



ANALYSIS OF LINUX FOR ENTERPRISE APPLICATIONS AND DATABASES

TABLE OF CONTENTS

- 2 Executive summary
- 3 Outstanding reliability, availability and serviceability of Linux
- 5 Market trends reflect Linux dominance
- 7 TCO, performance, and ease of migration from Solaris favor Linux
- 8 Linux: The strategic operating system for Oracle databases
- 10 Linux security is unparalleled
- 11 Red Hat: The open source software company
- 11 Reference
- 12 Appendix 1

EXECUTIVE SUMMARY



Linux[®] has become the operating system of choice for a broad community of software and hardware vendors. All over the world, thousands of developers and end users of standards-based systems choose Linux because of the open collaboration that occurs. The governance model for Linux is—unlike Solaris[®] 10 x86 or OpenSolaris[®]—extremely transparent and open, making vendor independence a reality. There are well-defined organizational and legal support structures in place to protect the long-term viability of the open source operating system. By making all vendors neutral, the industry is the beneficiary of low price, better performance, and service delivery as the basis for competitive differentiation. Linux distribution vendors like Red Hat maintain a demanding quality assurance (QA) process to ensure that the open source packages integrate and perform properly.

With Open Solaris and Solaris 10 x86, the governance model is controlled by Sun Microsystems. Neither the OpenSolaris community nor Sun participate in the Open Invention Network. The Open Invention Network is an intellectual property company that manages and shares critical patents that affect Linux development and innovation. Red Hat, Novell, and Oracle participate in this effort, which safeguards open source intellectual property.

Driven by a large developer community, the technical merits of Linux include many contemporary and emerging technologies which have been (or will be) accommodated into the upstream kernel. Operating systems like Solaris or OpenSolaris lack the community needed to maintain such an aggressive pace of innovation. Enhancements like Real Time extensions to contain latency spikes complete with IBM's Real Time Java Virtual Machine, a common architecture across all mainframe and distributed platforms, or the Oracle Data Integrity Initiative add to the performance, reliability, and security of the Linux platform.

During the course of the past six years, Linux-based platform deployments have grown dramatically at the expense of proprietary UNIX operating systems like Solaris. The acceptance of Linux by mainstream hardware providers like IBM, HP, Dell, AMD, and Intel and its embrace by a very vibrant software ecosystem has contributed to this change. The ever-increasing stability and maturity of the Linux kernel has pushed the operating system deeper into the enterprise, moving beyond web services and infrastructure. The Linux platform now supports mission-critical applications and databases across various industries.

OUTSTANDING RELIABILITY, AVAILABILITY, AND SERVICEABILITY OF LINUX



Linux maintains its relevance for low-level tasks while remaining highly useful elsewhere in the data center. In a study called “Get the Truth on Linux Management,” released in February 2006 by the Enterprise Management Associates[1], survey participants revealed the following:

Reliability - Linux environments are extremely stable, in many cases more so than many proprietary environments. Many enterprises experience 100% uptime for their Linux environments, and most experience greater than 99.99% availability. When problems occur, Linux administrators repair them quickly, in most cases faster than industry averages.

Provisioning - Deployment and provisioning of Linux systems and applications can happen very quickly. Over half of surveyed enterprises can provision a Linux system in less than 1 hour. This process is even more efficient with sophisticated management tools.

Management - Most respondents spend less than 5 minutes per server per week managing Linux systems - including patch management, system migration and repurposing, security management, and virus/spyware protection. Most respondents manage their Linux systems centrally, many use sophisticated management tools, and across the board, these tools made Linux system management easier and faster.

Storage Management - Sophisticated storage management technologies such as SAN and NAS have made storage management and utilization a platform-independent discipline. In these mature environments, most Linux administrators spend very little time managing storage, yet have high utilization rates.

Support - Customization and support of multiple versions does not have any discernible impact on the management effort for Linux systems. In many cases, administrators use the same management tools across many versions of multiple heterogeneous systems including Linux, UNIX®, Microsoft® Windows®, and z/OS.

System Availability - Older research claimed Linux has lower reliability and availability, specifically citing longer repair times caused by less skilled resources and difficulties with problem diagnosis, documentation, and repair. However, our research showed that Linux environments were extremely stable, with high availability and short problem resolution times.

