Introducing IBM Z System Recovery Boost

Jannie Houlbjerg

Kazuhiro Nakajima

Paul Schouten

IBM Z

Redpaper

Introducing IBM Z System Recovery Boost

This IBM® Redpaper publication introduces System Recovery Boost (SRB), which is a firmware feature that was introduced with IBM z15™ and enhanced on IBM z16™. SRB delivers substantially faster system shutdown and restart, short duration Recovery Process Boosts (RPBs) for sysplex events, and fast catch-up of an accumulated backlog of mission-critical workloads after specific events (such as system restart).

> **Note:** SRB is a firmware feature that is available on IBM z15 and IBM z16 central processor complexes (CPCs), and it requires operating system support for use. It is *not* available on older systems. Always check the latest firmware and software for available SRB functions.
>
> **IBM Z® naming**: The IBM Z server generations that are described in this paper are available as the following machine types and models (MTMs):
>
> ► Machine Type 3931 (M/T 3931) Model A01, which is identified as IBM z16 Model A01.
> ► 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.

System Recovery Boost overview

SRB is a firmware feature that was introduced with IBM z15 and enhanced on IBM z16. It delivers improved overall system and application availability by reducing the downtime that results from system shutdown and restart operations, and by using short-duration RPBs.

SRB realizes the following benefits:

► During a planned or unplanned system restart:

– It shuts down the system substantially faster than any earlier IBM Z platform.

– It accelerates the initial program load (IPL) and performs restart and recovery of the middleware environment and client workloads substantially faster than on any earlier IBM Z platform.

– It delivers extra processor capacity so that mission-critical workloads can catch up and work through a backlog after a downtime.

- ▶ During a Stand-alone Dump operation, SRB delivers extra capacity to shorten the duration of the dump.
- ▶ During specific (IBM defined) events, SRB delivers extra processor capacity for a short duration so that client mission-critical workloads can catch up and work through a backlog after a downtime. For some of the predefined events, a client can control whether a boost is initiated, but for most events, the boost is on by default.

System Recovery Boost delivery stages

SRB is enhanced with each new firmware level. At the time of writing, all SRB functions were delivered in three stages:

- ▶ SRB Stage 1 was delivered with IBM z15 (GA September 2019 - driver 41). Stage 1 provided extra processing capacity and parallelism for an image that was shutting down or being restarted by using the following methods:
  - – Central Processor (CP) Speed Boost: Running subcapacity CPs at full-capacity speed during boost periods in boosting images only.
  - – IBM System z® Integrated Information Processor (zIIP) Boost: Making general-purpose work available to run on zIIP processors during the boost periods in boosting images only.

  Also provided in SRB Stage 1 were the following features:
  - – SRB Upgrade: A priced temporary capacity offering to activate up to 20 zIIP processors for up to 6 hours per activation. It is on a 1- to 5-year subscription basis that is used during image shutdown or startup.
  - – IBM Geographically Dispersed Parallel Sysplex (IBM GDPS®) scripting and performance and parallelism enhancements.

- ▶ SRB Stage 2 was delivered with MCL P46602.005 for IBM z15 Driver 41C (Bundle S29) or later in addition to APARs for IBM z/OS® V2.R3 and V2.R4. SRB Stage 2 extends Stage 1 boost technologies to provide short-term $RPBs$, which include:
  - – Sysplex Partitioning Recovery

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

    Boosts all systems participating in Coupling Facility (CF) structure recovery processing, CF structure rebuild, duplexing failover, and reduplexing.
  - – CF Data Sharing Member Recovery

    Boosts all systems participating in recovery from the termination of a CF data sharing member.
  - – IBM HyperSwap® Recovery

    Boosts all systems participating in a HyperSwap recovery process.

- ▶ SRB Stage 3 is delivered with the IBM z16 (driver 51). New IBM z16 only use cases for SRB Stage 3 maximize service availability by using tailored, short-duration boosts to mitigate the impact of these recovery processes, which include:
  - – IBM SAN Volume Controller (IBM SVC) Dump Boost
  - – Middleware Region Startup Boost
  - – HyperSwap Configuration Load Boost

The new support for Stage 3 depends on z/OS PTFs that are only available on z/OS V2.R4 and z/OS V2.R5.

This support includes a new IEASRB proc that you can use to enable or disable RPBs on a system.

SRB Stage 3 also includes the following features:

– Displays output enhancements for `D IPLINFO,BOOST,STATE`.

– Improved monitoring for "potential" boosts, that is, boosts that were requested but could not happen because boosts were disabled or boost time was unavailable.

System Recovery Boost use cases

SRB 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 after 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.

► Stand-alone Dump processing

SRB provides boosted processor capacity to shorten the duration of the dump.

► RPBs (short duration) for helping with middleware recovery (in sysplex and non-sysplex environments)

SRB 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

SRB is a firmware feature that is designed to provide extra temporary processing capacity to logical partitions (LPARs) to accelerate shutdown, IPL, restart and recovery of the middleware environment and client workloads; perform Stand-alone Dump operations; and provide short-duration RPBs without increasing IBM software costs.

By default, SRB capacity is provided in the following ways:

► Speed Boost

By converting subcapacity CPs to full-capacity CPs for opted-in images[1] during the boost period.

► zIIP Boost

By dispatching general-purpose workloads to zIIPs (for IBM z/OS LPARs with allocated zIIPs in the LPAR profile) during the boost period.

