Kohki Ishikawa (kixs@jp.ibm.com)
IBM Japan, Ltd.

8th June, 2012

A Linux Application Tool to Leverage the Full Capability of Hardware

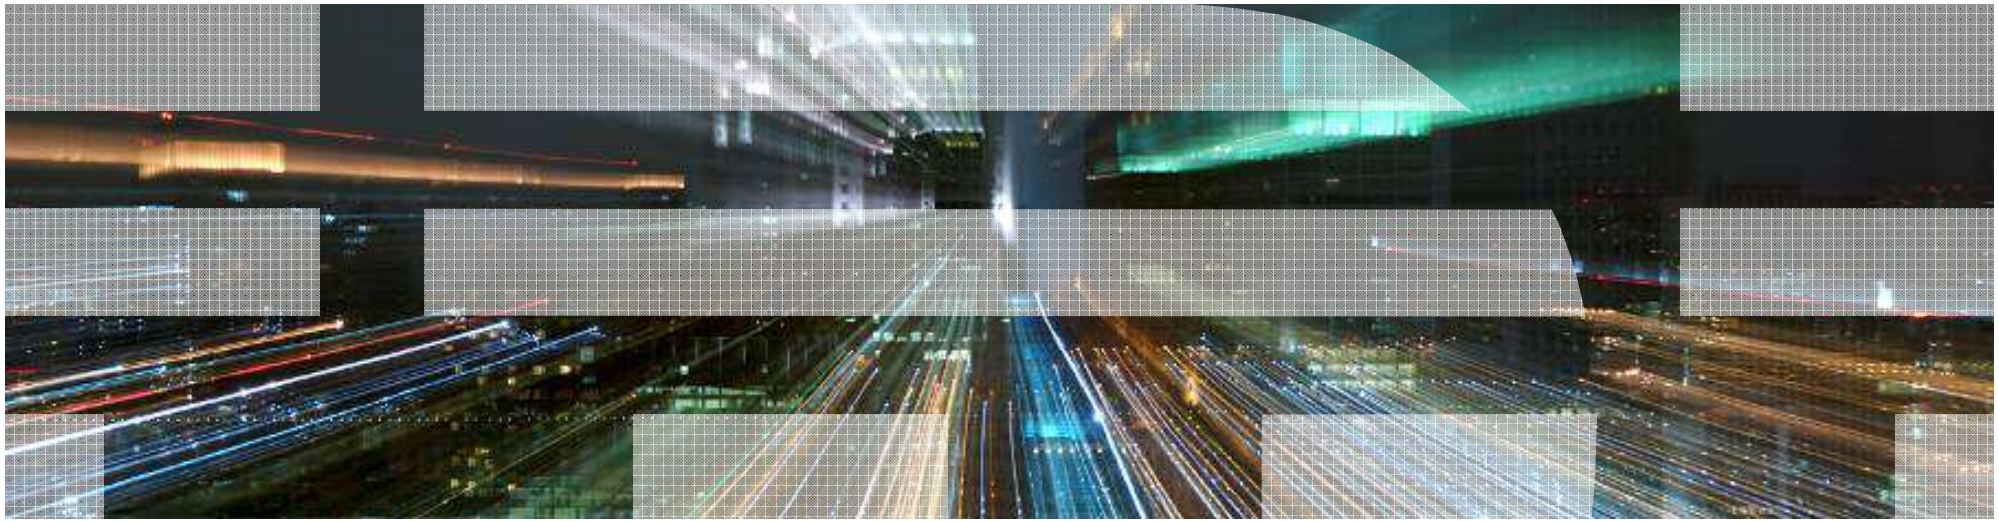

Agenda

- Linux and Open Source Software Trends
- POWER processor & Power Systems benefit to Linux
- Advanced Toolchain
  - Overview
  - Components
  - How to use
  - Example Adaption for PostgreSQL 9
- SDK for PowerLinux
  - Overview
  - Features
  - How to use
- Conclusion

Linux and Open Source Software Trends

- Becomes more “enterprise”
  - This report says
    - 73% companies now place open source equal to or ahead of proprietary software
    - 68 % companies use Linux, which is the most popular open source package than others

* please refer to the source, original website

Source : OpenLogic http://olex.openlogic.com/wazi/2011/survey-shows-enterprise-open-source-usage-ubiquitous/

From Edge to Business Critical Workloads

- Adoption area of Linux and Open Source Software seem to expand from edge towards backend workloads

Edge and Web Infrastructure

Characteristics:
- Community Driven
- Internet Enabled
- Worldwide Volunteers

Application and Data Serving

Characteristics:
- Open Industry Driven
- Open elements of IT industry join existing community
- Linux adoption in the enterprise accelerates

Business-Critical Workloads

Characteristics:
- Competition driven
- Accepted as mature, open, lower-cost platform for DB, BI, ERP, CRM
- Cornerstone of datacenter strategies
- Steady adoption through downturn

Next Generation Workloads

Characteristics:
- Innovation driven
- Fully established for business-critical use
- Workload allocation by platform capability
- IT-led cloud adoption
- Utility billing models
- Accelerated adoption post-downturn

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

PPC and PPC64 architecture

- Embedded, Super Computers, Game Consoles, Appliance and Enterprise Servers

IBM Power Systems

Blue Gene

* Please refer to endor website

OpenBlocks Micro Server

http://openblocks.plathome.com/

Pathfinder

Spirit

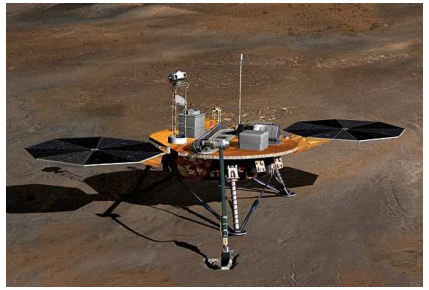

Phoenix

POWER processor & Power Systems benefit to Linux

- POWER is
  - a RISC microprocessor architecture developed by IBM
  - a microprocessor implementation of the POWER ISA

- IBM Power Systems
  - uses IBM POWER processor
    - Now the latest processor generation is POWER7
  - supports running Linux
    - Red Hat Enterprise Linux, SUSE Linux Enterprise Server
      Fedora, Debian…
    - IBM calls these Linux environments on Power Systems as “PowerLinux”

- POWER & Power Systems can provide to Linux users and market
  - additional choice of server hardware
  - more Scalability & Reliability