The report goes on to state that “some organizations had exceptional uptime, with 17% of respondents reporting 100% availability. Over half of all respondents (58%) reported availability of over 99.99%, and 89% reported availability over 99.95%. At the other end of the scale, only 2.5% of respondents reported availability worse than 98%.”

The Enterprise Management Associates report also cites that “a large international bank with over 1,500 Linux servers reported the only downtime in 2 years was from a CPU failure, and was not even software-related, adding that their Linux environment ‘just runs

and runs and runs.' A large hosting site reported 100% availability across 230 managed servers since October 2001. A large city university reported 100% availability for systems with sophisticated management tools (across both Intel and z/Series), with unplanned downtime only affecting a handful of unmanaged Linux systems."



The report added: "When problems do happen in Linux systems, they are quickly resolved. The average time to fix a server reporting a fault (i.e. Mean Time To Repair, or MTTR) for over 60% of respondents was less than 30 minutes. This is well below the average found in a separate EMA study, which showed most enterprises averaged MTTR of around 4 hours."

The report also provided examples about organizations who had achieved extremely high MTTR with alternative storage and system architectures under Linux. For example, "a large city university using a sophisticated provisioning solution keeps all their data on NAS, so they can simply switch out a failing system and diagnose it later, resulting in MTTR of around 5 minutes." A large financial institution running Linux on IBM z/Series hardware reported effective MTTR of 0 minutes, because they have "never run into any bugs from Linux," noting this as "one of the reasons we can support it with so few people - that and the reliability of mainframe hardware."

Several other interesting findings were also identified in the analysis, including:

- Problem resolution - In over 60% of cases, when problems occur in Linux environments they are diagnosed and repaired in less than 30 minutes, over 8 times faster than industry average.
- Management and support - 88% of enterprises with Linux and Windows spend less effort managing Linux; 97% believe it is, at worst, the same for both systems. Respondents with sophisticated management tools all report Linux management is the same or easier than Windows management.
- Storage management - Enterprises with sophisticated management tools did not find any significant difference in storage management effort or utilization for Linux.
- Resource costs - Most administrators, for either Linux or Windows, earn under \$60k. Salaries for combined Linux/Windows administrators are only marginally higher than for Linux-only administrators. Linux skills are readily available.
- Consulting and training costs - 79% of enterprises spent nothing on Linux consulting, and 63% spent nothing on training. Only 4% spent over \$10K on consulting or training.
- Security Management - 75% of Linux administrators spend less than 10 minutes per server per week managing security. With sophisticated management tools, this goes up to over 85%.

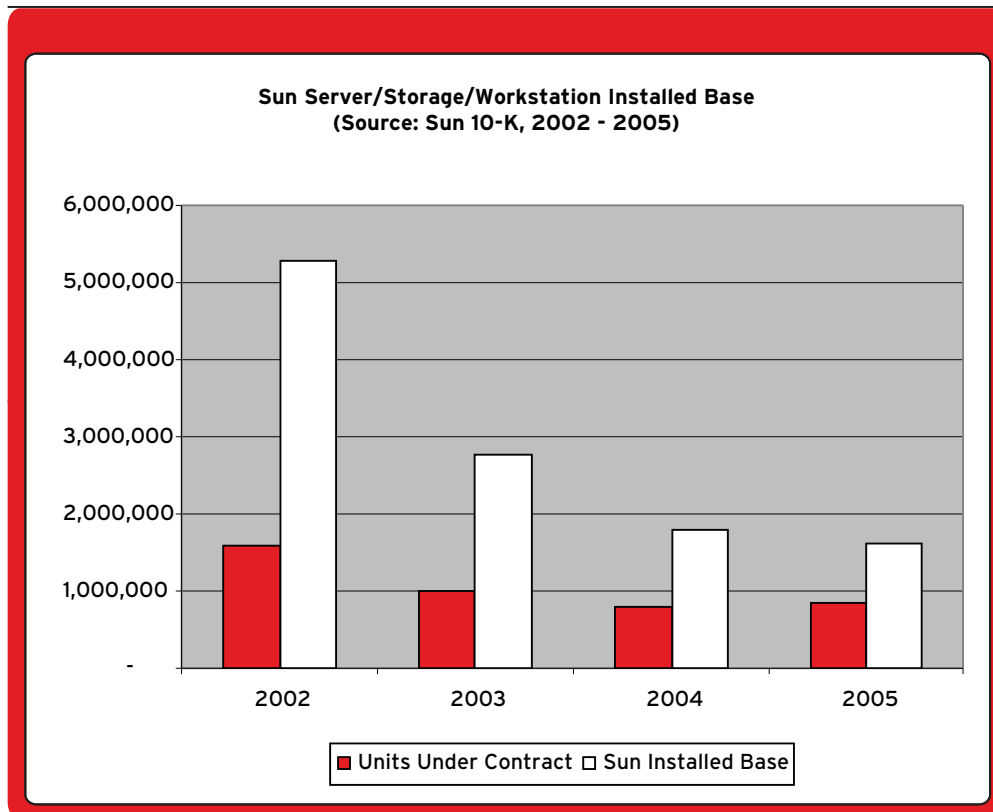
New tools have recently been introduced that enhance the serviceability of Linux. The first is called Frysk. Frysk provides enterprise customers with an open source, system-wide debugging and monitoring solution. The set of tools is advancing rapidly since its introduction in March 2007 with the launch of Red Hat Enterprise Linux 5. Please see www.sourceware.org/frysk for more details.

