- z/OS 1.11 At A Glance -

•Failure avoidance: Predictive failure analysis is designed to help provide early warning about system trends that can cause system or application impacts, in many cases before they impact your business. The output from predictive failure analyses can be used for reporting and logging or to drive responsive automation intended to prevent serious problems.
• Simplified management: Updates to z/OS health checks, new migration health checks, new autonomics for defining and managing I/O, simplified application of network security, as well as other operational improvements are planned to help simplify systems management, improve administrator, operator, and developer productivity, and ultimately provide less opportunity for error.
• Responsive networking: New z/OS Communications Server has designs intended to improve networking in a Parallel Sysplex®, thereby enabling more efficient workload distribution in a sysplex and helping to improve the quality of the load balancing in multitiered z/OS server and application environments.
• Trusted system: z/OS V1.11 has numerous designs for implementing and integrating new security technologies. New key generation and archival functions are planned to enable you to generate and recover private keys from PKI Services. Support for new SSL functions is planned to be integrated in z/OS and updates to IBM® Tivoli® Directory Server for z/OS are intended to help simplify the migration of security-related LDAP-based applications to z/OS, integrate them with RACF®, and ultimately enable unified enterprise-wide identity and access management.
• Accountability: Superior measurement and data collection and reporting capabilities are updated and can be used for comprehensive risk management, auditing, and compliance plans. For example, a new identity propagation function can allow z/OS subsystems to associate distributed identities to RACF for improved cross-platform interoperability and auditing capabilities.
• Storage scalability: An additional data set type is supported in the extended addressing space (EAS) on an extended address volume (EAV). This allows better exploitation of the capacity of an EAV, which can be as large as 223 GB. This helps to relieve storage constraints as well as simplify storage management by providing the ability to manage fewer, large volumes as opposed to many small volumes.
• Improved resource optimization and economics: z/OS is planned to be updated to enable z/OS CIM (Common Information Model) server processing to run on the System z® Integrated Information processor (zIIP). This means CIM client applications that use the CIM server on z/OS for basic information interchange (such as RMF, WLM, DFSMS, and BCP) or for cross-platform system management such as the System z Capacity Provisioning Manager and the z/OS Management Facility problem determination capability (see Statement of general direction) can benefit.
• z/OS release-to-release performance improvement: For the past several releases z/OS has provided a release-to-release performance improvement by systematically reducing or eliminating constraints and inefficiencies within the base operating system. z/OS V1.11 has designs to continue release-to-release improvements. Imagine migrating to z/OS Version 1 Release 11, getting all the new function, and not needing more system resource -- z/OS is probably the ***only*** operating system that can deliver this.
•New z/OS Communications Server function is planned to simplify the deployment of network policies. For example, now applying network security for IPSec, IDS, and AT-TLS is made easier with step-by-step tasks complete with setting up RACF security to defining other policyr related configurations.
• New health checks are designed to help check the health of active systems. Use IBMprovided checks to implement the latest IBM hints, tips, and best practices, or you can write your own checks to address application-specific needs. The output from these checks is useful for proactively resolving problems before they occur and could be used as additional documentation for your own audit and compliance needs.
• Migration health checks are designed to reduce the effort required to migrate to new levels of z/OS from prior levels.
• Improvements to Message Flood Automation are designed to make it easier to implement, helping you improve availability by preventing runaway message traffic from causing system and application impacts.
• New RMF Spreadsheet Reporter provides a fast way to display postprocessor data (used for z/OS reporting, logging, and auditing) using a spreadsheet application.
• Additional ease-of-use enhancements are available for Allocation, z/OS UNIX® System Services, IBM Configuration Assistant for z/OS Communications Server (a downloadable tool), OAM, DFSMS, NFS, DFSMSrmmtm, DFSMShsmtm, and Group Capacity Limits reporting.

Ease of use and platform simplification enhancements

• z/OS Communications Server plans include:
– The IBM Configuration Assistant for z/OS Communications Server is planned to be enhanced to:
> Provide a simplified dialog for definition of AT-TLS and IPSec policies.
> Provide expanded support for setting up z/OS Communications Server policy infrastructure.
> Provide new AT-TLS policy options in support for new AT-TLS functions, such as support for TLS V1.1.
> In z/OS V1.11, the Advanced Communications Facility Trace Analysis Program (ACF/TAP) is planned to be made a part of z/OS Communications Server element. ACF/TAP provides functions to format trace information, including VTAM® buffer traces and VTAM internal traces. ACF/TAP will also continue to be included in ACF System Support Program (ACF/SSP). This change in ACF/TAP packaging is intended to help you reduce cost and ordering complexity if you do not use the Advanced Communications Facility Network Control Program (ACF/NCP).

• In z/OS V1.11, z/OS planned to simplify management
-- Improving the z/OS UNIX efficient and to provide
-- Providing a new ISPF-based use to browse and UNIX file system and
-- Enhancing monitoring of the Policy Agent start, such as the IKE daemon

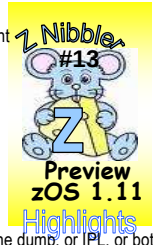

Communications Server designs are of the networking policy infrastructure by:
syslogd daemon to make it more automatic archive functions.
syslogd daemon browser you can search syslogd files in the z/OS archive data sets.
policy-related components by having monitor, and stop dependent functions, and NSS daemon.

