



Preview: z/OS V1.12 - Heralding a new generation of smart operating systems

Table of contents

| | | | |
|----------|---|-----------|--|
| 1 | Overview | 3 | Description |
| 2 | Key prerequisites | 26 | Statement of general direction |
| 3 | Planned availability date | 27 | Product number |
| | | 29 | Corrections |

At a glance

IBM® previews z/OS® Version 1.12. With this latest release of z/OS, IBM heralds a new area of smart operating systems by creating an environment that can proactively work for you to help promote improved operations, availability, manageability, and security through innovative self-learning, self-managing, and self-optimization capabilities. Enhancements include:

- 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.
- Avoiding data fragmentation and planned outages for data reorganizations - With the new CA (Control Area) Reclaim capability, applications that use VSAM key-sequenced data sets (KSDSs) can benefit from improved performance, minimized space utilization, and improved application availability through the avoidance of planned outages that used to be required to defragment and reorganize this data.
- Workload driven provisioning - Capacity Provisioning is planned to use CICS® and IMS™ 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.

Overview

Imagine, an IT system that knows your priorities and can make suggestions - even decisions - that can benefit your business. IBM previews z/OS V1.12. With this latest release of z/OS, IBM heralds a new direction of smart operating systems by creating an environment that can proactively work for you to help promote improved

operations, availability, and manageability through innovative self-learning, self-managing, and self-optimization capabilities.

z/OS is designed to learn heuristically from its own environment to anticipate and report on system abnormalities, predicting problems before they occur with its innovative Predictive Failure Analysis (PFA) capability. For example, PFA can be used during application testing to identify previously unknown potential problem areas before the application is put into production or can be used in production systems to help identify issues before they become serious.

In the event of an issue, z/OS can help you be responsive with real-time decision making assistance. A new z/OS Run Time Diagnostics function is planned to analyze key indicators on a running system quickly, and help identify the root causes of system 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 and help you maintain high levels of system and application availability.

In some situations, your operations may be so critical that human analysis and intervention may not be fast enough, and the system must have the ability to act quickly and decisively. In a Parallel Sysplex®, the GRS and XCF components are planned to have the ability to automatically initiate actions to preserve sysplex availability so as to help reduce the incidence of sysplex-wide problems that can result from unresponsive critical components. The system can take action to fence, or stop and start critical members automatically, preventing small problems from becoming major problems.

And in still other situations, z/OS keeps your system available automatically and transparently. z/OS will be designed to avoid data fragmentation and planned outages for data reorganizations. With the new CA (Control Area) Reclaim capability, applications that use VSAM key-sequenced data sets (KSDSs) can benefit from improved performance, minimized space utilization, and improved application availability through the avoidance of planned outages that used to be required to defragment and reorganize this data.

z/OS V1.12 can save you time and money. This ability to discover, decide, and resolve issues automatically and in a fraction of the time can keep your organization nimble and responsive to changing business needs.

Key prerequisites

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.

For a complete description of z/OS V1.12 software prerequisites, refer to *z/OS V1R12 Planning for Installation (GA22-7504)*, when available.

Planned availability date

September 2010

Previews provide insight into IBM plans and direction. Availability, prices, ordering information, and terms and conditions will be provided when the product is announced.

Description

Ease of use

Simplification of an IT system has many rewards. It can address the need for skills by making existing personnel more productive and by reducing the time needed for someone new to gain proficiency on the platform. It can address overall operational efficiency by reducing the components and steps for tasks, and by streamlining existing processes. And it can address quality of service and availability by reducing the time involved with addressing a problem, or by reducing the probability of the error even being introduced.

Ultimately, simplification can make your IT organization more responsive in meeting business needs because IT systems and processes are less apt to get in the way when action and agility are needed. This is IBM's long-term goal for mainframe simplification. More than a "screen scraper" or a pop-up installation shield and more than new layers of management processes, IBM has taken the long-term outlook by simplifying a mainframe system from the inside out and from end to end. IBM technologies are truly efficient and can help drive down the cost of complexity, reduce the cost from risk, and drive up user productivity and overall system agility.

IBM's commitment to mainframe simplification has been vast and has been delivered integrated into the platform stack. CICS Explorer provides CICS architects, developers, system programmers, and administrators a common tooling environment, with integrated access to a wide range of data and control capabilities. DB2® Data Studio provides an integrated set of tooling to support all phases of the data management life cycle. IMS is planned to provide a new integrated development environment and operational console to accelerate the development time for new IMS applications and optimize collaboration between IMS DBAs, system programmers, and application developers. Rational® Developer for System z helps simplify collaboration, development, and delivery of business applications and integrate existing core business applications with Web services and SOA. Tivoli® Service Management Center provides a set of integrated solutions and building blocks that allows a business to implement an enterprise-wide service management and process automation hub on System z. These technologies have the power to reduce application development, deployment, and management times significantly.

z/OS has had many improvements in the area of simplification as well. The past several releases of z/OS delivered improvements in the areas of simplifying diagnosis and problem determination; network and security management; and overall z/OS I/O configuration, sysplex, and storage operations. These improvements are designed to help simplify systems management; to improve application programmer, system programmer, and operator productivity; and to provide fewer opportunities for the introduction of human errors.

- The z/OS Management Facility (z/OSMF, 5655-S28) is the new face for z/OS and it provides support for a modern, Web-browser-based management console for z/OS. Automated tasks and wizards can guide users through tasks and help provide simplified operations. In z/OSMF V1.11, for example, tasks taking up to 20 minutes, such as collecting and packaging dump data, can now take as little as 30 seconds.

For its next release, z/OSMF V1.12 is planned to be expanded with the addition of z/OS Workload Manager Policy Editor functionality, enhancements to the already valuable Incident Log and Configuration Assistant for the z/OS Communications

Server functions, and the ability to add non-z/OSMF application launch points and links.

- 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, DFSMS™, I/O Supervisor, TCP/IP IPv4 and IPv6 usage, HFS to zFS migration, and still others.

- 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 RMF™ processing, improve performance, and create a customized view for the Library Server.

Details on the ease-of-use and platform-simplification enhancements intended for z/OS V1.12:

- 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.
 - 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.
 - An interface to allow the addition of non-z/OSMF launch points and links to the navigation tree.
 - z/OSMF is planned to support Microsoft® Internet Explorer 7 and Internet Explorer 8, Mozilla Firefox 3.0, and Firefox 3.5.
- 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.
- The Health Checker framework is planned to provide headers to enable you to write health checks using Metal C, in addition to existing support for High-Level Assembler and REXX™. Providing high-level language support can make it easier

to write complex health checks. Additionally, sample health checks written in METAL C are planned.

- 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. These checks can help you correct and prevent common sysplex management problems.
- SMB server is planned to add two health checks. The first is designed to detect SMB running in a shared file system environment and alert you that SMB cannot export zFS sysplex-aware read-write file systems in this environment, and the second to determine whether SMB is configured to support the RPC protocol (DCE/DFS) and display a message to remind you that IBM plans to withdraw support for this protocol in a future release.
- DFSMS is planned to add new health checks for the communications and active configuration data sets (COMMDS and ACDS). One new check is designed to alert you that the COMMDS and ACDS are on the same volume. The other is intended to identify COMMDS and ACDS data sets that were defined without the REUSE attribute, which is recommended. These new checks can help you manage your SMS environment.
- 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).
- z/OS Communications Server is planned to enhance the z/OS Health Checker for z/OS by adding two new checks: one check for IPv4 routing and one check for IPv6 routing. The checks determine whether the total number of indirect routes in the TCP/IP stack routing table exceeds a maximum threshold (default 2000). When this threshold is exceeded, OMPROUTE and the TCP/IP stack can potentially experience high CPU consumption from routing changes.

Two new maximum threshold parameters are planned to override the default values for the total number of IPv4 and IPv6 indirect routes in a TCP/IP stack routing table before warning messages are issued.

- 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.
- When two or more jobs having the same job name begin to process within the same system clock second and specify the same temporary data set names, the second and subsequent jobs will fail with JCL errors while attempting to allocate data sets with duplicate names. In z/OS V1.12, you will be able to use a new parmlib option to specify that the system use the data set naming convention for unnamed temporary data sets instead, which substantially reduces the probability of this JCL error without the need to change JCL.
- In z/OS V1.7, Language Environment allowed overridable run-time options to be defined in a new CEEPRMxx member of parmlib. In z/OS V1.12, this is extended to add support for non-overridable (NONOVR) options. This new support can allow you to specify the options for Language Environment without user modifications, eliminating a repetitive migration task.

