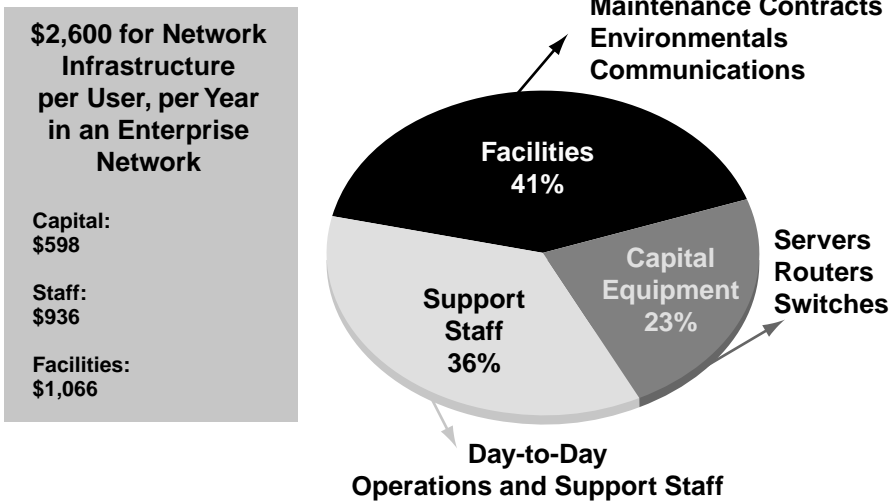
To illustrate this case, let's look at some of the elements of a cost justification case. As you can see in Figure 1, the total cost of ownership for a network-based solution falls into the following three distinct categories:

**Time-to-market is a deciding factor for most Internet service provider (ISP) and Internet data center (IDC) customers — it represents how quickly they can deliver information content to their potential customers; a matter of hours can make the difference between catching or missing a new trend or providing additional capacity in a crunch vs. failing to meet customers' needs.**

- hardware — the capital investment in boxes and infrastructure — things you can touch and photograph
- facilities — not just floor space to house the boxes, but also the environmental requirements such as power, air conditioning, maintenance contracts, connectivity and communications bandwidth, and other ongoing management and operational costs to keep the solution up and running
- staffing — the people required to deploy and operate the environment

Based on these factors, empirical observation in the enterprise and ISP industry reflects the percentages of Total Cost of Operation shown in Figure 1. It may come as a surprise that currently the cost of hardware is by far the smallest portion of the operation cost. The differentiating elements are the cost of staffing and facilities, which make up more than two-thirds of the total. By concentrating on these two elements, it is possible to provide additional value by attempting to minimize or

FIGURE 1: NETWORK COST OF OWNERSHIP



consolidate the management and ongoing operations cost. This can be accomplished by either providing economies of scale or by optimizing delivery in ways that reduce the requirements for operations support systems (OSS) and the accompanying installation, configuration, or enterprise staffing requirements.

In the case presented here, an additional element is added that represents a significant competitive advantage in today's marketplace: time-to-market. This term reflects the speed at which a service can be set up, configured, managed, and billed (note that billing is a critical point — in most cases we are not operating charities) to a specific customer. The time-to-market depends on a number of different elements, including time to purchase, mount, and configure the system for its initial purpose; integration of the server into its management infrastructure; setting and programming any triggers or

other alerting functions; and delivery of necessary access information to the customer in a timely manner. In most cases, large, established hosting providers have delivery times averaging five to seven days — some larger providers have used economies of scale and onsite inventory to reduce this time to three to four days at the cost of maintaining a stock of equipment subject to depreciation and aging. Time-to-market is a deciding factor for most ISP and IDC customers — it represents how quickly they can deliver information content to their potential customers; a matter of hours can make the difference between catching or missing a new trend or providing additional capacity in a crunch vs. failing to meet customers' needs. Empirically, customers often opt for services that can be deployed quickly over lower price, relying on the "first with the most" strategy to offset the financial impact. To borrow a phrase from a

Cisco staffer, "The competition is only a click away."

## CONSTRUCTING THE CASE

When discussing cost of ownership, it's important to present a realistic comparison — credibility is everything when talking to financial and senior executives, so it is necessary to ensure that the same common elements are included in each comparison case. The next few sections outline some common elements that can be used to demonstrate how Linux on the S/390 provides an interesting solution.

### Compare Hardware Costs

In most cases, this is the first element that comes to mind when comparing large system-based and workstation-based solutions, and most discussions of total cost of ownership begin here. In the Test Plan Charlie case, the original plan presented by the consulting firm represented a significant number of relatively inexpensive machines — in the \$6,000 to \$7,500 USD range — plus a number of more expensive systems for I/O-intensive applications averaging \$30,000 USD. The table shown in Figure 2 summarizes the hardware requirements for the initial deployment using the discrete system design.

