XtreemOS Plans for the Kerrighed Project

Christine Morin, INRIA

XtreemOS scientific coordinator

February 1, 2008

Kerrighed Summit

Paris, France

Information Society Technologies

XtreemOS IP project is funded by the European Commission under contract IST-FP6-033576

LinuxSSI Component in XtreemOS

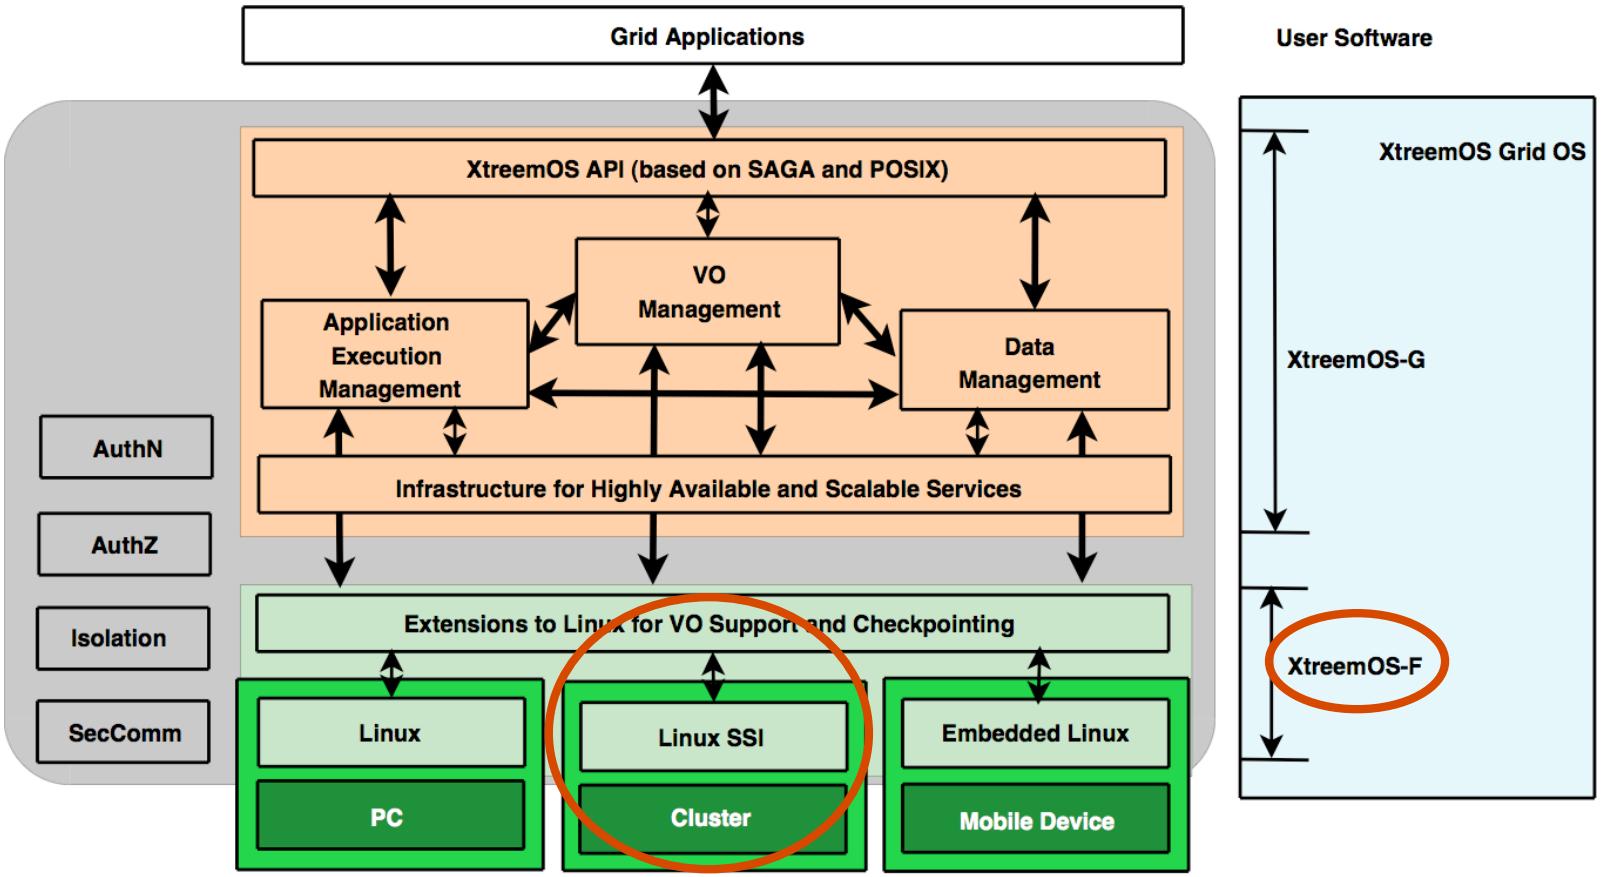

XtreemOS Activities on LinuxSSI

- **Development of LinuxSSI**
  - Services
  - Scalability
  - Pushing Kerrighed patches into Linux mainstream development
- **Integration of LinuxSSI in XtreemOS software**
- **Packaging**
  - RPM based Linux distributions (Mandriva, Redflag Linux)
  - Live-CD
- **Testing**
- **Experimentations with real applications**
- **Promotion of LinuxSSI**

Development of LinuxSSI

- **Based on Kerrighed open source technology**
  - Close cooperation with Kerrighed key developers
- **LinuxSSI: an R&D flavour of Kerrighed**
  - **Services**
    - Distributed IPC
    - Reconfiguration
    - Checkpoint/restart
    - Customizable scheduler
    - kDFS distributed and parallel file system
  - **Scalability**
    - Support for high performance NIC (IB)
    - Identification and removal of limitations to scalability
- **Pushing LinuxSSI into Kerrighed mainstream development**
  - IPC, checkpoint/restart, reconfiguration already in Kerrighed trunk

