the Availability Digest

www.availabilitydigest.com
@availabilitydig

The Dawn of Fault-Tolerant Computing

April 2016

In 1980, I published a four-part series in Computerworld entitled "Survivable Systems." The articles described the state-of-the-art fault-tolerant systems at the time. The need for systems that never (at least, hardly ever) failed was just being recognized. Several companies jumped in with their own versions of fault-tolerant systems, including Tandem, Stratus, Synapse, Auragen, August, NoHalt, Parallel Computers, and Tolerant Systems.

A lot has changed over the 36 years. Systems have become more "open," with Linux-like operating systems and x86-based hardware architectures. However, what hasn't changed is the need for systems that never fail. Applications that were hardly in use in the 1980s now are becoming mission-critical. The use of email is a perfect example. With the advent of social media, systems promoted as 24x7 can't risk a failure. As soon as a system is under distress, the Twitter universe explodes with complaints and comments, often causing irreparable harm to a company's reputation for reliability.

Some early products are still in use, for instance, Tandem and Stratus systems. Others have been incorporated into newer products. Still others simply have disappeared. In this article, we visit the dawn of fault-tolerant computers and the various architectures that were being promoted as such at the time.

Tandem Computers, Inc.

The Tandem computer was the granddaddy of fault-tolerant systems. Tandem's first system was delivered in 1976. Forty years later, its original architecture remains the dominant fault-tolerant technology. Then and now, a Tandem system was a loosely coupled multiprocessor system that contained anywhere from two to sixteen independent processors in a node. Up to 255 nodes could be included in a single network, linked via Tandem's Expand communication network.

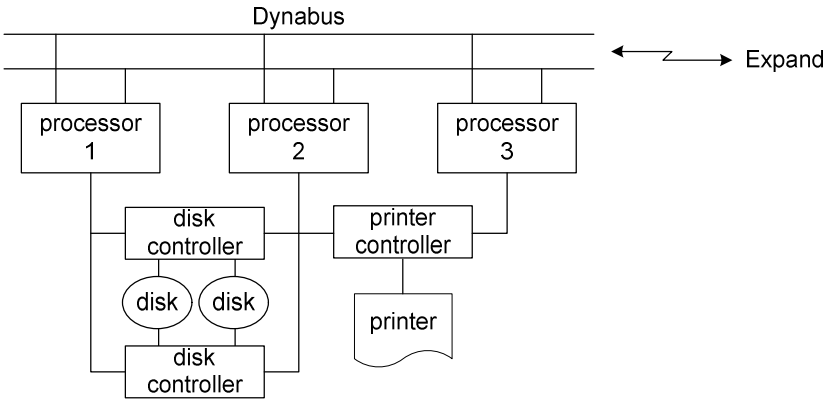

Tandem Computers

The processors in a node were linked via a duplexed, interprocessor messaging bus called the Dynabus, capable of a 26 megabyte/second data rate.

All device controllers were dual-ported so that there was always a path to a device even if a processor failed. All critical processes ran as process pairs in two different processors. One process was the primary process, and one was the backup process. The primary process kept its backup process synchronized via checkpointing messages. Should the primary process fail (presumably due to a processor failure), the backup process took over and continued processing with no apparent interruption to the user. (Tandem's later inclusion of the Pathway process monitor eliminated the need for application programmers to write checkpointed process pairs.)

With Tandem's second release of its product, the Tandem NS2, each processor could be configured with two megabytes of memory. Each mirrored disk pair could provide 128 megabytes of memory (yes, that's megabytes, not gigabytes).

Tandem was acquired by Compaq in 1997, which then was acquired by HP in 2002. Tandem computers are now known as HPE NonStop computers. (HPE is HP Enterprise, one of the two companies that resulted from the split in 2015 of Hewlett Packard into HP, Inc., which sells HP personal computers and printers, and HPE, which markets HP server, storage, and networking systems.)

Stratus Computer, Inc.

Founded in 1980, Stratus Computer, Inc. (now Stratus Technologies, Inc.[1]) marketed a system that was similar to the Tandem system in that it was an expandable multiprocessor system. However, the similarity ended there. Stratus achieved with closely-coupled hardware what Tandem achieved with loosely-coupled software.

Each Stratus processing module comprised four Motorola 68000 microprocessors. Introduced in 1967, the M68000 architecture is still in use today, almost four decades later.

