

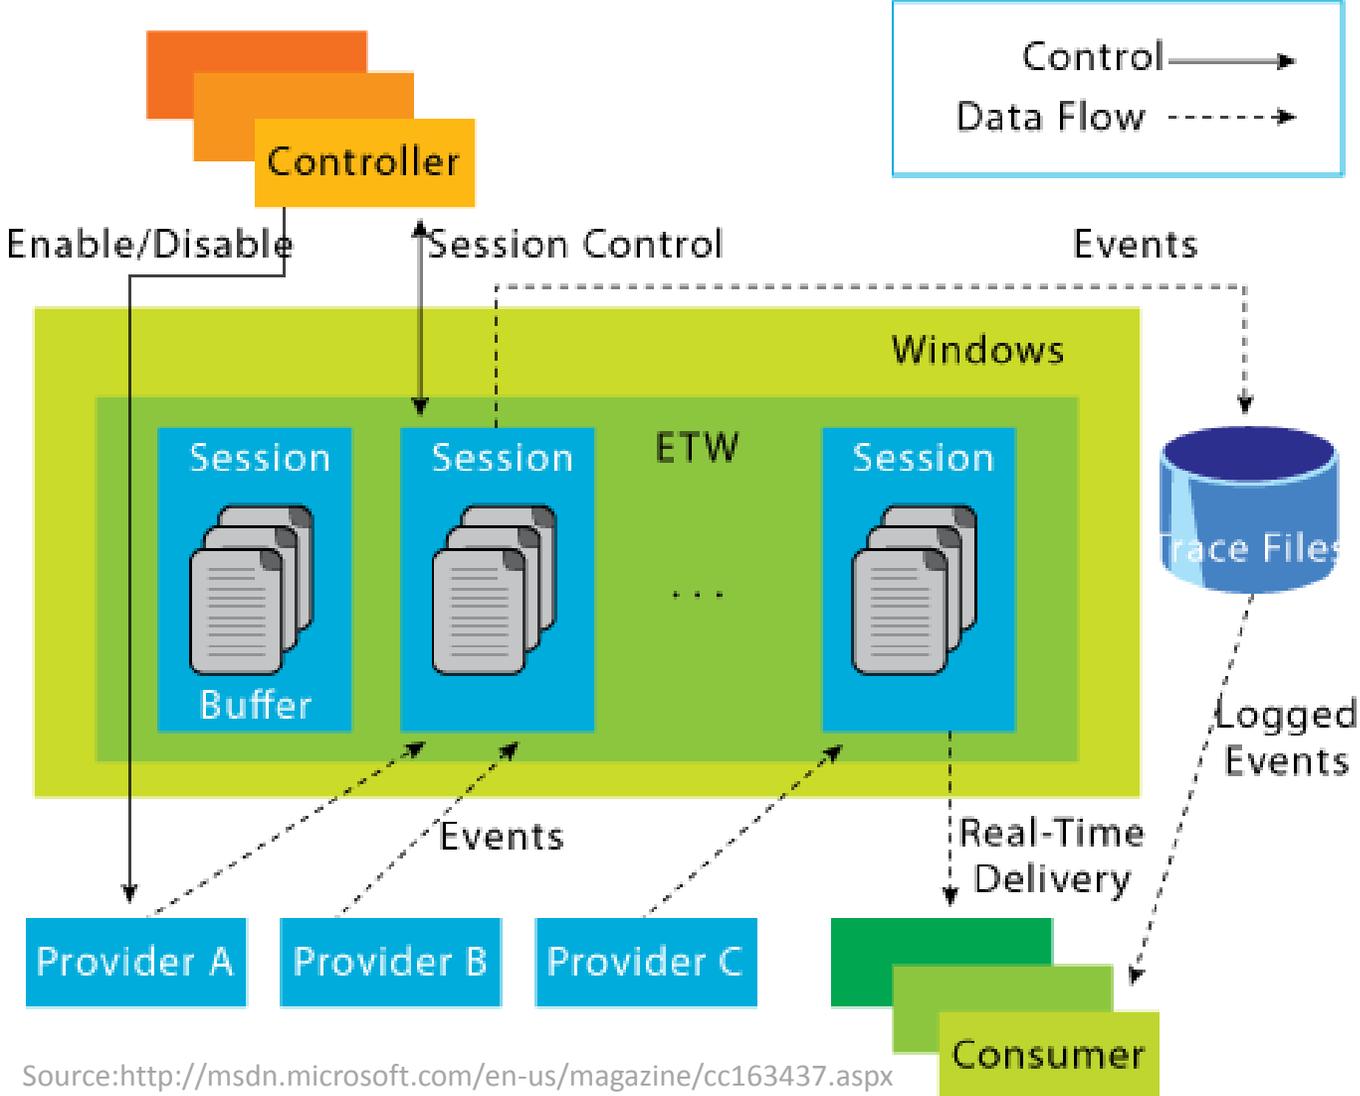
# Event Tracing for Windows (ETW)

