CERN openlab

Newsletter

Winter 2009-2010 - Issue 1

Yan!

'Yet another newsletter'

This thought might come to your mind when receiving this!

The openlab team decided to issue this newsletter twice a year to keep you updated on our activities, without overwhelming you with information.

Please do not hesitate to pass it on to your colleagues and to send us some feedback at openlab.newsletter@cern.ch: we would like it to develop according to your interests.

For more regular updates, feel free to visit our website news section: www.cern.ch/openlab-news

The entire openlab team wishes you a marvellous new year 2010!

Education Corner

A word from the CERN openlab Manager 1
openlab summer student programme invites 2010 applications! 2
Intel and openlab continue to deliver training to CERN's programmers 2

News in Brief

openlab tests PVSS Oracle Archiver 3
CINBAD team analyses 100GB of data per day 3
CERN openlab visits IDF 2009 San Francisco 3
13 Intel ISEF students visit CERN 4
The openlab presentations are available to you 4

Technical Articles

CERN openlab tests Oracle's Advanced Compression Option with Exadata 5
The Platform Competence Centre conducts multi-core scalability tests with CERN software 5
The Siemens openlab team lays emphasis on cyber security analysis for industrial control systems 6
The Intel openlab team evaluates the energy consumption and performance of Intel's Xeon 5500 DP servers 7
CINBAD keeps an eye on the CERN network 8
The Oracle openlab team makes major advances 10

Education Corner

A word from the CERN openlab Manager

The openlab is a unique undertaking at CERN. The only large-scale structure for developing industrial R&D partnerships, it comprises three dimensions: the Creation of Expertise, the Communications activities, and the Education dimension. The three are interlinked. Expertise is created through the evaluation of solutions as well as genuine research and development on IT technologies. It is then disseminated through multiple communications channels including this newsletter. The third dimension, the openlab education programme, complements the dissemination of knowledge, and constitutes an essential pillar of the openlab-III activities.

The openlab education programme provides active dissemination. This is where it differs from passive dissemination such as publication of reports or articles. This active dissemination is currently achieved through three lines of actions.

Firstly, workshops or seminars are regularly organized at CERN on advanced topics directly connected to openlab projects, such as Database Performance, Many-Core Programming or Optimized Programming of Modern Processors. Several of these workshops combine hands-off theory with hands-on practice. Secondly, openlab experts contribute to off-site education activities such as the CERN School of Computing. Thirdly, openlab runs a special summer student programme which in itself is a genuine education undertaking. Indeed, every year, a dozen young undergraduate students join CERN for 2 months. Their stay is financed by an innovative tripartite funding scheme (Industry, Universities and CERN). When at CERN, not only do they "learn by doing" but they attend a programme of lectures focusing on advanced IT technology, and combining fundamentals and CERN specific hot topics.

We will regularly report on recent achievements and plans on openlab education in this education corner of the newsletter.

François Flückiger
CERN openlab

ISSN 1663-845X

CERN openlab summer student programme invites 2010 applications!

The openlab student programme is open for applications from bachelor, master and PhD students in computer science and physics. Successful applicants will spend two months at CERN, during the period June to September 2010, to work with some of the latest hardware and software technologies.

The programme is more than just a summer at CERN: it can lead to follow-on projects at the home institute and may even inspire entrepreneurs in cutting-edge computing technologies. A series of lectures by experts in various domains of CERN related to high-throughput computing, and study tours to universities and CERN facilities are also part of the programme.

All nationalities are welcome to apply, as long as the home institute supports the application. Note that partial funding is available and normally the home institute is expected to co-fund the stay at CERN.

CERN openlab summer students 2009 and Sverre Jarp, CERN openlab CTO and coordinator of the student programme. *Photo by Andreas Hirstius.*

Please visit www.cern.ch/openlab-students for more information. Candidates should send a CV and letter of support from a supervisor to openlab.students@cern.ch.

The closing date is 31 March 2010.

Thanks for passing this information on to students, institutes and contacts that you may have in this area.

Mélissa Le Jeune
CERN openlab

Article published in CERN Computer Newsletter.

Intel and openlab continue to deliver training to CERN's programmers

Jeff Arnold, from Intel, teaching at one of the workshops. *Photo by Andrzej Nowak.*

CERN openlab has been teaming up with Intel to organize regular training for CERN's programmers. Even though four regular workshops relating to two major computing topics were established more than 2 years ago, it is not unusual for the sessions to be oversubscribed. Thanks for openlab's collaboration with Intel and Hewlett-Packard, each of the participants at the mentioned workshops has the comfort of working on a dedicated 8-core machine with a suite of tools, including Intel Software tools such as compilers and multi-threading optimizers.

One of the two workshop series, held twice a year for two days, aims to prepare attendees for the multi-core future, and relates to multi-threading technologies and parallelism. The other series, also held twice per year for two days, teaches programmers how to write efficient code for today's computing cores based on the Intel Architecture. CERN openlab lecturers Sverre Jarp and Andrzej Nowak, assisted by Jeff Arnold from Intel as well as a guest speaker, have recently led a session of this type for 35 participants. The topics covered ranged from computer architecture, through benchmarking, optimization and compilers to high level issues seen in specific programming languages.

