White Paper

Cisco IOS Software Modularity for the Cisco Catalyst 6500 Series and Cisco Nonstop Forwarding with Stateful Switchover

Cisco Nonstop Forwarding (NSF) with Stateful Switchover (SSO) leverages hardware redundancy and provides the software infrastructure to increase system-level availability. Cisco IOS® Software Modularity for the Cisco Catalyst® 6500 Series Switch delivers a new software infrastructure with support for process independence, process restart, and subsystem level in service software upgrades. These capabilities merge to improve system and network availability, while providing infrastructure for further development.

CISCO IOS SOFTWARE MODULARITY FOR THE CISCO CATALYST 6500 SERIES AND CISCO NSF WITH SSO TOGETHER IMPROVE HIGH AVAILABILITY

Cisco IOS Software is improving high availability by evolving from a statically modular system to a dynamically modular architecture that allows for runtime modification and dynamic software updates. Cisco IOS Software Modularity for the Cisco Catalyst 6500 Series switch ushers in significant new capabilities and integrates with previously available high availability features. This paper describes Cisco NSF with SSO functionality within the new modular software environment, the relationship of the various components, and the resulting advantages.

NO DOWNTIME

The challenge is to create networks and systems that provide access to applications, data, and content—anytime, anywhere. Of course, enterprises and service providers must temper that requirement based on individual business needs and the need to maximize ROI. Availability is an important consideration to every enterprise. Systems must be available or the benefit subsides.

Various redundancy techniques, fault detection methods, network designs with alternate paths, best practices, protocols, and operational procedures are all applied to ensure service availability by lowering Mean Time To Repair (MTTR) or increasing Mean Time Between Failure (MTBF). Using proven methods, most IT groups, network administrators, and staff are achieving higher and higher levels of availability. Network service is beginning to approach the vaunted 99.999% availability (See Table 1), also known as “five nines”. However, it remains difficult to achieve this high availability. Systems must be made more resilient and allow for better recovery from fault while minimizing the probability of failure if more organizations are to achieve service levels of 99.999% or better.

Table 1. General Availability Percentage Equation

The general formula for availability is: $$\text{Availability} = \frac{\text{MTBF}}{\text{MTBF} + \text{MTTR}} \times 100$$
The formula derives the percentage of the time service is available. The inverse is the amount of downtime and can derived using the formula: $$\text{Downtime} = \left(1 - \frac{\text{Availability}}{100}\right) \times 365 \times 24 \times 60$$
As you can see from the general availability formula, Availability is a function of the total service time, the mean time between failure (MTBF), and the mean time to repair (MTTR). Therefore, we endeavor to increase MTBF and reduce MTTR.

There seems to be no end to the requirements for new features, functionality, and performance. System enhancements continue to improve return on investment through increased value and improved price/performance. As that occurs, systems tend to increase in complexity. Software must therefore evolve to handle and contain faults that occur and recover from them without impact to service.

The challenge, therefore, is to create self-monitoring systems that can identify and contain faults, react automatically, and preserve service throughout an event.

CISCO CATALYST 6500 SERIES—LAYERED RESILIENCY

The Cisco Catalyst 6500 Series already offers multiple layers of redundancy. Redundant hardware components eliminate single points of failure and reduce MTTR. Components have embedded fault detection and error correction mechanisms to increase MTBF. Cisco IOS Software infrastructure and high availability features such as Cisco NSF with SSO leverage optional redundant control plane hardware and use data plane separation to keep packets flowing in case of failures.

As the software environment on the Cisco Catalyst 6500 Series evolves from a statically modular system to a dynamic one with support for process restart, coordination between hardware and software resiliency mechanisms becomes even more important. Not only does the system need to identify, contain, and recover from faults but also it must take the appropriate action automatically.

CISCO IOS SOFTWARE EVOLUTION

With the introduction of Cisco IOS Software Modularity for the Cisco Catalyst 6500 Series, several new infrastructure capabilities combine to increase system availability, including:

- Process independence
- Process restart
- Embedded event management

Previously, Cisco IOS Software was statically modular, meaning the functional software components were compiled and linked at build time and operated as tasks using a cooperative, run-to-completion model within a single address space. The system is redesigned by extracting several processes, which now operate cooperatively but independently. A microkernel is employed to support these POSIX-style processes. The microkernel manages common resources and provides preemptive scheduling. The processes run in their own protected address space.

System monitor functions catch any process faults and can automatically restart individual processes. This adds another level of resiliency beyond the control plane switchover function provided by SSO. Previously, any process failure resulted in a switchover to the redundant control plane processor in a redundant system. Now individual processes may be restarted. The Cisco IOS Embedded Event Manager (EEM) can further customize the failure policy to use for various event conditions.