SystemTap provides a means to simplify the gathering of information about a functioning Linux system. This assists in the diagnosis of a performance or functional problem. It also eliminates the need for the developer to go through the challenging instrument, recompile, install, and reboot sequence that may be otherwise required to collect data. SystemTap also offers a simple command line interface and scripting language for writing instrumentation for a live running kernel. Sample templates have been published, and there is an ongoing effort to broaden the internal "tapset" script library to aid in reuse and abstraction. New features are being developed and incorporated, like user space probing which will be available shortly. See Appendix 1 for a comparison of the features and functionality of SystemTap and Sun Microsystems's Dtrace.



MARKET TRENDS REFLECT LINUX DOMINANCE

The market share shift has greatly reduced Sun's installed base. This base has underwritten Sun's research and development effort for years, keeping it current and influencing technology trends. According to the company's 10-K filings from 2002 to 2005, Sun server, storage, and workstation installed base has declined by over 70%. Furthermore, their quarterly shipments are roughly half the levels of 1999 to 2000 which indicates that the installed base is not being replenished. Independent software vendors can no longer anticipate the Solaris volume of previous years. As a result, the software community's interest in the Solaris platform is diminishing. Oracle is a prime example. They have very publicly chosen Linux.





Faced with shrinking revenues and lost relevance, Sun modified its platform strategy. Though Sun still clings to its proprietary SPARC-based server architecture (well-known as the industry's weakest performer), Sun embraced the x86 platform, first with AMD® Opteron processors and more recently with Intel processors.

Though Sun Microsystems has blurred the reality of the market's acceptance of Solaris 10 x86 and OpenSolaris, it is clear that Linux has a tremendous lead in market share—a lead that is growing rapidly. Recent IDC statistics show that Sun continues to ship many more x86 systems with Linux than they do with any variant of Solaris:

CPU TYPE	OPERATING SYSTEM	Q4 2006 (REVENUE)
X86, 64	LINUX	\$15,101,000
	UNIX	\$5,520,000

After two years general availability, OpenSolaris and Solaris 10 x86 still lack the critical mass needed to be relevant to the industry. In contrast, the IDC's Linux server unit shipment numbers for 2005 show how massive the Linux installed base has become. The table below summarizes IDC's data.

CPU TYPE	OPERATING SYSTEM	2005 (UNITS)
X86, 64	LINUX	266,737,360

Although hardware providers like HP or IBM may quietly support Solaris 10 x86 on a limited number of platforms, the compatibility of various third party I/O and peripherals like fibre channel adapters, network interface cards, and RAID controllers is suspect. For example, Infiniband adapters for Solaris 10 x86 are not certified—however, several adapters are listed on Sun's website as 'verified.' This means that a product has undergone testing in a Sun environment but that support is not guaranteed by Sun Microsystems. With Linux, there is widespread support for all hardware platforms and peripherals.

Linux is clearly a Tier One platform for software vendors. OpenSolaris or Solaris 10 x86 is not. Certain core components may be available on Solaris 10 x86 from leading application providers. However, critical pieces of software needed to complete a standard build for any enterprise may or may not exist with the appropriate version. Given the relatively small market for Solaris 10 x86 or OpenSolaris, the availability of these software elements will be an ongoing problem for Sun's customers.

