Boomer: Next Generation Solaris Audio

Garrett D'Amore
PSARC 2008/318, Inception 11/5/2008
Sun Microsystems, Inc.

Boomer Overview

- New Kernel Audio Framework
- Supports multiple audio APIs
  - Sun legacy audio(7I) & mixer(7I)
  - Open Sound System
- New features
  - Multichannel (5.1, 7.1) audio
  - Higher sample rate & resolution
- Developer friendly DDI
- Full support for legacy devices

Boomer Goals – Primary

- Compatible
  - Solaris Applications
  - OSS (Linux and FreeBSD) Applications
  - Existing Solaris-supported audio hardware (no regressions)
  - Linux Branded Zones
  - New Hardware (e.g. Audigy, AudioPCI, etc.)
- Capable
  - Multi-channel support (at least 7.1)
  - High definition audio (192kHz, 24- or 32-bit)

Boomer Goals – Secondary

- Performant
  - Highly optimized mixing, format and sample rate conversion algorithms
  - Low latency – suitable for multimedia and VoIP
- Scalable
  - Future Sun Ray support
- Extensible
  - Future APIs & features (ALSA, Dolby Digital)
- Developer-friendly
  - Simple, straight-forward DDI

Boomer Non-Goals

- No MIDI or Wave-table Synthesizer
- Remove certain legacy features
  - No “compatible” (exclusive only) mode
  - Multistream codecs not exposed to apps
- No latency guarantees
  - Consumer audio applications only
- Not a porting layer
  - Porting drivers from other platforms takes effort
- Phase 2 Projects
  - Sun Ray support
  - Dolby Digital (AC3) Pass-Thru

Boomer Subsystems

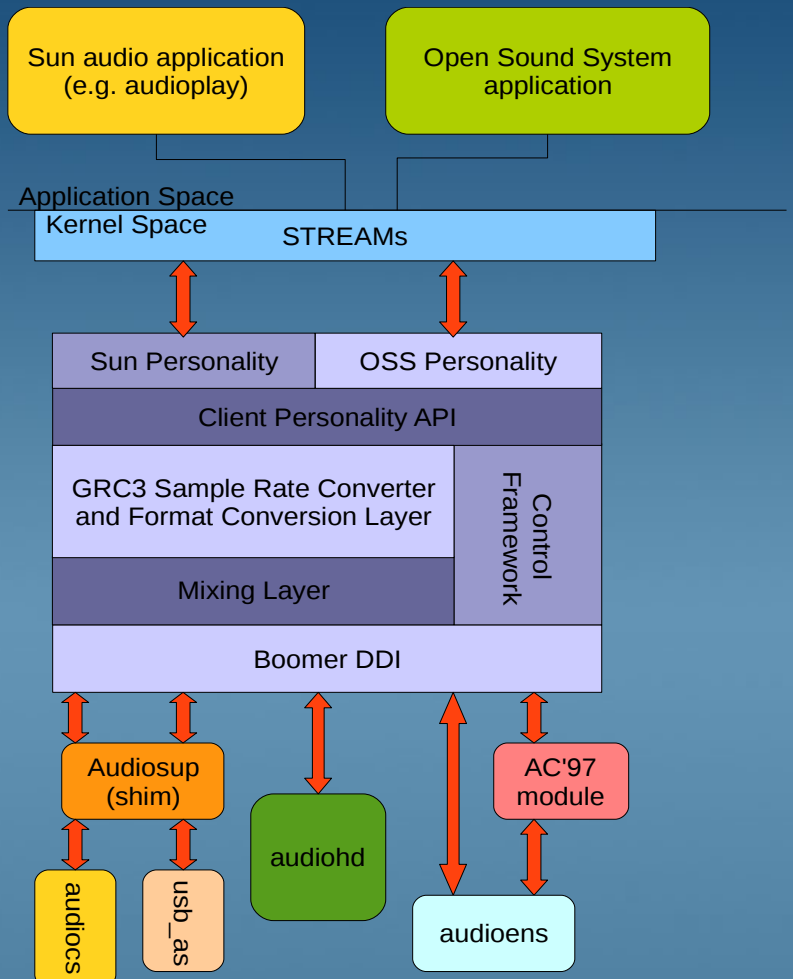

Audio Core

- /kernel/drv/{,sparcv9,amd64}/audio
  - Depends only on genunix
- Functionality needed by all audio drivers
  - Software mixing
  - Sample rate and format conversion
  - “Client personalities” - Sun & OSS
  - Driver DDI
  - Control framework
  - DDI glue (cb_ops, STREAMs srv, put, etc.)

