

# Introducing IBM Z System Recovery Boost

Octavian Lascu

Kazuhiro Nakajima

Paul Schouten



**IBM Z**





## Introducing IBM Z System Recovery Boost

This IBM® Redpaper publication introduces System Recovery Boost, which is a new function of the IBM z15™ system. System Recovery Boost delivers substantially faster system shutdown and restart, short duration recovery process boosts for sysplex events, and fast catch-up of an accumulated backlog of work after specific events (such as system restart).

**Note:** System Recovery Boost is a firmware feature that is available on z15 central processor complexes (CPCs) (M/T 8561 and M/T 8562), and it requires operating system support for use. It is *not* available on older systems. Always check the latest firmware and software for available System Recovery Boost functions.

**IBM Z® naming:** The IBM Z server generations that are described in this paper are available as the following machine types and models:

- ▶ Machine Type 8561 (M/T 8561) Model T01, which is identified as IBM z15 Model T01.
- ▶ Machine Type 8562 (M/T 8562) Model T02, which is identified as IBM z15 Model T02.

In the remainder of this document, IBM z15 (z15) refers to both machine types (8561 and 8562) unless otherwise specified.

# System Recovery Boost overview

System Recovery Boost is a new feature that is implemented in the z15 CPC firmware. It delivers improved overall system and application availability by minimizing the downtime that results from system shutdown and restart operations, and by using a short duration recovery process boost for sysplex environments.

System Recovery Boost realizes the following benefits:

- ▶ During a planned or unplanned system restart:
  - It shuts down the system substantially faster than any prior IBM Z machine.
  - It helps restart and recover the middleware environment and client workloads substantially faster than on any prior IBM Z machine.
  - It delivers higher processor capacity for a limited time following an initial program load (IPL) (a determined boost period) so that the operating system can start faster and client workloads can catch up and work through a backlog after a downtime.
- ▶ After an event in a parallel sysplex it delivers extra processor capacity for short duration recovery process boosts.

## System Recovery Boost use cases

System Recovery Boost provides value for many use cases, including the following ones:

- ▶ Single-system IPL (planned and unplanned):
  - Planned or rolling IPLs (for example, to install software maintenance and disruptive system maintenance).
  - Unplanned IPLs to recover after an operating system failure or “sick but not dead” occurrence that required a system shutdown or restart.
- ▶ Multi-system IPL (planned and unplanned):
  - Restart all images on a CPC after a planned CPC initial microcode load (IML) or Power-on Reset (POR) (CPC non-concurrent upgrade).
  - Restart all images on a CPC after an unplanned CPC failure following a CPC IML or POR.
  - Start a sysplex after a sysplex-wide (or sysplex multi-system) failure or “sick but not dead” occurrence that required a sysplex shutdown or restart.
- ▶ Disaster recovery (DR) or site switch:
  - Planned DR test: Start the test systems at the DR site.
  - Planned or unplanned site switch: Start the systems at the DR site.
- ▶ Recovery process boosts (short duration) for helping with sysplex recovery

System Recovery Boost provides boosted processor capacity to mitigate the impact on workload processing after short-term recovery events in a sysplex, and it restores normal steady-state sysplex operation as quickly as possible.

# System Recovery Boost functions

System Recovery Boost is a feature that was introduced with the IBM z15 firmware (Driver 41) that is designed to provide higher temporary processing capacity to logical partitions (LPARs) to speed up shutdown, IPL, and stand-alone memory dump (SADMP) operations, and to provide short duration recovery process boosts (following an event in a parallel sysplex) without increasing IBM software costs.

By default, the System Recovery Boost capacity is provided in the following ways:

- ▶ By converting subcapacity central processors (CPs) to full-capacity CPs, also known as a *speed boost*, for the opted-in images<sup>1</sup> during the boost period.
- ▶ By dispatching general-purpose workloads to z Integrated Information Processors (zIIPs) (for IBM z/OS® LPARs that are enabled for boost and with allocated zIIPs in the LPAR profile) during the boost period (zIIP boost).
- ▶ By using firmware enhancements that support greater parallelism and performance improvements in the hardware API services. These enhancements are used by IBM Geographically Dispersed Parallel Sysplex® (IBM GDPS®) to speed up the orchestration of shutdown and restart activities.

This support is available in GDPS V4.2 for workloads running on z15 CPCs.

The boost capacity does not contribute to other IBM software license charges.

Figure 1 shows a typical System Recovery Boost timeline for z/OS.




Figure 1 z/OS typical System Recovery Boost timeline

## Boost timeline

For z/OS, the current implementation enables a 60-minute boost period for IPL and 30 minutes for shutdown events.<sup>2</sup> The 60-minute boost period for performing an IPL of the z/OS system also enables catching up with the backlog work.

<sup>1</sup> Supported operating system images that are enabled for boost and running in an LPAR.

<sup>2</sup> Boost periods are specific to the operating system. Different operating systems can have different boost periods.

For SADMPs, the boost period extends to the duration of the event (but no more than 60 minutes) and uses subcapacity CP speed boost only.

**Note:** SADMP does not use zIIP for boost purposes.

For recovery process boosts, boost is provided up to 30 minutes (the sum of boost times for multiple processes) over a period of 24 hours per system image (z/OS LPAR).

## Delivering extra capacity by using System Recovery Boost

This section describes the ways in which extra capacity is delivered for System Recovery Boost.

