



# Datacenter Fabric Workshop

## Windows IB



## Windows Core SW

---

**Fab Tillier**

*SilverStorm Technologies*

[ftillier@silverstorm.com](mailto:ftillier@silverstorm.com)

**August 22, 2005**



# Agenda

---



- Overview
- SVN Layout
- Architecture
  - Access Layer
  - Bus Drivers
  - Driver Load Order



# Goals

---



- Industry standard IB SW stack
  - Increases adoption rate
  - Decreases development costs
- Broad range of upper level protocols
- Tight integration into the Windows OS
- Windows Server 2003
  - x86, x64, and Itanium
- Windows XP
  - x86, x64



# Features

---



- Derivative of IBAL SF Project
- InfiniBand 1.1 Compliant
- Support for 32-bit app on 64-bit kernel
- Integration into OS PnP framework



# Components

---



- HCA driver
- Bus Drivers/Access Layer
- IPoIB
- SRP Initiator\*
- Winsock Direct Provider\*
- SDP
- uDAPL
- OpenSM
- MPI

\*Available only on Windows Server 2003



# Agenda

---



- Overview
- **SVN Layout**
- Architecture
  - Access Layer
  - Bus Drivers
  - Driver Load Order



# SVN Layout

---



- User- and Kernel-Mode trees overlap
- Typical structure is:
  - component ← files shared between UM and KM
  - component\user ← UM-only files
  - component\kernel ← KM-only files
- Include tree slightly different
  - inc\complib ← common complib headers
  - incliba ← common IB headers
  - inc\user\complib ← UM complib headers
  - inc\user\iba ← UM IB headers
  - inc\kernel\complib ← KM complib headers
  - inc\kernel\iba ← KM IB headers



