

ALTAIR BREEZE™ - HPC APPLICATION I/O PROFILING & DEPENDENCY ANALYSIS

Breeze makes every engineer an I/O expert. It is a tool for profiling HPC applications, collecting data on dependencies and I/O across the whole process tree. With Breeze, you can quickly solve software deployment problems and resolve file and network dependencies. With detailed data for storage exports and summary reports for sharing, Breeze identifies good and bad I/O for easy wins.

Even when programs are designed well, it's easy for local settings to affect performance. Mistakes such as storing data in the wrong place, small I/O, and excess metadata can overload compute, network, and storage. By using Breeze and following our best practices you'll fix short-term problems, understand storage requirements, and optimize future planning. Breeze profiles application file I/O so you can ensure files are stored in the right place.



Breeze identifies good and bad I/O for easy wins

Understanding the Way Applications Access Data

Performance - Understanding I/O patterns for debugging, resource sharing, and tuning

Portability - Understanding application dependencies for migration and containerization

Planning - Understanding what you are using today so you can plan for tomorrow as well as deploying system-wide best practices such as user education and chargeback

Profiling application I/O means you can make better decisions about how to deploy applications on-premises and in the cloud.

How Much Time Are You Wasting on Bad I/O?

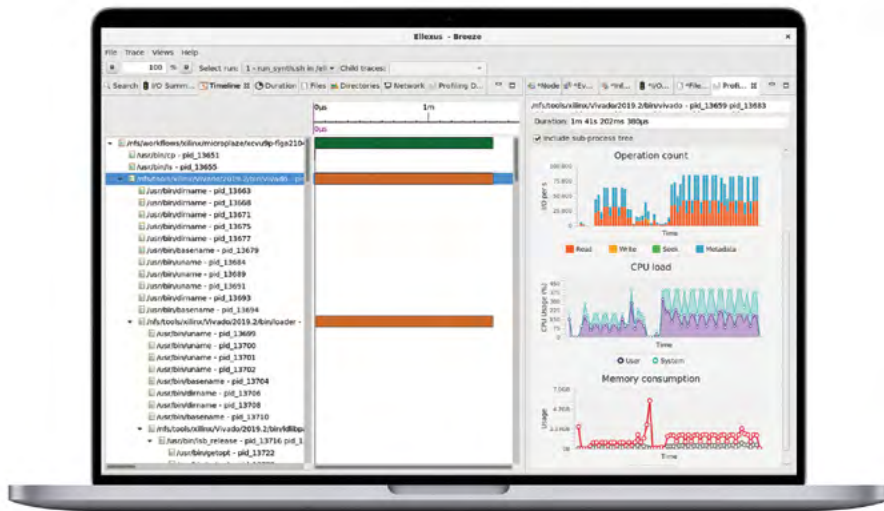
The Breeze I/O summary breaks down good and bad I/O. Users can easily identify bottlenecks and areas for optimization and can understand the impact of bad I/O patterns. Bad I/O patterns are things like small reads and writes, which harm the performance of shared storage, as well as failed I/O or opening a file that isn't used, which is always a waste of time.



Triage for System Administration: Troubleshooting Third-party Applications

“Why does this application work for you and not for me?” Deep dive into application arguments, environment, libraries, and dependencies to debug deployment and performance issues. Breeze provides everything you always wanted from strace, but with easily accessible graphical and machine-readable reports.

Although third-party applications can be hard to change, often the decisions about how to configure storage use is within your control. Placing data on the right volumes, removing legacy dependencies from workflows, and designing the right infrastructure for problem applications all become easy with Breeze.



Screenshot showing the Breeze UI with a timeline of the process tree on the left and I/O, CPU, and memory for the selected process shown on the right



Improving run time often doesn't require extensive rewrites. Knowing where to look is key.