Uwe Hartmann

# ETW Classification/General

- ETW = tracing facility
- efficient, general purpose, holistic performance analysis
- kernel-level
- Purpose: Debugging, performance profiling, monitoring/logging
- Enable/Disable tracing: dynamically, (without application restart)
- Live production system analysis
- Originally developed for windows development performance analysis

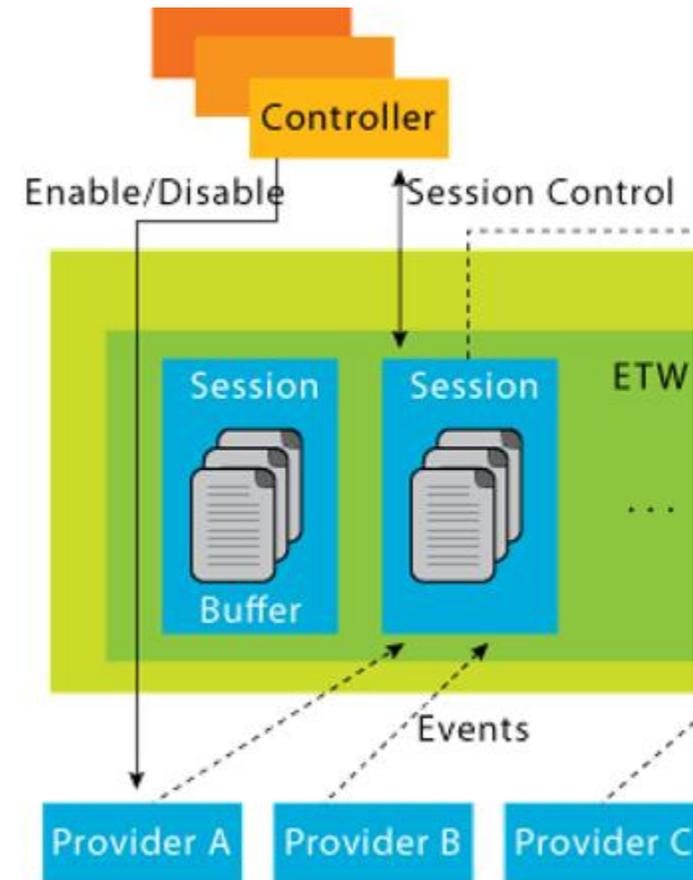
# ETW Architecture Overview



# ETW Architecture Deep Dive

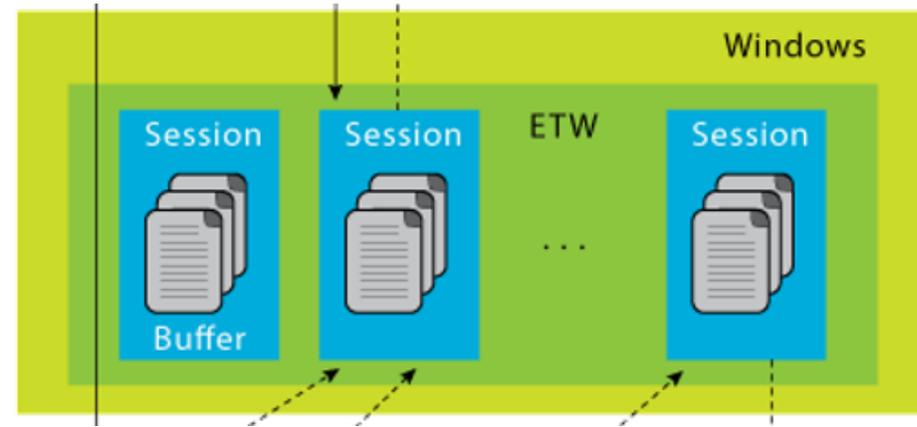
# Controller

- Any Process
- Functions
  - Start/stop Session
  - Configure Session
    - Size/type/location of log file/process for real-time delivery
    - Enable providers
    - Manage size of buffer pool
    - Resolution of time stamp
    - Verbosity, filtering
    - Update Session
