z/Architecture (64-bit) Migration Checklist

The IBM @server zSeries 800 (z800) and zSeries 900 (z900 and z990) servers support the 64-bit architecture, z/Architecture. With z/OS IPLed in z/Architecture mode on a z800, z900, or z990 server, you can take advantage of functions such as 64-bit real and virtual storage support, Intelligent Resource Director (IRD), and HiperSockets.

__ Decide how many servers and LPARs you need. 64-bit real storage support might allow you to consolidate your current systems into fewer LPARs or to a single native image.

__ Determine whether any locally-developed code is affected by z/Architecture. For information about 64-bit virtual addressing support, review:
- *z/OS MVS Programming: Authorized Assembler Services Reference ALE-DYN*
- *z/OS MVS Programming: Authorized Assembler Services Reference ENF-IXG*
- *z/OS MVS Programming: Authorized Assembler Services Reference LLA-SDU*
- *z/OS MVS Programming: Authorized Assembler Services Reference SET-WTO*

__ Check with your ISVs to ensure they support z/Architecture. Also check the ISV web site: http://www-1.ibm.com/servers/eserver/zseries/zos/zos_products.html -- scroll down to vendor software products for z/OS and click on the appropriate z/OS release.

__ Obtain PSP bucket for 64-bit service requirements:
    UPGRADE: 2084DEVICE SUBSET: 2084/ZOS (z990)
    UPGRADE: 2064DEVICE SUBSET: 2064/ZOS (z900)
    UPGRADE: 2066DEVICE SUBSET: 2066/ZOS (z800)

__ Install **required** Real Storage Manager service:
    OW55209 - available August 2002
    OW55255 - available September 2002
    OW54938 - available August 2002
    OW55729 - available October 2002
    __ Continue to review & install RSM HIPER PTFS

Note: It is important to read the APAR cover letters and apply circumvention's as appropriate. For example, the circumvention for OW55729 and OW55902, which is documented in OW55729, recommends you set the IEAOPTxx MCCAFCTH thresholds based on the amount of central storage on your system. Once you have the fixes for OW55729 and OW55902 installed, it is recommended you use the default MCCAFCTH thresholds (and remove the circumvention), and remove the MCCAFCTH statement from IEAOPTxx altogether.

__ Standalone Dump Planning
    __ Plan for additional resources for stand-alone dump. In z/Architecture, all of central storage is dumped rather than just the 2 GB of central storage in an ESA/390 system. In order to support this much larger requirement, review the use of multivolume standalone dump data set support, which was provided in z/OS V1R2. Refer to
    - *z/OS MVS Diagnosis: Tools and Service Aids*
    - WSC White Paper WP100269, *z/OS Performance: Managing Processor Storage in a 64-bit Environment - V1.1* (Chapter 3) at
      http://www.ibm.com/support/techdocs/atsmastr.nsf/PubAllNum/WP100269

z/Architecture (64-bit) Migration Checklist

__ *For performance considerations, review Washington Systems Center Flash 10143 at* http://www.ibm.com/support/techdocs

__ Configure expanded storage to all Central Storage. (Expanded Storage is ignored in z/Architecture mode.)

__ Understand how to evaluate an all central storage environment. See WSC White Paper WP100269 at http://www.ibm.com/support/techdocs/atsmastr.nsf/PubAllNum/WP100269

__ Understand Load Real Address (LRA) instruction considerations

- cannot return an address larger than 2GB
- can be used if the virtual is known to be backed below 2GB

Note: Some additional information: When z/OS is running in z/Architecture mode, authorized applications issuing the LRA instruction against unfixed storage might receive an abend 0D3, reason code 13. The results of an LRA instruction against unfixed storage has always been unpredictable. In z/Architecture mode, the LRA instruction might require a 64-bit result when only a 32–bit result can be returned; in that case, the hardware causes a program interrupt. If the program is using the LRA instruction to validate that the virtual address is backed by real storage, then use the TPROT instruction instead. If a valid real address is required, the storage must be properly fixed in real storage below 2 GB before issuing the LRA instruction. The STRAG or LRAG instruction can be used to obtain a 64-bit real address if the storage can be backed anywhere. You should review your usage before migrating to z/Architecture, and make any required changes.