EXISTING HIGH AVAILABILITY OPERATION PRESERVED

The renewed software infrastructure, which continues to support existing Cisco IOS Software features for enterprises, will retain support for—and capitalizes on—the high availability infrastructure. For example, the current implementation of all SSO infrastructure components and their clients does not change with the introduction of Cisco IOS Software modularity.

For additional information about High Availability, including Cisco NSF with SSO, visit:
http://www.cisco.com/go/availability/

Various software components or features that are HA-aware (also known as SSO clients) use the HA infrastructure to preserve state information and synchronize it with the redundant supervisor. If a fault occurs, the system can immediately switch to the standby supervisor with little or no impact to the traffic flowing through the system.

Preserving the existing function was a key design requirement. Building on the new capabilities is the next step. The HA infrastructure will provide several services that offer additional advantage in the renewed environment.

INFRASTRUCTURE ASSIST

Routing protocols now run within a separate process with Cisco IOS Software modularity for the Cisco Catalyst 6500 Series. As such, the routing protocols support process restart and can be restarted. After a restart, the routing protocols leverage Nonstop Forwarding (NSF) enhancements to re-acquire routing information from their neighbors.

Note: IS-IS is unique in that it is also an SSO client and can be configured for "nsf cisco" to synchronize its state information.

Optimally, a process will restart in the same place where it left off prior to the fault. To make this happen, processes are assisted by the HA infrastructure. A process may **checkpoint** dynamic state information while it executes. If it restarts, it can use this information to reconnect with the other system components faster.

Consider what happens if an SSO client needs to support both process restart and supervisor redundancy simultaneously. SSO and processes that can be restarted both use infrastructure services to save their dynamic state. The dynamic state is received by the Checkpoint Facility and preserved for process restart and for supervisor redundancy.

Prior to the introduction of Cisco IOS Software modularity, a failure in a Cisco IOS Software task would cause failure in the entire process. A failure on the active supervisor would invoke an SSO switchover to the Hot Standby supervisor. Now, with Cisco IOS Software modularity, a failure in a modular process can be restarted. It does not have to result in an SSO switchover since the fault is isolated and contained.

Process Synchronization and Restart

A process that has restarted may also need to re-synchronize dynamic state with the peer process on the Standby supervisor. Each SSO client process must support restart synchronization on the Standby; otherwise, a process restart on the Active supervisor will result in a reset of the Standby supervisor.

The first release of Cisco IOS Software Modularity for the Cisco Catalyst 6500 Series will have two SSO client software features: IS-IS (Cisco NSF mode) and VRF. A single routing process contains both of these features, along with other routing protocols that regain state information with the help of neighbors via NSF. Zero packet loss is possible, even if the routing process is restarted when the routing protocols that have been configured to take advantage of NSF. Support for process state synchronization and restart synchronization makes zero packet loss possible for protocols that maintain dynamic state such as when IS-IS nsf cisco is configured.

Note: Full support for stateful process synchronization will be available in a future release of Cisco IOS Software Modularity for the Cisco Catalyst 6500 Series. Cisco IOS Software Release 12.2(18)SXF2 supports routing protocol process restart for NSF only. This includes BGP NSF (or Graceful Restart), OSPF NSF, EIGRP NSF, and IS-IS NSF (IETF option).

Modular Process Restart Behavior

The System Manager software component is responsible for restarting failed processes according to the applicable parameters.

There are two cases when a process restart will not be performed:

1. The process is marked as non-restartable.
2. The process fails repeatedly after being restarted. The System Manager "gives up" when it reaches the "Process Restart Count."

When a mandatory, non-restartable process fails, the System Manager resets the Supervisor or the RP. This action will cause a system reload on a single RP system or a switchover in a redundant system with two Supervisors or RPs.

When the System Manager declares a restartable process “dead” it may be restarted and an associated failure counter is incremented. If the process continues to run for more than four minutes, it is declared to have successfully restarted and the failure counter is reset to zero. If the process crashes again in less than four minutes, then the failure counter is again incremented. If the process is mandatory and the failure counter reaches the defined process restart count, the process is declared dead. Otherwise, the process is restarted. This provides protection from repeated failures.

Note: The process restart behavior for the initial release of Cisco IOS Software Modularity for the Cisco Catalyst 6500 Series takes a conservative approach. If a process crashes a second time within four minutes, the Supervisor is reset. All processes are mandatory and the process restart count parameter is two for all processes.