### Subcapacity CPs speed boost

When the z15 is configured as a subcapacity model<sup>3</sup>, LPARs that are running in a boost period can access the speed boost. This feature requires operating system opt-in and support. At the time of writing, IBM z/OS (Versions 2.3 and 2.4 with PTFs), IBM z/VM® V7.2 and V7.1 with PTFs, z/VSE® V6.2 with PTFs, and z/TPF V1.1 with PTFs can use the subcapacity boost.

**Note:** Consider the following points:

- ▶ Speed boost applies to general-purpose processors (CPs) only. All other engines run at full capacity (Integrated Facilities for Linux (IFLs), zIIPs, and Internal Coupling Facilities (ICFs)).
- ▶ The following subcapacity models are available:
  - IBM z15 T01 subcapacity CP is available for up to 34 CPs with three subcapacity levels, 4xx, 5xx, and 6xx, where xx is 1 - 34. Model 7xx is a full capacity CP model.
  - IBM z15 T02 subcapacity CP is available for up to six CPs with 26 subcapacity levels, A0x - Y0x, where x is 1 - 6. Model Z0x is a full capacity CP model.

### Speed boost example

Figure 2 on page 5 shows an example of a subcapacity CP speed boost.

<sup>3</sup> For z15 T01: 4xx, 5xx, and 6xx, where xx is 1 - 34. For z15 T02: A01 - Y06.

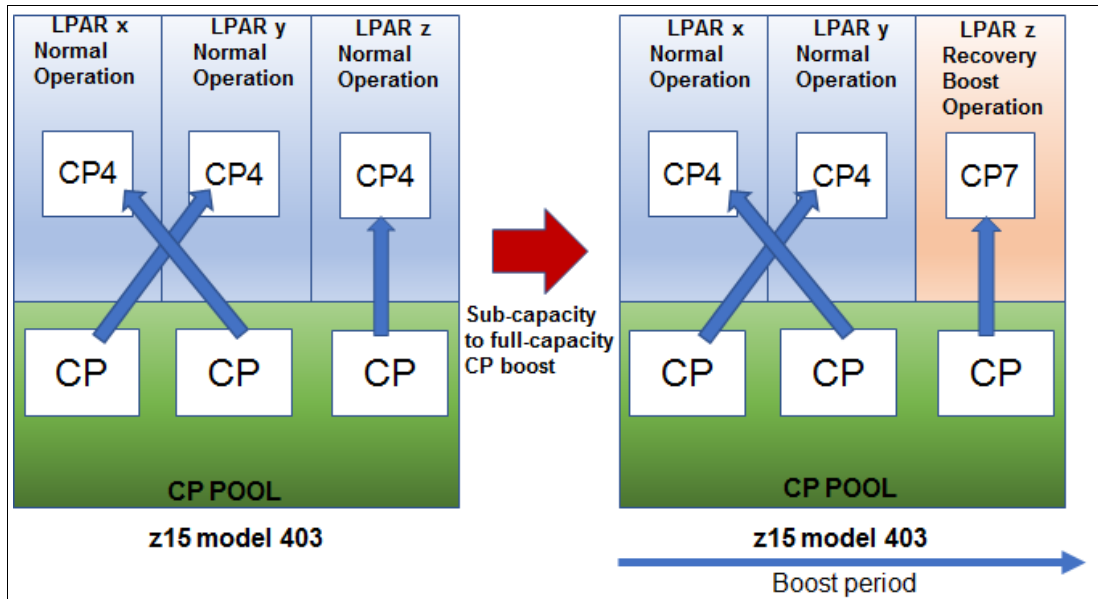


Figure 2 Sub-capacity to full-capacity boost example (z15 T01 shown)

In this example, three LPARs are defined in the z15 T01 model 403. In the normal operation, all work is dispatched on subcapacity CPs.

When LPARs enter a boost period, work that is dispatched from LPAR z runs at CP7 (full capacity). Other LPARs continue to be dispatched at CP4 (subcapacity). One boost period is started at LPAR z shutdown, and a new boost period is started at IPL (of LPAR z). At the end of the IPL boost period, LPAR z returns to normal operation at CP4 (subcapacity).

## zIIP processor capacity boost (zIIP boost)

Normally, only zIIP-eligible work (such as Distributed Relational Database Architecture (DRDA) and IBM Db2® Utilities) is dispatched to zIIPs. During the System Recovery Boost period, zIIP eligible and general CP work is dispatched to available zIIPs for the boost opt-in z/OS images (running in an LPAR).

**Notes:** Consider the following points:

- ▶ As of this writing, z/OS uses the zIIP boost feature.
- ▶ At least one zIIP entitlement must be available to use the zIIP boost.

In this period, the system can use following processors to run a CP workload:

- ▶ Entitled purchased CPs.
- ▶ Entitled purchased zIIPs.
- ▶ Extra zIIPs<sup>4</sup>, which can be added by using the (optional) System Recovery Boost Upgrade, are temporary boost capacity records (Feature Code (FC) 9930, FC 6802, and FC 6799 on z15 T01 only).

If more logical zIIPs are available and configured in the LPAR profile (whether they were added by the temporary boost capacity record or other temporary capacity activation) while in the boost period, images bring more logical zIIP processors online to use the extra physical zIIP capacity.

<sup>4</sup> The System Recovery Boost Upgrade feature (FC 6802, FC 6799, and FC 9930 - enablement) is available for z15 T01 only. For z15 T02, zIIP boost uses only customer existing (purchased) zIIPs.

After the boost period ends, the z/OS dispatching of work on CPs versus zIIPs returns to normal.

**Important:** With proper configuration, z/OS can automatically bring extra logical processors on at the start of a boost period (based on resources availability) and automatically take them off again at the end of the boost period.

