

# Red Hat Applies to the Enterprise Club

## Research Note

Gordon Haff  
22 April 2003

With increased respect and acceptance for Linux come higher expectations and new requirements. It's a given that any OS in an important role must be robust and speedy, but those qualities are not enough. Both IT managers and independent software vendors (ISVs) also demand that OS developers ship predictable and consistent product updates—and that they do so on a schedule that is attuned to conservative and deliberate enterprise IT processes, rather than the frenetic pace of open source development. Enterprise customers also want support that is both responsive and cognizant of their need to control how and when updates and patches are installed. Linux distributions as a class have not lived up to either requirement.

Enter the enterprise-oriented Linux distro, which takes on more traditional operating systems on their own terms. Sure, a big part of its pitch is that it saves users money,<sup>1</sup> but its development, distribution, and support processes are starting to look more like those in commercial Unix than traditional Linux. And they'll be an increasingly big part of Linux' ongoing success in supplanting more established OSs for an increasing number of datacenter tasks.

Copyright © 2003  
Illuminata, Inc. Licensed to  
Red Hat for web posting.  
Do not reproduce.



Red Hat is the Linux distributor that has most emphatically defined itself as a distributor of business-class Linux—something complementary to, but clearly distinct from, less buttoned-down Linux distributions. And it makes no bones about having Windows and Unix squarely in its sights.

### The Lessons of its Forebears

The quality and capabilities of the code in the Linux kernel and major applications that run on it have matured considerably in recent years. However, in spite of improvements, the processes and tools associated with many Linux distros often fall far short of the highly systematic framework that the major Unix vendors have put

1. But, these days, what vendor doesn't pitch such savings?

in place over years, and which they continue to refine. To those vendors and their customers, “enterprise OS” means things like solid and simple installation, thorough integration and configuration testing, and processes for the distribution of patches and upgrades that balance the need to respond to critical issues with the need to keep the platform stable. By contrast, open-source software has historically been most accurately viewed as a bunch of components, each of which gets rapid-fire updates at arbitrary intervals; if any component doesn’t work right the developer just fixes it and rolls out another update (or users can even fix it themselves).

The next logical step for Linux is toward a distribution with stability and reliability as its priorities. Other Linux distributors now also have distros with an enterprise focus, notably SuSE Linux Enterprise Server 8 and Mandrake Linux Corporate Server 2.1. But Red Hat led the way, first with Advanced Server, then with a lineup of three Enterprise Linux products and a clear differentiation between its “enterprise-class” products and its more traditional Linux releases.

### Pieces of the Puzzle

Red Hat’s Enterprise Linux 2.1 lineup includes Red Hat Enterprise Linux WS for workstations, Enterprise Linux ES for smaller servers, and Enterprise AS<sup>2</sup> typically for medium-to-large servers. In many respects, these versions are more alike than different. All are open source, as are Red Hat’s “traditional” editions—Red Hat Personal and Professional—which the company now pitches as consumer products. All the “enterprise” versions are built from a common code base and ship on a common release schedule—one that is less rapidfire and more deliberate than is the case with Red Hat’s traditional products, and *much* more measured than Linux’ as a whole.

The current multi-product lineup has its roots in Advanced Server, which Red Hat rolled out as its first enterprise distribution in May, 2002. However, the company soon realized that a point product wasn’t enough, especially when that product was touted most

for clustering and other scalability and reliability features that positioned it more as a database server than for more general purposes, like Web serving. Thus, like other commercial OSs, Red Hat has evolved to offering multiple products in order to better meet a variety of customer requirements in performance, support, and pricing. In some respects, the structure of its current x86 product line is most directly comparable to Windows Server, which—following the 2003 release—comes in four variants. However, unlike Microsoft, the Red Hat products differentiate primarily on the level of support and the size of system supported, rather than on licensing policies. For example, only AS comes with high-availability clustering, a 24x7 service option, and support for more than two processors.<sup>3</sup> By contrast, Unix vendors typically offer a single product rather than a line of products—in large part because their OS usually ships as part of a complete bundle with hardware, thereby offering them the opportunity to set different implicit pricing in different parts of their product lines.

Today’s Red Hat Enterprise Linux 2.1 lineup includes versions for 32-bit x86 and 64-bit IPF (Itanium Processor Family)<sup>4</sup>—essentially rebranded editions of, respectively, Red Hat Advanced Server 2.1 which first shipped in May, 2002 and the versions of Red Hat Advanced Workstation and Advanced Server for Itanium that shipped in November, 2002. The various IBM architectures for which Red Hat also offers distributions aren’t yet part of its formal Enterprise Linux line. Red Hat plans to change that with Enterprise Linux 3, however, which is scheduled to ship in the second half of 2003. At that point, all Red Hat’s enterprise-focused products will be folded under the Enterprise Linux 3 brand—including x86 and IPF, as well as

3. The two CPU limit for Enterprise Server WS and ES is essentially a service level agreement; there are no restrictions in the code to keep these distros from running on more processors, but Red Hat won’t support them on larger configurations. This reflects both a pricing strategy and an observation that larger hardware configurations inevitably exercise software more—thereby uncovering more bugs and generating more support calls.
4. Currently, Red Hat sells its IPF versions only through Hewlett-Packard, the system vendor selling the majority (80 percent by some reports) of IPF hardware. There is also no ES version for IPF, only WS and AS.