IBM Health Checker for z/OS enhancements are planned:
– AutoIPL health check: AutoIPL support is designed to automatically take a standalone dump, or IPL, or both when the system is about to enter any number of non-restartable disabled wait states. Two new checks are designed to report on your AutoIPL policy specification:
> One check will be intended to make sure that you have an AutoIPL policy established as a "best practice" whenever running on hardware that can support AutoIPL, when not in a GDPS configuration.
> The second check is designed to validate the devices specified in the DIAGxx for SADMP and MVS™ AutoIPL policies. Devices will be checked to make sure they are available, are DASD, and are not defined as secondary devices in Metro Mirror (PPRC) pairs.
•The Message Flood Automation message processing, introduced with z/OS V1.9, will be designed to be part of internal message processing and eliminate its use of the IEAVMXIT and system command exits. These changes are planned to eliminate the customization and exit integration steps required to use Message Flood Automation and allow Message Flood Automation to take action against all messages, including those sent through specific Message Processing Facility (MPF) exits. Making it easier to use Message Flood Automation is intended to make z/OS systems more available by responding to a wider array of system and workload errors.
• In z/OS V1.11, RMF is planned to introduce new Overview reports and support in the RMF Spreadsheet Reporter. The Spreadsheet Reporter provides a fast way to display Postprocessor data graphically using a spreadsheet application such as Lotus® 1-2-3® or Microsoft® Excel. These new extensions will completely replace the Postprocessor PLOTS control statement, which provided a character-based graphic view of system performance.
• The IDCAMS DELETE command can be used to delete multiple entries by using a wildcard character as part of the entry name. In z/OS V1.11, IDCAMS will be designed to provide more selective criteria on the DELETE command. A new MASK keyword will be intended to allow you to specify data set name selection criteria using a mask entry-name, or key filter, with the new keyword.
• The Common Information Model (CIM) is a standard data model developed by a consortium of major hardware and software vendors (including IBM) called the Distributed Management Task Force (DMTF) as part of the Web Based Enterprise Management (WBEM) initiative. WBEM includes a set of standards and technologies that provide management solutions for a distributed network environment. Interoperability is a major focus of WBEM, and using WBEM technologies can help you develop a single set of management applications for a diverse set of resources. With support for the CIM server on systems running z/OS, your applications can gain access to z/OS resources using an extensible, industry-standard model. In z/OS V1.11, an upgrade to CIM is planned to support the latest version of the OpenPegasus CIM Server and DMTF CIM Schema. This includes support of the "CIM Operations over HTTP" specification version 1.3 from the DMTF.
• SDSF added support for JES3 in z/OS V1.10. In z/OS V1.11, SDSF is planned to remove the requirement for the HASPINDX data set in JES2 environments. This is designed to allow you to remove the HASPINDX data set once all sharing systems are running z/OS V1.11 JES2.
• TSO/E and z/OS Communications Server are planned to support a new option that enhances the LOGON RECONNECT function of TSO/E when using TN3270 connections. A new LOGONHERE parameter in parmlib member IKJTSOxx controls whether the RECONNECT option can be used to reestablish a TSO/E session unconditionally. Without this support, an attempt to reconnect is rejected if the system detects that the user ID is already in use, either on the same workstation or a different one. This new feature is intended to help preserve work that might be in progress under TSO/E sessions and avoid operator intervention.
• Resource Recovery Services (RRS) includes ISPF panels you can use to browse the RRS log stream data. In z/OS V1.11, a new selection on the Global Options panel will be designed to allow you to specify whether the timestamps displayed on other panels are in local time or GMT. This is intended to make the dialog more usable.
• In z/OS V1.11, the Program Management Binder is planned to add support for a new IEWPARMS DD statement, to be used to specify a data set (which can be an inline or SYSIN data set) that includes parameters to be passed to the binder. The binder will be designed to process these parameters in addition to any parameters specified a different way (for example, by using the PARM keyword of the JCL EXEC statement). This support will be intended to allow options to be added to those passed by an existing program without requiring program changes, and to allow more than 100 characters of parameter data to be specified. Also, it is expected to help simplify debugging for programmers developing applications that call the binder API.
• In z/OS V1.11, z/OS UNIX System Services shell commands **oedit and obrowse** are planned to be changed to use ISPF Edit and Browse. This makes editing and browsing z/OS UNIX files work consistently in both environments and adds support for ASCII data when these commands are used from the shell.
• z/OS V1.11 will be designed to refresh the volume table of contents (VTOC) information related to volume size following a dynamic volume expansion in an IBM System Storage DS8000. This new function will be intended to allow you to use new space on volumes immediately without the need for manually refreshing the VTOC.

Scalability & Performance by Increased Efficiency & Innovation

• An additional data set type is supported in the extended addressing space (EAS) on an extended address volume (EAV). This allows better exploitation of the capacity of an EAV, which can be as large as 223 GB. This helps to relieve storage constraints as well as simplify storage management by providing the ability to manage fewer, large volumes as opposed to many small volumes.
• A new prefetch capability is designed to provide performance improvements for XL C/C++ applications running on IBM System z10 servers.
• Memory management improvements are planned for AMODE 64 XL C/C++ Language Environment applications that exploit large (1 MB) pages, this in addition to the current exploitation by the 64-bit SDK for z/OS, Javatm Tech. Edition, V6.
• HiperDispatch design is planned to be changed to improve the performance for large-scale z/OS systems that include zIIP processors. These changes will be intended to improve system performance for LPARs with a large number of zIIPs.
• New z/OS Communications Server designs are intended to improve sysplex networking by allowing the Sysplex Distributor to use new and improved WLM interfaces for more efficient workload distribution in a sysplex and to improve the quality of the load balancing decisions made by Sysplex Distributor in a multitiered z/OS server environment.
• New z/OS Communications Server designs are intended to provide networking Performance improvements for name resolution, TCP-based bulk transfers over high-latency networks, and Enterprise Extender workloads when using IPSec.
• Virtual Storage Constraint Relief removes constraints within the base z/OS operation System and can allow more work to be processed on a single z/OS system. In z/OS V1.11, VSCR improvements are planned for STAR mode sysplex-wide GRS query requests, XML code page support, and TCP/IP sockets processing. In addition, new SMF fields can record additional information about an address space's use of virtual, real, and auxiliary storage above 2 GB and can help with capacity planning, performance management, and accounting.
•In z/OS V1.11, support is planned to be added to several components to help ease storage constraints by allowing extended format sequential data sets to be placed in the extended addressing space on extended address volumes (EAVs) in addition to VSAM data sets, for which support was introduced in z/OS V1.10. This is intended to help simplify storage management by allowing you to manage few larger volumes.
•In z/OS V1.11, Allocation design will be changed to significantly reduce the amount of storage required by the Eligible Device Table (EDT) for many I/O configurations. This new function will provide common area virtual storage constraint relief, both for the EDT itself and for use of Dynamic Activation for I/O configuration changes.
• In z/OS V1.11, DFSMS support is planned for solid state drives (SSD, also called flash memory) on DS8000. This new DS8000 feature introduces a powerful new technology to IBM enterprise storage by using solid state flash memory arrays in place of hard disk drives (HDD). Because it eliminates disk latency, SSD-based storage can provide a substantial improvement in storage response times compared to HDD for high-use data that is written infrequently and processed randomly. Also, support is planned to allow you to define new SMS policies for the allocation of new data sets on volumes backed by SSD technology and to gather usage information using SMF that is designed to help you manage data placement and take the best advantage of these new features.
•IBM plans additional System z10 hardware exploitation in the z/OS V1.11 XL C/C++ compiler through the addition of the PREFETCH option. This option is intended to enable heuristics to automatically generate System z10 prefetch instructions. This is expected to help reduce the effects of memory latency by beginning to fetch data Before it is known to be needed. Additional optimization and tuning improvements are also planned for the compiler. These changes are intended to improve the performance of generated code without the need for changes to the source code.
• New function is planned for AMODE 64 XL C/C++ Language Environment applications to enable them to exploit large (1 MB) pages, in addition to the current exploitation in the 64-bit SDK for z/OS, Java Technology Edition, V6. Language Environment will be designed to allow AMODE 64 applications to request large page-backed memory objects using the __moservices() interface. Appropriate use of large (1 MB) pages can help reduce memory management overhead and increase translation lookaside buffer (TLB) hit ratios for exploiting programs.
• z/OS V1.11 has designs for continued improvement to cache and memory management; this will continue to be true for the foreseeable future. The time it takes to retrieve data from memory, while progressively shorter in the absolute sense on newer server models, has become progressively longer when measured in processor cycle time increments. In other words, though memory access is much faster than it used to be, processors spend more cycles waiting for it than they ever did. HiperDispatch helps address the system's management of cache. In z/OS V1.10, HiperDispatch was directed toward improving cache management in multi-processing LPARs. For z/OS V1.11 HiperDispatch design is planned to be changed to improve the performance for large-scale z/OS systems that include zIIPs. These changes will be intended to improve system performance for LPARs with a Large number of zIIPs.
• In z/OS V1.11, when IEFBR14 is used to delete a data set that has been migrated by DFSMShsm, Allocation will be designed to pass the delete request to DFSMShsm so the data set can be deleted without first recalling it. This is expected to take significantly less time and consume fewer system resources than waiting for the data set to be recalled before being deleted, thereby helping speed batch processing when a migrated data set is deleted using IEFBR14.

