

IBM Security Privileged Identity Manager (ISPIM) Virtual Appliance Performance Tuning Guide Version 2.1

This edition applies to version 2.1.x of the IBM Security Privileged Identity Manager Virtual Appliance and Data Tier

Copyright International Business Machines Corporation 2017

US Government Users Restricted Rights

Use, duplication, or disclosure restricted by GSA ADP Schedule Contract with IBM Corp

Table of Contents

About This Publication.....	4
Intended Audience.....	4
IBM Security Privileged Identity Manager Library.....	4
Related Publications.....	4
IBM Security Privileged Identity Manager Performance Tuning Introduction.....	5
Privileged Identity Management Overview.....	5
IBM Security Privileged Identity Manager Components.....	6
ISPIM Environmental Components.....	7
VA Tier.....	7
Data Tier.....	7
IBM Security Privileged Identity Manager Tuning Focus Areas.....	7
DB2.....	7
Directory Server.....	8
Load Balancer.....	8
Session Gateway.....	8
IBM Privileged Identity Manager Data Tier.....	9
PIM Session Recorder RDB: Priming.....	9
PIM Credential Management RDB: Priming.....	9
IBM Security Directory Server RDB: Priming.....	9
Virtual Appliance Tier Recommendations.....	9
Data Tier Recommendations.....	10
Physical and Logical Processors.....	10
Performance limiting factors to take into consideration.....	11
VA Tier and Data Tier Storage.....	11
Virtual Machines: Thick Clients.....	11
Session Recorder Disk Consumption.....	12
Session Recorder User Actions.....	12
DB2 Performance Optimization.....	13
Updating Table and Index Statistics.....	13
Disk I/O Performance Recommendations.....	14
Performance variables can be set in DB2 in order to improve DB2 related processes such as:....	14

Operating Resource Policies.....	14
Memory Tuning.....	14
Access Plan Optimization.....	14
DB2 Registry Variables.....	14
Environment variables can be set in DB2 in order to modify configuration values of particular applications.....	14
DB2 System Environment Variables.....	14
Load Balancer Performance Optimization.....	15
Core Functionality Module.....	15
ngx_http_core_module.....	15
ngx_http_upstream_module.....	16
PIM Session Gateway.....	16
Concurrent Sessions.....	17
Uploader Threads.....	17
Docker Storage Drivers.....	17
Notices.....	18
Copyright License.....	20
Trademarks.....	21

About This Publication

This edition includes sections associated with troubleshooting the IBM® Privileged Identity Manager Virtual Appliance (PIMVA) tiers, best practices for configuring, tuning, and managing the PIMVA, and regular recommended maintenance tasks for ensuring optimum PIMVA performance.

Intended Audience

This document is designed for the system administrators, security administrators, and operations teams which are responsible for an organization that uses IBM Privileged Identity Manager Virtual Appliance. In addition, readers are expected to understand system administration and information systems security concepts. Additionally, the readers must understand administration concepts for the following types of products:

- Database Server
- Virtual Appliance
- Directory Server

IBM Security Privileged Identity Manager Library

You can obtain the product documentation from the IBM Privileged Identity Manager information center. The information center is available here: [IBM Security Privileged Identity Manager Version 2.1.0 documentation](#).

Related Publications

- IBM Software Support Home Page

[http://www.ibm.com/support/entry/portal/overview/software/software_support_\(general\)](http://www.ibm.com/support/entry/portal/overview/software/software_support_(general))

- IBM Publications Center

<http://www-05.ibm.com/e-business/linkweb/publications/servlet/pbi.wss>.

- IBM Developer Works

<http://www.ibm.com/developerworks/>

IBM Security Privileged Identity Manager Performance Tuning Introduction

The IBM® Security Privileged Identity Manager is an appliance-based solution that provides privileged identity management, application identity management, and session recording.