- DFSMSrmm™ is planned to be enhanced for z/OS V1.12. The reason why a DFSMSrmm retention limit was reached is planned to be added to the ACTIVITY file. This function is also available now for z/OS V1.10 and z/OS V1.11 with the PTF for APAR OA30881. New reports created from the ACTIVITY and extract files are planned to help you see why retention limits were triggered. Also, OPENRULE ignore processing is planned to be available for duplicate tape volumes and support is planned to allow you to set a volume hold attribute to prevent expiration and to search and report on volumes which have the hold attribute. It is also planned that the DFSMSrmm ISPF dialog search results can be bypassed when using the CLIST option.
- DFSMS plans to provide a system option to control how the system handles multivolume tape label anomalies. This means that you can now prevent applications processing tape volumes out of sequence without coding an installation exit.
- 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.
- The JESXCF component is changed to allow you to log on to multiple systems within a sysplex using the same TSO/E user ID.
- DFSMSdfp™ is planned to allow a zFS data set to be recataloged with an indirect volume serial or system symbol. This is designed to allow the zFS file systems used for z/OS system software files (called version root file systems) to be cataloged using an indirect volume serial or a system symbol the same way as non-VSAM data sets to make cloning and migration easier.
- 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. This new function will be designed to automatically create catalog entries with correctly ordered volume lists while eliminating any duplicate volumes that might have been specified. This will make it easier to recatalog multivolume and striped VSAM 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.
- The Capacity Provisioning Control Center is planned to support displaying provisioning reports supported by the Capacity Provisioning Manager. This is intended to simplify the investigation of Capacity Provisioning reports and operation of the Capacity Provisioning server.
- 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 postprocessing.
- The Capacity Provisioning Manager client is planned to be updated to provide support for Windows® Vista.
- Library Server is designed to improve performance when building new catalogs and supporting multiple users on heavily loaded systems. A new Personal BookCase function in Library Server is planned to allow you to create, use, and share your own subset of the documents from a Library Server catalog. This function is designed to allow you to configure a Personal BookCase that includes the shelves and documents, as well as the infocenters and topics, that you are interested in so you can have the reference documents you routinely use available quickly. Also, indexing is planned to capture the author's intended definition of primary nodes for an InfoCenter's Table of Contents, and planned administrative improvements include long filename support, programmatically checking for the

required level of Java™, and generation of a new Test and Diagnostics page for use by Library Server administrators and IBM support personnel.

Library Server usability enhancements are planned for user interfaces, including improved navigation between certain dialogs, modernized icons, and descriptive hover popups for documents on a shelf.

- SDSF is planned to augment the CK panel by displaying recorded checks on a new health check history panel. The default will be to display up to ten prior iterations of each check from the log stream, and support is planned to allow you to browse and print check output from the history panel as you can on the primary CK panel.
- SDSF is designed to support displaying information about printers for JES3, and to eliminate the requirement for WebSphere® MQ when displaying JES2 MAS-wide data on the initiator panel for JES2 once all systems in the MAS are at z/OS V1.12 JES2. Also, displaying MAS-wide data on the printer panel for JES2 is planned not to require WebSphere MQ when all systems in the JES2 MAS are at or above z/OS V1.11 JES2.
- In z/OS V1.12 a new DISPLAY XCF,REALLOCATE,TEST option is planned to simulate the reallocation process and provide information about changes the REALLOCATE command would attempt to make, and any errors that might be encountered if an actual REALLOCATE process were to be performed. This capability is intended to provide information you can use to decide when to invoke the actual REALLOCATE process, and also whether you may need to make any coupling facility configuration changes before issuing the actual REALLOCATE command. A new DISPLAY XCF,REALLOCATE,REPORT command is also planned, to provide detailed information on the results experienced by a previously executed REALLOCATE command. This capability is intended to help you find such information without searching through the system log for REALLOCATE-related processing and exception messages.
- A number of enhancements are planned to be made to the processing of PROGxx parmlib 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 REPLACE option for exits to assure there is no window during which an exit is unavailable; and a new SVCNUMDEC keyword to specify the SVC number to be added.
 - Additionally, a new DEFAULTS statement is planned, so you can specify processing defaults intended to help prevent common errors. This includes allowing you to specify that LNKLIST DEFINE always require COPYFROM, that it default to COPYFROM(CURRENT), and that it automatically process aliases for modules added to Dynamic LPA.
 - LLA processing will be designed to support the use of dynamic LLA exits and to process multiple MODIFY commands in parallel.
- A new SUMMARY keyword of the DISPLAY SYMBOLS command is designed to provide summary information about symbols used on the system, including how many are in use. This can help you determine how many additional symbols can be defined.
- 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. This support is

planned to be made available for z/OS V1.9, z/OS V1.10, and z/OS V1.11 with PTFs for APAR OA30548 in February 2010.

Scalability and performance

The traditional view on scalability and performance has been to throw more hardware at something, or to wait and upgrade to faster hardware. This hardware-centric approach has worked for many years with the introduction of ever-larger distributed clusters and storage arrays, and higher-speed and denser chip designs. But the industry has begun to hit fundamental physical limits for chip design. Large-magnitude CPU speed increases with each generation of chip are a thing of the past and capacity increases will increasingly come not from raw hardware capabilities, but from a deeper type of technical alignment.

IBM System z has long understood the balance between scalability and performance and efficiency of the platform. The major components of the system, the processors, storage, I/O, and software, work together and help manage system resources. Essentially, z/OS and its subsystems provide for scalability not only based on faster chip speeds, but also via efficient single-image n-way processor growth, highly scalable sysplex clustering for horizontal growth, and scalable storage and data management as well.

z/OS has had many scalability/performance improvements over the past several releases. For example, z/OS V1.9 HiperDispatch can provide significant performance gains for large LPARs through smarter dispatching of workloads on higher n-way systems, and with z/OS V1.10 XL C/C++ workloads gained up to 8% performance improvements with new compiler options and System z10® prefetch capabilities.

- Parallel Sysplex is many clustering solutions in one. A single cluster can be used for scalability, performance, availability, software migrations, and disaster recovery. While other platforms are just beginning to grasp the cloud concept, Parallel Sysplex has been providing a dynamic cloud-like environment, where resources and workloads can seamlessly move to where they are needed, for over a decade. Parallel Sysplex provides a large single system image, dynamic load balancing, fault tolerance, and automatic restart capabilities. No other technology can compare -- other coupling capabilities are implemented in software, or loosely linked with non-integrated tools. With z/OS V1.12, Parallel Sysplex technology is planned to be updated with support for larger coupling facility structures.
- The scale and efficiency of System z do not end with the server. The amount of data being stored by organizations is going up exponentially. Much of this has to do with the wide variety of data formats and streams that are available, but a good part of the explosion of data is probably from management (or mismanagement) of a tremendous amount of data. The more data there is, the greater the need for availability, scalability, security, and networking, and the higher the risk from storage outages. Data on z/OS can help alleviate these problems. Data Facility Storage Management Subsystem (DFSMS) is a software suite that automatically manages data from creation to expiration and provides a consistent, policy-driven approach to storage management across the storage hierarchy. DFSMS provides allocation control for availability and performance, backup/recovery and disaster recovery services, space management, tape management, and reporting and simulation for performance and configuration tuning. DFSMS can help you drive storage utilization and efficiency up to well over 90%. With z/OS V1.12, DFSMS supports additional data set types in Extended Address Volumes (EAVs). EAVs can help relieve storage constraints as well as simplify storage management by providing the ability to manage fewer, large volumes as opposed to many small volumes.
- z/OS V1.12 also is planned to have updates for constraint relief for large volumes of DASD and tape data sets and concurrently open data sets, with new designs in the Program Management Binder, TSO/E, RACF, OAM, DFSMS, XCF, and InfoPrint® Server for z/OS. In addition, numerous improvements to dump management are planned to address the continued growth in diagnostic data that comes from larger systems and larger programs using ever-larger amounts of memory. These improvements can help you keep dump time and dump transmission time under control.

Details on the performance and scalability enhancements intended for z/OS V1.12:

- 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 (VVDSs). 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 is planned to make all data sets used by DFSMSrmm eligible for allocation in the extended addressing space of an EAV. This includes the DFSMSrmm journal and dynamically allocated temporary files.
- In z/OS 1.12, DFSMSrmm 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.
- JES2 will be designed to allow both spool and checkpoint data sets to reside in the EAS on an EAV, making it possible to place both spool and checkpoint data sets anywhere on an EAV and to define spool data sets up to the maximum size of 1,000,000 tracks (approximately 56 GB).
- JES3 will be designed to allow spool, checkpoint, and Job Control Table (JCT) data sets to be placed anywhere 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.
- The SNAP/SNAPX services and dump processing (including that for SVC, SYSABEND, SYSMDUMP, and SYSUDUMP dumps), and the AMASPZAP program are planned to support XTIOT.
- The Program Management Binder will be designed to support data sets having XTIOT entries.
- TSO/E will be designed to XTIOTs, uncaptured unit control blocks (UCBs), and DSABs above 16 MB for data sets allocated by programs.
- RACF will be designed to support XTIOTs, uncaptured UCBs, and DSABs above 16 MB for data sets allocated by programs.
- DADSM and CVAF changes are planned to support XTIOTs, uncaptured UCBs, and DSABs above the 16 MB line. This is intended to help you take advantage of those functions to allow more concurrently open data sets and provide for virtual storage constraint relief.
- 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. Additionally, the Volume Recovery utility will be designed to improve performance in certain situations when recovering object data stored on optical and tape media. Improvements are expected to be most noticeable when recovering a backup volume containing objects with primary copies in a large number of different collections on a large number of different volumes.
- 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.
- In z/OS V1.7, support was introduced in DFSMSdfp for large format sequential data sets (DSNTYPE=LARGE). In z/OS V1.8, Language Environment added support for these data sets using noseek (QSAM). Support for seek (BSAM) was limited to data sets no larger than 64K tracks on any volume when opened for read. In z/OS V1.12, seek (BSAM) support is planned to be extended to data sets

up to the maximum size when using record I/O. Binary and text I/O with seek continue to be supported for data sets up to 64K tracks in size on any volume when opened for read.

