

Linux, Inc.

A Survey on Open Source Software

October 2002

The Dravis Group

San Francisco, CA

info@dravis.net

415-665-8745

Table of Contents

Executive Summary	3
Introduction	4
Open Source 101: What is it about?	4
Open Source: More than the Linux Operating System	7
Commercial Use is Growing	9
Systems Hardware Providers Address the Market	10
The Linux Desktop: A Challenging Opportunity	11
Enterprise Software Firms Recognize the Trend	12
Embedded Linux Could Drive Consolidation	13
Grid Computing: The High-End of the Market	14
Governments Enter the Debate	15
About The Dravis Group	16

Tables

Table 1: A Sampling of Linux Distributions	6
Table 2: Commercial Deployments of Linux	9
Table 3: Enterprise Software Support	12
Table 4: Examples of Embedded Linux Solutions	13
Table 5: Government Actions	15

Charts

Chart 1: Linux can Provide Scalability	5
Chart 2: Apache: An Open Source Market Leader	7

Appendices

Appendix 1: Terms to Know	17
Appendix 2: US National Security Agency Comments	18
Appendix 3: GNU General Public License	19

This document is copyrighted © by The Dravis Group LLC and is protected by U.S. and international copyright laws and conventions. This document was developed on the basis of materials and sources believed to be reliable. This document is to be used “as is”. Opinions reflect judgment at the time and are subject to change.

Executive Summary

We believe both Linux and open source software are changing the agendas of technology providers and users alike. Interest is global and, in some countries, the discussion of open source versus proprietary software has entered into the realm of political debate. Fueling confusion is a technology industry coming to terms with the implications of open source software, along with competitive rhetoric. Some suggest there is significant risk in pursuing open source strategies, while others believe it increases the democratization of computing.

The convergence of multiple factors are driving Linux and open source software adoption. They include 1) Linux's similarity to the UNIX operating system, along with leverage gained by its availability on the Intel platform, 2) an increased focus on lowering technology expenses, 3) access to a global community of software developers, and 4) desires to reduce Microsoft's ability to leverage its desktop software monopoly into other areas of technology.

Regardless of the rationale, major software and hardware firms are supporting the use of Linux in products as diverse as home entertainment centers, digital video recorders, PDA's and communications equipment. **Google** runs its Internet Search business on it with over 10,000 systems. **Amazon.com** is migrating its infrastructure to it. Wall Street firms are incorporating it into trading systems. **Walt Disney**, **DreamWorks**, and **Industrial Light and Magic** are pursuing film production with it. **L.L. Bean** supports its e-mail on it. Even the **United States White House's** web site runs it and a security enhanced version is available from the **National Security Agency's** web site.

Our report highlights initiatives by leading hardware and software providers, while addressing the breadth of commercial use. Much work remains to be done as open source software is explored and exploited. This survey intends to broaden the discussion of this important topic.

Paul J. Dravis
October 2002
San Francisco, CA.
paul@dravis.net

Introduction

Interest in the Linux operating system and open source software are increasing globally. Technology firms are introducing product offerings to address these demand trends and technology users are assessing its impact as well. At the same time, the strength of these dynamics can be difficult to gauge. Adding to the confusion are concerns of an unproven business proposition, the belief that the software is unreliable, along with technology terms, acronyms and products which may seem unfamiliar to decision makers.

The emphasis on “openness” in open source software has fostered the growth of a world-wide community of developers contributing to the evolution and improvement of a variety of software programs. While such a diffused structure for software development may seem chaotic, this approach is being considered as a more democratic alternative to monolithic single vendor initiatives. Our report provides background on these efforts and intends to bring the discussion to a wider audience.

Open Source 101: What is it about?

Open source software is differentiated from proprietary software because the programming code used is available for inspection, modification, re-use and distribution by others. It is often assumed that open source software is free of charge. While this can be the case, open source software can be purchased for a fee as well. The concept of “free”, in this context, emphasizes what can be done with the source code rather than its cost.

While open source is associated with Linux and the GNU licensing agreement, there are many types of open source software programs available and the Open Source Initiative¹ has certified over 30 different open source licensing agreements.

GNU: Where it Started

The work which lead to the open source movement is derived from efforts at MIT in the early 1980s when Richard Stallman started the GNU project to prove that an operating system could be developed and shared freely. In 1996 the project added a kernel providing a complete GNU system. All programming code from this project adheres to the GNU General Public License (see Appendix 3).

Linux: Where Much of the Focus is

In 1991 Linus Torvalds, then a student at the University of Helsinki, started work on Linux as an open source operating system. In 1994 version 1.0 was released, based on a kernel he developed and code from the GNU project. As a result, Linux falls under the GNU General Public License. The current Linux version, 2.4, was released in January 2001.