A Stratus CPU board contained two such microprocessors running in lockstep. The outputs of the two microprocessors were continuously compared. As long as they behaved identically, their outputs drove a pair of high-speed buses that communicated with other modules in the system. However, should a difference be detected in the outputs of the two microprocessors, the board was shut down immediately. Thus, the CPU would not propagate any bad result.

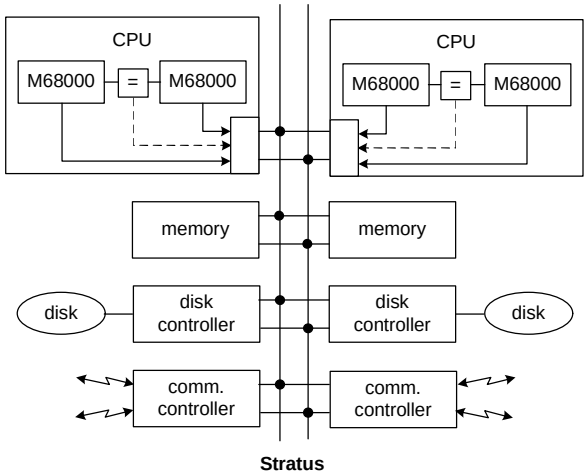

Stratus

[1] In 1998, Stratus Computer, Inc. was purchased by Ascend Communications for its communication products. The enterprise server portion of the business was spun off to the original founders in 1999. The new company was named Stratus Technologies, Inc.

To provide fault tolerance, two CPU boards drove the pair of buses. As long as both boards functioned properly, the buses were driven with the same signals from each board. However, if a fault was detected in one of the boards, the board would be shut down. The surviving board continued to provide the processing functions for the system.

This philosophy of hardware self-checking and redundancy was carried throughout the rest of the system via dual memories, dual disk controllers, and dual communication controllers. Multiple processing modules could be linked together via a simplexed or duplexed 2.8 megabyte/second interprocessor bus called the StrataLink. Each Stratus processor could be configured with four megabytes of memory and sixty megabytes of mirrored disk.

Over the years, Stratus' lockstep technology evolved to become its ftServer. Stratus also introduced a software solution, Avance, which used two general-purpose x86 servers that were kept synchronized with synchronous replication. Avance was replaced later by Marathon's everRun system when Stratus acquired Marathon in 2012.

In addition to these systems, in the early 1980s Stratus introduced the VOS operating system running on its high-performance Continuum hardware. VOS was a Multics-like operating system, and Continuum fault tolerance was achieved with dual processors running in lockstep. Stratus since has ported VOS to its ftServers. Interestingly, my company, The Sombers Group, was contracted by Tandem in the early 1980s to compare the performance of VOS with the Tandem TXP system. . Their performances were substantially identical. However, Stratus' focus then and now remains within the market for smaller x86-based solutions.

Synapse Computer Corporation

Synapse Computer Corporation took an approach strikingly different from that taken by Tandem and Stratus. Rather than providing virtually instant recovery from a failure, Synapse took the view that a short delay in recovery was acceptable provided the database was not corrupted. Synapse argued that it was perfectly acceptable to have the users wait a minute or two while the system recovered and then to require the users to reenter the transactions that were in progress at the time of failure. Back then, this was called a "mission critical solution." Can you imagine trying to sell it in 2016?

The Synapse architecture was a closely coupled configuration. Up to 28 general-purpose or I/O processors could be connected to a dual 32 megabyte/second bus. The processors communicated via a shared memory also connected to the dual buses. Thus, interprocessor communication was much faster than for the Tandem and Stratus systems, both of which relied on interprocess messaging.

Also, processes were not assigned to processors. Rather, all processors shared a common task queue. When a processor became idle, it began processing the next task on the queue. A highly efficient caching mechanism allowed processors to access data anywhere, even if it were in the cache of another processor.

The Synapse system was certainly not fault-tolerant in the sense of Tandem and Stratus. If a memory module failed, every executing process in the system could be affected. If a general-purpose processor failed, every process running in that processor failed.

Synapse depended upon a transaction model to recover from failure. Transactions were checkpointed as each new screen was presented to a user. If a failure occurred, the entire system was brought to a halt and was reconfigured around the failure. The database was recovered to ensure its consistency. Database recovery was accomplished via a History Log, which contained all before and after images of data changes made to the disks. Transactions then were restarted from their last checkpoints.