Continued…

• A number of service unit count fields in the Performance section of SMF Type 30 records are incremented as an address space runs. To contain these counts for very long-running address spaces that can consume a larger number of service units than the original fields can represent, new larger fields are added to SMF Type 30 records. These new fields can be used in place of the original fields to gather information about the service units consumed by long Running address spaces.
• A new XES design is planned to provide algorithmic enhancements for Coupling Facility subchannel operation completion, recovery, and notification processing, and is intended to improve efficiency for processing a large number of Coupling Facility subchannels. Coupling Facility subchannel information provided by console DISPLAY commands is also plannedto be changed to accommodate a large number of Coupling Facility subchannels. This new function is intended to support larger CF and CF link configurations that require more concurrent I/O capability.
• In z/OS V1.11, GRS design is planned to change to provide storage constraint relief for GRS GQSCAN and ISGQUERY star-mode global ENQ processing. Work buffers used by parallel requests are planned to be moved to 64-bit pools in the GRS address space. This will be intended to increase the amount of data that can be processed by concurrent system-wide GQSCAN and ISGQUERY requests.
• In z/OS V1.11, support is added for data tables contained in load modules and program objects to be placed above 2 GB. A new ADDR64 parameter on the LOAD macro supports loading load modules and program objects containing only data tables above the bar. This function is intended to help relieve virtual storage constraints by allowing data tables to be loaded above the bar.
• z/OS XML System Services is planned to provide virtual storage constraint relief in z/OS V1.11 by moving XML code page dependent tables above the bar. Implementation changes intended to improve z/OS XML System Services performance for validating parsing are also planned.
• In z/OS V1.11, IBM Tivoli Directory Server for z/OS is planned to exploit Work Load Manager (WLM) functions. This enhancement will be designed to allow performance goals to be set for work running in IBM Tivoli Directory Server for z/OS and to take advantage of WLM's health notification service. This support will be intended to allow classification of particular IP addresses or distinguished names (DNs) into exception enclaves so their system resource consumption can be throttled. Use of the health notification service is expected to help prevent work from being directed to malfunctioning servers when their response time is very short.
• If you have a large number of RACF users who log on to various applications many times every day, the recording of logon statistics can cause considerable RACF database I/O activity. In z/OS V1.11, RACF will be designed to allow you to specify that logon statistics be recorded only once per day for certain applications. This will be intended to limit the I/O and serialization activity for the RACF database and help improve RACF performance.
• A new asynchronous (asyncio) form of the accept_and_recv socket API planned for z/OS V1.11 will be intended to allow z/OS applications and middleware to process short-lived client TCP connections more efficiently by combining the functionality of multiple socket calls into one. This support is expected to help improve performance for TCP server applications that exploit the new API.
• z/OS Communications Server design will be changed in z/OS V1.11 to help improve performance and reduce CPU cost.:
•The performance of VSAM System Managed Buffer (SMB) Direct Optimized Access Bias can be impacted by too few index buffers being allocated when opening a data set that is small at first but grows as it is processed. In z/OS V1.11, the system's use of the SMBVSP parameter is planned to be enhanced to make the amount specified apply to both data and index buffers.
• In z/OS V1.10, OAM implemented 2 GB Object Support to enable applications to store objects up to 2000 MB (2097152000 bytes) in size using DB2 on direct access storage. In z/OS V1.11, OAM will be designed to extend that support to the tape tier of the OAM storage hierarchy. This support will be intended to provide full support for objects up to 2000 MB in size on both DASD and tape and is expected to reduce the need to separate large binary strings into multiple objects and simplify the application interface by eliminating the need for applications to materialize entire objects before storing them.
• The TCP transport layer in z/OS Communications Server TCP/IP is planned to be updated to detect and transparently react to selected common sockets programming errors that can cause severe application performance problems. This new function will be intended to help improve application availability and performance.
• z/OS Communications Server plans to add a QDIO accelerator function. The accelerator is planned to provide accelerated forwarding at the Data Link Control (DLC) layer for IP packets being routed between any combination of QDIO and HiperSocketstm network interfaces. The QDIO accelerator function will be intended to improve performance when IP packets are routed through z/OS using any of the supported network interfaces.
• Device Allocation is planned to be changed in z/OS V1.11 to improve its tape load balancing algorithm. This new support will be designed to consider the number of devices in each tape library, and how many are online, when allocating tape devices within libraries. For example, a library with 256 online devices would be preferred over a library with 16 devices. You will be able to specify whether this new algorithm is used by using the new SYSTEM parameter for the ALLOCxx parmlib member. This new design is intended to provide better balancing of tape library requests across the system.