More Scalable

- More numbers of CPU Core are supported

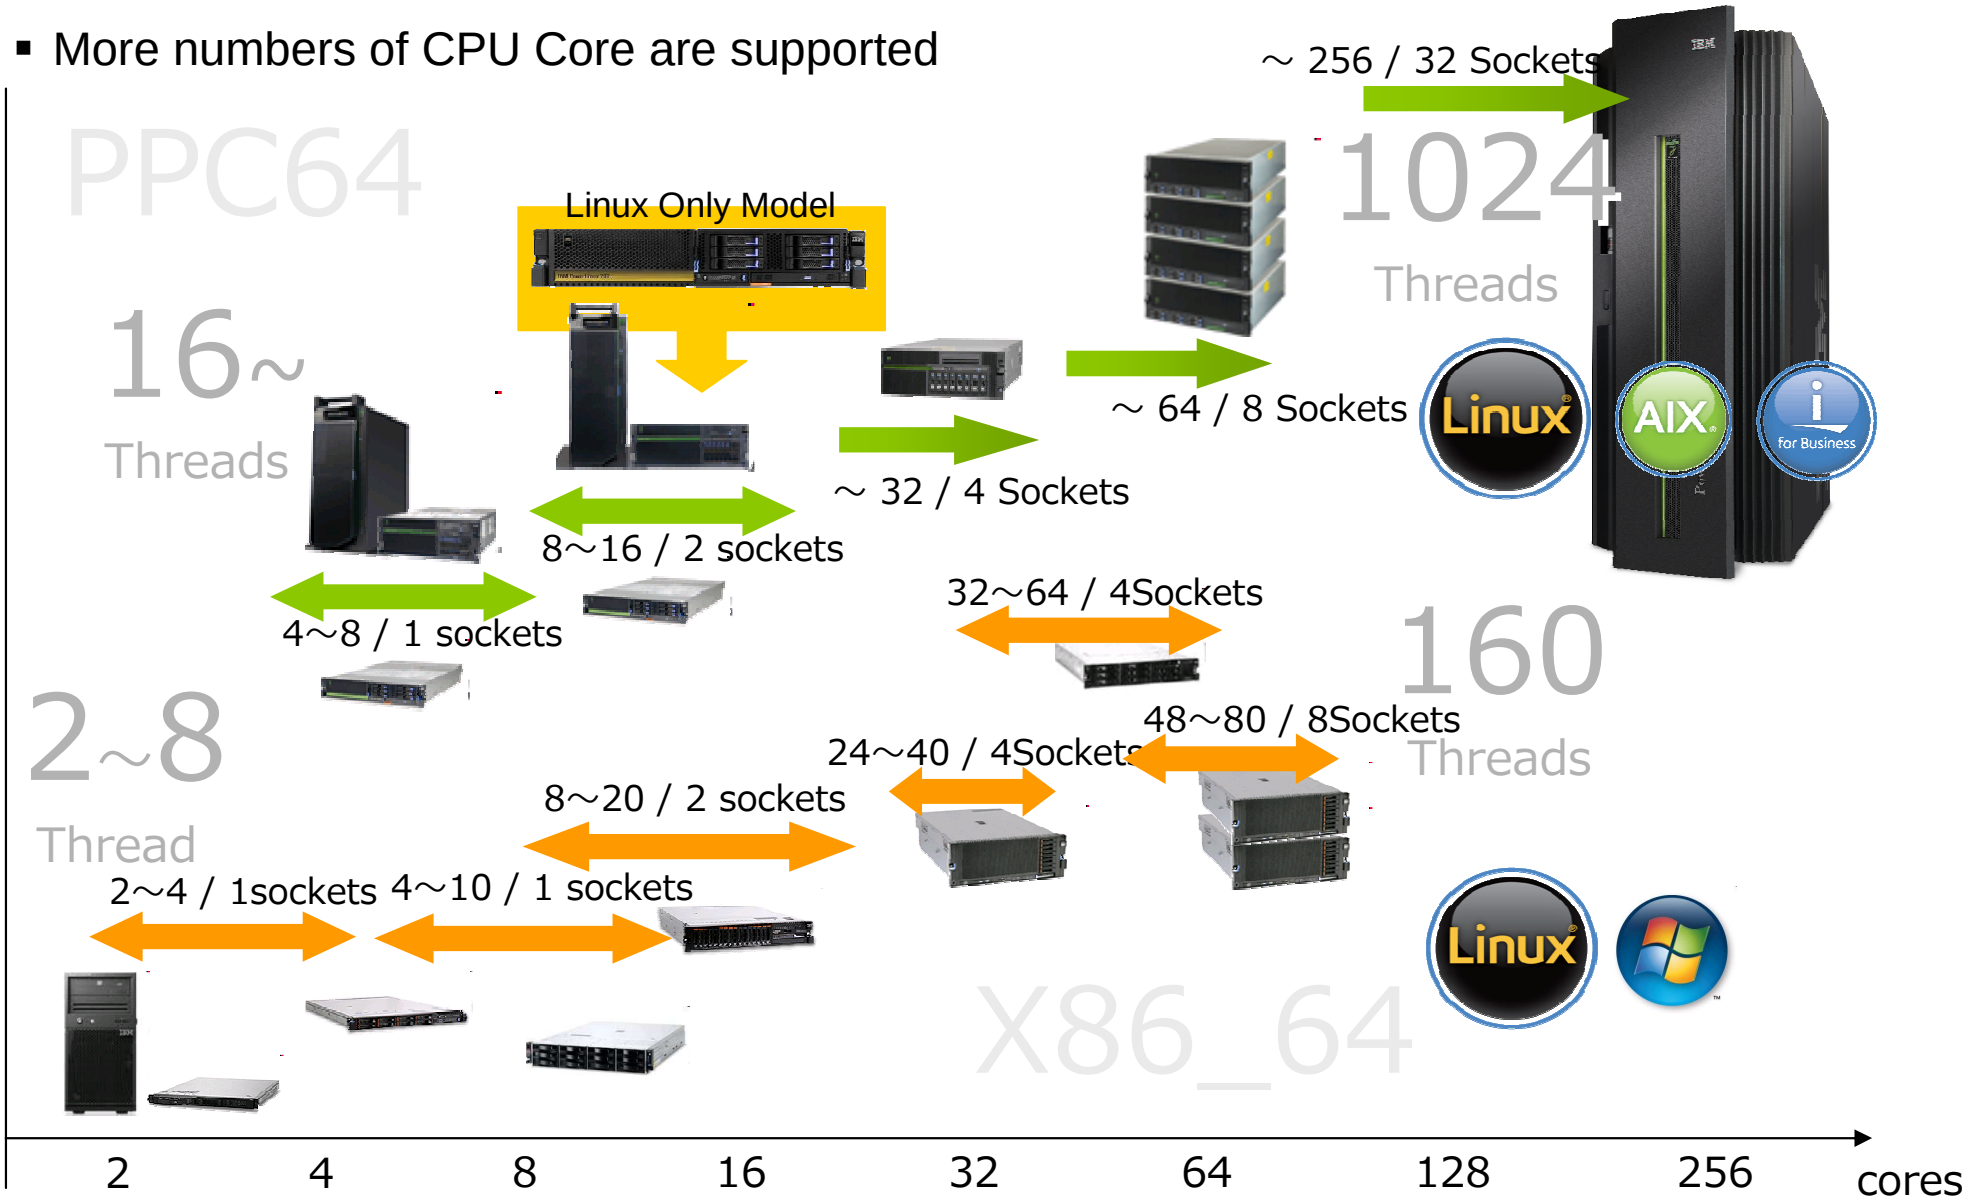

Source : http://www.ibm.com/systems/jp/x/product/ & http://www.ibm.com/systems/jp/power/hardware/ at Mar. 2012

More Reliable

- Trying to run continuously when an error occurs

Processor Recovery

Transparent error recovery by Hardware

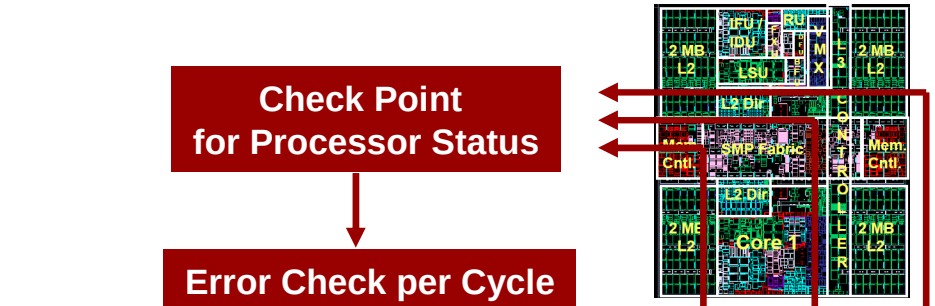

PCI Bus Error Recovery

Capability to improve the reliability of PCI/PCI-E bus peripherals

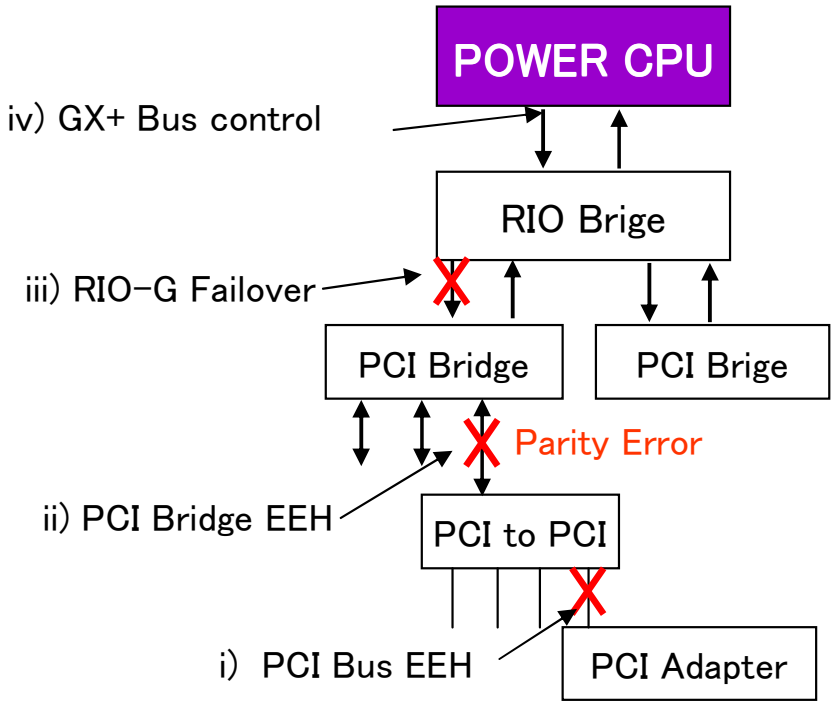

EEH is introduced in the competing session.
Please follow-up after finish of conference

Application is much portable

- Linux Application is portable although between different architectures.
  - On commercial distribution of Linux, such as RHEL or SLES, although different architecture,
    - Linux kernel is built from same version of source code, so available almost completely same function on kernel
    - Bundling middleware and libraries are also same version, so available almost completely same API for middleware or libraries

      ex) RDBMS, Language Runtime, Graphic Libraries…etc.
  - Recent application is coded by “portable” Language
    - Script Language, such as Perl, PHP, Ruby
    - Java
  - Web-App-RDB 3 Tier Application depends on only middleware layer
- Only C/C++ Applications or Libraries modification should be careful.

Two Major helpers for more easily C/C++ application development

- Advance Toolchain Linux on Power Systems

- Software Development Kit for PowerLinux

Advance Toolchain Overviews & Component

- Available from University of Illinois Web Site
  - A set of open source software development extensions and tools allowing users to take leading edge advantage of IBM latest hardware features:
    - POWER6 enablement
    - POWER6 Optimized scheduler
    - POWER6 Native DFP instruction support
    - POWER6 VMX enablement with auto-vector
    - POWER7 enablement
    - POWER7 Optimized scheduler
    - POWER7 Native DFP instruction support
    - POWER7 VMX/VSX enablement with auto-vector
    - ppc970, POWER4, POWER5, POWER5+,POWER6, POWER6x, POWER7 optimized system and math libraries
    - libhugetlbfs 2.0 support

AT5.0

http://www.ibm.com/developerworks/wikis/display/hpccentral/How+to+use+Advance+Toolchain+for+Linux+on+POWER

How to use AT

- The repository Information is available from following URL:

  http://www-304.ibm.com/webapp/set2/sas/f/lopdiags/yum.html

- Easy to install
  - Recent versions are available through online repository, such as yum and zypper.
  - Prior to install AT, install locally only 1 package
    - For example on RHEL

```
# rpm -ivh ibm-power-repo-1.1.6-5.ppc.rpm
```

  - After that, Just execute online installation command
    - For example on RHEL

```
# yum install advance-toolchain-at5.0-runtime
or
# yum install advance-toolchain-at5.0-devel
```

- Flexible, easily co-exist and switch the multiple versions of tool chains

```
$ ls /opt
at4.0  at5.0
```

AT5 Example Adaption for PostgreSQL 9

- PostgreSQL is
  - a famous open source RDBMS
  - developed by PostgreSQL Development Community
- PostgreSQL contains a self benchmarking tool, pgbench
  - pgbench execute serveral types of workloads
    - Offline Batch Transaction (pgbench default, TPC-B)
    - Read Only Queries (pgbench -S)
    - Online Mix Transaction (pgbench -N)
- pgbench -S mainly depends on CPU and memory

Results of Adaption of AT5.0

pgbench –S , scale factor 100, 500 sec,
32-core POWER7 on Power 750 (128 threads)
max_connections = 32, shared_buffer = 8GB