► GDPS Enhancement

By using firmware enhancements that support greater parallelism and performance improvements in the hardware API services. These enhancements are used by IBM GDPS to speed up the orchestration of shutdown and restart activities. The boost capacity does not contribute to other IBM software license charges.

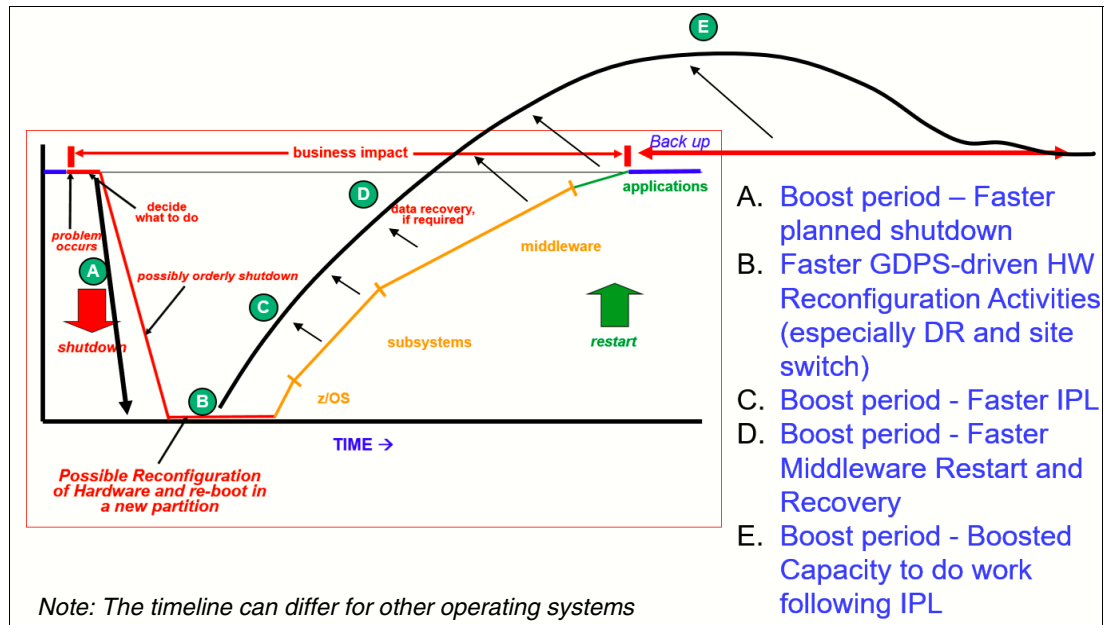Figure 1 shows a typical SRB timeline for a z/OS system restart.

A. Boost period – Faster planned shutdown
B. Faster GDPS-driven HW Reconfiguration Activities (especially DR and site switch)
C. Boost period - Faster IPL
D. Boost period - Faster Middleware Restart and Recovery
E. Boost period - Boosted Capacity to do work following IPL

Figure 1   z/OS typical System Recovery Boost timeline for z/OS restart

---

[1] Supported operating system images that are enabled for boost and running in an LPAR.

System Recovery Boost exploitation and duration

The current exploitation and duration per system image is listed in Table 1.

Table 1   System Recovery Boost exploitation and duration

Item	z/OS	z/TPF[a]	IBM z/VM®	IBM z/VSE®
Function	Speed Boost zIIP Boost GDPS Enhancement	Speed Boost	Speed Boost	Speed Boost
Shutdown boost	30 min.	30 min.	30 min.[b]	30 min.
Startup boost	60 min.	60 min.[c]	60 min.[b]	60 min.
Stand-alone Dump Boost	Up to 60 min.[d]	N/A[e]	N/A[e]	Up to 60 min.
RPB	Up to 30 min. (the sum of boost times for multiple processes) over a period of 24 hours.	N/A[e]	N/A[e]	N/A[e]

a. Existing support for a function (called TPF Dynamic CPU) can be used to provide more CP capacity when needed.
b. Second-level guests of z/VM (other than z/OS images) "inherit" the Speed Boost from z/VM during these z/VM startup and shutdown periods. This situation applies to second-level z/TPF, z/VSE, and potentially Linux images. However, individual startups and shutdowns of guests (in isolation from z/VM image startup and shutdown) are *not* boosted. Integrated Facility for Linux (IFL) always runs at full capacity, so no boost is available for IFLs.
c. A maximum of 20 minutes might be spent in the catch-up phase after the system restarts.
d. SADMP does *not* use zIIP Boost.
e. N/A: Not available.

Delivering extra capacity by using System Recovery Boost

This section describes the ways in which extra capacity is delivered for SRB.

Speed Boost

When the CPC is configured as a subcapacity model, 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, IBM z/VM, z/VSE, and z/TPF can use the subcapacity boost. For more information, see "Software support" on page 19.

> **Note:** Consider the following points:
>
> ► Speed Boost applies to general-purpose processors (CPs) only. All other engines run at full capacity (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.
>
> – IBM z16 A01 subcapacity CP is available for up to 39 CPs with three subcapacity levels, 4xx, 5xx, and 6xx, where xx is 1 - 39. Model 7xx is a full capacity CP model.

