2012 中国 Linux 内核开发者大会邀请函

一年一度的中国 Linux 内核开发者技术交流盛会，将于 10 月 13 日（周六）在北京拉开帷幕。本届大会由 AKA 阿卡信息技术小组、Intel 中国开源技术中心、南京富士通南大软件技术有限公司、清华大学 共同主办，并得到了 Intel 中国开源技术中心、IBM、富士通南大、华为、淘宝 的大力赞助，以及 Linux 基金会和中国开源软件推进联盟的指导。

作为国内信息技术和开源运动的先锋组织，AKA 一直致力于增进 Linux 开发者之间的交流和协作，致力于推动 Linux 在国内的普及和发展。本次大会继承历届会议传统，聚焦、关注技术性话题，增加交流讨论机会，旨在促进 Linux 内核社区的技术交流，诚邀广大 Linux 内核开发者和学习爱好者参加本次大会。

中国 Linux 内核开发者大会，是 Linux 内核方面的纯技术研讨会。2006 年至 2011 年，在 Intel 开源技术中心、富士通南大、IBM、华为、Canonical、淘宝、ADI、趋势科技等企业的大力支持下，已成功举办了六届 Linux 内核开发者大会，为社区的开发者提供了良好的技术交流平台。

会议时间：2012 年 10 月 13 日(周六)上午 9 点 （8 点开始签到）

会议地点：北京 中国科学院文献情报中心（北京市海淀区 北四环西路 33 号）

报名方式：大会官方网站 http://www.ckernel.org（网站内容尚在紧急制作中，但报名系统可用）

- ✧ **请完整填写资料，出现“报名成功”的对话框即表示报名完成。**
- ✧ **审核通过后，组委会将向您发送确认函。**
- ✧ **收到确认函（邮件形式）的人员，方可参加本次大会。**

官方微博：http://e.weibo.com/ckernel （中国 Linux 内核开发者）

会议技术主题：（请参见 P2 大会详细议程）

联系邮箱：guojinfen@gmail.com 参会咨询：郭锦盼（北京） 15811276440

2012 中国 Linux 内核开发者大会组委会


报告厅（全天会场）


多功能厅（下午会场）


媒体接待室

【大会议程】

2012年CLK大会(10/13)详细议程

时间	议题设置	
8:00 ～ 9:00	签到	
9:00 ～ 11:50	**CLK2011 Keynote** 会场：文献情报中心 报告厅 主持人：郭锦盼	
	大会欢迎辞、嘉宾介绍	郭锦盼
	领导致辞1	邱善勤主席（中国开源软件推进联盟主席）
	领导致辞2	倪光南院士
	Keynote1： Open Source - More than just code	Hans Peter Anvin (HPA), Intel
	Keynote2： Btrfs support for large pages	Mingming Cao, IBM
	Keynote3： State of the Linux Kernel Security Subsystem	James Morris, Oracle
13:30 ～ 16:50	**Session 1** 会场：文献情报中心 报告厅 主持人：郭锦盼	**Session 2** 会场：文献情报中心 多功能厅 主持人：吴峰光
	On The Way to a Healthy Btrfs Towards Enterprise 缪勰, Fujitsu	OpenSource, Linux and Standards Daniel Veillard, Red Hat
	基于ACPI的系统设备动态配置框架和实现 刘奖, HUAWEI	Libvirt Towards Enterprise Use 胡涛, Fujitsu
	KVM MMU implementation 肖光荣, IBM	Linux内核在互联网行业的优化 马涛, 淘宝
	ACPI 5.0 introduction 张锐, Intel	**ext4 snapshot机制的介绍与发展** 杨勇强，百度

【演讲嘉宾介绍】

Intel) Hans Peter Anvin(HPA)

演讲主题：Open Source – More than just code

主题介绍：In Open Source, it is often said that "code is king." However, becoming an active contributor to an Open Source project often involves much more than writing code and "throwing it over the wall." This talk will look at the relationship between contributors and maintainers, and the process of advancing from code contributor to becoming a subject expert to becoming an integral part of a large Open Source community like the one around the Linux kernel.

