Cray Research Software Division Report

Mike Booth, Cray Research, Inc., 655E Lone Oak Drive, Eagan, Minnesota 55121

ABSTRACT: *This paper describes the new Software Division organization, the strategy for acquiring new technology and differentiating appropriately, and the status of the three technical software groups.*

1 Introduction

Cray continually advances its software technology with acquisitions and investments, deploying resources to differentiate and add value appropriately, while pushing old technology into the market place. As part of this strategy, Cray recently reorganized the Software Division along industry trends of more highly integrated software. This paper describes the new software organization, the development strategy, and the status of the projects in each of the new technical software groups.

2 New Software Organization

Cray reorganized the Software Division along industry trends for better integration with the software development strategy described below. In the old model of the organization, the Groups were aligned on separate functions: peripherals (disk, tape) & mass-storage, front-ends (stations) & networks, OS & OS specific file systems, compilers, and libraries & tools.

The new model emphasizes integrated functions. For example, all I/O uses common GigaRing nodes. The new model reflects the increasing overlap and interaction among disks, tapes, networks, file systems, mass-storage software, etc. It provides a structure to match Cray's integrated Programming Environments with compilers, libraries, programming models, and tools all integrated into one Group.

This reorganization produced three technical groups: System Software, Programming Environments, and I/O. It also produced two support groups: Common Services and Business & Product Strategies.

3 Software Development Strategy

Cray will acquire open technologies to meet standard market requirements. This is an appropriate strategy for undifferentiated features. Examples are HPF, NFS, and Kerberos.

Table 1. Software Division Organization

Group	Leader	Group Type
System Software	Dave Wallace	Technical
Programming Environments	Dave Judd	Technical
I/O	Dave Thompson	Technical
Common Services	Janet Lebens	Support
Business & Product Strategies	Dianna Crawford	Support

Cray will invest where Cray differentiates and adds value. Cray often encounters unique customer requirements before general market demands. When this happens, Cray implements advanced features ahead of the market. Examples are Cray's 64-bit OS and file systems, fast-file I/O, parallel programming methods, scalable OS, DMF, SFS, and NQE. This differentiation requires a constant turnover of technology.

Investing in new technology requires pushing old technology into the marketplace. One lets go of the old while reaching for the new, as if climbing a moving ladder that is constantly sinking into soft ground. Over time, Cray negotiates with industry and customers to transition to open standards and release technology to the market.

Cray begins a transition when the industry forms similar standards with similar functionality. This preserves Cray added performance while reducing costs. Examples are transitions from stations to FTA and NQE, HSX to HIPPI, and COS to UNICOS.

4 Operating Systems Group

4.1 UNICOS/mk Migration

To users, the UNICOS to UNICOS/mk migration is a much smaller visible change than was the COS to UNICOS transition. When Cray users moved from COS to UNICOS they saw major differences in the application program interface (API) and many feature changes that were a result of an almost complete code rewrite.

UNICOS to UNICOS/mk has only small visible changes. The CRAY T3E environment will be a comfortable and familiar UNICOS environment, but with more processors than possible with the vector-SMP (PVP) systems. The UNICOS API is preserved while adding the major feature of scalability. Almost all UNICOS features are carried forward, leaving behind only obsolete features, such as stations.

This commonality is evident in the OS development methods and the new common I/O technology. Most GigaRing features are developed first on vector-SMPs on UNICOS and then moved to UNICOS/mk on CRAY T3E systems. The common code among these two versions of UNICOS is high—over 80%, with most of the differences in the platform-specific low-level code. This eases the process of developing code in UNICOS and then moving it to UNICOS/mk.

4.2 UNICOS/mk Status

UNICOS/mk is running on CRAY T3D testbeds 24 hours a day with many users. It is running on CRAY T3E platforms for initial checkout of machine dependent code and GigaRing peripherals. Some applications run on multiple CRAY T3E PEs—for example CHARM and climate models.