System Recovery Boost by using Speed Boost example

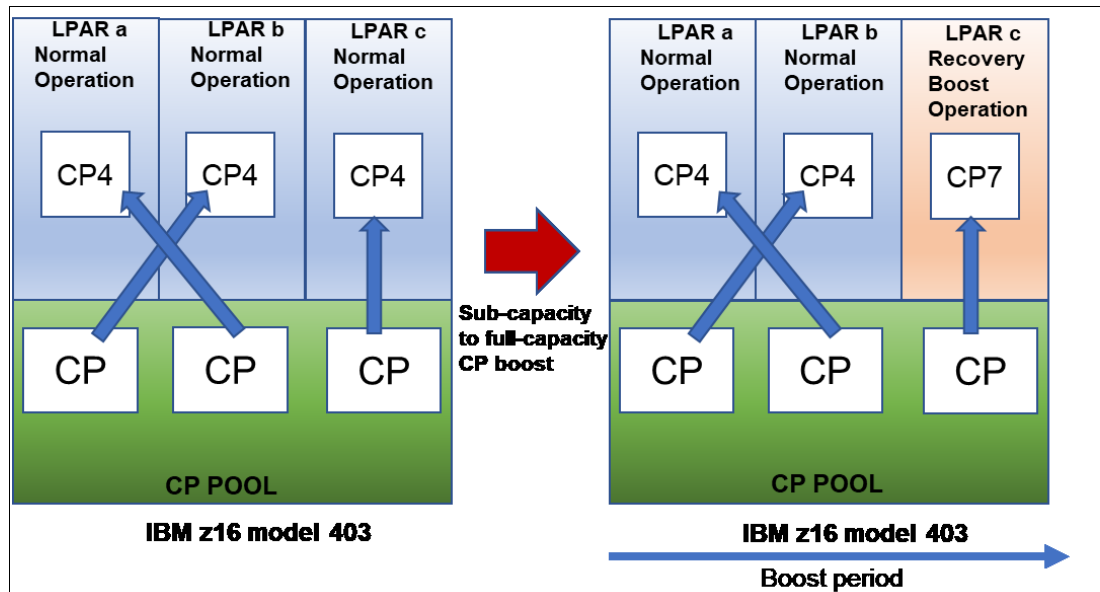Figure 2 shows an example of Speed Boost.

Figure 2   Subcapacity to full-capacity boost example (IBM z16 A01 shown)

In this example, three LPARs are defined in IBM z16 A01 model 403. In normal operation, all work is dispatched on subcapacity CPs.

When LPAR c enters 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 c shutdown, and a new boost period is started at IPL (of LPAR c). At the end of the IPL boost period, LPAR c returns to normal operation at CP4 (subcapacity).

zIIP Boost

Normally, only zIIP eligible work (such as Distributed Relational Database Architecture (DRDA) and IBM Db2® Utilities) is dispatched to zIIPs. During the SRB 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:
>
> ► At the time of writing, z/OS uses the zIIP Boost feature.
>
> ► At least one zIIP entitlement must be available to use zIIP Boost.

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

► Entitled purchased CPs.

► Entitled purchased zIIPs.

► Extra zIIPs, which can be added by using the (optional) SRB Upgrade, are temporary boost capacity records. For more information, see "System Recovery Boost Upgrade capacity record" on page 8.

If more logical zIIPs are available and configured in the LPAR profile while in the boost period, the images bring more logical zIIP processors online to use the extra physical zIIP capacity.

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.

System Recovery Boost by using zIIP Boost example

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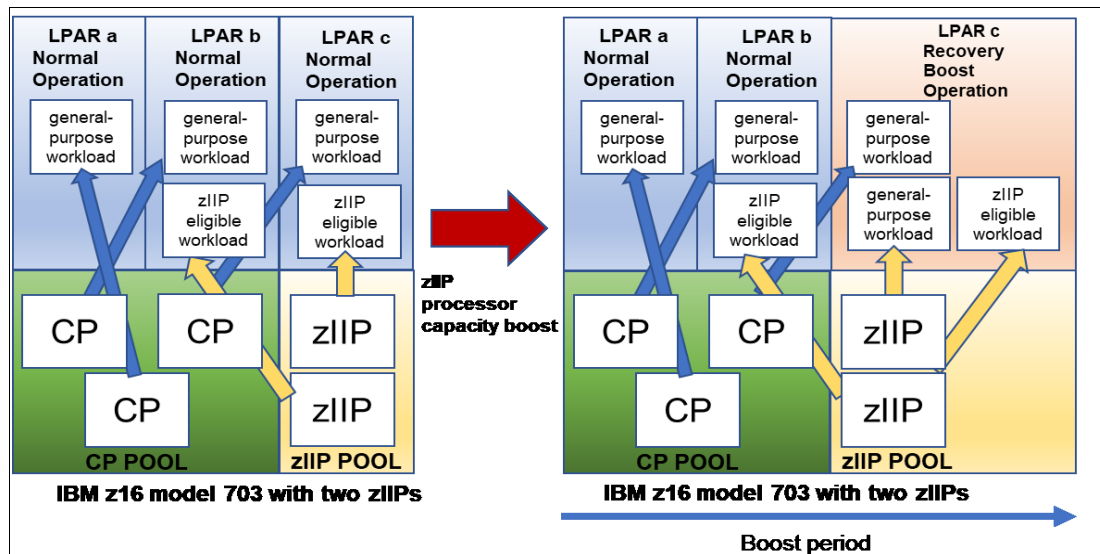

Figure 3   zIIP Boost (z/OS LPAR): IBM z16 A01 example

In this example, three LPARs are defined on the IBM z16 A01 model 703 with two zIIPs. Two zIIPs are shared between LPAR b and LPAR c.

