1 Common

1.1 Importance of Linux at Intel - Imad Sousou, Director of OTC at Intel

- Imad Sousou, Director of OTC, Intel pic.twitter.com/WFFKQmCH
- Intel heavy invests in W3C & APIs - for example NFC API
- Intel OSS - http://01.org

1.2 Forward-Looking Development: The Next Evolution in Enterprise Linux Software Development - Ralf Flaxa, VP of Engineering at SUSE

- Rafl Flaxa, VP of Enineering, Suse pic.twitter.com/33nbLd9u
- Ralf Flaxa: Talking as a Linux community dinasour - sharing his experience in Linux
- 21 years ago communications: E-Mail, FTP, newsgroup
- Rafl Flaxa: "We work for different companies but we have the same spirit - we are one big Linux family"
- Rebase with legacy support
- Ralf Flaxa: "Forward looking development is a new model for enterprise Linux"
- Research Into Open Hardware - Catarina Mota, Founder at openMaterials

1.3 Not all would be software. Catarina Mota speaking about Open Hardware on Electronics

- "The Global Village Construction Set" - using open hardware to build a city. Interesting concept presented by Mota.
- OSHWA - Open Source Hardware Association
- "Arduino is the main brain of open hardware" used by "backyard brains" to change the world @openmaterials Catarina Mota at
- OSHW: It's more compliated than software, it's not only distribute the 2d schema. . .
- OSHW: 3d Printers printing pieces for make 3d printers awesome!! They make himselfs

1.4 Linux: Where Are We Going - Linux Creator Linus Torvalds and Intel's Chief Linux and Open Source Technologist, Dirk Hohndel

- Somebody tries to improve the world through linux - but not I (c) Linux
- If company says that they want make billion of $ - I don't care (c) Linus
- Most interesting things - that something what I didn't expect (c) Linus
- "I do Linux for other people in the sense that I like the community," says Linus. Europe
- Linus Torvalds: "I don't know what you guys are interested in." Audience member shouts: "NVIDIA!"
- In linux kernel development we like argue.
- For the kernel the direction has always come from outside. Linus.
- Linus Torvalds: Im not worring about the future but mantain it open all this. Legal Situations, Patents.
- The patent system is just broken, says Linus @
- Linus: "Occasionally I worry about hardware companies..."
- Linus Torvalds: I mostly worry about kernel developers conflicts and the broken patent system
- Linux: We had huge issues initially with the embedded software companies
- Linus Torvalds: Daily builds and tests been making but we don't have all the hardware to test it.

- Linus Torvalds: The lockdown of the mobiles systems make difficult to work with this embedded systems kernels.
- Linus: "I'm so pleased with what has happened in the ARM community in the last two years"
- Many questions resolving around the needs of embedded systems and supposedly too high frequency of changes in kernel
- Linus: "I don't think there's a lot of innovation in OSes that's worth doing"
- " not a lot of innovations that need doing in the OS space" says Linus ;)
- Linus - "sometimes the old ways are the correct ways" re op systems at
- There is a high threshold for removing code from the kernel, says Linus. Europe
- Linus: "In general we do not remove code [from the kernel] unless it has major problems"
- android situation is better now, than a year ago (c) Linus
- Linus Torvalds: There are some vendors making wrong forks to work
- The Android "situation" is no different than the Red Hat or SUSE "situation" 10 years ago, says Linus. Europe
- Linus Torvalds: Google people are given the code from android tree but not all are been getting
- Google has one hundred million devices running their code. They must be doing something right. Linus.
- Linus Torvalds: we are all older but there is a lot of people from the university comming and working as professional
- Linus: "Actually performance vs. power is often not a versus"
- Linus at : "We are 99% male and 1% female in the community. Nobody knows why".
- We have no plan, but everybody is happy (c) Dirk
- You can read all the Q&A between Linus Torvalds and the Europe audience from today's keynote here: http://bit.ly/SB2uzj

2 System

2.1 FreeIPA, Jakub Hrozek / Alexander Bokovoy, Red Hat Open Source

- FreeIPA is an integrated security information management solution
- FreeIPA integrated with Microsoft Active Directory "out-of-the box"
- FreeIPA heavy rely on working DNS during install
- FreeIPA identification server should be heavy secured - its your Security SPOF
- Kerberos deployment is a fully management Kerberos realm
- allow AD user to connect to FreeIPA services, todo: allow FreeIPA users to interactively log in into AD machines
- missing now: CLDAP plugin to FreeIPA to respond on AD discovery quries, KDC backend to generate MS PAC
- FreeIPA provides smbd, not for sharing (you still can), but for LSA (Local Security Authority)
- FreeIPA provides nice Web UI configuration/management tool