In addition, Intel and openlab organized a separate two day workshop for CERN's key developers in mid-September. This special event was focused on getting the best performance out of applications and hardware by using tools such as the Intel Performance Tuning Utility. Another event of this type is foreseen for 2010. It will be targeting an expert audience as well.

Andrzej Nowak
CERN openlab

Article published in CERN Computer Newsletter.

News in Brief

openlab tests PVSS Oracle Archiver

PVSS is a supervisory control and data acquisition (SCADA) system used extensively at CERN to protect the experiment equipment. The ETM's PVSS related activities of the openlab started in February 2009 and have been focused on training while handling some development tasks.

Besides achievements on the Installation Tool, a tool which allows automated deployment of PVSS-based applications in the CERN environment, and porting the PVSS CVS plugin to SVN, a major contribution was already delivered on the Oracle Archiver code.

The purpose of the PVSS Oracle Archiver is to store/retrieve historical data and alarms. It is widely used at CERN and by other ETM clients. With the introduction of the version 3.8.1 of PVSS, some stability and performance issues have been identified. These issues have been solved in the context of the openlab in collaboration with ETM's developers in Eisenstadt. The fixes have been distributed through an official patch release for the benefit of the whole ETM's users' community.

Daniel Rodrigues
CERN openlab

CINBAD team analyses 100GB of data per day

CINBAD project is now focusing on providing enhancements for CERN Network Monitoring. These new improvements will facilitate day-to-day operations, the diagnosis of network problems and extend the understanding of the network evolution and design.

The CINBAD team is currently working on a visualisation model of this information and the promising prototype has already been presented. At the same time, the team has been collecting around 100GB of data per day from the CERN network and is analysing it, searching for different anomalies. Both statistical and pattern matching anomaly detection approaches by the CINBAD team led to the discovery of a number of misbehaviours, including Conficker worm infections, spammers and non-legitimate network scans.

In the next weeks we can expect a more complete toolkit for network operation and troubleshooting as well as a comprehensive report about anomaly detection techniques that were being investigated by the CINBAD team. More details about current achievements and activities are available in recent publications.

Milosz Hulboj, Ryszard Jurga
CERN openlab

For more information, please see page 8 and: www.cern.ch/openlab-cinbad

In September the CINBAD team presented a poster at the 12th International Symposium on Recent Advances in Intrusion Detection.

The team presented its system for handling sFlow data for anomaly detection as well as the tools that facilitate the network troubleshooting. The poster abstract is available in the conference proceedings.

http://www.rennes.supelec.fr/RAID2009/index.html

CERN openlab visits IDF 2009 San Francisco

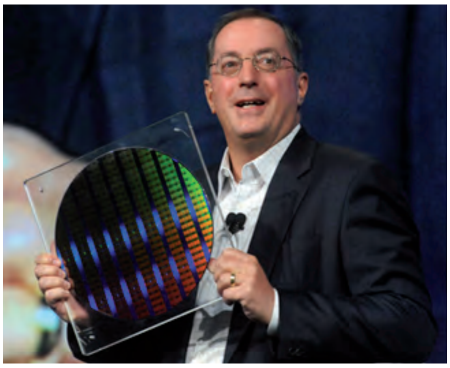

Intel President and CEO Paul Otellini displays a silicon wafer containing the world's first working chips built on 22nm manufacturing technology. *Photo courtesy of Intel.*

As in 2008, CERN openlab paid a visit to the annual Intel Developer Forum in San Francisco. A wide range of topics was presented, from news on Intel's silicon process to software technologies. Paul Otellini from Intel demonstrated a wafer in 22nm technology, along with interesting statistics.

In addition, Nehalem-EX and Larrabee hardware was demonstrated in test setups. Both are of significant interest to openlab, due to their many-core capabilities, Hyper Threading and other benefits.

Other interesting subjects included 32nm technology status, data centre efficiency, future solid state drive innovations and technologies accelerating application scaling to multi- and many-core systems.

Also of interest was the news of new Atom cores being on their way to consumers, along with updated development software.

openlab conducted an in-depth study of the Atom last year, and is looking forward to finding out more about the new members of this power-optimized family. Moreover, IDF 2009 was a good opportunity to hold valuable private meetings between Intel and openlab, which have enabled better understanding of the strategic and expansion plans of both sides.

Andrzej Nowak
CERN openlab

13 Intel ISEF students visit CERN

This summer, CERN hosted the visit of 13 pre-college students who won 'Best of Category' awards at the Intel International Science and Engineering Fair (ISEF) this spring in Reno, USA.

The young students spent four days at CERN and visited the Large Hadron Collider (LHC) facility, the world's largest particle accelerator and most complex machine ever built. They enjoyed presentations from various prominent scientists whose research is predicted to unearth evidence of new fundamental particles that will provide better insight in the fundamental laws of nature and the origins of our Universe.

