Rob Shook, Linux Sales Enablement and Marketing, IBM Software Group
September 2010

# Open Source Software: Asking the Right (Hard) Questions

# Why do IBM consider open source software important?

- OSS can be a major source of innovation
  - Innovation can happen anytime, anywhere
  - Development through “open communities” leads to potentially broad ideas and creativity
- OSS is a good approach for developing emerging standards
  - Popular open source projects can become *de facto* / open standards
  - Wide distribution deployment
- OSS is a source of competition in the marketplace
  - Office productivity applications (word processing, spreadsheets, presentation)
  - Operating systems (Linux for servers, desktops)
  - In some areas, perhaps the only growing competitor to a single, established vendor

# Some of IBM's history with open source software

## 1999 - 2001

- IBM forms Linux Technology Center
- Leads Apache projects Xerces (XML4J), Xalan, SOAP
- Creates OSI-approved IBM Public License
- Strategic participation in Mozilla
- IBM becomes founding member of OSDL
- Founder of Eclipse.org *and Eclipse Consortium*
- Creates internal bazaar using OSS methodology

## 2002 - 2003

- Linux contributions to scalability (8-way+), reliability (stress testing, defect mgmt, doc)
- Leads Apache projects: Web Services (WSIF and WSIL), Pluto (Portlet API) and WSRP4J (Remote Portal)
- Leads Eclipse projects GEF (editing), EMF (modeling), XSD (XML Schema), Hyades (testing), Visual Editor, AspectJ, Equinox rich client
- Globus Toolkit contributions for OGSA, OGSI

## 2004 - 2006

- IBM and Novell/SUSE achieve security milestone (EAL4+ and COE compliance)
- *Eclipse Foundation, Inc.* becomes independent - IBM contributes UML2, Voice Tools, *Aperi, COSMOS, Ajax Toolkit Framework*
- Globus Toolkit 4 is WS-I compliant
- Pledged 500 patents to open source
- Partner with Zend PHP
- IBM enhances Apache partnership
  - Contributes Derby database
  - Contributes voice recognition
  - Supports Geronimo J2EE project
  - Acquires Gluecode for skills
- IBM contributes accessibility code to Firefox

## 2007 - 2009

- IBM leads Open AJAX initiative and announces partnership with the Dojo Foundation
- Leads Apache Tuscany project
- Contributes code for security management to Eclipse Higgins
- IBM contributes to Apache Lucene project and announces OmniFind Yahoo! Edition
- IBM joins OpenOffice.org
- IBM joins Open Health Tools, moving code for medical record management from Eclipse OHF
- Contributes Open Services for ALM – Jazz REST Services samples
- IBM creates ODF Toolkit Union
- Contributes Lotus Notes/Domino app templates to OpenNTF.org
- Incubates Apache Wink (JAX-RS)

**More than 1000 IBM developers involved in OSS projects**

**IBM leads 80+ OSS projects**

**IBM contributes to 150+ OSS projects**

# What does it mean when we say that IBM supports open source?

- IBM has more than 800 developers working full time on open source projects, especially Linux, Eclipse, and Apache projects.
- IBM has invested billions of USD on open source.
- IBM participates in and support pro-open source organizations such as the Linux Foundation, the Software Freedom law Center, and the Open Invention Network.
- IBM collaborates with many other companies and people around the world on open source projects.
- However …
  - IBM does not “bless” every open source project in the world
  - Some open source projects compete with our products and we compete against them.

# How does IBM use open source?

- To run our business
- In our hardware
- In our software
- As part of service engagements
- As an R&D collaboration vehicle
- As a way of influencing the direction of the IT industry
- As a way of “leveling the playing field”
- As a way of invigorating stagnant market categories

**Double compute capacity with no increase in consumption or impact by 2010**

***IBM’S PROJECT BIG GREEN SPURS GLOBAL SHIFT TO LINUX ON MAINFRAME***

ARMONK, NY, August 1, 2007

# Is software good software, just because it is open source?

- It depends of your definition of “good,” but by most definitions, the answer is “no.”
- As of three days ago, a popular code repository listed 164,297 open source projects.
- Statistically, you might imagine that some are better than others.
- Your definition of “good” is critical.