2.2 New Challenges for Linux Network Support - Marcel Holtmann

- Enteprise wants better Ethernet, People wants better wireless
- on mobile side, we have small tools for too many tasks
- linux now good at ethernet, wifi personal/enterprise, GSM..LTE, DHCP, IPv4, IPv6, DNS, NTP
- including tethering, firewall setips, network switching, WISPr
- New Challenges: Am I actually online? (locally connected vs connected to Internet)
- ConnMan do "online checking" on global level
- Every application has different connectivity needs
- Waking all apps at the same time (when you becomes online) is a really bad idea
- Even more: mobile devices want different policies for each application
- Currently Linux haven't grained usage statistics (per app, per user, per interface).
- It will be nice, to be able say application to use internet, only on wifi connection, and ban mobile just for one app
- system should not only sync time, but also timezone
- Marcel Holtmann, Intel pic.twitter.com/5tYwxonT
- sometimes even per application routing tables are required
- what about Exclusive usage of VPN? one more thing, which will be nice to have
- Network connections are global, VPN, Bluetooth tethering - are global
- netfilter subsystem is too static, its hard to implement firewall policies depending on user/place/app
- Wifi offload - its hard to automate cases with switching connections to public network with security
- Bluetooth support pairing via NFC, Wifi support is coming
- Sensors Network: 802.15.4 and 6LoWPAN is a reality
- systemd and cgroups will be mandatory for a good mobile networking experience

2.3 Linux Kernel Report - Jon Corbet, Editor at http://LWN.net

- Jon Corbet, Editor, http://LWN.net pic.twitter.com/yBC3vYZO
- Linux Kernel: Volunteer participation going down, enterprise participation up
- Companies are very good for hiring kernel hackers...
- Very positive trend that embedded companies are contributing back to the Linux kernel ecosystem
- High level of kernel contribution from mobile and embedded players
- All what is happened in a world with network - happened in Linux
- Linus main role - is showing attention
- big.LITTLE - few slow, but power efficient and fast, but big power consumer - cpus on one SoC
- ext4 still the workhorse filesystem and getting new features
- btrfs continued stabilization plus send/receive, raid5/6 still waiting
- F2FS - Flash-friendly filesystem (high performance on NAND flash)
- CPU scheduler: NUMA scheduling, Power-aware scheduling
- Kernel problems tend to be: hardware/workload dependent =¿ hard to find automatically =¿ That is what users are for
- kernel regressions are most important bugs
- I'm not worried about... The health of Linux as a whole (c) Jon

2.4 Video4Linux2: Path to a Standardized Video Codec API - Kamil Debski, Samsung

- Kamil Debski, Samsung pic.twitter.com/cDB1uNHV
- Linux SoC codec libraries OpenMAX, VA API by Intel, VDPAU by Nvidia
- V4L2 supports TV tuners, webcams, video capture, output devices
- memory-to-memory (m2m) devices presents - enables support for video filters, codecs as filters
- V4L2 Codec API: m2m, multiplanar API, Videobuf2; extensions: h264, h264, mpeg4
- How to get started with a video coder driver: exisiting codecs in V4L2, flick through other m2m devices
- successes: API merged in mainline kernel, more hardware video uses V4L2, ARM Chromebook uses V4L2 codec API

2.5 Status of Linux Tracing - Elena Zannoni, Oracle

- The Tree of Tracing: ftrace, SystemTap, LTTng, perf, GDB, DTrace
- Uprobes: Dynamic Userspace Tracing
- Elena Zannoni, Oracle pic.twitter.com/sXFGx63x
- Uprobes described as inode, offset in file, list of associated actions, arch specific info
- Uprobes status: x86, x86_64(3,5/3.6), powerpc(3.7), ARM(wip)
- perf/ftrace/systemtap support for uprobes presents
- Ftrace have many plugins: Function*, Wakeup, Irqoff, Nop...
- /sys/kernel/debug/tracing/kprobes_events and /sys/kernel/debug/tracing/uprobes_events - for CLI control
- Perf list, annotate, lock, sched, kvm
- Ftrace-cmd - user space tool, many options, very flexible: record, report, start, stop, extract, list...
- LTTng - included in embedded Linux distributions
- SystemTap (Red Hat, IBM, others) - can use debuginfo, dynamic probes and tracepoints, has well defines rich scripting language...
- DTrace on Linux: solaris tool, porting with goal to support DTrace scripts for Solaris
- DTrace integrated with Oracle Unbreakable Enterprise Kernel
- DTrace provider, syscall prvider, SDT, Profile provider, Proc provider, x86_64 only, Kernel changes - GPL, Kernel modules (CDDL)
- Tracing: General Open Issues: KABI, Scalability, Code Injection, Embedded community

