

Performance and Availability This release of z/OS delivers real-time capabilities for application performance and data availability. With the new VSAM Control Area (CA) Reclaim capability, applications that use VSAM key-sequenced data sets (KSDS) benefit from improved performance, minimized space utilization, and improved application availability. For example, IBM system tests indicate VSAM workload performance could improve by up to 44%+ or more, in addition to helping you avoid outages that used to be required to defragment and reorganize this data. Performance improvements are anticipated for many applications using CICS®, VSAM, IMS™ VSAM, and Catalog processing.

- There are significant performance enhancements expected for other workloads as well.
- For example, DB2 9 for z/OS (5635-DB2®) startup time could be reduced# and z/OS XMLSystem Services validation parsing performance is anticipated to improve by 30-50% †.
- A revolutionary new fragment validation parsing capability helps DB2 10 for z/OS (5605-DB2) pureXML™ avoid costly re-validation of whole XML documents.
- SVC dump capture time can be reduced by up to 50-90%+.
- Network performance is improved. Processing overhead for Application Transparent - Transport Layer Security (AT-TLS) can be improved by 30% †. Network throughput for interactive workloads can be improved by 30-50% † using the new Inbound Workload Queuing (IWQ) function, which is exclusive to OSA-Express3 on IBM zEnterprise™ 196 (z196) and System z10 servers. It is anticipated that IWQ will provide benefit for streaming workloads by reducing the number of costly network retransmissions due to out-of-order packets, and will benefit Sysplex Distributor traffic by providing a more streamlined processing path.

Smart Technologies z/OS is designed to learn from its own environment and can predict system abnormalities and alerting you about potential problems before they occur with its innovative Predictive Failure Analysis (PFA) capability.

- Updated monitoring of SMF data rates for z/OS V1.12, PFA can be used to help identify abnormal trends of your running system, allowing you to act before they become serious.
- In the rare event of a system issue, a new base component, z/OS Run Time Diagnostics, is designed to analyze key indicators on a running system quickly and help identify the root causes of system problems and degradations.
- The Run Time Diagnostics function is anticipated to run in as little as one minute to return results quickly enough to help you choose between alternative corrective actions more quickly and effectively.

The new z/OS The z/OS Management Facility (5655-S28) is the new face of z/OS →

- It can shave hours of time from system management tasks, and with embedded guidance and best practices built in, new system programmers can be productive in weeks capturing and analyzing incident data, configuring TCP/IP network settings and z/OS Workload Management (WLM) policies, and monitoring z/OS performance data have never been easier.

- z/OS Management Facility V1.12 is in "lock step" with z/OS V1.12 and is ready to help you use many of its latest features (this is a no-fee feature).

z/OS V1.12 provides **simplified configuration of FICON®-connected disk and tape** through the revolutionary **new z/OS FICON Discovery and Auto Configuration (zDAC)** capability.

- zDAC can automatically discover new or updated devices and propose channel configurations based on high availability best practices and your existing configurations, helping preserve Parallel Sysplex® symmetric configurations.
- This function is sophisticated and takes advantage of the robust I/O subsystem that is only available with the new IBM zEnterprise 196 server.

Connected and Ready IBM introduces the new **IBM zEnterprise System** a revolutionary direction in multi-tier, multi-system, multi-workload management and coordination. <- See #64 z/TidBits (zEnterprise - The Debut)

- The combination of z/OS and the zEnterprise System can provide workload management integration and insight (through integration with the zEnterprise Unified Resource Manager), designed to provide both high speed and highly secure networking (through the use of the physically isolated data and service networks), and infrastructure simplification (with fewer hops, devices, servers, and processors that complicate today's data centers).
- z/OS is designed to integrate *seamlessly* with the zEnterprise system - only a handful of new configuration settings are needed to integrate and execute in the new environment. In addition, DB2 9 for z/OS and the IBM Smart Analytics Optimizer for DB2 for z/OS (5697-AQT) are the first exploiters of the IBM zEnterprise BladeCenter® optimizers.
- The IBM Smart Analytics Optimizer is designed to help mitigate database administration costs while speeding up the performance of typical data warehouse or OLAP-specific workloads by an order of magnitude†.