Proper certification of both hardware and software may not seem compelling when operations are running smoothly and without incident. However, when a problem occurs, those unsupported components can be costly—or impossible—to fix. Many products are identified on Sun's web site as 'verified' but not 'certified' on Sun hardware, making the product not technically supported on Sun systems by the vendor. Enterprises that deploy non-supported software or hardware elements in a production environment face a significant amount of risk.

The overall supportability of OpenSolaris is an interesting question. Recent analysis shows that there are only fifteen kernel engineers supporting OpenSolaris. By contrast, in

1999 there were already over 1000 contributors to the Linux kernel. The lack of community involvement may relate to the fact that OpenSolaris uses a CDDL License instead of the GPL, a favored license in the open source community. The CDDL License allows Sun Microsystems to control the pace of change and technology direction.




In a recent white paper, Robert Frances Group concludes that:

“...the CDDL licensing of Solaris 10 makes it legally open source; however, it does not provide the basis for Solaris users to benefit from the broad open source community. It particularly will not benefit from contributions from a larger hardware vendor community (the source of billions of dollars of Linux investment), which is discouraged from seriously investing in a product that will remain controlled by a single direct competitor.”

Furthermore, if an enterprise adopted either OpenSolaris or Solaris 10 x86 as an additional operating system, the engineering effort needed to manage the build and the ongoing support process would be considerable. That staff would need to ensure hardware and peripheral suppliers had certified their devices with those operating systems and that Independent Software Vendors had done the same with their software. In addition, there would be a need for testing and certifying the new platform. Conceivably the release schedule and patches for various platforms would not be uniform, making the internal certification process even more cumbersome.

TCO, PERFORMANCE, AND EASE OF MIGRATION FROM SOLARIS FAVOR LINUX

Sun Microsystems has touted the higher performance of OpenSolaris over Linux, the ease of migration from SPARC based Solaris, and the comparable TCO of OpenSolaris to Linux





as key reasons for customers to choose Sun's solutions. However, there is sufficient evidence[2] to show that Solaris of any type is not higher performing than Linux and that the effort required to move to Linux from Solaris SPARC is less than the other way around. Additionally, the total cost of ownership for Solaris is almost double that of Linux. The following link demonstrates that Linux is the performance leader with Oracle, not Solaris of any type.

http://tpc.org/tpch/results/tpch_perf_results.asp?resulttype=cluster

The migration of Solaris-based applications to a Linux platform is very straightforward since both operating systems are based on UNIX standards. Although Sun would like its customers to think that the migration from Solaris SPARC to OpenSolaris or Solaris 10 x86 is accomplished with ease, in reality that migration is similar to or more complex than moving to Linux. Common tools, commands, and overall environment between Solaris and Linux help facilitate the migration. As depicted above, the lower TCO of a Linux solution makes the migration attractive to an enterprise, which will benefit from an immediate return on investment.

According to migration documents like Hewlett Packard's "Migrating from Solaris to Linux," there are certain challenges associated with a migration from Solaris, but they're not difficult to overcome. For example, the guide explains that "like Solaris, Linux supports both 32-bit and 64-bit architectures. Linux supports 32-bit applications on 64-bit operating systems through a number of means, including emulators and native binaries. A significant difference when migrating applications from Solaris to Linux is endianness. Endianness refers to the way in which data is stored, and defines how bytes are addressed in multibyte data types. This is important because if you try to read binary data on a machine that is of a different endianness than the machine that wrote the data, your results will be different. OpenSolaris and Solaris 10 x86 share the same approach with regard to endianness."

LINUX: THE STRATEGIC OPERATING SYSTEM FOR ORACLE DATABASES

Through a series of public endorsements, Oracle has demonstrated to its customer base that Linux is very strategic to its longterm efforts. One recent announcement indicated their desire to ensure data consistency throughout a transaction life cycle. Unlike OpenSolaris or even Solaris 10, Oracle will be providing a multi-vendor data-checking specification for Linux alone. This new enterprise data integrity checking scheme is designed for mission-critical databases. The Data Integrity Initiative, led by Oracle and partners Emulex, LSI, and Seagate, plans to implement a standard data-checking mechanism. This will allow data integrity to be verified all the way from the application through the storage network to the disk drive. The development effort for this project will be incorporated into the upstream Linux kernel. This is recent evidence that Oracle is beginning to put some distance between its relationship with the oldline UNIX providers and its huge Linux installed base.