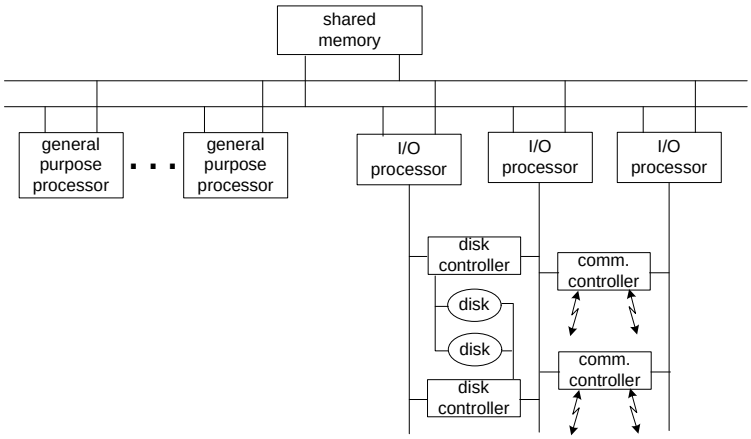

Synapse

A Synapse system used Motorola 68000 microprocessors. It could be configured with six megabytes of memory and 300 megabytes of disk storage.

The closely coupled architecture of Synapse systems never was accepted as a fault-tolerant solution, and Synapse systems no longer are sold.

Auragen Systems Corp.

Auragen's fault-tolerant offering was in many respects a hybrid combination of the loosely coupled architecture used by Tandem and the closely coupled architecture used by Stratus. It comprised from 2 to 32 processor clusters connected to a dual-system bus, each with a 16 megabyte/second capacity. Devices connected to the processor clusters via dual-ported controllers so that there was an access path to every device even in the event of a processor failure.

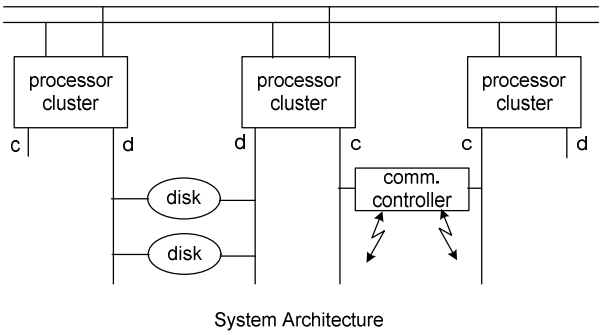

System Architecture

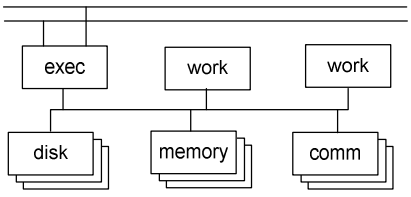

Processor Cluster

Auragen

A processor cluster was a closely coupled system comprising an executive processor and two work processors, all of which were Motorola 68000 microprocessors. The processors in a processor cluster were interconnected with shared memory and device controllers via a 20 megabyte/second bus. The

executive processor was responsible for all operating-system functions. The work processors were independent from each other and ran the applications.

The processor clusters were not fault-tolerant. Rather, fault tolerance was achieved in a manner similar to Tandem systems. Each process had a backup process running in another processor cluster, and the backup process was kept synchronized with the primary process via checkpointing. However, checkpointing was taken care of automatically by the processor clusters. All communication between processes was via interprocess messages. All interprocess messages directed to the primary process also were queued to the backup process. Periodically, the backup process would be synchronized with the primary process; and its queue of messages would be erased.

Should the active process fail, the backup process replayed all of the messages in its queue. During the replay, outgoing interprocess messages were inhibited to prevent duplicates. Following the completion of the interprocess-message replay, the backup process was ready to take over processing where the active process left off.

Auragen was involuntarily dissolved in 1985.

August Systems

The August system was unique in that it provided fault tolerance for both digital and analog signals. It comprised a loosely coupled, triplexed voting system using triple modular redundant (TMR) technology. Three independent paths were provided through the system, and their results were compared. If one result was different from the other results, it was discarded. The common result was passed on as the correct output.