- 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. This function is planned to be automatically enabled.
- 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 to help speed recovery processing. Also, this is designed to allow you to restore Fast Recovery copy pools from tape using DFSMShsm.
- DFSMShsm Space Management performance improvements are planned. A new option will be designed to allow you to specify that Primary Space Management, Interval Migration, and Command Volume Migration be done in parallel.
- 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. A new MODIFY CATALOG,CONTENTION command is planned to allow you to specify a different interval than the 10-minute default or to disable CAS contention detection. This new function is intended to warn about tasks that take excessive time to complete, or never complete, from affecting Catalog performance.
- Language Environment will be designed to improve performance for string manipulation intensive applications, such as certain applications written in the Perl language.
- InfoPrint Server for z/OS is planned to enhance Extended Mode processing to support more SYSOUT data with similar attributes, the maximum number of active jobs allowed by the job entry subsystem (JES2 or JES3), and Line Printer Daemon (LPD) support for file sizes up to 4 GB. Support for large file sizes is available on z/OS V1.9 and higher with the PTF for APAR OA28795. Also, InfoPrint Server will be designed to prioritize spooling and printing for existing jobs higher than receiving new work. These changes are intended to help relieve constraints and reduce spool occupancy for InfoPrint Server jobs.
- 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 two new services, IXCMGOX and IXCMGIX, are planned to be the 64-bit counterparts of the existing IXCMGO and IXCMGI services. These new services are intended to make it easier for exploiters to achieve virtual storage constraint relief by removing the need to copy message buffers and associated storage structures from 64-bit addressable virtual storage to 31-bit storage and back.
- 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. This is intended to help you take better advantage of large-capacity tape cartridges.

- 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 ECSA 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, which reduces the overall amount of ECSA (and Telnet private) storage allocated to support Telnet sessions.
- 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

The platform's classic strengths of availability, security, reliability, scalability, and management have made the mainframe the de facto standard for data serving and OLTP. It is logical to extend z/OS to Business Intelligence and Data Warehousing solutions as well, where large amounts of reports can be generated in a timely manner using source data -- all with a simplified reconciliation and restatement process. But it is also logical to deploy new or extend existing applications that leverage data on z/OS.

What sets z/OS apart is the ability to operate both new and existing applications within the same system, and in close proximity to your corporate data residing on z/OS. New applications based on Java, WebSphere Application Server, Perl, PHP, XML, C/C++, Unicode, HTML, HTTP, SOAP, z/OS UNIX System Services, and other Web services can operate side by side and integrate with classic applications based on CICS, IMS, DB2, Enterprise COBOL, Enterprise PL/I, REXX, System REXX, JCL, TSO/E, ISPF, Assembler, and Metal C. These applications can be colocated with relational (DB2) and non-relational (IMS) databases as well as record-oriented data. With such proximity to the data, applications on z/OS have a reduced need for expensive communications and networking infrastructure and can offer fewer opportunities for security breaches due to tight integration with traditional z/OS security, audit and resource access.

Businesses with applications on z/OS understand the value of those applications, as well as understanding that replacing these systems with standard packages or other custom-built alternatives is not needed and may, in fact, be unnecessarily risky. Modernizing z/OS applications can lower costs and drive business agility with significantly enhanced levels of usability and integration.

- z/OS V1.12 is planned 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.

Details on the application integration improvements planned for z/OS V1.12 include:

- 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.
- Language Environment is planned to provide Euro currency support for Slovakia in the C/C++ Run-time Library. Both Euro and pre-Euro support will be provided and the default locale for Slovakia will be changed to use the Euro symbol.
- Calendar times, represented by time_t, will overflow in January 2038. In z/OS V1.12, the Language Environment C/C++ Run-time Library is planned to include new services including time64_t, that will support constructed calendar times up to and including 23:59:59 UTC on December 31, 9999.
- In z/OS V1.12, Program Management Binder is planned to complement the existing Binder C/C++ API DLL functions (iewbndd.so, iewbndd.x) with an XPLINK

version (iewbnndx.so, iewbnndx.x). This is designed to offer XPLINK applications improved performance by eliminating expensive XPLINK to non-XPLINK transitions when the binder functions are called. Also, a C/C++ header is planned to be provided to map the IEWBMP structure (`__iew_modmap.h`). For C and C++ users, this will simplify the task of processing the module map, which the binder creates in programs when the MODMAP option is used. A number of smaller Binder enhancements are also planned:

- Sample programs planned to illustrate the use of both standard and Fastdata Binder APIs in High-Level Assembler and C
- Character translations in AMBLIST LISTLOAD output for load modules
- Improved AMBLIST header information for z/OS UNIX files
- Support for long names for AMBLIST LISTOBJ for object modules in z/OS UNIX files
- 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 planned 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).
- 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 `tsocmd` 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 ISGENQ 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.

- JES2 and JES3 are planned to provide function you can use to specify, using the SYSOUT application programming interface (SAPI), that a program receive ENF 58 notifications when SYSOUT data sets have been deleted. This new function is designed to help applications to monitor the progress of print data sets through the system.
- The System Data Mover (SDM) component is planned to provide a REXX interface for many of the functions of the SDM programming interface (ANTRQST). This new function is designed to provide interfaces to FlashCopy®, Global Mirror (XRC), and Metro Mirror (PPRC) SDM services.
- The CIM Server is planned to be upgraded to a newer version of the OpenPegasus CIM Server. Also, the CIM Servers Schema repository is planned to be updated to CIM Schema version 2.22. This is intended to keep the z/OS CIM Server and schema current with the CIM standard from OpenGroup and DMTF, and to allow z/OS management applications manage z/OS in an enterprise environment. It is planned to include CIM providers for the Host Discovered Resources (HDR) and Host Bus Adapter (HBA) profiles from the SMI-S standard.
- In z/OS V1.9 the C/C++ Run-time Library iconv() family of functions began to use Unicode Services to perform most character conversions. In z/OS V1.12, the ucmmap source or genxlt source for character conversions is planned to be removed from the C/C++ Run-time Library. You can create customized conversion tables using Unicode Services to replace these functions.
- The WLM service for requesting LPAR-related data (REQLPDAT) is planned to be enhanced to include character-based data about the machine model, a Model-Permanent-Capacity Identifier, a Model-Temporary-Capacity Identifier, the Model-Capacity Rating, the Model-Permanent-Capacity Rating, and the Model-Temporary-Capacity Rating. This new data is intended to be used for reporting.

Security

Security is often a moving target. New security-related capabilities are often followed by ever-more sophisticated and creative attempts to circumvent them. Yet as vital as security is, sometimes it may be difficult to get funding for the latest security features as it is difficult to show a return on investment (ROI) on security solutions.

z/OS has a huge breadth of security capabilities built into the base of the operating system at no extra cost. Many z/OS security functions, such as data encryption, encryption key management, digital certificates, password synchronization, and centralized authentication and auditing, can be deployed as part of enterprise-wide security solutions and can help mitigate risk and reduce compliance costs, while accelerating time to and reducing cost of implementation.

- 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). 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.

z/OS V1.12 is planned to be updated with many cryptographic capabilities, such as support for new smart card formats, new cryptography standards and algorithms (such as DSA, DH, EC, AES GCM, BLOWFISH, RC4, Galois/Counter Mode encryption for AES (GCM), Elliptic Curve Cryptography (ECC), Elliptic Curve Diffie-Hellman key derivation (ECDH), Elliptic Curve Digital Signature Algorithm (ECDSA), and Hashed Message Authentication Mode (HMAC)), as well as z/OS Communications Server support for IKEv2 and Federal Information Processing Standard (FIPS) 140-2.

- Digital certificates are used in managing and working with private key/public key encryption and are often required as part of security and compliance guidelines. They can be used by applications to establish secure communication sessions or to configure virtual private network (VPN) sessions, and to authenticate users and objects. z/OS PKI Services is a complete digital certificate authority included in the

base of z/OS at no additional charge. Relatively few z/OS resources can be used to generate thousands, even hundreds of thousands of digital certificates. Reduce risk and reduce cost by generating and managing your own digital certificates from z/OS.

