



**Enabling**

**Government and Business Transformation**

**With**

**Mandatory Access Control Security**

*December 2006*



---

## Executive Summary

---

### The “Conundrum”



The convergence of faster technology innovation and a global digital economy has created a fundamental change in the marketplace. Consumers are aware of the expanding service and product offerings which have become available to them and have raised their services expectations accordingly. Customer allegiances can no longer be taken for granted.

This shift is influencing the way businesses and governments interact with their consumers and constituents. Retaining mindshare today largely depends on an organization’s ability to consistently deliver innovative, high value services faster than the competition. Organizations are looking to take advantage of more efficient and effective web channels to deliver these higher value services. This kind of integrated service delivery is more feasible today with the wholesale adoption of open standards computing. There is a conundrum however which is becoming apparent to those responsible for administering security.

**“Web-based services are complex systems that, of necessity, have to be open for access by a large group of – possibly uncontrolled – users. Providing assurance of security mechanisms is hard because the user application relies on a large set of middleware components – web servers, databases and so on – that may reside on multiple machines. Some of these middleware components may have privileged access to system resources providing a potentially large prize to any malicious user who is able to breach the application security.”**

**Dr. Steve Marsh  
Intelligence and Security Advisor  
UK Cabinet Office**

This White Paper examines the risks associated with standard infrastructure security when protecting today’s increasingly sophisticated application infrastructures. The paper then introduces a security model which can be tailored to meet exact, in depth security requirements (wherein a breach in one area doesn’t imply breaches elsewhere). This capability is available in Linux (the first commercial operating system to do so) under the name SELinux (or Security Enhanced Linux). Also covered are findings from a partnership effort undertaken by IBM, the UK Cabinet Office and Tresys Technology wherein an effective trial of this technology was conducted using IBM’s WebSphere Application Server. This paper addresses further how SELinux uniquely stems and contains the escalating threat posed by cyber crime resulting in faster, more effective service delivery opportunities for businesses and Governments.

---

## 2. Cyber crime impacts

---

Many Governments are focused on transforming their service delivery to better support the needs of the public they serve. Today, constituents are demanding access to higher value services through the use of new web channels. As Public Sector organizations seek to deliver higher value solutions which incorporate a broad value chain contributing to a service offering, the infrastructure complexity and security exposures grow rapidly. This is creating a major challenge for Governments who are also charged with safeguarding citizen privacy and national security.

Illustrating this point are recent breaches being reported across many Government organizations. "Parts of the UK's Critical National Infrastructure (CNI) are being targeted by an ongoing series of email-borne electronic attacks. While the majority of the observed attacks have been against central Government, other UK organizations, companies and individuals are also at risk. The attackers' aim appears to be a covert gathering or transmitting of commercially or economically valuable information." The UK's National Infrastructure Security Co-ordination Centre (NISCC) warned in June 2005.

Security breaches have been acknowledged since May of 2006 in the US Navy, USDA, US Energy Dept., DoD and US Department of Health and Human Services. Social security numbers, security clearance levels, place of employment data and financial information on tens of thousands of Government personnel and private citizens have been stolen. "Recent Government Security Breaches" - Associated Press Article (June 22, 2006).

Military systems have also been known to experience serious breaches raising questions about national security. One such example resulted when a group called "Titan Rain" from China was credited with wreaking havoc on the U.S. security infrastructure. The group cracked in to multiple military computers and stole several significant key military software programs. "Chinese Hackers Attack U.S. Military", Security Pro News, John Smith (November 28, 2005).

The private sector has also been subjected to a major rise in cyber crime. The Financial Services sector has seen increased attacks targeting identity fraud where IT attacks have tripled in the last year (according to Deloitte). 78% of the Top 100 global Finance firms have experienced external IT attacks during the last year (as compared with 26% in 2005). "Finance firms face surge in security breaches", Computer Weekly (June 28, 2006). U.K. banks reported a 55% increase in losses from fraudulent online transactions for the first half of the year, mostly from phishing scams. "Online banking fraud dramatically jumps in U.K.", IDG News Service, Jeremy Kirk (November 8, 2006).