During normal operation, only zIIP eligible work is dispatched to the zIIPs. When LPAR c 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

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

> **Note:** Consider the following points:
> - ▶ This feature is available on the following machines:
>   - – IBM z15 T01.
>   - – IBM z16 A01.
> - ▶ SRB Upgrade consists of the following *prepaid* features:
>   - – Boost Authorization.
>   - – Sys Recovery Boost record: Enables temporary activation of extra physical zIIP processors.
>   - – Total SRB Years Ordered: Specifies the number of SRB Upgrade subscription years (1 - 5).

After SRB Upgrade 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.

Each activation includes an entitled number of zIIPs, which can be up to 20.

For CPCs that were ordered with the new SRB 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. 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).

The boost record features an expiration date of 1 - 5 years.

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.

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

System Recovery Boost Upgrade capacity record example

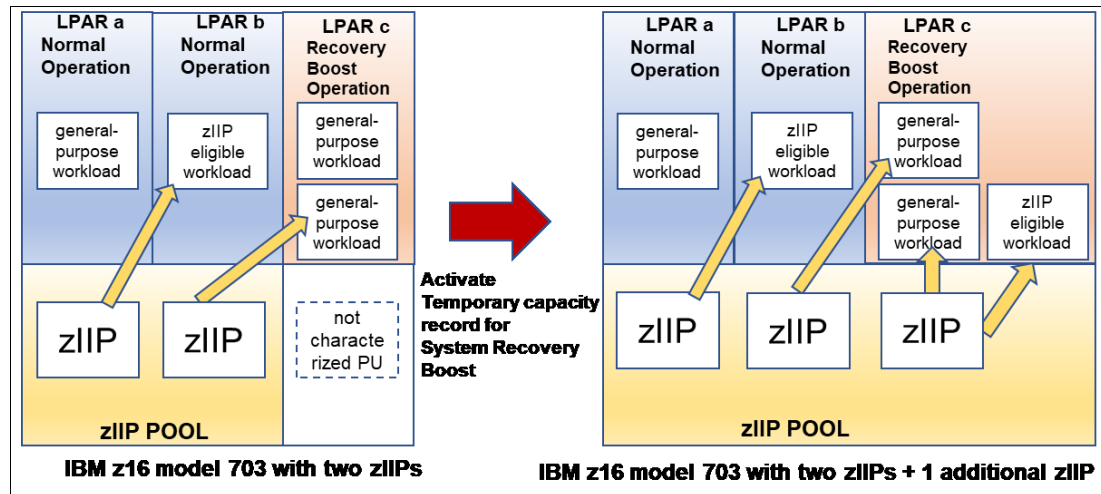Figure 4 shows an example of using an optional SRB Upgrade capacity record.

Figure 4   Example of a boost that uses a temporary capacity record (IBM z16 A01 only)

In this example, three LPARs are defined in an IBM z1n6 A01 model 703 with two zIIPs. Two zIIPs are assigned to LPAR b and LPAR c. LPAR b is in normal operation.

When LPAR c is in a boost period, both general-purpose work and zIIP eligible work can be dispatched to the zIIPs. LPAR c 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 c and brought online. In this boost period, general CP work for LPAR c 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 IEA Shut Down Boost Start (IEASDBS) PROC. Consider the following points:

► To use SRB to speed up shutdown, you start the proc IEASDBS to tell z/OS that a shutdown is in progress and to begin the shutdown boost. If for some reason you want to stop a shutdown boost before the 30-minute time limit, you can run proc IEABE to do so. However, stopping a shutdown boost by using the IEABE proc will not allow you to use any remaining boost time in the future, that is, the IEASDSBS proc cannot be re-invoked to start the shutdown boost once again.

► 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 (event notification facility (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.
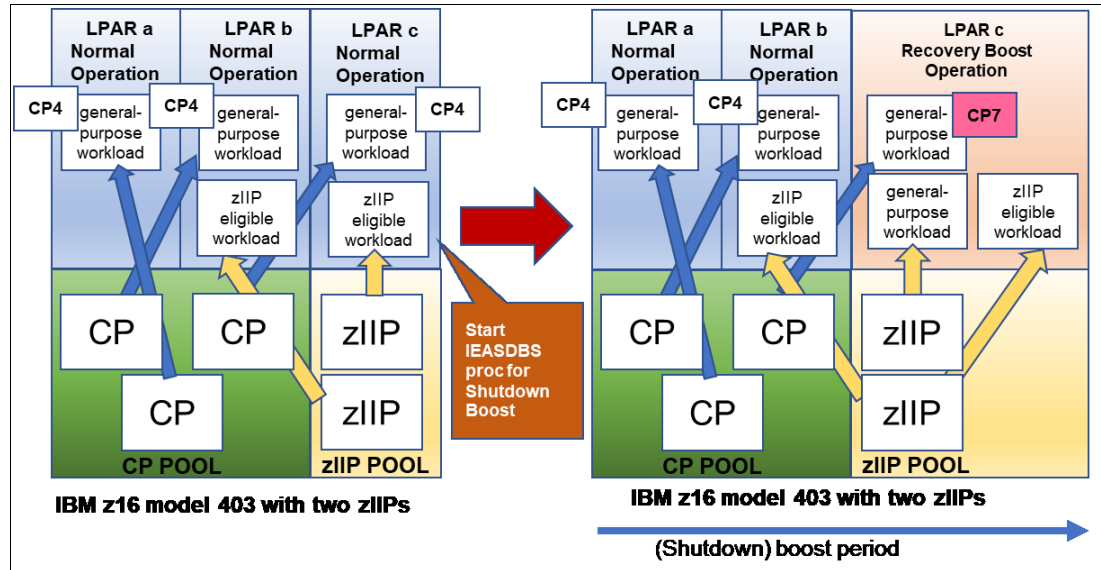

