

Cisco ASR 1000 Series Aggregation Services Routers: ISSU Deployment Guide and Case Study

In most networks, a significant cause of downtime is planned maintenance and software upgrades. The Cisco ASR 1000 Series Router offers In Service Software Upgrade (ISSU) capabilities to allow a user to upgrade Cisco IOS XE Software while the system remains in service. This in service upgrade dramatically reduces downtime due to planned software upgrades and thereby increases service availability. (Note: One of the key original and ongoing design principles with the Cisco ASR 1000 is the criticality of ISSU and high availability.)

This white paper details the In Service Software Upgrade (ISSU) capabilities of the Cisco ASR 1000 Series Routers. This paper further examines the WAN Edge design and deployment options that can be employed to minimize service interruption during an in service upgrade or downgrade. This paper will step through these options with particular reference to a typical WAN aggregation topology illustrated in figure 1.

Cisco ASR 1000 Series Software and ISSU Introduction

In order to truly understand the ISSU capabilities of the Cisco ASR 1000, it is important to provide a high-level overview of the Cisco ASR 1000 series and the various ISSU options. For the case study and topologies examined in this paper only the Cisco ASR 1006 will be covered. However, for reference, all platforms are reviewed.

The Cisco ASR 1000 Series Router family consists of three different models:

- The Cisco ASR 1002 Router is a 3-SPA, 2-rack-unit (RU) chassis with one Embedded Services Processor (ESP) slot that comes with an integrated Router Processor (RP), integrated Cisco ASR 1000 Series Shared Port Adapter Interface Processor (SIP), and integrated four Gigabit Ethernet ports.
- The Cisco ASR 1004 Router is an 8-SPA, 4-RU chassis with one ESP slot, one RP slot, and two SIP slots.
- The Cisco ASR 1006 Router is a 12-SPA, 6-RU, hardware redundant chassis with two ESP slots, two RP slots and three SIP slots.

For the single Route Processor Cisco ASR 1000 platforms, the Cisco ASR 1002 and Cisco ASR 1004, the Route Processor has a dual Cisco IOS Software option that allows these routers to use Cisco IOS software redundancy. For more details on this capability please see the ASR homepage at http://www.cisco.com/go/asr1000

The Cisco ASR 1006 Router supports fully redundant Route Processors that allow for full Route Processor hardware redundancy and control plane redundancy using Cisco high-availability features, Nonstop Forwarding with Stateful Switchover (NSF/SSO), Graceful Restart (GR), and In Service Software Upgrades (ISSU). The Cisco ASR 1006 also has the capability to support two ESP's for data plane redundancy. This ESP data plane redundancy enables a sub 50msec forwarding interruption during either an HA failover or an ISSU procedure.

The Cisco ASR 1000 Series Routers utilize Cisco IOS XE Software and introduce a distributed software architecture that moves many operating system responsibilities out of the IOS process. In this architecture, Cisco IOS, which previously was responsible for almost all of the internal software processes, now runs as one of many Cisco IOS XE processes while allowing other Cisco IOS XE processes to share responsibilities for running the router.

The In Service Software Upgrade (ISSU) process allows software to be updated or otherwise modified while packet forwarding continues with minimal interruption.

For the Cisco ASR 1000 Series Routers, it is important to realize that ISSU-compatibility depends on the software sub-package being upgraded and the hardware configuration. This paper will consider only a fully redundantly configured Cisco ASR 1006 such that both the control plane (RP) and data plane (ESP) can be upgraded with sub 50msec of traffic interruption. Importantly, the SPA and SIP software sub-packages must be upgraded on a per-SPA or per-SIP basis, during any SIP/SPA upgrade traffic on the SIP or SPA being upgraded will be "out of service" during the upgrade process. This is an important note, in a Cisco ASR 1006 it is the SIP/SPA upgrade and network topology and design that determines the service downtime. This will be discussed further in this paper.

When updating multiple sub-packages, it is important to understand that the sequence of the upgrade has an impact on router operation during the software upgrade. The specific procedure in this document represent supported and tested installation sequences for a sub-package mode fully redundant Cisco ASR 1006.

Cisco ASR 1000 Series Software ISSU Architecture

The Cisco ASR 1000 Series Routers can be run using a complete consolidated package or using individual sub-packages, in this case study the routers are configured to use sub-packages.

Each Cisco IOS XE consolidated package contains a collection of software sub-packages. Each software sub-package is an individual software file that controls a different element or elements of the Cisco ASR 1000 Series Router.

The In Service Software Upgrade (ISSU) process allows software to be updated or otherwise modified while packet forwarding continues with minimal interruption.

The Cisco ASR 1000 Series Routers uses the Cisco IOS ISSU infrastructure in order to provide customers with software upgrade procedure that can result in minimal service interruption.