Half of the firms have suffered internal attacks as well during the last year. The sophistication of the attacks implies the involvement of professional hacking and organized crime.

---

### **3. What is the challenge with traditional security methods in today's business environment?**

---

#### ***Security hardening limits business functionality***

Traditional methods of hardening systems and applications are to disable certain functions that are detrimental to security; for example by denying internet users access to any system resources (applications and data) which reside in a restricted environment. By disabling them the system designer attempts to address the security problem; however this potentially limits the functionality of the system impacting on business and service delivery. This negative impact on business functionality would not be acceptable to most businesses.

#### ***Zero-day attacks***

These attacks refer to security vulnerabilities lurking undiscovered in most software yet to be exploited and for which patches are not available yet. These could be serious exposures which hackers can exploit and potentially use to take over the enterprise.

#### ***Poor patching of middleware and applications***

Most enterprises do a good job at keeping up to date on patches for their operating systems. However, when it comes to middleware and applications they are less stringent. These are the component weaknesses that hackers exploit. This is because patching at this level represents considerable work for overloaded IT staff with other priorities and application providers themselves will put new release functionality as first priority, rather than security patches. From a security standpoint, patching middleware and applications is as important as patching operating systems.

#### ***Assurance of complex systems and networks***

The ultimate goal is to secure and assure all transactions flowing through the network. Enterprise IT infrastructures are becoming increasingly complicated with hundreds of components connected to one another in complex information supply chains. This distributed nature of the enterprise makes it very difficult to provide a high level of security assurance. Hackers will get into any weak part of the system and then move to other parts of the network. It is proving ever more difficult to assure systems in a distributed IT infrastructure, as opposed to the traditional centralized enterprise IT systems.

#### ***Internal attacks***

There is considerable threat today within an enterprise from its own employees and contractors. Studies have repeatedly shown that around 70% of the attacks against an organization are internal. How does an organization protect against a disgruntled employee, a naïve operator or a malicious administrator? Normally system administrators are granted the highest system privileges and have access to all electronic assets in an enterprise. Organizations need to find techniques for limiting the access rights of these privileged users.

#### ***Exploiting software vulnerabilities***

It is a known fact that all software contains bugs and this will be an ongoing issue for organizations. Some of these bugs are security vulnerabilities. Studies have shown that on an average there is one bug per 1000 lines of code (Information Week 2002, Reasoning Inc. 2003). In large software systems with millions of lines of code, these represent a vast number of software vulnerabilities. Hackers exploit these vulnerabilities to penetrate into systems. These problems are more serious in homogeneous systems where a hacker can exploit vulnerability repeatedly and move fairly rapidly to other parts of the enterprise and cause further damage.

### ***Internal networks are neither robust nor resilient***

Most organizations protect their internal networks by putting many security products around the perimeter (e.g. firewalls, proxy servers, Intrusion Detection Systems Antivirus, Antispam, content filtering, VPN, etc). These are necessary but not sufficient. As a result, most organizations put a considerable amount of trust in their internal networks. Internal networks also need to have sufficient protection mechanisms in place, in the event that an attacker is able to penetrate through all the outer layers of defense.

### ***Privilege escalation***

Considerable studies have been done to understand the characteristics of computer attacks. A standard pattern has been observed whereby hackers get into a weak part of the system and by exploiting software and configuration vulnerabilities are able to escalate their system and administrative privileges. This is a standard pattern hackers use to exploit systems. It involves the hacker accessing the system through a weak insecure point such as an application, then escalating their privileges to perform actions above the remit of the original application..

### ***Lack of secure application containment***

In order to save costs, many enterprises run multiple applications on the same platform. However, in most cases these applications are not adequately isolated from each other. If hackers can break into one application, they are able to penetrate into others with the potential of taking over the entire enterprise. Damage to one application should not be allowed to propagate into others.