2.6 DRM/KMS, FB and V4L2: How to Select a Graphics and Video API - Laurent Pinchart, Ideas on Board

- Origins: fbdev, Linux 1,3,94 in 1996, Blanking in 2000, 4CC Formats in 2012
- Origins: DRM Linux 2.2.16 in 2000, GEM in 2008, KSM/TTM in 2009, Dumb Buffers on 2011, Planes, DMABUF in 2012
- V4L2 in Linux 2.4.0 in 2002, V4L2 subdev in 2008, Media contoller videobuf2 in 2011, DMABUF(?) in 2013
- there is no any reason to write fbdev drivers nowadays, DRM have full superset of features now

2.7 Optimizing File System Performance When Memory is Tight - Theodore Ts'o, Google

- ext4 new features: Punch system call; Metadata checksumming
- ext4 - "Modern" file system that still reasonable simple
- ext4 - Incremental devlopment instead of "rip and replace"
- ext4 well understood performance metrics
- ext4 disadvantages: fixed inode table, bitmap allocations, 32bit inode numbers
- ext4 RAID support is extremely weak, lack of sexy new features
- Ext4 - Default File System for Desktop/Server (distros can change it in near feature)
- Android devices (Hoeycomb / Ice Cream Sandwich) - use ext4
- ext4 "is common" in cloud storage servers
- The economics of cloud computing - run many tasks on single server - memory becomes a bottleneck
- Restricted memory means less caching available
- CPU can be a problem too, for PCIe attached flash / trasnscoding video
- ZFS Open Solaris list suggest 8Gb RAM jusr *for filesystem*
- Avoiding latency makes the users happy, a few slow requests slow requests behind them
- Google tests FS with focus on latency, not on avg performance
- Don't pay for features you don't need - no journal mode for ext4 in Google
- A lot of metadata caching have done in ext4
- ext4 sparce files improvements are coming; inline data; raid stripe awareness; atomic msync
- Conclusion: Remember to optimize the entire storage stack: thin-provisioned snapshots, dm-cache/bcache; optimize userspace

2.8 Rygel: Open Source DLNA, ready for Customer Products? - Jens Georg, Openismus GmbH

- .. smth good which comes from MeeGO Disaster
- UPnP(-AV) Technology & Devices and Services
- DLNA - vendor independent. set of media formats, on top UPnP-AV
- Rygel - UPnP Rendereer and AV-server
- DLNA picky on calling Rygel DLNA-compatible. Funny guys, my certified Hifi equipment fails quite often
- 9h Pieter Hollants @pfhllnts
- Rygel aims to: Be small, works out-of-the-box
- Nokia's N9 Smartphone - certified
- Rygel+DLNA have upload feature
- Rygel have Media server and renderer plugins
- SDK: librygel-server/renderer/render-gst

3 Resource Management

3.1 Resource Isolation: The Failure of Operating Systems & How We Can Fix It - Glauber Costa, Parallels

- Glauber Costa, Lead Software Developer, Parallels pic.twitter.com/fJtUqjmw
- history books tell us that back in the day, a computer ran a single program (c) Glauber
- VM may page out or give less CPU for your app, that you need
- "forking" app (in sum) gets more CPU than single-thread app
- Easy DOS: $ whole true; do mkdir x; cd x; done Will consume all memory before any disk quota can kick in
- hupervizors(KVM as example) fixes this resource isolation: 1. fair memory 2. N:1 CPU mapping
- But typical user wants run web, mail servers, databases - to many Virtual Machines for fair resource isolation
- but sometimes users wants run different versions of Linux for some apps. Solution - containers
- containers, network namespaces: 40 processes connected to port 80? No problem
- unique IP, raw devices for each group, per app packet filtering
- mount namespaces, user namespaces (more that one "root" user in the system)
- cgroups allows you logical grouping of processes. systemd is heavy uses it
- CRIU - Checkpoint/Restore complex processes mostly with userspace code, with the aid of some new infrastructure in the kernel
- cpu/memory controller can do not only fair resource isolation, but also isolation
- containers, today: openvz - stable, secure and mature Open Source, alternative - LXC, but its still not production ready
- Work Status: all works: cgroups, CPU, network; need improvements: user/mount/pid/blockIO/CR namespaces
- To be merged: fork bomb prevention, group-aware kernel memory, filesystem; specialized loop device
- containers tools: libvirt, vzctl, systemd, etc
- Note vzctl ¿= 4.0 works with both OpenVZ and upstream kernels RT containers tools: libvirt, vzctl, systemd, etc
- kmem controller is already in Andrew Morton's -mm tree