Process status and information can be viewed using show process CPU and show process CPU detail CLI commands.

```
snow#sh proc cpu | inc iprouting
24619    0%      0%       0% iprouting.iosproc
snow#
snow#sh proc det 24619
                  Job Id: 71
                     PID: 24619
         Executable name: iprouting.iosproc
         Executable Path: sbin/iprouting.iosproc
             Instance ID: 1
                 Respawn: ON
           Respawn count: 1
Respawn since last patch: 1
  Max. spawns per minute: 30
            Last started: Wed Nov  2 22:03:41 2005
           Process state: Run
            Feature name: iprouting
                    Core: SHAREDMEM MAINMEM
               Max. core: 0
                   Level: 100
               Mandatory: ON
     Last restart userid:
       Related Processes:
PID    TID  Stack pri state       Blked  HR:MM:SS:MSEC FLAGS    NAME
24619  1      32K  10 Receive     1       0:00:00:0000 00000000 iprouting.iospro
24619  2      32K  10 Receive     1       0:00:00:0000 00000000 iprouting.iospro
24619  3      32K  10 Receive     1       0:00:00:0680 00000000 iprouting.iospro
24619  4      32K  11 Nanosleep           0:00:00:0000 00000000 iprouting.iospro
24619  5      32K  10 Receive     1       0:00:00:0304 00000000 iprouting.iospro
24619  6      32K  10 Receive     1       0:00:00:0104 00000000 iprouting.iospro
----------------------------------------------------------------
```

The command output above shows the *iprouting* process. It is a mandatory, restartable process that has been restarted one time.

Process Restart and Customizing HA Policy Using EEM

Systems with redundant control plane hardware—two RPs or Supervisors operating in SSO mode—have multiple levels of redundancy for restartable processes. The default behavior for a process crash on the Active RP is to restart the process if it is restartable. The default behavior can be modified with the help of the Embedded Event Manager and the System Manager Event Detector. The System Manager Event Detector "Process Abnormal Termination" event will invoke the aforementioned default behavior unless modified specifically. An EEM action script can override the default behavior and take other action like invoke an SSO switchover.

The script would function as follows:

- **Trigger event**: System Manager Process Abnormal Termination
- **Process to monitor**: Routing Process (could be a different one)
- **Action**: Run the CLI command redundancy force switch-over
- **Return code**: skip default action

For more information on the Cisco IOS Embedded Event Manager, visit:
http://www.cisco.com/go/availability

The System Manager Process Abnormal Termination Event is a synchronous event by default. This means that the return code of zero following the EEM applet or script indicates to the System Manager to skip the default action.

Example 1 illustrates an EEM applet that is invoked when a process restart happens, takes an action, and allows the System Manager to continue with the default action to restart the process.

Example **1**: EEM Applet to Preserve Default Action

```
event manager applet OVR_PR_1
 event process abort path "iprouting"
 action 1.0 syslog msg "OVR PR 1 applet invoked"
 set 9.0 _exit_status 1
```

Example 2 illustrates an EEM applet to override the default action. Here the $_exit_status is set to 0 (the set could also have been omitted).

Example 2: EEM Applet to Override Default Action

```
event manager applet OVR_PR_2
 event process abort path "iprouting"
 action 1.0 syslog msg "OVR PR 2 applet invoked"
 set 9.0 _exit_status 0
```

Note: There is no real reason to implement this type of policy. Example 2 is used for illustrative purpose only.

Example 3 illustrates an EEM applet to override the default action and force a switchover.

Example 3: EEM Applet to Force A Switchover

```
event manager applet OVR_PR_3
 event process abort path "iprouting"
 action 1.0 syslog msg "OVR PR 3 applet invoked,
forcing switchover"
 action 2.0
 set 9.0 _exit_status 0
```

Example 3 illustrates the policy action for customized high availability. Cisco IOS Embedded Event Manager provides a unique capability to automate actions and define the behavior based on events detected by the Cisco IOS Software. This does not require an outside network management system.

CONCLUSION

Cisco IOS Software is improving **high availability** by evolving from a statically modular system to a dynamically modular architecture that allows for runtime modification and dynamic software updates. Cisco IOS Software Modularity for the Cisco Catalyst 6500 Series switch ushers in significant new capabilities and integrates with previously available high availability features, thereby providing a more comprehensive and granular high availability solution. Multiple levels of resiliency are now available to improve the availability of networks and access to applications, data, and content—anywhere, anytime.

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on **the Cisco Website at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

205393.BK_ETMG_SH_12.05

Printed in the USA