For z/OS V1.12 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. New standards, such as Certificate Management Protocol (CMP), mean devices can now request, revoke, suspend, and resume certificates from z/OS PKI Services directly and automatically. Certificates generated by z/OS PKI Services can also be customized for use with Microsoft Exchange and smart card readers.

- Authentication, auditing, and compliance are growing concerns. Many laws and standards have been recently refined, enacted, or created, governing the protection and access of data. z/OS has a long history of resource access and reporting capabilities built into the platform that can be useful for administering z/OS security, monitoring for threats, and auditing usage and policy compliance. z/OS V1.12 is planned 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.

Details on the security enhancements intended for z/OS V1.12:

- ICSF is planned to provide support for translation of external RSA tokens wrapped with key encrypting keys into one of three smart card formats. A new callable service, PKA Key Translate (CSNDPKT), is designed to translate an existing RSA private key in CCA external format into a specified smart card (SC) format in support of VISA, or the common ME or CRT format. To use this new function, you will need an IBM System z9® or System z10 server with the Crypto Express2 feature with a minimum driver and microcode level. This function is also available on z/OS V1.8 and higher with the z/OS V1.8, z/OS V1.9 or z/OS V1.10 with the Cryptographic Support for z/OS V1R8-V1R10 and z/OS.e V1R8 Web deliverable and PTF UA46713.
- An enhancement to Central Processor Assist to Cryptographic Function (CPACF) on IBM System z10 servers with the CEX3C feature is designed to help facilitate the continued privacy of cryptographic key material when used by the CPACF for high-performance data encryption. Leveraging the unique z/Architecture®, protected key CPACF is designed to help ensure that key material is not visible to applications or operating systems when used for encryption operations. Protected key CPACF is designed to provide significant throughput improvements for large volumes of data and low latency for small blocks of data. In z/OS V1.12, ICSF is planned to exploit the enhancements made to the CPACF in support of separate key wrapping keys for DES/TDES and AES. This is designed to provide the same functions available using the PCI card, but with the advantage of CPACF performance.
- There are a number of improvements planned for PKI Services.
 - In z/OS V1.12, PKI Services is planned to allow you to create and sign certificates with ECC keys, in addition to RSA and DSA keys.
 - RACF and PKI Services will be designed to support longer distinguished names in digital certificates. This is intended to support your use of certificates with very long distinguished names.
 - Certain events, such as restoring a prior level of the security database, or removing and reinstalling the Certificate Authority (CA) certificate, can cause the security manager to return serial numbers to be used for new certificates that have been used before. PKI Services will be designed to detect this and find the first unused serial number before issuing a new certificate, to avoid attempting to issue two certificates with duplicate serial numbers. Also, a new utility is planned to allow you to post existing certificates in LDAP, avoiding the need to post them manually. Additionally, another new utility will be designed to allow you to post updates to Certificate Revocation Lists (CRLs) immediately when you need to, rather than waiting for the interval you have specified. Last, PKI Services performs certain tasks, such as removing old or expired certificates and requests, and processing certificate expiration notification

- warning messages, once a day. These housekeeping tasks have historically consumed considerable processing time when you have a large number of certificates. A new PKI Services design is intended to markedly improve the performance and reduce the processing time of these tasks and additionally allow you to specify the time of day and days of the week this task will be run.
- PKI Services is planned to support passing the reason a certificate request was rejected from the administrator to the requester in the rejection e-mail. Also, PKI Services will be designed to support custom extensions to X.509 Version 3 certificates; for example, creating a Domain Controller certificate with an extension called Certificate Template Name, with an OID, and with BMP data "DomainController" for use with Microsoft Exchange or Smart Card Login. Last, PKI Services is planned to allow you to create a certificate with a Subject Alternate Name that contains multiple instances of each of the General Name forms support. For example, more than one IP address may be specified where only one was allowed before.
 - Certificate Management Protocol (CMP) is an Internet protocol used to manage X.509 digital certificates described by RFC 4210, which uses the Certificate Request Message Format (CRMF) described by RFC 4211. In z/OS V1.12, PKI Services is planned to provide support for parts of the CMP standard, allowing CMP clients to communicate with PKI Services to request, revoke, suspend, and resume certificates. This is intended to allow you to use CMP in a centralized certificate generation model.
 - Elliptic Curve Cryptography (ECC). See more information below.
 - RACDCERT enhancements include:
 - The RACF RACDCERT command is planned to be enhanced to support the creation of certificates with expiration dates in the far future to give greater flexibility on certificate validity period for customers.
 - RACF and PKI Services will be designed to support longer distinguished names in digital certificates. This is intended to support your use of certificates with very long distinguished names.
 - Elliptic Curve Cryptography (ECC). See more information below.
 - In 2009, the U.S. National Institute of Standards and Technology (NIST) published an IPv6 profile that requires support of certain cryptographic suites as defined in RFC 4869, Suite B Cryptography Suites for IPsec. One of the technologies referenced was Elliptic Curve Cryptography (ECC), which is regarded as providing stronger cryptography with smaller key sizes than RSA cryptography. This type of cryptography is expected to be attractive for use with small devices such as mobile devices and smart cards, that have limited computing power. In z/OS V1.12, PKI Services is planned to allow you to create and sign certificates with ECC keys in addition to RSA keys. In z/OS V1.12, System SSL is planned to provide support for ECC-related data structures, signing data, and verifying signed data using ECDSA (Elliptic Curve Digital Signature Algorithm). This is intended to allow exploiters of z/OS System SSL to import ECC style certificates and private keys into key database files or PKCS#11 tokens and use ECDSA certificates in signing and verifying operations. In z/OS V1.12, the RACF RACDCERT command is planned to allow you to create and sign certificates with ECC keys, in addition to RSA and DSA keys.
 - A discrete general resource profile with generic characters (*,%,&) in its name, defined in a class enabled for generics (GENCMD or GENERIC), is often called a "ghost" profile. Such profiles are not referenced by RACF for authorization checking. However, when defined, they can confuse and annoy RACF administrators and system programmers. In z/OS V1.12, RACF is planned to provide a new NOGENERIC keyword for the RDELETE command to enable you to delete these profiles. Also a GENERIC=N option is planned for R_admin DELETE. cms.
 - The Command Prefix Facility (CPF), which you can use to route commands from one system to another within a sysplex, is planned to support security checking similar to that provided for the ROUTE operator command. Defining a new MVS™.CPF.ROUTE.CHECK profile in the RACF OPERCMDS class will specify that the system use the MVS.ROUTE.CMD profile in the RACF OPERCMDS class to determine whether the operator is allowed to send a command to the specified system. This is intended to add the same level of checking to CPF that exists for the MVS ROUTE command.

- The Network Authentication Service for z/OS is planned to utilize RACF function to help improve the availability of applications that use Kerberos or GSSAPI services when deployed in a DVIPA environment. This new support is designed to allow you to remove the dependency on which image of the Sysplex a Kerberos or GSSAPI application request is routed to. This can help improve application availability by enabling transparent failover for improved application availability and improved workload balancing between images in a Sysplex.
- 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. Support is planned for automatic password revocation, password expiration, formatting checks, history, and a password change mechanism that can be enforced on an individual, group, or directory basis. This new function is intended to help you ensure that:
 - Users change their passwords periodically
 - New passwords meet your password requirements
 - Recently used passwords not be reused
 - Users can be locked out after a defined number of failed attempts

In addition, when a password policy control has been received, native or SDBM authentication will map RACF response codes to password policy response codes where possible, and the password policy response control will be returned.

- IBM Tivoli Directory Server for z/OS is planned to support continuous activity logging. This new function will be designed to close the current log file or generation data set and open a new one based on the time of day or the size of an activity log file you specify. The console command will be designed to allow initiation of an activity logfile switch. Also planned in this support is a new function that will allow specification that log entries be filtered by IP address.
- 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, to add or remove permissions based on:
 - Bind distinguished name (DN)
 - Alternate DNs
 - Pseudo DNs
 - Groups a bind or alternate DN belongs to
 - IP address of the client connection
 - Time of day that directory entry was accessed
 - Day of week that directory entry was accessed
 - The bind mechanism used
 - Whether bind encryption was used

This function is designed to provide additional flexibility in access controls for LDAP connections.