The Linux operating system is used in many ways including support for networking, software development, servers and desktop platforms and is considered as a low cost alternative to other operating systems.

¹ A non-profit organization dedicated to manage and promote open source licensing certification.

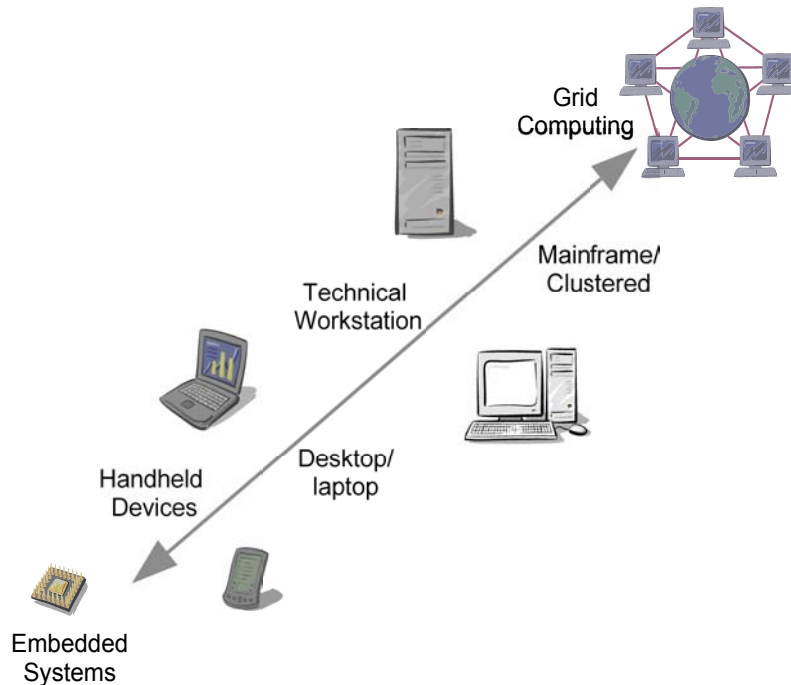
Why a Penguin?

The Linux Penguin logo was selected by Linus Torvalds to represent the image associated with the operating system he developed. When version 2.0 was announced Linus said: "Some people have told me they don't think a fat penguin really embodies the grace of Linux, which just tells me they have never seen a angry penguin charging at them in excess of 100 mph. They'd be a lot more careful about what they say if they had."

Software Scalability is an Advantage

Linux was first available for the Intel x86 platform, providing a viable low cost UNIX development environment and an alternative to Microsoft's Windows. Today, Linux support includes the Compaq Alpha AXP, Sun SPARC and UltraSPARC, Motorola 68000, PowerPC, PowerPC64, ARM, Hitachi SuperH, IBM S/390, MIPS, HP PA-RISC, Intel IA-64, DEC VAX, AMD x86-64 and CRIS architectures. This broad support provides the potential for Linux development efforts to scale from low-end embedded systems to high-end mainframe and distributed environments as well.

Chart 1: Linux can Provide Scalability



What is a Linux Distribution, and Who are the Players ?

The Linux community consists of dozens of different software “distributions”² which can be downloaded for free or purchased from commercial organizations. A Linux distribution is a collection of software programs which include the operating system, along with tools to install the operating system and set-up the user interface. Word processors, e-mail and web server programs can be included in a distribution as well.

With Linux development involving multiple teams in the open source community, the distributor works to assure that the collected components work properly together. Distributors can provide different levels of product support resulting in a variety of pricing schemes in commercial offerings.

Table 1 highlights a few of the dozens of distributions which are available.