The Best Solution Get the best of all worlds with z/OS V1.12 and the new **IBM zEnterprise System**.

Together, this revolutionary platform can extend the strengths of z/OS by deploying multi-tier workloads in a consolidated, simplified, and highly integrated environment bringing many new applications up close to z/OS.

Response Time Improvements based on IBM laboratory results:

- VSAM performance improvement is through the use of VSAM Control Area Reclaim, actual benefit may be *more or less* and will depend on the degree of VSAM data fragmentation. It is anticipated that VSAM KSDS that are severely fragmented or rarely reorganized will see more benefit.
- DB2 9 for z/OS startup time reduction is through z/OS Allocation, DFSMSdpp™, and GRS improvements, actual benefit will depend on number of data sets opened. It is anticipated that address spaces opening up many thousands of data sets will see more benefit.
- z/OS XML System Services validation parsing performance improvement will depend on the amount of data being parsed and the degree of complexity of the schema. Actual SVC dump time will depend on amount of data being captured and the amount of that data dumped from auxiliary storage.
- Actual AT-TLS CPU consumption improvement will depend on amount of data being transmitted and whether the workload is interactive or streaming. Throughput gain due to this improvement in CPU consumption is likely, but would vary depending on overall utilization of the z/OS image.
- Interactive networking response time improvement is via z/OS Communication Server exploitation of the new OSA Express 3 Inbound Workload Queuing. Actual benefit will depend on amount of data being transferred and if transfer is z/OS to z/OS, or z/OS to distributed system.
- Actual IBM Smart Analytics Optimizer CPU-intensive query performance improvement will depend on the specific query workload.

Predicting problems - z/OS Predictive Failure Analysis® (PFA) is planned to monitor the rate at which SMF records are generated.

- When the rate is abnormally high for a particular system, the system will be designed to issue an alert to warn you of a potential problem, potentially avoiding an outage.

Real-time decision making in the event of a system problem - A new z/OS Run Time Diagnostics function is planned to help you quickly identify possible problems in as little as one minute.

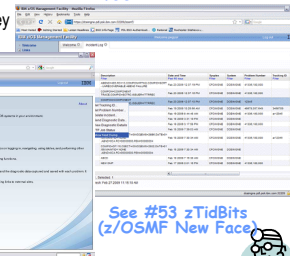
Automatic partitioning - GRS and XCF components are planned to automatically initiate actions to preserve sysplex availability to help reduce the incidence of sysplex-wide problems that can result from unresponsive critical components.

† IBM Lab Analysis



Sysplex Distributor is a combination of the high availability features of DVI/PA and the workload optimization capabilities of WLM.

z/OSMF



See #53 z/TidBits (z/OSMF New Face)

These checks can help you correct and prevent common sysplex management problems

- z/OS V1.12 is planned to run on these IBM System z® servers:
- z10™ EC
 - z10 BC
 - z9™ EC
 - z9 BC
 - z990*
 - z890*
 - z900*
 - z800*
- *These products are withdrawn from marketing

See #08 z/Nibbler Predictive Failure Analysis



Workload driven provisioning - Capacity Provisioning is to use CICS® and IMST™ monitoring data to determine if additional resources are needed to meet service-level requirements for these workloads.

Storage management and scaling - Extended Address Volumes are planned to support additional data set types, including sequential (both basic and large) data sets, partitioned (PDS/PDSE) data sets, catalogs, and BDAM data sets. Overall, EAV helps you relieve storage constraints as well as simplify storage management by providing the ability to manage fewer, large volumes.

Advanced cryptography - z/OS is planned to support **Elliptic Curve Cryptography (ECC)**, which is regarded by the U.S. National Security Agency (NSA) as a faster algorithm that requires a smaller key than RSA cryptography. This function is embedded into z/OS and is not a separately chargeable product.