IBM Security Privileged Identity Manager Virtual Appliance features:

- A configuration wizard for the first time configuration of the IBM Security Privileged Identity Manager solution in a stand-alone or a cluster mode.
- A dashboard for viewing system status such as system notifications, cluster status, component and application status, deployment statistics, and disk usage.
- Analysis and diagnostics tools such as memory statistics, CPU utilization, and troubleshooting log files.
- Control of system settings such as host name, date, time, and network settings.
- A graphical management interface for configuring the IBM Security Privileged Identity Manager features.

Privileged Identity Management Overview

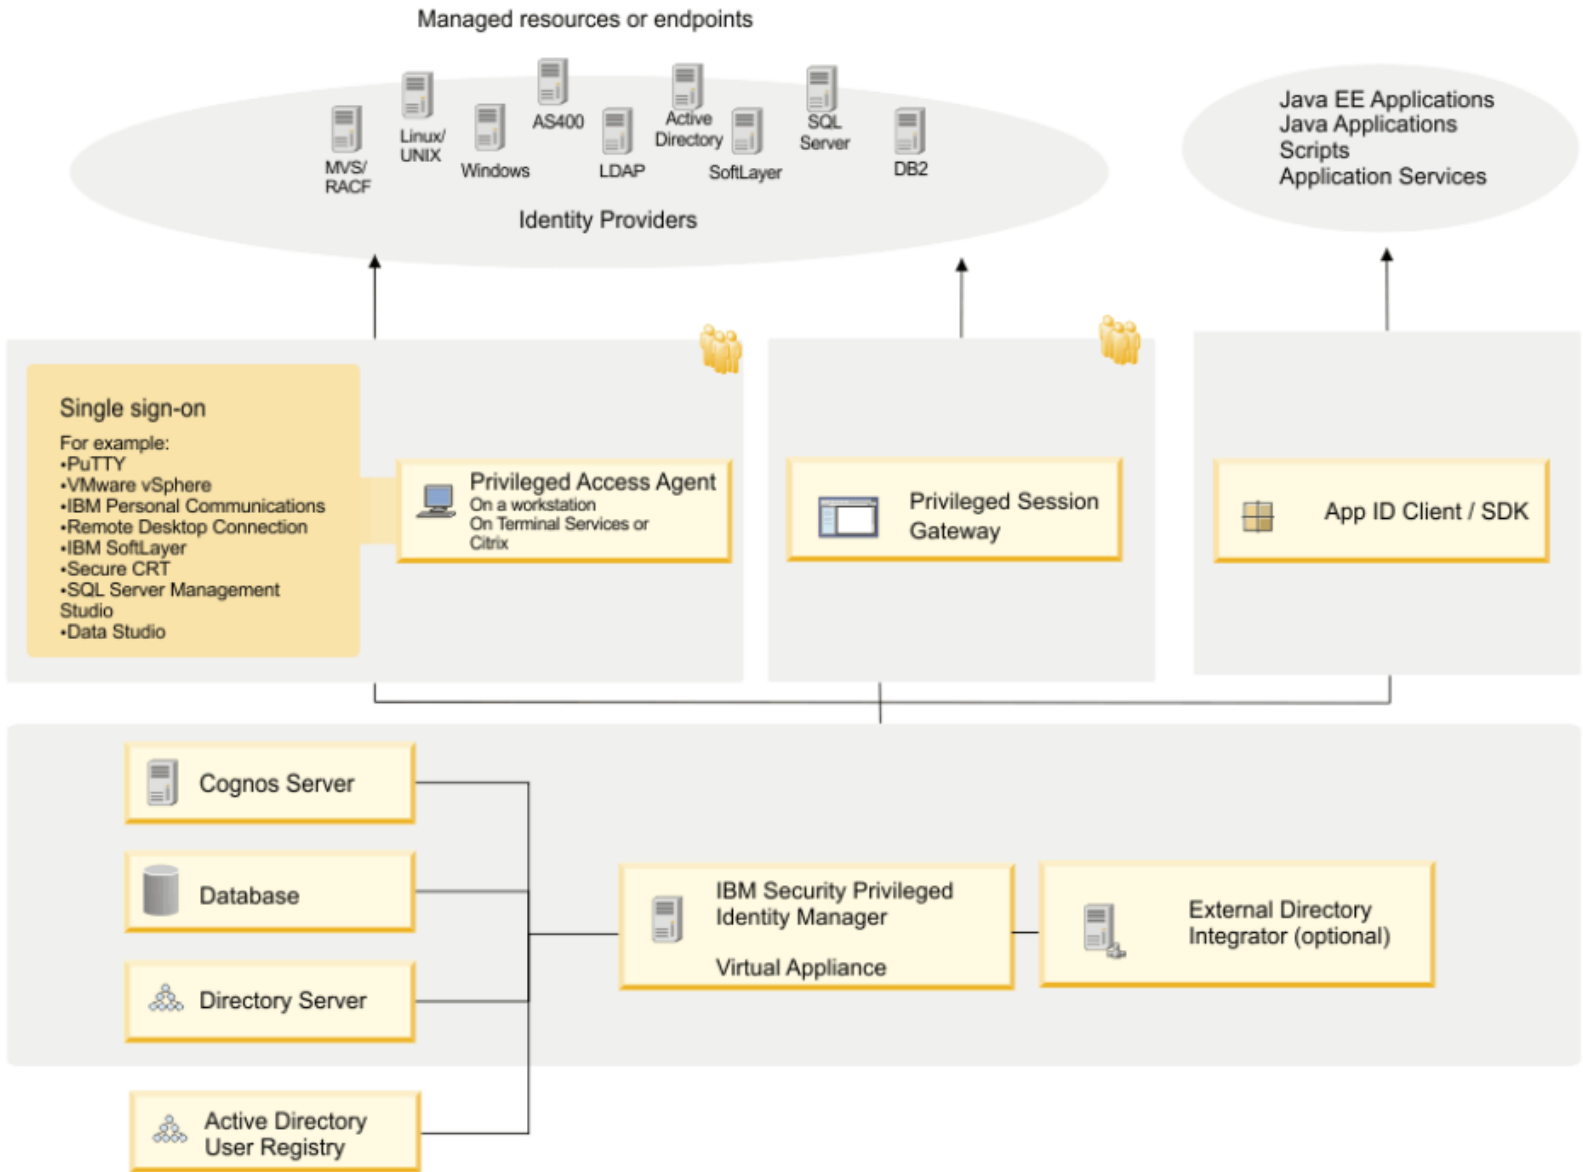
IBM® Security Privileged Identity Manager helps organizations manage, automate, and track the use of shared privileged identities.