Application integration - Leverage your data and assets

• SYSREXX[tm] functionality to make writing system-level and automation applications much easier.
• METAL C functions intended to make it easier to integrate C code with assembler code, to help facilitate reuse of existing assembler assets.
• New, more efficient, XL C/C++ debugging capabilities, additional Unicode support for ease of porting applications to z/OS, and enhanced support for numerous IEEE, ISO/IEC, and GNU Compiler Collection standards. Also planned is a new prefetch capability, designed to provide further performance improvements for XL C/C++ applications.
• In z/OS V1.11, a new z/OS Communications Server mail client for sending mail from z/OS to a mail server is planned. The new mail client will be designed to support the existing user interface for sending mail from z/OS that was defined by the old SMTP/NJE mail gateway server, and many new mail enhancements such as secure connections, the ESMTP protocol, and IPv6 connections. The new z/OS Communications Server mail client will be designed to support sending mail only; it will not support receiving mail from the Internet. This support will be designed to provide secure connections to a mail server, provide simplified setup and operation, use less resources than the existing SMTP/NJE mail gateway server when sending mail from z/OS, and support existing applications that write mail messages via the JES spool.
• In z/OS V1.11, DFSMSdfp[tm] processing is planned to be changed to indicate end-of-file (EOF) during the allocation of data sets on DASD that are not SMS-managed and have either sequential or an undefined data set organization. This is intended to make this processing for both SMS-managed and non-SMS-managed data sets consistent, to make it unnecessary to open data sets solely to indicate EOF, and to help prevent programs from reading old data when a data set is read immediately after being allocated.
• PKI Web pages: In z/OS V1.11, the PKI Services Web pages that are provided to help request certificates and perform certificate administration tasks are planned to be enhanced. The Web pages are planned to be provided in Java in addition to REXX CGI. This change is intended to make it easier to integrate PKI Services Web pages with other Web-based applications you might have by allowing you to define them in XML files with XML schemas and then customize them using modifiable Java Server Pages (JSP) files.
• In z/OS V1.11, Language Environment is planned to be enhanced to allow AMODE 64 applications to use the Preinitalization routine (CELQPIPI) to specify that they will do their own storage management using GETSTORE and FREESTORE, and have their own message service routines (MSGRTN). This function is already available for AMODE 31 applications using CEEPIPI. Also, Language Environment will be designed to provide additional IEEE decimal floating-point (DFP) support for XL C/C++ applications.
• In z/OS V1.11, the dbx debugger will be designed to allow you to load a combined set of debug information (in an .mdbg file) to help improve dbx performance during startup and variable evaluation processing.
• New function is planned for METAL C, including the addition of the strtof(), strtod(), and strtold() functions, enhancement of the sprintf() and sscanf() family functions in the run-time library with additional format specifiers that will allow use of floating-point numbers, and the creation of a static library that users can bind with directly instead of using the system vector. Also, support is planned for the "extern" and "all" keywords to the Metal C PROLOG and EPILOG compiler options. This will be intended to allow you to specify whether the PROLOG/EPILOG applies to external functions only or to both internal (static) and external functions within the source file, and to make it easier to integrate C code with assembler code.
• In z/OS V1.11 XL C/C++, customer requirements planned to be addressed include Unicode Literal Support, which will be designed to add support for the char16_t and char32_t types through the use of typedefs in C and as native types in C++. This is designed to make it easier to port code exploiting these types to z/OS. Also included will be a new SKIPSRC option intended to allow programmers to specify that the compiler exclude source from the listing which has been #ifdef-ed out. This will be designed to improve programmer productivity by improving readability of the compiler listings.
• New support will be designed to provide GRS ENQ contention notification (ENF 51) enhancements that will allow listeners to be notified for specific resources, including those for conditional requests that are not satisfied because resources are not available. The support will allow ENQ monitors to report on unsatisfied conditional requests and for ENQ exploiters to be able to track and react to contention for their resources.
• z/OS Communications Server FTP support is planned for transferring files to and from z/OS UNIX System Services named pipes. z/OS applications that support reading from named pipes can process data transferred into a named pipe while FTP is still writing data into that pipe. Likewise, FTP can initiate transfer of data written to a named pipe while an application is still writing to that pipe. This support is intended to help reduce overall latency when transferring data to and from z/OS systems in combination with pre- or post-processing by other z/OS applications or subsystems.
• In z/OS V1.11, new support is planned for Language Environment intended to make it easier to write programs that can be ported between operating systems such as z/OS and z/VSE[tm]. These additions are planned: a CEEGLOB assembler macro, RMODE and AMODE support in the CEEENTRY assembler macro, a new FTCHINFO keyword on the CEEFETCH assembler macro that allows for the load of both Language Environment and non-Language Environment modules, and SERVICE keyword support in CEEPPA assembler macro.
• CIM Client for Java Version 2 support is planned to be provided for the new "CIM Operations over HTTP" specification version 1.3 from DMTF.
• z/OS V1.11 XL C/C++ is planned to support parts of the ISO/IEC TR 19768:2007, Technical Report on C++ Library Extensions. This support is intended to make it easier to port C++ code from other platforms to z/OS.
• The XL C/C++ compiler is also planned to be extended to support parts of the ISO/IEC JTC1/SC22/WG21 draft of a new C++ language standard, commonly referred to as C++0x.
• New z/OS Unicode services support is planned for the new conversion table between CCSID 1390/1399 and Unicode for JIS X0213. This support will be added for the 2002 version of the 1390/1399 CCSIDs, which contain JIS X0213 characters. The JIS X0213 standard promoted by the Japanese government is expected to conform with the family registration law for names of people that has been established by the Ministry of Justice and the Kanji glyph style standard specified by the Ministry of Education, Culture, Sports, Science and Technology.
• In z/OS V1.11, z/OS Communications Server is planned to add support for a new programming interface designed to allow applications to determine the MTU (maximum transmission unit) for an IPv4 TCP/IP network interface and cross-platform portability.
• z/OS Communications Server is planned to be changed to deprecate support for IPv6 Type 0 Routing Headers. This new design is intended to comply with IPv6 RFC5095, "Deprecation of Type 0 Routing Headers in IPv6."
• In z/OS V1.11, Language Environment will be designed to provide a consistent implementation of the getcontext(), makecontext(), setcontext(), and swapcontext() functions between AMODE 31 and AMODE 64 applications. This will be intended to allow applications to be common for AMODE 31 and AMODE 64 and make it easier to enable applications for 64-bit programming.

Security - Your secure data vault

• IBM Tivoli Directory Server for z/OS (IBM TDS, LDAP) for z/OS is planned to provide new function that is intended to simplify migration of LDAP server workloads to z/OS. Many security-related applications use LDAP, and migrating them to z/OS and integrating them with RACF can ultimately help unify disparate, disconnected processes for identity and access management.
• RACF administration is planned to be simplified with the availability of more RACF function and capability via IBM TDS.
• A new identity propagation function allows z/OS subsystems to associate distributed Identities to RACF for improved cross-platform interoperability and auditing capabilities
• New key generation and archival functions plan to enable users to generate and recover private keys from PKI Services.
• New SSL functions and standards are planned to be integrated in z/OS and can be applied to z/OS applications transparently by Application Transparent – Transport Layer Security.
• WebSphere Application Server for z/OS V7 and IBM Tivoli Key Lifecycle Manager (TKLM) both add SMF audit records which can help your business comply with corporate standards and government regulations. In addition, SMF records for RACF can be loaded into XML format for use and analysis on other platforms.
• As the number of applications, particularly Web-based applications, increases, your vulnerability to attack may go up. For z/OS V1.11, z/OS Communications Server has designs to support some of the latest security standards and RFCs, the latest SSL updates to Application Transparent-TLS, and many simplification enhancements to its sophisticated centralized policy-based networking processes.
• Drive-level encryption and advanced key management support. New DS8000 features, along with z/OS V1.11 and IBM Tivoli Key Lifecycle Manager (TKLM) functionality, allow the DS8000 to encrypt sensitive data when drives are physically removed. Drive-level encryption is designed to have no performance impact, to be transparent to the application and to the server, and to help minimize the costs associated with encrypting data at rest with a solution that is simple to deploy and simple to maintain..
• z/OS Identity propagation will be designed to provide a way for z/OS transactional Subsystems to associate users' distributed identities with RACF user IDs under z/OS security (RACF) control while maintaining the users' original identity information for audit purposes. This will be intended to improve cross-platform interoperability and provide value for both host centric and heterogeneous application environments. CICS Transaction Server for z/OS is planned to be the first exploiter of this new function.
• National Institute of Standards and Technology (NIST) is the U.S. federal technology agency that works with industry to develop and apply technology, measurements, and standards. One of the standards published by NIST is the Federal Information Processing Standard Security Requirements for Cryptographic Modules, referred to as FIPS 140-2. In z/OS V1.11, System SSL is planned to provide a mode of operation designed to meet the NIST FIPS 140-2 Level 1 criteria. This mode is intended to restrict a System SSL application to using FIPS-approved algorithms, key sizes, and SSL protocols.
• RACF, the Program Management Binder and Program Management loader are planned to add generalized functions that enable program objects to be digitally signed and verified. These functions are intended to allow you to sign and verify applications you have written, to enable software vendors to sign their applications, and for the system to verify signatures of these program objects. This includes Program Management Binder support for including digital signatures in program objects, and loader support to verify program object signatures before using them when you specify that signatures are to be verified. This is intended to provide another capability that can help you ensure that change control procedures are enforced.
• In z/OS V1.11, IBM Tivoli Directory Server for z/OS (LDAP) is planned to enhance the SDBM backend to allow administration of RACF general resource profiles in a manner similar to that currently used to manage RACF User, Group, and Connect profiles. The new RACF R_admin callable service design will be intended to allow general resource profile data to be extracted from the RACF database with good performance, and the SDBM backend of the IBM Tivoli Directory Server for z/OS server design extended to manage RACF Resource profiles. This new function is intended to allow the use of LDAP operations to create, modify, delete, and display discrete and generic RACF Resource profiles and access lists. In addition, it is planned to allow setting RACF options that pertain to classes through an SDBM modify command. Enhancements are also planned to be made to the change logging Extended operation interface to provide change log entries for these RACF profiles.
• IBM Tivoli Directory Server (LDAP) for z/OS will be designed to provide enhanced Replication function equivalent to that currently provided by the IBM Tivoli Directory Server for distributed platforms. This will include: replication of subtrees of the Directory Information Tree (DIT) to a specific server or, multitiered topology.