Table 1: A Sampling of Linux Distributions

Company/ Distribution	Description
<i>Red Hat, Inc.</i>	- Supports Linux from the mainframe to the server to the embedded device. Focused on the enterprise market and supports third party software applications such as the Oracle Database. Partnering with software firms such as BMC , Borland , CheckPoint , Computer Associates , IBM (WebSphere, DB2, Lotus, Tivoli), Oracle (Oracle 9i Database, Oracle 9i Application Server, Oracle E-Business Suite), Synopsys , Tibco , and VERITAS .
<i>UnitedLinux</i>	- A joint effort of The SCO Group (US), Conectiva S.A. (Brazil), SuSE Linux AG (Germany), and Turbolinux, Inc. (Japan) to streamline Linux development and provide a uniform global Linux distribution. UnitedLinux will be available in English, Japanese, Simplified and Traditional Chinese, Korean, Portuguese, Spanish, Italian, German, French, and Hungarian.
<i>Lycoris</i>	- Targeting the Linux desktop user with development efforts focused on improvements in the graphic user interface.
<i>Lindows.com</i>	- In addition to addressing desktop usability, the firm is promoting its “Click-N-Run” technology, allowing users to easily add software programs to their desktop system.
<i>Debian</i>	- Started in August 1993 as a new distribution which would be made openly and available, Debian is typically used by a technical audience.
<i>MandrakeSoft</i>	- Founded in 1998, this distribution provides solutions for desktops and servers (web server, e-mail server, print server, etc.). Mandrake supports over 40 international languages.
<i>Redflag</i>	- Beijing based Red Flag Software focuses on the Chinese Linux user. Their solutions include Red Flag Linux Server, Red Flag Linux Desktop, Red Flag Linux operating system for IBM S/390/Alpha/Intel/PA RISC/Sun Sparc host, Red Flag Linux server and Embedded Linux OS for Set Top Boxes, PDAs, and thin clients.

² Refer to DistroWatch, www.distrowatch.com, for many of the Linux distributions available.

Open Source: More than the Linux Operating System


Open source initiatives extend to areas such as web servers, browsers, media servers, e-mail systems and work group collaboration. Some efforts are research-oriented, while other applications have moved into commercial franchises. The interest in Linux could increase the acceptance for other open source programs as well.

Open source programs are not exclusive to the Linux platform. Some open source applications are available for multiple operating systems. *Apple Computer* provides open source support for its Rendezvous networking technology and its Darwin 6.0.1 operating system under the Apple Public Source License.

The Apache Web Server: Where Open Source is the Market Leader

In the market for web server software, used to manage web sites pages, Apache is the market leader. According to Netcraft³, of 15.3 million active public web sites, 67% use this open source program, with Microsoft's Internet Information Server (IIS) in second place at 25%. Apache is available as part of all major Linux distributions. To address the support needs of the Apache market, firms such as *Covalent Technologies, Inc.*, founded by Randy Terbush⁴, deliver products to address security, reliability and manageability requirements. The recently released Apache 2.0 runs on Linux, Solaris, UNIX, and Windows 2000.

Chart 2: Apache – An Open Source Market Leader



Source: Netcraft

³ Offers services in the areas of World Wide Web Publishing, Internet Security, and Systems & Network Management.

⁴ One of the developers on the open source Apache project.

E-Mail: A Critical Application

E-mail, messaging and collaboration continue to be central to Internet growth. An early entrant into the e-mail market was the open source program Sendmail, with its origins in the early 1980s. Today its developer, Eric Allman, is a member of the management team at *Sendmail, Inc.* which builds e-mail systems for large enterprises and service providers. The firm's customers include *Pfizer, Lowes, Farmers Insurance, Harvard University, L.L. Bean, Korean Airlines, Sprint* and *UPS*.

Targeting Linux and Unix desktop users, *Ximian, Inc.'s* Evolution provides e-mail, calendar, contact and task list management while integrating to Microsoft Exchange, Lotus Notes, and other messaging systems and standards. The firm's Mono Project focuses on building an open source version of Microsoft's .NET development platform for Linux and UNIX.

Positioned as an alternative to Microsoft Exchange, *Bynari Inc.'s* product line targets the Linux market with messaging services running on Intel and IBM S/390 platforms. The firm has announced support for an open source model with its own open licensing terms.

Databases: A Large Opportunity with Low Penetration

MySQL is an open source database created by Europeans David Axmark, Allan Larsson, and Monty Widenius. They developed the offering in 1996 and started *MySQL AB*, which provides fee-based technical support for it. MySQL is used by *Yahoo, Cisco Systems, NASA, Lucent, Google, Hewlett-Packard, Xerox* and *Sony Pictures* among others, and has an installed base estimated to be over 4 million users.

While both *Oracle* and *IBM* have embraced Linux as a way to fend off *Microsoft*, these firms are likely to view open source databases as less feature-rich than their own. There are however many uses for open source databases such as assisting in the management of web pages. MySQL may not have as many features as offerings from Oracle, IBM or Microsoft, but it now provides transaction processing support and we assume the larger players are keeping a watchful eye on the dynamics of this market.

Red Hat provides an open source database offering with object oriented technology from *PostgreSQL*.

Commercial Use is Growing

Most initial Linux implementations support departmental file/print, mail or Web servers. With firms such as *Dell Computer*, *Hewlett-Packard*, *IBM*, and *Sun Microsystems* increasing their involvement in Linux, its large scale enterprise deployment should expand. Already implementations are extending beyond its roots in the education and government sectors, and being adopted in many industries as an alternative to both UNIX and *Microsoft's* Windows platforms.