The solution provides the following features:

- Centralized administration, secure access, and storage of privileged shared account credentials
- Role-based access control for shared account credentials
- Lifecycle management of shared accounts' password
- Single sign-on through automated check-out and check-in of shared credentials
- Auditing of shared credentials access activities
- Session recording and replay
- Integration with the broader Identity and Access Management Governance portfolio
- Application identity management

Privileged IDs are general user IDs that are distinguished by the assignment of security, administrative, or system privileges. These IDs include pre-built administrative accounts found in operating systems and applications, such as root, administrator, sa, or db2admin. In an enterprise environment, multiple Administrators might share access to a single privileged ID for easier administration. When deployed with its Single Sign-On feature, IBM Security Privileged Identity Manager allow privileged users to log on to a system without any knowledge of the password for the privileged identity.

IBM Security Privileged Identity Manager Components



ISPIM Environmental Components

VA Tier

Install IBM Security Privileged Identity Manager Virtual Appliance either as a single server or as clustered server environment with member nodes. Clustered virtual appliances are recommended due to the distinct advantages of load balancing, fault tolerance and task scalability.

Data Tier

For performance tuning/troubleshooting purposes, it is recommended that both the database server and directory server are located on separate *physical machines*. In addition, these machines should be *configured for optimal disk drive I/O performance* (example: Security Performance labs utilized SATA drives with these specifications: **RAID-10 (PERC H200 (6Gb/s)) 15K RPM.**)

IBM Security Privileged Identity Manager Tuning Focus Areas

DB2

- *Instance Memory* – This parameter specifies the maximum amount of memory that can be allocated for a database partition if you are using DB2® database products with memory usage restrictions or if you set it to a specific value. Otherwise, the AUTOMATIC setting allows instance memory to grow as needed. *In most cases, the recommended value is AUTOMATIC.*
- *IO Cleaners* – This parameter allows you to specify the number of asynchronous page cleaners for a database. This value should match the number of physical cores . When set to automatic, the formula to calculate num_iocleaners is $max(ceil(\# CPUs / \# local logical DPs) - 1, 1)$. *Recommended value is AUTOMATIC*
- *Bufferpools* – A bufferpool is an area of main memory that has been allocated by the database manager for the purpose of caching table and index data as it is read from disk. *In most environments, it is recommended to set all applicable bufferpools to automatic for each of the three PIM databases.*
- *Indexes* – Ensure that all three PIM databases contain the latest indices to improve searches.
- *Automatic Runstats Configuration* – With automatic statistics collection, you can let DB2 determine whether database statistics need to be updated. DB2 will automatically run the RUNSTATS utility in the background to ensure the most current database statistics are available.

Directory Server

- *Worker Threads* – You must determine the number of worker threads in use to assign an appropriate value for database connections. To find the number of worker threads, use the `available_workers` threads value in the monitor search result. To increase the workers threads or to increase the back-end connections, you must set `ibm-slapdDbConnections` under the `cn=Directory, cn=RDBM Backends, cn=IBM Directory, cn=Schemas, cn=Configuration` DN entry
- *LDAP Caches* – The LDAP caches include entry cache, filter cache, group member cache, and group member cache.
- *Bufferpools* – A bufferpool is an area of main memory that has been allocated by the database manager for the purpose of caching table and index data as it is read from disk. *It is recommended to set all applicable bufferpools to automatic*
- *Indexes* – Ensure that the ISDS database contain the latest indices to improve searches.
- *Automatic Runstats Configuration* – With automatic statistics collection, you can let DB2 determine whether database statistics need to be updated. DB2 will automatically run the RUNSTATS utility in the background to ensure the most current database statistics are available.

Load Balancer

- *Maximum Simultaneous Connections* (`max_conns`) – limits the maximum number of simultaneous active connections to the proxied server. Default value is zero, meaning there is no limit. If the server group does not reside in the shared memory, the limitation works per each worker process. *Security Performance labs tested with this parameter set to 0.*
- *Fail Timeout* – sets the time during which the specified number of unsuccessful attempts to communicate with the server should happen to consider the server unavailable and the period of time the server will be considered unavailable. *Security Performance Labs tested with this parameter set to 20s*

Session Gateway

- *Uploader Threads* – These are the threads associated with copying recordings from the session gateway stash to the session recorder. *Recommended value: 80*

IBM Privileged Identity Manager Data Tier

It is recommended to run the data tier on a *physical machine* rather than a virtual machine. This suggestion is based on notable overhead associated with the virtualization of I/O bound activities. In addition, when tuning a database in a newly deployed environment, it is very important to prime your database statistics. *Failing to prime the database can result in poor performance or transaction rollbacks.*

PIM Session Recorder RDB: Priming

In order to prime the PIM Session Recorder relational database, first capture a relatively small number (15K frames) of session recordings. Finally, update the table/index statistics via runstats/reorg invocation. *These database statistics tuning tasks are a vital part of the IBM Privileged Identity Manager product performance.*

PIM Credential Management RDB: Priming

Next, in order to prime the PIM Credential Management (PCM) DB, load 500 - 1000 resources via the PIMUI. Finally, update the table/index statistics via runstats/reorg invocation. *These database statistics tuning tasks are a vital part of the IBM Privileged Identity Manager product performance.*