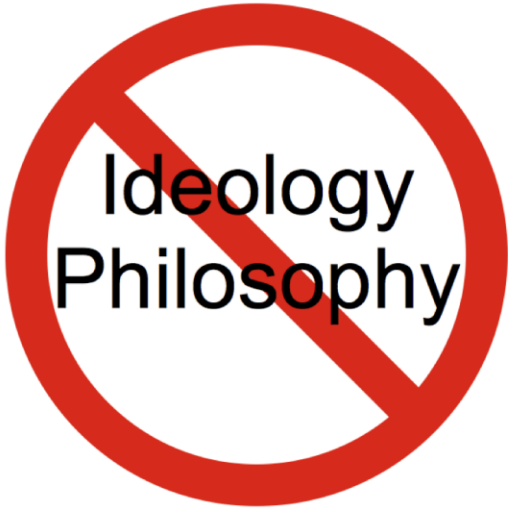

# Is the code well architected and implemented?

- Great code may start with the germ of a fantastic idea, but it eventually gets rewritten one or more times to be faster, more reliable, more secure, and more extendable.
- If you are not an expert yourself, seek independent assessments of the quality of the code.
- The quality of the documentation and user interface are important considerations in their own rights, but may also give you an idea of how well-designed the core of the software is.

## Who are the founders, contributors, and users?

- *People* write code and drive software projects and products.
- Unreliable people may place the future of the software … and thus your investment … in jeopardy.
- Work out “what if” scenarios for what you will do if the code gets abandoned, forked, or acquired.
- Learn what other users have done with the code and about the quality of their experiences with the software and those who created it.

# What is the form and governance of the community?

- Find out if the open source code you are considering is being developed by a healthy, democratic, and meritocratic community or if it is really just a controlling company “coding in public.”
- Learn if the community also includes documenters, graphic designers, and evangelists in addition to coders.
- Look at the project forums, Facebook, Twitter, and other social networking tools to get a sense of the health of the community.
- Don't ignore warning signs of trouble in the community and things that may make you uneasy about it.

# Are there intellectual property issues involving copyrights or code provenance?

- Ignoring legal issues with software can be one of your most expensive mistakes and can literally put you out of business.
- Learn about open source licenses and consider hiring an intellectual property attorney as a consultant when you are considering use of software or negotiating a contract.
- Don't mix open source licenses unless it is legal.
- Make sure the developers of the software you want to adopt played by the legal rules.
- Don't pretend to be an attorney if you are not.

# Does the license suit all your future plans for the code?

- Some open source licenses can be combined and others cannot.
- Some open source licenses allow for free use in commercial, “closed source” applications and others do not.
- Some open source licenses specify some restrictions when you host software-as-a-service.
- Be especially careful if you want to use open source code libraries.
- Understand if the software you plan to use can be hosted on either a private or public cloud.

# Do you have proper legal controls and business processes in place to deal with open source software?

- That is, what is your open source governance strategy?
- Five years ago, it was not uncommon for that strategy to be defined as “you shall use no open source software.”
- You need to understand the legal risks and responsibilities for any software you use, and weigh those against the business value.
- Work out a plan that specifies what business and legal controls are in place to approve use of open source in your organization or in your products, and make sure you have a well defined escalation path.

# Is the software enterprise-ready?

- There's been a lot of discussion about whether open source software is more secure than proprietary software.
- Which open source software and which proprietary software?
- In addition to security, you need to look at reliability, availability, scalability, interoperability, and performance.
- Make sure the software is available on the right hardware platform so you can optimize the environment for your workload.