Table 2: Commercial Deployments of Linux

Company	Description
<i>Amazon.com</i>	- Migrating most of its technology infrastructure to Linux platform to help reduce costs.
<i>Bohnacker GmbH</i>	- Implement integrated sales planning, using software from <i>SAP</i> running on <i>IBM's</i> xSeries.
<i>Credit Suisse First Boston</i>	- Migrated a RISC/Unix based application to a <i>Red Hat</i> Linux/Intel platform to support its Agora' enterprise notification system which resulted in both server consolidation and improved performance.
<i>Google</i>	- Runs over 10,000 Linux based Intel PCs in support of their search engine service.
<i>Industrial Light and Magic</i>	- Converting its workstations and renderfarm to Linux in support of the production of Star Wars, Episode II: Attack of the Clones.
<i>L.L. Bean</i>	- Uses the Sendmail e-mail system with Linux on an IBM zSeries mainframe.
<i>Merrill Lynch</i>	- Using Linux supporting Web Services, messaging and application server needs. Reviewing the benefits of Linux applications scalability between the Intel and IBM Z-series platforms.
<i>Regal Entertainment Group</i>	- The movie theater operator to use Linux based kiosks to sell popcorn, soft drinks and other snacks at their movie theaters.
<i>Reuters</i>	- Working with <i>HP</i> , <i>Intel</i> and <i>Red Hat</i> to make Reuters Market Data System available on Intel Xeon and Itanium servers.
<i>Royal Dutch Shell</i>	- The oil exploration unit is working with IBM to build a Linux based supercomputer linking 1,024 servers to analyze seismic data and other geophysical information as part of its efforts to find new supplies of oil.
<i>Verisign, Inc.</i>	- Implemented Linux servers when upgrading their Internet Domain Name Service used for Web site lookups. This architecture, called ATLAS, uses low cost Linux/Intel boxes on the front-end and higher-end Linux/IBM systems on the back-end.
<i>WesternGeco</i>	- Supporting their business of seismic imaging requiring intense computations with a cluster of 256 IBM eServer xSeries systems running Linux as an alternative to a high-priced supercomputing system.

Systems Hardware Providers Address the Market

In the past, *IBM*, *Hewlett-Packard* and *Sun Microsystems* provided their own variants of the UNIX operating system to support their workstation and server product offerings. Each firm has committed to the Linux market, but to varied degrees. *Dell Computer* has entered this market as well.

IBM: Linux is Strategic

IBM leads in its Linux commitment identifying it as a strategic focus for the firm. Linux is supported on the firm's Intel xSeries servers, its mid-range iSeries servers and its mainframe zSeries, with plans to put it on its pSeries servers, which typically run IBM's AIX. Additionally, IBM is funding research such as its 250 person Linux Technology Center and has initiatives to provide Linux support for their *DB2* database, *WebSphere* e-commerce software, *Tivoli* systems management and *Lotus* e-mail and calendar software product lines

Hewlett-Packard: Increasing its Commitment

Hewlett-Packard is going where its customer needs are. Recognizing the Linux trend, the firm plans to train 5,000 staff in their services organization on the platform. Significant Linux deployments for the firm include *Walt Disney Feature Animation* where HP's Linux-based workstations and servers are components in next-generation digital animation production, *Amazon.com* where HP helped in the transition to Linux servers, and *DreamWorks* by supplying technical workstations, servers, printers, networking and Linux technologies. *DreamWorks'* Shrek was developed on Linux workstations.

Sun Microsystems: Adaptive, Proactive, Reactive

Sun Microsystems' success is built upon its workstations, servers and the Solaris operating system. Its open source commitment is evident with the focus on the StarOffice software. This open source office suite is available for Linux, Solaris and Windows. Sun recently introduced the LX-50 Linux based server targeted at the low-end of the market and plan to release a low-end desktop system in 2003 bundled with Linux, the *Mozilla* Web browser, StarOffice, and *Ximian, Inc.'s* Evolution open source e-mail program. The desktop offering will be targeted at specific business opportunities such as call centers. It remains difficult to gauge the strength of Sun's commitment to Linux relative to their Solaris-oriented product line.

Dell Computer: Moves with the Direction of the Market

Unlike the other systems providers mentioned, *Dell Computer* has no UNIX legacy. Its franchise was built around the Wintel platform. Like others, Dell is addressing customer needs and provides *Red Hat* Linux as an option for its server product line. At least 13% of Dell servers are running Linux. Dell's strategy has been to focus on the Unix-to-Linux migration market and supporting the *Oracle* 9I RAC database.