IBM Security Directory Server RDB: Priming

In order to prime the IBM Security Directory Server relational database, first load 5K persons via DSML file. Next, update the table/index statistics via runstats/reorg invocation. *These database statistics tuning tasks are a vital part of the IBM Privileged Identity Manager product performance.*

IBM Security Privileged Identity Manager Resource Allocation

Resources such as memory, CPU (virtual sockets/cores), and storage (HDD/SSD) should be efficiently allocated at virtual machine creation time. Allocating adequate resources to the Virtual Appliance/Data tier ensures timely completion of workload and system. It is recommended to dedicate cores and memory to the virtual appliance by not using resource pools or over committing memory.

Virtual Appliance Tier Recommendations

- *CPU*
4 Socket, 2 Cores per socket (dedicated)
- *Storage*
100GB HDD/SSD

- *Memory*
16-20 GB Memory

Data Tier Recommendations

- *DB2 CPU*
4-8 Sockets, 2 Cores per socket (dedicated)
- *DB2 Storage*
1TB HDD/SSD
- *DB2 Memory*
16-20 GB Memory
4-6 GB for each database associated with ISPIM (essoDB, pimsrDB, isimDB)
- *ISDS CPU*
4-8 Sockets, 2 Cores per socket (dedicated)
- *ISDS Storage*
300GB HDD/SSD
- *ISDS Memory*
12-16 GB Memory

Physical and Logical Processors

When running the ISPIM Virtual Appliance, it is very important to dedicate actual physical CPUs on the hypervisor to the logical CPUs on the virtual machine. (ex: If the hypervisor has 'x' total sockets, one or more sockets should be dedicated to a particular VA. This ensures that the VA is not included in resource pool sharing. This would also apply to the data tier if running on a virtual machine (running the data tier on a virtual machine is not recommended for I/O bound applications and activities). Since DB2 is enabled for multi-threaded applications, it will perform best on a multi-processor server which has allocated dedicated sockets to ISPIM operations. Keep in mind, even in a well-tuned environment, system bottlenecks might vary between the processor, memory, and disk on the IBM Security Privileged Identity Manager Session Recorder server.

Performance limiting factors to take into consideration

- network throughput constraints
- firewall throttling
- network intrusion prevention systems
- network intrusion detection systems

VA Tier and Data Tier Storage

System administrators and database administrators can adjust the amount of disk space available for IBM Security Privileged Identity Manager Session Recorder and DB2. Each of the middleware components uses different amounts of disk space for various purposes. IBM Security Privileged Identity Manager Session Recorder uses local file indexes on each application server where the product is installed. Ensure that the file systems for the virtual appliances and for the physical machines on the data tier have adequate storage space to accommodate these indexes. *VA Tier recommendation: 100-120 GB.*

Virtual Machines: Thick Clients

Virtual machine thick provisioning is based on the concept of reserving all necessary space on the hard drive at the time of virtual machine creation.

Thick Provisioning Advantages

- Restricts allocation of virtual data stores based on physical HDD constraint thus preventing error cases in which virtual capacity exceeds physical capacity
- Under certain conditions, benchmarks seem to indicate better I/O performance because all of the blocks on the disk will be pre-zeroed, thus removing the need to zero the blocks at write time. This performance benefit only applies to a latency sensitive, IO bound applications such as PIM Session Recorder .

Thick Provisioning Disadvantages

- Storage space is allocated at a substantially faster rate.
- Inefficient usage of storage if initial size estimates are incorrect.

Thick Provisioning Options

- **Lazy Zeroed Thick:** The hypervisor allocates the space on the VMFS at virtual machine creation time. However, blocks of data are utilized on the back-end data store at **write-time** on the virtual machine.