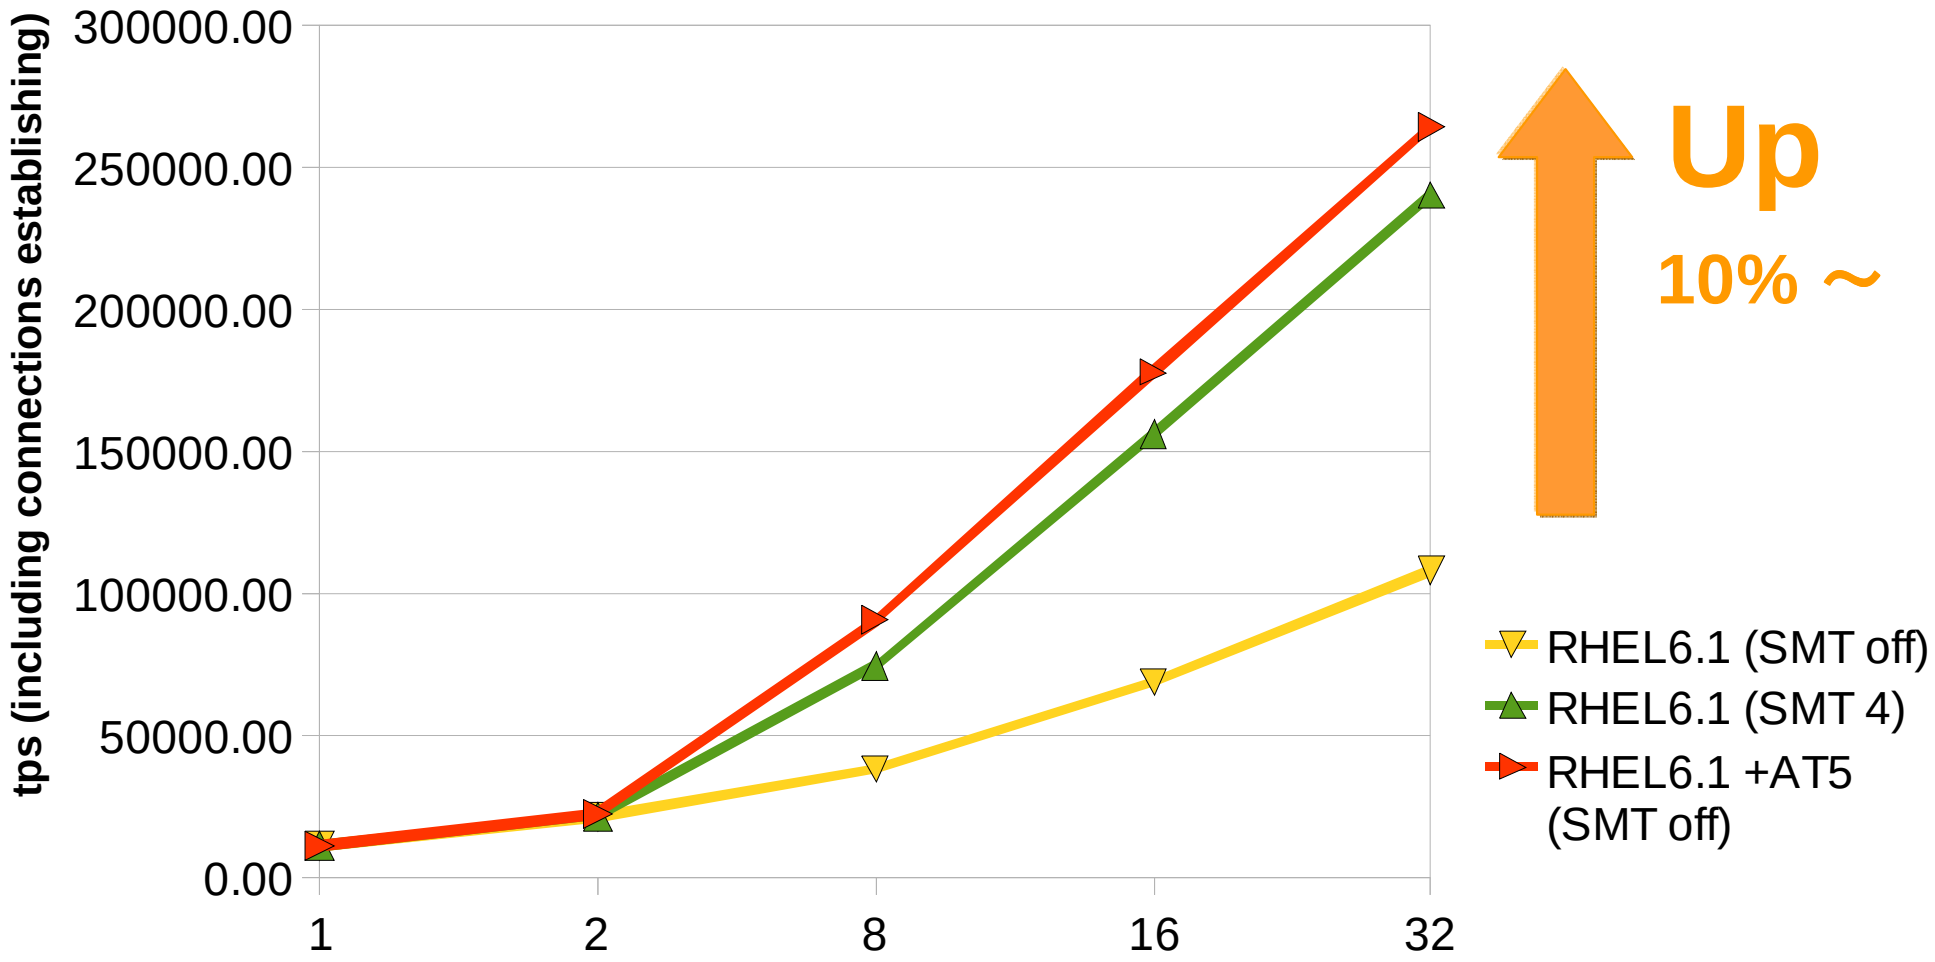

SMP scalability is not automatically realize

- AT5.0 also provides other functions:
  - Decimal Floating Point Library (libdfp)
  - GNU Binary Utilities (ld, ldd, objcopy, objdump, nm, and others)
  - GNU Debugger (gdb)
  - performance analysis tools (Oprofile, Valgrind, gprof, mtrace, xtrace, iTrace)
  - The AUXV Library (libauxv)

- But does not automatically covers
  – SMP scalability

http://www.ibm.com/developerworks/wikis/display/hpccentral/How+to+use+Advance+Toolchain+for+Linux+on+POWER

Example: PostgreSQL 9.2dev Modification on POWER

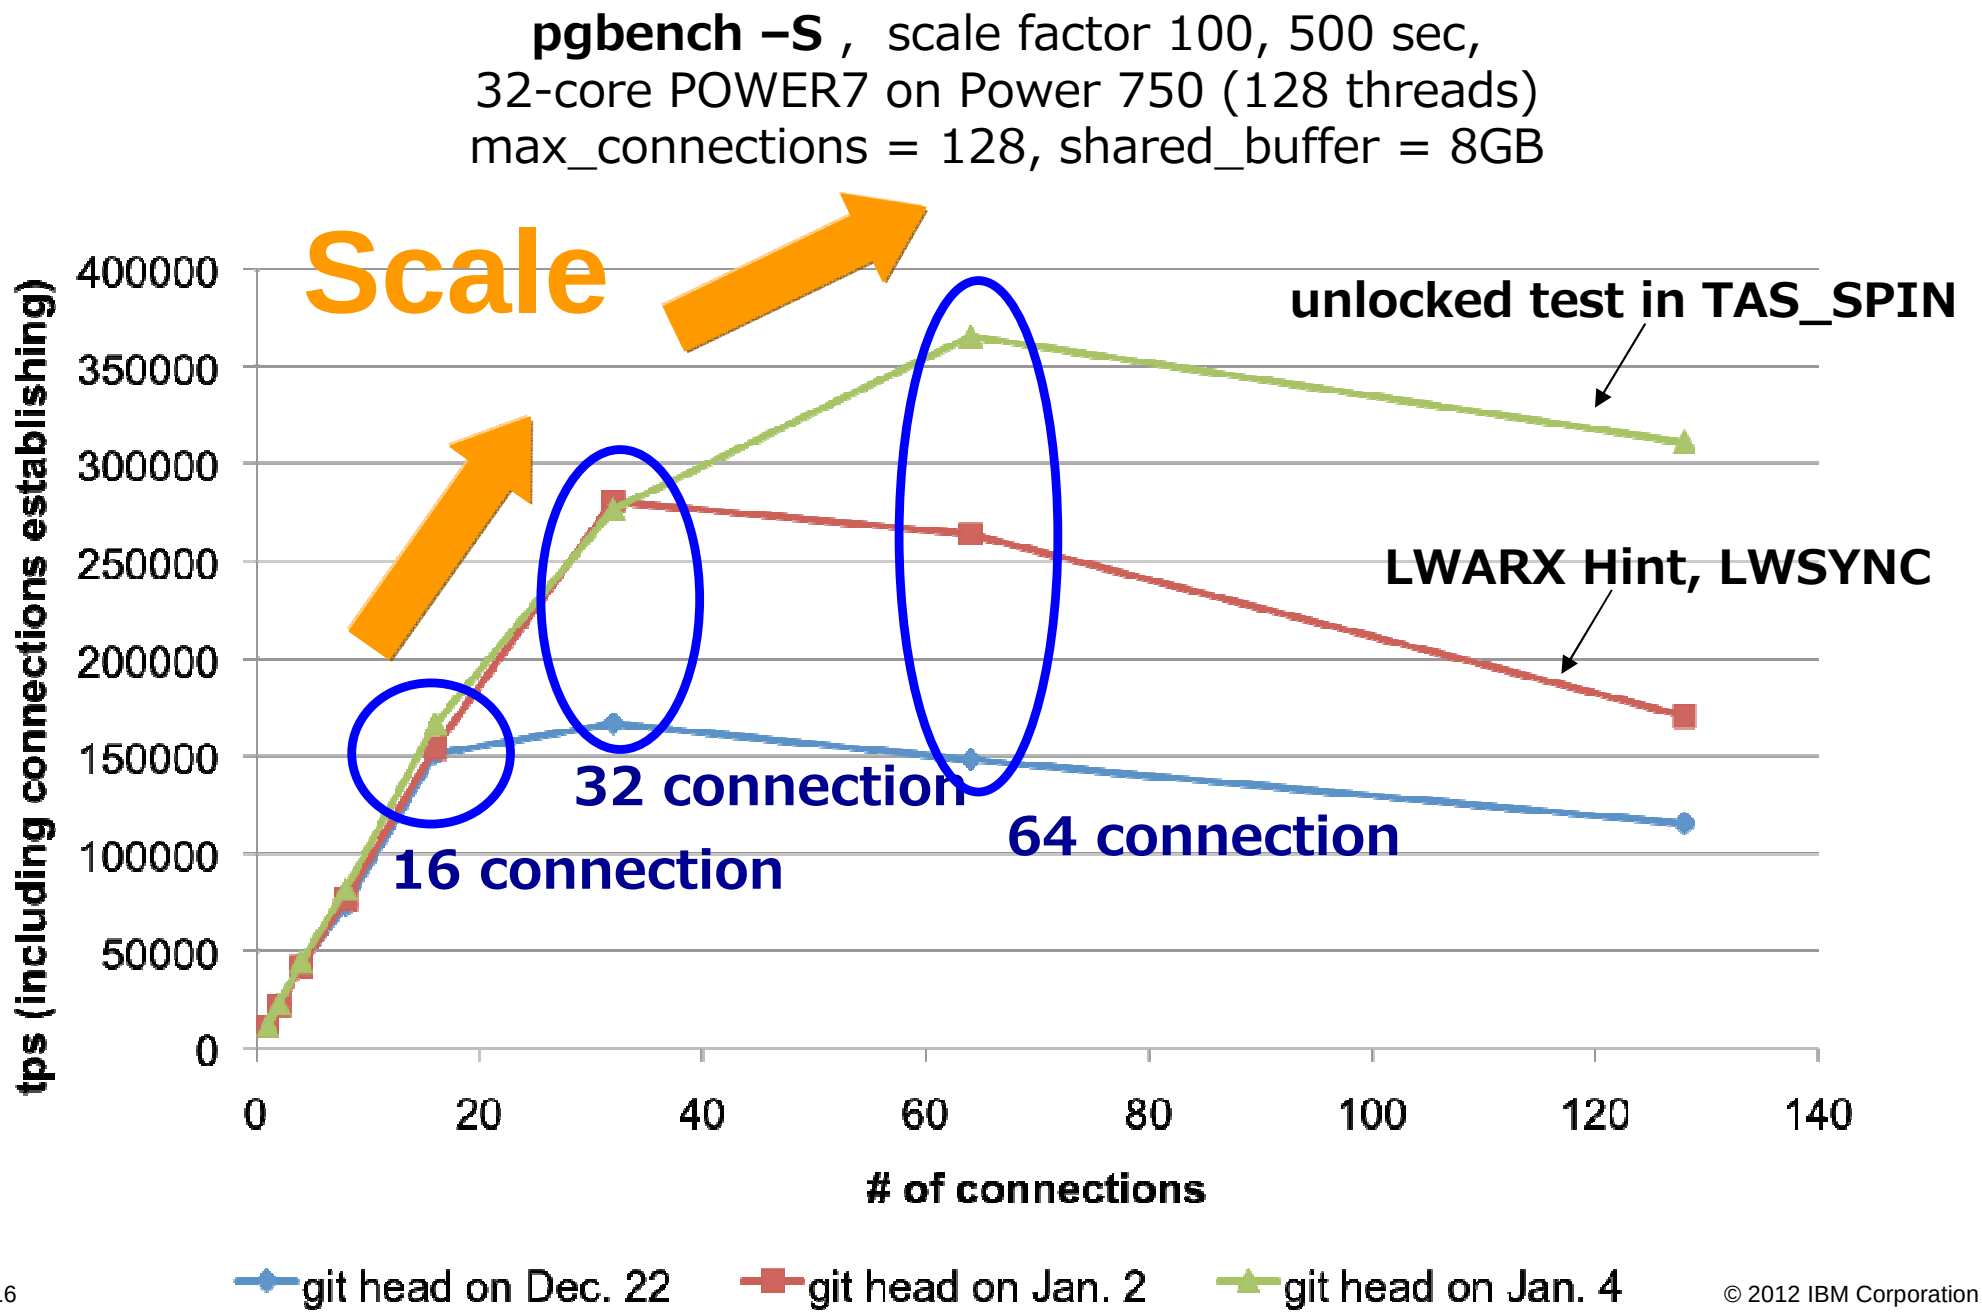

To Leverage the Full Capability

- SMP scalability improvement is important
- Analytics and Suggestion may help your development
  - System Analytics
  - Code Suggestion