A recent Dell/Linux deployment at the State University of New York (SUNY), Buffalo launched a supercomputer consisting of over 2,000 Dell PowerEdge servers running *Red Hat Linux*. The system will be used to conduct drug research to combat cancer, Alzheimer's disease and AIDS.

The Linux Desktop: A Challenging Opportunity

Much of Linux's success has been on the server side of the computing landscape because of its similarity to UNIX, which dominates the server market. However, desktop dynamics are different from those of the server. The diversity of application software and hardware components are greater, and user expectations are quite varied. Issues for Linux include the challenges of establishing a desktop support infrastructure when Windows support is already in place and compatibility with Microsoft desktop applications. Some desktop users may become confused with software available from multiple distributors, new nomenclature and unfamiliar product offerings.

Minimal support by Systems Providers is an Issue

Mainstream access to Linux desktop systems is limited. While major systems providers are increasing their Linux support for servers, their limited commitment on the desktop seems surprising. Desktop users must seek out specialty providers or purchase a Windows system, remove the software and then install Linux. There are some initiatives however. *Sun Microsystems* will enter the market in 2003, but with targeted initiatives, such as addressing call centers with low-end systems. Retail distributor *Wal-Mart* sells systems on their website with distributions from *Lindows.com* and *MandrakeSoft SA*.

Desktop Software: Is Microsoft Compatibility Needed?

Access to the Internet has changed the way desktop systems are used. This has increased the value of networked access, but for many users, an Office suite remains a critical component of desktop system use. While mainstream desktop solutions for Linux are limited, users with a more technical background can find numerous offerings to meet their needs.

Sun Microsystems provides StarOffice, a Microsoft Office equivalent product for Linux, Windows and Solaris. Many of the commercial Linux distributions incorporate it into their offering. *Codeweavers, Inc.'s* CrossOver Office 1.0 provides *Microsoft* Office compatibility with technology from the WINE project⁵. However, there can be performance and compatibility trade-offs with this approach.

In 1999 *Corel Corp.* was an early software entrant into the market with its WordPerfect office suite, but has since withdrawn support. A re-entry would position the firm as the sole provider of an office suite available on the Windows, *Apple* Macintosh and Linux platforms.

Accessing Shared Content is Improving

Allowing Linux users to share content with and having access to media services available to Windows and *Apple's* Macintosh users is important. While Internet and World Wide Web standards help in establishing some of this foundation, more is needed. Commercial software providers which support the Linux desktop include *Real Networks* with its RealOne Player and *Adobe Systems* with an Acrobat PDF Reader. *IBM* provides support for *Lotus* Notes/Domino and *America Online* supports its Instant Messenger on Linux. Additionally, numerous web browsers are available, with some supporting multiple platforms such as *Mozilla* and *Opera Software*⁶.

⁵ Approach uses Windows APIs and native DLL support.

⁶ The firm's browser supports Windows, Linux and Macintosh; technology available in the embedded market with IBM, AMD, Symbian, Ericsson, Sharp, Lineo, and others.

Enterprise Software Firms Recognize the Trend

Much of Linux's success has been in Internet applications, technical workstations and clustered processing environments. Among the next challenges is support for more traditional enterprise application environments. While information technology spending remains constrained in today's economic environment, enterprise software firms have initiated efforts to address the Linux market.

The following table overviews various initiative by major technology software providers.

Table 3: Enterprise Software Support

Company	Description
<i>BEA Systems, Inc.</i>	- Offers a hardware independent Application Server Platform and Java Virtual Machine optimized for a variety of Linux versions, including Red Hat and SuSE , and has alliances with the key Linux platform providers including HP, IBM, Intel, and Sun Microsystems.
<i>BMC Software, Inc.</i>	- Provides varied infrastructure and service management needs for Linux, while providing a single console addressing these application components on Intel and IBM eServer zSeries platforms.
<i>Borland Software Corp.</i>	- Released Kylix, a Linux version of their Delphi application development tool.
<i>Computer Associates, Inc.</i>	- Rolling out over 50 products addressing the Linux market including updates to its UniCenter line, storage management enhancements to BrightStor, and security management solutions for eTrust. Through its subsidiary, ACCPAC International, Inc. , it provides a mid-market accounting application for Linux servers and desktops.
<i>IBM Corp.</i>	- Pursuing a variety of initiative to bring Linux support to its broad software product line. In addition, funded an industry group, called Eclipse, to provide open source software development environments. Other firms involved include Borland , Merant , QNX Software , Rational Software , Red Hat , SuSE , and TogetherSoft .
<i>Oracle Corp.</i>	- One of the first commercial software providers to focus on Linux with support for the Oracle9i Database, Oracle9i Application Server, Oracle9i Developer Suite and Oracle E-Business Suite. Worked with Red Hat in the development of their Advanced Server. Customers using Oracle on Linux include the Federal Aviation Administration with a 5-node RAC database and CERN (European Center for Nuclear Research) which tested 9-node RAC databases and is plans for additional deployments.
<i>SAP AG</i>	- Established the LinuxLab to focus on the development and release of mySAP, their core solution platform, on Linux and address related support problems. Currently Linux support included distributions from Red Hat and SuSE, IBM DB2, Informix, Oracle and SAP DB databases and hardware from Dell Computer, Fujitsu-Siemens, HP, IBM and Bull. SAP Linux solutions have been deployed at Bohnacker GmbH and DSD Dillinger Stahlbau GmbH
<i>Veritas Software, Inc.</i>	- In storage management, is addressing the need for data backup/ recovery and high availability. Development partners include Red Hat and Intel and have distribution agreements with Dell Computer , IBM and Hewlett-Packard .

