


IBM FlashSystem 9200R Product Guide

Matt Smith



Storage



IBM FlashSystem 9200R Rack Solution

The FlashSystem 9200 combines the performance of flash and end-to-end Non-Volatile Memory Express (NVMe) with the reliability and innovation of IBM® FlashCore technology, the ultra-low latency of Storage Class Memory (SCM), the rich features of IBM Spectrum® Virtualize and AI predictive storage management, and proactive support by Storage Insights. All of these features are included in a powerful 2U enterprise-class, blazing fast storage all-flash array.

For more information about the FlashSystem 9200, see the IBM FlashSystem 9200 Product Guide, [REDP-5586](#), and [IBM Spectrum Virtualize FAQ](#). It is suggested that you familiarize yourself with the FlashSystem 9200 before reading this IBM Redpaper publication.

Scale up and scale out

IBM Spectrum Virtualize technology allows the FlashSystem 9200 to scale up (add capacity) and scale out (linearly scale performance, connectivity, and capacity) to create a much larger system for customers who want to grow large or consolidate their workloads.

For the FlashSystem 9200, scale up is achieved by adding FlashSystem 9000 expansion enclosures to the controller enclosures. Two models are offered: an AFF, which is a 2U 24-drive system, and an A9F, which is a 5U 92-drive system. Both enclosure models support SSD drives up to 15.36 TB.

Scale out is achieved through clustering. This process takes two, three, or four model AG8 controller enclosures and clusters them to operate as a single system with a single point of control. Each controller enclosure is effectively operating in an active-active node pair (known as an *I/O group*) where write cache is mirrored across both nodes for redundancy. These node pairs come together to operate as a single system with the ability to move workloads between I/O groups to manage performance and balance workloads.

Each model AG8 controller enclosure supports 24 NVMe drives (up to 38.4 TB) with more capacity possible by way of hardware compression. Each controller enclosure features dedicated CPUs, ports, and cache to serve that storage. Therefore, if you cluster controller enclosures, you can scale the performance, connectivity, and capacity linearly. You can expect 2, 3, and 4 times the capability of a single model AG8.

FlashSystem 9200R overview

The FlashSystem 9200R provides a way of ordering a clustered FlashSystem 9200 with a dedicated Fibre Channel network and optional expansion enclosures. The system is delivered ready-assembled in a rack.

The FlashSystem 9200R can be ordered with 2, 3 or 4 FlashSystem 9200 model AG8 controller enclosures. These different configurations are referred to as IBM FlashSystem® 9202R, IBM FlashSystem 9203R, and IBM FlashSystem 9204R, with the last number denoting the number of AG8 controller enclosures in the rack.

Each rack features a pair of 32 Gb Fibre Channel switches with all ports enabled. This configuration allows a dedicated, redundant Fibre Channel network to be created, which allows inter-cluster traffic to flow between the canisters in the AG8 controller enclosures. By having this dedicated network, cluster traffic is prevented from interfering with host or copy service traffic and provides reliable, consistent communications.

For customers wanting more storage, a single 2U 24 drive model AFF or 5U 92 drive model A9F expansion enclosure can be added to each model AG8. To allow the expansion to be attached to the model AG8, a 12 Gb SAS card is required in one of the I/O card slots, reducing the number of slots available for host and clustering connectivity.

Compatibility

Although the FlashSystem 9200R is a separately orderable product, it remains equivalent to a clustered FlashSystem 9200.

All of the functionality that is available through Spectrum Virtualize (such as data reduction, copy services, and DRAID) is available in the FlashSystem 9200R. The Spectrum Virtualize GUI shows the system as a clustered FlashSystem 9200.

The FlashSystem 9200R features the same licensing scheme as the FlashSystem 9200. Therefore, all-inclusive licensing is available for all functions except encryption (which is a country limited feature code) and external virtualization.

The configuration rules and best practices that are associated with the FlashSystem 9200 are also the same. Other considerations are involved when clustered FlashSystem products are used, but these considerations are the same when the FlashSystem 9200 is supplied in FlashSystem 9200R rack form. Any limitations (such as number of volumes and the number of RC partnerships) also match the limitations of the FlashSystem 9200.