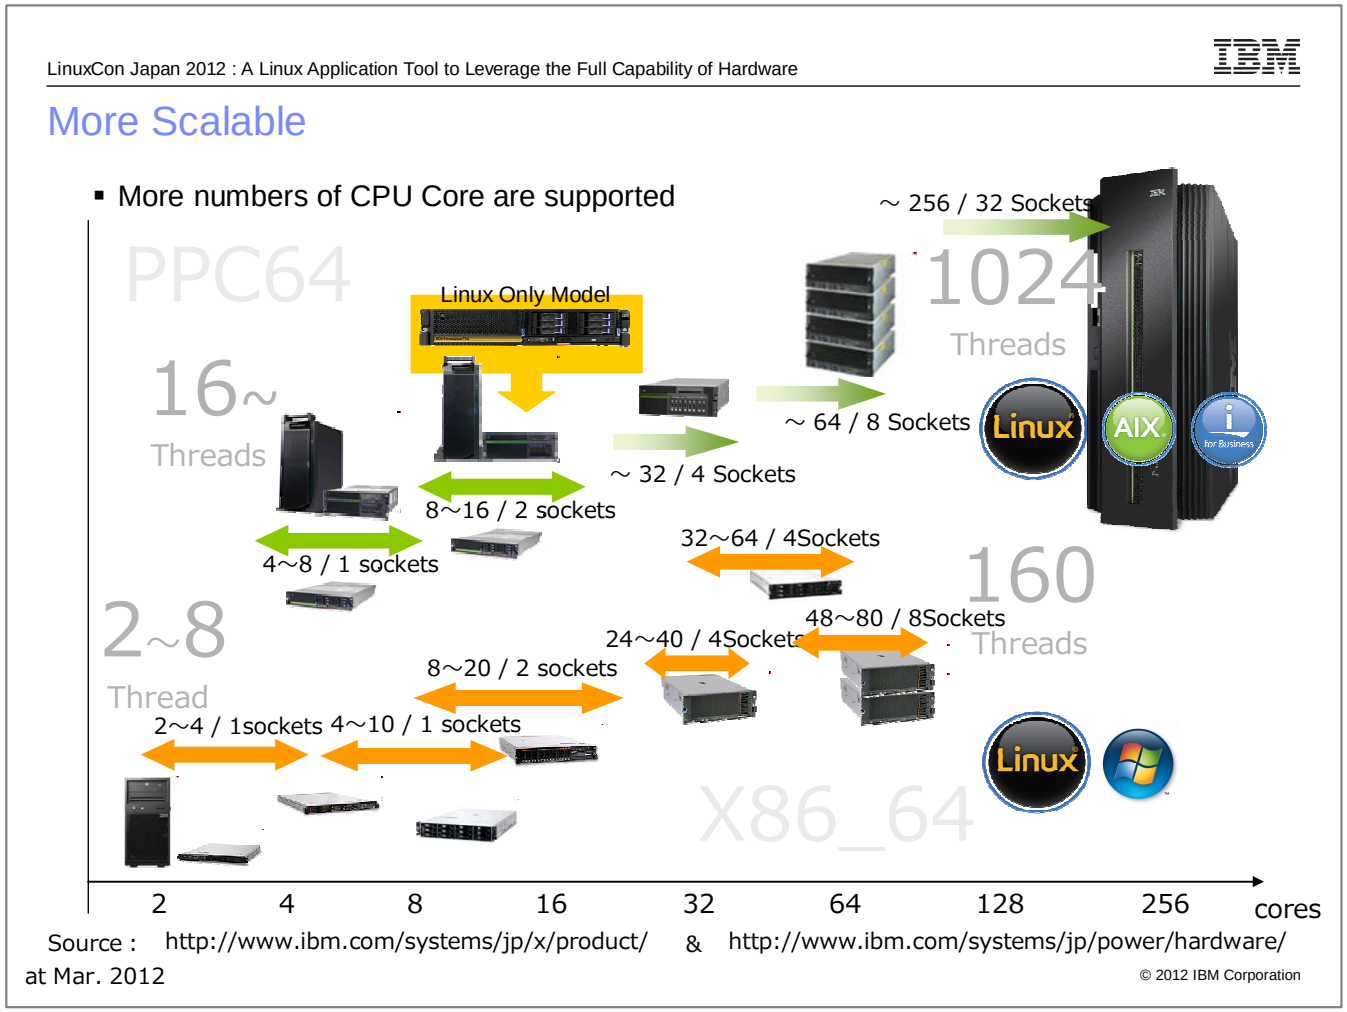

* please refer to slide #7 in this presentation

SDK for PowerLinux

- New released in 2012
- Eclipse-based IDE
  - CDT
  - Code Analytics Tool
  - GFE
  - PTP
- All-in-one solution for developing softwares on PowerLinux (Linux running on Power Systems)

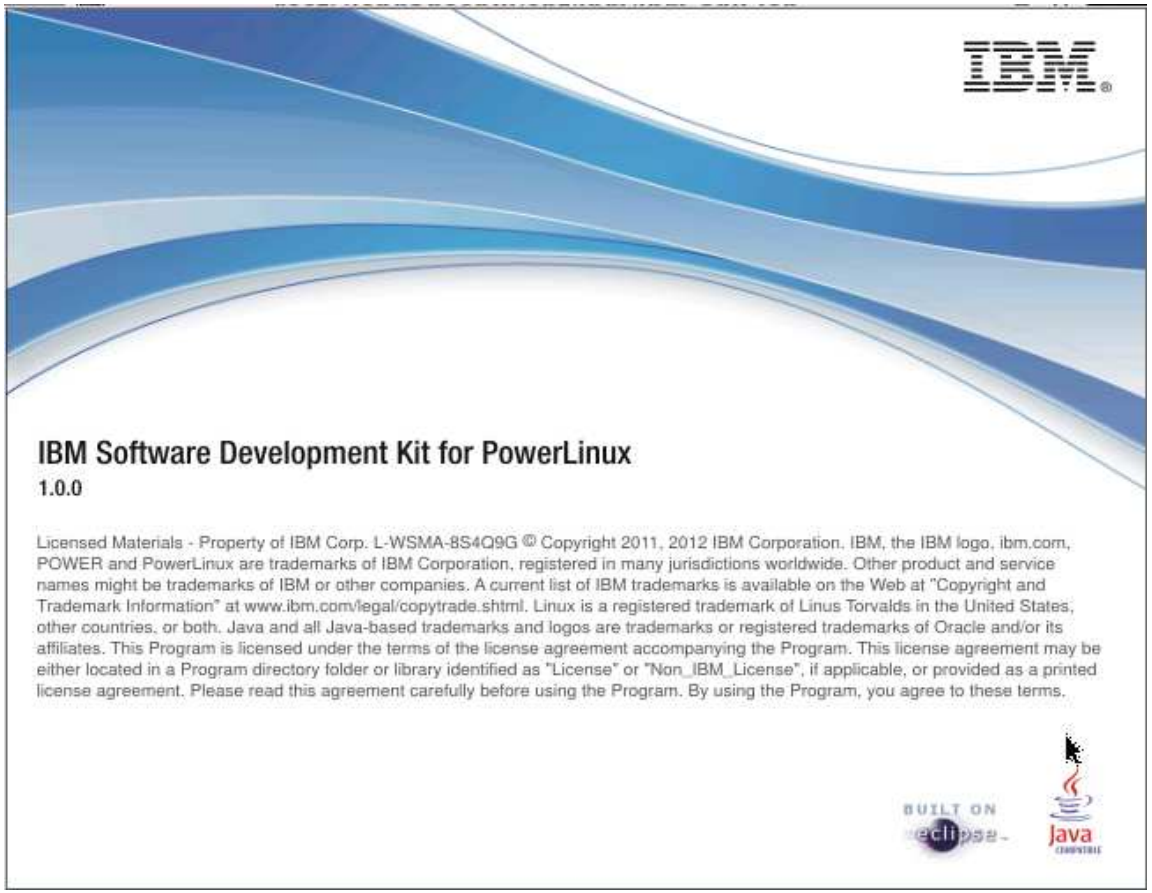

http://www.ibm.com/webapp/set2/sas/f/lopdiags/sdklop.html

How to use

- Available from following URL;

  http://www-304.ibm.com/webapp/set2/sas/f/lopdiags/sdklop.html

- Easy to install (same as AT)
  - Recent versions are available through online repository, such as yum and zypper.
  - Prior to install SDK, install locally only 1 package
    - For example on RHEL

```
# rpm -ivh ibm-power-repo-1.1.6-5.ppc.rpm
```

  - After that, Just execute online installation command
    - For example on RHEL

```
# yum install ibm-sdk-lop
```