Embedded Linux Could Drive Consolidation

Similar to changes occurring in the server market, Linux may reduce fragmentation in the embedded systems industry. Increased Linux use could standardize low-level technology development and allow OEMs to increase their focus on higher-level product differentiation. These dynamics could result in Linux becoming the embedded platform of choice in many industries.

The embedded systems market delivers solutions which integrate computer hardware and software resulting in products as diverse as communication routers, factory automation equipment, medical equipment, test and measurement devices, entertainment systems, set-top boxes, and video recorders, cameras, phones, and PDAs.

Table 4: Examples of Embedded Linux Solutions

Company	Description
<i>Hewlett-Packard Co.</i>	- HP Digital Entertainment Center has Linux 2.4 and the X Window system, a 566 MHz Celeron, with 64MB RAM and a 40GB hard disk, USB, Ethernet, HPNA, a built-in V.90 modem, stereo audio, and TV video out.
<i>IBM/Citizen Watch</i>	- A prototype watch running Linux 2.4.18, X11 graphics, and a Bluetooth wireless protocol stack. Includes a Cirrus ARM EP7312 processor, 16 MB of Flash, 8MB RAM, a 320x240 monochrome reflective LCD display, Bluetooth wireless communication hardware, a fingerprint sensor, a speaker, a microphone, and infrared module.
<i>Intel Corp.</i>	- A new PC peripheral called the digital media adapter, which provides a link between PCs, TVs, and stereos. Based on an XScale microarchitecture PCA210 'applications processor' and runs an embedded Linux operating system.
<i>Linksys Group Inc.</i>	- Wireless Presentation Gateway: Wireless mobile PC users can project presentations and other data onto VGA multimedia projectors, monitors, and LCD devices.
<i>Sharp Corp.</i>	- Zaurus SL-5500: A PDA with Lineo's Embedix embedded Linux, a 206 MHz Intel StrongARM processor with 64MB system RAM and 16MB built-in flash storage, a color 320 x 240 pixel TFT LCD, plus built-in keyboard.
<i>SONICblue, Inc.</i>	- SONICblue Rio Central: a home stereo component with a 40GB hard drive, a 206MHz Intel StrongARM processor along with 16MB of system RAM running an Embedded Linux OS derived from Debian/ARM
<i>Sony Corp.</i>	- Plans to release a digital video recorder as part of new audio-visual products taking advantage of broadband networks. The system runs Montavista Linux version 2.4.17, on a MIPS processor running at 350 MHz.
<i>TiVo Inc.</i>	- Linux used in their Digital Video Recorder product line.
<i>Toshiba Corp.</i>	- Toshiba Wireless Mobility Server - Magnia SG20 is intended to make it easy to pack up an entire network and take it on the road. Device contains a Red Hat Linux along with Apache (which provides web-based administration and application support).

Grid Computing: The High-end of the Market

The Internet and web browsers illustrated the power of networked resources to access information. Another major wave of technology innovation may come from Grid Computing. Technically, this type of computing pulls together the processing power of servers and workstations into a single resource. This should allow large tasks to be distributed across many computers delivering supercomputer speed from “off-the-shelf” hardware.

Lawrence Livermore National Laboratory selected *Linux NetworX* to deliver what could be the largest and most powerful Linux supercomputer in the fall of 2002. Multiple programs at the labs plan to use the clustered supercomputer to support the Laboratory's national security mission. When delivered, the cluster is expected to be one of the five fastest supercomputers in the world clustering 1,920 *Intel* Xeon processors at 2.4 GHz with a theoretical peak of 9.2 teraFLOPS.