Security - your secure data vault...continued

• A new RACF option is planned to allow you to specify that the system automatically create OMVS segments with unique UIDs or GIDs when users who do not have an OMVS segment first attempt to use z/OS UNIX System Services. This new function is designed to save the newly-assigned UIDs and GIDs in the OMVS segments of the user and group profiles.
• In z/OS V1.11, RACF RACDCERT and PKI Services enhancements are planned. RACDCERT services are planned to be updated to extend the current multibyte character support. The RACDCERT support will be intended to allow installation, retrieval, and authentication functions on certificates to specify multibyte characters outside of the 1047 code page, allowing characters used in additional languages to be used in the Subject Distinguished Name. PKI Services will be designed to support the "SHA256 with RSA encryption" signature algorithm. This enhanced support will be intended to allow the SHA256 signature algorithm to be used for signing certificates, for certificate and authority revocation lists (CRLs/ARLs), and for OCSP responses.
• In z/OS V1.11, new designs for key generation and key archival/recovery capabilities are planned to be introduced for PKI Services. This new support will be intended to give certificate requesters the option to generate public/private key pairs themselves as they can today or to have PKI Services generate key pairs. This new design will also enable a requester to recover its private key which was generated by PKI Services.
• WebSphere Application Server V7 introduces an audit infrastructure designed to enable audit records to be created for security-related events. Similarly, the Tivoli Key Lifecycle Manager (TKLM), 5698-B35, will also provide SMF audit records on z/OS. The RACF SMF Unload Utility is planned to be updated to support both the SMF type 83 subtype 5 records planned to be written by WebSphere Application Server and the type 83 subtype 6 records to be written by TKLM.
• ICSF provides services for major credit card vendors, such as VISA and MasterCard, to generate and verify the verification values. The verification values are used in the process of validating the authenticity of the cards. In z/OS V1.11, the VISA Card Verification Value (CVV) Service Generate (CSNBCSG) and VISA Card Verification Value (CVV) Service Verify (CSNBCSV) callable services are planned to be extended to support 14-digit Diners Club primary account numbers (PANs). Support for additional lengths (15-digit, 17-digit, and 18-digit PANs) is planned to be included for future applications.
• ICSF is planned to introduce a Key Store Policy, providing a set of policy controls designed to allow you to specify further limits on application control of key material and provide a central point of control. The new set of policy controls planned includes Key Token Authorization, Key Label, and Duplicate Key Label Checking, Granular Key Label Access, and Symmetric Key Label Export Control. This new set of policy controls is planned to extend ICSF's use of the z/OS Security Manager and is intended to provide additional policy-based protection for key material stored in the CKDS and PKDS data sets.
• ICSF is planned to provide a new ICSF Query Algorithms service that returns a summary of available crypto algorithms, providing information applications and middleware can use to determine whether to use system services or provide their own cryptographic implementations.
• The Advanced Encryption Standard (AES) is a National Institute of Standards and Technology specification for the encryption of electronic data. It is expected to become the accepted means of encrypting digital information, including financial, telecommunications, and government data. AES is the symmetric algorithm of choice, instead of Data Encryption Standard (DES) or Triple-DES, for the encryption and decryption of data. ICSF will support the AES encryption algorithm with secure (encrypted) keys of 128, 192, and 256 bits..

Availability - Resilience that helps reduce risk from outages

• System z servers are 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. For example, the System z10 server provides additional microcode driver enhancements, and dynamic segment sparing for memory as well as a fixed Hardware System Area (HSA). Also, new features in the IBM System Storage DS8000 enable Metro Mirror (PPRC) support for FlashCopy® for improved storage availability.
• With every release, z/OS continues to refine its error checking, fault tolerance, isolation, error recovery, and diagnostic capabilities. z/OS V1.11 availability enhancements include designs for improved diagnostic data processing and advances in Mean Time To Recovery.
• z/OS V1.11 plans to extend predictive failure analysis. Beyond error checking, first failure data capture, and recovery routines, predictive failure analysis means your z/OS system heuristically learns from its own environment and is able to anticipate and report on potential system issues (however rare) before they are an impact to your business.
• Parallel Sysplex clustering is used for scalability, availability, software migrations, and disaster recovery and may even be considered the first cloud computing environment ever. For z/OS V1.11, Parallel Sysplex technology is enhanced with updates planned for improved recovery management and improved networking performance and workload management.
• Support for IBM Remote Pair FlashCopy, also known as Preserve Mirror. This new function, supported by new DS8000 features, will be designed to allow you to perform FlashCopy operations where Metro Mirror volumes are the destination while preserving the duplex state of the volumes. This is expected to help you improve storage availability and achieve more consistent recovery point objectives (RPO) by removing this reason for interrupting Metro Mirror duplexing. This function is planned to be made available on z/OS V1.8, z/OS V1.9, and z/OS V1.10 with the PTFs for APARs OA24809, OA24816, and PK64715. For more information about the new DS8000 features, refer to Hardware Announcement 109-119, dated February 10, 2009.
• In z/OS V1.11, design changes are planned for System Logger and the I/O Supervisor (IOS) to reduce IPL and restart time. The changes to System Logger will be intended to reduce concurrent log stream connection time involving staging data set allocations, and IOS design will be changed to streamline the channel programs used during initialization. These changes are intended to help improve IPL and Logger restart times.
• In z/OS V1.11, support is planned for a new ALTROOT parmlib statement you can use to specify an alternate sysplex root file system if the current sysplex root becomes unavailable (unowned). z/OS UNIX System Services will be designed to switch to the alternate sysplex root. This new support is intended to help improve availability by recovering from sysplex root file system failures.
• Dynamic LPA is planned to be enhanced to accept the full path and file names for modules passed using BYAddr and save the path name as part of the information associated with each entity so that CSVINFO and CSVQUERY can return those values for these modules. This support will be intended to help you determine the source for an LPA module that originally came from a z/OS UNIX file system.
• SETALLOC will be designed to allow you to change all the settings specified in ALLOCxx parmlib members with the exception of 2DGT_EXPDT, which specifies how 2-digit years in expiration dates are to be processed. The command is planned to use keywords similar to those used in ALLOCxx parmlib members.
• DISPLAY ALLOC will be designed to show the currently active Allocation settings as set from the ALLOCxx parmlib members used during IPL and any subsequent SETALLOC commands. Also, this command is planned to allow you to display allocation information that can help IBM Service diagnose problems.
• The AutoIPL support introduced in z/OS V1.10 is extended to multisystem-capable sysplex configurations with active Sysplex Failure Management (SFM) policies in z/OS V1.11, and available for z/OS V1.10 with the PTF for APARs OA26993 and OA26995. New LPAR firmware support is also required. This new support is designed to perform the requested AutoIPL actions when an SFM policy is active in the sysplex. For information about the required microcode levels, refer to the text of APARs OA26993 and OA26995 for z/OS V1.10, or *z/OS V1R11.0 Planning for Installation*, when available.
• The dump management roadmap is intended to keep pace with the growth in diagnostic data that results from larger systems and larger programs using ever-increasing amounts of memory. These improvements are meant to help you keep dumping time and dump transmission time under control.
• z/OS UNIX System Services will be designed to allow you to capture a System Call (Syscall) Trace for a specific application or set of applications. This is intended to allow you to gather more information about program processing history to facilitate application debugging. You will be able to specify that syscall tracing be turned on or off using a new SIGTRACE signal and bpxtrace command.
• Improved recovery times and reduced manual operator intervention for events that impact System z Sysplex primary time source, using new z/OS V1.11 and GDPS function. Improvements intended to allow a GDPS Controlling system (K-system) in a GDPS/PPRC configuration to continue processing, even when a failure has caused the server on which the K-system runs to no longer be synchronized to the sysplex timing source (Server Time Protocol (STP) or External Time Reference (ETR)) are planned for z/OS V1.11. Allowing a K-system to continue processing in a failure scenario can allow it to continue to fulfill its role as the driver of GDPS's recovery actions, such as coordinating Metro Mirror disk swap or disk recovery, removing failed systems from a sysplex, and restarting failed systems and workloads at a recovery site. A loss of synchronization with the timing source in a sysplex previously would have required operator intervention. For more information, see Software Announcement 309-002, dated February 24, 2009.
• Enhanced Contention Analysis (ECA) for GRS Latches will be designed to provide improved analysis for latch contention, including a new latch identity service that can be used by exploiters to provide information for display commands so that the holder of a latch is easier to identify. A new DISPLAY GRS,ANALYZE command, similar to the one available for enqueues, is planned for latches. Latch identification information and better real-time diagnostic data are planned to be displayed. In z/OS V1.11, System Logger and Resource Recovery Services (RRS) are planned to provide identification information about their latches. The new functions for analyzing latch waiters, blockers, and dependency contention analyses will be intended to make it easier to identify deadlocks and to find the root cause of latch contention.
• In z/OS V1.11, XCF design is planned to change to automatically adjust the failure detection interval (FDI) to use for systems in the sysplex when needed. The intent is for the system's effective FDI to be the longer of the two intervals resulting from the FDI you specify and a value based on the system's excessive spin parameters, making the system's processing of excessive disabled spin conditions, the sysplex's handling of missing system heartbeats, and the initiation of sysplex partitioning to remove unresponsive systems more consistent.