# Who will maintain your installation of the software?

- If you are planning for your IT staff to install and maintain your software, make sure it doesn't get orphaned when you have personnel turnover.
- When software updates come along, you will need a plan to decide which ones to install and when, especially if major releases come along every six months or so.
- If you customize open source code for your organization, are you prepared to propagate those changes into newer versions of the code?

# What is different about providing services around open source?

- Much less than you may think, compared with proprietary software
- Most IBM services today somehow involve open source in the tools used or the software deployed.
- How do you get open source software fixed?
  - –You do it
  - –You engage with vendors such as Red Hat, Novell, or others to do it
  - –You outsource support
  - –You engage with the community directly

## How easy is it to integrate the software with your data or other software you already use?

- Does your software use recognized industry standards or does it have its own way of formatting data?
- Are the developers of the software involved in creating the standards that will allow interoperability?
- Are the standards used consortium, nationally, or internationally approved?
- Do the standards provide the interoperability needed in your organization, government, or region?
- If you adopt the software, who will do the integration tasks?
- Is the software certified for use on the operating system and hardware platform you plan to use?

# Are benchmarks available to allow performance evaluations of the software with comparable products/projects?

- Benchmarks are used to measure the performance and behaviour of software in a defined environment so that you can easily compare offerings from different providers.
- While benchmarks can be abused, they can be important in learning if particular software is really usable in your business.
- You might worry less about published benchmarks and more about proofs of technology and head-to-head comparisons among the software choices you are considering.
- Consider your software provider's response to such requests for “bake offs” when making your adoption decision.

# In summary: areas to consider

- Quality and performance of the code
- Community history, health, and governance
- Legal issues
- How open source and proprietary software will work together
- Standards
- Code maintenance, deployment, and service management
- Your governance of creation and use of open source

# Conclusion

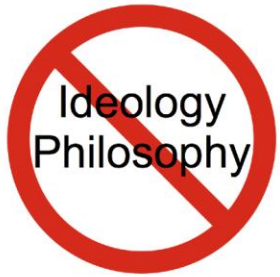


- First and foremost, open source software is *software*.
- When it comes to business and especially enterprise use, open source software should get no immediate free pass because it happens to be open source.
- Conversely, proprietary software should also be measured on a level playing field with open source, and get no special initial treatment.
- All those things that you worried about when choosing proprietary software—security, performance, reliability, availability, interoperability, support, maintenance—are also areas to investigate when considering open source software.

# Legal

**Trademarks and Disclaimers**

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

IBM, the IBM logo, BladeCenter, Calibrated Vectored Cooling, ClusterProven, Cool Blue, POWER, PowerExecutive, Predictive Failure Analysis, ServerProven, Power Systems, System Storage, System x , System z, WebSphere, DB2 and Tivoli are trademarks of IBM Corporation in the United States and/or other countries. For a list of additional IBM trademarks, please see http://ibm.com/legal/copytrade.shtml.

The following are trademarks or registered trademarks of other companies:

Java and all Java based trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries or both
Microsoft, Windows,Windows NT and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries, or both.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries or both.
Linux is a trademark of Linus Torvalds in the United States, other countries, or both.
Cell Broadband Engine is a trademark of Sony Computer Entertainment Inc.
InfiniBand is a trademark of the InfiniBand Trade Association.

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

**Notes**

Linux penguin image courtesy of Larry Ewing (lewing@isc.tamu.edu) and The GIMP

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Users of this document should verify the applicable data for their specific environment.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Information is provided “AS IS” without warranty of any kind.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the 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.