Table 1 provides a list of the Cisco ASR 1000 Series Routers sub-packages and how they relate each of the Cisco ASR 1000 models and ISSU. (Note: It is critical to understand what components can potentially impact networking services when being upgraded/downgraded. The topologies utilized have a direct impact on the availability and operational uptime of the network when performing an ISSU or when performing HA functions.)

The data in table 1 is examined in detail in the WAN aggregation case study and topology design/deployment section in this paper.

Software Sub-Packages	Cisco ASR 1002 / Cisco ASR 1004	Cisco ASR 1006
RPBase	This contains the underlying Linux kernel so can not be upgraded "in service" Requires reboot	The standby RP in the Cisco ASR 1006 chassis may be upgraded and then switched over to active mode "in service" Requires RP switchover; No transit packet loss
RPControl	Can be upgraded "in service" No transit packet loss	Can be upgraded "in service" on both active RP and standby RP No transit packet loss
RPAccess	Can be upgraded "in service" No transit packet loss	Can be upgraded "in service" on both active RP and standby RP No transit packet loss
RPIOS	Can be upgraded "in service" if the system is running in dual IOS mode Requires IOS processes switchover; No transit packet loss	Can be upgraded on standby RP and switched over o active in service Requires RP (IOS) switchover; No transit packet loss
ESPBase	Upgrade causes complete loss of local state (i.e. statistics, stateful FW/NAT) on the ESP and in service affecting Forwarding Interruption until upgrade is completed. Router is still accessible. No reboot required	Upgrade causes complete loss of local state (i.e. statistics, stateful FW/NAT) on the ESP being upgraded and will result in small traffic interruption when switching to the standby ESP Minimal transit packet interruption < 50ms; no RP switchover required
SIPBase	Upgrades cause complete loss of local state on the affected SIP, however the other SIP in Cisco ASR 1004 Router is unaffected by this activity. Hitless for the other SIP in Cisco ASR 1004 Router not being upgraded	Upgrades cause complete loss of local state on the affected SIP, however other SIPs are unaffected by this activity. SIPBase upgrades only take place if initiated from the active RP. Hitless for other SIPs not being upgraded
SIPSPA	Upgrades cause specific SPA to completely reboot and is service affecting. SIPSPA can be upgraded on per SPA basis. Hitless for other SPAs not being upgraded	Upgrades cause specific SPA to completely reboot and is service affecting. SIPSPA can be upgraded on per SPA basis. SIPSPA upgrades only take place if initiated from the active RP. Hitless for other SPAs not being upgraded

Table 1	This Table Details t	he Effect of Each	Package Ingrade
			rackage opgrade

ISSU Upgrade Procedures for Cisco ASR 1006 with dual RPs and ESPs—subpackage

The following section details the procedure followed to upgrade the Cisco ASR 1006. This is the procedure used in the case study upgrade referenced throughout this paper. A detailed explanation of this procedure and other approved ISSU procedures can be found at the following: http://www.cisco.com/en/US/docs/routers/asr1000/configuration/guide/chassis/issu.html#wp1081374

The goal of this install sequence is to use IOS XE ISSU to facilitate upgrading the hardware elements one-by-one to optimize system availability at each stage of the process. Both the RP and ESP will failover to their standby components during this software upgrade. To reiterate, the RP failover will not cause dataplane packet loss and the ESP failover will have a sub 50ms failover.

The standby RP is upgraded, reloaded and allowed to reach SSO ready state. This RP upgrade and reload has no effect on the ESPs, SIPs or SPAs in the system. (Note: While the standby RP is upgraded/reloaded it is out of service. If the active RP fails during this window then a stateless failover does not occur.) The SIP and SPA sub-packages are upgraded a slot at a time to minimize interface flapping. Slot-by-slot upgrades of the ESPs are performed allowing the retention of data plane state such as NAT and stateful firewall flows. Finally, the active RP is provisioned with the new software resulting in a stateful failover to the standby RP. By upgrading all non-RP entities before the active RP, no reload of those entities takes place on a RP switchover. The system is running all of the software that is provisioned for all live elements as soon as the switchover is initiated, the former active is still running old software however it is in the process of restarting and will come up with the software provisioned in the last step.