The start and end of the boost period is signaled by a console message, ENF signal (84), and cutting an IBM System Management Facility (SMF) record. Also, the start and end of boost period starts a new SMF interval. A system command or PROC (IEABE) is provided to allow for early opt-out of the boost if wanted.

### System Recovery Boost by using zIIP boost example

Figure 3 shows an example of recovery boost by using a zIIP capacity boost.

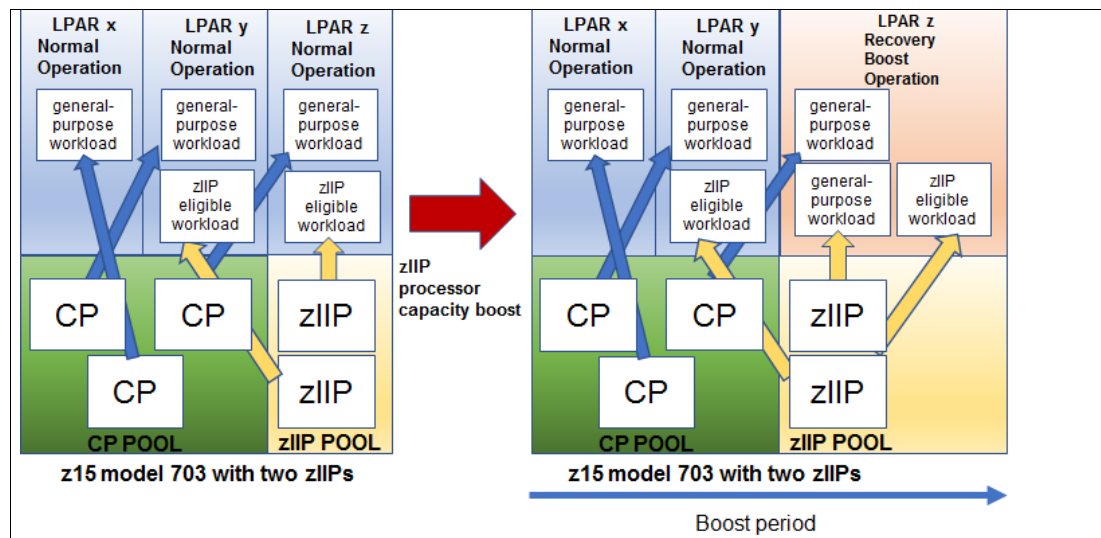


Figure 3 Example of a zIIP recovery boost (z/OS LPAR): z15 T01 example

In this example, three LPARs are defined on the z15 T01 model 703 with two zIIPs. Two zIIPs are shared between LPAR y and LPAR z.

During normal operation, only zIIP eligible work is dispatched to the zIIPs. When LPAR z enters a boost period, general-purpose work and zIIP eligible work might be dispatched to the zIIPs.

When the boost period ends, only zIIPs eligible work is dispatched to the zIIPs.

### System Recovery Boost Upgrade capacity record (optional and z15 T01 only)

You can temporarily increase the number of physical zIIPs to use for System Recovery Boost. This feature is the new System Recovery Boost Upgrade record that you can activate from the **HMC/SE Perform Model Conversion** menu, or by using automation that calls the hardware API services.

System Recovery Boost Upgrade consists of the following features:

- ▶ FC 9930 - System Recovery Boost Upgrade authorization contract (terms and conditions)
- ▶ FC 6802 - System Recovery Boost record - Temporary zIIP records
- ▶ FC 6799 - System Recovery Boost Upgrade Years Ordered (for multi-year SRB Upgrade)



After it is activated, zIIPs are added to the zIIP pool, and the LPAR shares of the extra physical zIIP capacity follow normal LPAR weight rules. You can set up the LPAR image profiles in advance with extra logical zIIP processors to enable the effective use of the extra physical zIIP processor capacity.

You should deactivate the temporary boost capacity record at the end of the IPL or shutdown actions or change windows for which they intend to provide a boost (the record auto-deactivates after 6 hours).

For z15 T01 systems that were ordered with the new System Recovery Boost Upgrade record, the zIIP:CP ratio of 2:1 can be exceeded during the boost periods for the boost opt-in images (LPARs). The boost record activates the zIIPs for up to 6 hours. The boost record features an expiration date of 1 - 5 years, which is controlled by FC 6799.

You must activate the boost record *before* a boost event.

Consider the following points regarding the optional System Recovery Boost Upgrade record:

- ▶ It is a priced feature.
- ▶ The subscription feature is valid for 1 - 5 years (as ordered and controlled by FC 6799).
- ▶ Each activation includes an entitled number of zIIPs, which can be up to 20 and might violate the 2:1 ratio rule between zIIPs and CPs.
- ▶ Activation of this record cannot cause the activation of the first or only zIIP on the machine. Therefore, at least one entitled zIIP must be present.

### Boost temporary capacity record example

Figure 4 shows an example of using an optional boost temporary capacity record.