- **Eager Zeroed Thick (recommended)** Hypervisor both reserves all the space on the VMFS and zeros out the disk blocks at creation time. Based on results in the laboratory, virtual machines created with eager zeroed thick provisioning will require more time to initialize. In terms of long term performance, this is the optimal deployment solution because the blocks have already been zeroed-out thus reducing the overhead associated with writing to disk.

Session Recorder Disk Consumption

IBM Security Privileged Identity Manager Session Recorder writes session recordings to the configured database. The amount of disk space required may vary depending upon your session activity and usage requirements. *Data Tier commendation: 1TB*

Session Recorder User Actions

The ISPIIM Session Recorder captures user actions in the form of frames. A frame can be defined as simply a captured snapshot of the terminal. For an action, this happens when ENTER is pressed after entering the command. If the terminal scrolls, it generates a second frame. The method for calculating frames from captured actions can be found below

GUI Applications

1 Mouse Action	1 Keyboard Action
Single Mouse Click	Filling in Text Field
Double Mouse Click	Typing 1 Sentence in Text Editor

Text Applications

1 Terminal Action
Execute 1 Terminal Command
Edit 1 Line in Text Editor
Scroll up/down 1 page in Text Editor

Converting from Action to Frame

Scenario Type	Frames per Action
Screen Recording in GUI Application	3.86
Screen Recording in Text Application	4.34
Text Recording in Text Application	2

PIM Session Recorder RDB – HDD Space Consumed: TEXTHDD Storage Used Before Frame Capture: **93915 MB (93.9 GB)**HDD Storage Used After Frame Capture: **97851 MB (97.9 GB)**Total Storage Consumed During Frame Capture: **3936000 KB/ 3936 MB/ 3.9 GB**KB per FRAME: **4.37 KB****PIM Session Recorder RDB – HDD Space Consumed: GREYSCALE**HDD Storage Used Before Frame Capture: **84505 MB (84.5 GB)**HDD Storage Used After Frame Capture: **93915 MB (93.9 GB)**Total Storage Consumed During Frame Capture: **9410000 KB/ 9419 MB/ 9.4 GB**KB per FRAME: **9.41 KB****PIM Session Recorder RDB – HDD Space Consumed: COLOUR**HDD Storage Used Before Frame Capture: **67887 MB (67.8 GB)**HDD Storage Used After Frame Capture: **84504 MB (84.5 GB)**Total Storage Consumed During Frame Capture: **16617000 KB/ 16617 MB/ 16.6 GB**KB per FRAME: **16.6 KB**

DB2 Performance Optimization

Updating Table and Index Statistics

In order to make use of the most efficient access plan for executing queries, DB2 requires accurate statistics on the exact number of rows in the tables and available indexes. Current DB2 versions can update the statistics automatically, and we recommend manually updating the statistics in certain situations after substantial changes have been made to the system.

These situations include:

- The capturing of significant number of multi-session recordings
- Creation, modification, deletion of users, resources, credentials, etc
- After an extended period of DB2 operations without updating table statistics

It is recommended to execute runstats/reorg utilities on an idle or lightly used database because it requires update locking on the system statistics table to update the database statistics. The system acquires locks on the tables that are used by the database optimizer

to fulfill queries. The locks might cause transaction rollbacks on a database with a heavy load. Also, it might be necessary to stop the directory server in order to complete runstats/reorg.

Disk I/O Performance Recommendations

Performance variables can be set in DB2 in order to improve DB2 related processes such as:

- Operating Resource Policies
- Memory Tuning
- Access Plan Optimization

DB2 Registry Variables

Environment variables can be set in DB2 in order to modify configuration values of particular applications

DB2 System Environment Variables

Due to the nature of PIM and ISIM related transactions (latency sensitive, I/O bound), IBM Security Performance recommends the Data Tier located on *physical machines* configured for speed and fault tolerance (RAID 5/RAID 10). DB2 registry variables can be modified to improve performance on the Data Tier related systems. For all systems, enable DB2_USE_ALTERNATE_PAGE_CLEANING . This variable specifies whether a DB2 database uses the alternate method of page cleaning algorithms or the default method of page cleaning. When this variable is set to ON, the DB2 system writes changed pages to disk, keeping ahead of LSN_GAP and pro-actively finding victims. Doing this allows the page cleaners to better utilize available disk I/O bandwidth. When this variable is set to ON, the *chnpggs_thresh* database configuration parameter is no longer relevant because it does not control page cleaner activity.