演讲人介绍：Hans Peter Anvin, known as H. Peter Anvin, or simply Peter Anvin, or even hpa, is a Swedish computer programmer who has distinguished himself by his contributions to Free and open source software projects. He is the originator of SYSLINUX,klibc, Linux Assigned Names and Numbers Authority (LANANA), various Linux kernel hacks such as: UNIX98 ptys，CPUID driver，The Linux kernel automounter，zisofs，RAID 6 support。

Anvin was previously maintainer of the linux.* Usenet newsgroup hierarchy and the Linux kernel archives at kernel.org, wrote the original Swap Space How-to, and the "Linux/I386 Boot Protocol" (file: linux/Documentation/i386/boot.txt)

Peter Anvin graduated in 1994 from Northwestern University, where he also was president of the Northwestern Amateur Radio Society (W9BGX); his amateur radio call sign is AD6QZ (formerly N9ITP).

In addition to his regular employment at Intel's Open Source Technology Center, Anvin is currently co-maintainer of the unified x86/x86-64 Linux kernel tree, chief maintainer of the Netwide Assembler (NASM). Previous employers include Transmeta, where he performed as architect and technical director; Orion Multisystems, working on CPU architecture and code morphing software; and rPath.

IBM) Mingming Cao

演讲主题： Btrfs support for large pages

主题介绍： Btrfs filesystem is a new filesystem that brings a lot of cool features for future. One of improvement done in btrfs is, it support variable blocksize, unlike other filesystem, which do not allow btrfs filesystem blocksize could be larger than pagesize. This will help btrfs improve large IO performance and reduce metadata fragmentation. But there is a issue with btrfs which by design it could not handle blocksize less than pagesize. This key issue prevents end users to freely move btrfs partition across different architectures with different page size support. This talk will go through the btrfs key structures to handling block and page cache mapping, how it support large size blocks, and the technique challenges we face to support smaller blocks on large pages size.

演讲人介绍： Mingming cao graduated from Oregon State University in 2000 and since then has been working in IBM Linux Technology Center. She has been working in the kernel team, involved many kernel scalablity work, in the area of IPC, block IO, filesystem. She co-started the ext4 filesystem and has contributed many. Right now she is focusing on improving btrfs performance and supporting btrfs for Power.

Oracle）James Morris

演讲主题： State of the Linux Kernel Security Subsystem

主题介绍： In this talk, we'll examine the current state of the Linux kernel security subsystem. Starting with a brief overview of existing features, we'll discuss recent developments, current efforts and future directions. We'll also discuss the evolving threat landscape, and the increasing need for mobile and cloud security. This will be a high-level technical discussion aimed at IT professionals. A good general knowledge of operating system and computer security concepts will be advantageous.

演讲人介绍： James Morris， James Morris is a the Linux kernel security subsystem maintainer. He is the author of sVirt (virtualization security), multi-category security (MCS), the kernel cryptographic API, and has contributed to the SELinux, Netfilter and IPsec projects. He works for Oracle in Sydney, Australia.

Fujitsu）缪勰

演讲主题： On The Way to a Healthy Btrfs Towards Enterprise

主题介绍：Btrfs has been on full development for about 5 years and it does make lots of progress on both feature and performance, but why does everybody keep tagging it with ""experimental""? And why do people still think of it as a vulnerable one for production use? As a goal of production use, we have been strengthening several features, making improvements on performance and keeping fixing bugs to make btrfs stable, for instance, ""snapshot aware defrag"", ""extent buffer cache"", ""rbtree lock contention"", etc. This talk will cover the above and will also show problems we are facing with, solutions we are seeking for and a blueprint we are planning to lay out. For this session, I'll focus on its features and performance, so for the target audience, it'd be better to have a basic knowledge base of filesystem.

演讲人介绍：Miao Xie is a linux kernel developer working for Fujitsu. He began to work in Linux field since 2007, and did some kernel test and bug fix (such as schedule, cgroup and so on) at the beginning. Now he devotes himself to Btrfs development.