The National Science Foundation has a \$53 million project called TeraGrid with the goal of building the world's largest distributed infrastructure for open scientific research. When the TeraGrid is completed, it will include more than 13 teraFLOPS of Linux-based computing clusters distributed over the four sites. Partners in this effort include *SDSC at UC San Diego*; the *National Center for Supercomputing Applications* at the University of Illinois; the *Center for Advanced Computing Research at Caltech*; and *Argonne National Laboratory* in Argonne, IL, with corporate sponsorship from *IBM*, *Intel*, and *Qwest Communications*, *Myricom*, *Sun Microsystems*, and *Oracle Corp.* Efforts to build software development tools to manage and leverage this environment are driven by the Globus Project.

Platform Computing Inc., a developer of software to manage distributed computing environments, is working with the Globus Project to provide a commercially supported version of the open Globus Toolkit.

Grid computing should play a key role in the evolution of on-line multi-user gaming. *Butterfly.net, Inc.* is providing computing grids and tools for Massively Multiplayer Games to connect players on PCs, consoles, and mobile devices. Linux systems support their universal gaming infrastructure where servers can be added, or replaced, without interrupting game-play. The firm is working with *IBM*, *Sony* and *Microsoft* in these efforts.

Governments Enter the Debate

Governments around the world are increasing their involvement in the open source software debate bringing varied actions as they assess its potential for minimizing their reliance on propriety software vendors such as Microsoft. *The MITRE Corp.* is concluding a study of Linux use in the US Federal Government and the *European Commission* released a report "Pooling Open-Source Software" to address cost savings by sharing software on an open source licensing basis.

International initiatives by technology providers include *IBM* establishing Linux design/support centers and *Sun Microsystems* providing free versions of StarOffice for education use in Europe, Africa and Asia. In response, *Microsoft* has provided financial contributions to governmental programs, such as in Mexico, and is renegotiating licensing terms in others.

Keeping abreast of the changes in the international debate can be a daunting task. Table 5 attempts to overview some of the current activity.

Table 5: Government Actions

Country	Action/Status
<i>China</i>	- Beijing government awarded Chinese software vendors contracts. Beijing-based Red Flag Software Co. Ltd. will provide government computers with its version of Linux. The Hong Kong government has installed more than 100 Linux servers in various departments in the past three years. Public pressure to avoid dependence on single-vendor products has prompted government interest in open source.
<i>Finland</i>	- Government testing of open source software. 13 government agencies completed a project in April 2002 to test Open Office open-source desktop productivity suite and Star Office, from Sun Microsystems Inc.
<i>France</i>	- The French Ministry of Culture and Communications replaced software on some of the government's servers, running Windows NT and IBM's AIX, with Red Hat Linux.
<i>Germany</i>	- The German Federal Minister of the Interior has a deal with IBM to promote Linux use in the public sector. IBM will use a distribution from <i>SuSE Linux AG</i> .
<i>Korea</i>	- Korea's <i>HancomLinux Inc.</i> working with Korea's Central Procurement Office to provide 120,000 copies of Linux desktop office software.
<i>Peru</i>	- A bill under debate in Congress would require government agencies to use open source software. Proprietary or commercial applications, such as those from Microsoft or IBM Corp.'s Lotus Development Corp., could only be used when no open source alternative is available ⁷ .
<i>Thailand</i>	- Government-subsidized National Electronics and Computer Technology Centre has developed its own Linux package for use on government desktop computers and servers based on <i>Red Hat's</i> distribution.

⁷ Correspondence between Microsoft's General Manager and the Congressman sponsoring the bill have become required reading in the debate. See <http://www.pimientolinux.com/peru2ms/index.html>

About The Dravis Group

The Dravis Group was founded to assist technology providers, users and investors in assessing and navigating through the changing technology landscape. Our research intends to bring clarity to market dynamics leading toward focused and practical decision-making.

Paul J. Dravis has participated in the technology sector for over 15 years from a number of perspectives. His background includes investment management (Dresdner RCM Global Investors), equity research (Banc of America Securities, Robertson Stephens and Co. and JP Morgan) and advanced technology strategy and application development (JP Morgan). He has published on topics as diverse as The Storage Management Market, Systems Management, The Year 2000 Challenge, Client-Server Technology, Data Networking and the World Wide Web. In addition to being a Wall Street Journal All-Star Analyst, his media exposure includes ABC (Nightline), BBC, CNN, CNET, CNBC, National Public Radio, New York Times, Business Week and Computerworld.