- IBM Tivoli Directory Server for z/OS is planned to provide Salted SHA-1 encryption support. Intended to make dictionary attacks against SHA-1 encrypted data much more difficult, stored Salted SHA-1 password values in LDAP will include a random 20-byte string so that encrypting the same password more than once will usually result in differing encrypted values. This is intended to make it much more difficult to determine the encrypted password value. This support is designed to be functionally equivalent to that currently provided by the IBM Tivoli Directory Server and can allow easier migration of LDAP server workloads to z/OS.
- IBM Tivoli Directory Server for z/OS is planned to provide support for the syntaxes and matching rules currently supported by IBM Tivoli Directory Server. This support will be designed to allow migration and replication of schema and directory entries using these syntaxes and matching rules from IBM Tivoli Directory Server on other platforms.
- 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. Partner security credentials can be retrieved if both endpoints of a TCP connection reside in the same z/OS image, z/OS sysplex, or z/OS subplex, and the endpoints are within the same security domain. In such a topology, partner programs can use trusted connections to authenticate each other as an alternative to using an SSL/TLS connection with digital certificates for client and server authentication.
- Internet Key Exchange version 2 (IKEv2) is the latest version of the Internet Key Exchange (IKE) protocol specified by RFC 4306. IKE is used by peer nodes to perform mutual authentication and to establish and maintain security associations (SAs). In z/OS V1.12 the Communications Server IKE daemon (IKED) is planned to be enhanced to support IKEv2, in addition to its existing IKEv1 support. The z/OS Communications Server support for IKEv2 is planned to include:
 - IPv4 and IPv6 support
 - A new identity type called KeyId
 - Authentication using pre-shared keys or digital certificates; certificates may use RSA or elliptic curve (ECDSA) keys
 - Re-keying and re-authentication of IKE SAs and child SAs
 - Hash and URL specification of certificates and certificate bundles
 - A new certbundle command which can create certificate bundles as specified by RFC 4306
- z/OS Communications Server is planned to introduce these enhancements to the network security services daemon (NSSD) IPsec Certificate Services:
 - IKEv2 support: X.509 certificate-based signature creation and validation for IKEv2
 - Elliptic Curve Digital Signature Algorithm (ECDSA) support: X.509 certificates that contain ECDSA keys may be utilized for IKEv2 digital signature creation and verification
 - X.509 certificate trust chain support: The entire X.509 trust chain will be taken into consideration during IKEv1 or IKEv2 digital signature creation and verification
 - Certificate Revocation List (CRL) support: CRLs may be retrieved via HTTP and consulted during IKEv1 or IKEv2 digital signature verification
 - Hash and URL support: Certificates and certificate bundles specified using the Hash and URL format specified in RFC 4306 may be utilized during IKEv2 digital signature creation and verification

The z/OS Internet Key Exchange daemon (IKED) is planned to be enhanced to use these new NSSD functions when a stack is configured as a network security client.

- z/OS Communications Server is planned to introduce these enhancements to IPsec and IKE support for cryptographic currency:
 - Support for the Advanced Encryption Standard (AES) algorithm in Cipher Block Chaining (CBC) using 256-bit keys, an addition to the previously existing 128-bit key support. You can use the longer key length for more-sensitive data.
 - Support for the Advanced Encryption Standard (AES) algorithm in Galois Counter Mode (GCM) and in Galois Message Authentication Code (GMAC) mode. AES in GCM is intended to provide both confidentiality and data origin authentication. AES-GCM is a very efficient algorithm for high-speed packet networks. AES in GMAC mode is intended to provide data origin authentication but does not provide confidentiality. AES-GMAC, like AES-GCM, is also a very efficient algorithm for high-speed packet networks. z/OS V1.12 Communications Server is planned to support both 128-bit and 256-bit key lengths for these algorithms.
 - Support for the use of Hashed Message Authentication Mode (HMAC) in conjunction with the SHA-256, SHA-384, and SHA-512 algorithms. These algorithms are intended to be used as the basis for data origin authentication and integrity verification. The new algorithms, HMAC-SHA-256-128, HMAC-SHA-384-192, and HMAC-SHA-512-256, are designed to help ensure that data is authentic and has not been modified in transit. Versions of these algorithms

- that are not truncated are available as Pseudo-Random Functions (PRFs). These algorithms are called PRF-HMAC-SHA-256, PRF-HMAC-SHA-384, and PRF-HMAC-SHA-512
- Support for an authentication algorithm, AES128-XCBC-96, that can help ensure data is authentic and not modified in transit.
 - Support for Elliptic Curve Digital Signature Algorithm (ECDSA) authentication.
 - Support for Elliptic Curve Diffie-Hellman (ECDH) key agreement
 - z/OS Communications Server IPsec and IKE support is planned to leverage z/OS cryptographic modules that are designed to address the Federal Information Processing Standard (FIPS) 140-2 security requirements for cryptographic modules. FIPS 140 defines a set of security requirements for cryptographic modules to obtain higher degrees of assurance regarding the integrity of those modules. FIPS 140-2 provides four increasing, qualitative levels of security intended to cover a wide range of potential applications and environments. z/OS V1.12 Communications Server support is planned to be configurable such that it will only utilize underlying security modules (System SSL and ICSF's PKCS #11 capabilities) that are operating in FIPS 140 mode. System SSL and ICSF's PKCS #11 capabilities are designed to address the requirements for FIPS 140-2 level 1.
 - RFC 4301 compliance for IPsec filter rules is planned to become mandatory. RFC 4301 "Security Architecture for the Internet Protocol" specifies the base architecture for IPsec-compliant systems, including restrictions on the routing of fragmented packets. Compliance enforcement may require minor changes to IP filters for IP traffic that is routed through z/OS. The Configuration Assistant will be designed to identify any non-compliant IP filters and policy agent will not install an IPsec policy that contains any non-compliant IP filters.
 - In prior releases, System SSL supported X.509 certificates with RSA key sizes up to 2048 bits for use in PKCS#11 tokens. In V1.12, System SSL gskkyman is planned to be enhanced to support the creation and management of X.509 certificates and keys within a PKCS#11 token that have RSA key sizes up to 4096 bits, DSA keys, and Diffie-Hellman keys. These X.509 certificates and keys are planned to be usable through the System SSL APIs.

Availability

Resilience that helps reduce risk from outages

There is more to "availability" than just the server being up -- the application and the data must be available as well. For the System z platform this means hardware, I/O connectivity, operating system, subsystem, database, and application availability, too. The System z hardware is designed to reduce planned and unplanned outages through the use of self-healing capabilities, redundant componentry, dynamic sparing, and the ability for concurrent upgrades and microcode changes. Data availability and integrity are upheld with capabilities such as address space isolation, storage protect keys, I/O channel redundancy, and I/O error checking.

Beyond the single system is z/OS Parallel Sysplex clustering (see also the [Scalability and performance](#) section). Parallel Sysplex clustering is designed to provide your data sharing applications and data with not only continuous availability for both planned and unplanned outages, but also near-linear scalability and read/write access to shared data across all systems in the sysplex for data sharing applications.

z/OS also has error checking, fault tolerance, isolation, error recovery, and Parallel Sysplex capabilities that it continues to enhance every year. Unlike other operating systems, z/OS plans to advance in a new, innovative direction for availability. z/OS is planned to extend its proactive learning, monitoring, and analysis, to enable the system to analyze a component or a subsystem that it suspects is failing and provide warnings and guidance for operators and system programmers.

- z/OS extends its high-availability characteristics by going beyond failure detection to predicting problems before they occur. With Predictive Failure Analysis (PFA), introduced with z/OS V1.11, your z/OS system is designed to learn heuristically from its own environment to anticipate and report on system abnormalities. It

compares present and past behaviors and models system behavior in the future, and is intended to notify you when a system trend might cause a problem.

For z/OS V1.12 PFA is planned to monitor 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. PFA can take into account the normal swings of daily, weekly, or monthly spikes and can learn the idiosyncrasies of your system, thus avoiding false warnings given by static monitors.

- In z/OS V1.12, a new capability, 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. Run Time Diagnostics is anticipated to run quickly to return results fast enough to aid you in making decisions about alternative corrective actions and facilitate high levels of system and application availability.
- In z/OS V1.12, a new Timed Auto Reply Function is planned to 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.
- z/OS availability is beyond the server as well. Parallel Sysplex can provide a large single system image, dynamic load balancing, fault tolerance, and automatic restart capabilities, so a single cluster can be used for scalability and performance as well as for availability and disaster recovery. With z/OS V1.12, 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 autonomies whereby the z/OS system can help identify CF structures and network connections that are unresponsive or in a degraded state. In addition, 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.

Details on availability improvements planned for z/OS V1.12 include:

- Over time, VSAM key-sequenced data sets (KSDSs) for which records are added and deleted have often become fragmented and have a significant number of empty Control Areas (CAs) that consume DASD space, increase the size of the indexes, and reduce performance. Performance and DASD space utilization can usually be improved for such data sets by copying, deleting and reallocating, and reloading them. This requires scheduled outages for applications using these data sets. In z/OS V1.12, DFSMSdfp is planned to allow you to specify that VSAM dynamically reclaim unused control areas for KSDSs, including those used for record-level sharing (RLS), and reclaim the associated index records as needed. This new function is intended to help you preserve performance, minimize space utilization for KSDSs, and improve application availability, and to allow you to discontinue the use of jobs whose sole purpose is to reorganize KSDSs.
- In z/OS V1.12, a new component named z/OS Run Time Diagnostics is planned. This function is planned to help you reduce the time spent deciding what actions to take to resolve a problem. It can be used to identify potentially related symptoms and causes when it appears a significant system problem might affect the systems ability to process your workloads. Often, you must quickly analyze these problems to preserve application availability. Run Time Diagnostics is designed to run using the START operator command and return results quickly to help you decide alternative corrective actions and maintain high levels of system and application availability. Run Time Diagnostics is planned to identify critical messages, search for serialization contention, find address spaces consuming a high amount of processor time, and analyze for patterns common to looping address spaces.
- A new Timed Auto Reply Function is planned to provide an additional way for the system to respond automatically to write to operator with reply (WTOR) messages. This new function is designed to allow you to specify message IDs, timeout values, and default responses in an auto-reply policy, and to be able to change, activate, and deactivate auto-reply with operator commands. Also, when

enabled, it is designed to start very early in the IPL process, before conventional message-based automation is available, and continue unless deactivated. An IBM-supplied auto-reply policy in a new AUTOR00 parmlib member that you can replace or modify is also planned. This new function is expected to help provide a timely response to WTORs and help prevent delayed responses from causing system problems.