Client Personality Interface

- Common layer used by multiple APIs
  - Sun /dev/audio and /dev/audioctl
  - Multiple OSS APIs (mixer, pcm, and sndstat)
- Narrow interface
  - #include “audio_client.h”, not “audio_impl.h”
- Uses circular FIFO for data exchange
  - FIFO contents format (incl. sample rate) chosen by client implementation
- Creates minor nodes
  - one for each personality
  - associated with hardware device node

Open Sound System Personality

- Popular Linux/FreeBSD audio API
  - Supported by many freeware applications
  - http://manuals.opensound.com/developer/
- Has some issues
  - See Boomer Specification for details
  - Might not implement certain diagnostic use portions
  - Mixer APIs have issues ... inconsistent in 4Front
  - May “constrain” certain names for 4.x mixer API
- Uncommitted (for now)
  - Until Sun Ray is supported, cannot raise commitment

Common Core

- Mixing layer operates on 24-bit PCM
  - Can detect overflows into 32-bits
- Each stream has own settings
  - Independent attenuation – relative to master
  - Independent sample rates
  - Independent formats
- GRC3 sample rate converter
  - Licensed from 4Front
- Exports data to hardware format
  - 16, 24, or 32 bit signed PCM (any endian)

Audio DDI

- Consolidation Private (for now)
- Very easy to use
  - Similar to Linux, NetBSD, FreeBSD driver API
  - Uses simple & efficient circular FIFO
  - Average driver about half size of legacy driver
- No DDI duplication
  - Drivers can make full use of any Solaris DDI
- Supports multiple “engines” per device
  - Unlimited, “multistream codec” support
  - Engines can be heterogeneous

Audio Control Framework

- Exposes hardware knobs to client personalities
  - Generally controls are “out-of-band”
- Well known (numeric) controls and driver private controls
  - Modeled on Brussels (NIC drivers)
- Examples
  - Monitor gain
  - 3D effects

Common AC'97 Module

- Separate misc/ac97 module
- Implements common AC'97 codec operations
  - Drivers supply low level register access routines
  - Similar to MII in principle
- Drivers can override methods for specific controls, or add their own controls
- Many drivers can use
  - Examples: audiots, audio810 ,audioens
  - Counter-examples: audiohd, usb_as, audiocs

Driver Compatibility Shim

- Replaces misc/audiosup
  - misc/mixer, misc/amsrc2 stubs or removed
- Converts SADA driver to Boomer
  - Not a complete emulation
  - Just “good enough” for current SADA drivers
- Drivers are limited by SADA
  - No multi-channel audio, for example
- Remove when drivers updated to Boomer
  - SADA interfaces become Obsolete

High Definition Audio (audiohd)

- Intel specification
  - Replaces AC'97
  - Full controller spec, including PCI class
- Rich features
  - 15 playback and 15 record voices
  - Up to 16 output channels
  - Extremely flexible routine & codec configuration
- Enhanced support in Boomer
  - 7.1 audio
  - High definition (24-bit, 192 kHz) audio
  - SPDIF digital output

New Drivers

- audiopci
  - Supports legacy ENS1370 (Creative AudioPCI)
  - Emulated in some environments (VMware)
- Creative Audigy, SBLive!, SB-XFi
  - Common commodity parts
  - Import drivers from 4Front (licensed)
  - Will depend on resourcing and demand
- Other drivers as demand exists

Boomer Inception
PSARC 2008/318
November 5, 2008

garrett.damore@sun.com