Interoperability is the same as the FlashSystem 9200 and is referenced as a FlashSystem 9200 in the [IBM System Storage Interoperability Center](#).

Capabilities

When purchasing any FlashSystem product, it is important to understand the workload requirements and the suitability of the product to meet those requirements. For more information, see the [FlashSystem Family Overview FAQ](#).

Table 1 shows a high-level comparison of a single FlashSystem 9200 and clustered (or racked) FlashSystem 9200s. The capacity shows only the NVMe capacity in the controller enclosure, with more capacity possible by using expansion enclosures.

Table 1 High-level comparison of single and clustered systems

FlashSystem	Controller enclosures	Physical capacity	Effective capacity 2:1	Effective capacity 5:1	IOPS (4K read miss)	Bandwidth
9200	1	921 TB	1.6 PB	4 PB	1.2 M	45 Gbps
9202R	2	1.8 PB	3.2 PB	8 PB	2.4 M	90 Gbps
9203R	3	2.7 PB	4.8 PB	12 PB	3.6 M	135 Gbps
9204R	4	3.6 PB	7.3 PB	18.4 PB	4.8 M	180 Gbps

Rack configurations

For the initial offering, the FlashSystem 9200R features a limited set of configurations, as listed in Table 2.

Table 2 FlashSystem 9200R configurations

FlashSystem	9848-AG8	9848-AFF ^a	9848A9F ^a	FC switches ^b
9202R	2	0 - 2	0 - 2	2
9203R	3	0 - 3	0 - 2	2
9204R	4	0 - 4	0 - 2	2

a. For each 9848-AG8, only one expansion enclosure can be ordered. AFFs and A9Fs cannot be mixed in the same system in the initial order.

b. Exactly 2 32 GB switches are supplied with each rack. These switches are used only for internal fabric connections. A choice is available between an identical pair of 8960-F24 switches or a pair of 8977-T32 switches.

Although this set of configurations falls short of the total number of configurations possible when clustering the FlashSystem 9200 with expansion enclosures, they represent the most popular configurations.

An alternative configuration can be created by using two methods. The first is to order one of the FlashSystem 9200R products for delivery that is fully assembled, and then order more expansion enclosures for installing into the rack at the customer site (the expansion enclosure is racked and cabled by an IBM engineer). The second approach is to order the rack, switches, enclosures, and controllers separately and then use an IBM Lab Service or Business Partner engagement to assemble the rack entirely onsite.

Component options

The FlashSystem 9200R makes ordering all of the components and feature codes that are needed much easier. Some options within the ordering system are restricted.

The model AG8 is the most configurable component and where possible retains all of the available options on the FlashSystem 9200 model AG8. The 32 Gb Fibre Channel connectivity is mandated, as is SAS connectivity if expansion enclosures are specified, but other I/O card options can be specified. Cache is defaulted to 1.5 TB maximum, but can be reduced to 768 GB per controller enclosure.

Models AG8, AFF, and A9F can all have different amounts of storage added to them. However, it is recommended that you balance the storage across the components to make most of the available performance.

The switch configuration is largely fixed, with all ports populated and enabled for future expansion.

The rack configuration is similarly fixed, with the PDU and power cords to each of the components automatically selected by the ordering system. The user is responsible for selecting a rack power cord from the rack to their data center power supply.

Models AG8, AFF, and A9F can have features added post initial order by using MES. Where available, more cache, different I/O cards, and more storage can be added.

Ordering

Contact your IBM representative (or Business Partner) to order a FlashSystem 9200R.

The FlashSystem 9200R is a larger unit than a single 2U FlashSystem 9200 and requires a more up-front planning. For more information, see [IBM Knowledge Center](#).

A FlashSystem 9200R specific Pre-sales Technical and Delivery Assessment (TDA) should also be reviewed with your IBM representative (or Business Partner) to understand the environmental considerations before ordering.

To order the FlashSystem 9200R, your IBM representative (or Business Partner) uses the e-config tool. Within the Storage product family, under FlashSystem, three options are available: FlashSystem 9202R, 9203R, and 9204R (see Figure 1 on page 7).