The students were particularly impressed to go 100 m underground to visit the LHC tunnel, a 27 km ring of superconducting magnets chilled later to a temperature of just 1.9 degrees above absolute zero (-271.3°C), colder than outer space and actually the coolest place in the Universe.

More information is available on:
www.cern.ch/openlab-news

Mélissa Le Jeune
CERN openlab

Intel ISEF students visiting CMS cavern.
Photo by Michael Hoch.

The openlab presentations are available to you

Most of the presentations given by the openlab team members are available to you on the openlab website.

We encourage you to have a look to the presentations section of the website:
www.cern.ch/openlab-presentations

In this section, you will find most of the presentations given at the Minor Review Meetings (monthly thematic updates on the work in progress), at the Major Review Meetings which are held twice a year, in the presence of the openlab partners, and at the openlab Board of Sponsors.

Presentations given at conferences are also accessible. In particular last quarter, the openlab teams working in the Database Competence Centre gave various presentations at the Oracle OpenWorld in San Francisco in October and at the UK Oracle User Group conference in Birmingham, in November (cf. articles page 5 and page 10 of this newsletter).

Mélissa Le Jeune
CERN openlab

Aliases to the key sections of the openlab website

News / Events

Press Corner
Press Releases
Press Coverage
Multimedia Corner
Partners' Spotlights

About CERN openlab
What is it?
Participants
Guiding Principles

Documents
Annual Reports
Technical Documents
Presentations

Management

Breaking News

9th Dec – openlab is Hiring one Network Engineer / Developer in the field of wireless networking

30th Nov – Dr. Rajeev Hazra will be giving a computing seminar on Exascale Computing at CERN on Friday 4 December

What CERN openlab is

CERN openlab is a framework for evaluating and integrating cutting-edge IT technologies or services in partnership with industry, focusing on future versions of the WLCG.
Through close collaboration with leading industrial partners, CERN acquires early access to technology that is still years from general computing market. In return, CERN offers expertise and highly demanding computing environment for pushing new technologies to their limits and provides a neutral ground for carrying out advanced R&D with various partners.

This initiative has been organised into successive three-year phases. In openlab-I (2003–2005) the focus was on the development of an advanced prototype called opencluster, and openlab-II (2006–2008) addressed a range of domains from platforms, databases and interoperability to security and networking. We are now entering the third phase – openlab-III (2009–2011) – which will not only capitalize on and extend the successful work carried out in openlab-II, but it will also tackle new crucial areas.

Annual Report 2009
(PDF)
(4 MB)

CERN openlab – III

The framework of the third phase of openlab is similar to the previous ones, with the aim of hosting a few major projects with a particular focus on technologies and services relevant to CERN as the WLCG Tier-0 as well as other nodes in the Grid. The technical activities are organized in different domains of competence:

Automation and Controls Competence Centre
Database Competence Centre
Networking Competence Centre
Platform Competence Centre

Forthcoming events	
15 December 2009	openlab Minor Review Meeting
28 January 2010	openlab Major Review Meeting
22-23 April 2010	Annual Board of Sponsors Meeting

DISCLAIMER:
This Web page contains pointers to material related to the management of openlab in the the Information Technology Department of the European Organization for Nuclear Research (CERN). These materials are publicly available on the [illegible] exclusively for the benefit of the staff and users of CERN. Their use and distribution are regulated by the CERN copyright [illegible]

- www.cern.ch/openlab-about
- www.cern.ch/openlab-direction
- www.cern.ch/openlab-education
- www.cern.ch/openlab-events
- www.cern.ch/openlab-news
- www.cern.ch/openlab-newsletter
- www.cern.ch/openlab-people
- www.cern.ch/openlab-presentations
- www.cern.ch/openlab-projects
- www.cern.ch/openlab-press
- www.cern.ch/openlab-reports
- www.cern.ch/openlab-students

CERN openlab events

› 28 January 2010
Major Review Meeting

› 9-10 February 2010
Computer Architecture and Performance Tuning Workshop

› 16 February 2010
Minor Review Meeting

› 16 March 2010
Minor Review Meeting

› 13 April 2010
Minor review Meeting

› 22-23 April 2010
Board of Sponsors

› 4-5 May 2010
Multithreading and Parallelism Workshop

Technical Articles

CERN openlab tests Oracle's Advanced Compression Option with Exadata

The Physics Database Services section provides Oracle-based database services for the Physics community at CERN. Currently we host ≈150 Oracle 10g RAC nodes and ≈850 TB of raw disk space to offer database services to the LHC experiments in the context of the World LHC Computing Grid. The physics community requires high-end database services in particular for high availability, reliability and performance. The current implementation at CERN deploys RAC on ASM, Streams and Data Guard, following the best practices of the Oracle Maximum Availability Architecture.