Serviceability Enhancements

• In z/OS V1.11, new function in the IPCS LEDATA verb exit is designed to provide a new high level trace mechanism for C FILE I/O that can help reduce the time it takes to understand and debug complicated I/O problems. Also, the IPCS function in combination with additional settings on the HEAPCHK option are intended to improve formatting of heap pools control blocks and extents to make it easier to find damaged cells.
• z/OS UNIX has been changed to provide a new option you can use to specify that all syscalls made during a kernel shutdown be failed for applications. Also, additional documentation and information will be provided for security system calls that end in error to help you diagnose security problems for z/OS UNIX applications.
• System Logger design will be changed to help protect the system from SQA shortages that can result from write requests for a log stream being issued faster than they can be processed. Thresholds are created for outstanding asynchronous IXGWRITE requests. When they are exceeded, requests will be failed with a new returncode to indicate to callers that too many outstanding requests for the specified log stream exist. An Event Notification Facility (ENF) signal will indicate that the backlog has been reduced enough for new requests to be accepted. This change is intended to help improve system availability.
• A new JCL keyword on the DD statement, SMSHONOR, is planned. It will be intended to be used to specify that Allocation should honor the device number or Esoteric device type specified by UNIT as it processes the tape library devices selected for the request by SMS. This change is intended to help improve service-ability by allowing you to specify a specific device within a tape library when gathering diagnostic information, such as when using GTF tracing, and improve device availability.
• Improvements are made to z/OS UNIX file system processing design. z/OS UNIX automount processing design is planned to discontinue automated UNMOUNT attempts when unmount failures occur. A new BPXF251I message is planned to be issued when unowned file systems are recovered and made active in shared file system configurations. The DISPLAY OMVS,WAITERS command is planned to provide date and time information, and the DISPLAY OMVS,PFS display to include physical file system status information; the address space name for colony file systems including zFS, NFS and TFS; and the user, date, time, system, and master configuration file for the automount command.
• New MODIFY DFStm commands will be designed to allow you to format the DFS trace table, restart the trace, display the size of and other information about the trace table, and change its size dynamically while the SMB server runs. These changes will be intended to help improve availability and serviceability for SMB server.
• Currently, some errors are returned to the z/OS UNIX System Services file system caller with reason codes that are the NFS "file-and-line" type, requiring that analysis be done by NFS development to identify the code point reporting the error. To make it easier to identify the root causes for these kinds of errors, NFS design will be changed in z/OS V1.11 to convert many of these notifications to published reason codes you can use to discover the reason for the error. Also, NFS Client Memory Management will be designed to provide more meaningful diagnostic information, including the module and offset of the calling function when an error occurs.
• In this release, VSAM serviceability improvements are planned. VSAM design for record management and internal tracing will be changed to add additional trace information, and corresponding formatting support is planned for IPCS. These changes are intended to improve first failure data capture for VSAM, and allow you to capture VSAM diagnostics data without having to re-IPL your system, simplifying VSAM problem diagnosis.
• z/OS Communications Server is planned to add improved internal monitoring and management of storage, which is used by sockets applications that for various reasons build up an excessive amount of data on their send or receive queues. In addition, z/OS Communications Server is planned to add support to allow the z/OS routing Daemon (OMPROUTE) to continue processing during temporary storage shortage conditions. This function will be intended to improve network and system availability during periods where network traffic volume peaks and storage is constrained.

Optimization and Management Capabilities

• A new z/OS component, the Base Control Program Internal Interface (BCPii) is designed to allow authorized programs to perform Hardware Management Console (HMC) functions and to communicate via an HMC network without the need for an IP network for connectivity.
• DFSMSrmm enhances its enterprise-wide tape management capabilities with designs for improved search capabilities, new health checks, and simplified monitoring and management functions.
• Support for enhanced WLM routing algorithms: The server-specific Workload Manager (WLM) recommendations which are used by the Sysplex Distributor to balance workload across systems when the SERVERWLM distribution method is chosen are planned to be enhanced to allow you to specify how zIIP and zAAP specialty processor capacity influences selection of eligible target systems. In addition, new configuration controls are planned to allow you to specify that WLM consider the importance of displaceable capacity when determining server specific recommendations. These enhancements are expected to improve the workload balancing functions provided by WLM and Sysplex Distributor.
• In z/OS V1.8, DFSMShsm introduced support to recover individual data sets from copy Pool backup versions. DFSMShsm provided support for preallocated, cataloged data sets residing on the same volumes they were on when backed up. Deleted or moved data sets had to be reallocated and cataloged to the same volumes on which they resided when backed up in order to recover them. In z/OS V1.11, DFSMShsm will be designed to capture catalog information for the data sets within a copy pool at the time of the backup, and use that information to recover the dataset.
• The existing DFSMShsm (H)BACKDS command allows you to create backup versions of specified data sets. In z/OS V1.11, new function is planned to allow you to specify a retention period on the data set backup command. The specified retention period will be intended to keep an individual backup copy for either a shorter or longer than normal period of time.