Figure 5   Example of a shutdown boost in a subcapacity model (IBM z16 A01 shown)

In this example, three LPARs are defined in the IBM z16 A01 model 403 with two zIIPs. Two zIIPs are assigned to LPAR b and LPAR c. 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 c, the IEASDBS proc is started by an operator or automation. This process starts the shutdown boost. CP work that is dispatched by LPAR c is run at full-capacity (CP7) and a general-purpose workload is dispatched to zIIPs. LPAR a and LPAR b 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.

Firmware changes on the Hardware Management Console (HMC) and Support Elements (SEs) support greater parallelism and performance improvements in the hardware API services, and GDPS has changed the processing and scripting to take advantage of the firmware improvement. GDPS itself also does more work in parallel when possible.

For z/OS and GDPS support and requirements, see "Software support" on page 19.

Recovery Process Boost (short duration)

> **Notes:**
> - All these capabilities are available for IBM z16 A01.
> - *Some* of these capabilities are available for IBM z15 with a specific (concurrently installable) LPAR Microcode Change Level (MCL). For more information, see IBM ResourceLink (an IBMid is required.)
> - At the time of writing, z/OS can exploit these features. For more information, see "Software support" on page 19.

With enhanced SRB support, IBM is extending boost technologies to provide short-term RPB acceleration for specific sysplex recovery events.

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

The RPB is designed to provide boosted processor capacity to mitigate short-term recovery impacts and restore 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.

When the events occur, sometimes an RPB is automatically initiated, and sometimes more configuration or setup is required to initiate the RPB, depending on what the event is.

When an event occurs, z/OS manages the RPBs 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. For certain events, additional configuration or setup is required to initiate RPB. If RPBs happen to "overlap" (a second RPB 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. For setup considerations, see "z/OS" on page 15.

z/OS supports a dynamic enable and disable capability for all the RPB candidates, which provides more control over the applicability and usage of RPBs, perhaps at different times of day or under different operating conditions. For more information, see "z/OS" on page 13.

Recovery Process Boost candidates

The SRB RPB 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 perform other tasks.

- 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.

► SVC Dump Boost

Boost the system on which SVC dump is taken to reduce system impact and expedite diagnostic capture.

By default, SVC Dump Boost is not enabled. You must set a threshold to enable it. For more information, see "z/OS" on page 13.

> **Note:** SVC Dump Boost is only available on IBM z16 A01.

► Customer-identified Middleware Region Boost for startup or restart

Boost the system on which a middleware instance is being restarted to expedite resource recovery processing, release retained locks, and other items.

This boost is applicable to planned restarts; restarts after failure; and automated or z/OS Automatic Restart Manager (ARM) driven restarts. SRB does not boost any system address space by default, and must be explicitly configured by the WLM service definition. For more information, see "z/OS" on page 13.

> **Note:** Customer-identified Middleware Restart Boost is available only on IBM z16 A01.

► HyperSwap Configuration Load Boost

Boost the system in which the HyperSwap configuration and policy information are being loaded or re-loaded.

This boost applies to both Copy Services Manager (CSM) and GDPS. HyperSwap Configuration Load Boost is enabled by default. There are no thresholds or criteria that are applied to the boost request based on the size or number of devices that are present in the HyperSwap configuration.

> **Note:** HyperSwap Configuration Load Boost is available only on IBM z16 A01.

Operational considerations

Here are several operational considerations:

► During an RPB 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 an RPB 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).

► SRB Upgrade record activation is *not* supported for use with RPB periods.

► z/OS starts and ends a new IBM System Management Facility (SMF) interval during an RPB 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 an RPB.

► For z/OS, a system command or PROC (IEABE) is provided to allow for early opt-out of an IPL or shutdown boost if wanted.

► z/OS does not permit overlap between the RPBs and the longer image-level startup and shutdown boosts:

  – RPBs are not initiated while an image-level startup boost is still in progress because the system is already boosted.

  – If an RPB is in progress when a system image-level shutdown is initiated, then z/OS "cancels" the in-progress RPB and initiates the shutdown boost period for system shutdown.

  – If more transient zIIPs were already online during the RPB, 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

SRB is a firmware feature for operating systems that are running in an LPAR, which requires operating system support.