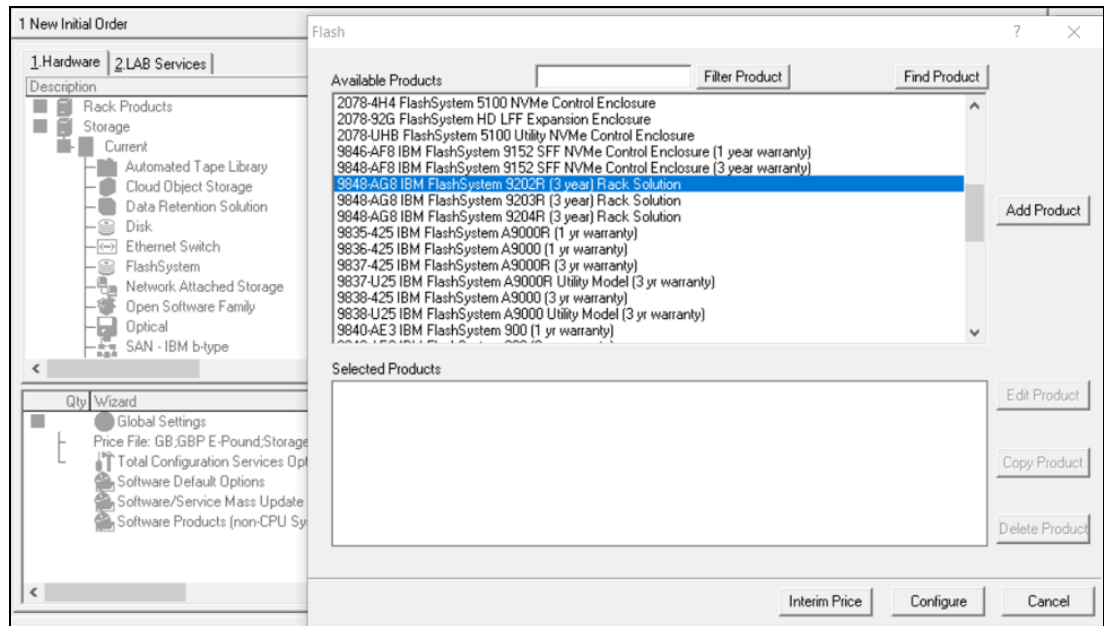


Figure 1 FlashSystem 9200R options

Selecting one of these products results in the minimum required components being selected. Expansion enclosures can then also be added as an option. Figure 2 shows a FlashSystem 9202R with expansion enclosures available to add.