The expected data growth is roughly ≈30 TB per year per experiment. The experiments need to have all data available at any time not only during the experiment lifetime (10-15 years), but also for some time afterwards, as the data analysis will continue. To meet this need we have to provide an efficient way of accessing and storing the few Petabytes of mostly read-only data. The answer to our challenge is the compression available in ORACLE 11G Release 2 on the Database machine.

We have tested the Advanced Compression Option with Oracle's Exadata 1. The half-rack was located in Reading, UK and accessed remotely from Geneva during two weeks. It consisted of four nodes and seven storage cells with 12 disks each. The tests focused mainly on OLTP and Hybrid Columnar Compression (EHCC) of large tables for various representative production and test applications used by the physics community, like PVSS, GRID monitoring and test data, file transfer (PANDA) and logging application for the ATLAS experiment. Tests on export datapump compression were also performed. The test results are rather impressive: we achieved 2-6X compression factors with OLTP compression and 10-70X compression factors with EHCC archive high. The EHCC can achieve up to ≈3X better compression than tar bzip2 compression of the same data exported uncompressed.

Oracle Compression offers a win-win solution, especially for OLTP compression as it shrinks the used storage volume while improving performance.

We were invited to present this to the Oracle OpenWorld 2009 conference and received excellent feedback. The following talks were given:

- "Research and Development OOW Workshop", Maria Girone
- "Database Machine Analyst Panel", Maria Girone
- "Oracle Advanced Compression: Stories from the Most Trusted Source? Customers!", Svetozar Kapusta

Svetozar Kapusta
CERN openlab

The team from the CERN IT Databases and Engineering Systems group, also gave a few talks related to the Oracle partnership. To find out more, please see page 10 and: www.cern.ch/openlab-presentations

The Platform Competence Centre conducts multi-core scalability tests with CERN software

It is quite common these days to hear news about the multi-core and many-core revolution. CERN openlab, always keeping pace with the latest developments, has been hard at work for quite a while in order to establish how the new multi-core architectures relate to High Energy Physics (HEP) software used today. Parallelization efforts, although not widespread at CERN, have started to bear fruit in recent months and are in a large part actively supported by openlab.

One such activity is carried out by Northeastern University researchers, Xin Dong, and his supervisor, Prof. Gene Cooperman. It relates to a complete multi-threaded conversion of a serial physics processing framework commonly used in HEP.

One of the prototypes resulting from this work has been passed to openlab for testing. Initial examinations have shown good scalability on 8-core Harpertown systems, prevalent in CERN's computing centre. Additional tests carried out in the summer have shown that the scalability on the 8-core Nehalem-EP platform is promising as well. Other runs, executed on a 24-core Dunnington system, provided by Intel, revealed some areas needing further optimization when moving to double digit core counts.

The openlab Platform Competence Centre team interfaced with the researchers from the US, as well as with local experts, in order to find ways to make the software more scalable. The work is still in progress, however significant advances and optimizations have already been made, and there are noticeable improvements introduced by the developing team based on suggestions from openlab.

In addition, in the light of the promise of 32-core Nehalem-EX systems being available soon, openlab is looking forward to expanding its operations to this new hardware, and to gaining an insight into the behaviour of High Energy Physics frameworks on a modern many-core architecture.

Andrzej Nowak
CERN openlab

The Siemens openlab team lays emphasis on cyber security analysis for industrial control systems

The growing usage of Ethernet and TCP/IP in industrial devices (replacing dedicated networks) has led to the necessity to reach a higher level of security against common threats on Ethernet cable. These threats can be deliberate (attackers), collateral (viruses and worms), or accidental (misconfigurations). Moreover the introduction of more and more IT functionalities into process control devices gives us more reasons to perform security analysis in order to find any possible weak points.

The collaboration between Siemens and CERN focuses on the robustness of automation devices (e.g. Programmable Logic Controllers) through a deep investigation of these devices' resistance against attacks. More specifically, the major aim of the project is the definition of a test bench and specific procedures which allow us to perform a security mapping of devices' architecture and to simulate common attacks originating from either the internal or the external network.

Once the security mapping is complete, it is necessary to generate a detailed vulnerability report. It specifies the security breaches that need to be analysed in order to possibly develop several practical and easy-to-apply solutions to fix those vulnerabilities.

Standards and guidelines can be used to help identify problems and reduce the vulnerabilities in a cyber security system. By knowing the problems and vulnerabilities, standards can be applied to cyber security systems in order to minimize the risk of intrusion. This is why at the beginning of our activities, we compared three cyber security standards: ISA-99 (and part of the ISA-95), NERC-CIP, IEC-62351.

During the analysis of these standards we have noticed lots of congruencies and some discrepancies in the specific approaches they suggest. At the end of this analysis, ISA-99 seems to be the most relevant standard, the only one able to face up to the wide heterogeneity of control systems (which is also relevant for CERN experiments)[1].