> **Important:** The base SRB capability is *built in to* the firmware and does not require ordering any other features.
>
> SRB Upgrade capacity record is an optional, orderable prepaid feature that is available for IBM z15 Model T01 and IBM z16 Model A01. SRB Upgrade capacity record provides more temporary zIIP capacity for use during boost periods. For more information, see "System Recovery Boost Upgrade capacity record" on page 8.

z/OS

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

```
BOOST=SYSTEM | zIIP | SPEED | NONE
```

You can configure the RPB - SVC Dump by using RPBMINSZ in the `CHNGDUMP` parameter. For more information, see Managing Recovery Process Boost.

You can configure the RPB - Customer-identified Middleware Restart Boost through your WLM Service Definition. WLM specification is required to designate any Middleware Restart Boost for started tasks (STCs) only. The boost is identified in WLM Classification Rules, with `Boost=Y | N`.

For all the RPB candidates, you can configure these boosts as `ENABLED` or `DISABLED` by using the IEASRB PROC, as shown in the following example:

```
S IEASRB,[CLASS=RP,]REQ=DISABLE|D to dynamically disable Recovery Process boost
S IEASRB,[CLASS=RP,]REQ=ENABLE|E to dynamically enable Recovery Process boost
```

For more information, see *z/OS 2.5: System Recovery Boost.*

If you want to use offline zIIPs or extra zIIPs that are provided by the SRB record, you must define reserved zIIPs in the image profile, as shown in Figure 6.
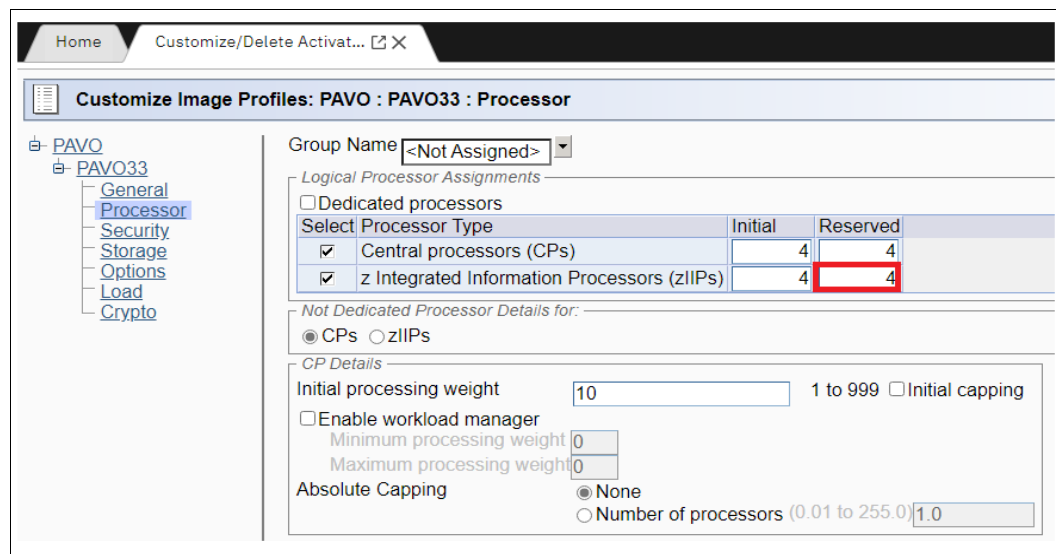

Figure 6   Reserved zIIPs definition window in the image profile

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

z/VSE

SRB is automatically enabled during IPL and Stand-alone Dump. The shutdown boost must be enabled by running the `SYSDEF SYSTEM` command, as shown in the following example:

```
SYSDEF SYSTEM,BOOST=START
```

z/VM

SRB is enabled by default. It can be controlled by the `FEATURES` statement in the system configuration file, and it can be `ENABLED` and `DISABLED`, as shown in the following example:

```
FEATURES ENABLE RECOVERY_BOOST
FEATURES ENABLE RECOVERY_BOOST
```

Monitoring System Recovery Boost

This section describes monitoring SRB for the following items:

► z/OS

► z/VSE

► HMC

z/OS

The `D IPLINFO,BOOST,STATE` command shows an image's current boost state, both for startup and shutdown boosts, and the RPBs, including details such as time used, remaining time, and other details within the last 24-hour window, as shown in Example 1.

Example 1   z/OS sample display for boost status

```
Sample 1:
IEE258I00.18.13 Boost State 899
 Boost class: IPL
 zIIP  boost: active with 5 transient zIIP cores
 Speed boost: active
Recovery Process Boosts: enabled at 09.34.54 on 03/24/2022
Recovery Process Boost Usage Information
  zIIP  boost: Total Avail:  30:00  Used:  00:00  Remaining:  30:00
  Speed boost: Total Avail:  30:00  Used:  00:00  Remaining:  30:00

  Total Available time will be refreshed in: 04:33:22

Recovery Process Boost Usage Since IPL
  Actual: 00:00:00 Potential_E: 00:00:00 Potential: 00:00:00
```