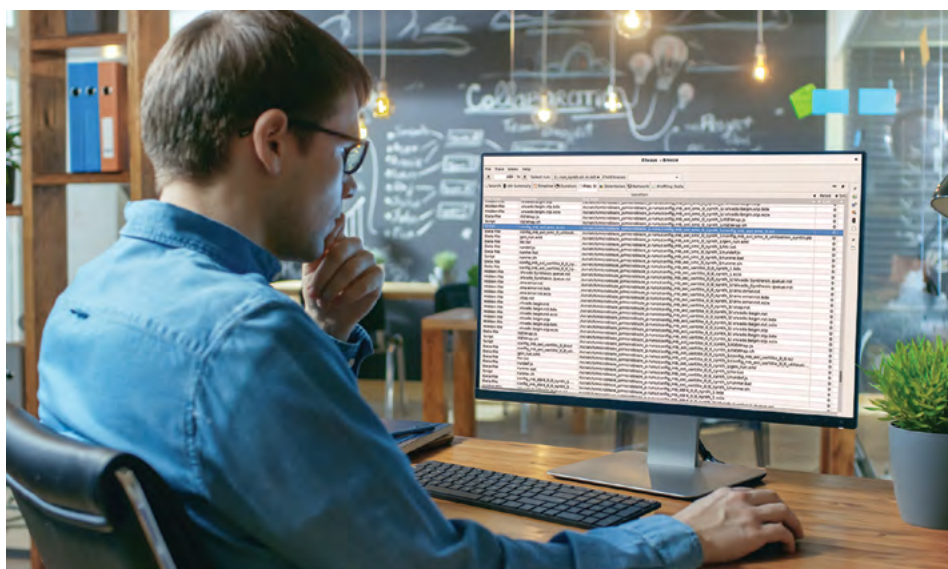
Keiran Raine, Cancer Researcher, Wellcome Trust Sanger Institute

Application Tuning and HPC User Education

“Why is my application slow today?” Automate performance and regression testing with detailed I/O profiling. Export files and dependencies for application correctness and migration. Breeze provides per-file and per-process I/O patterns so you can quickly dive in and see where the bottlenecks lie. Make every engineer an I/O expert with an I/O Healthcheck Report — a straightforward I/O overview and performance recommendations based on common I/O issues.

Automate Your Go-to-cloud Strategy With the Breeze Command-line APIs

Discover dependencies in legacy environments for the lift and shift. Containerize workloads, taking only the data you need. Profile application I/O for right-sizing CPU, memory, and storage resources to manage costs in the cloud.



Files view in Breeze showing every file, script, and library used by the application and its I/O patterns

Storage Benchmarking

It is vital to have accurate performance metrics across a variety of I/O patterns to troubleshoot storage issues and procure new solutions. Breeze gives you the information you need to see where bottlenecks lie. You can use real applications to test different storage solutions because Breeze will let you deep-dive into performance patterns. Breeze captures detailed filesystem performance metrics from within the application being profiled, helping you select right technology for your applications.

Technical Requirements

Breeze works out of the box on any Linux machine without modification to the machine or the application being profiled. There is no need to recompile the application or to install kernel modules. Breeze works in user space, collecting data on every process and every I/O operation.

Breeze uses an offline analysis tool so you can profile an application on one machine, then analyze the results on another. It comes with an interactive GUI with dashboards for understanding I/O patterns. You can also generate an HTML Healthcheck report as a summary and you can use the command-line APIs to automate dependency generation and benchmarking.

Distributed Workloads and MPI

For distributed applications, such as MPI jobs or workloads that have a master-worker architecture spanning multiple machines, Breeze will automatically profile the workload across each machine, generating a separate log for each host or job. Breeze understands job submissions to Altair® PBS Professional® and Altair® Grid Engine®, as well as to many third-party workload managers. Breeze can run on systems with or without a scheduler.

Containerized Workloads

For applications running in containers, Breeze has native support for Singularity containers. Docker containers and other container technologies can be profiled with Breeze with minor modifications to the launch configuration.



Breeze gives you the information you need to find bottlenecks