Note that while each discrete system is relatively inexpensive, this costing does not include the additional infrastructure required to make these systems usable and manageable, nor does it provide for additional expansion in capacity without incurring additional system cost. Although the larger, up-front cost of an S/390 machine may deter some users, careful analysis indicates that the less expensive per-box cost may not be the best deal, after all. The comparable S/390 design is summarized in Figure 3.

While the total amount of CPU power (in terms of raw cycles available to do work) is greater in the discrete system solution, consider that in most cases idle cycles in large server farms are wasted due to the inability of idle cycles in one physical machine to be applied to a lack in another area. In the virtual server environment, cycles not needed by one virtual image can be applied to another image, allowing a smaller (on paper) CPU to actually perform a larger amount of useful work and to expand for supporting additional function without additional investment in equipment (the

FIGURE 2: DESCRIPTION OF DISCRETE SOLUTION

	Number of Systems	Avg Cost per System
Sun UE2, 2G, 2x9.1G disk, 1 quad Ethernet	500	\$6,500
Sun SC1000, 2G, 1 quad Ethernet, 30 bay RAID enclosure, 11x18.2G	250	\$30,000
Total	750	\$10.75M

FIGURE 3: DESCRIPTION OF LINUX ON S/390 SOLUTION

	Number of Systems	Avg Cost per System
IBM 9672-R46, 8GB storage	1	\$945,000
IBM ESS DASD (3.4TB)	1	\$545,000
Total	2	\$1.49M

9672 scales from one virtual system to tens of thousands of virtual images using VM/ESA). Thus, we can conclude that basing a decision completely on hardware cost is often misleading.

### Compare Facilities Costs

To build on that idea, we turn next to the largest portion of the cost: facilities. Costs of facilities include a large number of elements, but can be broken down into three major categories:

- environmental costs — such as floor space, power, air conditioning, cabling, and rack space
- deployment costs — such as the time required to build the system to order, or to ensure deployment in a supportable environment
- management costs — which include the cost of operating the server in terms of maintenance, capacity planning, system command and control infrastructure

It's important to consider all three elements in the analysis, as they tend to be difficult to quantify in absolute terms due to a general industry-wide lack of experience in operational management techniques.

For illustration, let's return to the telco case, which was presented in Part I. In this specific example, the customer would have required 10,000 square feet of floor space including rack space, accessways, control center space, and utility access spaces. Each set of customer servers consumed one-half of a rack, thus requiring 125 4x4 foot racks at more than \$3,000 each, plus an additional 30 racks for network and LAN equipment. The 10,000 square feet of space (at \$27 per square foot per month) required a \$750,000 battery backup installation and more than 12 miles of network and power cabling to support the initial 250 customers.

Large-scale system management suites such as Tivoli or CA-Unicenter promise significant management and control efficiencies; however, the implementation cost of these systems is also a substantial investment. In the telco's case, to provide adequate network and system support management of the proposed environment, an additional 20 Sun SC1000 servers were proposed to support deployment of the Tivoli TME 10 enterprise management suite. The Tivoli

software costs more than \$30,000 per management server, with an additional license cost of \$350 per customer server to provide visibility, console management capability, applications management, and backup capability for the original 250 servers.

Again, let's contrast the same solution based on a virtual server environment. An S/390 G5-class CPU and an IBM Shark disk cabinet occupy approximately 400 square feet of space. Only four external racks were required to host the Cisco 7513 routers connecting the G5 to the incoming IP network. The G5 server and VM/ESA provided virtual system replacements for the LAN and server infrastructure, reducing the total cost of the hardware to \$375,000 for Cisco routers and hardware, and to \$676,000 for the G5 and Shark systems.

As detailed in Part I, all systems within the complex ran virtual machines with Linux for S/390 serving all system and support functions for routing, WWW service, and other system management and configuration functions. The VM/ESA operating system provides system management and resource control tools as integral components of the operating system. No external control systems were required. In addition, the customer was able to implement a resource accounting system to ensure that customers exceeding their target performance goals were notified and prompted to upgrade their service.

### Compare Staff Costs

We now turn to the element least considered in most cases: who's going to do the work? The availability of trained staff in the data center environment is an acute problem in most places; experienced staff with operational experience command substantial premiums and salaries. Without competent staff, much of the facilities problem is exacerbated. In our case, the original customer staffing requirement estimate was an increase in network and systems operations staff of a minimum of 25 persons per shift for a 24x7 operations facility, a total of 100 full-time employees, once managers and support staff are included. Comparably, the integrated nature of the VM-based solution also allows the configuration and performance management of the complex to be completely automated, thereby reducing the operations staff to an incremental increase of three full-time employees.

Keep in mind that there are significantly more organizations and universities producing Unix- and specifically Linux-literate graduates than S/390 trainees, so the average price of a qualified S/390 systems programmer is considerably higher than the price for an average Linux-literate college graduate. There may be no cost savings here; however, the virtual environment lends itself to substantial automation impossible in the discrete environment, thus making a smaller number of resources reach substantially further. Note also that there is a cost associated with equipping staff that needs consideration as well — fewer staff require less equipment.