Note: Utilizing Gigabit Ether Channel (GEC) or other mechanisms allow interface level redundancy. Specifically, arranging SPAs so that they reside in separate SIPs allows traffic to flow uninterrupted during the installation of the SIPBase and SIPSPA sub-packages and subsequent restart of the SIPs. This is further discussed in the case study below

ISSU Upgrade Procedure Steps

1. Install all sub-packages on the standby RP

a. Using the IOS ISSU command "issu loadversion rp 1 file stbybootflash:asr1000rp*02.02.01.122-33.XNB1.pkg" to provision the software on the standby slot. The 'force' option may be needed

- 2. Use the "hw-module slot r1 reload" command to reload the standby RP
- 3. Wait for the standby RP to reload and the system reach SSO ready state
- 4. On the active RP install the SIPBase and SIPSPA sub-packages on each SIP

The following steps describe how to perform the installation on a per-slot basis. The 'slot' argument may be omitted to install all SIPs and SPAs in one step. As a SIP is upgraded, it resets. This installation phase upgrades the SIPs+SPAs on the active RP and takes effect immediately. The standby RP has already been fully provisioned with the new version and reloaded.

a. If performing a slot-by-slot installation, repeat the above steps for all present SIPs updating the "slot" parameter for each. Again, note that the provisioning is being performed on the active RP.

5. Install the ESP sub-package on the ESPs

The ESPs may be installed individually or manually, per slot. It is recommended that the installation be performed manually for maximum control. This installation phase upgrades the ESPs on the active RP and takes effect immediately. The standby RP has already been fully provisioned with the new version and reloaded. In the steps below, F0 is assume to be the active ESP and F1 is assumed to be the standby ESP

- 6. Install the RPBase sub-package and all remaining sub-packages. The RPBase sub-package, which will be matched by the command below, will automatically set the on-reboot flag but it may be specified manually. This step will not take effect until the next reload. BY using a wildcard pattern, we help ensure that any and all sub-packages that have not been provisioned, including SIP slots other than those manually upgraded, are set to use the new software. This is an important step. Use "file bootflash:asr1000rp*02.02.01.122-33.XNB1.pkg" to provision the software. The 'force' option maybe needed. The 'on-reboot' flag is set automatically.
- Use the "redundancy force—switchover" command to perform an RP switchover and trigger a reload of R0

Cisco ASR 1000 Series Software ISSU Upgrade Case Study

WAN Aggregation Topology

Figure 1. Logical WAN Edge Topology Used to Profile ISSU



In this case study the Cisco ASR 1006 is acting as a WAN aggregation router, with 200 OSPF peerings out to the WAN and five core OSPF routing adjacencies to the Core. The topology represents the typical Enterprise WAN aggregation topology with common routing services configured; such as dynamic routing protocol (OSPF), routing protocol authentication, Netflow, uRPF, BFD, IPv4 multicast, ACL's and QoS per sub-interface (i.e. per site).

Figure 1 shows the services, number of IGP routes and the number of routing adjacencies maintained during an ISSU procedure. There are a large number of routes and adjacencies in this topology to "stress test" the control plane high availability features, NSF/SSO, during the RP upgrade. This helps ensure that no issue with the testing can be masked, as all control plane graceful restart operations must complete successfully to meet the expected hitless RP upgrade goal. In this scenario the forwarding plane (FIB) is frozen during the NSF restart interval. The process to upgrade the 6RU platform from IOS XE release 2.1.2 to IOS XE release 2.2.1 was as detailed in the section: **ISSU Upgrade Procedures for Cisco ASR 1006 with dual RPs and ESPs.**

This upgrade proved the control plane HA with zero packet loss for both unicast routing and IPv4 multicast, while the ESP ISSU upgrade (failover) resulted in sub 50 millisecond traffic loss. The physical topology options and expected downtimes are detailed below:

 Table 2.
 Service Downtime During Upgrade—No SIP Redundancy

Component Upgraded	Service Downtime
RP Upgrade	0
ESP Upgrade	39 microseconds
SIP Upgrade (upgraded together)	~450 seconds

Figure 2. Physical Topology (20 router peers/GigE) No SIP Redundancy



Figure 3. Physical Topology (20 router peers/GigE) With SIP Redundancy



Table 2 shows the SIP/SPA upgrade element in the non-redundant SIP configuration is the element which impacts the service uptime in this non-redundant SIP topology. The cause of this service downtime is that all routing adjacencies are lost on that SIP/SPA during the SIP/SPA upgrade. This topology, as shown in figure 2, has a large number of peers and routes. Since all routes need to be installed in the RIB and the FIB the loss is in the order of minutes.

It must be noted that one of the goals of ISSU is to provide a staged and controlled upgrade of the system that may or may not be hitless but is completed while the platform is "in service". Even during a SIP upgrade, the operator is still connected to the RP via either the Ethernet Mgmt interface or via a non-upgrading SPA interface. The ability to roll back at any stage is also of great value as no longer is a whole system reload required to upgrade the router, as is the case with traditional Cisco IOS routing platforms,.