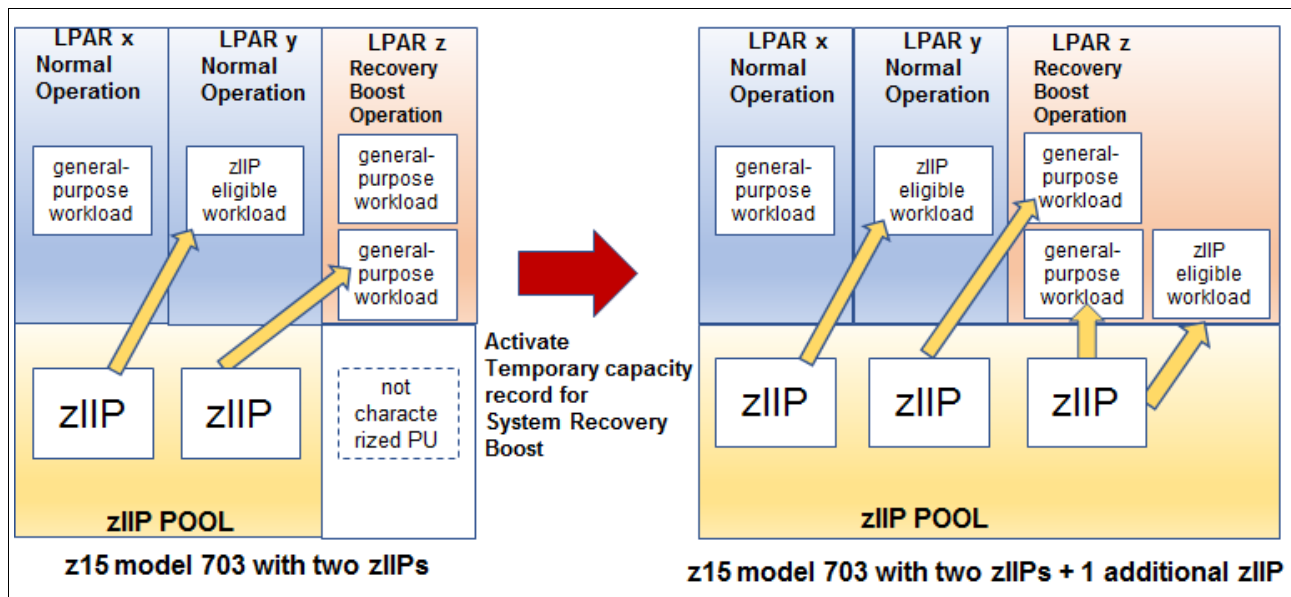


Figure 4 Example of boost that uses a temporary capacity record (z15 T01 only)

In this example, three LPARs are defined in a z15 T01 model 703 with two zIIPs. Two zIIPs are assigned to LPAR y and LPAR z. LPAR y is in normal operation.

When LPAR z is in a boost period, both general-purpose work and zIIP eligible work can be dispatched to the zIIPs. LPAR z includes a reserved zIIP that is specified in its image profile.

By activating a temporary boost capacity record, one zIIP is added to the zIIP pool and automatically allocated to LPAR z and brought online. In this boost period, general CP work for LPAR z is dispatched to zIIPs and CPs.

## Planned shutdown boost

A z/OS system can signal that it wants to enter a boost for a planned shutdown by starting the IEASDBS PROC. Consider the following points:

- ▶ In response to starting the PROC, which is driven manually or by way of automation, z/OS opts in to the allowed boosts that are permitted by using a parmlib.
- ▶ The start and end of the boost period is signaled by a console message, ENF signal (84), and cutting an SMF record. The start and end of the boost period starts a new SMF interval.

In a sysplex, Workload Manager (WLM) sysplex routing starts to route work away from a system after the shutdown PROC is started to accelerate shutdown.

All z/OS and middleware processing during the shutdown boost period benefits from higher capacity CP processors or extra parallelism that is provided by zIIPs and allows CP work to run on zIIPs.

## Shutdown boost example

Figure 5 shows an example of a shutdown boost by using a subcapacity CP speed boost and a zIIP capacity boost.