Prices are suggested US list prices and are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

The 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

Rob Shook, Linux Sales Enablement and Marketing, IBM Software Group
September 2010

# Open Source Software: Asking the Right (Hard) Questions

# Linux as a Catalyst for a Smarter Planet

Rob Shook
Manager, Linux Sales Enablement and Marketing
IBM Software Group

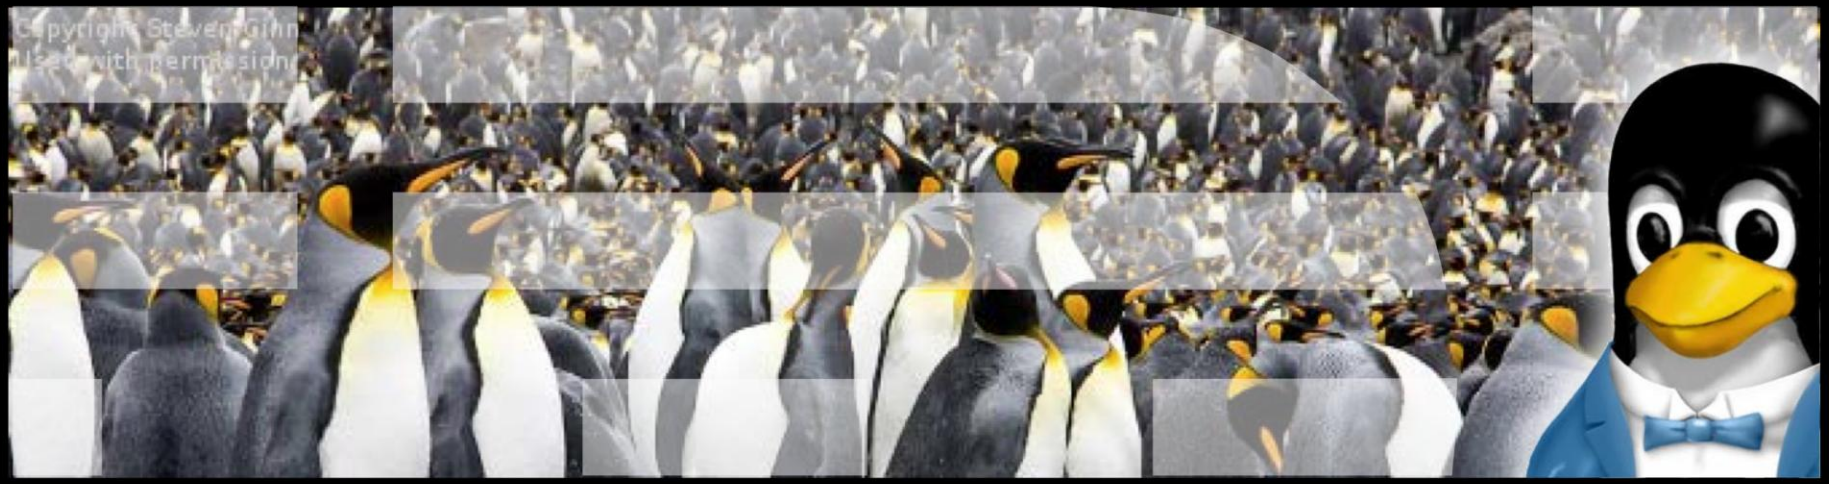

# Global forces are driving a fundamentally different world

- The global financial crisis changed business priorities – and the IT that supports it
- The business landscape is evolving, and IT must evolve with it
- Technology has enabled solutions that weren't feasible in the last downturn
- Fast-developing communities drive constant technology change

# Something meaningful is happening

Our world is becoming
**INSTRUMENTED**

Our world is becoming
**INTERCONNECTED**

Virtually all things, processes and ways of working are becoming
**INTELLIGENT**

# The opportunity for progress is clear

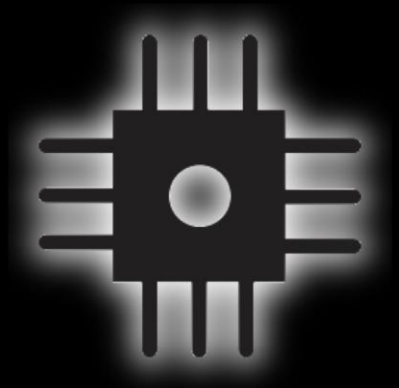

## INSTRUMENTED

**We now have the ability to measure, sense and see the exact condition of everything**

- There are 1 billion transistors for each person on the planet[1]