IBM Health Checker for z/OS has a long history of helping to simplify and automate the identification of potential configuration problems before they impact system availability by comparing active values and settings to those suggested by IBM or defined by your installation. The z/OS Health Checker is extremely valuable not only in identifying exceptions to z/OS configurations, but also in identifying migration actions and checking that these migration actions are completed accurately. In addition, output reports from the z/OS Health Checker may be used to support your corporate compliance. For example, z/OS Health Checker reports can help identify unsecured resources that should be RACF® protected, can help validate the redundancy in a Parallel Sysplex configuration, and could be used as part of risk assessment exercises.

- For z/OS V1.12 z/OS Health Checker is planned to be updated with the ability to write checks in Metal C, and with the addition of checks for Parallel Sysplex (such as best practices for coupling facility structure size, couple data set specification limits, Sysplex Failure Manager policies, and coupling facility allocation), SMB server, DFSMSTM, I/O Supervisor, TCP/IP IPv4 and IPv6 usage, HFS to zFS migration, and still others.

Duplicate Temporary Datasets - There are additional ease-of-use enhancements planned to help prevent JCL errors from duplicate temporary data set names, simplify Language Environment® and zFS migration, simplify RMFTM processing, improve performance and create a customized view for the Library Server.

Ease-of-use and platform-simplification enhancements:

The following functions are planned for **z/OSMF V1.12:**

- The **z/OSMF Configuration Assistant for z/OS Communications Server** is planned to:
 - > Support the configuration of IKE version 2.
 - > Enforce RFC4301 compliance for IPSec filter rules.
 - > Support the configuration of certificate trust chains and certificate revocation lists.
 - > Support the configuration of new cryptographic algorithms for IPSec and IKE.
 - > Support the configuration of FIPS 140 cryptographic mode for IPSec and IKE.

* **WLM Policy Editor** functionality to be integrated into z/OSMF V1.12 will facilitate the creation and editing of WLM service definitions, installation of WLM service definitions, and activation of WLM service policies, and monitoring of the WLM status of a sysplex and the systems in a sysplex.

* **Incident Log** A number of improvements for the Incident Log function including support for encryption of all incident files, including dumps, to be sent to IBM, breaking dumps into multiple data sets that can be sent via FTP in parallel to reduce transmission time, specifying additional data sets to an incident and adding free-form comments to new fields for problem descriptions and FTP destinations. Incident Log will also support the creation of diagnostic log snapshots based on the SYSLOG and LOGREC data sets, as well as the OPERLOG and LOGREC sysplex log streams. These enhancements are intended to help you manage problem data more easily.

The **Health Checker** framework is enhanced to allow health checks to be registered without a message table and for them to issue messages directly without using a message table. This makes it easier to write health checks quickly.

- **New health checks** are planned for the **Parallel Sysplex** components, XCF and XES. They are designed to warn you when a coupling facility structure's maximum size as specified in the CFRM policy is more than double its initial size, when any couple data set's (CDS's) maximum system limit is lower than the primary sysplex CDS's system limit, when shared CPs are being used for coupling facility partitions, when the CFRM message-based event management protocol can be used for CFRM event management but the policy-based protocol is being used instead, when your Sysplex Failure Management (SFM) policy does not specify that automatic actions are to be taken to relieve hangs caused by the unresponsiveness of one or more of a CF structure's users, and when a CF does not have a designated percentage of available space to allow for new CF structure allocation, structure expansion, or CF failover recovery.
- **New health checks** are designed for the **I/O Supervisor (IOS)**. IBM recommends using the relatively new MIDAWs and Captured UCB Protection functions introduced in recent releases, and locating eligible I/O-related control blocks above the 16 MB line. These health checks are designed to notify you when these functions are not being used, to help you manage system performance and the use of virtual storage.

The **Health Checker** started task is planned to support running with an assigned user ID that has access to the **BPX.SUPERUSER** profile in the FACILITY CLASS. This will make it unnecessary to run the Health Checker address space with a user ID having UID(0).

IBM recommends that you **use zFS file systems for z/OS UNIX® System Services**. In z/OS V1.12, a migration health check is planned to identify HFS file systems you should consider migrating to zFS file systems. This is intended to help you easily obtain and track the list of remaining file systems to be converted.