To significantly reduce the realized service outage during a SIP/SPA upgrade the Cisco ASR 1000 could employ Gigabit Ether Channel as detailed in figure 3. Instead of having the routing peers configured on sub-interfaces across the Gigabit interfaces these can now be configured as sub-interfaces on a port-channel, which has two Ten Gigabit Ethernet interfaces configured as group members The connecting switch now has a layer two port-channel configured allowing appropriate VLAN on either side of the Cisco ASR 1000. This enables the Cisco ASR 1000 to still peer directly with the tester in figure 3, but now has link redundancy via the port-channel. If the ISSU procedure is now repeated and during the SIP/SPA upgrade the traffic will re-route over the remaining port-channel link serviced from the still active SIP. Note: GEC over a WAN link is a service that has been offered by a few service providers world wide to enable the redundancy detailed in the figure 3 scenario.

Figure 4 takes the resilient infrastructure approach one step further and integrates the portchannel into a Catalyst 6500 virtual switching (VSS) environment. The service impact of a complete system upgrade, or downgrade now with the port-channel in place is now sub-second, as the staged SIP upgrade is hitless and the OSPF peerings now do not drop and thus the routing need not re-converge. With this resiliency in place the only realized service impact is the 39usec interruption on ESP failover and additionally the packets that were lost because they were on the SIP at the time of the SIP being upgraded. The total service interruption in this case is now in the order of tens of milliseconds.



Figure 4. Physical Topology (200 Router Peers/TenGigE) SIP Redundancy and Switching VSS

 Table 3.
 Service Downtime During Upgrade—With SIP Redundancy

Component Upgraded	Service Downtime	
RP Upgrade	0	
ESP Upgrade	39 microseconds	
SIP Upgrade (in a staged approach)	<100msec	

Understanding the ISSU procedure and system operation as detailed above allows one to translate this test case into a real world topology already employed by customers using the Cisco ASR 1000. The following topology diagrams outline some of the options allowing a full system upgrade including SIP/SPA's while maintaining all networking services. Table 3 showcases the ISSU service downtime for both figure 3 as well as figure 4. When careful consideration is made for the entire network topology and ISSU capabilities of the Cisco ASR 1000, very minimal downtime can be achieved while maintaining constant service availability.





Figure 5 shows one deployment option where there are two equal cost routes to the remote site more suited to a campus/Metro topology—such that one can rely on equal cost multi-path routing (ECMP) to provide SIP redundancy. Many larger customers who have deployed self managed fiber with a DWDM wavelength provided by a passive optical solution can employ ECMP. This enables geographical redundancy with two P2P WDM circuits. Importantly the connection to the local campus is also provided over a port-channel.





Another option for resiliency is where a Service Provider offers an L2 service in the form of a portchannel. This is a form of MEC that is transparent to the Cisco ASR 1000. This topology is most similar to the test topology detailed in figure 3. This topology is popular for some customers as the service is priced as one managed vpn service with edge redundancy.





Figure 7 shows a popular customer deployment scenario where the datacenter connectivity is redundant, and WAN services are split across SIP/SPA's. This topology allows an operator in the datacenter to maintain connectivity to the router during the entire system upgrade. In the case of the remote site connectivity the numerous services are split across the SIP/SPA's such that not all services are affected at any point during the upgrade and for the datacenter there is never any connectivity lost.

All of these sample topologies achieve the goal of reducing and in most cases nearly eliminating any service outage during the upgrade. These topologies are examples of deployed WAN edge networks deployed with the Cisco ASR 1000, where the goal of the topology was to maximize service uptime during system in service upgrades.

Conclusion

The stated development goal of the Cisco ASR 1000 Series was to achieve a platform that can provide the highest level of service availability even during software upgrades and failover events while utilizing the services and configurations that are typical for our customer deployments. As the data demonstrates, the ESP and RP upgrade have minimal to no impact on the service uptime of the router and any services and data traversing the router. With the Cisco ASR 1000 and a focus on end to end network design you can design a network incorporating the Cisco ASR 1000 that enables your corporation to deliver unprecedented network uptime for all of your services and applications.



Americas Headquarters Cisco Systems, Inc. San Jose, CA Asia Pacific Headquartera Cisco Systems (USA) Pix Ltd. Singacore Europe Headquarters Cixco Systems International BV Amsterdam, The Notherlands

Clace has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Clace Website at www.clace.com/go/offices.

CODE, COENT, Class Eas, Class HealthPrezence, the Class Jopa Class Lumin. Class Nexus, Class Studium/Valon, Class TelePrezence, Class WebEx, DCE, and Welcome to the Human Network are trademarke: Changing the Way We Work, Live, Hay and Learn and Class Store are service marks, and Acaeve Registrat, Aronet, AsyncolS, Bringing the Meeting To You, Catalyst, CCEN, CCEP, CCEP

All other trademurbs mentioned in this document or website are the property of their respective ownere. This use of the word partner does not imply a partnership relationship between Cisco and any other company, (381216)

Printed in USA

C11-517184-00 01/09