### ***Network topology enforcement***

Connectivity is perhaps the biggest enemy of security. Most enterprises today are highly connected, sometimes via the internet. This enables a hacker to move from one part of the enterprise to another, if proper controls are not in place. Enforcing proper connectivity has been a very difficult goal. Confining requirements for each component must be defined and adequate controls must be in place to enforce that requirement.

These factors whether taken individually or collectively represent a major challenge for systems managers. Analysts are now saying that organizations must find new methods to efficiently handle inter-enterprise and interagency information flows across security boundaries while meeting the ongoing operational challenges of their organizations. Traditional infrastructure security has proven to be inadequate in protecting against the emerging threats in high connectivity environments. As Governments and businesses seek to enhance their service delivery capability to citizens and customers, the traditional border controls of standard infrastructure security are no longer 'good enough' in a world of increased cyber threats.

Fortunately *there is* an exciting new security innovation which can play a key role in the resolution of these challenges ...

**"We consider Mandatory Access Control to be a key enabling technology to aid government and businesses alike, in being confident they can deliver more services, more quickly and with better function, without compromising security"**

**Dr. Steve Marsh  
Intelligence and Security Advisor  
UK Cabinet Office**

---

## 4. What is Mandatory Access Control Security?

---

To understand Mandatory Access Control security it is important to first understand standard infrastructure security which is referred to as Discretionary Access Control (or DAC) security.

The DAC model restricts access to computer objects (programs, data sets, application configuration files etc.) based on user identity and the privileges assigned to each class of user. The DAC approach allows a user to “own” objects and grant access to these objects to other users or applications at their (the owning user’s) discretion. Every program or system function which is invoked by that user runs with that individual’s privileges, including mail clients, web browsers, and any downloaded programs from the internet. These programs are empowered with the user’s privileges and are capable of accessing all objects to which the user has been granted access. The reality is that most users of computer systems will be completely unaware of this privileged set of capabilities that they have been assigned or have inherited. Not so the malicious and sophisticated hacker who, having found any weakness in the IT environment can cause software to run that the user did not intend and of which he is probably blissfully unaware. That software will call other software from the network which can escalate privileges and gain overall system control. This is the problem of the so-called Trojan horse and other viruses so common in today’s software systems. Perhaps the simplest example is that of the innocent looking e-mail that contains a malicious attachment with a suitable benign name and file type. Launch that malicious file just once and within a few milliseconds, a whole system can be destroyed – or information stolen. Since administrators are also users, they too are vulnerable to the weakness of DAC. Once the attacker’s illicit software runs as administrator, it will gain unfettered access to the system and in all probability, the rest of the overall network infrastructure. The attacker is then in a position to launch malicious programs across an enterprise or to access the most sensitive information in system data sets.

Mandatory access control (MAC) security resolves the weaknesses of DAC by adding multi-layered additional access restrictions that are not subject to the discretion of users (including administrators). Once installed and configured by the security administrator the MAC security rules remain constant. MAC security limits user access to resources, even those that they create and own, according to a system wide policy defined by the organization’s security administrator. The MAC security policy determines which access rights are granted, typically by what the program does rather than who is running it. Illicit software is contained to the privileges of the vulnerable program and not all the privileges of the associated user.