The **Interactive Storage Management Facility (ISMF)**, used to manage your SMS configuration, allows you to copy storage group definitions from one control data set (CDS) to another. In z/OS V1.12, ISMF is extended to allow you to specify that the volume list for pool-type storage groups be copied at the same time. This allows you to copy entire storage groups from one configuration to another without having to add their volumes to the destination CDS afterward.

In prior releases, partial release operations for **VSAM data sets supported releasing space only** on the last volume containing data for each data set. In z/OS V1.12, partial release is planned to be extended to support releasing unused volumes in addition to releasing space on the last volume of a multivolume VSAM data set that contains data.

The **IDCAMS DEFINE RECATALOG** command is planned to be enhanced for multivolume and striped data sets.

IDCAMS is planned to be enhanced to allow you to **delete all members of a partitioned data set in a single operation** by specifying a wildcard character (*) as the member name for a data set when using the DELETE command. This new support allows you to remove all members of a PDS or PDSE data set in a single command.

In z/OS V1.9, support was added to write SMF data to log streams. In z/OS V1.12, **RMF** is planned to be enhanced to read SMF records directly from a log stream. This is intended to allow you to eliminate any intermediate steps you currently use to unload SMF data from a log stream to a sequential data set for RMF post processing.

SDFS is planned to augment the **CK panel** by displaying recorded checks on a new health check history panel.

A number of enhancements are made to the processing of **PROGxx parmib members** and to **Link List Lookaside (LLA)** processing. These include support in **PROGxx** for passing a specified parameter to a dynamic exit, automatically including alias names for modules to be placed in Dynamic LPA, and specifying volumes on **SYSLIB** for data sets so they need not be cataloged in the master catalog.

A new **SUMMARY** keyword of the **DISPLAY SYMBOLS** command is designed to provide summary information about symbols used of the system, including how many are in use.

When a corrupt PDSE is detected in the link list during IPL, the system enters a wait state. In z/OS V1.12 the system will be designed to issue a message identifying the corrupt PDSE prior to entering the wait state. This allows the user to attempt to restore the corrupt PDSE and re-IPL the system and avoid taking a standalone dump to debug the problem.

System Logger is planned to be enhanced to correct the VSAM SHAREOPTIONS for new log stream data sets when it detects that they are not correctly set. Messages are planned to indicate that Logger has detected and corrected a data set's SHAREOPTIONS settings. This new function is intended to prevent data set access problems from arising when SHAREOPTIONS(3,3) has not been set in the data class used to allocate log stream data sets.

System Logger is planned to be enhanced to support log data set sizes up to 4 GB (from the previous 2 GB limit). This applies to both OFFLOAD and STAGING data set types. As part of this support, System Logger is planned to add messages to show key data set characteristics at allocation and deletion time.

Scalability

DFSMS is planned to support additional data set types, including sequential (both basic and large) data sets, partitioned (PDS/PDSE) data sets, catalogs, and BDAM data sets in the extended addressing space (EAS) on an EAV. Support is also included for generation data groups (GDGs) and VSAM volume data sets (VVDSSs). Overall, EAV helps you relieve storage constraints as well as simplify storage management by providing the ability to manage fewer, large volumes as opposed to many small volumes.

Support to make all data sets used by DFSMSmm eligible for allocation in the extended addressing space of an EAV. This includes the DFSMSmm journal and dynamically allocated temporary files.

In z/OS 1.12, DFSMSmm support for IPv6 is also planned.

Language Environment provides support for C/C++ to access alternate indexes (AIXs) for extended format VSAM key-sequenced data sets (KSDSs) that reside in the EAS on an EAV. Some workloads require an increasing number of open data sets. In z/OS V1.12, the BSAM, QSAM, and BPAM (basic and queued sequential, and basic partitioned access methods) and EXCP (execute channel program) processing will be designed to support the use of an extended task I/O table (XTIOT) with uncaptured UCBs, and support data set association blocks (DSABs) above the 16 MB line. This is expected to allow more data sets to be allocated by an address space and to provide virtual storage constraint relief for DASD and tape data sets.