### Commercial vs. Open Source Applications?

Turning from operational requirements, we need to also explore the relatively slow adoption of Linux (on any platform) by the commercial software vendors. While in the ISP and IDC industry, this is less likely to be an issue, as most Internet services and utilities are already based on open source tools, this can be a drawback to the Linux-based solution. In our ISP case, the applications used (bind, the standard Unix Internet domain name server, and INN, the standard Usenet News server software) were already open source. There are additional examples of business-oriented tools such as the PostgreSQL database or the Apache WWW server that are certainly ready for production use.

Fortunately, the advent of Linux for the S/390 has prodded a number of large vendors such as Software AG, BMC Software, and IBM itself, to announce enterprise middleware such as MQSeries, BMC Patrol, DB2, and Tamino for Linux on a number of platforms.

Consider also that the richness of the open source community has produced some commercial-quality tools and capabilities that may in some cases provide an auxiliary source or replacement for an expensive commercial application. As an additional bonus, defect remedy tends to move much faster in the open source community than in the proprietary software world; serious flaws in sendmail or the Linux kernel, for example, are typically fixed in a matter of hours rather than days or weeks. Each environment has differing requirements, but in the cost case, you may be able to gain a substantial savings by

using open source applications in some portions of the environment.

### Compare Time to Market

Last, let's look at how an S/390 solution helps a company compete by delivering a solution more quickly than the comparable discrete solution. In the discrete solution, deploying a new service with significant service level agreements (SLAs) requires the deployment of additional physical systems to guarantee reliable service without the opportunity for one user of the server to "starve out" other users in denial-of-service attacks. Corresponding with the need for additional hardware is the requirement to configure, rack, and connect that hardware in a systematic fashion — a process that takes an average of five to seven days for most discrete solutions.

In the virtual server environment, in contrast, much of the configuration and set up can be automated to a very high degree, providing some significant advantages in all of the areas discussed so far:

- staff — fewer people are required
- facilities — fast deployment of preconfigured images
- hardware — virtual servers can take advantage of idle resources (vs. idle resources being wasted)

In both scenarios with our example, some integration is required, however, the tightly coupled nature of the virtual server environment reduces the time required to deliver a service dramatically — in this case, to 90 seconds per server image deployed.

### SELLING A LINUX SOLUTION TO MANAGEMENT

At times, the best rational case and decision doesn't make the case airtight for management to appreciate yet move beyond what seems to be "common" wisdom. To get the case across to a manager, often there are two types of justifications that need to be

made: a financially-oriented case focusing on cost savings and competitive advantage, and a technology-enabling case that focuses on the technical capabilities and additional possibilities the solution provides. In each case, the following considerations help generate the necessary elements to demonstrate the power of the solution.

## Each environment has differing requirements, but in the cost case, you may be able to gain a substantial savings by using open source applications in some portions of the environment.

#### Technical Case

The technical case might be more comfortable and easier to use for opening a persuasive argument. In our example, the technical focus revolved around the flexibility of the virtual system environment — being able to create, destroy and reconfigure servers quickly, and to react equally quickly to customer capacity and software requirement demands. Use of open source applications allowed for fast response to customer problems and queries, and the level of automation possible in the virtual system environment provided a significant technical advantage to minimize operational complexity and increase the number of services available with minimal additional hardware.

#### Financial Case


Correspondingly, the financial case focused on the total cost of ownership and acquisition. As presented in the hardware section, the cost of each individual dis-

crete server is quite a bit lower than the S/390. However, only one S/390 is required to replace 750 discrete servers, with substantial savings in environmental and additional management infrastructure. We also demonstrated a much shorter time to market, which offers a significant competitive advantage to the operator by allowing both lower prices and higher availability to the customer.

### CONCLUSION

As this two-part series has shown, Linux for S/390 and the virtual server concept can be used to dramatically reduce the cost of ownership for a large server farm. A similar deployment could easily be done for an enterprise application or other solution requiring similar scalability and manageability for LAN services or other tools. In general, the case for this deployment relies on two maxims:

- Compare apples to apples.
- Don't forget the cost of services and deployment.

These elements provide a strong basis for questioning the "common wisdom" and exploring how virtual servers can provide flexible and scalable service environments in your own organization. 

---

*Adam Thornton was first exposed to the System/390 as David Boyes' junior systems programmer at Rice University in the early 1990s. Since then he has spent nearly a decade in the penguin's clutches, having used Linux since version 0.09. Subsequent to graduate studies at Princeton, he volunteered as system administrator for penguinvm.princeton.edu, Princeton University's Linux/390 virtual machine. Recently, he became a founder and principal engineer of Sine Nomine Associates. Adam can be contacted via email at adam@sinenomine.net.*