Optimization and Management capabilities ..continued

• In z/OS V1.11, DFSMShsm will be designed to allow data sets larger than 64 K tracks in size to be migrated and backed up to disk and to enable Migration Level 1 (ML1) overflow volumes to be selected for migration processing and backup processing.
• In z/OS V1.11, DFSMS volume selection for striped data sets will be changed to be more consistent with volume selection for unstriped data sets. This new design will be intended to prefer enabled volumes over quiesced volumes and normal storage groups over overflow storage groups, allow volumes that are above the high threshold after a stripe is allocated to be eligible for selection with lower preference, and honor the multitiered storage group attribute and additional volume preference attributes.
• In z/OS V1.11, the SMF log stream dump utility, IFASMFDL, is planned to allow you to specify a range of dates relative to the date on which the program is started. For example, you will be able to specify that the SMF records created yesterday, or those created for six days starting nine days ago, be processed. This design will allow date ranges to be specified by days, weeks, or months. Also, IFASMFDL is planned to support new options to manage the retention of the data in the SMF log stream and allow you to specify that data by archived (dumped and deleted) or deleted from the log stream. These new IFASMFDL functions are intended to make it easier to exploit log stream-based SMF data management.
• OSA-Express connection isolation will be designed to prevent OSA-Express from internally routing packets between TCP/IP stacks sharing the OSA-Express port. This support is intended to provide an improved ability to enforce IP filter processing in outboard IP routers for traffic that flows between System z LPARs that share OSA-Express ports.
• z/OS Communications Server has designs to provide updated network management interfaces to provide additional details to network management applications.
• Many installations use products or functions that employ sampling techniques to help them determine which programs are in use and provide a view of the components of application response time and each component's contribution to response time. z/OS Communications Server code is planned to be changed to make it easier for such products and functions to map load modules and entry points. This is intended to make it easier for you to gain a better view of how the network-related components of your workloads contribute to application response time.
• Extend number of reporting classes to 2048: With systems and environments becoming larger, more reporting classes are needed in WLM. In z/OS V1.11, the number of report classes is planned to be increased from 999 to 2,048. This increase is expected to allow more fine-grained reporting of your workloads.
• Make REQLPDAT callable from unauthorized users: In z/OS V1.11, the service used to request data about a logical partition, REQLPDAT, which can return LPAR and capacity information from the system, will be enhanced to support unauthorized applications. This is expected to make it easier to obtain information about the system.
• New function in z/OS Communications Server will be intended to improve the accuracy of workload balancing decisions for server applications that use the Fast Response Cache Accelerator (FRCA) function with persistent HTTP connections, such as Web serving applications. This new function is intended to improve workload management for workloads that use the Fast Response Cache Accelerator function.
• In z/OS V1.8 and later releases, you can define capacity groups for LPARs to help you manage the 4-hour rolling average utilization of your servers while allowing capacity unused by one LPAR to be used to satisfy workload demands on other LPARs on the same server. In z/OS V1.11, the RMF Postprocessor Group Capacity report will be designed to include information about the available capacity in each capacity group and the percentage of time each partition was actually capped during the reporting period. This is in addition to existing fields that report the capacity of each group and whether each partition was eligible to be capped during the same period. This new support is intended to make it easier to manage LPAR capacity groups.
• In z/OS V1.11 SMS is planned to support the allocation of critical data sets (such as DB2 partitions) on different volumes to help reduce I/O contention. This new function will be designed to expand the existing data set separation function, to allow you to specify that critical SMS-managed data sets be separated across extent pools and volumes that are not used by other data sets specified in the separation group.