SDK sample screen shots

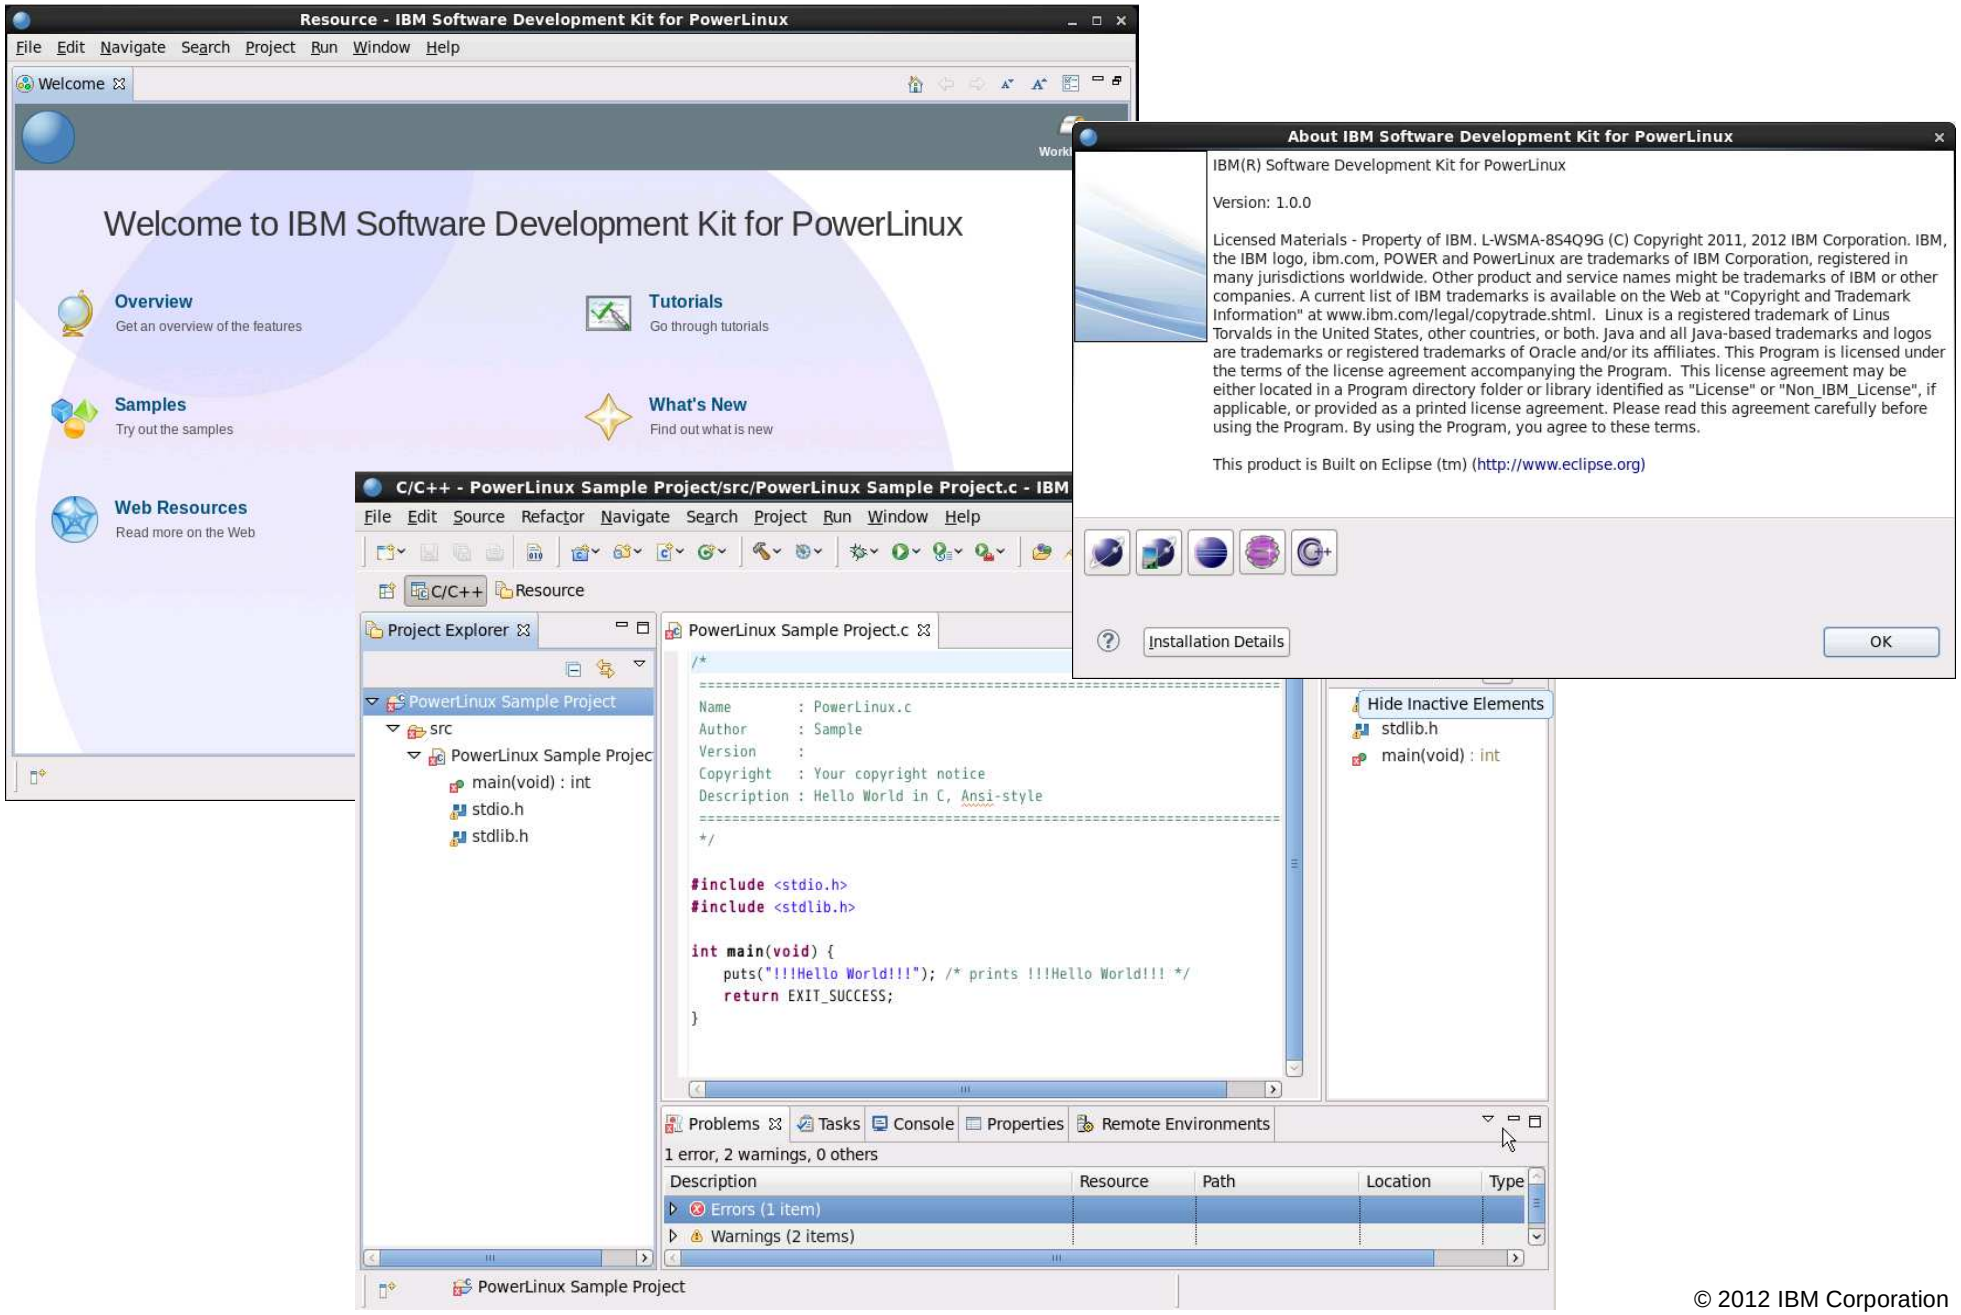

SDK for PowerLinux extends value-add plugins

- C/C++ project support of IBM Advance Toolchain
  - Power optimization wizard
- Linux Tools OProfile plugin
  - Launch and analysis integrated with code development
  - Configurable for HW specific event profiling
  - POWER6/7 PMU events
- Linux Tools Valgrind plugin
  - Launch and analysis integrated with code development
  - Open framework for dynamic analysis
    - Memcheck, detects memory leaks and malloc/free errors
    - Cachegrind, cache and branch miss analysis
    - Helgrind, thread and data race analysis
  - PowerISA features for POWER6/7
- Linux Tools RPM plugin
  - Build RPM packages from source code

SDK for PowerLinux includes additional Power-unique features

- FDPR (Feedback Directed Program Restructuring)
  - Integrated with Eclipse/CDT for ease of use
  - Works on both executable programs and shared libraries
  - Provides post-link global code optimization step
  - Tunes program to a representative workload

- Source Code Advisor
  - Leverages the FDPR dynamic inter procedural analysis capabilities
  - Provides interactive feedback to the developer
  - Identifies hot spots in source code that need rework
  - Specific suggestions for
    - Source code structure improvements
    - Compiler/linker options to use

- Code Migration Assist plugin
  - Integrated with Eclipse context sensitive source tooling
  - Scan/Analyze application source for common migration issues
    - Data Endian dependent unions and structures
    - Cast with potential endian issues
    - Non-portable data types
    - Non-portable inline assembler code
    - Non-portable or arch dependent compiler builtins
    - Proprietary/Arch specific APIs
    - Performance degradation