__ Determine if you want to control whether IP applications programs should be passed 64-bit backed storage. Controlled by using API64R start option.

Note: TCP/IP 64-bit real addressing support is automatically enabled. TCP/IP exploits real storage in excess of 2 gigabytes by allowing z/OS to back most fixed CSM data space pages above the 2-gigabyte real storage bar. Information on API64R is located in *z/OS Communications Server: SNA Resource Definition Reference.*

__ Understand z/Architecture provides a different format for architected parts of the first page of storage (the PSA).

The ESA/390 format is mapped by macro IHAPSA. The z/Architecture format is mapped by macro IHAPSAE. In particular, the "old PSW", "new PSW" and "interruption parameter" fields have moved. Programs should not use any of these fields. But, if they do, they will not work correctly when running in z/Architecture (64-bit) mode. None of the programming interface fields in the PSA have changed. For information about IHAPSA and IHAPSAE, look under the name PSA in *z/OS MVS Data Areas, Vol 3 (IVT-RCWK).*

__ Understand capacity impacts of migrating to 64-bit mode.

__ WSC Flash 10086 outlines the capacity planning methodology

__ Use the SOFTCAP Tool to evaluate the effect on z/Architecture and S/390 processor capacity when migrating to newer levels of software. Obtain from: http://www.ibm.com/support/techdocs/atsmastr.nsf/PubAllNum/PRS268

__ Evaluate and plan your paging environment

z/Architecture (64-bit) Migration Checklist

__ Review WSC White Paper WP100269, *z/OS Performance: Managing Processor Storage in a 64-bit Environment - V1.1* (Chapter 2) at
http://www.ibm.com/support/techdocs/atsmastr.nsf/PubAllNum/WP100269

__ Understand enhancements in RMF. All metrics related to expanded storage are obsolete.

__ Understand 64-bit Virtual Restrictions
- Data spaces > 2GB not supported
- Hiperspaces > 2GB not supported
- Subspace capability not extended to virtual above 2GB
- DIV not extended to virtual above 2GB
- Sharing virtual above 2GB not supported
- Copy-on-Write not supported for virtual above 2GB
- DAT tables for virtual above 2GB are not pageable

__ Verify virtual storage limits are set properly

__ Review IEFUSI exit for 64-bit virtual considerations.
Historically, users could limit program storage below 16 megabytes in virtual storage by using IEALIMIT. IEALIMIT can still be used to limit program storage in the non extended region; however, IEFUSI is the preferred exit routine. See *MVS Installation Exits*, SA22-7593-02, for details on the advantages over IEALIMIT.

__ Use of MEMLIMIT parameter of SYS1.PARMLIB member SFMPRMxx.
A new parameter has been added to the SMFPRMxx parmlib member. MEMLIMIT specifies the default value used by SMF jobs not having an explicit memory limit. See Washington Systems Center Flash10165 at:
http://www.ibm.com/support/techdocs/atsmastr.nsf/PubAllNum/Flash10165

__ Move to new High Level Assembler (High Level Assembler for MVS, VM &VSE V1R4) which supports 64-bit instructions.

__ Review RAS APAR OW54022, which describes changes to the thresholds at which SRM detects a common storage shortage below the line.