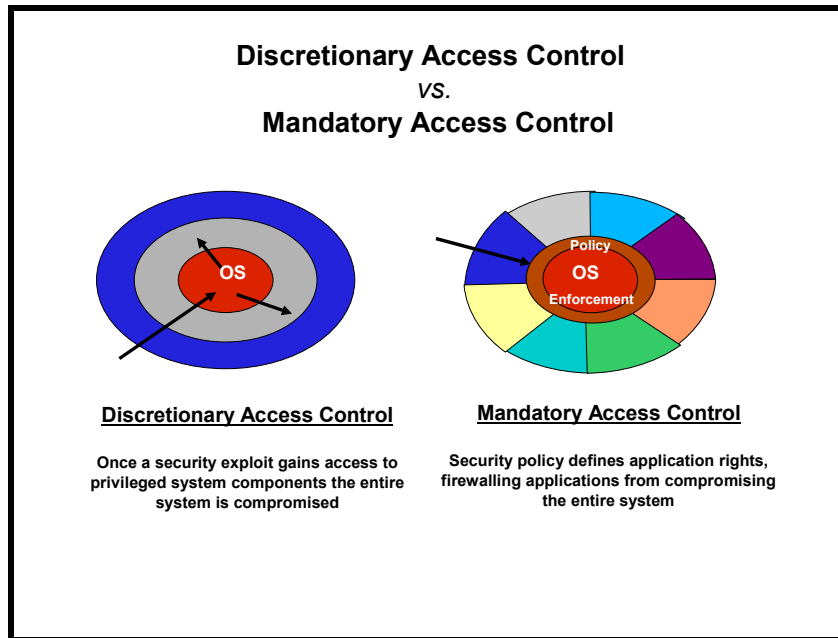
MAC security prohibits the escalation of unwarranted access privileges that are so common in software attacks today. Security Enhanced Linux (SELinux) implements a flexible MAC technology called Type Enforcement. SELinux is developed and maintained by the Open Source community and is included in the mainline Linux 2.6 kernel (and fully supported in Red Hat Enterprise Linux 4). It is the first time that MAC security has been made widely available in a mainstream operating system. Type Enforcement has proven its ability to address a broad range of security problems, from the most sensitive government security challenges to the everyday enterprise security concerns.

**“SELinux is the culmination of over 35 years of operating system security research. For the first time we have MAC security widely available in a mainstream operating system in a form that is flexible enough to meet a wide range of security challenges.”**

**Frank Mayer**  
Chief Technology Office, Tresys Technology  
Co-author, SELinux by Example

## 5. How does Mandatory Access Control Security address the threat of cyber crime while improving business delivery?

The primary capability that Mandatory Access Control security gives us is that of confinement. It allows us to place boundaries between applications such that it is much harder for a hacker to cross those boundaries. This section addresses how MAC security addresses shortcomings associated with traditional infrastructure security practices.



### ***Removes the need to limit business functionality***

It is no longer necessary to secure the infrastructure by limiting application and business functionality. Application functions that bring business benefit may remain enabled after hardening (i.e. they can be explicitly permitted). The SELinux policy can be used to enforce a specific mode of operation that cannot be broken by the applications. Programs are given only the privileges they need to perform their function and no more. Since SELinux has *flexible* MAC technology, the security policy can be adapted for any type of organizational security challenge. As much or as little security as desired can be implemented without impacting program functionality.

### ***Limits the threat posed by zero-day attack***

Applications are primarily used and tested for functionality, and even the most stringent code review will not highlight all potential security flaws. By using the flexible SELinux MAC technology to limit individual applications to only the privileges they require, unknown security flaws will be limited in their ability to access system resources. This way, even if the application is compromised, it will be unable to breach the bounds set upon it by SELinux. Any newly publicized application flaws may be taken advantage of on “zero day,” but the attacker will be restricted to only those resources permitted within the SELinux MAC policy. Most current operating systems do not provide this type of application confinement; running an SELinux enabled system is one of the few, and currently the best, ways to achieve it.

### ***Limits the threat posed by poor patching***

Rather than just protecting against unknown application flaws, SE Linux can confine an attack against known flaws. Applications or middleware that have not been patched to the latest level pose the same threat as unknown flaws, but in this instance there may be more potential attackers due to the greater awareness of the flaw. Again, an attacker is confined to only those resources for which the application was intended. This restricts the attack to a single application and stops the spread across to other applications, throughout the operating system, or even across the network to other machines. While not a substitute for proper patch management, SELinux provides added protection against known vulnerabilities.