Sample Screenshot : Profile application performance with ease

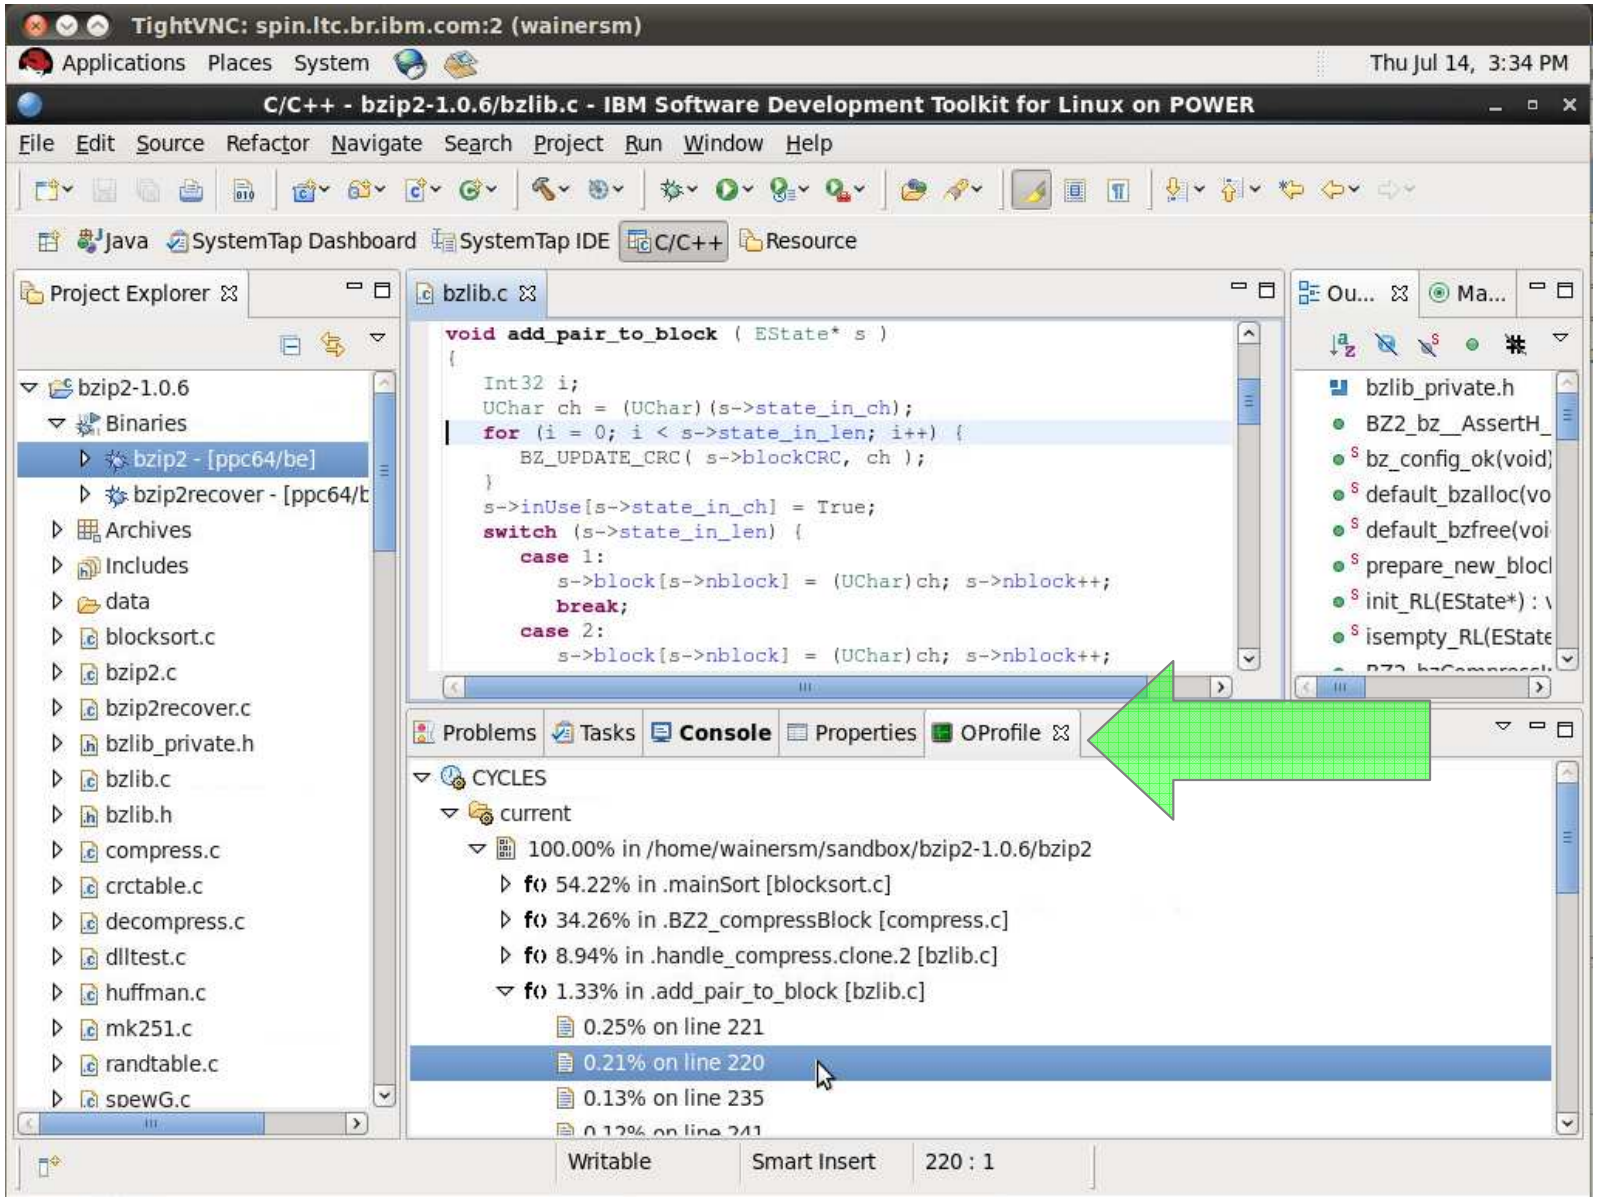

Sample Screenshot : Source Code Advisor

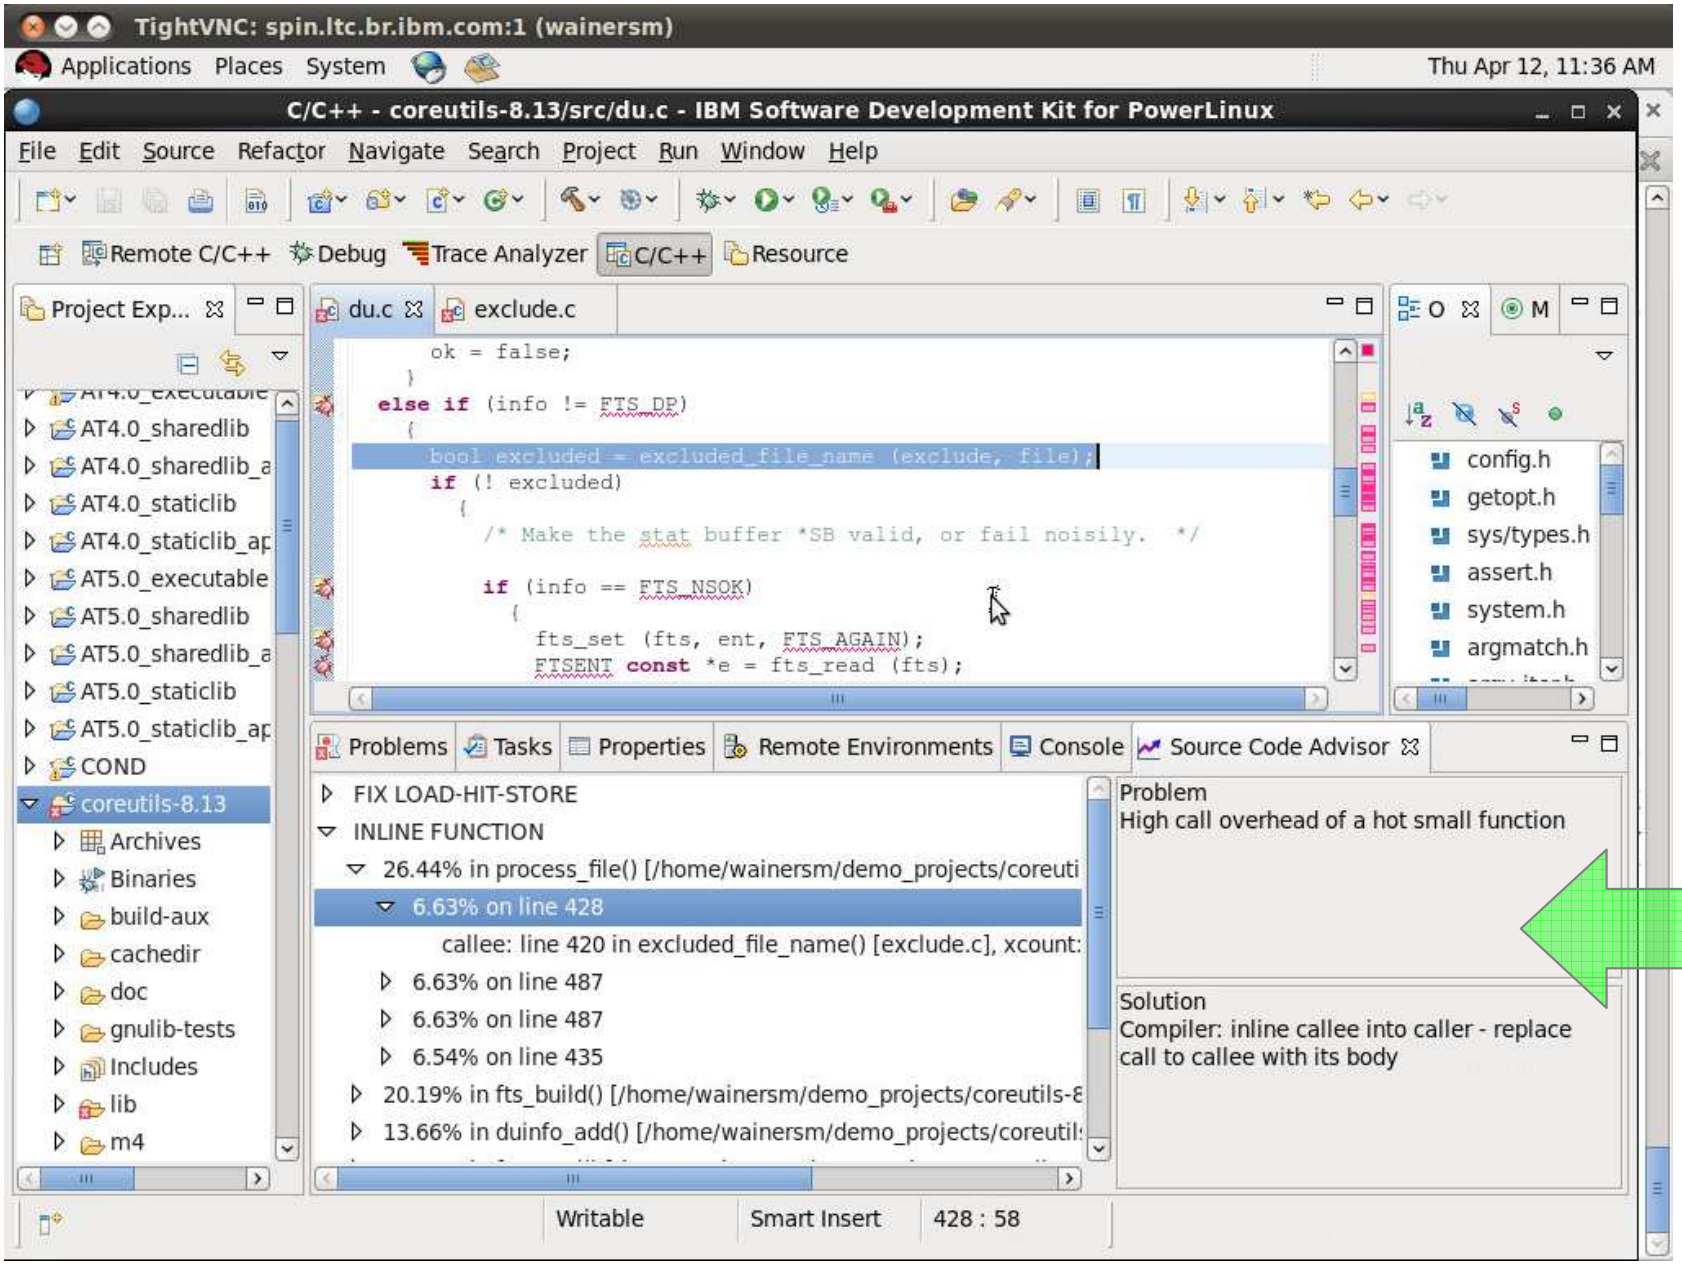