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
```

`DISPLAY DUMP,OPTIONS` can provide information about the RPBMINSZ threshold value that is set to control boosts.

In addition, `DISPLAY DUMP,INFO` can provide information about dump sizes.

A number of IEA-messages are introduced with SRB. Messages are issued when a boost is started and when a boost ends. These messages include information such as which type of boost (Shutdown, IPL, or RPB) is activated. In an RPB situation, more information is available, such as requester name (for example, 'Sysplex partitioning'). Messages also are issued if SRB cannot be activated due to an error or limits having been reached, or if SRB or RPB is `DISABLED`. For more information, see *z/OS 2.5: System Recovery Boost.*

A new SMF interval starts whenever a boost starts or ends. SMF 90-40 is generated when a boost event occurs (starting or stopping a boost). SMF 90-40 contains information about the boost. SMF 90-40 is generated every 24 hour, and it contains summarized information about RPB, such as enablement or disablement of state and usage. SMF types 30, 70-1, 72-3, 89-1, 89-2, and 99-1 all receive a boost information byte that informs you whether a boost is active, and if so, which type. During the boost, SMF records that contain capacity information, such as 30, 70-1, 72-3, and 89, reflect the unboosted capacity of the image, not the effective boosted capacity. For more information, see *z/OS 2.5: MVS System Management Facilities (SMF)*.

z/VSE

The **SIR** and the **QUERY SYSTEM** commands show an image's current boost state.

HMC

When an LPAR is in the boost period, you can confirm the status of the SRB in the HMC or SE Partition Image Details window, as shown in Figure 7. During the boost period, Processor Boost is shown as `Active`.
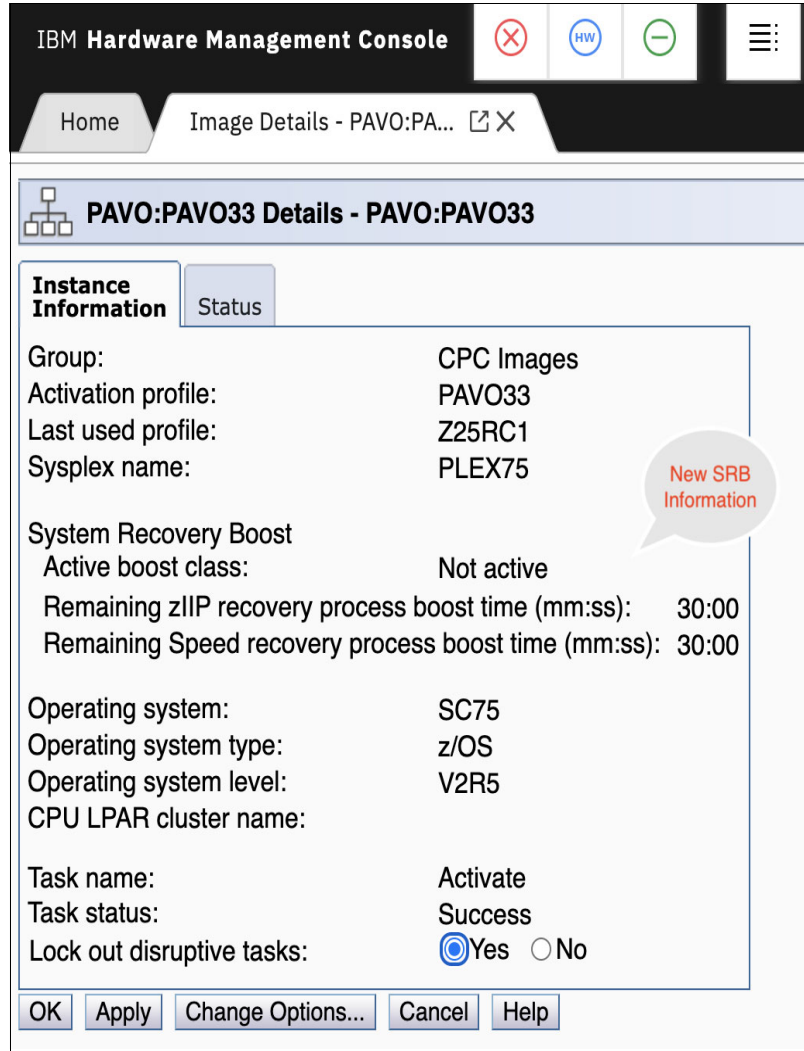

Figure 7   HMC Partition Image Details window

Also, the processor boost status is shown in HMC Monitors Dashboard, as shown in Figure 8.
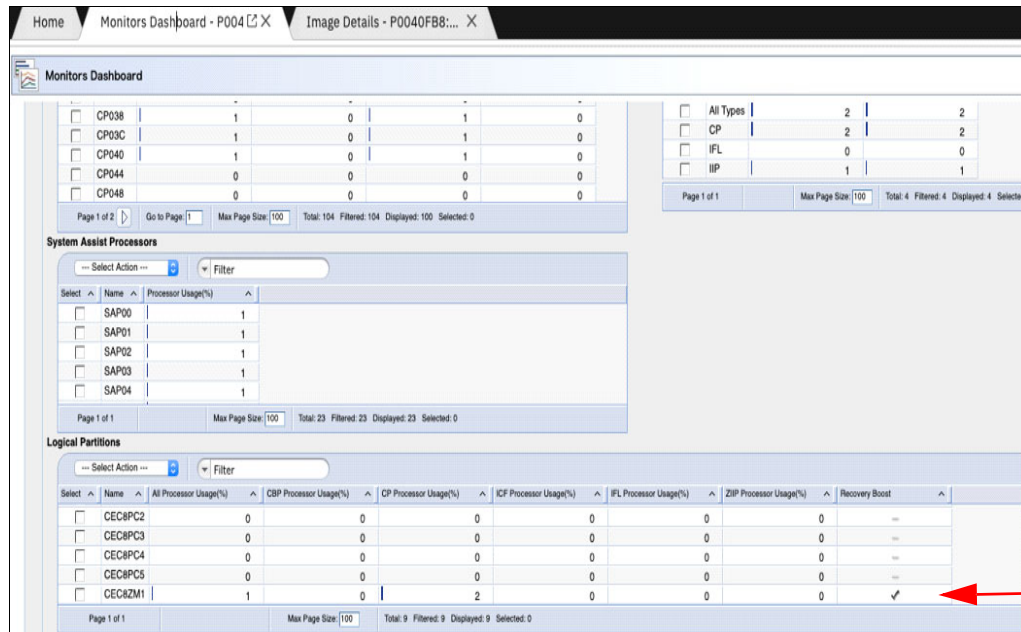

Figure 8   HMC Monitor Dashboard Window

Automation

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

► Activate (and deactivate) the eBod SRB Upgrade record to provide extra physical zIIPs for an IPL or shutdown boost.

► Dynamically modify LPAR weights, which might be needed to modify or "skew" the sharing of physical zIIP capacity during a boost period.

► Drive the invocation of IEASDBS that indicates the beginning of a shutdown process.

► Use new composite hardware API reconfiguration actions.

► Control the level of parallelism that is present in the workload at start (for example, starting middleware regions) and shutdown. Automation can pace or throttle these activities to varying degrees. With boost, less pacing or more parallelism might be wanted.

► Automate the new z/OS messages that are issued at the start or end of boost periods to take whatever actions are needed.

► Drive IEASRB proc invocations to enable or disable RPBs at certain times of day or as otherwise needed.

Pricing

In this section, the available pricing options are described.

No additional charge functions

The following standard no additional charge IBM z15 and IBM z16 hardware facilities are available:

► Subcapacity to full-capacity boost for CPs
► zIIP Boost that uses 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 a temporary record also is available.