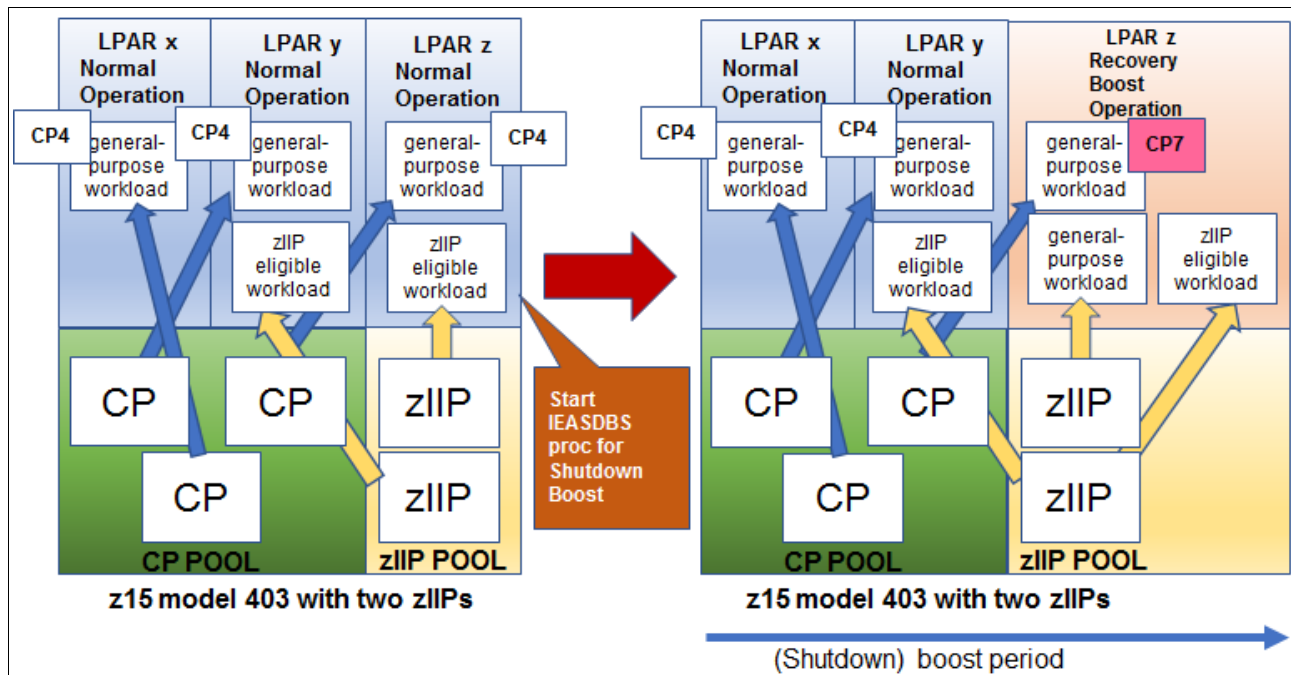


Figure 5 Example of a shutdown boost in a subcapacity model (z15 T01 shown)

In this example, three LPARs are defined in the z15 T01 model 403 with two zIIPs. Two zIIPs are assigned to LPAR y and LPAR z. During normal operation, all CP work is dispatched at subcapacity (CP4), but only zIIP eligible work is dispatched to zIIPs.

Before the planned shutdown of LPAR z, the IEASDBS proc is started by an operator or automation. This process starts the shutdown boost. CP work that is dispatched by LPAR z is run at full-capacity (CP7) and a general-purpose workload is dispatched to zIIPs. LPAR x and LPAR y continue in the normal operation at subcapacity (CP4), and the only zIIP-eligible workload is dispatched to zIIP.

## IBM Geographically Dispersed Parallel Sysplex actions, performance, and parallelism

GDPS drives Base Control Program internal interface (BCPii) hardware APIs for orchestrating Capacity BackUp (CBU) activations, image activations, resets, and IPLs for multiple images in many planned and unplanned DR site-switch scenarios.<sup>5</sup>

Firmware changes on the z15 HMC and SEs support greater parallelism and performance improvements in the hardware API services.

GDPS uses the following changes to take advantage of available parallelism in the following underlying hardware services:

- ▶ GDPS usage that requires GDPS V4.2.
- ▶ Firmware support that is delivered on the z15 machine.

## Recovery process boosts (short duration)

**Important:** The capabilities that are described in this section are available for both IBM z15 T01 and IBM z15 T02 machines running a z/OS images member of a parallel sysplex. These capabilities require the following items:

- ▶ A specific (concurrently installable) LPAR Microcode Change Level (MCL) for z15 Driver 41 or later (check IBM ResourceLink [IBM ResourceLink](#) (an IBM ID is required.)
- ▶ Fixes for supported z/OS versions (V2R3 and V2R4)

**Note:** z/OS APARs are associated with the IBM Fix Category (FIXCAT) for System Recovery Boost (FIXCAT name: IBM.Function.SystemRecoveryBoost).

With enhanced System Recovery Boost support, IBM is extending the boost technologies to provide short-term *recovery process boost* acceleration for specific sysplex recovery events in z/OS.

Currently, these sysplex recovery events often cause short-duration workload impacts or workload spikes while the sysplex is busy recovering for a recovery event. Recovery affects the normal execution of the client workload in the sysplex until the recovery processing completes.

The recovery process boost is designed to provide boosted processor capacity to mitigate short-term recovery impacts and restore normal sysplex steady-state sysplex operation as quickly as possible following specific recovery events, and to provide boosted processor capacity for a short period following restoration of steady-state operation to help with workload “catch-up” from the recovery event.

<sup>5</sup> z/OS exploitation requires z/OS V2.4 with rollback to V2.3 with PTFs (for subcapacity and zIIP boost on z15 CPC).

Each z/OS image can receive boosts as follows:

- ▶ One long-duration boost for image startup (60 minutes) of each type (CP speed boost, zIIP boost), and one long-duration boost for image shutdown (30 minutes) of each type.
- ▶ Several short-duration recovery process boosts, of each type, each of less than 5-minutes duration, with a total usage of no more than 30 minutes of recovery process boost time in any given consecutive 24-hour period per z/OS image (LPAR).

z/OS manages the recovery process boosts internally with the operating system initiating the boosts as these recovery events take place, and only on the images that are affected by these events. If recovery process boosts happen to “overlap” (a second recovery process boost occurs before a first one uses its entire boost period), then the overlapping boosts are merged and the boost period may be extended to allow the full boost period for the second recovery process.

### **Recovery process boost candidates**

The System Recovery Boost recovery process boost provides boosted processor capacity and parallelism to accelerate the following events:

- ▶ Sysplex partitioning

Boost all surviving systems in the sysplex as they take on the extra workload of sysplex partitioning-related recovery after a planned or unplanned removal of a system from the sysplex.

When a system in the sysplex is removed, the surviving systems must do a large amount of recovery processing to clean up after the failed system, free up resources that were held on the failed system, and other tasks.

- ▶ Coupling Facility (CF) structure recovery

Boost all systems participating in CF structure recovery processing, such as CF structure rebuild, duplexing failover, and reduplexing.

Recovering failed CF structures and their data can be a laborious process that requires the participation of all systems that were using those CF structures, and it can apply to many structures in cases like the loss of a CF image.

- ▶ CF data sharing member recovery

Boost all systems participating in recovery from disconnection from a CF lock structure while lock resources were still held.

When a connected user is disconnected from a lock structure while holding lock resources or is disconnected implicitly as the result of address space termination or a system being removed from the sysplex, the other surviving members must do much recovery and cleanup processing to free up locks and other data sharing resources that are held by the failed member.