**Everything will become instrumented:**

- Supply chains
- Healthcare networks
- Cities
- Natural systems

[1] Sam Palmisano speech, November 12, 2008

# The opportunity for progress is clear

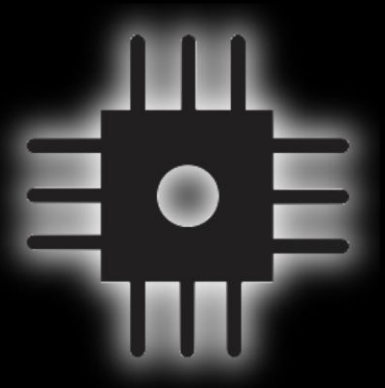

## 100%

**Analog electric meters replaced with smart devices in Malta**

### Monitoring quality and optimizing consumption

The Maltese National Electricity and Water Utilities — Enemalta Corporation (EMC) and Water Services Corporation (WSC), use Red Hat Enterprise Linux to implement a nationwide smart grid for electrical and water service

# The opportunity for progress is clear

## INTERCONNECTED

**People, systems and objects can communicate and interact with each other in entirely new ways**

- Almost one third of the world will be on the web by 2011[1]

**The Internet of connected devices…**

- Cars
- Appliances
- Roadways
- Pipelines
- Livestock

**…is headed to 1 trillion.**

[1] Sam Palmisano speech, November 12, 2008

# The opportunity for progress is clear

## 70%

**PERCENT OF SAO PAULO'S EMISSIONS DUE TO 6,000,000 CARS**

**Improving quality of air (and life) with vehicle testing**

As the 7th largest city in the world and the 5th most polluted, Sao Paulo decided to control vehicle pollution through a new distributed vehicle inspection system, enabled by a scalable, robust, and secure infrastructure built on Red Hat Enterprise Linux

# The opportunity for progress is clear

## INTELLIGENT

**We can respond to changes quickly and accurately, with better results**

- 15 petabytes of new information are generated every day[1]
- An average company with 1,000 employees spends $5.3 million a year to find information stored on its servers[1]

**New computing models and advanced analytics use the massive amounts of generated data to become smarter**

[1] Sam Palmisano speech, November 12, 2008

# The opportunity for progress is clear

# 2.2 Million

**DISPENSING ERRORS FROM HANDWRITTEN PRESCRIPTIONS[1]**

**Improved patient care and research through informatics**

A North American research hospital uses Linux to align internal processes and create a robust, closely governed data warehouse, improving quality of patient and research data, reigning in costs, and improving efficiency

[1] Sam Palmisano speech, November 12, 2008

# The progress to date is inspiring…

Smarter cities like Dubuque, which will be enabled by a Linux-based cloud and analytics from IBM

Smarter medicine: Astellas Pharma reduces drug discovery with Linux clusters from System x

Smarter oil and gas exploration with Shell, simulating more and speculating less

Smarter traffic systems in Stockholm, driven by Linux

In 1910, only 16 cities had 1M or more people.

In 2010, the number has grown to 450.

«—»

In 2007, half of the human population lived in cities.

In 2050, the concentration will increase to 70%.

http://www.ibm.com/smarterplanet/us/en/sustainable_cities/ideas/
http://www-03.ibm.com/press/us/en/pressrelease/28420.wss

# The progress to date is inspiring…

Smarter cities like Dubuque, which will be enabled by a Linux-based cloud and analytics from IBM

Smarter medicine: Astellas Pharma reduces drug discovery with Linux clusters from System x

Smarter oil and gas exploration with Shell, simulating more and speculating less

Smarter traffic systems in Stockholm, driven by Linux

In a major hospital, stents costing $25k often disappear

«—»

In a clinic, a patient receives repetitive tests due to out-of-date medical records

«—»