OAM is planned to provide API support for the Object Storage and Retrieval function (OSR) to run in a CICS threadsafe environment. This is intended to allow exploiters to take advantage of the improved multitasking and throughput capabilities provided by threadsafe programming.

Large (1 MB) pages were introduced in z/OS V1.10. In z/OS V1.12, the nucleus data area is planned to be backed using 1 MB pages. This is intended to reduce the overhead of memory management for nucleus pages and to free translation lookaside buffer (TLB) entries so they can be used for other storage areas. This is expected to help reduce the number of address translations that need to be performed by the system and help improve overall system performance.

DFSMS support for catalogs with extended addressability (EA) is planned. This will be designed to make it possible to define and use Integrated Catalog Facility (ICF) Basic Catalog Structures (BCS) with EA, allowing catalogs larger than 4 GB.

z/OS Communications Server AT-TLS processing will be designed to provide reduced CPU usage for encryption and decryption of application data while improving throughput for some types of workloads.

VSAM record level sharing (RLS) is planned to support striped data sets. This will be designed to bring the benefits of VSAM striping, such as allowing single application requests for records in multiple tracks or control intervals (CIs) to be satisfied by concurrent I/O requests to multiple volumes. Using striped data sets can result in improved performance, by transferring data at rates greater than can be achieved using single I/O paths.

DFSMSdss™ will be designed to use larger blocks when possible for DUMP, COPYDUMP, and RESTORE operations, and to support Extended Format Sequential dump data sets on DASD for DUMP, RESTORE, and COPYDUMP. The use of larger block sizes is intended to improve performance for these operations, and using Extended Format dump data sets is intended to support striping and compression.

DFSMShsm™ will be designed to support parallel processing for recovery from dump tape volumes when the dumps reside on multiple tape volumes and multiple tape drives are available. This new function is intended to allow you to specify that up to 64 concurrent tasks be used.

PDSE processing is planned to be changed to reduce delays that can occur when two systems are accessing a PDSE concurrently while it is being updated. PDSE will be designed to improve its cross-system sharing capabilities, including member-level sharing, within a GRG complex but outside a Parallel Sysplex. These changes are intended to make PDSEs more usable outside single-system and Parallel Sysplex environments.

The Catalog address space (CAS) will be designed to check for SYSZTIOT enqueue contention periodically. Based on an interval you specify and the reason for contention, CAS will be designed to write a logrec record and a notification message when tasks have waited longer than the specified interval and contention checking is active.

Two new services based on existing XCF signaling services are planned to be introduced to support the use of 64-bit addressable virtual storage message buffers and associated input and output parameters.

The DFSMSHsm DUMP function used to copy source disk volumes to a target tape volume is planned to be enhanced. The dump stacking function will be designed to allow up to 255 source volumes to be dumped to a single tape volume, up from the prior limit of 99.

z/OS Communications Server TN3270E Telnet server plans to provide access method control block (ACB) sharing for logical units (LUs) as a way to help reduce ECDSA usage. Prior to z/OS V1.12 Communications Server every Telnet LU name opened its own ACB to VTAM®. You can code a new SHAREACB statement to allow multiple Telnet LUs to share a single ACB.

Standalone Dump is designed to support extended format dump data sets in the extended addressing space (EAS) on Extended Address Volumes (EAVs).

Superzap (AMASPZAP) is planned to support dumping and altering data for sequential, partitioned, and direct data sets placed in EAS on EAVs.

Application Integration

z/OS V1.12 to include the following updates: enhancements to C/C++ in support of Euro currency, new standard time services, and Unicode; and enhancements to z/OS XML System Services in support of schema extraction and fragment parsing.

SDSF will be designed to make Java classes available, to provide a new means of accessing SDSF. Classes will be provided for each of the SDSF panels that can be used by applications to request SDSF functions. This new support is designed to allow Java-based applications to easily access SDSF.

SDSF is planned to introduce a new ISFLOG command for SDSF REXX. It is designed to read the system log and return its records in stem variables, and to support options to limit the number of records returned and specify start and end times. This new function will simplify access to the system log for SDSF REXX.