Sample Screenshot : Analyze thread using the Trace Analyzer

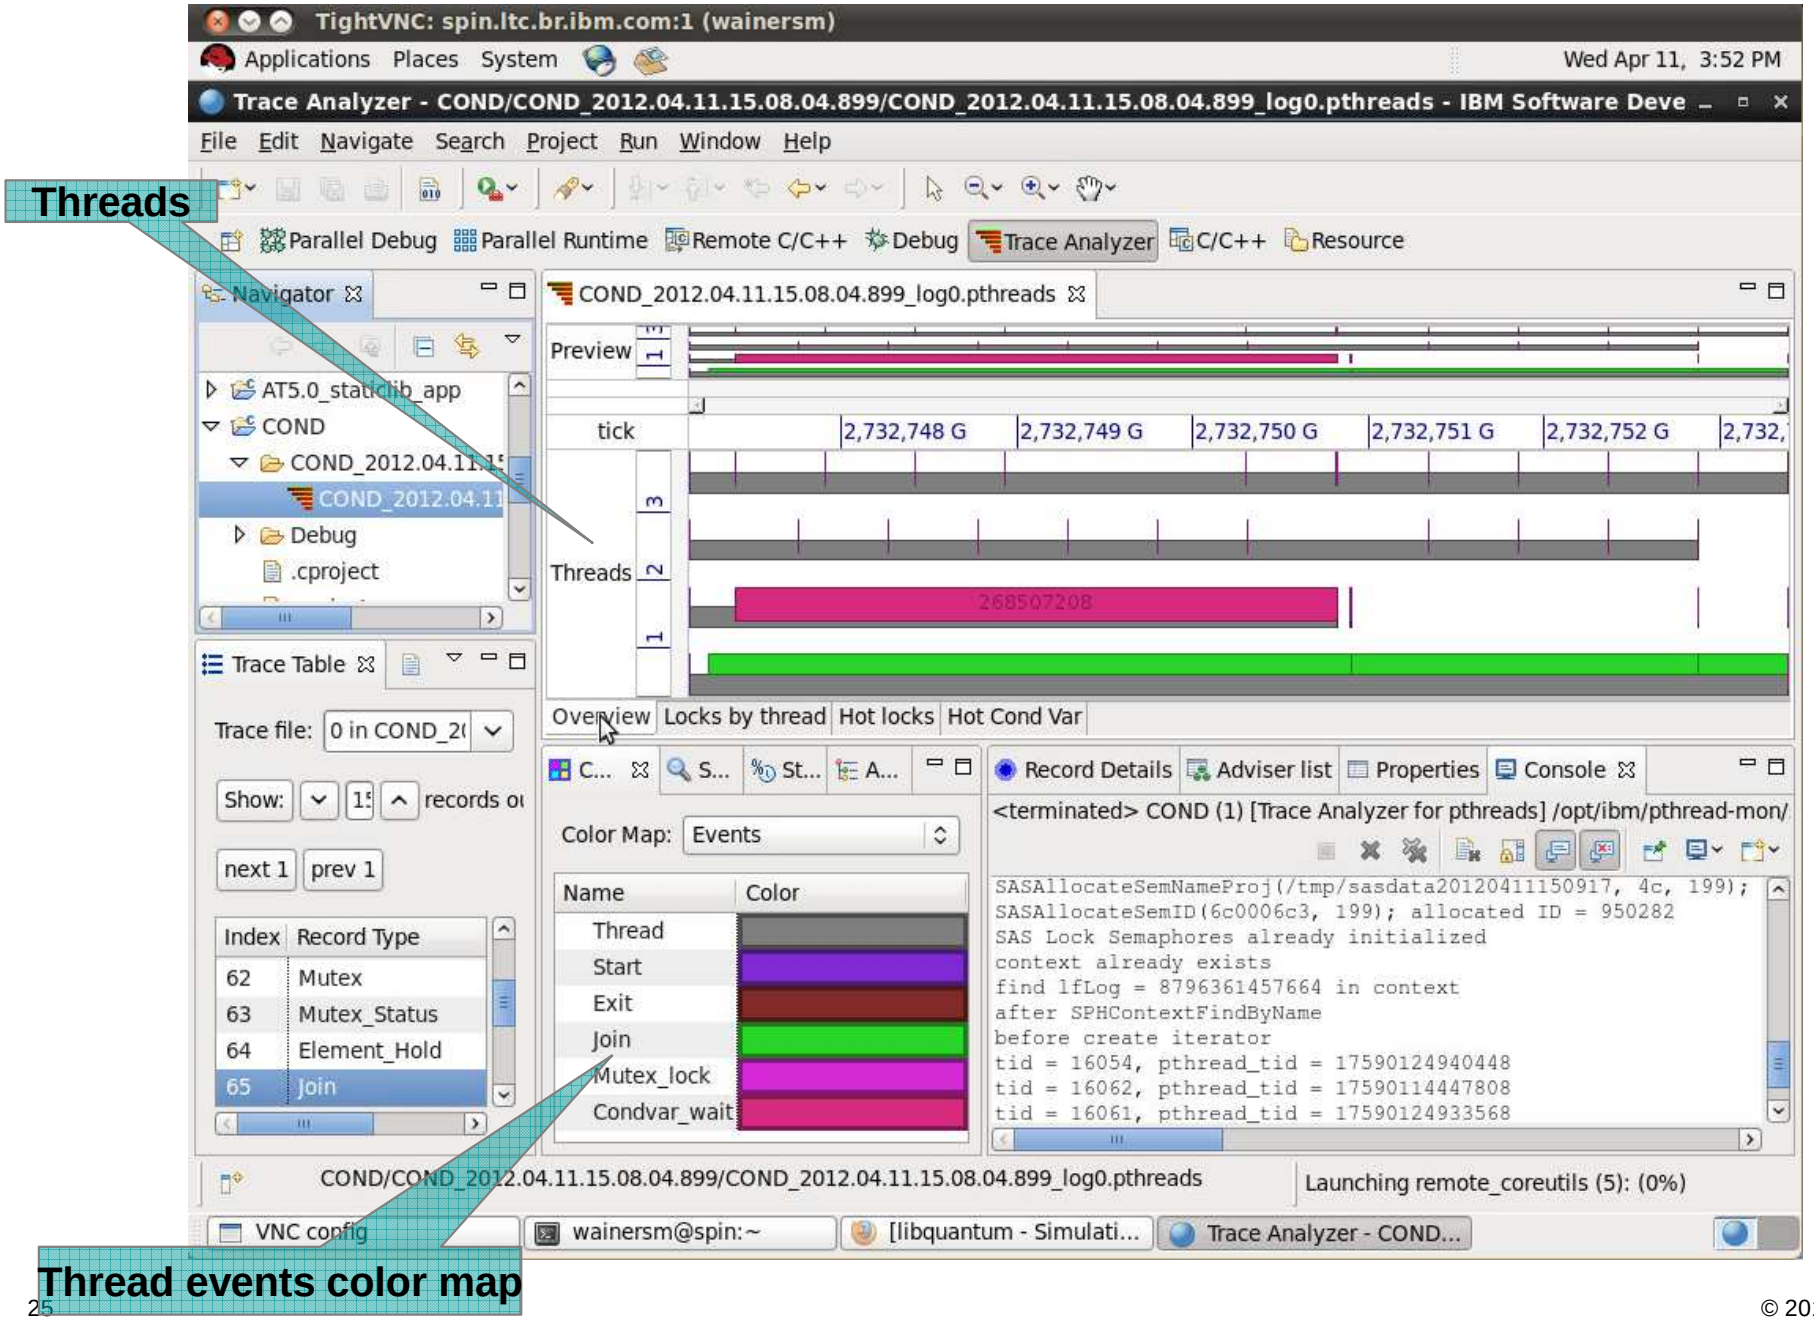

Sample Screenshot : Execute Migration Advisor (1/2)