Table 2. 1996 OS Releases

Level	Date	Features
UNICOS 9.0 on J90/EL	April 1996	9.0 features
UNICOS 9.2 on T90 and J90se	3Q96	GigaRing Support
UNICOS/mk 1.2 on T3E	April 1996	Initial multiple PE OS
UNICOS/mk 1.2.1 on T3E	July 1996	Additional I/O
UNICOS/mk 1.3 on T3E	Sept. 1996	Increased Functionality
UNICOS/mk 1.4 on T3E	Dec. 1996	Increased functionality and resiliency

5 Programming Environments

In 1996 Cray is supporting CF90 2.0 across its vector-SMP (PVP), CRAY T3E, and SPARC-SMP platforms while supporting Cray C/C++ 2.0 on vector-SMP and CRAY T3E platforms. Cray provides explicit programing-model support with the Message Passing Toolkit (PVM, MPI, get/put). For additional explicit-model support, Cray is evaluating the F--language. Cray continues to support the implicit programming models of AutoTasking® and HPF_CRAFT.

Table 3. 1996 Programming Environment Releases

Release	Date	Features
CF90 2.0 PE	Feb. 1996	CF77 Equivalence
C/C++ 2.0 PE	Feb. 1996	Native C++ compiler
PE 1.3 for T3D	Feb. 1996	CF77 6.2.2 Performance with CrayTools 2.0 and CrayLibs 2.0
CF90 1.0.4 for SPARC	Feb. 1996	Bugfixes
CF90 2.0 for SPARC	May 1996	Performance and PVP/MPP Compatibility
PE 2.0.1 for PVP and T3E	May 1996	PVP Bugfix and T3E support
PE 2.0.2 for PVP, T3E, and SPARC	4Q96	Bugfixes

Future directions include adding selected Fortran 95 and Fortran 2000 features. Cray will track PVM and MPI, championing the proposed MPI-2 single-sided communication (get/put) standard. Cray will track C++ and influence the explicit and implicit distributed-parallel programming models.

6 I/O Group

The I/O Group supports disks, tapes, network interfaces, and SSDs. This includes

- GigaRing on T3E, T90 and J90se
- Model-E on T90, C90, Y-MP
- VME on J90 and EL
- Consoles: SWS, OWS/MWS

The I/O Group also supports distributed computing, including network protocols, distributed file systems, NFS,

DFS, SFS, NQE, clustering, and MPT (PVM, MPI, SHMEM). Mass storage support is in the I/O Group: file systems, data migration, backup, archive, and tape subsystems.

Table 4. Message Passing Toolkit (MPT)

Release	Date	Features
1.0 for PVP	Feb. 1996	Optimized PVM, Initial MPI, SHMEM for PVP
1.1 for T3E & PVP	2Q96	Optimized MPI for PVP, Homogeneous PVM for T3E
1.2 for T3E & PVP	3Q96	Heterogeneous PVM for T3E, Homogeneous MPI for T3E

The I/O Group directions include support for the GigaRing scalable I/O on UNICOS and UNICOS/mk with higher performance and lower cost and easy incremental expansion. (With GigaRing I/O, host and channel bandwidths are rarely a limit.)

The I/O Group is developing and supporting fast network interfaces such as HIPPI, ATM OC-3, OC-12, and eventually OC-48. It supports distributed computing with open standards for high throughput capabilities, fast distributed file systems, message passing, and network job-queue load-balancing. This Group will also transition Cray's mass storage capabilities to open standards while preserving performance and capabilities. This includes new plans for backup, archive, data migration, and file systems.

7 Summary

Cray continues investing in Open Supercomputing with scalable and standards-compliant technologies. In 1996, Cray reorganized the Software Division along industry trends for better integration with the software development strategy described in this paper.

The three new technical groups in the Software Division (Programming Environments, System Software, and I/O) will deliver major new software releases in 1996 to support the new CRAY T3E and GigaRing I/O products.

This strategy will allow Cray to invest in new technologies where Cray differentiates its products while moving older technologies towards standards to lower costs.