- ▶ IBM HyperSwap®

Boost all systems participating in a HyperSwap recovery process.

HyperSwap processing is a coordinated, sysplex-wide recovery process that restores access to direct access storage devices (DASDs) following the failure of a storage controller. Its recovery time is sometimes limited by slow processing on one or more participating systems.

## Operational considerations

- ▶ During a recovery process boost period, WLM does not route work away from the system (as it does during a shutdown boost) or toward the system (as it does during a startup boost). WLM ignores short-duration recovery boosts for workload routing purposes.
- ▶ When bringing reserved logical zIIP processors online and offline at the start and end of a recovery process boost period, z/OS limits the number of “transient” zIIPs brought online and offline automatically to at most two (more transient zIIPs during IPL and shutdown boost periods may be configured).
- ▶ System Recovery Boost Upgrade record activation (used to activate additional temporary zIIP capacity on the z15 T01) is *not* supported for use with recovery process boost periods.
- ▶ z/OS starts and ends a new SMF interval during a recovery process boost period, but when two or more recovery processes boost “overlap”, they are merged into a single boost period and a single SMF interval.
- ▶ z/OS issues ENF signals and console messages when starting, extending, or stopping a recovery process boost.
- ▶ z/OS does not permit overlap between the using recovery process boosts and the longer image-level startup and shutdown boosts:
  - Recovery process boosts are not initiated while an image-level startup boost is still in progress because the system is already boosted.
  - If a recovery process boost is in progress when a system image-level shutdown is initiated, then z/OS “cancels” the in-progress recovery process boost and initiates the shutdown boost period for system shutdown.
  - If more transient zIIPs were already online during the recovery process boost, z/OS potentially must bring more online for the shutdown boost, up to the full quota of reserved logical zIIPs.

## Setting up the System Recovery Boost

System Recovery Boost is a firmware feature of the z15 CPC for operating systems that are running in an LPAR, which requires operating system support.

**Important:** The base System Recovery Boost capability is *built into* z15 firmware and does not require ordering any other features.

System Recovery Boost Upgrade (consisting of FC 9930, FC 6802, and FC 6799) is an optional, orderable feature that is available for IBM z15 Model T01 (machine type 8561) *only*. System Recovery Boost Upgrade provides more temporary zIIP capacity for use during boost periods.

Consider the following points:

- ▶ FC 9930 (enablement) is *not* required to use the base System Recovery Boost capability. It is needed only if more zIIP temporary capacity is required.
- ▶ FC 6802 enables temporary activation of extra physical zIIP processors (SRB Upgrade records).
- ▶ FC 6799 specifies the number of SRB Upgrade subscription years (1 - 5).
- ▶ By default, System Recovery Boost is enabled for z/OS (opt-in), and z/OS must run on a z15 CPC.
- ▶ For more zIIP temporary boost capacity, FC 9930, FC 6802 (System Recovery Boost Upgrade), and FC 6799 (number of SRB Upgrade subscription years) must be ordered with the z15 T01 CPC (not available for z15 T02).

You can configure a z/OS system-level parameter (IEASYSxx) to control whether a specific z/OS image opts in for the zIIP processor Boost, as shown in the following example:

```
BOOST=SYSTEM | ZIIP | SPEED | NONE
```

No hardware configuration setup is required.

If you want to use offline zIIPs or extra zIIPs that are provided by the System Recovery Boost record, you must define reserved zIIPs in the image profile, as shown in Figure 6 on page 13.

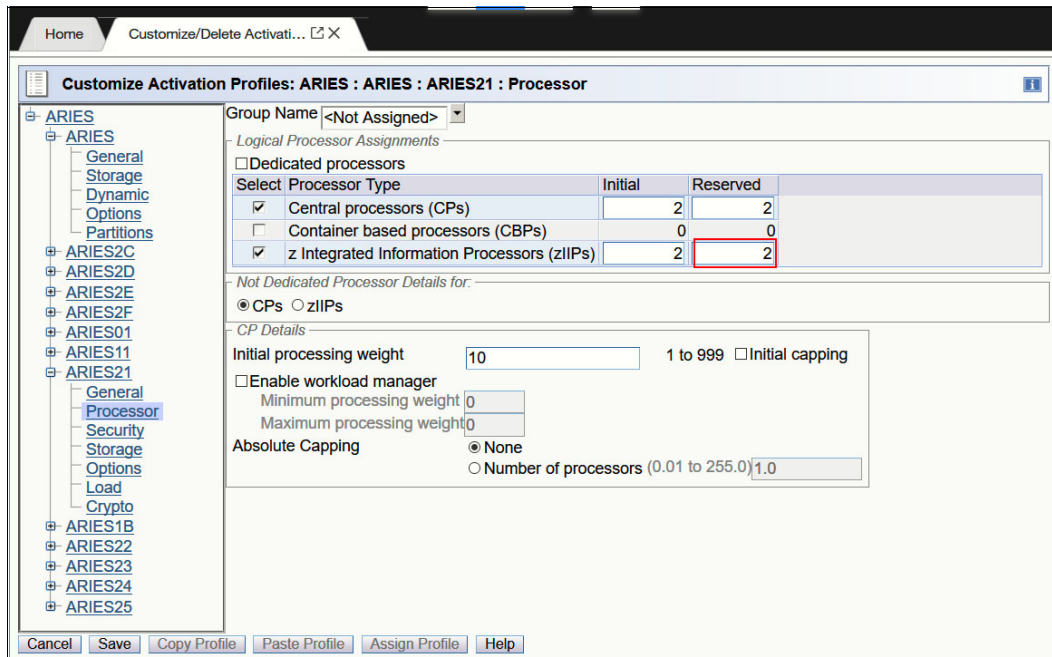


Figure 6 Reserved zIIPs definition window in the image profile

You also should review LPAR weights and storage allocation to ensure that they meet your system requirements.

