

Performance Co-Pilot and Large System Performance

Ken McDonell
Strategic Software Organization
Silicon Graphics Inc.

kenmcd@sgi.com


Understanding Performance - What's the Problem?

- **Application-centric versus system-level performance**
- **Big systems and the tyranny of scaling**
 - lots of performance data, need all of it some of the time, cannot afford to process all of it all of the time
- **Who needs to understand system performance?**
 - most people in the software life-cycle
- **Where is the performance data hiding?**
 - everywhere
 - multiplicity of data semantics, access methods and evolution across releases


Issues for System Level Performance Management

- Scalable infrastructure, protocols and tools
- Intelligent filtering of performance data and automated reasoning
- Flexible archive logging and retrospective analysis
- Extensible framework for accessing performance data from disjoint domains
- Customization of tools
- Distributed processing with centralized monitoring and management

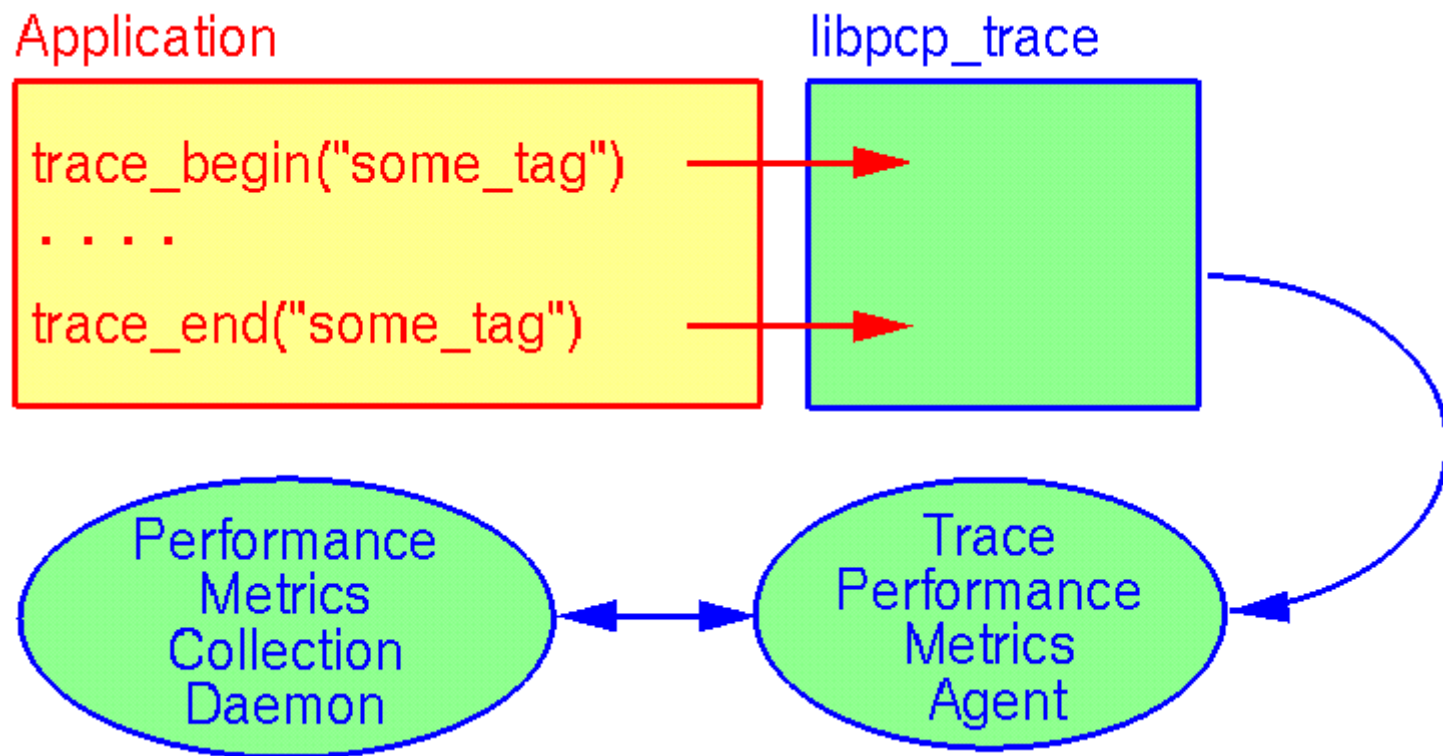
Scalability and Information Overload (1)



Scalability and Information Overload (2)




Getting Performance Data Out of an Application



Customization


Problem Solving Rate

[moomba] trace.point.count Large [moomba] trace.point.count Medium [moomba] trace.point.count Small




CPU Utilization

[moomba] irix.kernel.all.cpu.user [moomba] irix.kernel.all.cpu.sys [moomba] irix.kernel.all.cpu.sxbrk
 [moomba] irix.kernel.all.cpu.intr [moomba] irix.kernel.all.cpu.wait.total [moomba] irix.kernel.all.cpu.idle