The vast majority of Oracle database customers have extremely stable implementations on Linux. Occasionally, there have been situations where the engineering team or the



operations team have created and deployed a Linux build which is not properly designed for the database function. For example, a Linux build for web services may have been reused for the Oracle database. When this occurs, stability can sometimes become a problem. There are several best practices which should be incorporated into the database platform:

If 64-bit Linux is not an acceptable option, use the hugemem kernel on 32-bit systems with more than 6GB of RAM.

Use huge pages. Huge pages can be configured in Linux by specifying `vm.hugetlb_pool=Size_in_MB` in `/etc/sysctl.conf`. They will be used by Oracle 9.2.0.6 and Oracle 10 automatically. Earlier versions of Oracle require patches (see Metalink Note: 262004.1. Patch 3318884 for 9.2.0.4 and patch 3566570 for 9.2.0.5).

The sum of all Oracle SGA segments should be equal to the size of huge pages pool. Do not create too large a pool of huge pages because this memory can't be used for other purposes.

Use direct I/O if your Oracle data files are on ext3 file system. The 2.4 kernel does not support direct I/O and async I/O on ext3 simultaneously. If you need async I/O for performance reasons put Oracle data files on raw devices or better, use the 2.6 kernel as the platform. The 2.6 Linux kernel supports direct I and async I/O on ext3. The implementation of async I/O in the 2.6 kernel is much more efficient than in the 2.4 kernel.

Using raw devices for Oracle data files and async I/O is a recommended configuration with the 2.4 kernel. If you prefer to put Oracle data files on a file system, you can consider using 3rd party file systems supporting async I/O and avoid data caching on operating system level. Also:

- OCFS supports only direct I/O operations and requires a special set of utilities to deal with files on this file system.
- VxFS can be mounted in direct I/O mode.

Pay attention to the NTP configuration for Oracle servers. You should have an odd number of time sources and the NTP drift file should be writable for NTP daemon.

Many Linux customers have begun to incorporate the native dynamic multi-pathing and cluster suite offering into their Oracle database environments to reduce cost, simplify the storage engineering effort, and add stability to the platform.

The Linux community has a commanding dominance over the Oracle database market. Early last year, Gartner Group announced that Oracle had achieved over \$1B in Linux-based database revenue. Major database consumers such as Yahoo! have standardized on Linux for their databases.

LINUX SECURITY IS UNPARALLELED



Designed with scale-up and scale-out architectures in mind, Linux is the platform that allows you to implement the secure, redundant, scalable, reliable, and ultimately available processes to meet your business' service level agreements (SLA). Since scaling up architectures for transactions processing limits your ability to adapt, a modular architecture not only removes processing bottlenecks from transaction execution, but simultaneously provides the flexibility to grow and adapt the architecture with business demands.

Increasing the volume of digital transactions also requires increasing levels of operational security--and Linux delivers. The Linux platform provides the first open implementation of Multiple Category Security (MCS), a derivative of Multiple Level Security (MLS). This is accomplished through Security Enhanced Linux (SELinux), which is included with every subscription at no additional charge. With these technologies, you can harden your environment in ways previously only available to the military and intelligence agencies, yet enable increased and protected sharing of information with easy-to-use graphical or command-line tools.

With increasing regulatory pressures on the industry, a clear understanding of the activities of your users and their use of platforms is imperative to ensuring the integrity of business operations. Linux extends traditional system access models with the use of strongly authenticated, fully integrated smartcards to regulate system access. Advanced auditing capabilities permit detail tracking of every aspect of a transactional system, allowing businesses to both control and review information on system use. These natively increasing abilities are necessary to adhere to regulatory requirements such as Sarbanes Oxley and Basel II.

RED HAT: THE OPEN SOURCE SOFTWARE COMPANY



At Red Hat, choice is about so much more than just Linux.

Over the past five years, UNIX-to-Linux migration has saved enterprises millions of dollars, delivering a choice in hardware and superior performance--but Red Hat is more than Linux. We are delivering the next major open source migration: Moving legacy applications to an Open Source Architecture (OSA).

We are accelerating the shift to service-oriented architectures and enabling the next generation of web-enabled applications. Our latest subscriptions simplify the use of JBoss across the application life cycle, from development to production.