This also implies that ISA-99 approach is quite general and can only provide a

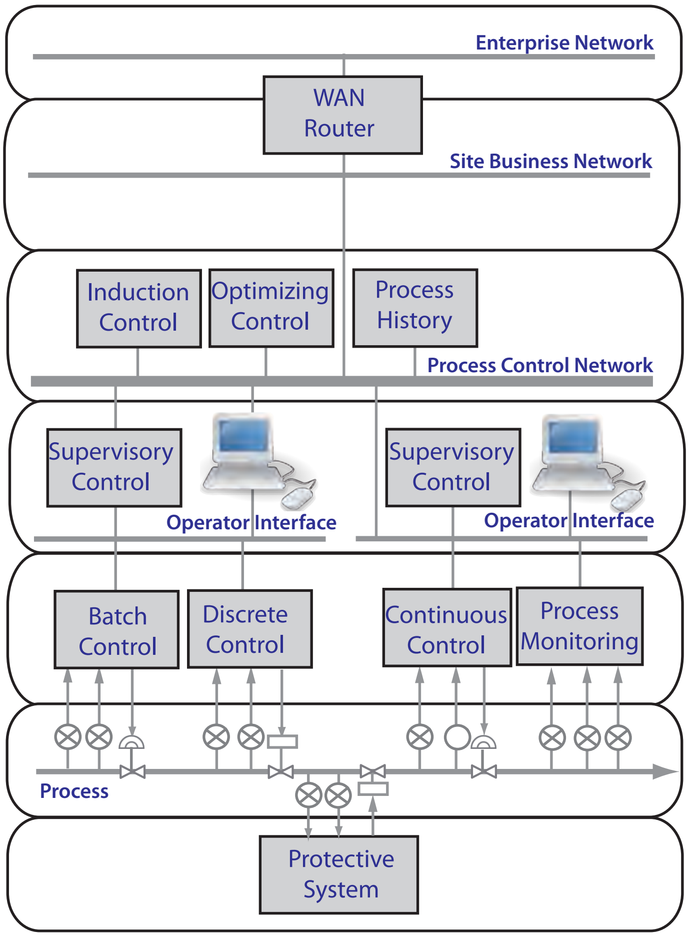

Level 5 - Enterprise
- Enterprise Financial Systems

Level 4 - Site Business Planning
- Site Production Scheduling
- Site Accounting

Level 3 - Site Manufacturing Operations
- Production Control
- Optimizing Control
- Process History
- Windows Domains

Level 2 - Area Supervisory Control
- Supervisory Controllers
- Primary Operators Interface

Level 1 - Basic Process Control
- Batch Controllers
- Continuous Controllers
- Discrete Controllers
- Process Monitoring

Level 0 - Field Instrumentation
- Sensors, Transmitters, Control Valves
- Field Networks (e.g. Foundation Fieldbus, Profibus)

Safety-Critical
- Protective Systems
- Safety Instrumented Systems

Distributed Control Systems (DCS) General Reference Model[2].

The Siemens openlab team lays emphasis on cyber security analysis for industrial control systems (continued)

theoretical (instead of practical) guidance and direction on how to establish and implement procedures (overall in the assessing phase, designing the security plan and defining the security policies). Moreover it is recognized that standards and guidance documents are living documents (the standard is not totally completed yet) that will continually evolve to meet the dynamic needs of industry and stay current with changing technology. The defense-in-depth model is sustained as customer's security scheme by the ISA-99 standard too, which recognizes that some attacks will inevitably penetrate the boundaries and thus requires further protections within the boundaries.

Programmable Logic Controllers (PLCs) represent the lowest level in the layers architecture of any control systems. As such, they are an essential link in any defense-in-depth strategy and must be considered as first class citizens in the chain of control[3].

Component Testing is finalized to assure that the specific component meets the required security specifications. To do this, we have defined some procedures and an entire test bench[4] which allow us to validate the confidentiality, integrity, and availability of every single process control device. In this context, one of the major problems is represented by the definition of the features and key cyber-security aspects (relevant to CERN) which must be tested, and of the minimum level of compliance which would allow us to identify whether a component is safe or not.

Unfortunately, at the moment, there are no standards able to provide any criteria or specific procedures which must be followed for procuring and implementing secure control systems. Almost all of the existing security software tools are finalized towards structural and architectural IT components and systems, but they are not designed specifically for industrial devices. For this reason, we have deployed and developed specific techniques and methodologies of attacks to evaluate the robustness of process control devices.

In the following phase, we are reporting all the discovered vulnerabilities that need to be fixed in order to improve the quality and security level of these control devices which are widely deployed at CERN.

Filippo Tilaro
CERN openlab

References
1 F. Tilaro, "Control system cyber-security standards, convergence and tools", CERN technical report, April 2009

2 Manufacturing and Control Systems Security dISA-99.00.01

3 B. Copy, F. Tilaro, "Standards based measurable security for embedded devices", CERN, 2009

4 F. Tilaro, "Testbench for Robustness of Industrial Equipments (TROIE)", CERN, 2009

The team evaluates the energy consumption and performance of Intel's Xeon 5500 DP servers

As soon as the new generation Intel Xeon processor 5500 ("Nehalem") servers came into production, Intel offered CERN openlab the opportunity to evaluate this new microarchitecture. We chose to evaluate three different flavours of the processor - L5520, E5540 and X5570 – since they provide different levels of performance and power consumption.