XtreemOS
Enabling Linux for the Grid

Plans for 2008

- **Complete reconfiguration framework**
  - Node addition or eviction
  - Call-backs in Kerrighed services
    - Service developers need to be involved
    - kDFS reconfiguration design
- **Customizable scheduler**
  - DRMAA interface
  - Implementation of probes and policies
    - State of art policies
    - I/O probes and I/O aware scheduling policies
    - Self adaptation of the scheduling policy

XtreemOS
Enabling Linux for the Grid

Plans for 2008

- **Checkpoint/restart**
  - Processes using IPC
  - Writable open files checkpointing
- **kDFS**
  - Efficiency (file striping)
  - File data redundancy
  - Reconfiguration
  - Support for checkpoint/restart
  - Interaction with the scheduler
- **Support of Infiniband**
- **Pushing KDDM into Linux kernel mainstream development**

Integration of LinuxSSI in XtreemOS

- **Building LinuxSSI-XOS**
  - Node level VO support mechanisms
    - Logging of Grid users
    - Kernel key rings cluster wide
  - Checkpoint/restart interface
    - Same interface for kernel checkpointers on individual PCs and clusters
      - BLCR based checkpointer for individual PCs
- **Executing XtreemOS Grid-level services on top of LinuxSSI-XOS**
  - Interaction with the application execution management service
  - Interaction with the XtreemFS Grid file system
    - Data transfers
    - Checkpoint storage

Packaging/Testing

- **Kerrighed/LinuxSSI releases packaged for RPM Linux distributions**
  - Packages available on Mandriva website
  - Links to packages from Kerrighed and from XtreemOS websites
  - Integration in Redflag Linux distro planned
- **Kerrighed packages for Debian & OSCAR**
  - XtreemOS plans to rely on the Kerrighed and OSCAR open source communities from now on
- **XtreemOS live-CD including Kerrighed**
  - Next versions will include LinuxSSI
  - Should be possible to contribute to a Kerrighed live-CD
- **Automated test platform (Mandriva)**

Experimentation with Real Applications

- **Features planned to be tested**
  - process migration
  - checkpointing
  - scheduling
  - SMP support
  - customizable scheduler
- **Features not tested**
  - scalability
  - I/O performance
  - Reconfiguration
- **Tests performed with real applications & test applications**
  - EADS, XLAB
- **Results published in D4.2.x deliverables available on XtreemOS website**

LinuxSSI Promotion

- **XtreemOS consortium promotes XtreemOS Grid OS**
  - Promotion of individual components that can be used stand-alone (eg. XtreemFS, LinuxSSI)
  - We always refer to Kerrighed when promoting LinuxSSI
  - Kerrighed 2.1.1 has been packaged and is available from Mandriva
- **Plans**
  - Make LinuxSSI release available on XtreemOS website very soon
    - 0.9 LinuxSSI internal release (Dec. 07) based on Kerrighed 2.2.0
  - New packages under construction for Kerrighed 2.2.1 and LinuxSSI 0.9
  - Promote LinuxSSI/Kerrighed
    - website, live-CD

Topics of Interest

- **Virtual clusters**
  - Deployment of Kerrighed/LinuxSSI on virtual clusters
- **Virtual machines on top of Kerrighed/LinuxSSI**
  - Portability of applications, isolation
- **cgroups and name spaces support in Kerrighed**
  - Isolation, accounting, checkpointing
- **Dynamic streams in Kerrighed**
  - checkpointing of processes using sockets and pipes
- **Service execution on top of Kerrighed**
  - XtreemOS Grid-level services
  - Business applications
- **Easy deployment on Grid'5000 clusters**
- **Pushing Kerrighed patches into Linux mainstream development**
  - Current XtreemOS strategy: pushing KDDM first

Kerrighed Community contributions to LinuxSSI

- Knowledge and code sharing
- Useful features implemented by Kerlabs
- Significant improvement in LinuxSSI stability

→ **Speed-up of LinuxSSI development**

XtreemOS Contributions to Kerrighed Community

- New features
  - kDFS, distributed IPC, checkpointing mechanisms, …
- Testing, debugging & user support
- Feed-back from experimentation with various real applications
- Kerrighed RPM packages produced by Mandriva & RedFlag, deployment tools
- Documentation
  - Contributions on Kerrighed web site
  - Deliverables on XtreemOS web site
    - Kerrighed & LinuxSSI internals
    - Packaging (RPM, Debian & OSCAR)
    - Feed-back from experiments with applications

Conclusion

- **Work in close collaboration with the Kerrighed open source community has been highly beneficial to XtreemOS consortium**
  - Wish to keep LinuxSSI as close as possible to Kerrighed
    - Push LinuxSSI services and mechanisms into Kerrighed official releases
  - Kerrighed developers phone conferences needed
    - Technical discussions, community organization, releases, roadmap
  - Reinforce collaboration to push Kerrighed patches in mainstream Linux kernel development

Conclusion

- **XtreemOS consortium will continue to contribute to Kerrighed open source community**
  - Development of LinuxSSI
  - Testing, debugging, packaging for **RPM based Linux distributions**, live-CD, experimentation with real applications, user support
  - Documentation
    - XtreemOS D2.2.x, D4.1.x, D4.2.x public deliverables
- **XtreemOS consortium will communicate on LinuxSSI and its contributions to Kerrighed**

Questions?

XtreemOS IP project is funded by the European Commission under contract IST-FP6-033576