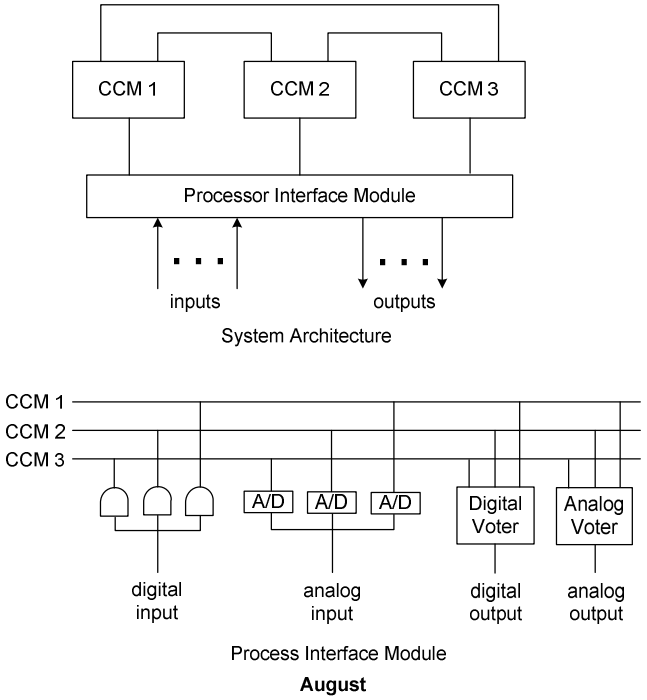

August

The August system comprised three 8086-based Control Computer Modules (CCMs) that communicated with each other over read-only interprocessor buses. Each CCM received inputs from the Processor Interface Module (PIM) and made its own calculations. The results were compared via the interprocessor buses. If a CCM found itself outvoted, it adjusted its results to comply with the other two modules. If a CCM found itself consistently outvoted, it declared itself out of service.

Each CCM returned its response to the PIM, which itself voted on the results and returned the majority result to the outside world. Analog inputs were digitized via analog-to-digital (A/D) converters. For analog output signals, the digitized values were reconverted to analog values; and the median value was returned as the result.

The PIM provided input and output redundancy for digital and analog signals. Input redundancy was provided by distributing input signals to all three CCMs. The output circuits were fault-tolerant. If any component opened or shorted, the voter still functioned properly.

Each CCM could be configured with 128 kilobytes of memory.

In 1997, August Systems was purchased by ABB, a Zurich, Switzerland-based company operating mainly in robotics and the power and automation technology areas. The August system is now known as the ABB Triguard TMR product. It has over 1,000 systems installed worldwide with over 10 million operational hours.

NoHalt Computers, Inc.

NoHalt Computers aimed at the low end of the market. Its system comprised a mirrored database with up to sixteen 8-bit Zilog Z80 work processors interconnected by a pair of 1.25 megabyte/sec. interprocessor buses.

Mirrored files were implemented via a pair of file processors that were independently powered and that managed their own disk units. They communicated with the work processors via the interprocessor buses.

Each work processor could support up to four peripheral devices such as terminals, modems, or printers. However, the work processors were not multitasking. They each could perform only one task at a time.

In the event of a file-processor failure, the surviving processor continued to support the system. In the event of a work-processor failure, all peripherals connected to that processor were out of service.

Each file or work processor could be configured with 64K of memory. The system could be configured with 140 megabytes of mirrored disk.

NoHalt Computers was acquired by TPC Logistics Services, Inc., and the NoHalt product was renamed the Reliant fault-tolerant computer system. In 1984, TPC announced the Reliant product had been enhanced to support 16-bit 8086 microprocessors.

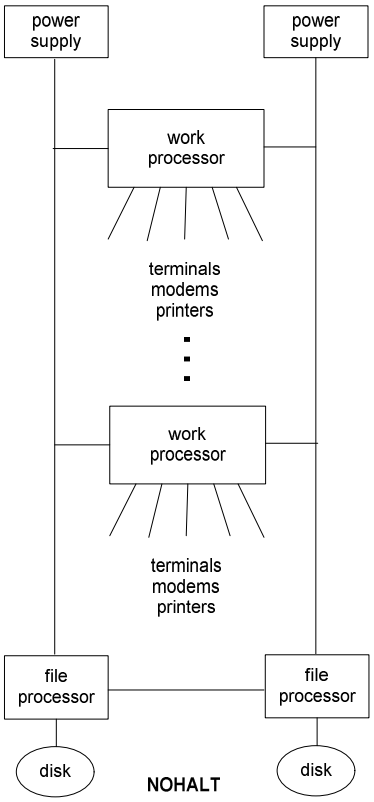

Parallel Computers, Inc.

Based on Motorola 68000 microprocessors, the Parallel Computer system could connect up to five parallel processing units (PPUs) to a simplexed interprocessor bus. Up to four disk controllers and two communication line controllers also could be connected to the bus to communicate with the parallel processing units. The disk controllers supported mirrored disk pairs.