Their efficiency was evaluated by measuring the typical power consumption, using standard benchmarks to stress the different subsystems in the server. We also evaluated the performance of the processors with the C++ subset of the SPEC2006 benchmarks.

According to our measurements, the L5520 is the most efficient overall, being 36% better than the previous Xeon 5400 ("Harpertown") servers, the other 5500 flavours reaching 30%.

Improved efficiency was not the only positive point, since the Nehalem introduces Turbo mode and reintegrates SMT (Simultaneous MultiThreading). In the new Xeon generation, SMT allows each processor to execute two threads simultaneously by sharing the execution pipelines. SMT was thoroughly evaluated, because it may offer a further advantage for a computer centre: increasing the throughput of processed jobs by 15 to 21% based on our tests.

This evaluation involved multiprocessing (using a Monte Carlo based benchmark, "test40") and a multithreaded benchmark ("tbb") based on the ALICE High Level Trigger and the Intel Threading Building Blocks as well as a real-world complex framework (from ALICE) and compared the efficiency of different global scheduling policies.

Julien Leduc
CERN openlab

CERN openlab

Evaluation of energy consumption
and performance of
Intel's Nehalem architecture

Axel Busch, Technical Student
Julien Leduc, CERN openlab Fellow

October 23, 2009

The report compiling these results, 'Evaluation of Energy Consumption and Performance of Intel's Nehalem Architecture', by Axel Busch and Julien Leduc, has been presented to the international press on 4th November 2009, by Sverre Jarp, CERN openlab CTO, during the event hosted by Intel at CERN's Globe of Science and Innovation. Related press coverage and the report are available at:
www.cern.ch/openlab-press
www.cern.ch/openlab-reports

CINBAD keeps an eye on the CERN network

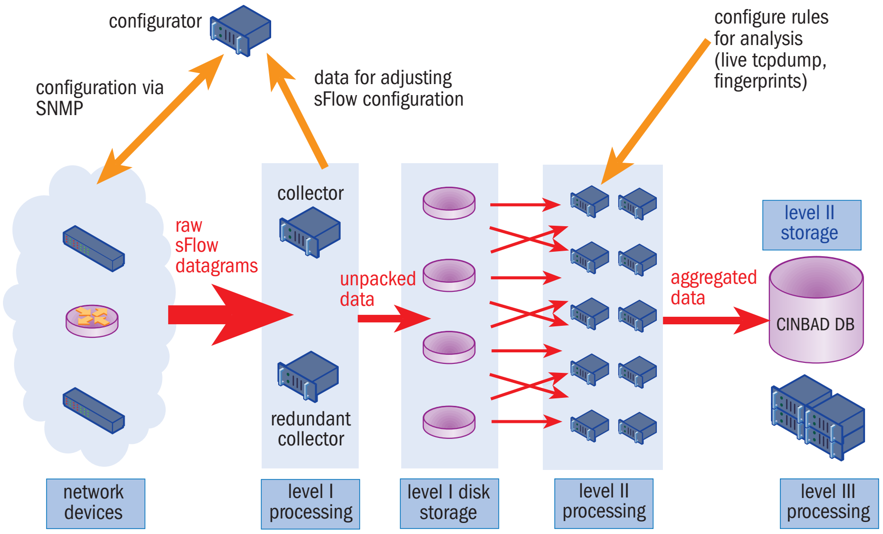

The CINBAD sFlow data collector receives and processes the CERN network traffic.
Fig. 1. Courtesy of IOP Publishing.

The CINBAD (CERN Investigation of Network Behaviour and Anomaly Detection) project was launched in 2007 as a collaboration between CERN openlab, IT-CS and HP ProCurve Networking. The project's aim is to understand the behaviour of large computer networks in the context of high-performance computing and campus installations such as those at CERN. The goals are to detect traffic anomalies in such systems, perform trend analysis, automatically take counter measures and provide post-mortem analysis facilities.

CERN's network

CERN's campus network has more than 50,000 active user devices interconnected by 10,000 km of cables and fibres, with more than 2500 switches and routers. The potential 4.8 Tbps throughput within the network core and 140 Gbps connectivity to external networks offers countless possibilities to different network applications. The bandwidth of modern networks is growing much faster than the performance of the latest processors. This fact combined with the CERN specific configuration and topology makes network behaviour analysis a very challenging and daunting task.

CINBAD in a nutshell

The CINBAD project addresses many aspects associated with the CERN network. First, it provides facilities for a better understanding and improved maintenance of the CERN network infrastructure. This includes analysing various network statistics and trends, traffic flows and protocol distributions. Other factors that might have an impact on the current network status or influence its evolution are also studied, such as connectivity, bottleneck and performance issues.