Memory Demand

[moomba] irix.swap.pagesout



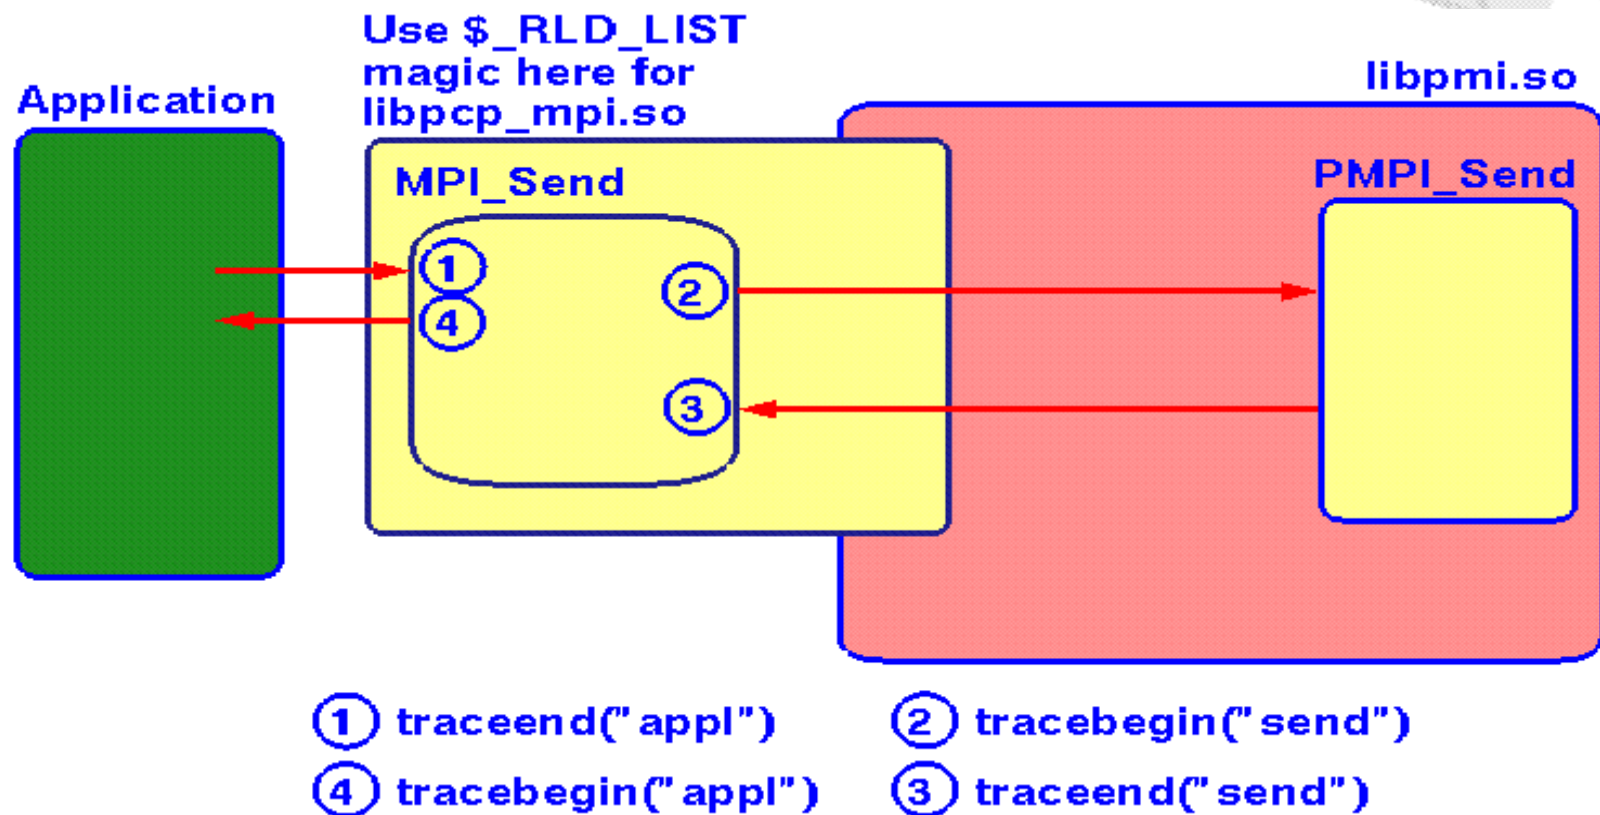
09:23:30 09:24:05 09:24:40 09:25:15 09:25:50 09:26:25 09:27:00 09:27:35 09:28:10 09:28:45 09:29:20

Fri Jun 12 1998 EST


Case Study - MPI Visualization

- Parallel applications built using the Message Passing Interface (MPI).
- All MPI services via calls to *libmpi.so*
- Dynamic behavior is hard to conceptualize
- Consider execution to be in several states: one for the application and one for each MPI routine
- Behavior may be characterized by frequency of state transitions and time spent per state
- Aggregate data over all parallel components

Capturing the Performance Data



MPI Visualization



Concluding Comments

- **Performance monitoring and management for large systems is a difficult problem**
- **Performance Co-Pilot offers some capabilities that help:**
 - visualization
 - automated reasoning
 - flexible archive logging
 - retrospective analysis
 - extensible collection framework
 - customizable tools