3.2 Systemd: The First Two Years - Lennart Poettering, Red Hat

- Lennart Poettering, Red Hat, pic.twitter.com/9BtOeFQ7
- systemd is a default in most distributions, included in others as optional
- systemd community quite a big 15/374 commiters/contributers
- systemd can now boot system shell-free
- systemd can boot system in less than 1 second
- systemd is a Init System; systemd is a Platform
- systemd: PID 1 does Unit Control Basic Set of Auxiliary Services do the rest
- systemd loves Embedded hackers, but Lennart personally won't hack embedded features (I asked Q: systemd busybox status)
- systemd intergrated with dracut, udev, D-Bus, Plymouth, Gummiboot
- systemd focus on movile, embedded, desktop, server

- when people brings patches to systemd to support uclibc, in most cases its better to implement such features in libc
- obsoltes: ConsoleKit, sysvint, initscripts, pm-utils, inted, acpid, syslog, watchdog, cgrulesd. Later on: cron, anacron, atd
- systemd control PM (power button, led) only util other program (Gnone/KDE) says: "No I control the buttons"
- systemd: Bring back Unix:Multi-Seat, Services as a File, Set of componenets that are integrated yet separate
- The Journal Structured, Indexed, Secure, Reliable, Modern, ...

3.3 Checkpoint and Restore: Are We There Yet? - Pavel Emelyanov, Parallels

- C/R helps with Live migration / reboot-less kernel update
- less popular scenario: start-up boost, working enviroment snapshots, HPC loadbalancing
- OpenVZ could do live migration since 2005 (implemented 100% in kernel)
- OpenVZ(Parallels?) failed to upstream their work - they working on it since 2008 to 2010
- Next idea - reimplement work in usespace, and improve kernel interfaces
- CRIU ultimate goal: any app -¿ dump -¿ restore
- Exiting kernel APIs: proc, system calls (about self, about anybody), netlink
- But sometimes kernel could not answer correctly. Solution is obvious but not easy - just add required interfaces.
- CRIU - a project by various mad Russians to perform c/r mainly from userspace (c) Linus Torvalds
- CRIU kernel features: 1. parasite code injection (works similar to trace api, which gdb for example use)
- CRIU kernel features: 2. kcmp - check which kernel objects are shared between processes
- CRIU kernel features: 3. sockets dumping via netlink - externdable sockets retrive engline
- Pavel Emelyanov, Parallels pic.twitter.com/IYqhBNm3
- CRIU kernel features: 4. TCP repair mode
- CRIU improved: virtual net devices indexes, proc map_files, socket peeking offset, more socket options
- CRIU features: x86_64, process tree linkage, multi-thread, mmap, terminals, groups, sessions, open files, lxc, ...
- CRIU tests every API separately, but also have tests some big applications (apache, mysql, nginx, mongodb, ...)
- CRIU plans: Full OS resources coverage, vanilla kernel support, crtools integration with LXC and OpenVZ, live migration, speedup
- You can easy find CRIU in internet http://criu.org/Main_Page
- CRIU shoild write a FAQ, and first answer should be about security

3.4 Lightweight Virtualization: LXC Best Practices - Christoph Mitasch, Thomas-Krenn.AG

- Christoph Mitasch, Thomas-Krenn AG pic.twitter.com/MCIamHin
- Hardware Virtualization (Full, Para), Software Virtualization(containers)
- cgroups - the hearth of all Linux container virtualization techs
- cgroups /sys have "posix" interface - create cgroup with mkdir, change and attach with echo
- cgroups could hide "fork" bomb, from the used
- Kernel Namespaces is a "hands" for containers

- LXC have nice(easy to understand) tools in userspace
- LXC networking: no entry, empty, veth(bridge), vlan, macvlan, phys
- LXC + CRIU support is coming soon
- LXC is unsecure - "Don't give container root to someone you don't trust"
- ha-cluster of LXC could be done, with DRDB with LCMC, and activate lxc heartbeat
- rumors says, that you should use vzctl now with upstream kernel support, instead lxc
- lxc - test kill -PWR 1 initiates a proper shutdown