When we have learnt and understood the network behaviour, CINBAD can help to identify various abnormalities and determine their causes. Because there are many factors that can be used to describe the network status, anomaly definition is also very domain specific and includes network infrastructure misuse, violation of a local network security policy and device misconfiguration. In addition, the expected network behaviour never remains static because it can vary with the time of day, the number of users connected and network services deployed. As a consequence, anomalies are not easy to detect.

Network sniffing

To acquire knowledge about the network status and behaviour, CINBAD collects and analyses data from numerous sources. Alarms from different network monitoring systems, logs from network services like Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), user feedback, etc. – all of these constitute a solid base of information. A naive approach might be to look at all of the packets flying over the CERN network. However, if we did this we would need to analyse even more data than the LHC could generate. The LHC data are only a subset of the total data crossing via these links.

CINBAD overcomes this issue by applying statistical analysis and using sFlow, a technology for monitoring high-speed switched networks that provides randomly sampled packets from the network traffic. The information that we collect is based on the traffic from around 1000 switches and routers and gives a representative sample of the CERN network traffic with more than 3 Terabytes of data per month. The multi-stage collection system was designed and implemented in consultation with experts from the LHC experiments and Oracle, to benefit from their data-analysis and storage experience. The system has now been up and running for more than a year (figure 1).

Network operation enhancements

The field of network monitoring and planning can greatly benefit from the CINBAD activities. We provide tools and data that simplify the operation and problem-diagnosing process. In addition, our statistics help in understanding the network evolution and design.

A very basic piece of information that is of interest for network operations is knowledge about the host's activity. CINBAD is able to provide detailed statistics about the traffic sent and received by a given host, it facilitates inference about the nature of the traffic on a given outlet/port and can thus identify the connected machine. This information could also be used to diagnose routing problems by looking at all of the packets

CINBAD keeps an eye on the CERN network (continued)

outbound or inbound to a particular host.

CINBAD is also able to provide information about the traffic at CERN. The sampled data collected by the project are sufficient to obtain the switching/routing/transport protocol information as well as gaining information about the application data. This provides valuable input for an understanding of the current network behaviour. Here the CINBAD team uses descriptive statistics. The potential set of metrics that we can provide to characterize the traffic at CERN is very extensive and specific needs are currently being discussed. For example, we can enumerate protocol-type distributions, packet size distributions, etc. Depending on the requirements, these statistics can be tailored even further.

Top n-list is another form of network summary that might be of interest. Such lists would allow the identification of the most popular application servers, either inside or outside CERN. Although this information might be available on each individual CERN server, CINBAD provides the possibility to collect these statistics for all servers of a given type, whether or not they are centrally managed by the IT Department. This information may be of value to both network engineers and application-server administrators.

These statistics can also be useful for network design and provisioning. The CINBAD project can provide valuable information about the nature of the traffic on the links. These statistics can also be used to detect the trunks with potential bottlenecks. This information can be compared with the service-level agreements that specify the conditions for link usage, enabling appropriate corrective actions to be taken.

With all of these improvements, CINBAD offers a comprehensive system to facilitate day-to-day operations, diagnose network problems and extend our understanding of network evolution and design. The CINBAD team is currently working in close collaboration with IT-CS on a visualization model of this information that is suitable for network operation and troubleshooting.

Security enhancements

Security is another area that benefits from the CINBAD project. The only safe computer is a dead computer, or at least one disconnected from the network (if no-one can get to it, no-one can harm it). Nowadays, we cannot avoid communicating with others and therefore we expose our machine to outside threats. Although CERN centrally managed desktops have up-to-date anti-virus software and firewalls, this does not guarantee that our machines and data are shielded from attacks. These tools are usually designed to detect known patterns (signatures) and there are also other machines (unmanaged desktops, PDAs, etc) connected to the CERN network that might be less protected.

Currently, detailed analysis is only performed at critical points on the network (firewall and gates between network domains). The CINBAD team has been investigating various data-analysis approaches that could overcome this limitation. These studies can be categorized into two main domains: statistical and signature-based analysis. The former depends on detecting deviations from normal network behaviour while the latter uses existing problem signatures and matches them against the current state of the network.

The signature-based approach has numerous practical applications, for example SNORT (an open-source intrusion-detection system). The CINBAD team has successfully ported SNORT and adapted various rules to work with sampled data. It seems to perform well and provides a low false-positive rate. However, the system is blind and can yield false negatives in cases of unknown anomalies.

This problem can be addressed by the statistical approach. Expected network activity can be established by specifying the allowed patterns in certain parts of the network. While this method works well for a DNS or web server that can only be contacted on a given protocol port number, for more general purposes this approach would not scale.

A second approach is to build various network profiles by learning from the past. The selection of robust metrics that are resistant to data randomness plays an important role in characterizing the expected network behaviour. Once these normal profiles are well established, the statistical approach can detect new and unknown anomalies.

