

UNICOS/mk and T3E status and Update

Jim Grindle

Silicon Graphics Inc.
Eagan, Minnesota, USA

June, 1998

ABSTRACT

This paper discusses the current hardware and software status of the Cray T3E and its operating system, UNICOS/mk. It covers recent hardware updates, Software Problem Report (SPR) status, reliability numbers, feature status and provides some recent benchmark performance information.

Introduction

This paper goes over the current (June, 1998) status of UNICOS/mk and the Cray T3E. It covers the recent hardware improvements, the UNICOS/mk software roadmap, including support plans and releases; and a number of software status issues including Software Problem Report (SPR) status, Mean Time To Interrupt (MTTI) numbers, and software feature status.

Questions regarding the content of this paper should be addressed to Jim Grindle, William White, or Steve Reinhardt.

Cray T3E Hardware Status

Recent Hardware Improvements

Recent hardware improvements include the Cray T3E 900, with a 450 MHz DEC Alpha based processor and the Cray T3E 1200, with a 600 MHz based DEC Alpha based processor. The most recent model improvement is the Cray T3E 1200E with an improved router performance.

Future Hardware Improvements

There are no further hardware improvements planned for the mainframe portion of the Cray T3E at this time. We continue to monitor the DEC EV-5.6 chip situation in the event that a vendor generates a viable 750 MHz version of this processor, but we do not see this at this time.

There are some planned improvements for the Gigaring I/O. These include an 18 Gigabyte Fiber Channel Node drive and an updated ESCON tape drive.

UNICOS/mk Roadmap

We plan to continue providing feature enhancements to the UNICOS/mk operating system through June of Calendar year 1999. Software Division Support for UNICOS/mk will continue through June of calendar year 2004. Subsequent to that, support would be provided through the Software Product Sup-

port organization.

UNICOS/mk Releases

Recent UNICOS/mk releases include 2.0.2 released in January, 1998 and UNICOS/mk 2.0.3 released in May, 1998. Future releases of UNICOS/mk will feature an extended release interval. We plan to go with a minimum 6-9 month release interval in the future. This interval could be extended to 12 months or more as the new feature content is reduced in the future. Current weekly archive updates could be altered to be bi-weekly instead of weekly archives if fix content allows.

SPR Plans and Priorities

This section on SPR plans and priorities is a reflection of our practice over the last year. It is an explanation of the priorities and practice we have been using over the past year and plan to continue into 1998 and 1999.

Priorities

Our top-most priority is Critical SPRs and Critical site situations. We will continue to deal with these and give them our utmost attention. The next priority will be urgent SPRs followed by verifying the severity of incoming SPRs to make sure the incoming urgents, majors, minors are classified correctly and don't need to be upgraded. Following that will be addressing the incoming major, minor, and design SPRs and the backlog of SPRs.

From viewing the SPR charts in the presentation, you can see that in UNICOS/mk we've been able to make some progress on the backlog even though there is still a reasonably high incoming rate. The only area that we have a consistent increase is in design SPRs.

Reliability/Resiliency

UNICOS/mk and the T3E have increased reliability by about 3 times in the last year. Software Mean Time to Interrupt(MTTI) went from 1000 hours in June of 1997 to 3059 hours in May of 1998. Hardware MTTI went from 600 hours to 1162 hours in the same time period. MTTI is measured through site reports using CRUISE ticket data. We are aware that we don't get 100% reporting with this method, but it is still useful in showing us the trends in MTTI.

Feature Content

We've made some big improvements in the last 6 months, especially in the area of resiliency. Here is a brief list of some of the biggest features, by release.

Recent Features in UNICOS/mk 2.0.2

With UNICOS/mk 2.0.2 we pretty much reached full UNICOS equivalence with features like parallel file systems, DMF, accounting, MLS, NFS, and Year 2000 verification. Other MPP specific features were remote mount which allows us to spread the file servers around the system and take that bottleneck out of the way. Pcache is an ldcache equivalent for the MPP which allows us to use local memory on certain OS PEs for caching data. The other big feature for 2.0.2. was the initial use of psched, allowing us to make better use of machine resources through political scheduling options.

Recent Features in UNICOS/mk 2.0.3

New features in UNICOS/mk 2.0.3 were improvements to swapping to make it more parallel and allow more swap configurations, the prime job feature which provides a mechanism to force a job from the GRM queue into the system overriding the application limits, and the express message queue which allows higher priority messages to be slated for an 'express queue' thus helping us spread the concept of priority off of a single PE. Up to this point, all messages were handled in first in, first out order by the OS PEs.

The other big improvement in UNICOS/mk 2.0.3 was in vastly improved application PE resiliency. This was started in UNICOS/mk 2.0.2 and completed in UNICOS/mk 2.0.3. This work provided much improved ability for the system to survive an application PE failure. Prior to this, a failure on one application PE could frequently cause the whole system to go down because of incomplete application clean-up.

Upcoming Feature Plans

Upcoming features include a Warm Boot capability which allows an administrator to reboot a single PE in the case of a software failure or intermittent hardware failure. This feature allows considerable improvement in the ability to keep a system up and not have to reboot just to bring one PE back in.

Another upcoming feature is the use of big page sizes. This will allow a specially compiled application to take advantage of a larger than normal page size which can improve the performance of certain codes.

Possibilities for the future that we are currently considering include improvements to the current scheduling capabilities, TotalView and PAT improvements, C++ and Fortran improvements and a feature called Persistent Objects. Persistent Objects will allow an administrator to configure an amount of memory on command PEs. to keep local copies of commonly used commands thus improving local performance and reducing the amount of usage of the root file system. Along with this feature comes an order of magnitude improvement in the averaged time for fork(2) and exec(2).

Other possibilities could include DCE/DFS, boot/dump speed-ups, and an initial investigation of system partitioning to allow hardware maintenance on one portion of the machine while running the system in a reduced configuration.

Summary

We are continuing to improve the reliability and resilience of the UNICOS/mk operating system and where we can we are making functional improvements that will benefit the largest numbers of customers. We have to be careful where we apply our resources, but we are continuing to make improvements that will benefit the majority of our customers.

We welcome your suggestions. You may reach the author by e-mail at:

jsg@cray.com

or by Post at:

Silicon Graphics Inc.
655F Lone Oak Dr.
Eagan, MN USA 55121