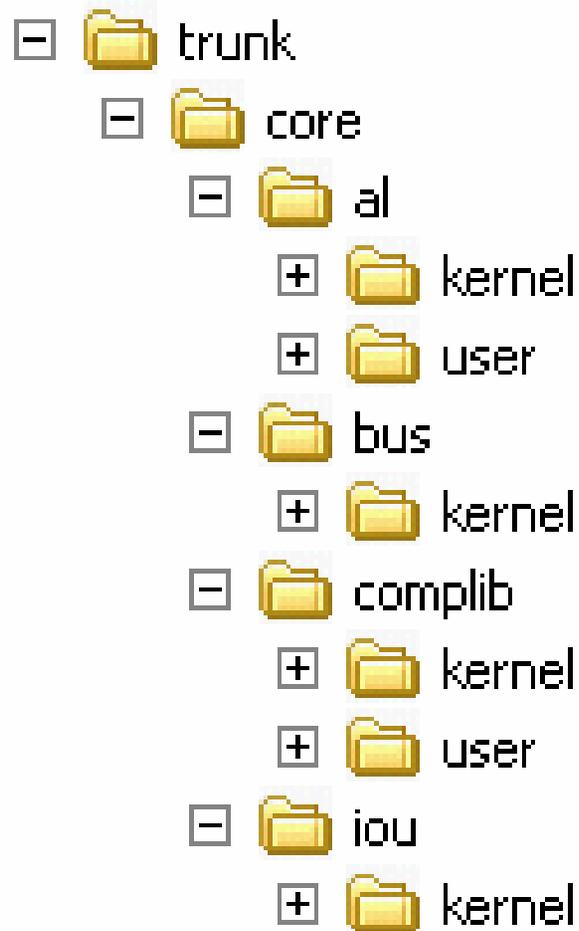
# SVN Layout - Trunk

---

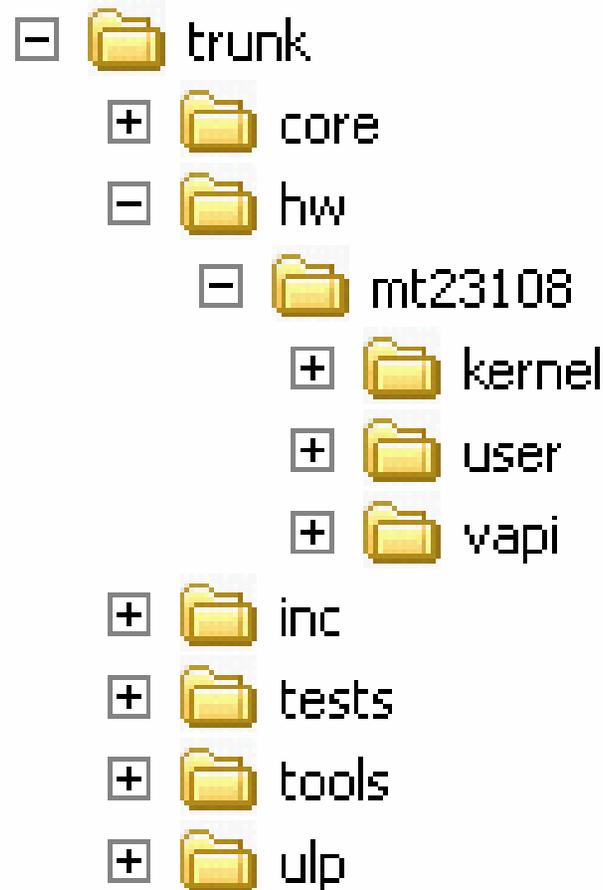


- [-]  trunk
  - [+]  core
  - [+]  hw
  - [+]  inc
  - [+]  tests
  - [+]  tools
  - [+]  ulp