HUAWEI）刘奖

演讲主题： 基于 ACPI 的系统设备动态配置框架和实现

主题介绍：计算机系统的 RAS 特性（可靠性、可用性和可服务性）是进入关键应用领域和电信领域的关键特性之一。Intel 高端 Xeon 服务器逐步实现了丰富的硬件可靠性特性，与高端 RISC 服务器之间的 RAS 特性差距逐渐缩小。但是 Linux 操作系统在 RAS 特性方面与传统 Unix 操作系统还有较为明显的差距，无法充分发挥 IntelXeon 服务器的硬件 RAS 能力。系统设备动态配置技术是可服务性的重要技术之一，本讲座重点介绍如下内容:

1）Intel Xeon 处理器与系统设备动态配置相关的 RAS 技术
2）Linux 系统设备动态配置历史
3）基于 ACPI 的系统设备动态配置框架
4）系统设备动态配置关键技术
5）基于 ACPI 的系统设备动态项目现状和展望

演讲人介绍：刘奖，于 2004 年从清华大学电子系毕业后加入 Intel，在 Intel SSG 工作了七年，一直致力于开源软件研发，在 Intel 平台的 RAS 特性研发方面拥有丰富的经验；并于 2011 年加入华为，负责 Linux 平台上的 RAS 技术研发。目前致力于 Linux 系统设备动态配置框架和实现方面的开发

IBM）肖光荣

演讲主题：KVM MMU implementation

主题介绍：Memory virtualization is one of the important components in virtualization and is the base infrastructure for the other components such as I/O virtualization, migration and so on. Memory virtualization is very complex, so we shall show you the way of memory virtualization approaches(Shadow page, Ept/Npt and nested npt) and some optimization techniques(unsync page, de-duplicated page, prefetch, thp...)patch set

演讲人介绍：Guangrong is a Linux Kernel Developer working for IBM's Linux Technology Center in Shanghai, China. He has worked on Ftrace, MM, BTRFS but his main interest is in KVM. He has introduced many features and performance optimization in KVM。

Intel）张锐

演讲主题：ACPI 5.0 introduction

主题介绍：ACPI 5.0 introduction ACPI is an abstraction layer between the OS and platform firmware and hardware. This abstraction allows the OS and the platform to evolve independently. Not only should a new OS be able to handle old hardware, but an old OS should be able to handle new hardware. This has been working well on IA platforms for quite a while. Now, a lot of changes has been introduced in ACPI 5.0. With these changes, we can take advantage of ACPI not only on traditional PC platforms, but also on phone/tablet platforms as well. In this session, I'll give an introduction about ACPI 5.0 and the status of our current developing work in Linux.

演讲人介绍：Zhang Rui joined Intel in 2006, after he got a master degree form Huazhong University of Science and Technology. He worked on Linux/ACPI subsystem developing and maintenance for 6 years, and he started working on Linux ACPI 5.0 enabling since 2011. Now he also works on Linux thermal management.

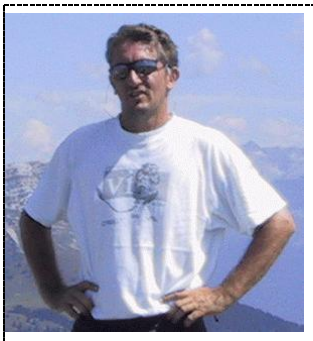

Red Hat）Daniel Veillard

演讲主题：OpenSource, Linux and Standards

主题介绍：This talk will focuse on the complex but important relationships that exists between various standards and related open source project the Linux kernel being a good example. It will show how both side can contribute efficiently toward a common goal of interoperability and large deployements. It will also expose some of the common problem that can bring such cooperation to a halt.

演讲人介绍：Daniel Veillard，Daniel Veillard has been working on Virtualization for Red Hat for the last 7 years, mostly on the libvirt project which he created in 2005. He is also the main author and maintainer for the libraires libxml2 and libxslt, and created the rpmfind.net software archives. He worked previously at the Web Consortium (W3C) on XML standardization.

Fujitsu）胡涛

演讲主题：Libvirt Towards Enterprise Use

主题介绍：kvm/qemu is maturing to be an enterprise virtualizaton solution. But as a low level virtualization implementation, it does not offer a user-friendly interface to manage a mass of virtual machines, which are common in enterprsie virtualization. Libvirt is the tool to fill the gap. This talk will introduce libvirt, and how libvirt can be used in enterprise virtualization.