__ Install the z/OS V1R2/3/4 Bimodal Migration Accommodation on z/OS V1R4 so that you're ready to fall back from z/Architecture mode to ESA/390 mode on your z800 or z900 server if necessary.
Note: The z/OS V1R2/3/4 Bimodal Migration Accommodation is software that can assist you in the transition from ESA/390 architecture to z/Architecture by providing a fallback capability. For example, if you encounter problems after migrating to z/Architecture with z/OS on a z800 or z900 server, you can fall back to z/OS in ESA/390 mode on the z800 or z900 server (while you are within the terms of the Accommodation). Without the Accommodation, z/Architecture mode is the only mode possible when running z/OS on a z800 or z900 server. The Accommodation software is available for six months for each z/OS license starting from the registration of a z/OS license to a z/Architecture-capable server. It is initially made available concurrent with the general availability of z/OS V1R4 but it applies to z/OS V1R2 and V1R3 as well as V1R4. There is no additional charge for it. You can get the Accommodation only from the z/OS downloads Web page: http://www.ibm.com/eserver/zseries/zos/downloads/

z/Architecture (64-bit) Migration Checklist

Restriction: When you use the z/OS V1R2/3/4 Bimodal Migration Accommodation to run z/OS in ESA/390 mode on a z/Architecture (z800 or z900) server, functions that require z/Architecture are not supported, such as:

- Intelligent Resource Director (IRD)
- HiperSockets
- 64-bit real storage
- 64-bit virtual storage
- IPv6

In addition, Workload License Charges on z900 servers is not available when running z/OS in ESA/390 mode. To use the Accommodation after you download and install it, specify a value of 1 on the ARCHLVL parameter of parmlib member LOAD xx before you re-IPL. This will bring up z/OS V1R4 in ESA/390 mode on your z800 or z900 server. For example:

```
IODF 01 SYS0 TEST01
ARCHLVL 1
NUCLEUS 1
NUCLST 00
SYSCAT MCATDS123SYS1..MCAT RES01
IEASYM (01,L)
INITSQA 0064K
SYSPARM 01
```

To help you understand the appropriate ARCHLVL value, refer to the following table. The table shows the specifiable and default values by release and server. Notice that the default for the z800 and z900 changed from OS/390 R10 to z/OS.

OS/390 or z/OS release	IBM Server	ARCHLVL value you can specify	ARCHLVL value default
OS/390 V2R9	Any supported by OS/390 R9	None (not supported)	None (not supported)
OS/390 V2R10	Multiprise ® 3000, G5, G6, and prior servers supported by OS/390 V2R10	1	1
	z800 or z900	1 or 2	1
z/OS any release	Multiprise 3000, G5, or G6	1	1
z/OS V1R1	z800 or z900	2	2
z/OS V1R4 with the z/OS V1R2/3/4 Bimodal Migration Accommodation installed and within the six-month term	z800 or z900	1 or 2	2
z/OS V1R4 without the z/OS V1R2/3/4 Bimodal Migration Accommodation installed or after the six-month term has expired	z800 or z900	2	2
Note: ARCHLVL 1 = ESA/390 mode, ARCHLVL 2 = z/Architecture mode			

z/Architecture (64-bit) Migration Checklist

__ Switch to z/Architecure

> When you're ready to switch to z/Architecture, reconfigure the LPAR to use only real storage (formerly called central storage) rather than central and expanded storage. If you run in z/Architecture mode and still have expanded storage defined, the expanded storage is not used and the system issues a message.

> **Fallback note**: ESA/390 mode can still be IPLed successfully in an LPAR defined with greater than 2 GB real storage (with OS/390 V2R10 and z/OS V1R4). However, only the first 2 GB of real storage will be used, and the system issues a message System services that used expanded storage in ESA/390 mode (such as hiperspaces) have been changed to use real storage in z/Architecture mode. Programs that use these system services should not require any changes. The amount of central storage you should need is the sum of your current central and expanded storage. No additional processor storage is required by z/Architecture itself. You should need more storage only if your workloads grow.

__ Remember to remove the ARCHLVL parameter from PARMLIB member LOAD xx once your migration to z/Architecture is complete. If you don't remove it, and if you've installed the z/OS V1R2/3/4 Bimodal Migration Accommodation on z/OS V1R4, and if you specify ARCHLVL 1, your system will come up in ESA/390 mode rather than z/Architecture mode. To prevent such a mistake, the safest thing to do after your migration to z/Architecture is complete is to remove ARCHLVL entirely and allow the system to default to z/Architecture mode.