# SVN Layout - Core

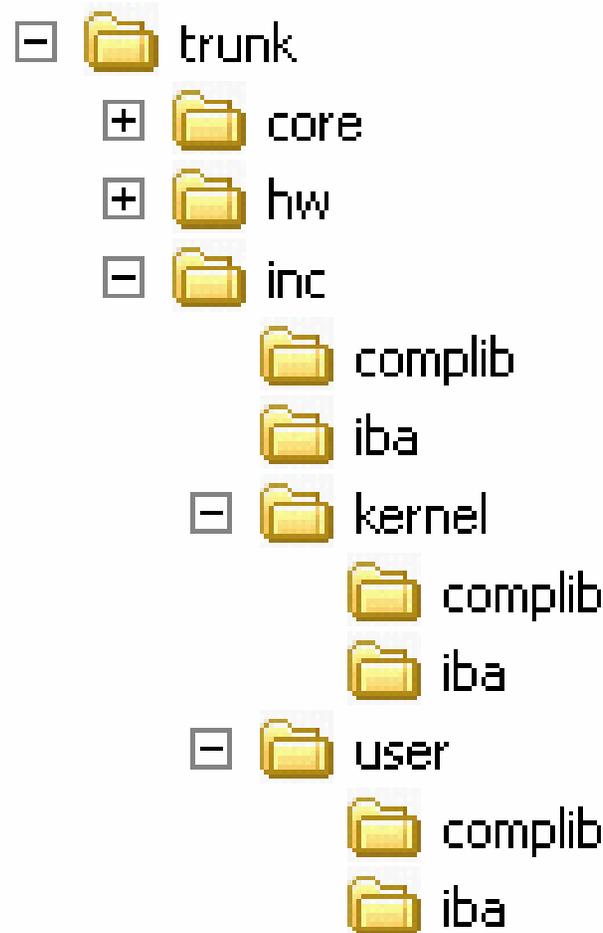


# SVN Layout - HW





# SVN Layout - Inc





# SVN Layout - ULP



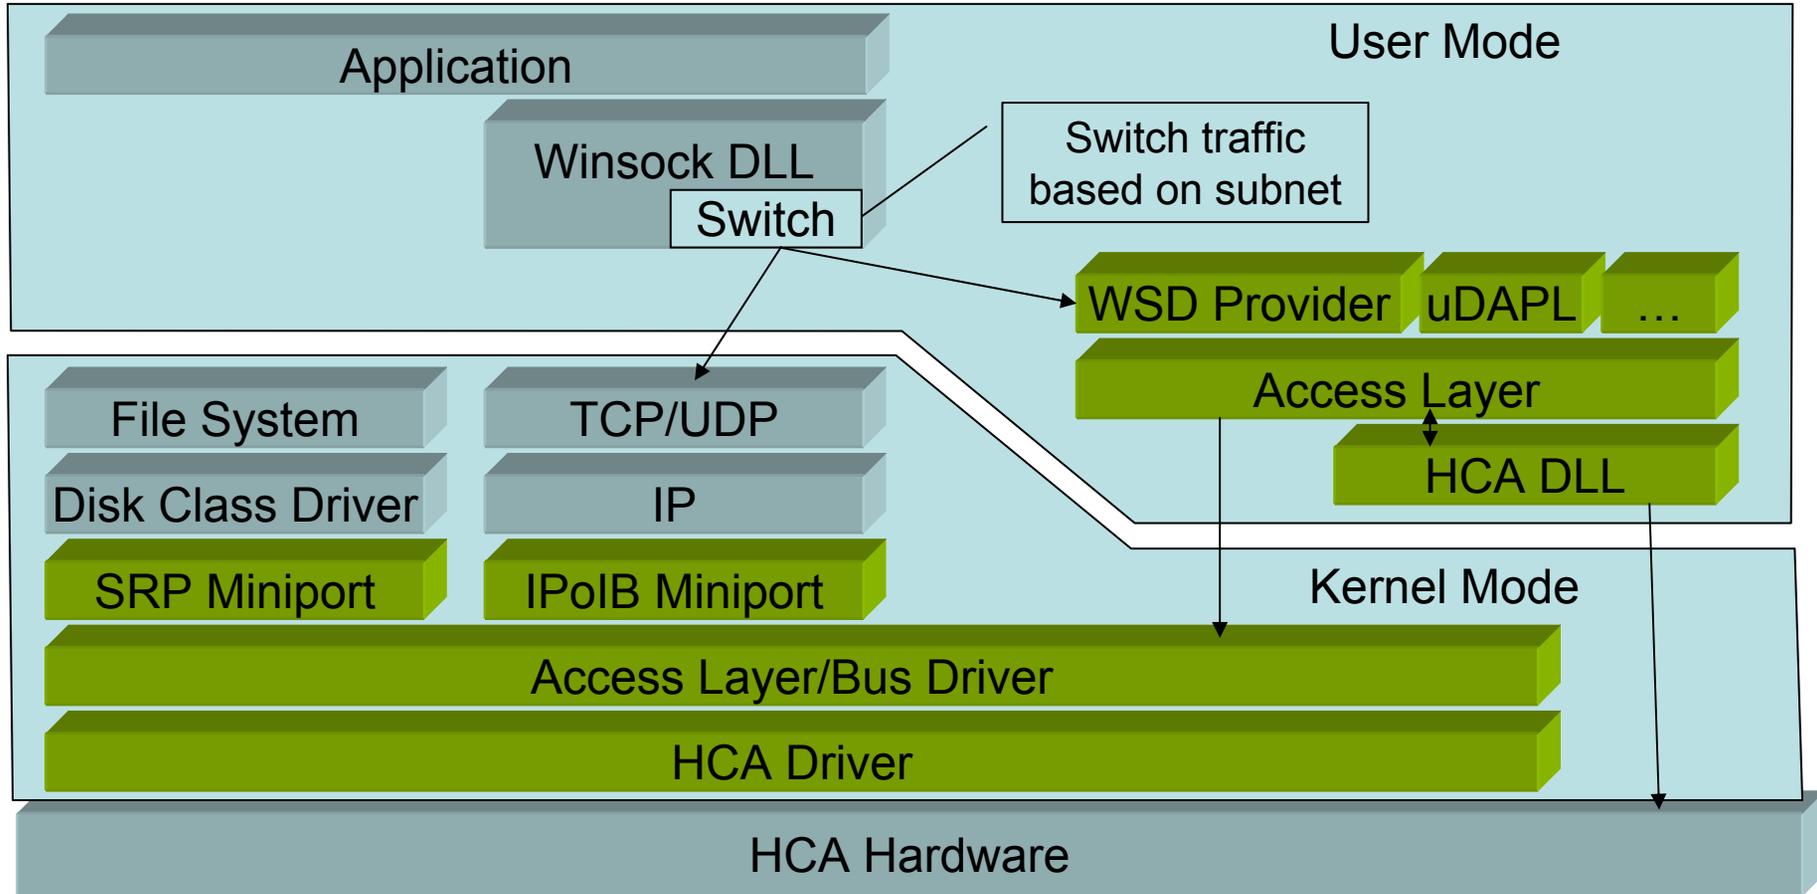


# Agenda

---



- Overview
- SVN Layout
- **Architecture**
  - Access Layer
  - Bus Drivers
  - Driver Load Order





# Agenda

---



- Overview
- SVN Layout
- Architecture
  - Access Layer
  - Bus Drivers
  - Driver Load Order