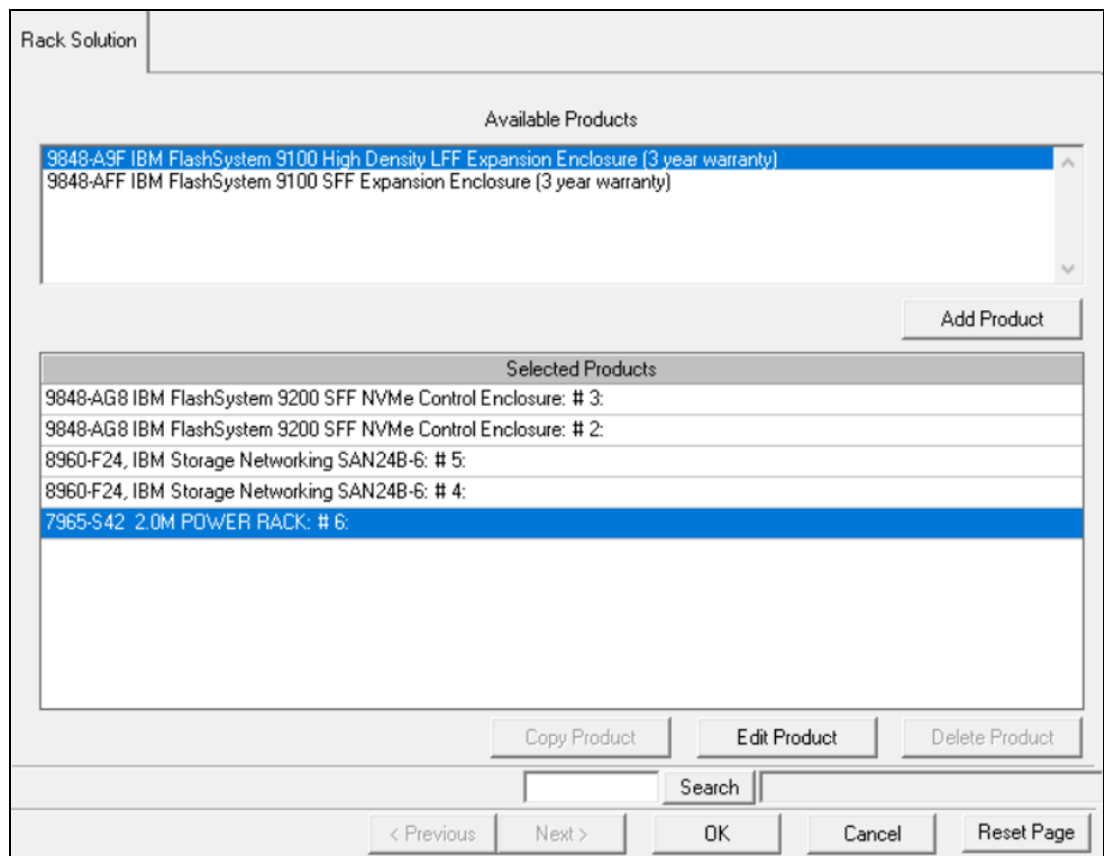


Figure 2 FlashSystem 9200R with expansion enclosures available to add

Editing one of the selected products is the same as ordering a non-rack FlashSystem 9200, with some of the options removed or defaulted. If you remove or add components that are required, an error results.

After the components are selected and configured, a diagram of the rack is displayed in e-config with the components displayed in their intended position. Figure 3 shows the FlashSystem 9202R with two model AFF expansion enclosures.

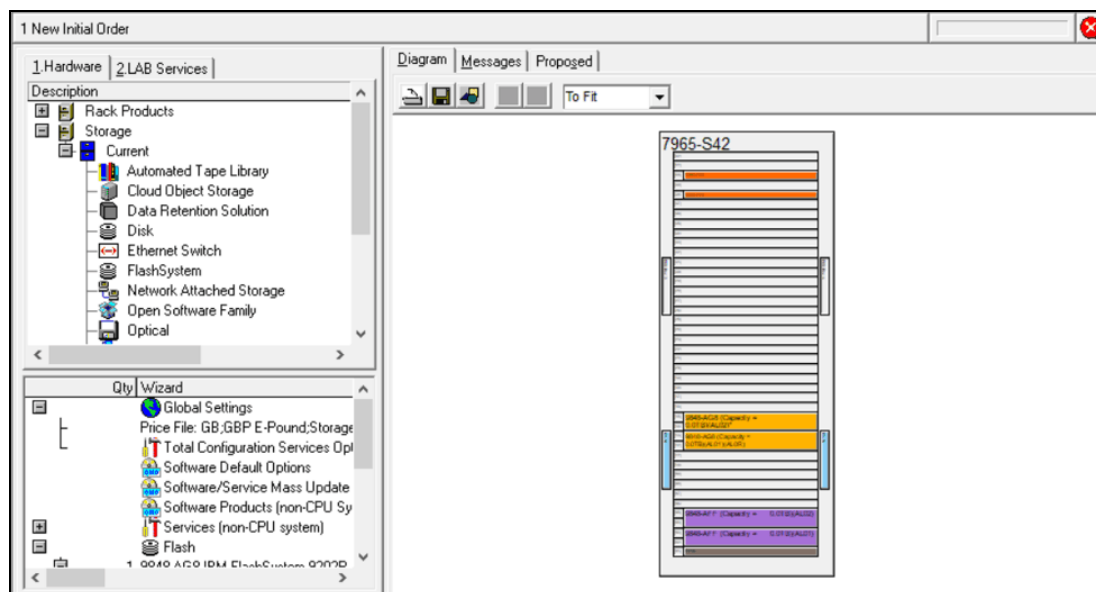


Figure 3 FlashSystem 9202R with two model AFF expansion enclosures

After the configuration is finalized, you can progress through e-config to complete the order.