- XCF Status monitoring will be designed to incorporate information about system-critical XCF group members that identify themselves, and initiate termination actions, including partitioning a system from the sysplex, if a monitored member fails to respond when polled for status or indicates impairment. This function is intended to help reduce the incidence of sysplex-wide problems that can result from unresponsive critical components. GRS is planned to exploit these XCF critical member functions in both ring and star modes. Additionally, GRS will be designed to monitor key tasks and notify XCF if it detects that GRS is impaired.
- A new Predictive Failure Analysis check is planned to detect and automate the system's response to tasks that are writing SMF records at unusually high rates. Another new function, SMF record flooding automation, is designed to allow you to define a policy for responding to these situations in the SMFPRMxx member of parmlib, by specifying whether record flooding automation is to be active, whether operators are to be warned, and the actions to take for specific SMF record types if record flooding occurs. This is intended to limit the impact of such problems by allowing less-important data to be discarded while keeping the data from critical SMF records intact. Additionally, new function is planned for the SMF dump program (IFASMF DL) to provide additional information to help you develop a record flooding policy.
- Function is planned for Predictive Failure Analysis (PFA) to allow you to specify that PFA ignore data related to certain jobs or address spaces when you expect their behavior to be atypical. This can help you improve the overall accuracy of PFA checks for logrec, message, and SMF record arrival rates. There are two types of machine learning, supervised and unsupervised. In z/OS V1.12 PFA will support both supervised and unsupervised learning. To support supervised learning, function is planned for PFA to allow you to specify that PFA ignore data related to certain jobs or address spaces when you expect their behavior to be atypical. By providing supervision (insight), you can help improve the overall accuracy of PFA checks for logrec, message, and SMF record arrival rates.
- Four changes to improve the quality of PFA modeling are planned for z/OS V1.12. PFA will be designed to:
 - Capture data when exceptions are issued to help you identify problems
 - Use dynamic modeling intervals based on system stability
 - Discard the last hour's LOGREC data from before a shutdown
 - Monitor smaller increments of common storage assigned to system
- New functions are planned for recovery and termination processing (RTM). These include a new option on ESTAEX to specify that SPIE or ESPIE exits be superseded by ESTAEX, a new option on ESPIE to request percolation to RTM, and passing information about held locks to ESTAE-type recovery routines.
- z/OS UNIX System Services file system processing will be designed to provide better information when a DISPLAY GRS,ANALYZE command is issued by identifying itself as the holder of held latches to GRS. This will be intended to help you diagnose and take corrective actions for latch contention problems that involve file system processing.
- Information about DFSMSrmm active and queued tasks is planned to be available via the DFMSMrmm API and via TSO/E subcommand, enabling storage applications to monitor and act on the available information.
- In z/OS V1.12, RSM and dump processing will be designed to improve capture performance for large amounts of data during SVC dump processing. This is expected to markedly reduce dump capture time when a large amount of data must be paged in during SVC dump processing. Internal IBM laboratory tests have shown that the capture time for SVC dumps can be reduced by over 60% for large dumps with a substantial percentage of data on auxiliary storage. The amount of improvement is expected to vary depending on how much data must be paged in during dump processing, how much real storage is available to the system, and the system's workload.

- Standalone Dump will be designed to better prioritize data capture for address spaces, and to dump a number of system address spaces first irrespective of their ASID numbers. This is intended to attempt to capture the data most often needed to diagnose system problems more quickly in case there is not enough time to take a complete standalone dump. Also, Standalone Dump will be designed to allow you to specify additional address spaces to be added to the predefined list using a new ADDSUMM option.
- Dump processing will be designed to act on a new option for the CHNGDUMP and DUMP commands, and in parmlib member IEADMCxx. The new DEFERTND option will allow you to specify that task nondispatchability for address spaces being dumped as a result of a DUMP operator command be delayed until after global data capture is complete. This is intended to reduce the amount of time tasks and address spaces being dumped are set nondispatchable to capture volatile data to reduce the impact of command-initiated SVC dumps.
- The existing XCF/XES CF structure hang detect support is planned to be extended by providing a new CFSTRHANGTIME SFM Policy option that will allow you to specify how long CF structure connectors may have outstanding responses. When the time is exceeded, SFM will be designed to drive corrective actions to try to resolve the hang condition. This is intended to help you avoid sysplex-wide problems that can result from an affected CF structure that is waiting for timely responses from CF structure connectors.
- One focus area in z/OS V1.12 is the time it takes to shut down and restart the z/OS system itself and major subsystems such as DB2. Substantial reductions in shutdown and restart times for DB2 systems that use a large number of data sets are expected in addition to improvements in the time required for some phases of z/OS initialization processing. Planned improvements include:
 - Design changes for Allocation intended to improve performance for address spaces that allocate a large number of data sets in a short time. These changes are expected to markedly reduce the startup time for these address spaces, such as DB2 address spaces and batch jobs that process a large number of data sets per job step.
 - Changing subsystem initialization from serial to parallel for initialization routines that are listed in IEFSSNxx parmlib members after a new BEGINPARALLEL keyword, to allow you to reduce system startup time by allowing many of these routines to run in parallel.
 - An XCF design change to help reduce IPL time when very large sysplex couple data sets are in use.
- z/OS Communications Server plans to introduce sysplex distributor support for a hot-standby server through the use of a new distribution method, HotStandby. You configure a preferred server and one or more hot-standby servers. The preferred server with an active listener receives all new incoming connection requests, and the hot-standby servers act as backup servers should the designated preferred server become unavailable. The hot-standby servers can be ranked to control which hot-standby server becomes the active server. You can also control whether the sysplex distributor automatically switches back to using the preferred server if it again becomes available, and whether the distributor automatically switches servers if the active target is not healthy.
- z/OS Communications Server sysplex problem detection and recovery is planned to be enhanced to detect when the TCP/IP stack has abended five times in less than a minute. Existing sysplex recovery logic is applied when this problem is detected.

Optimization and management capabilities

With the ability to intelligently manage workloads, reprioritize work, dynamically reallocate system resources between applications quickly and efficiently, and help meet business priorities, z/OS and System z can handle unexpected workload spikes and help improve your system's efficiency and availability.

- The z/OS Workload Manager is a cornerstone to z/OS leadership in on demand computing. With workload management, you define performance goals and assign a business importance to each goal. You define the goals for work in business terms, and the system decides how much resource, such as CPU and storage, should be given to it to meet each goal. Workload Manager will constantly monitor

the system and adapt processing to meet the goals. The scope of the Workload Manager extends from helping the management of incoming TCP/IP and SNA traffic, to managing requests for I/O. z/OS middleware like DB2, CICS, IMS, WebSphere MQ, and other WebSphere products can take advantage of WLM to manage the priority and execution of transaction requests across the z/OS system. For z/OS V1.12, WLM is planned to be updated with enhancements to improve batch management.

- Batch processing windows can be shortened and optimized through several other applications. For example, in z/OS V1.11, job streams using the IEFBR14 program during the batch window can be run faster by enabling Allocation to delete data sets without first recalling them. z/OS V1.12 is planned to be updated so IDCAMS can avoid DFSMSshm recalls when deleting generation data groups (GDGs).
- Extending the scope of z/OS and System z management, the Capacity 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.

Details on the optimization improvements planned for z/OS V1.12 include:

- Initiator address spaces consume processor time on behalf of starting and ending job steps that in prior releases are not associated with a particular batch job. There can be considerable variation in the processor time consumed by an initiator for different jobs. To help you 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.
- The creation of new VSAM data sets with IMBED and REPLICATE attributes has been unsupported since z/OS V1.3. These attributes, originally introduced to improve performance on older DASD, typically act only to occupy additional space and slow performance on modern cached DASD. In z/OS V1.12, the system will be designed to remove these attributes automatically from VSAM data sets logically dumped using DFSMSdss and migrated using DFSMSshm when DFSMSdss is used as the data mover during restore and recall processing. An informational message is planned to confirm that newly restored data sets no longer retain these attributes.
- DFSMSdss and DFSMSshm are planned to exploit the Fast Reverse Restore feature of the IBM System Storage™ DS8000™ Series. This function will be designed to allow recovery to be performed from an active, original FlashCopy target volume to its original source volume without having to wait for the background copy to finish when the volume pair is in a full-volume FlashCopy relationship. DFSMSdss will be enhanced to create full-volume copies using a new keyword in order to support a Fast Reverse Restore function. DFSMSshm FlashCopy backup and recovery operations will be designed to create full-volume FlashCopy relationships when the devices support it. The Fast Reverse Restore function will support the recovery of volumes associated with copy pool backups including Space Efficient and Incremental FlashCopy targets. A new DFSMSshm SETSYS parameter is planned to allow you to specify whether extent or full-volume FlashCopy relationships are to be established between volume pairs when DFSMSshm invokes DFSMSdss to perform fast replication backup and recovery.
- It is planned that DFSMSrmm will help with reporting of data sets and logical volumes which are copy exported from a TS7700 virtualization engine.
- 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. The Type 99 SMF record is also planned to be extended to show when the number of initiators to be started was limited for this reason. These changes are intended to improve WLM batch management.
- Changes to the dispatching of discretionary work are planned. 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.