# Access Layer

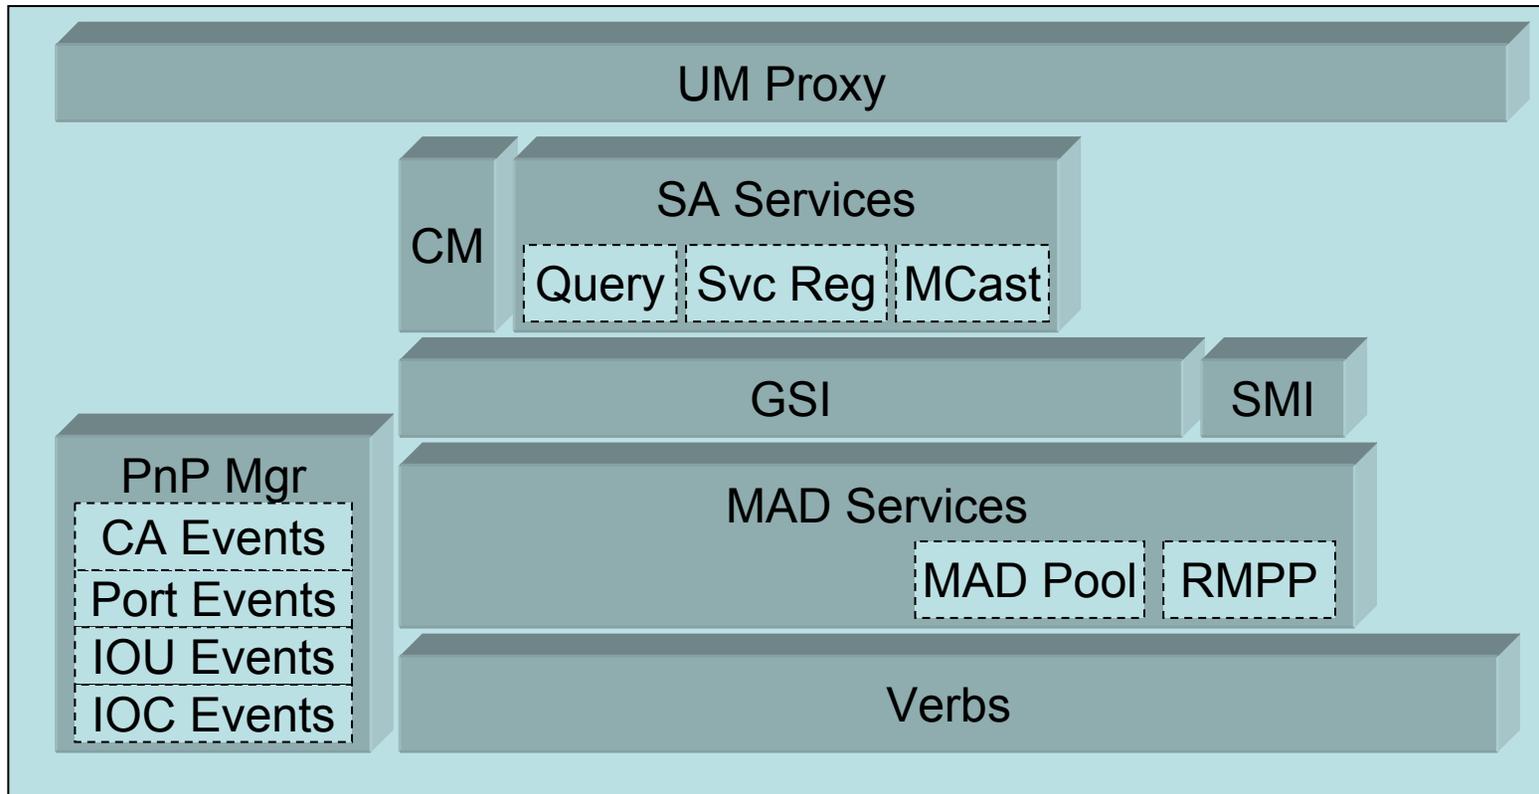
---



- Provides all IB services
  - Verbs
  - MAD services & pools
  - SMI
  - GSI
    - CM
    - Subnet Query
    - Service Registration
    - Multicast Management
  - IB Plug and Play notifications
  - User-mode proxy
  - Client management
- Consistent API look and feel for all operations