- Obtain execution statistics of the Session
  - Number of events lost



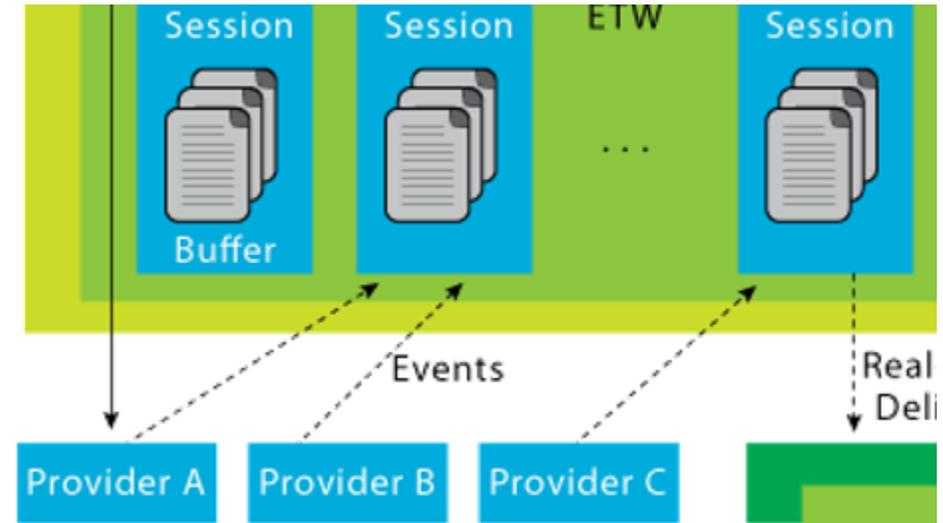
# Session

- Record events from (multiple) providers.
- Managing and flushing the buffers
- Buffers per processor, in nonpaged pool kernel memory
- Not statically tied to providers
- Log to
  - Disk
  - Memory
  - Real time delivery
  - Ring buffer ...
- NT Kernel Logger Session
  - The NT Kernel Logger event tracing session records predefined system events generated by the operating system, for example, disk IO or page fault events.



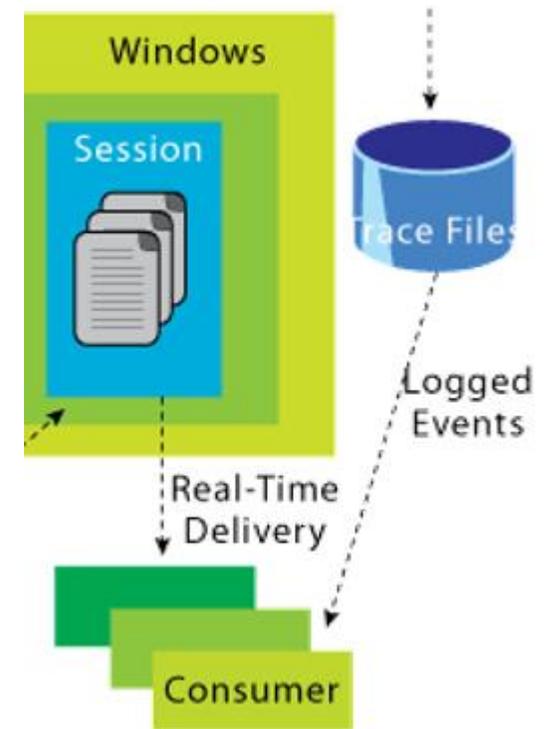
# Provider

- User-mode or kernel-mode
- for capturing log data



# Consumer

- Uses data collected by the sessions
- Can register callback for real-time consumption
- Reads recorded data from binary file