### ***Provides platform level assurance***

Since SELinux can isolate and protect individual applications, the operating system and middleware can be assured separately from the applications. This makes the assurance process much shorter, cheaper, and less time consuming when changing or upgrading an application. Using SELinux, the operating system and middleware are protected from applications. Thus, when changing/upgrading applications, only the updated code is put through a new assurance process. This usage will also protect against malicious code updates during upgrade by ensuring the application cannot gain additional access to resources. During development, this may also aid application debugging from a security point of view. Another way of looking at this is to say that middleware and conventional business applications are all “applications” as far as the operating system is concerned. By tightly and explicitly defining, within the kernel of the operating system, what application functions are permitted (and only those functions), a significant level of protection against misbehaving application and middleware is achieved.

### ***Minimizes exposure to internal attacks***

With SELinux, it is possible to define access for individual applications and not depend upon the privileges of the user running the program. If such a SELinux policy is implemented, it is possible to entirely remove the concept of a super-user account on a system. This would restrict the amount of damage an internal attacker would be able to inflict by restricting access to key resources such as files and network connections. Any internal attacker would be isolated to the environment defined by SELinux and unable to infiltrate other applications or systems, locally or across a network link.

### ***Limits damage caused by software bugs***

Applications can have a well understood and defined run-time environment. As a result, if the application were to start behaving in a different way than expected, this may be a security risk. By enforcing the defined run-time environment using SELinux policies organizations can protect again this security threat. This would apply to applications that have any known/unknown bugs which may or may not constitute a security exposure. If the application is upgraded at any point it can be run within the same SELinux MAC policy, thus protecting against any new software bugs.

### ***Improves internal network security***

Today network security at best is controlled at the system level with some type of firewall technology. In general, all applications within a given system can access all the network resources (ports, hosts, network interfaces) that the system is allowed to access. This leaves many opportunities for network and privilege escalation attacks. With SELinux, the network resources can be limited down to the level of the individual applications. For example, an application on a middle-tier system may be allowed to access all hosts on the Internet but be restricted to which hosts it may access on the internal network even though other applications on the same system can access the internal network. This approach vastly

improves the security of the network by extending the same protection SELinux provides for system resources to network resources. Simply breaking into one network service no longer gives the attacker access to the internal network.

### ***Constrains administrative accounts***

Today, if an attacker is able to penetrate a network and compromise an application gaining administrative privilege over it, or worse on an entire system, they would be able to read any data or cause any damage they wished. A system running SELinux has the capability to restrict all user login accounts, including administrative accounts. In this way, even if an attacker gained access to a system, privilege escalation would either not be possible or the user would be vastly restricted once escalated. SELinux can be used to prevent onward attack through other applications and across the network to other systems. Another advantage here is for any middleware that requires administrative privilege in order to run; such privilege can be granted to the middleware but restricted precisely to those areas necessary to function correctly.

### ***A provision for secure application containers***

With SELinux applications can run at the operating system or middleware level with the least privileges required. This is the central security design concept of least privilege, which SELinux can enforce with or without the cooperation of the applications. Least privilege prevents applications from interfering with each other either by accident, or by malicious intent. This would allow, for example, an application server to run multiple applications on the same system that could not interact with each other under normal or any unusual circumstances. Isolating applications in this way by placing them in a container with least privilege is the best way of ensuring ultimate system security between applications on a per system basis and across an organization.

### ***Enforces network topology***

All multi-tier network applications intend that certain applications access back-end resources and others do not. However, with the ability to control network resources only at the system level there is no way to enforce this goal. SELinux, with its ability to restrict network resource access for an individual network resource, can strongly enforce the network topology and information flow the application developer intends. This finer level of control provides much stronger security architecture for the entire network enterprise, and again avoids many possible zero day attacks.