Capping, accounting and measurement, capacity reporting, workload management and workload routing, capacity management, operating system messages, ENF signals, control block APIs, SMF data during the boost period, and so on, are identical for recovery process boosts as they are for image startup and shutdown boosts.

A new “boost class” is available for recovery process boosts, which appears in various system messages, ENF signals, SMF fields, and other z/OS APIs.

## Monitoring System Recovery Boost

The **D IPLINFO,BOOST,STATE** command shows an image’s current boost state, both for startup and shutdown boosts, and the new recovery process boosts, as shown in Example 1.

*Example 1 z/OS sample display for boost status*

### Sample 1:

```
IEE257I Boost State
Boost class: IPL
zIIP boost: active with 5 transient zIIP cores
Speed boost: active
```

### Sample 2:

```
IEE257I Boost State
Boost class: Recovery Process
Requestor: Hyperswap
zIIP boost: active with 2 transient zIIP cores
Speed boost: active
```

In addition, the **DISPLAY M=CPU** command is enhanced with the following parameters (see Example 2):

- ▶ **I** indicates zIIPs.
- ▶ **B** indicates (transient) boost zIIPs. This CPU was configured online at the start of the boost period, and is configured offline when the boost ends.

*Example 2 CPU information for transient zIIPs*

---

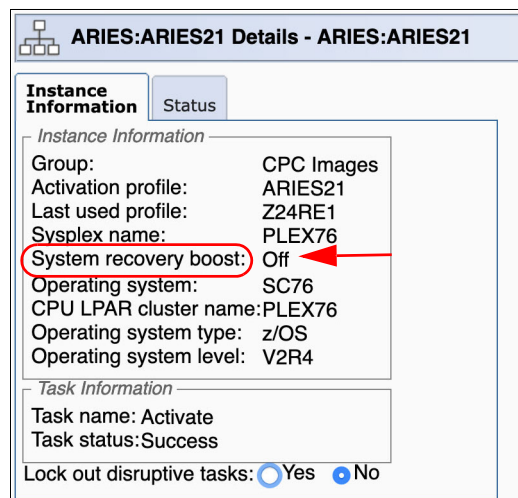
```

SY1 IEE174I 09.58.10 DISPLAY M 328
PROCESSOR STATUS
ID CPU SERIAL
00 + 0449D74381
01 + 1449D74381
02 +I 2449D74381
03 +B 3449D74381
04 +I 4449D74381

```

---

When an LPAR is in the boost period, you can confirm the status of System Recovery Boost in the HMC/SE Partition Image Details window, as shown in Figure 7. During the boost period, Processor Boost is shown as 0n.



*Figure 7 HMC Partition Image Details window*

Also, the processor boost status is shown in HMC Monitors Dashboard.

## Automation

Your automation product can be used in the following ways to automate and control System Recovery Boost activities:

- ▶ To activate (and deactivate) the eBod System Recovery Boost Upgrade record to provide extra physical zIIPs for an IPL or shutdown boost.
- ▶ To dynamically modify LPAR weights, which might be needed to modify or “skew” the sharing of physical zIIP capacity during a boost period.
- ▶ To drive the invocation of the PROC that indicates the beginning of a shutdown process (and the start of the shutdown boost).
- ▶ To use new composite HW API reconfiguration actions.



- ▶ To control the level of parallelism that is present in the workload at start (for example, starting middleware regions) and shutdown (for example, performing an orderly shutdown of middleware).

Automation can pace or throttle these activities to varying degrees. With boost, less pacing or more parallelism might be wanted.

- ▶ To automate the new z/OS messages that are issued at the start or end of boost periods to take whatever actions are needed.

## Pricing

In this section, the available pricing options are described.

### No extra charge functions

The following standard no-charge z15 hardware facilities are available:

- ▶ Subcapacity to full-capacity boost for CPs
- ▶ zIIP boost that use a client's entitled zIIPs
- ▶ GDPS scripting and firmware enhancements

### Priced feature: System Recovery Boost Upgrade record

A charge function is a priced activation of extra zIIP capacity for boost usage. An annual subscription model for entitlement to renewals for zIIP boost with activation of unassigned processor units (PUs) by using temporary record also is available.

System Recovery Boost Upgrade records on z15 T01 require Boost Enablement contract (FC 9930) and temporary pre-paid zIIP boost records (FC 6802 and FC 6799).<sup>6</sup>

### Software pricing

Boost should *not* increase a customers' IBM software costs regardless of whether the client is using 4HRA Pricing, Solution Pricing, or Consumption-based Pricing.

## Software support

The following software is supported by System Recovery Boost:

- ▶ z/OS: z/OS V2R4 with rollback to V2R3 with PTFs and extra PTFs for recovery process boost.

**Note:** z/OS APARs are associated with the FIXCAT for System Recovery Boost (FIXCAT name: IBM.Function.SystemRecoveryBoost).

- ▶ GDPS: Version 4 Release 2.
- ▶ SADMP: SADMP uses subcapacity to full-capacity boost for CPs during memory dump processing (zIIP capacity boost is *not* used for SADMP).

<sup>6</sup> The priced feature System Recovery Boost Upgrade (which provides extra temporary zIIP capacity) is available for z15 Model T01 only.