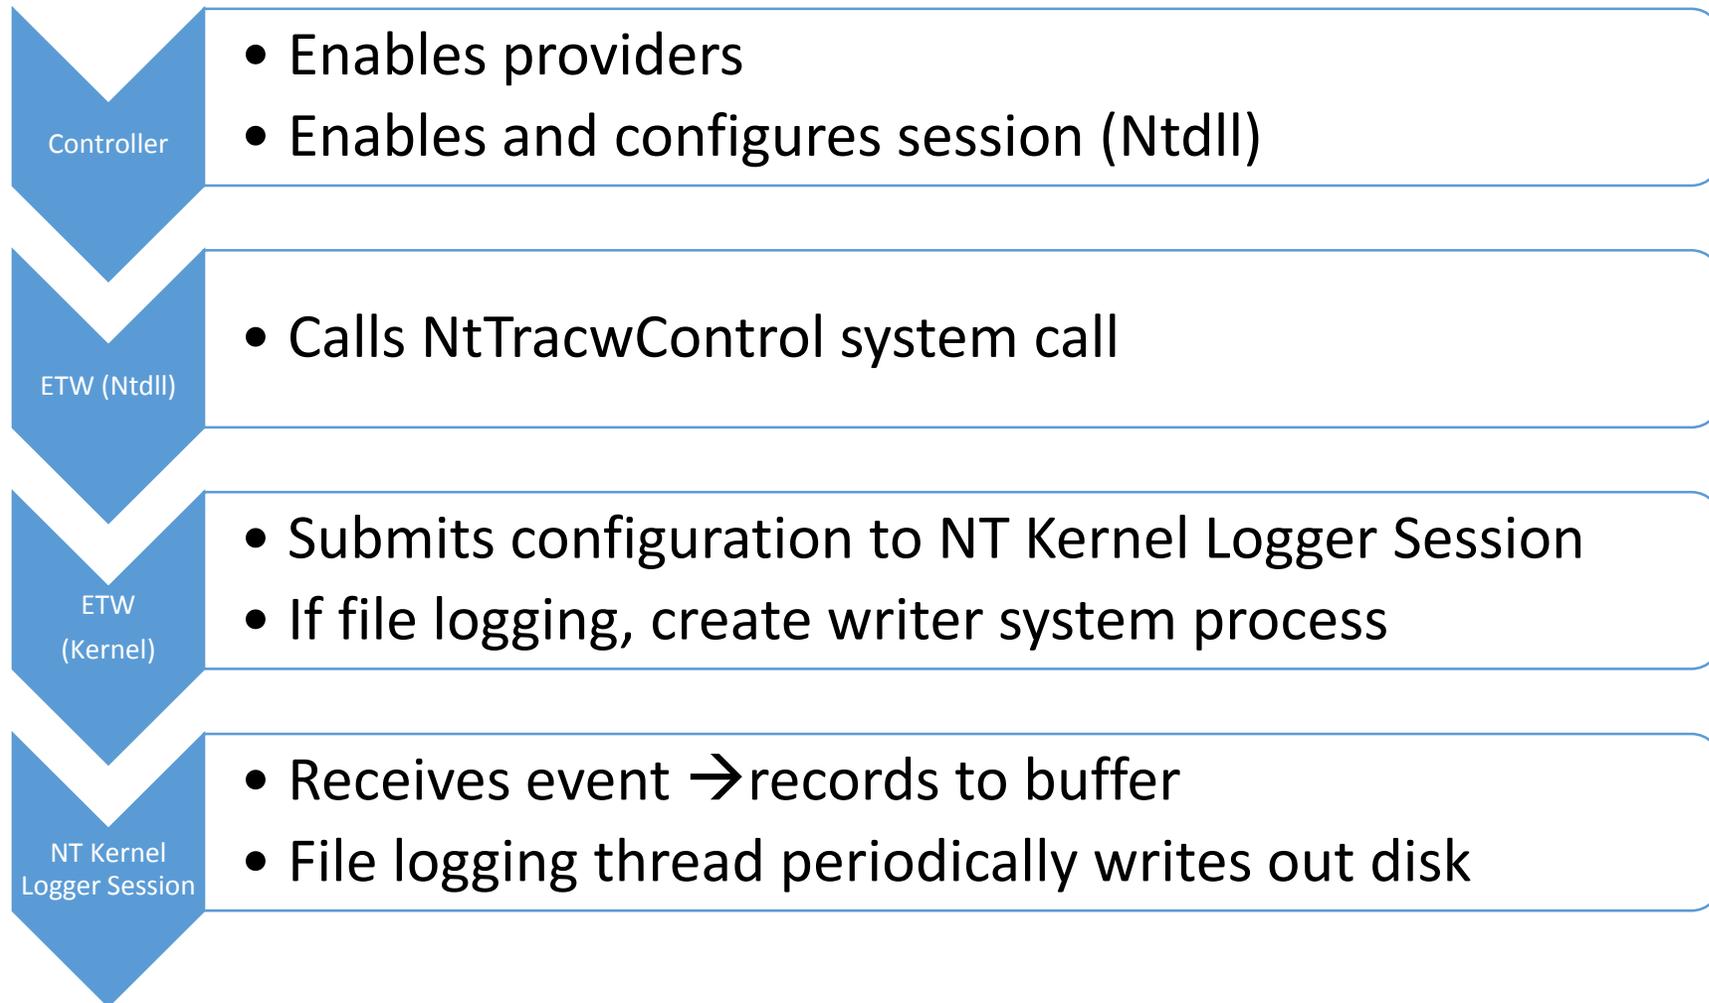
# Special Purpose Sessions

- Global Logger Session
  - The Global Logger event tracing session records events that occur early in the operating system boot process, such as those generated by device drivers.
- Private Session
  - user-mode event tracing session that runs in the same process as its event trace provider

# ETW Library

- Contained in Ntdll.dll
- Parts in ntoskrnl.dll
  
- ETW kernel-mode API
  - For drivers
  - Mirror ETW Library principles, slightly different functions
  - Completely Contained in ntoskrnl.dll

# Control Flow Kernel logging



Consumer	Controller	Provider
XPerfView	Xperf	Windows kernel
Windows Performance Analyzer (WPA)	Windows Performance Recorder (WPR)	Ressources (CPU, Memory, Disk, File IO, Networking, Power, GPU ...)
PerfMon	Process monitor, explorer	.Net, XAML
Resource Monitor	Resource Monitor	SQL Server, ASP.net
...	Microsoft Message Analyzer (Network Monitor)	Audio, Video
	PerfMon	...
	...	

# ETW vs. DTrace

## **Dtrace**

- Application predefines places to log
- Probes
- D Language
- Selection and transformation at recording Time
- Can be dynamically enabled/disabled
- Low Overhead, usable for production environment, lockless