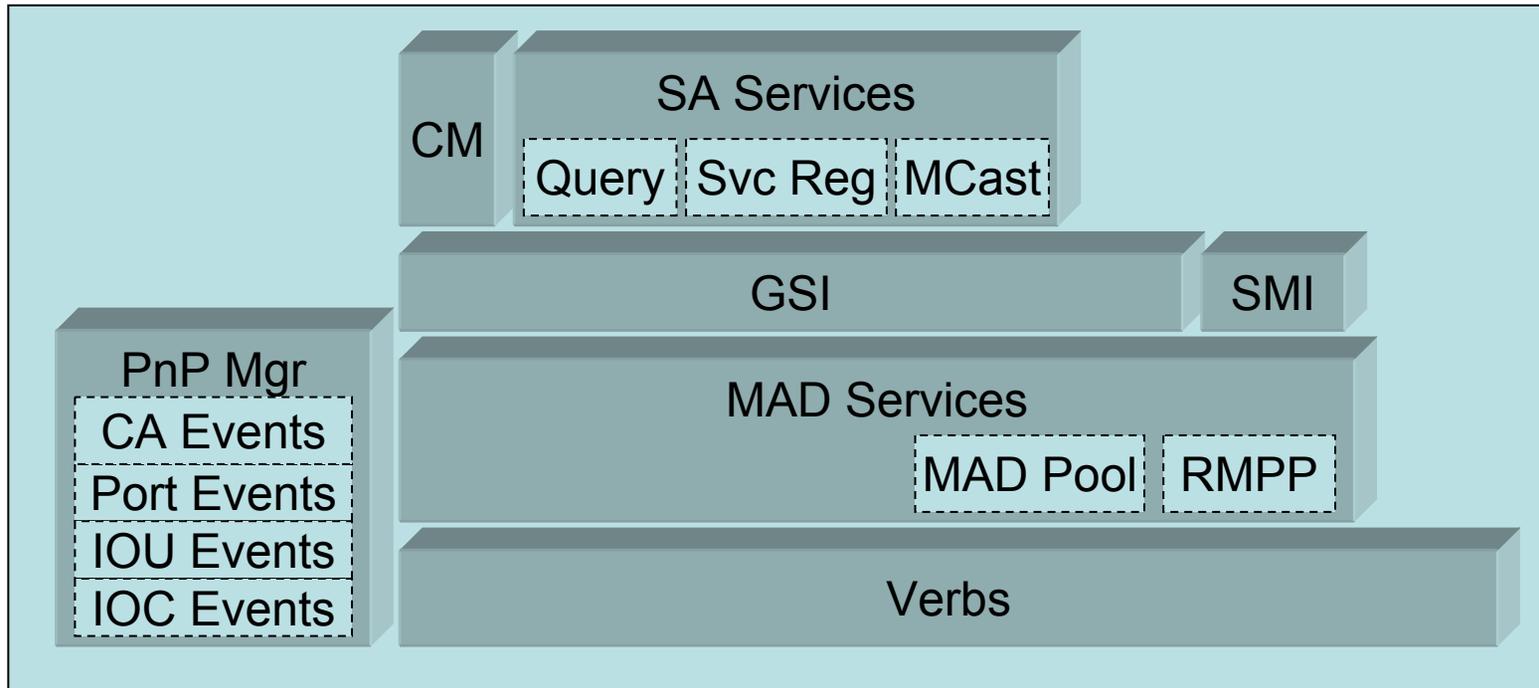


# Kernel Access Layer





# User Access Layer





# 32-bit Apps on 64-bit Kernel



- All IOCTL structures defined to be constant between architectures
  - Embedded pointers in IOCTLs are 64-bits
    - `__ptr64` modifier does the right thing
- It's that simple!

```
struct _ual_open_ca_ioctl_in
{
    ci_umv_buf_t          umv_buf;
    ib_net64_t           guid;
    void* __ptr64        context;
};
```



# Agenda

---



- Overview
- SVN Layout
- Architecture
  - Access Layer
  - **Bus Drivers**
  - Driver Load Order



# Bus Drivers

---



- Reports devices to PnP Manager
  - IPoIB
  - I/O Units
  - I/O Controllers
- Manages PDO for devices
  - Forwards IRP to parent
- Devices created as children of HCA
  - PCI Interface comes from HCA devnode