Appendix 1: Terms to Know

- Apache** - An open source Web server available on most UNIX systems (such as Linux, Solaris, Digital UNIX, and AIX), and Windows NT/2000.
- Application Server** - A server program in a computer in a distributed network that provides the business logic for an application program. The application server is frequently viewed as part of a three-tier application, consisting of a graphical user interface (GUI) server, an application (business logic) server, and a database and transaction server.
- Free Software Foundation(FSF)** - Founded in 1983 along with its demonstration GNU project by Richard Stallman at MIT to prove that an operating system could be developed and shared freely. "Free" does not mean at no charge, but refers to the use the person who acquires the software has with it. FSF believes that individuals the right to study and make changes to program's source code that improve the program , and redistribute and sell improved versions of the software, as long as it is "open" to others.
- GNOME (GNU Network Object Model Environment)** - A graphical user interface and set of applications including word processor, spreadsheet program, database manager, presentation graphics, Web browser, and e-mail. GNOME comes from work of the Free Software Foundation.
- GNU** - A UNIX-like operating system built with source code that can be copied, modified, and redistributed. The GNU project was started in 1983 by Richard Stallman and the FSF. Linux consists of GNU components and a kernel developed by Linus Torvalds.
- Java** - A programming language introduced by Sun Microsystems in 1995, designed for use in the distributed environments. IBM, Microsoft, and others offer Java compilers.
- Kernel** - The center of an operating system providing basic services for to its other parts.
- KDE** - K Desktop Environment is an open source graphical desktop environment. KDE includes a file manager, a window manager, a help system, a configuration system, tools and utilities, and several applications. The KDE project was started in October 1996 by Matthias Ettrich.
- Linux** - An operating system designed to provide Intel PC users with a low-cost alternative to UNIX systems. Linux's kernel was developed by Linus Torvalds at the University of Helsinki. To complete Linux, Torvalds used of components from the Free Software Foundation's GNU project. Linux is available on major hardware platforms.
- Mozilla** - Mozilla was Netscape Communication's nickname for Navigator Web browser, and more recently, the name of an open source Web browser project.
- .NET** - Microsoft's strategy and programming efforts to address Web services. Their goal is to provide seamless interaction between applications and computers.
- Open Source** - Any program whose source code is made available for use or modification as users or other developers see fit. Historically, proprietary software developers have not made source code available.
- Operating System** - The program loaded into the computer by a boot program (BIOS) that manages the other programs. Linux, Windows 2000, VMS, OS/400 and AIX are examples of operating systems.
- Source Code** - A computer program before it is compiled and run on a computer. Purchased operating system or application software, is usually in compiled object code while the source is not included.
- Web Server** - A program that manages the files that form Web pages which are presented Web users. The leading Web servers are Apache and Microsoft's Internet Information Server (IIS).

Appendix 2: US National Security Agency Comments

While the National Security Agency's web site www.nsa.gov does not run on Linux, the Agency does provide a security-enhanced Linux system, which can be downloaded from the site. Accompanying the download, the agency provides the following information:

“As part of its Information Assurance mission, the National Security Agency (NSA) has long been involved with the computer security research community in investigating a wide range of computer security topics including operating system security. Recognizing the critical role of operating system security mechanisms in supporting security at higher levels, researchers from the NSA's Information Assurance Research Group have been investigating an architecture that can provide the necessary security functionality in a manner that can meet the security needs of a wide range of computing environments.

End systems must be able to enforce the separation of information based on confidentiality and integrity requirements to provide system security. Operating system security mechanisms are the foundation for ensuring such separation. Unfortunately, existing mainstream operating systems lack the critical security feature required for enforcing separation: mandatory access control. As a consequence, application security mechanisms are vulnerable to tampering and bypass, and malicious or flawed applications can easily cause failures in system security.

The results of several previous research projects in this area have been incorporated in a security-enhanced Linux system. This version of Linux has a strong, flexible mandatory access control architecture incorporated into the major subsystems of the kernel. The system provides a mechanism to enforce the separation of information based on confidentiality and integrity requirements. This allows threats of tampering and bypassing of application security mechanisms to be addressed and enables the confinement of damage that can be caused by malicious or flawed applications.

Linux was chosen as the platform for this work because its growing success and open development environment provided an opportunity to demonstrate that this functionality can be successful in a mainstream operating system and, at the same time, contribute to the security of a widely used system. Additionally, the integration of these security research results into Linux may encourage additional operating system security research that may lead to additional improvement in system security.”

Appendix 3: GNU General Public License

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.
675 Mass Ave, Cambridge, MA 02139, USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

PREAMBLE

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of

having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it. Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you

received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