The order is assembled in a rack by IBM and then shipped to the customer site. There, an IBM representative completes the configuration and makes the rack ready for use.

Rack installation

To keep the configuration simple and predictable, the components in the rack are always installed in the same location. Figure 4 on page 9 shows the position of each of these components.

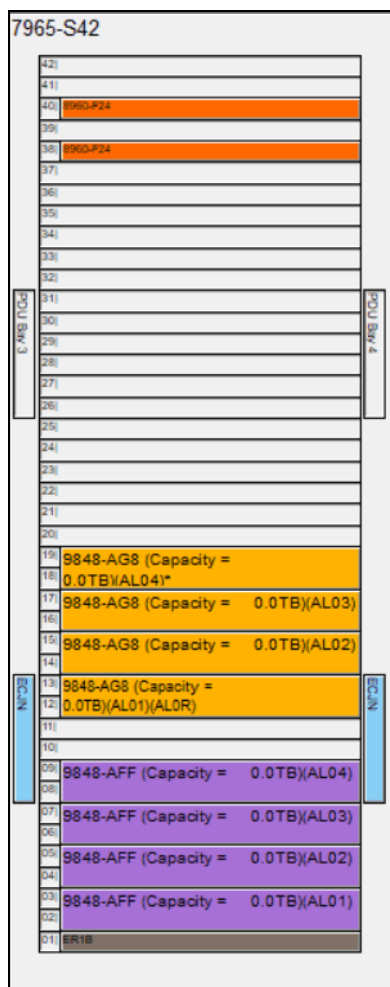


Figure 4 Component position

The two Fibre Channel switches (shown in orange in Figure 4) are at the top of the rack with a 1U gap between them.

Two or four PDUs are included in the side of the rack (shown in blue in Figure 4). If more PDUs are required, they are installed horizontally above the model AG8 controller enclosures.

The model AG8 controller enclosures (shown in yellow in Figure 4) are installed at U12 upwards, and ordered upward starting with the largest capacity at the bottom of the rack to the smallest capacity at the top. The feature codes AL01 - AL04 denote the rack position and are assigned automatically.

The I/O group ordering of the model AG8 controllers matches the rack position with the first I/O group at U12 and indicated by AL01.

The model AFF and A9F expansion enclosures (shown in purple in Figure 4) are installed with a 1U gap at the bottom of the rack. Similar to the model AG8, they are ordered upward starting with the largest capacity. The feature codes AL01 - AL04 are also used with the expansion enclosure to denote rack position.

The free rack space can be used for other components. However, an assessment must be made to ensure that suitable power and cooling is available for all of the rack's components.

Fibre Channel cabling

Figure 5 shows the rear of a model AG8 controller. It shows two PSUs on the left and right, with a pair of canisters between them. Although both canisters are identical, the top canister is inverted. The shaded yellow boxes indicate the location of the first I/O card slot. These two I/O card slots have 4 32 Gb Fibre Channel ports each and are dedicated to clustering.




Figure 5 Rear of model AG8 controller

To ensure redundancy, ports 1 and 3 from each card are connected to the first Fibre Channel switch, and ports 2 and 4 are connected to the second Fibre Channel switch. The model AG8 controllers are cabled to the switches from left to right. For more information about the Fibre Channel cabling, see [IBM Knowledge Center](#).

The outlined yellow boxes that are shown in Figure 5 indicate the location of the other available I/O card slots. You can install 25 Gb Ethernet cards in these slots, or more 32 Gb Fibre Channel cards. These cards can be used for host connectivity and remote copy to other clusters.

Fibre Channel configuration

The switches are configured to create a zone that is used for inter-cluster traffic only. This configuration creates a separation from host and remote copy Fibre Channel traffic. Maintaining this separation ensures that the inter-cluster traffic (and ultimately the operation of the cluster) are not affected by Fibre Channel congestion.