The Program Management Binder will be designed to allow you to specify that a specific residency mode (RMODE) be applied to all initial load classes of a program object, rather than the classes in the first segment containing the entry point. This new function is intended to offer application programmers more flexible options for program storage residency.

The Program Management Binder is to make program object attribute data (PMAR data) available to programs using the fast data interface, and to support programs loaded using the z/OS UNIX System Services load service (loadhfs).

† IBM Lab Analysis



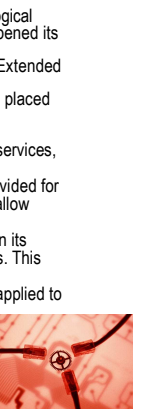
z/OS has a huge breadth of security capabilities built in to the base of the operating system at no extra cost

ICSF is unique and could be considered more secure than other cryptographic solutions because it can manage the encryption and decryption of sensitive material without exposing the keys in clear.

The new Application Transparent Transport Layer Security (AT-TLS) support in z/OS can allow enterprises to enable network security on behalf of TCP applications without requiring application modification. The z/OS IPsec support provides additional options which include the ability to encrypt data end-to-end, or across just a portion of the network.

The z/OS Com Server is there to meet new challenges with a wide Array of networking technologies supported (including both TCP/IP and SNA).

z/OS Com Server is enhanced for the performance of fast local sockets for TCP connections.



z/OS XML System Services will be updated to enhance XML schema validation support by allowing applications to extract schema location information from an XML instance document without the application first performing a separate parse. This is planned to improve the usability of the validating parsing interface and intended to reduce the processing cost of obtaining this information.

z/OS XML System Services is planned to be updated to allow you to validate part of an XML document when performing validating parsing, rather than the entire document. Called fragment parsing, this capability is intended to reduce the processing cost of performing validation by allowing you to validate only a portion of a document rather than requiring the validation of the entire document. For example, this can be useful when only a subset of a large document containing multiple fragments has changed.

z/OS XML System Services will be updated to provide a new validating parse capability that allows applications to restrict the set of element names to be accepted as valid root elements to a subset of those allowable in an XML schema. This is intended to provide an additional level of validation capability beyond that provided by the W3C schema language.

Previously, the tsocmd shell command was available only from the Tools and Toys section of the z/OS UNIX System Services Web site. In z/OS V1.12, z/OS support for this function is planned. Unlike the existing tso command, the tso cmd command can be used to issue authorized TSO/E commands.

Support is planned in z/OS UNIX System Services for the record file format in the cp, mv, ls, pax, and extattr shell commands as well as the ISHELL command. In addition to binary and text format, files can be handled in record file format. z/OS applications accessing these files by using QSAM, BSAM, VSAM, or BPAM and coding FILEDATA=RECORD will be able to take advantage of the record file format to read and write data as records.

z/OS UNIX System Services supports the memory mapping (mmap) function for files in zFS and HFS file systems. In z/OS V1.12, support is planned to allow applications to use memory mapping for NFS Client files. This will enable NFS-mounted file systems to be used by applications that use memory mapping.

A new option is planned for the ISENGQ service that can be used to serialize resources. This new support will be designed to allow an unauthorized program to interrupt serialization processing and opt not to continue to attempt to obtain control of a resource when the resource is not available or to do other work asynchronously while waiting to obtain an ENQ resource. For example, a programmer might wish to set a time limit for obtaining control of a resource. This is expected to help programmers to better manage contention delays and remove pending enqueue requests in recovery.

Security

Encryption obscures information and is intended to make it unreadable to unauthorized parties. It can be used to protect the confidentiality, integrity, and availability of both data at rest and data being transmitted, and in general remains one of the strongest aspects of IT security. z/OS is the logical choice for cryptography and storing and managing the cryptographic keys due to the nature of key handling by z/OS Integrated Cryptographic Service Facility (ICSF).

z/OS PKI Services is planned to be enhanced with several usability enhancements which are expected to reduce the amount of time and number of manual tasks associated with finding certificate serial numbers, and issuing renewal and revocation e-mails.