演讲人介绍：Hu Tao has been working in Linux area since 2007. He jioned Fujitsu in 2010, working at virtualization technology for 2+ years. He's now in charge of Libvirt development and becomes the core developer in Libvirt community. He gave a talk about Libvirt last year at LinuxCon Japan, Yokohama.

淘宝）马涛

演讲主题：Linux 内核在互联网行业的优化

主题介绍：阿里集团核心系统部内核团队一直努力推进让 Linux 系统在阿里集团内部的服务器系统上更稳定高效的运行，同时还和开源社区紧密合作，基于阿里集团丰富的应用业务在 Linux 内核的稳定性、性能、可扩展性等多方面持续改进 Linux 内核。本次分享中我们将会对团队最近一年中在 Linux 内核中的文件系统，内存管理，块设备，cgroup，网络等方面的工作展示给大家。主要内容：

1. 文件系统方面， 我们在 ext4 方面做了大量的优化， 主要包括 bigalloc,inline data, dio overwrite nolock, io tree 等。
2. 内存管理方面，我们对 Page reclaim 的算法做了优先级的调整，从而能够动态控制 mmap 页和 read 页不同的置换策略。
3. 块设备方面,我们不仅调整了 io 调度器的某些属性,同时对 io throttle 等做了修正，从而能够更好的反应各个组的拥塞情况。另外我们还对 flashcache 做了修改，使他能够支持 cgroup io controller。
4. Cgroup 方面，我们优化了 cpu cgroup 的统计，从而实现了 instance aware 的 cpu 监控，同时做到了对 mem cgroup 的动态调整。
5. 网络方面，我们通过开源 skbtrace,从而能够做到对网络收发包的动态监测和分析，对评估和分析网络有很大帮助。

演讲人介绍：马涛， 阿里集团核心系统部内核组技术专家，国内最活跃的 Linux 文件系统开发人员之一，主要专注于文件系统和存储系统。曾在 2009 年和 2011 年 CLK 大会上做过演讲,分别介绍了 ocfs2 以及 ext4 这两个文件系统方面的研究和开发工作。

百度）杨勇强

演讲主题：ext4 snapshot 机制的介绍与发展

主题介绍 ext4 snapshot 由 ext3 snapshot（NEXT3）发展而来，它使 ext4 支持文件系统级的内建快照，不需要文件系统自身以外的额外存储空间，对 ext4 性能影响很小。目前，NEXT3 已用于 OpenNode。

大纲：1、Ext4 snapshot 发展；2、Ext4 snapshot 实现；3、Ext4 snapshot 的目前状态；4、Ext4 snapshot 正在进行的工作

演讲人介绍：我参与开发的 ext4 snapshotkernel 的代码在 https://github.com/YANGYongqiang/ext4-snapshots/tree/ext4-snapshots,还有一些是用户态工具的支持，代码在 https://github.com/YANGYongqiang/e2fsprogs-snapshots-patch-queue。我的一些工作申请了 Google Summer Of Code 2011 和 2012，2012 的工作正在进行中 J.下面是主题、摘要和大纲。

【会场地图】

【签到及餐饮】

1. 现场签到

签到时间：**10/13 8:00~9:00AM**

签到注意事项：

1） 如您已在网站上注册，请告知工作人员以便统计。（网上报名并审核通过人员将获得优先入场机会，敬请谅解。）

2） 在签到时，请尽量提供您的名片、工作单位、邮件地址、手机号码 等信息。

3） 签到时领取会议资料、参会胸牌和餐券，请您在签到后随身佩带好参会胸牌和餐券。胸牌将作为您进入会场、餐券将座位就餐的有效凭证，请务必妥善保管。

2. 餐饮

1） 10/13 中午，提供所有与会人员午餐，就餐餐厅当日通知。

2） 10/13 晚上，提供特邀嘉宾、主办/赞助单位、演讲人员餐叙，就餐餐厅当日通知。

【组委会友情提醒】

- **北京 10 月份会议较多，请一定要提前预定宾馆。**
- **请提前规划好地铁路线。**

2012 中国 Linux 内核开发者大会组委会