The CINBAD project combines the statistical approach with the signature-based analysis to benefit from the synergy of the two techniques. While the latter provides the detection system with a fast and reliable detection rate, the former is used to detect the unknown anomalies and to produce new signatures. The CINBAD team constantly monitors both the campus and internet traffic using this method. This has already led to the identification of various anomalies, e.g. DNS abuse, p2p applications, rogue DHCP servers, worms, trojans, unauthorized wireless base stations, etc. Some of these findings have resulted in refinements to current security policies.

The future

The CINBAD project offers many opportunities to improve CERN's network operation, and it also provides a unique opportunity for the CERN Computer Security Team to identify (and protect against) incidents that might not be seen otherwise. It also enables other groups concerned with varying network applications, such as web services and mail servers, to understand their behaviour.

Milosz Hulboj, Ryszard Jurga
CERN openlab

For more information about CINBAD, please see page 3 and: www.cern.ch/openlab-cinbad

Article published in CERN Computer Newsletter.

The Oracle openlab team makes major advances

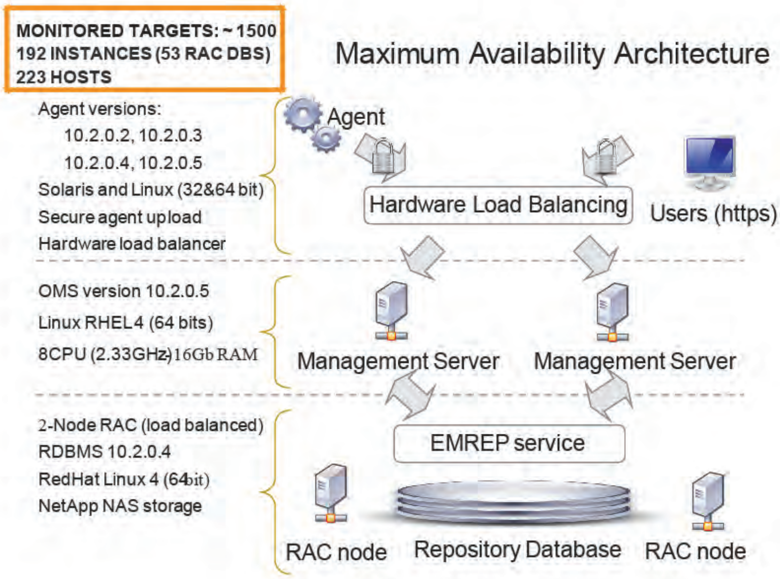

Our Enterprise Manager Grid Control Environment.

During the last six months, a large amount of work has been done within the Oracle openlab framework.

The main activities were focused on Security, Oracle Database Virtualization, work in the Software Testing Program and Monitoring.

In the Security field, the work was focused on setting up the new rules for iptables firewall in the RACs. Real Application Cluster technology requires specific iptables rules, the work consisted in developing a scheme for handling the iptables rules for large scale deployment. A report was produced that describes this work so that it can be used by other RAC users.

The main goal in Virtualization during this period was integrating OracleVM into the CERN ELFms (Extremely Large Fabric management system). Thanks to the collaboration of the CERN FIO group and Chris Barclay, Adam Hawley and Madhup Gulati from Oracle, this phase of the project has now been completed.

CERN also participated in the Software Testing Program. It has focused on the 11gR2, specifically in the Dataguard technology, with very satisfying results.

The Monitoring area has also been a very important activity. Oracle Enterprise Manager has been migrated to version 10.2.0.5, in which the virtualization console has been intensively tested, and the use of web transaction monitoring has been investigated.

The database policies have been defined, resulting in a presentation which was given by Manuel Guijarro at Oracle OpenWorld 2009.

During these six months, following outreach was presented:

- 11.2 testing result,
- Phone conferences on OracleVM and EM,
- CERN Computer Newsletter article on OracleVM,
- Oracle Press release about OEM management of OracleVM,
- Presentations at the OOW conference, the Swiss Oracle User Group Special Interest Group meeting, as well at the UK Oracle User Group technology conference.

Carlos Garcia Fernandez
CERN openlab

The team from the CERN IT Data Management group also gave a few talks related to the Oracle partnership. To find out more, please see page 5 and: www.cern.ch/openlab-presentations and www.cern.ch/openlab-reports

PARTNERS

CERN, the European Organization for Nuclear Research, is the world's leading laboratory for particle physics. For more information: www.cern.ch, http://twitter.com/cern, www.youtube.com/cern

CERN openlab is a framework for researching, developing and evaluating cutting-edge IT technologies or services in partnership with industry. Through close collaboration with leading industrial partners, CERN acquires early access to technology that is still years from general computing market. In return, CERN offers expertise and highly demanding computing environment for pushing new technologies to their limits and provides a neutral ground for carrying out advanced R&D with various partners. The technical activities are organized in different domains of competence:

- Automation and Controls Competence Centre,
- Database Competence Centre,
- Networking Competence Centre,
- Platform Competence Centre.