---

## 6. Mitigating procurement fraud in the UK National Health Service

---

**“In the WebSphere prototype, for the first time we were able to demonstrate that a multi-system SELinux MAC policy can significantly improve the security of a N-tier application without undue impact on the application’s functionality or administration”**

**Frank Mayer**  
Chief Technology Office, Tresys Technology  
Co-author, SELinux by Example

IBM undertook a proof of value project with the UK Cabinet Office and Tresys Technology (Columbia, Maryland) to validate the merits of SELinux as an enabler of secure Government Transformation. Tresys developed a commoditized set of SELinux MAC policies and tools called Tresys Razor and adapted them for the WebSphere Application Server running in a standard N-Tier architecture configuration on IBM Intel xSeries technology. This solution used SELinux to enhance the security of the N-Tier application in such a way as to be configurable to an end-user environment without requiring the user to be an expert with SELinux.

The goal of the proof of concept was to demonstrate that MAC security could be enabled in a mainstream commercial off-the-shelf middleware application. The SELinux-based solution was applied to Belmin Group’s ARIES eProcurement application which runs on top of IBM’s WebSphere Application Server. The ARIES (automated reconciliation & invoicing) application replaces existing manual methods of invoice processing to reduce costs and increase efficiency. ARIES provides an internet facing service, utilizing IBM WebSphere to automate reconciliation and invoicing processes between the UK National Health Service (NHS) and external suppliers. The ARIES proof of concept was then tested in a production environment providing the service to County Durham & Darlington Acute Hospitals NHS Trust in the UK. Penetration tests showed that all attempts to escalate privileges in the application environment were successfully prohibited by the WAS SELinux solution. When the SELinux MAC policy enhancements were enforced there was no noticeable performance degradation of the ARIES application and from the client end user perspective, SELinux was invisible.

The coupling of MAC security and WebSphere Application Server demonstrated both the system assurance and web application value necessary to deliver innovative applications across an inter-enterprise network infrastructure.

**“Security is paramount to Belmin and its users, who comprise many UK public sector organisations, including 94 NHS Trusts and The Department for Work & Pensions. Belmin solutions currently process some £920 million per annum of sensitive UK public sector purchasing expenditure. The beauty of the SE Linux implementation from our perspective as a solutions developer, was that when we enabled the extra security layer, nobody was aware that it was even running. In the past security has often compromised application functionality but SE Linux has had minimal impact whilst providing a whole new tier of assurance. To be able to apply such granularity of control without any detriment to service delivery is both progressive and compelling.”**

**Richard Stoneman**  
Marketing Manager, Belmin Group Ltd.

---

## 7. Conclusion

---

Until now, the market adoption of Mandatory Access Control security has been limited to proprietary trusted systems. These environments have been prohibitively expensive for most organizations and limited solely for the use of bespoke applications (e.g. mostly niche requirements for National security related initiatives). We have seen that Mandatory Access Control Security is now fairly easy to set up and once in place provides a far more secure environment for all applications.


The advent of SELinux enables Governments and businesses to experience the power of Mandatory Access Control security in a mainstream, commercially available operating system. The same information assurance that has been sought after for deep custom security requirements is now generally available to mainstream organizations who are serious about addressing the increasing threat posed by cyber crime.

SELinux policies have already been built for WebSphere Application Server and are available for reuse for those wishing additional information assurance. IBM and Tresys are now also progressing MAC policy security solutions for other core IBM middleware applications including DB2 and Tivoli.

Organizations seeking to improve their agility and business innovation potential should strongly consider the merits of MAC security as a key component of an overall 'defense in depth' security strategy. This capability can mean the difference in an optimized delivery of innovative business services while providing reliable system-wide enterprise network security.