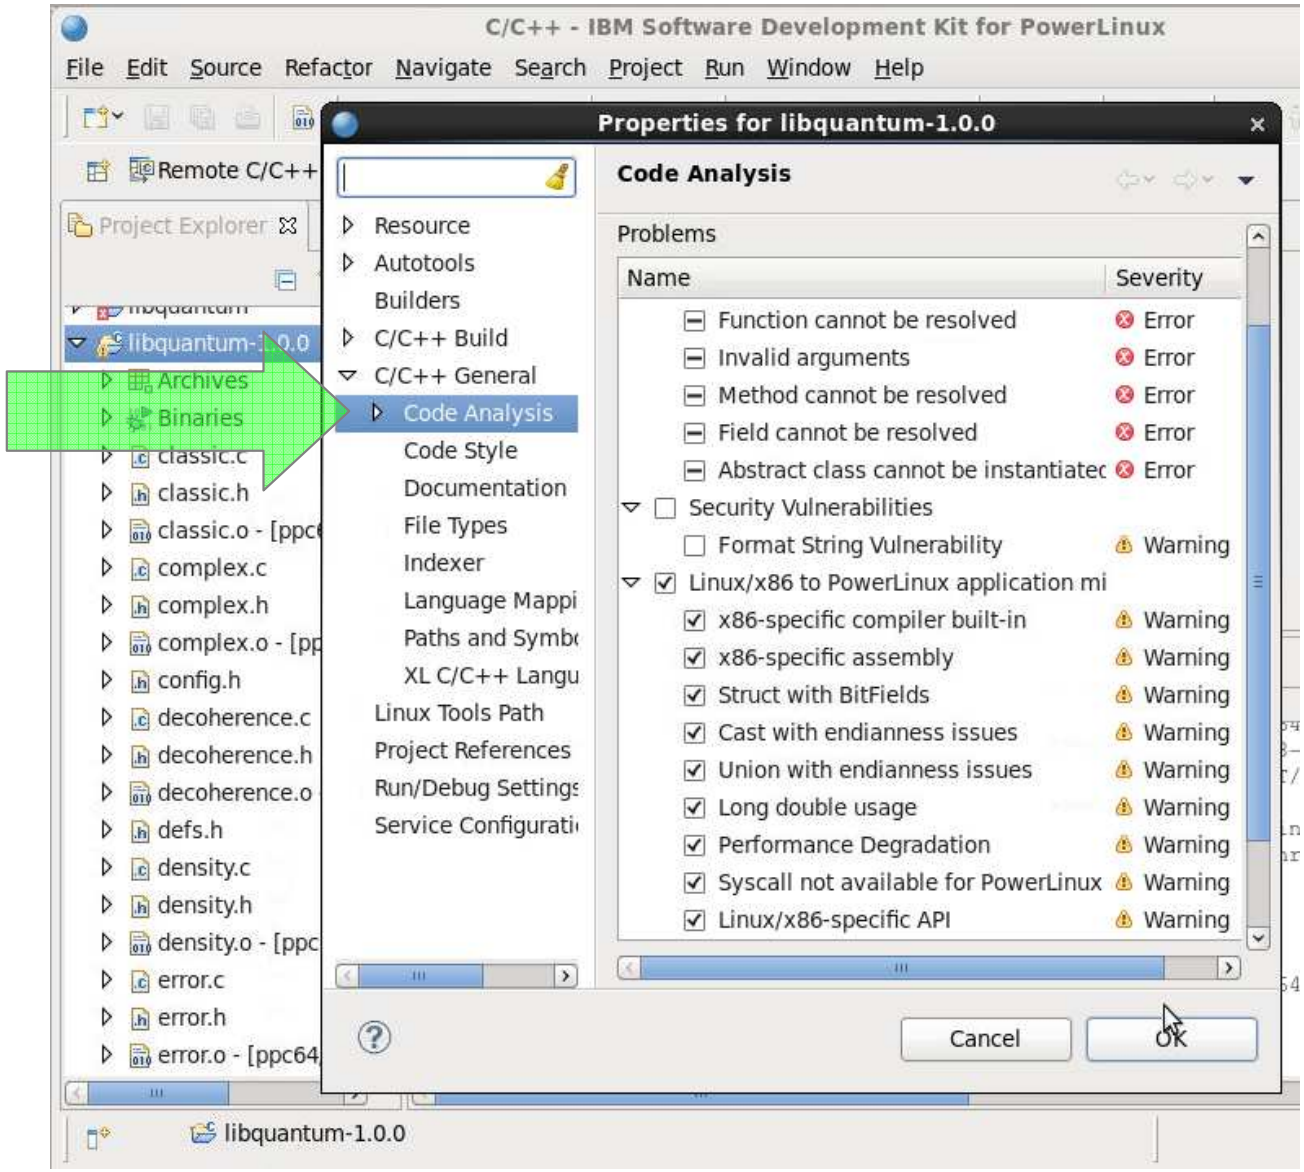

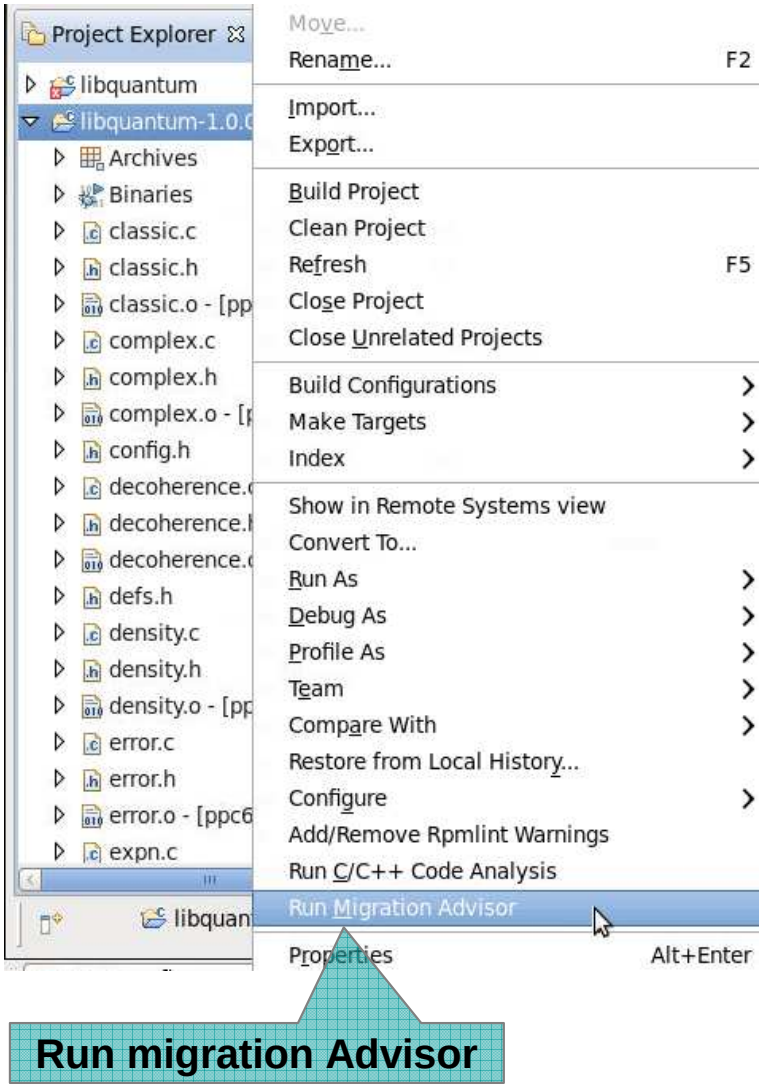

Sample Screenshot : Execute Migration Advisor (2/2)

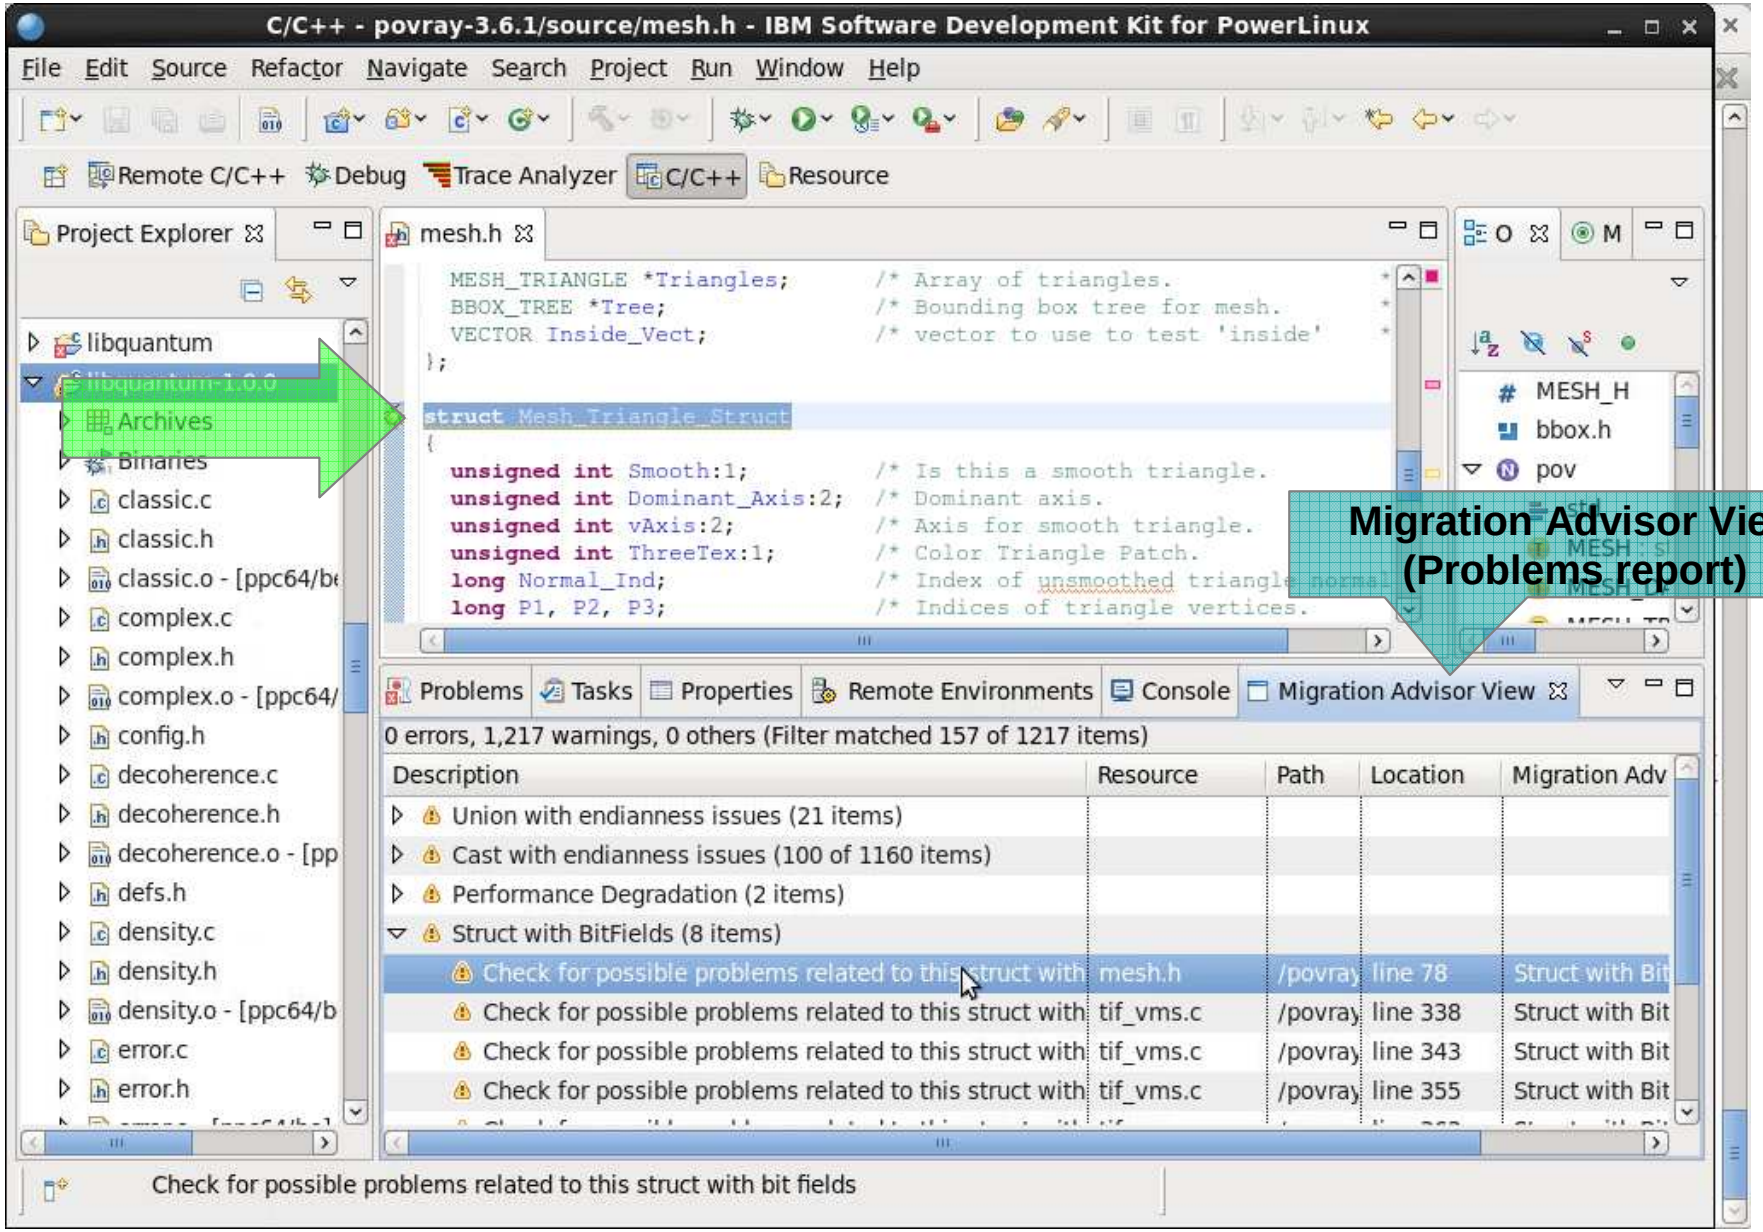

Conclusion

- Linux and Open Source Software becomes
  - to be adopted deeply inside of enterprises.
- POWER processor & Power Systems can provide
  - More scalability and reliability for Linux
  - Another hardware choice for Enterprise Linux Server Market
- PPC64 Linux can use Free Application development tools
  - Advanced Toolchain
    - Easy to improve application running performance
  - SDK for PowerLinux
    - Help your system analysis and C/C++ application development

If you are interested in PowerLinux, please feel free to contact to IBM.

Special notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

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

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

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

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. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

Revised September 26, 2006

Special notices (cont.)