- Capacity Provisioning is planned to use the delay data for transaction service classes provided by RMF to help determine whether a provisioning action is required for servers on which CICS and IMS are running. Monitoring delay data for CICS and IMS transaction classes is intended to help improve capacity provisioning decisions for servers with LPARs running CICS and IMS.
- The Capacity Provisioning Manager will be designed to allow you to specify that it is to use rolling performance intervals to determine whether a provisioning action should be taken rather than fixed intervals. This is intended to help improve the responsiveness of capacity provisioning.
- In z/OS V1.12, z/OS Communications Server is planned to use new TCP/IP callable NMI requests to provide TCP/IP stack network interface information and network interface and global statistics. Network management applications can use the requested output to monitor interface status and TCP/IP stack activity. z/OS V1.12 Communications Server is planned to provide the following new requests:
 - GetGlobalStats - Provides TCP/IP stack global counters for IP, ICMP, TCP, and UDP processing
 - GetIifs - Provides TCP/IP network interface attributes and IP addresses
 - GetIfStats - Provides TCP/IP network interface counters
 - GetIfStatsExtended - Provides data link control (DLC) network interface counters
- z/OS Communications Server plans to provide enhancements to improve the management of the CSSMTP application by adding the following new SMF 119 record subtypes:
 - 048 - CSSMTP Configuration data records
 - 049 - CSSMTP Target server connection records
 - 050 - CSSMTP Mail records
 - 051 - CSSMTP Spool records
 - 052 - CSSMTP Statistics records

It is intended that applications that want to process the new SMF 119 subtypes can obtain them from a traditional MVS SMF exit routine or in real time from the z/OS Communications Server Network Management Interface (NMI) for SMF, SYSTCPSM.

CSSMTP issues the SIOCSAPPLDATA ioctl to add application data (appldata) to the TCP connections used to connect to target mail servers. You can see the application data (appldata) displayed in the Netstat All/-A, AllConn/-a, and Conn/-c reports.

- z/OS Communications Server plans to introduce sysplex event notification through new SMF 119 event records (subtypes 32 - 37) that describe the following events:
 - DVIPA status change (subtype 32)
 - DVIPA removed (subtype 33)
 - DVIPA target added (subtype 34)
 - DVIPA target removed (subtype 35)
 - DVIPA target server started (subtype 36)
 - DVIPA target server ended (subtype 37)

The new SMF 119 event records are planned to be written to the MVS SMF data sets, and can also be obtained from the real-time TCP/IP network monitoring NMI (SYSTCPSM).

- DFSMS enhancements are planned for storage group management and volume selection performance. As volume sizes increase, one percent of a volume represents an increasingly large amount of storage. For example, on a 223 GB volume, 1% is over 2 GB of storage. In z/OS V1.12, the limit on the high threshold you can specify for space utilization for pool storage groups is planned to be increased from 99% to 100%. In most cases, IBM recommends a high

threshold value less than 100% for storage groups. This allows data sets to expand without an increased risk of encountering out-of-space abends. The 100% specification is intended to be used to make more storage capacity available for storage groups that hold static data. Also, SMS processing of volume lists is planned to be changed in a way intended to improve allocation performance for large volume lists.

- The Integrated Storage Management Facility (ISMF) includes a Data Collection application, DCOLLECT, which provides storage-related measurement data that can be used as input to the DFSMSrmm Report Generator to create customized reports or to feed other applications such as billing applications. In z/OS V1.12, DCOLLECT data class (DC) records are planned to be updated to include information about all data class attributes. Also, data set (D) records are planned to include job names, and storage group (SG) records are planned to include information about OAM Protect Retention and Protect Deletion settings.
- z/OS Communications Server planned improvements include:
 - The ability to learn indirect prefix routes from IPv6 Router Advertisement messages
 - The ability to associate preference values with all routes that are learned from IPv6 Router Advertisement messages

Use of these functions is expected to reduce the number of IPv6 static routes that must be defined and the ability to route around network failures when not using OMPROUTE to install routes learned via a dynamic routing protocol, such as OSPF.

- RMF is planned to include information in the CPU Activity Report about how many units of work, represented by work element blocks (WEBs), are running or waiting for a processor (CP, zIIP, or zAAP). Additionally, this function will be designed to provide this information in SMF Type 70 records. This new information is expected to be helpful for determining how much latent demand there is for processor time for multitasking address spaces.
- IDCAMS is planned to be enhanced to avoid DFSMSHsm recalls for any generation data sets that are migrated when deleting entire generation data groups (GDGs). Instead, IDCAMS will call DFSMSHsm to delete such data sets without recalling them. This is expected to reduce processing time, particularly when one or more generation data sets have been migrated to tape.

Networking

Where would we be without computer networks? Explosive growth in Web-based services, applications, appliances, and mobile devices is fueling a need for increased network performance, scalability, security, and management capabilities.

The z/OS Communications Server is there to meet the challenges with a wide array of networking technologies supported (including both TCP/IP and SNA). System and data security technologies, fault tolerance, autodetection and autorecovery capabilities all mean the z/OS Communications Server can provide reliable and trustworthy networking services. With intelligent configuration, dynamic optimization, self-tuning, and network routing, it adapts to different networking conditions and is capable of shifting workloads and traffic to meet quality of service and business needs. Designed for the largest enterprises in the world, z/OS provides network scalability, supporting both IPv4 and IPv6.

- It has been said "z/OS is not just a node on the network, it IS the network," and in some cases this is no exaggeration. What sets z/OS apart from other technologies is its sophisticated networking in a cluster (Parallel Sysplex). In a cluster, the z/OS Communications Server supports multiple applications, tools, databases, operating system images, partitions, servers, locations, and remote locations, with the ability to support multiple TCP/IP stacks, to provide different security and networking characteristics for these TCP/IP stacks, to automatically fail over a network, to dynamically manage networking traffic routing it by security, workload priority, or other quality of service characteristics, and to apply TCP/IP security capabilities centrally from an attractive, easy-to-use graphic user interface (the Configuration Assistant for the z/OS Communications Server). This is all integrated into and included with z/OS; the networking, its dynamic routing, and its policy-based security are not an optional add-on, but a vital part

of the system. z/OS V1.12 is planned to support new trusted TCP connections in a sysplex, providing a faster, simpler method for members in a sysplex to communicate. The next release is planned to have the ability to automatically add TCP/IP stacks to a sysplex at a later time, when you need it.

- Many data security breaches arise from data being plucked from an unsecured network connection. The Internet Protocol Security (IPSec) standard is just one of the industry standards useful for encrypting packets of a data stream. The z/OS Communications Server already allows for simplified and centralized configuration of IPSec security through its Configuration Assistant and allows most IPSec encryption and decryption to be eligible for the zIIP specialty engine. IPSec encryption on z/OS has the value of encrypting data right at the source. z/OS V1.12 is planned to support Internet Key Exchange version 2 (IKEv2), which is a more streamlined and efficient method of IPSec dynamic key exchange than the currently available IKEv1. Also for z/OS V1.12, z/OS Communications Server IPSec and IKE support is planned to leverage z/OS cryptographic modules that are designed to address the Federal Information Processing Standard (FIPS) 140-2 security requirements for cryptographic modules. Additionally, z/OS Communications Server IPSec and IKE are planned to support a variety of new cryptographic algorithms, enhanced X.509 digital certificate support, and more. Details on the latest on IPSec and IKEv2 can be found in the [Security](#) section.