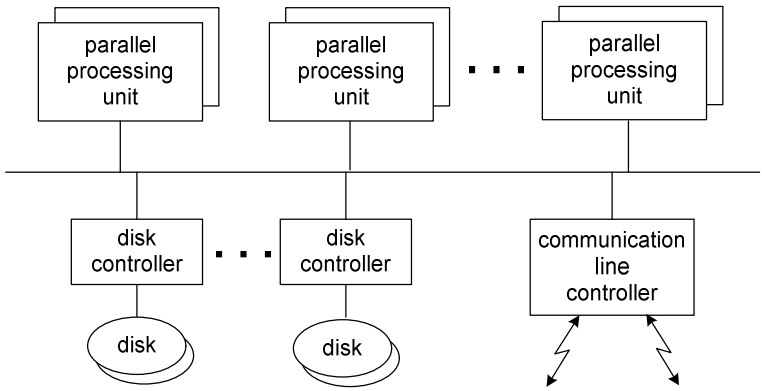

Parallel Computers

Each PPU was a duplexed fault-tolerant computer. A process in a PPU ran independently in each half. Periodically, the process pairs were synchronized. Thus, if one side of a PPU failed, the other side could carry on the processor functions in a manner transparent to the user.

Each PPU could contain one megabyte of memory.

In 1988, Parallel Computers was acquired by IMP, a British computer company. According to an IMP cofounder, IMP acquired Parallel Computers to gain entry into the U.S. computer marketplace.

Tolerant Systems

Tolerant Systems provided a series of System Building Blocks (SSBs) based on the National 32032 microprocessor. The SSBs could be arranged in a variety of ways by the user to achieve desired functions and levels of redundancy.

In a typical system, SSBs were used as communication servers to process requests from the users and to return responses, as application servers to process the transactions, and as file servers to access and update the database. In general, multiple servers of each type were available for redundancy.

Fault tolerance was based on transactions. The communication servers maintained a log of all incomplete transactions, and the file servers maintained the before images of all incomplete updates. Should a server fail, its load was transferred to a like surviving server. All transactions being handled by the failed server were aborted, and the corresponding partial database updates were rolled back via

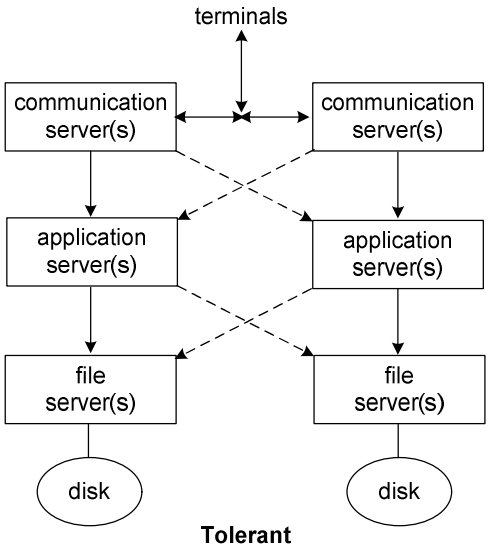

Tolerant

the before images maintained by the file servers. Those transactions then were replayed from the communication server logs and enabled full recovery transparent to the user except for a time delay.

Tolerant Systems was renamed Veritas Software Corporation in 1989. After developing several new applications, Veritas went public in 1993 at a value of USD $64 million. Veritas subsequently was acquired by Symantec Corporation in 2005. Symantic produces software for security, storage, backup and availability

Summary

Following Tandem Computer's successful entry into the fault-tolerant server field, numerous companies attempted to follow. Of them, Stratus was the most successful. Others were absorbed or no longer exist.

After the success of Tandem and Stratus, IBM introduced its Parallel Sysplex fault-tolerant system in 1994. Today, most large mission-critical systems are powered either by Tandem (now HPE NonStop) or by IBM. Mission-critical systems on the edge (such as in sales offices) are largely managed by Stratus.

Much of the fault-tolerant technology we discuss here may seem woefully outdated. However, forty years after it was introduced, the Tandem NonStop architecture remains the same. That is simply amazing. Back when Tandem and similar systems were first introduced, there was no Internet. There was no Big Data. Cloud computing? Huh? From that perspective, what should we infer will be the scope of IT four decades from today? Will fault tolerance remain relevant or even exist? What we do know for certain is that the past is an indicator of future innovation and evolution. As such, appreciating the dawn of fault-tolerant computing will shape our anticipation of what lies ahead.