The remaining host ports that are available on each of the model AG8 controllers should be connected to a SAN network. Port masking or switch zoning can be used to determine what traffic flows on each port. For more information, see *IBM FlashSystem 9200 and 9100 Best Practices and Performance Guidelines*, [SG24-8448](#).

SAS cabling

Figure 6 shows the rear view of a model AG8 controller. Although a Fibre Channel card is shown, the purple boxes indicate the location of the SAS cards *if* model AFF or A9F expansion enclosures are used.




Figure 6 Rear of AG8 controller showing location of SAS cards

Figure 7 shows the rear of the model AFF. It has two canisters with two SAS ports each. The left side SAS port on each canister, indicated with a purple box, is connected to a SAS port in each of the model AG8 canisters. The connectivity of the A9F is similar, but because the canisters are installed inverted, the two right side SAS ports of each canister are used.



Figure 7 Rear of model AFF

Because of the restricted configurations where you can have only 0 or 1 expansion enclosure per controller enclosure, a pair of SAS cards is only included in a model AG8 by default if an associated model AFF or A9F is included in the rack. If a SAS card is not installed, a 32 Gb Fibre Channel or 25 Gb Ethernet card can be installed as an option.

Power cabling

Generally, the components in the rack use a C13 to C14 power cord to connect to the Power Distribution Units (PDU) outlets. The model A9F expansion enclosure requires C19 to C20 power cords. A pair of redundant PDUs is included in the rack with the right set of outlets to support the selected components. Where a model A9F is included, a PDU is included that features outlets for both types of power cords.

If a component has dual power supply units (PSUs), such as the model AG8 controller, one of the PSUs is connected to one PDU, with the other PSU connected to the other PDU. The switches have only a single PSU; therefore, one switch is connected to one PDU and the other switch is connected to the other PDU.

All of the power cords that are required to connect the components to the rack PDUs are automatically specified during ordering and cabled to the PDUs before the rack is delivered. The exception is the power cords that are required by the rack. These cords must be selected and ordered based on the power requirements of the environment in which the rack is installed.

Rack management

All FlashSystem 9200 clusters present themselves as a single system (that is, a single GUI with a single command line). All resources can be managed from within the GUI, with each of the model AG8 controller enclosures in the cluster shown as an I/O group. Workloads can be moved between I/O groups to balance the performance across all the available resources in the cluster.

The switches have their own separate GUI, which is used to set up the zoning, among other tasks. After the switches are set up, either switch GUI should not need to be accessed. Any issues with ports or connections to the model AG8 controllers are detected and reported from within the FlashSystem 9200 event log.

The use of SNMP or email error reporting, along with enabling IBM Call Home, improves the support experience, as does the use of Storage Insights to give you an overview of your storage estate. Storage Insights alert you to any device (or cluster) that needs attention, and highlight any performance issues. You can also give IBM service representatives access to your Storage Insights data to improve problem resolution and aid in remote debugging.

Conclusion

The FlashSystem 9200R products offer a convenient way of ordering a pre-packaged cluster with a dedicated Fibre Channel SAN for clustering. The assembly, cabling, and pre-configuration of the cluster suits customers who want to get to first I/O as quickly as possible.

For customers wanting to scale out an FlashSystem 9200 installation, or customers who want to use their own data center infrastructure, clustering can still be the answer and is fully supported outside of the FlashSystem 9200R environment. The components that are needed can be ordered separately, but consideration should be given to rack space, SAN infrastructure, power, cooling, and cabling.

Authors

This product guide was produced by a team of specialists from around the world working at IBM Redbooks, San Jose Center.

Matt Smith is a FlashSystem Offering Manager.

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

Rob MacIndoe
IBM Systems, Hursley, UK

This project was managed by:

Jon Tate
IBM ITSO

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an IBM Redbooks® residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:
ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:
<http://www.facebook.com/IBMRedbooks>
- ▶ Follow us on Twitter:
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.


COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

Redbooks (logo) ®
IBM®

IBM FlashSystem®
IBM Spectrum®

Redbooks®

The following terms are trademarks of other companies:

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



REDP-5593-00

ISBN 0738458821

Printed in U.S.A.

Get connected