Details on the networking improvements planned for z/OS V1.12 include:

- z/OS Communications Server V1.12 is planned to provide 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. In addition to the notification, statistics regarding the number of queries attempted and the number of queries that received no response are displayed for each currently unresponsive name server at five-minute intervals. This can alert you to a possible problem with your DNS name server configuration that may be adversely affecting applications on your z/OS system. The default value for the TCPIP.DATA RESOLVERTIMEOUT configuration statement, which controls the timeout value for UDP requests sent to a name server, is planned to be modified to be five seconds instead of 30 seconds.
- z/OS Communications Server plans to extend the VARY TCPIP,,DROP command to allow the dropping of all established TCP connections for servers that match the specified filter parameters. When issued, each server that is found to match the specified filter parameters will have all its established TCP connections dropped. You can filter by port, jobname, or server ASID. This function is expected to make it easier to move workload from one application instance to another application instance.
- z/OS Communications Server is planned to provide the option of keeping a TCP/IP stack isolated from the sysplex; you can use a new configuration parameter to prevent a stack from automatically joining the sysplex group at startup. You can have the stack join the sysplex group at a later time by issuing the VARY TCPIP,,SYSPLEX,JOINGROUP command.
- z/OS Communications Server is planned to enhance the performance of fast local sockets for TCP connections. This function is planned to be automatically enabled.
- z/OS Communications Server provides local and path MTU discovery to learn the correct MTU size for Enterprise Extender (EE) connections. The MTU size is used to modify the link size for EE connections. In z/OS V1.12, the link size is planned to be updated at the RTP pipe endpoints in addition to the EE endpoints when the MTU size changes.
- z/OS Communications Server packet trace filtering is planned to be enhanced to support:
 - Including the next hop IP address on the trace output. This can be obtained from the fully formatted packet trace using IPCS. The next hop IP address is also available to applications that consume the real-time packet trace through the real-time TCP/IP networking monitoring API.
 - Making packet trace filtering available to encapsulated packets that are used in VIPAROUTE traffic.

- z/OS Communications Server is planned to provide 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 is planned to reduce CPU utilization for TCP/IP Callable Network Management Interface (NMI), EZBNMIFR, GetConnectionDetail. All the filters that are specified for the request must contain the complete identification (4-tuple) of established TCP connections. The 4-tuple of a TCP connection consists of the local IP address, local port, remote IP address, and remote port for the connection.
- The z/OS Communications Server Netstat function is planned to provide support for verifying that message catalogs being used are at the correct level when they are opened. This function is intended to prevent Netstat from abending or not functioning correctly when the message catalog is out of synch with the Netstat program.
- z/OS Communications Server enhances TCP/IP data tracing (DATTRACE) to provide two new trace records:
 - A Start record with State field "API Data Flow Starts" that indicates the first data sent or received by the application for the associated TCP or UDP socket
 - An End record with State field "API Data Flow Ends" that indicates the socket has been closed
- z/OS Communications Server is designed to support RFC3484 by providing a configurable policy table for default address selection for IPv6. The source address selection algorithm and destination address selection algorithm are planned to be enhanced to support additional address selection rules in conjunction with the configured or default policy table. For example, you might choose to prefer IPv4 communication over IPv6 by providing a custom policy table for default address selection.
- z/OS Communications Server is also planned to support RFC5014 by providing IPv6 socket API for source address selection. Applications can indicate they prefer temporary IPv6 addresses over public IPv6 addresses or public IPv6 addresses over temporary IPv6 addresses.
- Additionally, z/OS Communications Server is planned to enhance the SRCIP configuration to allow an administrator to indicate that the TCP/IP stack should prefer public IPv6 addresses over temporary IPv6 addresses. This will allow you to override the preferences specified by an application using the IPv6 socket API for source address selection.
- z/OS Communications Server is planned to allow the system resolver to send requests to Domain Name System (DNS) name servers using IPv6 communication. This function is planned to allow you to use the existing NSINTERADDR and NAMESERVER resolver configuration statements in the TCPIP.DATA dataset to define the IPv6 address of the name server.
- z/OS Communications Server allows the coding of MULTIPATH in the TCP/IP profile that enables multipath support for IP packets. You might want this behavior for TCP connections but not for Enterprise Extender (EE) connections. In z/OS Communications Server V1.12, the multipath function is planned to be disabled by default for EE connections regardless of the value specified in the TCP/IP profile. You can use the VTAM start option MULTIPATH to control the multipath function for EE.
- z/OS Communications Server plans to enhance the digital certificate access server (DCAS) to allow modification of the debug level without restarting the application.

Statement of general direction

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 parmlib member to change the default Language Environment run-time options for the system.

IBM plans to pursue an evaluation to the Federal Information Processing Standard (FIPS) 140-2 using National Institute of Standards and Technology's (NIST) Cryptographic Module Validation Program (CMVP) for the PKCS #11 capabilities of the Integrated Cryptographic Service Facility (ICSF) component of the Cryptographic

Services element of z/OS. The scope of this evaluation will include algorithms provided by the CP Assist for Cryptographic Functions (CPACF) that are utilized by ICSF. This is intended to help satisfy the need for FIPS 140-2 validated cryptographic functions when using z/OS Communications Server capabilities such as the IPsec protocol.

Plans related to Extended Address Volume (EAV) larger volume sizes as described in 5694-A01, Preview: z/OS V1.10, announced in Software Announcement 208-042, dated February 26, 2008, will be communicated at a later date.

All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice.

Product number

5694-A01

Business Partner information

If you are a Direct Reseller - System Reseller acquiring products from IBM, you may link directly to Business Partner information for this announcement. A PartnerWorld® ID and password are required (use IBM ID).

<https://www.ibm.com/partnerworld/mem/sla.jsp?num=210-008>

Important Web sites

- z/OS Web site
<http://www.ibm.com/systems/z/os/zos/>
- General Q & A
<http://www.ibm.com/systems/z/faq/>
- Previously announced statements of direction
http://www.ibm.com/systems/z/os/zos/zos_sods.html
- z/OS Internet Library
<http://www.ibm.com/systems/z/os/zos/bkserv/>
- z/OS Basic Skills Information Center
<http://publib.boulder.ibm.com/infocenter/zoslnctr/v1r7/index.jsp>
- Descriptions of courses worldwide
<http://www.ibm.com/services/learning>
- z/OS downloads
<http://www.ibm.com/systems/z/os/zos/downloads/>
- CustomPac
<http://www.ibm.com/services/custompac>
- ShopzSeries
<http://www.ibm.com/software/shopzseries>
- z/OS Communications Server
<http://www.ibm.com/software/network/commserver/zos/>
- IBM Open Class® Library Transition Guide
<http://publibz.boulder.ibm.com/epubs/pdf/cbcit100.pdf>

Customized offerings

z/OS product deliverables are shipped only via CBPDO, ServerPac, and SystemPac.

Software delivery for z/OS and z/OS platform products on DVD is planned for September 10, 2010. This expands the delivery options available to include Internet,

DVD, 3590, and 3592 tape. Installation will require network connectivity between your z/OS system and a workstation having a DVD drive.

CBPDO, ServerPac, and SystemPac are offered for Internet delivery in countries where ShopzSeries product ordering is available. Internet delivery can reduce software delivery time and allows you to install software without the need to handle tapes. For more details on Internet delivery, refer to the ShopzSeries help information at

<http://www.software.ibm.com/ShopzSeries>

You choose the delivery method when you order the software. IBM recommends Internet delivery. In addition to Internet and DVD, the supported tape delivery options for CBPDO, ServerPac, and SystemPac include:

- 3590
- 3592

Note: Product delivery on all 3480 and 3490 tape media is planned to be discontinued October 26, 2010.

Most products can be ordered in ServerPac and SystemPac the month following their availability on CBPDO. z/OS can be ordered via all three offerings at general availability.

Production of software product orders will begin on the planned general availability date.

- CBPDO shipments will begin one week after general availability.
- ServerPac shipments will begin two weeks after general availability.
- SystemPac shipments will begin four weeks after general availability due to additional customization and data input verification.

Trademarks

IMS, z10, z9, DFSMS, RMF, REXX, DFSMSrmm, DFSMSdftp, DFSMSdss, DFSMShsm, MVS, System Storage and DS8000 are trademarks of IBM Corporation in the United States, other countries, or both.

IBM, z/OS, Predictive Failure Analysis, CICS, Parallel Sysplex, System z, DB2, Rational, Tivoli, RACF, Language Environment, WebSphere, System z10, VTAM, FlashCopy, System z9, z/Architecture, PartnerWorld and Open Class are registered trademarks of IBM Corporation in the United States, other countries, or both.

Microsoft and Windows are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

InfoPrint is a registered trademark of Ricoh Co., Ltd. in the United States, other countries, or both.

Ricoh is a registered trademark of Ricoh Co., Ltd. and its affiliated companies.

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

Terms of use

IBM products and services which are announced and available in your country can be ordered under the applicable standard agreements, terms, conditions, and prices in effect at the time. IBM reserves the right to modify or withdraw this

announcement at any time without notice. This announcement is provided for your information only. Additional terms of use are located at

<http://www.ibm.com/legal/us/en/>

For the most current information regarding IBM products, consult your IBM representative or reseller, or visit the IBM worldwide contacts page

<http://www.ibm.com/planetwide/us/>

Corrections

(Corrected on April 5, 2010)

The Description section was revised.

(Corrected on February 19, 2010)

A parmlib member was corrected, and information on Customized Offerings was revised.

(Corrected on February 12, 2010)

The list of key prerequisites was revised.