## **ETW**

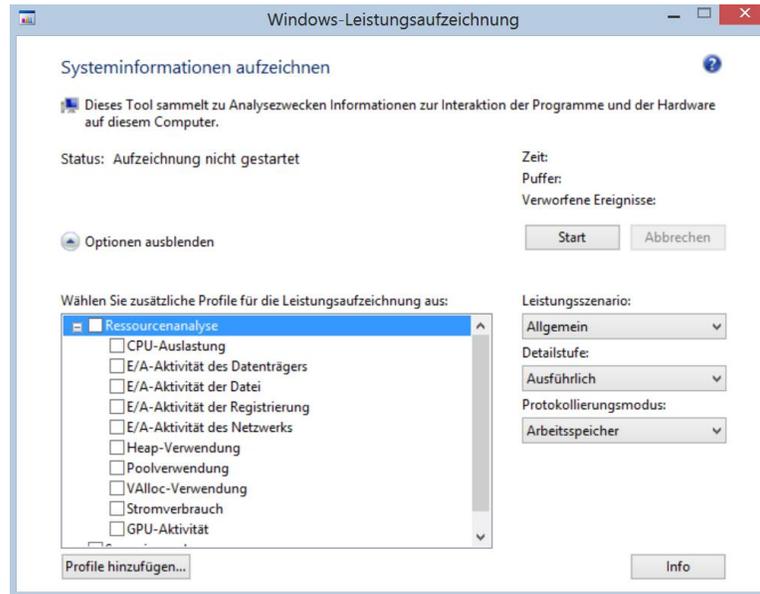
- Application predefines information and places to log
- Events
- Static typed
- Selection and transformation at analysis Time
- Can be dynamically enabled/disabled
- Low Overhead, usable for production environment, lockless

# Dropped Events

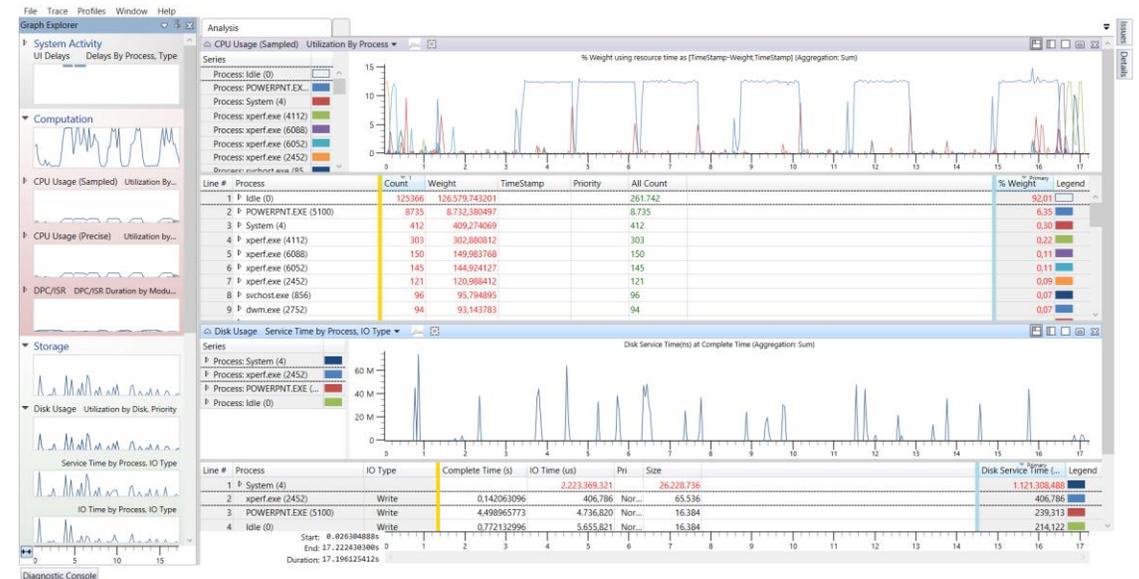
- Events too large
- Real-time consumer not fast enough
- When logging to file disk too slow

# Windows Performance Toolkit (WPR + WPA)

- Used within Microsoft to build Windows, Office etc.



Windows Performance Recorder  
ETW Controller



Windows Performance Analyzer  
ETW Consumer

# Powerpoint Demo

# Purpose

- Visual Studio Profiling Tools
  - During development in particular for CPU and memory issues
- EWT
  - For everything else in windows
  - Pick your tool or write your own
  - In particular for production environment profiling and debugging, operating system profiling
  - Advanced issues

# Scenarios

- Debugging
- Monitoring
- Diagnosis
- Capacity Planning
  
- Scanning
- Delta Analysis
- Statistical Analysis
- State Machine and Resource Tracking
- End-To-End Tracing

End