HealthGrades estimates 195,000 patient deaths are avoidable[1]

http://www.ibm.com/smarterplanet/us/en/healthcare_solutions/ideas/index.html
http://www-01.ibm.com/software/success/cssdb.nsf/CS/JSTS-7NRRHY

# The progress to date is inspiring…

Smarter cities like Dubuque, which will be enabled by a Linux-based cloud and analytics from IBM

Smarter medicine: Astellas Pharma reduces drug discovery with Linux clusters from System x

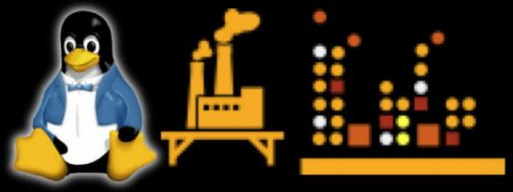

Smarter oil and gas exploration with Shell, simulating more and speculating less

Smarter traffic systems in Stockholm, driven by Linux

Projected demand for energy is expected to increase by 50% from 2005 to 2030

«—»

Only 33% of oil is pulled from existing reservoirs

«—»

A 1.5% improvement in oil recovery can yield 6 months of consumable oil

http://www.ibm.com/smarterplanet/us/en/oil_exploration/ideas/
http://www-03.ibm.com/press/us/en/pressrelease/29538.wss

# The progress to date is inspiring…

Smarter cities like Dubuque, which will be enabled by a Linux-based cloud and analytics from IBM

Smarter medicine: Astellas Pharma reduces drug discovery with Linux clusters from System x

Smarter oil and gas exploration with Shell, simulating more and speculating less

Smarter traffic systems in Stockholm, driven by Linux

In the US, urban populations grew by 20% from ’82 – ’01

«—»

Traffic grew 236%

«—»

Building new roads and new lanes is no longer feasible in many urban areas

http://www.ibm.com/smarterplanet/us/en/traffic_congestion/ideas/index.html
http://www-03.ibm.com/press/us/en/pressrelease/29903.wss

# Yet the reality can seem daunting…

Numerous system integrations are required to make *"anything"* smarter

Tremendous analytic power is needed to discover *"new treatments for cancer"*

Massive amounts of data flowing from hundreds of thousands of *"smart meters"* must be read multiple times per hour

Staggering numbers of images must be captured, stored, managed and linked to *"billing and collection systems"* in real time

# Especially in light of today's challenges

**41%** of data center managers claim their data centers will max out their energy capacity within one to two years[1]

Processor power doubles every 18 months, but **85%** of this power sits idle[2]