Plans to have significant updates for Tivoli Directory Server (LDAP) in support of new password policy rules, improved logging, and new extensions for access control lists.

There are a number of improvements planned for PKI Services. IBM Tivoli Directory Server for z/OS is planned to provide support for configurable password policy rules that can be applied to user passwords in the directory.

IBM Tivoli Directory Server for z/OS is planned to support continuous activity logging. IBM Tivoli Directory Server for z/OS is planned to provide an extension to access control lists (ACLs) to provide the ability to dynamically transform base ACLs using filter ACLs you specify.

TSO/E will be designed to accept passwords that include one or more special characters. This is intended to leave the checking for acceptable password characters to an external security manager such as RACF.

z/OS Communications Server is planned to introduce trusted TCP connections, to enable sockets programs to retrieve sysplex-specific connection routing information and partner security credentials for connected sockets.

Availability

Predictive Failure Analysis (PFA) monitors the rate at which the system is generating SMF records. When the rate is abnormally high for a particular system, the system will be designed to issue an alert warning you of a possible problem, potentially avoiding an outage.

z/OS Run Time Diagnostics, is planned to help when the need for quick decision-making is required. With Run Time Diagnostics, your z/OS system will be designed to analyze key system indicators of a running system. The goal is to help you identify the root of problems that cause system degradation on systems that are still responsive to operator commands.

A new Timed Auto Reply Function will enable the system to respond automatically to write to operator with reply (WTOR) messages. This new function is expected to help provide a timely response to WTORs and help prevent delayed responses from causing system problems.

Parallel Sysplex technology is planned to be updated with new health checks; improved command routing; and improved network traffic routing, security, availability and reporting. There are also plans to provide autonomic whereby the z/OS system can help identify CF structures and network connections that are unresponsive or in a degraded state.

Optimization and management capabilities

WLM is planned to be updated with enhancements to improve batch management. Provisioning Manager for z/OS enables z/OS and the System z10 server to add temporary capacity automatically when necessary, with or without operator intervention. Capacity Provisioning for z/OS V1.12 is planned to use CICS and IMS monitoring data to determine if additional resources are needed to meet service level requirements for these workloads. What has taken minutes or hours to discover, identify, decide, and resolve, now can happen automatically in seconds.

z/OS V1.12 is updated so IDCAMS can avoid DFSMSHsm recalls when deleting generation data groups (GDGs). To help better understand the resources consumed by batch jobs and improve the accuracy of chargeback programs, z/OS V1.12 will be designed to record the CPU time consumed for job steps in initiator address spaces using new fields in SMF Type 30 records.

WLM will be designed to consider resource group maximums and the projected increase in system or sysplex demand before starting initiators during resource adjustment and policy adjustment processing when the service class has been assigned to a resource group and a resource group maximum has been defined.

WLM changes to the dispatching of discretionary work. The system will be designed to run discretionary work for a longer period of time before dispatching other discretionary work, while still interrupting it after short periods for nondiscretionary work. This change is intended to help improve the throughput for systems with a high percentage of discretionary workloads.

New PDSE functions. A new utility will be designed to verify that the structure of a PDSE is valid, and programming services will be designed to perform similar checking to help programs verify the state of a PDSE before and after critical operations.

z/OS Communications Server provides the option to check the health of an Enterprise Extender (EE) connection during the activation of the connection. The health of active connections can also be verified.

z/OS Communications Server provides notification to the operator console when a Domain Name System (DNS) name server does not respond to a certain percentage of resolver queries sent to the name server during a sliding five-minute interval.

z/OS Communications Server provides the option of keeping a TCP/IP stack isolated from the sysplex; you can use a new config parameter to prevent a stack from automatically joining the sysplex group at startup and join the sysplex group at a later time. "SOD" In a future release, IBM plans to remove the capability to change the default Language Environment run-time options settings via SMP/E installable USERMODs. IBM recommends using the CEEPRMxx parmib member to change the default Language Environment run-time options for the system.