2. Replacing Advanced Server.

IBM's pSeries, iSeries, zSeries, and S/390<sup>5</sup> and AMD's new AMD64 architecture. Red Hat plans to offer versions of all three products (WS, ES, and AS) for most of its supported architectures, with the exception of combinations that wouldn't make sense for most people—for example a workstation version designed to run on a mainframe server.

The things that set *all* these enterprise-oriented variants apart from Red Hat's more consumer-oriented line fall into three major categories: the product itself, the way it's supported by Red Hat, and its more deliberately-paced release cycle.

### What's Different about the Actual Product?

Enterprise Linux's kernel code spun off from the stock Linux version 2.4.x kernel at about the time Red Hat 7.2 shipped in 2001. Since that time, the code's received significant enhancements to improve its scalability—especially Oracle database scalability.

Many of these enhancements relate to I/O—the movement of data to and from disk—which is so important to transaction processing and other commercial computing tasks. Because the latencies associated with moving data in and out of a system are so long compared to the time it takes to actually operate on that data, efficiency and economy of movement are essential to overall system performance. If processors end up sitting idle waiting for data to come in from disk, the performance penalty is huge. Crunching numbers that are already in the processor's closest data store might take a modern CPU a billionth of a second or less. But retrieving those numbers from disk will take about 10 milliseconds—10 *million* times longer.<sup>6</sup>

Asynchronous I/O is one Red Hat addition that attacks that "latency penalty." A simple application, which will typically have only a few code streams running at any given time, must necessarily leave CPUs idle because they have no other tasks for the

processor to perform while it waits for data. But large and complex applications can much more readily fill in the "dead time" by shuffling from one task to another, rather than waiting for a single task's data. The wait staff in restaurants does the same thing, taking the order of a second party while waiting for the first party's meal to be prepared. If there's only customer in the restaurant, the technique doesn't save much time. But if there's a crowd, total throughput is vastly higher if the order taker can move from task to task during the delay in preparation. Similarly, Asynchronous I/O provides an Application Programming Interface (API) that lets applications issue an I/O request, then immediately continue processing, rather than waiting for a result to come back.

Red Hat has also improved the granularity of the locks it uses to coordinate access to critical I/O code in multiprocessor systems.<sup>7</sup> The more processors there are in a system, the more discriminating they have to be in deciding how much of a common resource pool each can grab. As Linux scales up, enhancements such as this one will inevitably be an ongoing effort. Big Iron vendors have spent years working at tasks like these because coarse-grained locks tend to keep additional processors from doing additional useful work.

Outside the kernel, the most significant addition is the high availability clustering—derived from Mission Critical Linux's "Kimberlite" technology—that's part of the AS product. This Red Hat Cluster Manager software coordinates "active-active" access to common shared SCSI or Fibre Channel storage: when both systems are functional, both can do useful work. By contrast, in "active-passive" arrangements a backup system sits idle waiting for the primary to fail—a much less efficient use of server hardware. Cluster Manager is a significant advance in an open source Linux world, which largely defines clustering as load balancing and failover among small servers that don't share storage—something suitable for Web serving farms, but not for critical database engines.

5. Red Hat has versions for both IBM's older 31-bit S/390 mainframes (which Red Hat refers to as its S/390 version) and well as newer 64-bit mainframes, which IBM now calls the zSeries.

6. See Illuminata note "Latency Matters!" (September 2002).

7. Red Hat refers to this enhancement as I/O Spinlock Contention Reduction.

But these products are not just about adding new features and code.

They're also about removing packages that aren't up to what Red Hat considers a necessary level of maturity and reliability—for example, the initial release of the logical volume manager (LVM) in the Linux version 2.4 kernel on which Enterprise Linux is based. Instead, Red Hat chose to steer those requiring the ability to manage their storage through a logical rather than physical view toward partner Veritas, which offers a closed source volume manager that is widely used and trusted by data-center operators.

While it might initially seem odd that a supposedly more advanced—and more expensive—product has *fewer* features, it's really not. IT managers in charge of large datacenters have little desire to mess with—and, even less so, place into production—new whiz-bang features if they're not ready for prime time. Thus the new role of Red Hat's standard Linux distros, as a sort of technology delivery vehicle that lets customers kick the tires of new packages and updates that may yet need a bit more testing and development before they're ready for serious data-center deployment.