**80%** of digital data is now unstructured[3] and requires greater effort to transform it into usable intelligence

# Requirements are changing for a more instrumented, interconnected, intelligent world

## Due to Changing Workloads

e.g. Real-time optimization of a Power Grid by reading meters every 15 minutes…

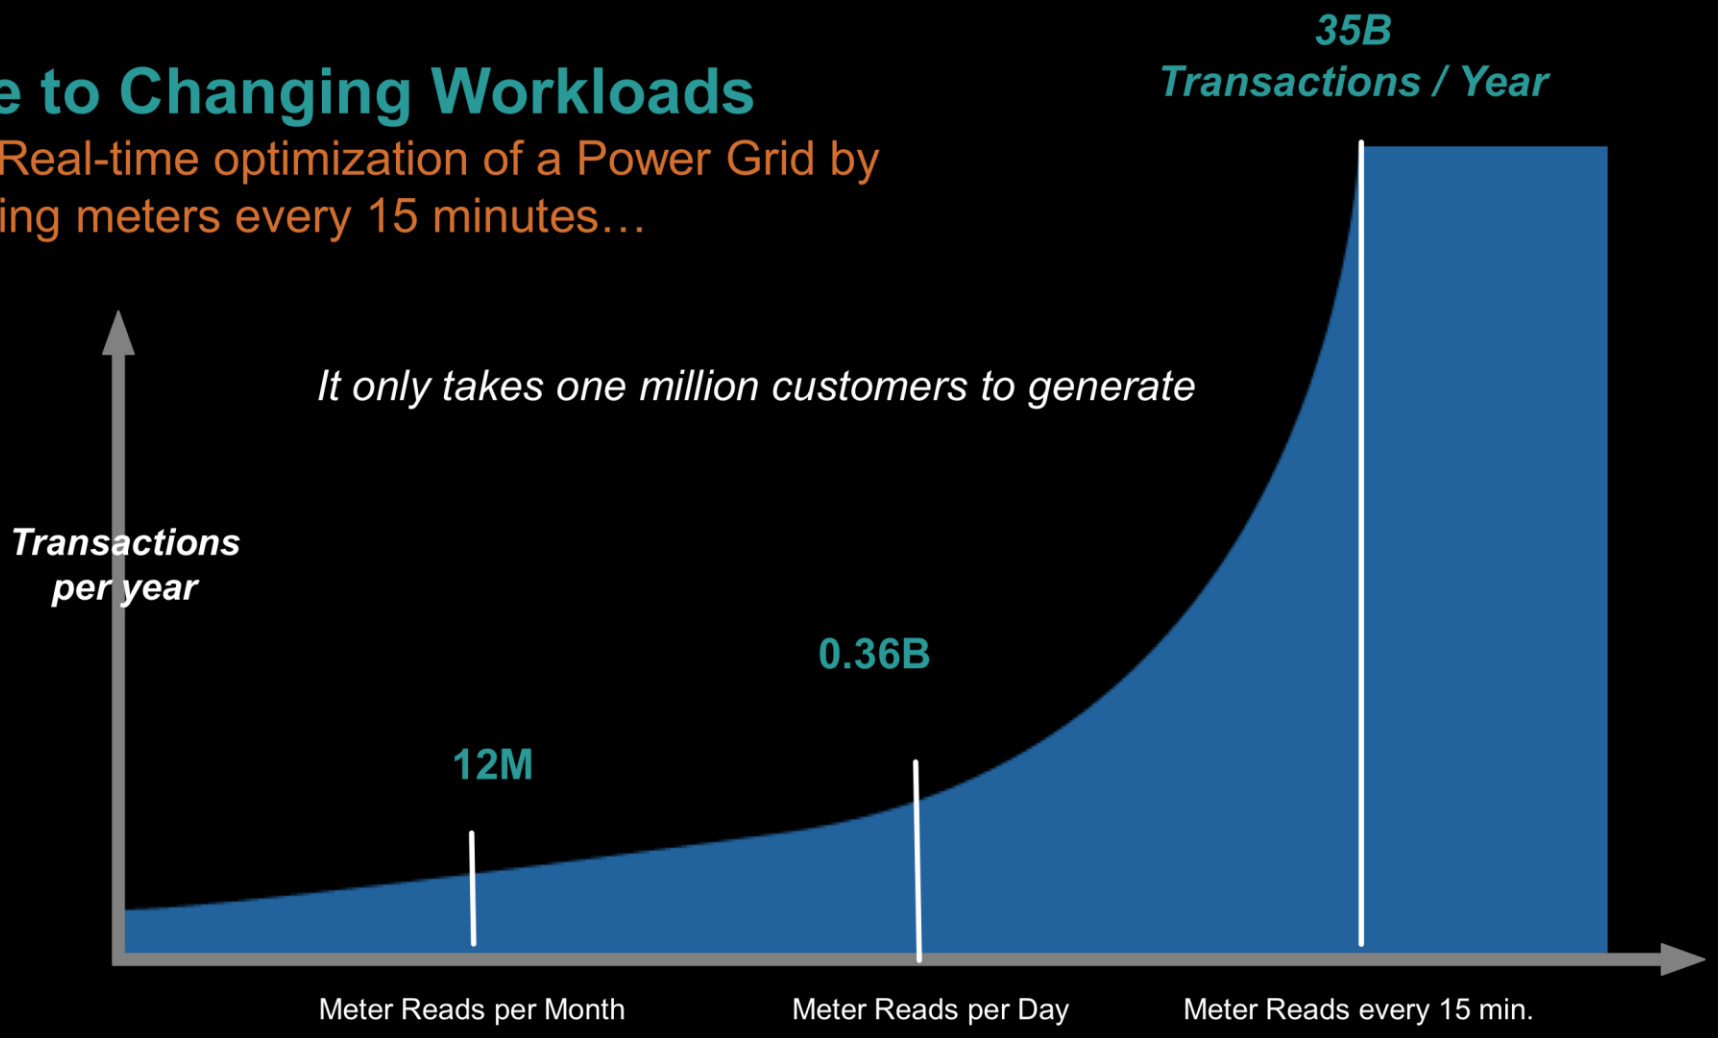

# Linux is at the core of the datacenter, and a smarter planet

***Linux*** *continues to enable new ways of doing business*

## Edge and Web Infrastructure

- Community focus
- Internet enabled
- Worldwide Volunteers

## Application and Data Serving

- Commercial focus
- Open elements of IT industry join community
- Enterprise Linux adoption grows

## Business-Critical Workloads

- Competitive focus
- Extensive use for DB, BI, ERP, CRM
- Cornerstone of IT strategies
- Steady adoption through downturn

## Next Generation Workloads

- Leadership focus
- Accelerated adoption
- Workloads drive platform decision
- Cloud / flexible allocation models
- Fully established for business-critical use

1991 – 2004 | 2005 – 2006 | 2007 – 2009 | 2010+

# IBM provides complete Linux solutions: top-to-bottom, end-to-end

IBM Services, Support, and Financing

Information Management | WebSphere® | Tivoli® | Lotus® | Rational®

IBM Systems Software

IBM System x | IBM Power Systems | IBM System z

IBM Systems Storage

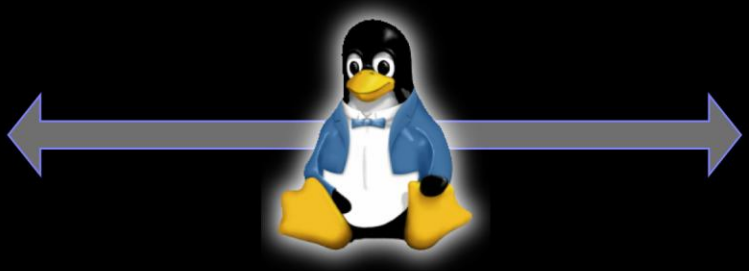

# Smarter solutions with Linux start with IBM


**IBM provides leadership solutions with Linux, top to bottom, end to end**

- Choice of supported hardware
- Enterprise-grade middleware
- Implementation and support services

# Legal

Trademarks and Disclaimers

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries:

IBM, the IBM logo, ibm.com, Smarter Planet and the planet icon, BladeCenter, Power, System Storage, System x , System z, WebSphere, DB2 and Tivoli are trademarks of IBM Corporation in the United States and/or other countries. For a complete list of IBM trademarks, please see www.ibm.com/legal/copytrade.shtml

The following are trademarks or registered trademarks of other companies:

Java and all Java based trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries or both
Microsoft, Windows,Windows NT and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries, or both.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries or both.
Linux is a trademark of Linus Torvalds in the United States, other countries, or both.
Cell Broadband Engine is a trademark of Sony Computer Entertainment Inc.
InfiniBand is a trademark of the InfiniBand Trade Association.

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

NOTES:
Linux penguin image courtesy of Larry Ewing (lewing@isc.tamu.edu) and The GIMP

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Users of this document should verify the applicable data for their specific environment.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Information is provided “AS IS” without warranty of any kind.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the 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.

Prices are suggested US list prices and are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

The 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.

# Thank you