SRB Upgrade records on IBM z15 T01 and IBM z16 A01 require a Boost Enablement contract and temporary pre-paid zIIP Boost records. For more information, see "System Recovery Boost Upgrade capacity record" on page 8.

Software pricing

A 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

At the time of writing, the following software is supported by SRB:

► IBM z/OS: V2R3, V2R4, and V2R5 with PTFs. Some of the RPBs that are provided in Stage 2 have support in z/OS V2R3. The new Stage 3 RPBs require z/OS V2R4 or later.

> **Note:** z/OS APARs are associated with the FIXCAT for SRB (the FIXCAT name is `IBM.Function.SystemRecoveryBoost`).

► IBM GDPS: V4R2, V4R3, and V4R4.

► IBM z/VM: V7.1 with PTFs, V7.2, and V7.3.

► IBM z/TPF: V1.1 with PTFs.

► IBM z/VSE: V6.2 with PTFs.

IBM Z processor characterization (engine types)

An overview of image modes, supported engine types, and operating systems on IBM z15 and IBM z16 is shown in Table 2.

Table 2   Image modes that are supported on IBM z15 and IBM z16

Image mode	Engine type (configurable)	Operating system (primary)	Second-level guest
General	CP and zIIP	► z/OS (CP and zIIP) ► z/VSE (CP) ► z/TPF (CP)	N/A
CF	ICF or CP	Coupling Facility Control Code (CFCC)	N/A
Linux only	IFL or CP	Linux on IBM Z (CP or IFL) Kernel-based Virtual Machine (KVM) (CP or IFL)	Linux on IBM Z under KVM (CP or IFL)
z/VM	IFL, CP, zIIP, and ICF	z/VM (IFL or CP)	► Linux on IBM Z (CP or IFL) ► CFCC (CP or ICF) ► z/OS (CP and zIIP) ► z/VSE (CP) ► z/TPF (CP)
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 IBM Z operating systems. It can also run CFCC.

– For IBM z15 T01 and IBM z16 A01, CP can run at subcapacity levels CP4, CP5, and CP6, or at full-capacity (CP7).

– For IBM 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 SRB periods for running a general-purpose workload.

► IFL is a processor for use by z/VM for Linux guests and Linux on IBM Z operating systems, and the KVM supervisor (that is based on the Linux kernel). The KVM hypervisor is part of supported Linux on IBM 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, Poughkeepsie Center.

Jannie Houlbjerg is a Systems Programmer working at JN Data in Denmark. She has more than 20 years of experience in the IBM® Z field. Her areas of expertise include IBM Z hardware and infrastructure, IBM Parallel Sysplex®, connectivity, performance, IBM GDPS, and technical project management and documentation. Jannie has co-authored several IBM Redbooks® publications.

Kazuhiro Nakajima is a Senior IT Specialist at IBM Japan. He has more than 30 years at 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 IBM zEC12 to 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 high availability (HA) solutions for IBM Enterprise customers.

The project that created this IBM Redpaper publication was managed by Octavian Lascu, Redbooks Project Leader, Poughkeepise Center.

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

Robert Haimowitz and Bill White
IBM Redbooks, Poughkeepsie Center

Dave Surman, Peter Relson, and 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

Stay connected to IBM Redbooks

- ► 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 z16™	z/VM®
GDPS®	Parallel Sysplex®	z/VSE®
HyperSwap®	Redbooks®	z15™
IBM®	Redbooks (logo) ®	z16™
IBM Z®	System z®	
IBM z14®	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.

Get connected

ibm.com/redbooks