3.5 Resource Management With Linux - Bruno Cornec, HP

- Bruno Cornec, HP pic.twitter.com/VN0lOKhC
- DEMO will fail during presentation, because every DEMO fails during presentations
- resource isolation on 160 threads, 4TB RAM, dozen of disks
- HP do a lot for more NUMA performance in Linux, first patches merged in 2.6.16
- cgroups - Dedicated subsystems to manage specific resources: CPU shares, memory, blkio; CPU sets, ns, freezer, c/r

4 Cloud

4.1 Mark Shuttleworth, Founder, Canonical pic.twitter.com/MTnjvvHu

- Mark speaking about importance of skale out
- Ubuntu is #1 Cloud OS
- "We should focus on operation part of the Cloud" (c) Mark
- Ubuntu decision to be on every cloud, regardless of business & philosophy obstacles, is smart.
- DevOps hell pic.twitter.com/KZ9vSdZu
- Juju - is a tool, which helps escape from devops hell
- juju charm includes puppet chef ... python, perl and much more
- dev@laptop test@cloud deploy@metal
- Device in your pocket - is truly personal computer (c) Mark
- the way he explains it it will introduce us to management hell .. which tool is manageing what hell
- Virtual Machine Management Costs becomes less than cost of managing Desktop PC
- Research.Design.Refine in order to lead, according to Mark Shuttleworth
- Ubuntu have thin clients for web and windows apps
- Final words from Shuttleworth: In very short term, a common version of Linux will be avail on any class of device.

4.2 Mostly Sunny: Why Evernote Runs Their Own Linux Servers Instead of "The Cloud" - Dave Engberg, CTO at Evernote

- Evernote have about 400 linux servers
- Dave Engberg, CTO, Evernote pic.twitter.com/T4oHKOZd
- Public Cloud is a nice to start business, but later (if success) you should have own servers, or private cloud
- Evernote use Amazon S3 for static, other staff based on their servers
- Evernote use SSD on metadata servers
- Evernote: 10 tb mysql. 10 tb lucene. 380 tb apache httpd.
- Evernote use WebDav (apache, mod_dav) as internal protocol for your data
- Amazon* is good enough for 90-95% sartups
- Evernote do calculations, going to private saves them about 75% of costs ($228k -¿ $60k for compute + metadata task)
- "Cleverness is the enemy of stability." -Dave Engberg
- Evernote's CTO making the case for your own private cloud ... it's cheaper .. wasn't that obvious already
- You should use cloud, unless you absolutely sure that you shouldn't (c) Dave, because of business risks

4.3 Linux: At the Forefront - Brian Stevens, CTO and VP Worldwide Engineering at Red Hat

- RedHat - 20 Years of Disruption
- The Red Hat Linux Model: fast upstream development, hardened into enterprise releases
- Is Open Source is a Business Model? - No, But i is the best *development* model (c) Brian
- Open Source defines architecture of IT (rails, mongodb, etc)
- Brian Stevens, CTO and VP Worldwide Engineering, Red Hat pic.twitter.com/sK4hgpOc
- Linux brings your application to ALL platforms
- The most successful opensource projects of the past and today created by users and developers, not just companies. - @addvin EU
- RHEL certifies 17 clouds (in additional to 3000+ "hardwares")
- with OpenShift you can develop in Java, Ruby, P*, Node.js. Use git, Eclipse, or .. Web IDE
- Linux Application Stacks: Security from SELINUX, QoS from CGROUPS, LXC
- 100,000+ Active applications runs on OpenShift
- Linux is redefined Storage and Data, think about Gluster, Hadoop, Mongo..
- 80% of data doesn't sit in a database. Volume of generated doubles every 3 years - Brian Stevens, Red hat cto
- KVM and Linux have a head start: oVswitch, OpenStack and Quantum, Emergning s/w network controllers

4.4 Open Source Cloud Platforms - Marten Mickos, CEO at Eucalyptus Systems

- 4 open source sitsers: OpenNebula, OpenStack, CloustStack, Eucalyptus
- Evolution of cloud types: Public -¿ Private -¿ Hybrid -¿ Mobile
- If Amazon is Starbucks, Eucalyptus is the espresso machine, says @martenmickos Europe
- people goes to private clouds from public clouds and data-center
- cloud customers look for: BIZ: agility, dependability; TECH: experimentation, participation; FIN: optionality&control
- Marten Mickos, CEO, Eucalyptus Systems pic.twitter.com/XchdEaKM
- open clouds saves 80% money, if compare with private clouds
- open cloud features: Innovation, Cross-breeding, Deployments, Contributions, Industry support
- open cloud: freedom of environment, scale, deployment