Our customers have asked us to provide options. These include replacing expensive and restrictive proprietary software like Veritas, Powerpath, Oracle, BEA, or WebSphere. Today with our integrated virtualization and Global File System technologies, Red Hat Identity Management Suite, AMQP-based Red Hat Messaging, and JBoss, we can deliver on an ever-expanding open source promise to replace your proprietary stack with lower cost options.

"It was no accident that when we brought out Red Hat® Enterprise Linux in 2003, we went to Wall Street first. Now, almost every Wall Street investment bank runs on Red Hat Enterprise Linux. And from those early customers, it spread to 70 percent of the Fortune 500, and for the last five quarters, we've been adding about 10,000 new customers every quarter." Charlie Peters, Red Hat CFO, NYSE Magazine, Second Quarter, 2007)

We know Linux. We know open source. We know the performance, reliability, control, and cost savings it delivers. And better than anyone else, we use this technology to build complete, cost-effective infrastructures for the enterprise. That's why over 900 US-based CIO's and other senior-level IT executives have ranked Red Hat number one in value three years straight[3].

[1] Enterprise Management Associates. "Get the Truth on Linux Management." February 2006.

[2] Robert Frances Group. "TCO for Application Servers: Comparing Linux with Windows and Solaris." August 2005.

<http://www-1.ibm.com/linux/whitepaper/robertFrancesGroupLinuxTCOAnalysis05.pdf>

[3] CIO Insight magazine's Vendor Value Survey: <http://www.redhat.com/promo/vendor/>

REFERENCE

http://www.redhat.com/f/pdf/rhel/dynamic_it.pdf

www.redhat.com/f/pdf/rhel/financial_services.pdf

www.principledtechnologies.com/clients/reports/Red%20Hat/Red%20Hat.htm

APPENDIX 1



	systemtap	dtrace
PROJECT		
license	GPL	CDDL
operating system support	Linux	BSD, Solaris
processor support	wide (as per kprobes: x86, x86_64, ppc64, ia64, s390, sparc?, ...)	narrow (x86_64, sparc)
kernel coupling (“upstream” lock-step)	little	lot
core developers	open community	Sun internal + external porting
development began	2005	2001
development status	ongoing	stable
target audience	developers, users, sysadmins	developers, users, sysadmins
target usage	debugging, tracing, profiling	tracing, profiling
LANGUAGE		
style	scripting	scripting
full control structures (conditionals, loops, functions)	yes	no
variable typing	implicit, inferred	implicit
complex report generation	yes	no (post-process with e.g. perl)
scalable aggregates	yes	yes
aggregate value readable by script	yes	no
thread-local variables	yes (from first principles via tid-indexed auxiliary arrays)	yes
speculative tracing	made redundant by available control structures	yes
binary tracing	yes	?
early boot tracing	not yet	yes
probe execution	optimized native code	interpreted bytecodes
PROBING CAPABILITY		
number of possible probe points	millions (statements, markers)	tens of thousands (functions, markers)
probe arbitrary points in code (function entry, exit, interior)	yes (using debugging information)	no



	systemtap	dtrace
probe dynamically loaded kernel objects	yes	yes
concurrent probes on multiprocessors	yes	yes
extract arbitrary data at probe point	yes (any context-visible variable)	no (fixed \$arg1 parameters, designated globals?)
context pointer type punning	not yet	yes
end-user extendable probe library	yes (script based tapsets)	no (compiled-in providers)
probe user-space programs	soon	yes
probing JVM	soon (special case of user-space)	yes
statically inserted probe points, kernel side	yes (STP_MARK)	yes (DTRACE_PROBE)
statically inserted probe points, user side, e.g. interpretive languages	not yet	yes (macros)
timer-based probing	yes (profiling interrupt, software timers)	yes
hardware performance counter based probing	soon	no
SAFETY		
time-limited probe handler execution	yes (statement counting)	yes (implied by language restriction)
non-blocking, atomic probe handlers	yes	yes
space-limited execution	yes (static allocation of all data)	yes
division-by-zero protection	yes	yes
null pointer dereferencing protection	yes	yes
means available to bypass protection for advanced users	yes (guru mode, embedded-C)	limited (predefined signal actions available)
safe use on production systems	soon	yes
translate-time error checks	many	many
run-time error checks performed by	automatically generated C code	bytecode interpreter
use by unprivileged users	for now, sudo all-or-nothing	graduated access by privilege level