**"For a long time it was believed that implementing & deploying Mandatory Access control was hard. Through this effort we have proved otherwise. In addition, we have observed that such strong security can be achieved in a very non-intrusive manner with almost no application performance penalties"**

**Doc Shankar, Executive IT Architect, IBM Federal Strategy/Architecture**



---

## 8. Acknowledgments

---

### **Dr. Steve Marsh, Strategic Advisor, UK Cabinet Office**



Steve is Strategic Advisor on Intelligence and Security in the Cabinet Office. He was Director of the Central Sponsor for Information Assurance from its formation in October 2002 until July 2006.

Prior to this, Steve was Director of Security Policy in the Office of the e-Envoy, responsible for establishing a common framework for the security of electronic government systems. This included the ways by which individuals and business users authenticate themselves when using electronic government services.

In April 2000 Steve joined the Central IT Unit in the Cabinet Office, which merged with the Office of the e-Envoy later that year. He has over 18 years experience in security and IT within the public sector, and was the winner of the European Information Security Award 2005 for excellence in the field of policy.

### **Marc Hocking, Senior Security Architect, UK Cabinet Office**

Marc is Senior Technical Security Architect in the Cabinet Office. Marc is responsible for determining the technical level approach for the Enterprise Security Architecture Strategy for the UK Government. Prior to this Marc worked as a Senior Technical Architect in the eGovernment Unit of the Cabinet Office.

Marc previously worked for Price Waterhouse Coopers' beTrusted Security Division in Australia as the Senior Security Architect for the Asia Pacific region.

### **Frank Mayer, Co-Founder, President and Chief Technology Officer, Tresys Technology**

Frank has 23 years of experience in the research, design and development of computer and network security technologies. Under his leadership, Tresys has become one of the industry's most trusted providers of high-end information security technology and engineering services. Mr. Mayer spearheads the business and R&D direction at Tresys. He is the chief architect of Tresys extensive involvement in Security Enhanced Linux (SELinux), initiating many SELinux open source projects and products, creating related commercial products, starting the annual SELinux Symposium, and co-authoring of the definitive book on SELinux.

### **Doc Shankar, Certified Executive IT Architect, Linux Security Lead**


Doc is a certified executive IT architect at IBM. Over the last 5 years, Doc has led IBM's security initiatives regarding the Linux Operating System. Since climbing the Linux bandwagon, Doc has focused on ways to improve security in Linux-based systems. When not in the office, he can be found convincing customers (or anyone) that Linux is secure! He has a PhD in Computer Sciences from UC, Berkeley.

### **Fred Logan, WW Program Director, IBM Software Group, Open Source and Linux**

Fred is a senior consulting sales specialist working in IBM's Open Source and Linux business. He works closely with Governments around the world to help position and implement the advantages of IBM's Open Source and Linux software technologies. Fred's previous IBM roles have been focused on managing regional sales execution in Europe and North America for IBM Software Group.

### **Graham White, IT Specialist, Emerging Technology Services - Linux Integration Centre**

Graham White works from the IBM lab in Hursley/UK where he has garnered deep Linux skills in various areas for almost 6 years. One of Graham's major focus areas is to develop, setup and proof security concepts of IBM Middleware on Linux as well as security concepts of the Linux operating system in general for IBM Business Partners and customers.



© Copyright IBM Corporation 2006  
All Rights Reserved.

IBM Canada  
8200 Warden Avenue  
Markham, ON  
L6G 1C7  
Canada

Printed in United States of America  
06/06

IBM, IBM (logo), AIX, DB2, DB2 Universal Database, eServer, Tivoli, Tivoli Enterprise Console, TotalStorage, and xSeries are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both. Windows is a trademark of Microsoft Corporation in the United States, other countries, or both.

Intel is a trademark of Intel Corporation in the United States, other countries, or both. UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates. The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

The information in this white paper is provided AS IS without warranty. Such information was obtained from publicly available sources, is current as of 01/30/2005, and is subject to change. Any performance data included in the paper was obtained in the specific operating environment and is provided as an illustration. Performance in other operating environments may vary. More specific information about the capabilities of products described should be obtained from the suppliers of those products.