4.5 Scaling an Open Source Community: How we Grew the OpenStack Project - Monty Taylor, Manager of Automation and Deployment at HP

- OpenStack: Compute, Networkoing, Storage, Dashboard
- Monty Taylor, Manager of Automation and Deployment, HP, OpenStack pic.twitter.com/cSVIEzSF
- current clouds is very similar to mainframes
- AIX, Ultrix, HPUX, Solaris, Irix - they all created just for binding customers to their own hardware - this is why they dead now
- OpenStack is done by 132 Companies
- OpenStack - Open Source, Design, Development Community
- OpenStack "goverment model" is a meritocracy (very popular in OSS btw)
- OpenStack Code Review/Code Standards/Consistent technology/Automated Testing/Automated Everything/Constant Vigence

4.6 Introduction to oVirt Virtualization Management Platform - Itamar Heim, Red Hat

- oVirt - focused on KVM, management tool with Web UI
- oVirt have classic development model: oVirt =¿ Fedora =¿ RHEL
- oVirt have small footprint (170Mb Fedora) - to run hypervizor only
- High Availability, Live Migration, System Scheduler, Power Saver, Maintence Manager, Image Management, Monitoring&Reporting...
- oVirt could add hosts just in Web UI, will install all required package themself
- Itamar Heim, Red Hat pic.twitter.com/RDtGMw21
- oVirt supports various types of storage: NFC, iSCSI, Gluster, and many more
- oVirt supports not only VNC, but also Spice
- oVirt have nice UI for creating pools of VMs
- oVirt support for http://spice-space.org/ remote desktop protocol. Looks great!
- oVirt also provides "heavy" permission functionality
- oVirt heavy integrates glusterfs
- oVirt also provider user "portal", where user see what (and how) he can use in cluster

- oVirt also provides a lot of reports about resource usage
- oVirt have nice RESTful Web Service
- even more, oVirt have python SDK, very nice CLI with documentation and autocompletions
- oVirt Engine written on Java, PostgreSQL as database, and can be integrated with AD/LDAP
- oVirt use libvirt/VDSM on Host/Note, Local/Shared Storage
- oVirt could run Linux/Windows - have guest agent for them (single sign on plugin for example)
- oVirt Host Agent - VDSM use KSM - "share" memory between different VMs
- oVirt High Level Architecture - oVirtWiki http://ow.ly/1Pp7Zg
- its easy to write hook, which do smth between VDSM calls libvirt

4.7 The OpenStack Project and the OpenStack Foundation - Eileen Evans, HP

- OpenStack started by Rackspace and NASA in July 2010
- There is a big transition, between use OSS internally and contribute to OSS projects, and start OSS project
- OpenStack do magor release every 6 month, plans for next release are made at OpenStack summit right after release
- OpenStack 4 opens: Source, Design, Devlopment, Community
- OpenStack code is available under Apache 2.0 license, and this is forced for all contributions
- OpenStack Foundation owns the OpenStack trademark
- Eileen Evans, HP pic.twitter.com/TN04v1Cc

4.8 Cloud Storage Reloaded: Distributed Filesystems (CephFS and GlusterFS) - Udo Seidel

- Shared FS - Normal business for Linux: network NFS/CIFS; shared disk OCFS2, GFS2; parallel/distributed: Ceph/GlusterFS
- recent attention on distributed storage - cloud hype and big data
- Operations - important part of the cycle; technical challenge; geeks vs. Enterprise
- Object based storage - partition, file, unique id
- Ceph and GlusterFS - sufficient community presence, picked up by Enterprise Linux vendors
- Ceph and GlusterFS - block storage =¿ storage server (OSD); (POSIX) file system -¿ Meta data; HA -¿ replication and distribution
- Client Part - (POSIX) file system, Storage level
- Parts of ceph: Cluster monitor, md5 cluster, [ceph] , OSD Cluster
- GlusteFS Storage/Brick server ¡=¿ (Meta data I/O, Data I/O). NFS included
- ceph(now) and glusterfs (soon) could be integrated with QEMU, to store vm images directly
- challnge one: server vs storage
- challenge two: infrastructure
- challenge three: support (service provider/application support)