- ▶ **z/VM:** When running on CP processors, z/VM LPAR uses subcapacity to full-capacity boost for CPs for start and shutdown (IFLs always run at full capacity, so no boost is available for IFLs).  
Second-level guests of z/VM “inherit” the subcapacity boost from VM during these periods, which accelerates the start and shutdown of hosted second-level guests (except for z/OS as a second-level guest).
- ▶ **z/TPF:** z/TPF uses speed boost for CPs for start and shutdown boost. Existing support for a function (called TPF Dynamic CPU) can be used to provide more CP capacity when needed.
- ▶ **IBM z/VSE:** General-purpose processors (CPs) running at subcapacity can be boosted to full capacity for a limited time. z/VSE support for System Recovery Boost support is available as of September 2020 for both z15 T01 and T02. z/VSE requires PTF for APAR DY47832.

During IPL and stand-alone dump processing, these processors are boosted for up to 60 minutes, and during system shutdown these processors are boosted for up to 30 minutes. The z/VSE system automatically enables the boost during IPL and stand-alone dump processing. To enable the boost for system shutdown, use the **SYSDEF SYSTEM** command.

## IBM Z processor characterization (engine types)

An overview of image modes, supported engine types, and operating systems on the z15 is shown in Table 1.

*Table 1 Image modes that are supported on the z15*

Image mode	Engine type (configurable)	Operating system (primary)	Second-level guest
General	CP and zIIP	<ul style="list-style-type: none"> <li>▶ z/OS (CP and zIIP)</li> <li>▶ z/VSE (CP)</li> <li>▶ z/TPF (CP)</li> </ul>	N/A
CF	ICF or CP	Coupling Facility Control Code (CFCC)	N/A
Linux only	IFL or CP	Linux on Z (CP or IFL) KVM (CP or IFL)	Linux on Z under KVM (CP or IFL)
z/VM	IFL, CP, zIIP, and ICF	z/VM (IFL or CP)	<ul style="list-style-type: none"> <li>▶ Linux on Z (CP or IFL)</li> <li>▶ CFCC (CP or ICF)</li> <li>▶ z/OS (CP and zIIP)</li> <li>▶ z/VSE (CP)</li> <li>▶ z/TPF (CP)</li> </ul>
Secure Service Container	IFL or CP	Appliance Software	IFL or CP

### ***Processor unit characterization***

Consider the following points:

- ▶ CP, which is the general-purpose processor, supports running the z/OS, z/VSE, z/VM, z/TPF, and Linux on Z operating systems. It can also run CFCC.
  - For z15 T01, CP can run at subcapacity levels CP4, CP5, and CP6, or at full-capacity (CP7).
  - For z15 T02, CPs (up to six) can run at subcapacity levels CP-A - CP-Y, with CP-Z for a full capacity CP.

All other engines run at full speed.

- ▶ zIIP is an “offload processor” for a workload that supports applications, such as Db2 and z/OS Container Extensions (z/CX). It can also be used for z/OS during System Recovery Boost periods for running a general-purpose workload.
- ▶ IFL is a processor for use by z/VM for Linux guests and Linux on Z operating systems, and the KVM supervisor (that is based on the Linux kernel). The KVM hypervisor is part of supported Linux on Z distributions.
- ▶ ICF is a processor for use by the CFCC.

## **Authors**

This paper was produced by a team of specialists from around the world working at IBM Redbooks, Tucson Center.

**Octavian Lascu** is an IBM Redbooks® Project Leader and a Senior IT Consultant for IBM Romania with over 25 years of experience. He specializes in designing, implementing, and supporting complex IT infrastructure environments (systems, storage, and networking), including high availability (HA) and DR solutions, and high-performance computing deployments. He has developed materials for and taught over 50 workshops for technical audiences around the world. He is the author of several IBM publications.

**Kazuhiro Nakajima** is a Senior IT Specialist in IBM Japan. He has 30 years with IBM Japan and has been active as an advanced subject matter expert (SME) of IBM Z products for over 20 years. His areas of expertise include IBM Z hardware, performance, z/OS, and connectivity. Kazuhiro has co-authored several IBM Z configuration setup IBM Redbooks publications from the IBM zEC12 to the IBM z14®.

**Paul Schouten** is an IBM Z Client Technical Specialist who is based in Sydney, Australia. During his 40 years supporting mainframe systems, he has performed several roles, including Certified IT Architect, Systems Software Developer, and Systems Programming. He has extensive experience developing and documenting HA solutions for IBM Enterprise customers.

Thanks to the following people for their contributions to this project:

Robert Haimowitz and Bill White  
**IBM Redbooks, Poughkeepsie Center**

Dave Surman, Peter Relson, Bill Lamastro  
**IBM Poughkeepsie**

## Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an IBM Redbooks residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

[ibm.com/redbooks/residencies.html](http://ibm.com/redbooks/residencies.html)

## Stay connected to IBM Redbooks

- ▶ Find us on Facebook:  
<http://www.facebook.com/IBMRedbooks>
- ▶ Follow us on Twitter:  
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:  
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:  
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:  
<http://www.redbooks.ibm.com/rss.html>

# Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.


## COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

# Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

Db2®	IBM z14®	z/VM®
GDPS®	Parallel Sysplex®	z/VSE®
HyperSwap®	Redbooks®	z15™
IBM®	Redbooks (logo)  ®	
IBM Z®	z/OS®	

The following terms are trademarks of other companies:

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Other company, product, or service names may be trademarks or service marks of others.





REDP-5563-01

ISBN 0738459208

Printed in U.S.A.

Get connected

