

PORTUS™ High Availability

Charles Manick Livermore
October 22th, 2004



The PORTUS Application Protection System (APS) is designed to protect Mission Critical systems that have severe consequences for failure. PORTUS provides predictable high levels of security, reliability and survivability required for financial organizations, safety-critical domains such as aviation, healthcare, infrastructure, first-responders and national defense. High Availability is a difficult to achieve standard that requires a system to be available 99.999% of the time, which limits unscheduled outages to less than 6 minutes per year.

PORTUS pioneered High Availability security systems in the mid 1990s. Some customers have experienced more than 6 years of uptime without any unplanned outages due to hardware or software failure. PORTUS is in use at large International financial organizations where outage costs exceed \$250,000.00 per hour. Our customers employing High Availability solutions have never reported an outage.

Introduction

There are two aspects to PORTUS High Availability. First, PORTUS provides High Availability for the application servers it protects. Server availability is improved by blocking both known and previously unknown forms of attack using multiple algorithms. PORTUS also detects server failures and redirects transactions to available servers thus preventing application disruptions. Details regarding these protections are discussed in other documents. Second, **PORTUS provides High Availability for itself**. This paper is focused on how PORTUS keeps itself up and running more than 99.999% of the time.

The PORTUS High Availability option uses fault-tolerant software and hardware to provide “five nines” (99.999%) availability. PORTUS High Availability configurations are available that have demonstrated Mean Time Between Failures in excess of 200,000 hours (22 years) with Mean Time To Repair times of less than ten seconds. Various levels of support are available including 9x5 or 24x7 software support and 9x5 or 24x7 on-site hardware repair.

Fault-tolerant Software

PORTUS is unique in that it was designed from the ground up as a high availability APS. Since all hardware and software is subject to failure the key to High Availability is in the architecture of the product. Hardware and software redundancy is an essential but insufficient condition for High Availability. The design must do more than eliminate single points of failure. The software must provide real-time detection, isolation and recovery from errors. In addition it must provide immunity from hostile attacks to prevent catastrophic system wide failures.

Error Isolation: PORTUS was designed to isolate the effects of both hardware and software errors to prevent error propagation. With PORTUS each transaction is handled by a separate process and there is strict separation between processes. This means a failure handling one transaction is isolated and can not propagate to other transactions. In contrast all stateful packet

filter (SPF) firewalls and intrusion detection systems (IDS) run as part of the kernel permitting errors to propagate. These errors can cause catastrophic system failure or worse yet allow the SPF or IDS to fail wide open permitting malicious traffic to pass undetected.

Immunity: PORTUS transaction processes run at the application level in a chrooted directory. This means the process handling the transaction is unable to access or modify any system files. As a result, a skilled hacker can not exploit any possible errors in PORTUS to access or modify the programs or configuration files on the system. This makes PORTUS immune to attack even if it contains errors.

Granular scalable redundancy provides multiple copies of selected software components that can be non-disruptively replaced if damaged. Regeneration techniques recognize damage and automatically recover operational capability. These features allow it to survive massive and extremely hostile attacks.

Cognitive immunity: PORTUS can recognize attacks that have never been seen before by detecting protocol anomalies designed to attack systems. This makes PORTUS as well as systems it is protecting immune to new forms of attack. Protocol anomaly detection will also block attacks which remain to be invented.

Triple Level Software Recovery provides for non-disruptive service in the event of a software failure. A parent process spawns child processes which do all the work. These child processes run without root privileges in a chrooted directory. If there is a bug in the PORTUS code, then these child processes can not be used to access or modify any part of the trusted computing base (TCB). This is possible because these child processes cannot actually access any part of the TCB.

Each transaction runs as an independent process. If an error occurs the process tries to correct the error. If first level process recovery fails the error is confined to a single transaction. If a child process terminates for any reason, the parent processes performs second level of error recovery. The parent process dynamically detects and recovers the failed process. Process self-regeneration, similar to biological response strategies, prevents process depletion and loss of function.

Process **self-regeneration** is also used to prevent failures caused by slow system resource depletion. Child processes are automatically retied after 1000 transactions. When a process is retired, all resources it acquired during its lifetime are returned to the system. Thus even if there were a memory leak in the code, all the memory would be returned to the system. When a child process is retired, the parent will automatically re-generates a new child process using the same mechanism that was used to recover from an error. This allows PORTUS to run for years without rebooting.

At the third level, specialized system monitors can detect, report and recover from lower level failures. All three levels are available on a single PORTUS system.

Automated System Management routines keep the system running without operator intervention. Specialized monitors alert the system administrators of pending problems allowing preventative action to be taken before the problem causes a disruption in service.

