

Operating Systems Carrier-Grade Linux

MITA, Vol 1, Technologies and
Standardisation, pp. 149-153

MITA Seminar presentation 5.11.2003

tommi.viitanen@nokia.com

Outline

- Background
 - Open Source Development Lab (OSDL)
- Support for Replication
- Real-Time
- Embedded
- Monitored
- Fault Resistant
- Conclusions
- References

Background

- Telephone network designed reliable
 - carrier-grade
 - requirement: 30 minutes downtime per year -> 99.999% availability
- Rules to calculate downtime
 - HW and SW failure count
 - soft real-time deadlines for processing
 - availability: call completion failures
 - drop possible before call attempt completed
 - ongoing calls should not be dropped
 - overload: graceful degradation in performance
 - emergency calls always through
- Above requirements have consequences for architecture
- Mobile Internet emergence: Web and Mobile domain converge
 - Linux usage
 - some changes needed

Open Source Development Lab (OSDL)

- Central body dedicated to accelerating use of Linux for enterprise computing
 - enterprise-class testing and technical support
 - Linux-industry resources to focus investments
 - practical guidance to members
- Carrier Grade Linux Working Group
 - to guide creation of Linux platform that supports carrier grade, commercial, off the shelf development
 - focus on better positioning Linux and Linux-based SW for highly available and robust products

Support for Replication

- Faulty HW slow to replace -> replication
- Models
 - 2n: pair
 - n+1: extra units
- Redundant bus
- Permanent names for units
- Computer unit with two NICs connected to two local networks
 - different IP address models
 - one IP address for client
 - each network iface with different IP address
 - SCTP multi-homing
- Kernel changes needed

Real-Time

- Soft real-time requirements for Linux kernel
 - 1 milliseconds
 - Intel x86: 20 microseconds
 - Linux 2.4.17 kernel response time 100 milliseconds
 - low-latency and preemption patches improves performance
- Real-time programming requires
 - special operating system interfaces
 - real-time scheduling
 - real-time primitives
- Predictable performance, hard disk I/O problematic
- Network driver API changes
 - too many packets received -> early drop
- Symmetric Multi-Processing (SMP)
 - processes should be locked to named CPUs

Embedded

- Carrier-grade Linux means clusters of computer units
- Soft real-time req's imply little use for virtual memory
 - permanent storage also not needed
 - hard disk only for loading SW in start-up -> one central storage
 - easier to manage and upgrade
- Fast start-up
 - stable -> most options fixed

Monitored

- Which HW units are working and which not?
- Watchdog usage
- Important applications running all the time
 - leaking resources to be prohibited
 - visibility to resources needed
- Kernel and device driver panics
 - system images analyzed after reboot

Fault Resistant

- Kernel and device driver panics should be rare
 - removed from code
 - replaced by analyze and remedy methods
 - hardening
- Fault-injection code to be included
 - testing error-correction SW
- System failing despite precautions
 - kernel panic -> log writing ability

Conclusions

- Similar level performance for Linux servers as to fixed line telephone networks?
- Requirements covered for carrier-grade Linux
 - none impossible
 - development effort ongoing, e.g. by ODSL
- Future will show: how well and fast carrier-grade features find their way to Linux kernels

References

- MITA, Vol 1, Technologies and Standardisation, pp. 149-153
- Open Source Development Lab, Carrier Grade Linux,
http://www.osdl.org/lab_activities/carrier_grade_linux/

Thank you!

Questions?