### But It's Process, Too

When discussing an OS, it's inevitable that discussions turn to the product attributes—its bells and whistles, speeds and feeds. However, those are not ultimately what most distinguish Red Hat's Enterprise Linux line from the company's traditional distributions. After all, the enterprise versions are still open source, and most of their best features will still end up in the standard kernel over time. Indeed, in theory at least, anyone could assemble his or her very own version by downloading the components from Red Hat's web site.<sup>8</sup> But that would be missing the point of the Enterprise Linux line.

8. Although doing so is not particularly easy or straightforward; Red Hat doesn't offer free CD image (aka ISO) downloads of its Enterprise Linux distros as it does for its Personal and Professional Editions—requiring a much more labor intensive download and build process.

Fundamentally, these products are sold on a per-server, per-year subscription model of computing. Yes, there's a CD (or a download) with an OS included in the price, but it's the ongoing support that's the real value, not the initial (open source) bits. Standard and Premium subscriptions are all-inclusive: they provide for unlimited incident support, upgrades are included, and it doesn't matter how many clients are attached to a server.<sup>9</sup> The service level varies by product: two tiers within each. The mainstream "Standard Edition" offers 12x5 support and 4-hour response; AS can also be bought in a pricier 24x7/1 hour version for servers operating under more stringent "no downtime" dictates.<sup>10</sup>

The "subscription" component of the products comes through an enhanced "Enterprise" version of the existing Red Hat Network (RHN) that delivers the updates and notifications needed to keep code current and secure. This sort of proactive systems maintenance has been a traditionally weak area for OS vendors of all stripes. Consider that even Sun—as sophisticated an enterprise Unix vendor as any—felt the need to release a brand new Patch Manager with its new Solaris 9 in May, 2002.<sup>11</sup> This explains why system administrators often have to fill the gap with homegrown scripts and other incomplete tools.

Red Hat Network is a service that notifies users of, and provides access to, updates over the Internet. It also handles package dependencies; i.e. determining whether a given update means that other pieces *also*

9. By contrast, Microsoft tiers its server pricing in part through the sales of client access licenses (CALs) and tiers its service pricing in part through limiting the number of incidents it will cover in the base price.
10. WS and ES can also be purchased in a "Basic Edition" which provides 90 days of Installation and Configuration support and a year subscription to Red Hat Enterprise Network (for access to updates/errata). Without ongoing support, however, Basic Edition doesn't really fall under the same subscription model as the other editions.
11. Although the terminology differs—in large part reflecting the different approaches to updating a largely closed source vs. an open source OS—Solaris's Patch Manager serves much the same "push out updates and fixes" function as Red Hat Network. See Illuminata note: "Solaris 9: Fortifying the Foundation" (July 2002).

need to be updated for the OS to continue running reliably and securely. Red Hat Enterprise Network's enhancements over the basic RHN are there primarily to deal with managing large numbers of systems by grouping together servers with similar functions and updating all of them at once—a big increase in ease and consistency. Red Hat Enterprise Network's big future thrust is value-add application and system monitoring—technology that has its roots in Red Hat's NOCpulse acquisition in 2002.

### **New Releases, Frugal ISVs, and the Harried IT Manager**

Red Hat Enterprise Linux is on a 12 to 18 month release schedule, rather than the more rapid four to six month schedule of its traditional distros.<sup>12</sup> This shift in part reflects increasing OS maturity; over a given period, more enhancements are now incremental improvements rather than critical fix-ups. It also reflects a loud “slow down!” cry from both IT managers and ISVs, who are simply and increasingly unwilling to expend the resources to deal with rapid churn.

In shifting to a year-plus release cycle, Red Hat is mirroring the actions of just about every major OS vendor from Sun to Microsoft—vendors who are in turn responding to what they're hearing from their customer bases. Indeed, the cycles are continuing to lengthen, with two to three years between releases not unheard of.<sup>13</sup> As Linux continues to mature, expect it to follow this continuously lengthening trend as well.

In Red Hat's case, ISVs are perhaps the biggest reason behind the shift to longer cycles. In recent years, ISVs have dramatically cut back on the number of platforms that they support—a cut-back that helped to accelerate the demise or planned

obsolescence of a host of second-tier Unixes. The issue isn't so much the cost of the port; major ISVs generally do a fairly efficient job of keeping most of their code portable across multiple architectures, the better to keep development costs down. Rather, it's the testing and support that chews up serious dollars for not only each additional platform, but for each operating system version on each platform. As a result, ISVs have gotten very particular about not only which OSs on which processor architecture they'll support, but about even how many OS *versions* they'll certify.