PORTUS automatically rotates logs, compresses and archives them onto another system. This prevents the file systems from filling up. As a precaution, there is a disk monitor that monitors the disk space utilization. When a file system passes a configurable threshold, a notice is

automatically sent to the administrator. There are four threshold levels for each file system and as higher levels are passed, a notice of higher severity is sent. If a file system reaches 99% utilization, the proxies will automatically be throttled so that no new transactions will be served as the system would not be able to log all the transactions.

Automated garbage collection routines remove expired files preventing loss of disk space.

The Process manager automatically scans the process table to ensure the processes that should be running are, and ones that should not are not. Alerts are generated when there are either too few or too many processes running of a specific type. This enables preventative administrator action to be taken before a disruption occurs. The parent processes automatically ensure there are sufficient child processes to handle the work load. However, in the unlikely case the parent process should die the process monitor will alert the administrator enabling them to recover the parent before a disruption occurs.

Automated procedures can determine if a process has gone into a loop and automatically terminate the process, allowing other work to run without degradation.

Some models of the PORTUS appliance record all hardware and software errors in a system error log. Error reports are automatically generated that inform the systems administrator of temporary error conditions that can be corrected before they become permanent errors.

As a result, PORTUS is far more robust and resistant to hardware and software failures and hostile attacks than any of its competitors. Its self healing architecture permits it to recover from errors that would cause lesser systems to fail.

Fault Tolerant Hardware

PORTUS High Availability systems use multiple fault tolerant hardware features to provide granular scalable hardware redundancy. Multiple levels of redundancy provide non-disruptive operations in spite of one or more component failures. Standard features include redundant hot swap power supplies, redundant hot-swap cooling fans and disk drives. The system components most likely to fail can be swapped out and be replaced without disrupting the system. Some models offer fault tolerance that far exceeds the standard features found on competitive product offerings. The extended availability features include:

Dynamic processor deallocation varies a failing processor offline (SMP configurations only), moves the transaction to a working processor, marks the failed processor as inactive and issues an alert. This prevents lost transactions and system crashes as a result of a processor failure. Other systems, including SMPs, will crash when a processor fails.

Chipkill memory can detect and correct multi-bit errors in a single byte, while ECC memory can only correct single bit errors within a byte. A system with multiple gigabytes of ECC memory is as likely to suffer a memory failure as a system with 64 MB of non-ECC memory. Since Chipkill memory is 100 times more reliable than ECC memory it virtually eliminates system crashes due to memory failure.

Dynamic network adapter reconfiguration automatically enables a spare adapter to takeover from a failed adapter without disrupting service.

PORTUS was the first firewall/IPS to share its workload across multiple systems. In addition to providing higher throughput, the system provides higher availability. When two PORTUS

systems are configured with the high availability option each one monitors its sibling. When one fails the other dynamically assumes the failed systems workload.

High Availability using Redundant Systems

The PORTUS High Availability option consists of multiple PORTUS Application Protection Systems configured in groups of two. Pairs of systems can be added to the High Availability configuration allowing for non-disruptive growth in capacity. In the configured pair example, both systems share the workload while they monitor their sibling for errors. The PORTUS systems have similar configurations, and both are connected to the same networks. This allows each PORTUS system to automatically test its siblings connectivity using every NIC. If system A detects that system B is not functioning (frozen, crashed or one or more inactive NIC's) then system A can automatically take over all the traffic of system B.

PORTUS uses daemons that run on both systems to monitor the pulse of its sibling. The fwpulse daemon will negotiate the takeover if one of the systems is partially functional. Fwpulse performs an intelligent takeover, taking care not to do an unnecessary takeover due to a transient network condition. However once a takeover decision has been made the process occurs in the blink of an eye.

During this takeover process the remaining system assumes the IP addresses of its sibling, the proxies are updated and gratuitous ARPS are sent to the routers, allowing them to change the MAC address associated with the switched IP addresses. Pre-takeover and post-takeover scripts can be configured to customize the take-over process. This enables two asymmetric systems to back each other up.

Unlike SPF Firewalls the PORTUS High Availability configuration allows the two systems to share the workload. Since many hardware and software failures are more likely to occur when the system is under a heavy load, the fact that PORTUS can balance the workload across two systems greatly reduces the risk of a stress induced failure. SPF Firewalls that offer a HA option run the second system as a hot standby. In this configuration one system takes the entire load, while the other system simply waits for the first system to fail.

For more information on the technology leader in redundant fault tolerant security systems contact LSLI at:

Livermore Software Laboratories, Intl.
Division of Fremont Avenue Software
1830 S. Kirkwood Suite 205
Houston, TX 77077
Phone: 281-759-3274
Fax: 281-759-8558
e-mail: portus@lsli.com
Web: www.lsli.com