Networking

• z/OS Communications Server designs for z/OS V1.11 include:
– IBM Health Checker for z/OS RFC4301 compliance
– IBM Health Checker for z/OS DNS server check
– New SMTP client for sending Internet mail
– AT-TLS enhancements
– Enhancements to Configuration Assistant - AT-TLS and IPSec improvements
– Enhancements to Configuration Assistant policy infrastructure simplification
– Enhancements to Configuration Assistant - new AT-TLS options
– Changes to make it easier for products and functions to map load modules and entry points
– The Advanced Communications Facility Trace Analysis Program (ACF/TAP) is made part of z/OS Communications Server element
– Changes to syslog daemon (syslogd) design to improve performance and add automated archival of log files
– New ISPF-based syslog daemon browser application and search facility
– Enhanced policy infrastructure management
– Improved responsiveness to storage shortage conditions
– Disabled use of moving DVIPA as source IP address
– Sysplex autonomics improvements for Fast Response Cache Accelerator (FRCA)
– TCP throughput improvements for high-latency networks
– Sysplex Distributor connection routing accelerator
– Enhanced resolver DNS cache
– Sysplex Distributor optimization for multitier z/OS workloads
– QDIO routing accelerator
– Improvements to TCP/IP pathlength
– Enhanced IPv6 stateless address autoconfiguration
– New API to obtain IPv4 network interface maximum transmission unit (MTU)
– Deprecated IPv6 type 0 route header
– Enhanced QDIO for WLM IO priority
– Enhanced network management interface - detailed Communications Storage Manager (CSM) usage
– New accept_and_recv socket API
– Enhanced SNA High Performance Routing (HPR)
– FTP access to z/OS UNIX named pipes
– Enhanced QDIO support for OSA interface isolation
– Sysplex Distributor support for enhanced WLM routing algorithms
– NSS private key and certificate services for XML appliances
– EE IPSec performance improvements
– IPSec enhancements
• In z/OS V1.11, a new z/OS Communications Server mail client for sending mail from z/OS to a mail server is planned. The new mail client will be designed to support the existing user interface for sending mail from z/OS that was defined by the old SMTP/NJE mail gateway server, and many new mail enhancements such as secure connections, the ESMTP protocol, and IPv6. The new z/OS Communications Server mail client will be designed to support sending mail only; it will not support receiving mail from the Internet.
•z/OS Communications Server Application Transparent SSL/TLS support (AT-TLS) is planned to be updated to support System SSL functions that have been added since z/OS V1.7.
• IBM Configuration Assistant for z/OS Communications Server is planned to be enhanced to provide a simplified dialog for definition of AT-TLS and IPSec Policies.
• IBM Configuration Assistant for z/OS Communications Server is planned to be enhanced to support new system SSL and TLS options when defining AT-TLS policies. In addition to creating and maintaining policy definition files, the Configuration Assistant is planned to lead administrators through other tasks needed to get z/OS Communications Server policy functions such as IPSec or AT-TLS active and running on the z/OS system. These tasks include setting up RACF security, configuring the Policy Agent, and configuring other daemons that might be involved in the function. This support will be intended to simplify deployment of the full networking policy environment and provide improved integration of the various z/OS Communications Server networking policy components, ranging from the Configuration Assistant to the daemons running on z/OS.
•Many installations use products or functions that employ sampling techniques to help them determine which programs are in use and provide a view of the components of application response time and each component's contribution to response time. z/OS Communications Server code is planned to be changed to make it easier for such products and functions to map load modules and entry points. This is intended to make it easier for you to gain a better view of how the network-related components of your workloads contribute to application response time.
• In z/OS V1.11, the Advanced Communications Facility Trace Analysis Program (ACF/TAP) is planned to be made a part of z/OS Communications Server element. ACF/TAP provides functions to format trace information, including VTAM buffer traces and VTAM internal traces. ACF/TAP will also continue to be included in ACF System Support Program (ACF/SSP). This change in ACF/TAP packaging is intended to help you reduce cost and ordering complexity if you do not use the Advanced Communications Facility Network Control Program (ACF/NCP).
• In z/OS V1.11, z/OS Communications Server syslog daemon (syslogd) is planned to be changed to make it a multithreaded application. This will be intended to speed processing of z/OS UNIX System Services log messages. Also, this design will allow the syslogd job name to match the JCL procedure name and provide a set of operator commands for starting, stopping, and controlling syslogd. Syslogd is also planned to be enhanced to archive z/OS UNIX files to sequential data sets automatically based on configurable options. These changes will be intended to improve processing performance and reliability for z/OS UNIX log messages and to simplify syslogd file space and archive management operations.
• In z/OS V1.11, z/OS Communications Server is planned to provide a new ISPF-based syslogd daemon browser application to support browsing and searching both currently active syslogd files (those currently being written to by the syslogd daemon) and syslogd archive data sets that were created using a new integrated syslogd archive function. This support will be intended to provide easier access to the information logged by syslogd.
• z/OS Communications Server is planned to add improved internal monitoring and management of storage used by sockets applications that for various reasons build up excessive amount of data on their send or receive queues. In addition, z/OS Communications Server is planned to add support to allow the z/OS routing daemon (OMPROUTE) to continue processing during temporary storage shortage conditions. This will be intended to improve network and system availability during periods where network traffic volume peaks and storage is constrained.
• In z/OS V1.11, z/OS Communications Server design is planned to be changed to enhance the networking Policy Agent to provide monitoring, automatic start, and restart for the Defense Manager daemon, Internet Key Exchange daemon, Network Security Server daemon, syslogd daemon, and Traffic Regulation Management daemon. This function will be similar to the AUTOLOG function in the TCP/IP stack, except that applications will not need to maintain a listening socket. It will be intended to provide simpler management and operations for a set of applications that are associated with the policy infrastructure, and to reduce the complexity of enabling and operating z/OS Communications Server policy infrastructure by providing a single point of control and monitoring for its most widely-used components.
• z/OS Communications Server is planned to add support for preventing a dynamic VIPA address (DVIPA) in moving status from being used as source IP address for outbound connections. This support will be intended to improve management of DVIPA address movement.
• New function in z/OS Communications Server will be intended to improve the accuracy of workload balancing decisions for server applications that use the Fast Response Cache Accelerator (FRCA) function with persistent HTTP connections, such as Web serving applications. This new function is intended to improve workload management for workloads that use the Fast Response Cache Accelerator function.
• z/OS Communications Server design will be changed in z/OS V1.11 to help improve performance for inbound streaming TCP connections over networks with large bandwidth and high latency by automatically tuning the ideal window size for such TCP connections. This support is intended to improve throughput for large file transfers over long distances.
• Performance of z/OS Communications Server Sysplex Distributor when forwarding inbound IP packets for distributed connections is planned to be enhanced via a new Sysplex Distributor connection routing accelerator function. The Sysplex Distributor connection routing accelerator is based on a new general QDIO routing accelerator function that also is part of z/OS V1.11. This is expected to reduce the CPU time and latency for Sysplex Distributor connection routing.
• z/OS Communications Server is planned to enhance the resolver function by introducing system-wide caching of domain name server (DNS) responses. The resolver cache will be enabled by default and shared across the entire z/OS system image. This new function will be expected to provide significant performance improvements for z/OS workloads that perform repetitive resolver queries, such as Web services workloads. These performance improvements are expected to be realized without the need to set up, monitor, and administer a name server for caching purposes on your system.
• z/OS Communications server Sysplex Distributor is planned to enhance its optimized local support in environments where multiple tiers of z/OS applications need load balancing. Prior to z/OS V1.11, you can configure Sysplex Distributor with OPTLOCAL to optimize connections when both connection endpoints reside on the same TCP/IP stack within the z/OS Sysplex. In this type of environment OPTLOCAL is deployed on the tier 2 Sysplex Distributor to optimize the connections between tier 1 and tier 2 server applications.
• z/OS Communications Server plans to add a QDIO accelerator function. The accelerator is planned to provide accelerated forwarding at the Data Link Control (DLC) layer for IP packets coming inbound over HiperSockets and being forwarded outbound over OSAExpress QDIO or HiperSockets and for packets coming inbound over OSA-Express QDIO and being forwarded outbound over OSA-Express QDIO or HiperSockets. The QDIO accelerator function is expected to improve performance when IP packets are routed through z/OS using any of the supported network interfaces.
• The TCP transport layer in z/OS Communications Server TCP/IP is planned to detect and transparently react to two common sockets programming errors that can cause severe performance problems. z/OS Communications Server will be designed to transparently guard against traffic stalls due to use of the Nagle algorithm and delayed TCP acknowledgments and stalls caused by TCP receive buffers that are too small. This new function will be intended to help improve application availability and performance.
• A new asynchronous (asyncio) form of the accept_and_receive socket API planned for z/OS V1.11 will be intended to allow z/OS applications and middleware to process shortlived client TCP connections more efficiently by combining the functionality of multiple socket calls into one. This support is expected to help improve performance for TCP servers.
•SNA High Performance Routing (HPR) performance is planned to be enhanced by implementing a new progressive mode Adaptive Rate-Based (ARB) pacing algorithm, which is expected to increase HPR performance in virtualized or CPU-constrained environments. In addition, unproductive path switches when an HPR endpoint is unresponsive are planned to be reduced or eliminated. The intent of this new function is to improve SNA Enterprise Extender (EE) performance when EE partners.
• z/OS Communications Server FTP is planned to implement a new design that will allow z/OS applications that support reading from named pipes to process data transferred into a named pipe while FTP is still writing data into that pipe.
• The server-specific WLM recommendations which are used by the Sysplex Distributor to balance workload across systems when the SERVERWLM distribution method is chosen are planned to be enhanced to allow you to specify how zIIP and zAAP specialty processor capacity influences selection of eligible target systems..
• z/OS Communications Server in z/OS V1.11 plans to enhance network security services (NSS) to further enable XML appliance security as a logical extension of z/OS security. Two new services are planned: a certificate service to provide certificate management operations, and a private key service to allow retrieval of private keys from RACF certificates, and to perform RSA signature creation and RSA decryption using ICSFprotected keys. This support will be intended to enhance z/OS centralized management of security for XML appliances.