# Bus Drivers





# Agenda

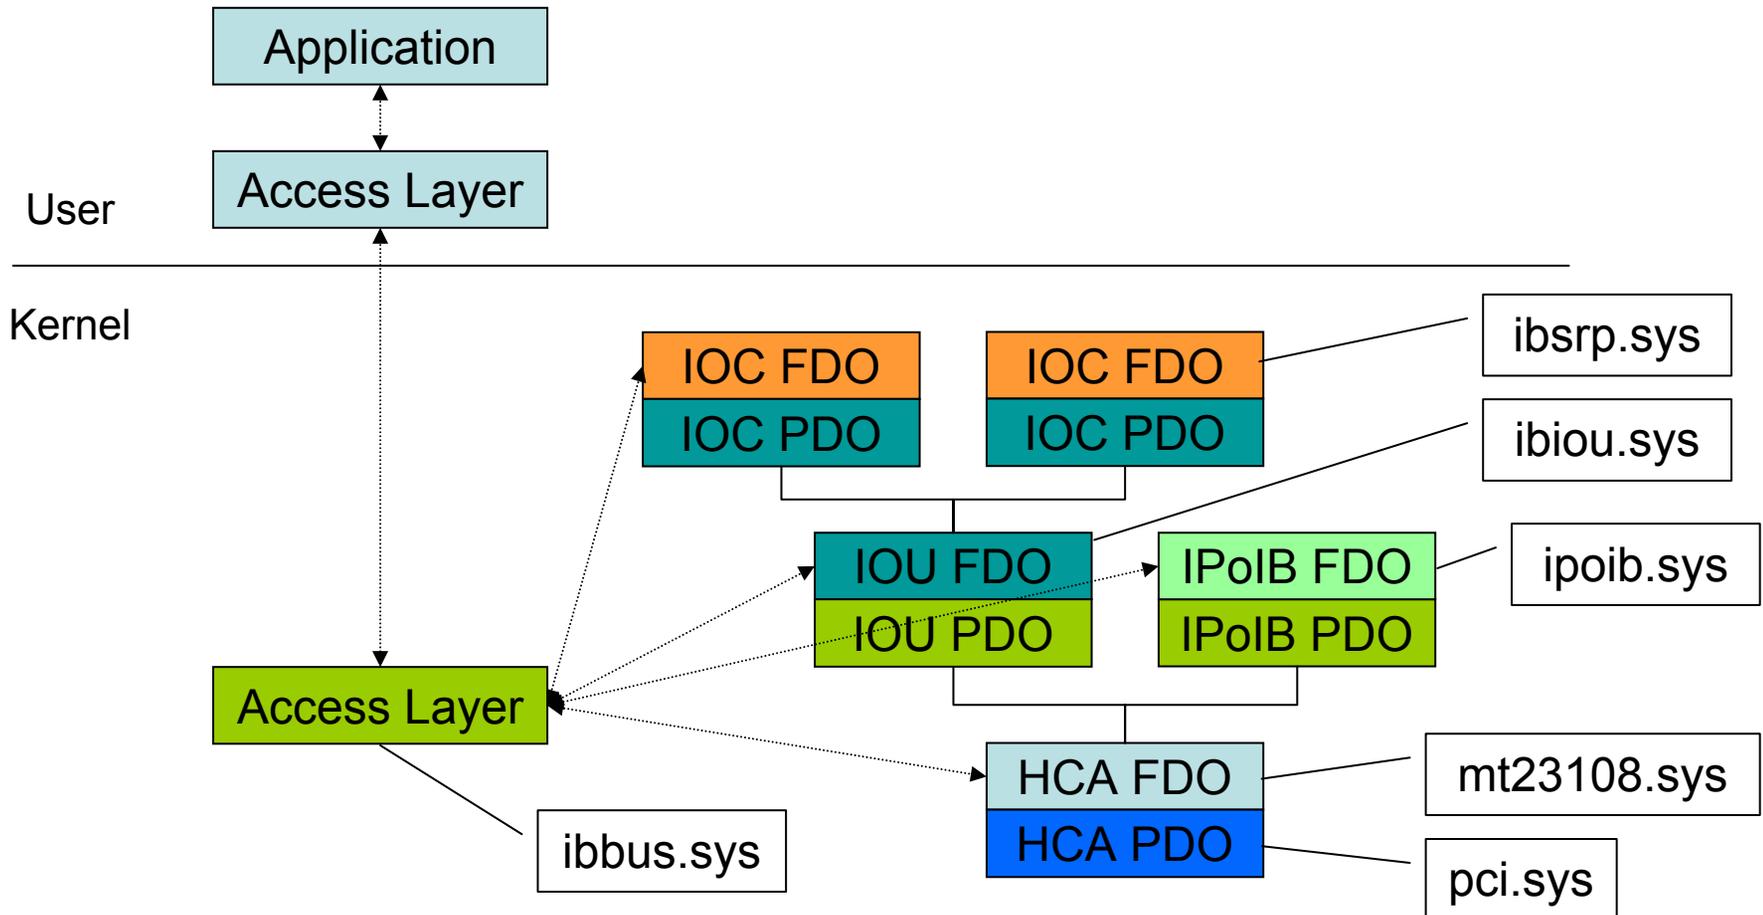
---



- Overview
- SVN Layout
- Architecture
  - Access Layer
  - Bus Drivers
  - Driver Load Order



# Driver Load Order





# Resources

---



- OpenIB Wiki
  - <https://openib.org/tiki/tiki-index.php?page=OpenIB+Windows>
- Openib-windows mailing list
  - <http://openib.org/mailman/listinfo/openib-windows>
- Sign up to contribute
  - <http://windows.openib.org/openib/contribute.aspx>



# Q & A

---