The real difficulty comes because datacenters don't run just a single application. They run many—and these many applications often need to be certified to work together by their respective vendors. But these are rarely blanket guarantees; they cover only specific software versions running on specific OS versions. Overly-frequent OS releases can end up causing individual ISVs to be selective about the OS versions that they match with particular application versions—resulting in a patchwork that can make it very difficult for an IT manager to assemble a complete certified software suite. By slowing things down to a release cycle that's more acceptable to ISVs, it will be easier for Red Hat to ensure that all of its releases,<sup>14</sup> and not just some of them, get included in ISV planning.

### **Duking It Out with the Big Boys**

Linux will continue to evolve; certainly it has lessons from its elders still to apply in areas like workload management. And it's unclear if Linux as a whole will ever bring the same level of focus to single system commercial application scalability as do the major Unixes and, increasingly, Windows. But in many respects it's already surprisingly sophisticated, both as a base operating system and in the specific form of Red Hat's product line.

For example, Red Hat Enterprise Network provides notifications and product updates with a degree of sophistication that was lacking even in major Unixes until fairly recently; it is more consistent

12. Although Enterprise Linux 3 will follow the March, 2003 Enterprise Linux 2.1 release by only six months or so, Red Hat reasonably considers the initial 2.1 version to be a rebranding of Advanced Server packages from 2001 and 2002, rather than a new release.

13. The exception is IBM, which has slightly shortened the time between AIX releases. But it's still a relatively leisurely (by traditional Linux standards) 18 months or so.

14. At least for popular processor architectures like x86.

and integrated than what Microsoft can do today.<sup>15</sup> Performance is creeping up as well. Although Linux' focus remains more on distributed computing than ultimate single system performance,<sup>16</sup> systems running Enterprise Linux (or its Advanced Server predecessor) have delivered very respectable results on TPC-C (commercial online transaction processing), Oracle Applications Standard Benchmark, and ECperf (a measurement of J2EE performance) relative to Unix and Windows systems.

Add to those capabilities the benefits brought by the base Linux OS and its associated software ecosystem. Indeed, in a sense, the fundamental purpose of enterprise Linux distributions is to harness the enormous potential benefits of open source in a way that's useful for the corporate datacenter.

One such benefit is certainly cost. Total Cost of Ownership studies give evidence that just about any vendor's products or services are cheaper over a several year period when all the "relevant" factors are taken into account. However, not even Microsoft disputes that Linux is almost always the acquisition-cost champ—even when, as in the case of a commercially sold distribution, it's not *quite* free. In combination with high-volume x86 hardware, it's hard to undercut Linux solutions on price. Even Sun, the most price-aggressive of the Unix vendors at the low end, explicitly compares itself with Wintel, never with Lintel—and, indeed, began a reluctant "if you can't beat them, join them" approach to Linux in 2002.

15. Microsoft's Chief Security Strategist, Scott Charney, frankly admitted at Microsoft's Industry Analyst Summit in March, 2003 that "Our patch management story is not pretty"—a problem he described as largely a business process failure stemming from Microsoft's decentralized nature.

16. See Illuminata note "Looking the Right Way at Linux" (November 2002).

But open source isn't *just* about cost. To be sure, in corporate IT, the ability to roll your own kernel fixes or to substitute the file system of your choice won't appeal to most. And those it does appeal to aren't likely to be the best customers for the certifications, code stability, and orderly processes of enterprise Linux. However, open source still brings benefits even to those who never write a line of code. It brings them flexibility. It's much easier to switch Linux distros than it is to move from one Unix to another, much less from Windows to anything else. And that flexibility in turn helps give customers more power; if one Linux distributor won't help, another doubtless will.

### Helping Linux Grow Up

Red Hat's business-focused lineup formalizes the evolution of Linux into a more buttoned-down enterprise role than its wild-haired, open-source siblings. Its enterprization will be an ongoing process, however, that will inevitably be the source of some tension with the "Linux should be free" crowd. Even though Enterprise Linux remains fully and completely open source, Red Hat's quite clear that its product line is about support and application certification as much as it's about the product itself. And support and certification cost money that self-service development does not.

In a way, it's not really surprising that Red Hat has evolved its enterprise-oriented distros to look more and more like the commercial Unixes that they can often quite capably replace. After all, enterprise IT requirements don't change because the OS is Unix, Linux, or Windows. Sure, big back-end systems often have more stringent needs than smaller distributed systems, but as Linux grows up it increasingly has to follow the same rules as everyone else. It will get no byes because of the cute little penguin or open-source wonderfulness. And that's where enterprise Linuxes come in.



Through subscription research, advisory services, speaking engagements, strategic planning, product selection assistance, and custom research, Illuminata helps enterprises and service providers establish successful infrastructure in five key areas:  
Server Technologies, Information Logistics, Application Strategies,  
Enterprise Management, and Pervasive Automation.