As the instance owner:

```
db2set DB2_USE_ALTERNATE_PAGE_CLEANING=ON
```

(instance restart required)

For SAN, RAID, and other advanced disk subsystems set the system environment variable to DB2_PARRALLEL_IO to *. This registry variable is used to change the way DB2 calculates the I/O parallelism of a table space. When I/O parallelism is enabled (either implicitly, by the use of multiple containers, or explicitly, by setting DB2_PARALLEL_IO), it is achieved by issuing the correct number of prefetch requests. Each prefetch request is a request for an extent of pages. For example, a table space has two containers and the prefetch size is four times

the extent size. If the registry variable is set, a prefetch request for this table space will be broken into four requests (one extent per request) with a possibility of four prefetchers servicing the requests in parallel. You can replace *TablespaceID* with an asterisk (*) to specify all table spaces. For example, if DB2_PARALLEL_IO=*, all table spaces use six as the number of disks per container

As the instance owner:

```
db2set DB2_PARALLEL_IO=*
```

(instance restart required)

Load Balancer Performance Optimization

[IBM HTTP Server Documentation](#)

[Nginx Documentation](#)

NGINX Parameter	Recommended Value
keepalive_timeout	100
worker_connections	1024
max_fails	3
fail_timeout	20s

Core Functionality Module

events directive

worker_connections

Sets the maximum number of simultaneous connections that can be opened by a worker process. It should be kept in mind that this number includes all connections (e.g. connections with proxied servers, among others), not only connections with clients.

Another consideration is that the actual number of simultaneous connections cannot exceed the current limit on the maximum number of open files, which can be changed by [worker_rlimit_nofile](#).

ngx_http_core_module

keepalive_timeout timeout [header_timeout]

The first parameter sets a timeout during which a keep-alive client connection will stay open on the server side. The zero value disables keep-alive client connections. The optional second parameter sets a value in the "Keep-Alive: timeout=*time*" response header field. The two parameters (keep alive/header) can be set to different values.

ngx_http_upstream_module

upstream directive

max_fails

Sets the number of unsuccessful attempts to communicate with the server that should happen in the duration set by the `fail_timeout` parameter. Once the number of unsuccessful attempts has been exceeded, the server will be considered *unavailable* for a duration also set by the `fail_timeout` parameter. By default, the number of unsuccessful attempts is set to 1. The zero value disables the accounting of attempts.

fail_timeout

- the time during which the specified number of unsuccessful attempts to communicate with the server should happen to consider the server unavailable;
- and the period of time the server will be considered unavailable.

By default, the parameter is set to 10 seconds.

PIM Session Gateway

The Privileged Session Gateway (PSG) grants privileged users the ability to log on to SSH managed resources from a web-based console without installing extra client software on workstations. In terms of deployment location, the PIM Session Gateway is situated between the privileged users end point and the managed resource.

PSG Features:

- *Total Control With Gateway.* Privileged access control is now centralized and resides on customer controlled appliance deployed between the privileged users' client end point and managed hosts.
- *Seamless Single Sign-On.* secure privileged session single sign on without exposing the privileged credentials to the client end point.
- *Agentless Deployment.* Privileged users are able to access the managed hosts without installing any of our software on their machine.
- *Clientless Access.* Privileged users can simply access the managed hosts using a web browser on any devices including mobile devices.

Concurrent Sessions

Depending on the exact characteristics of the workload, each PSG instance can safely support up to 800 - 1000 concurrent sessions. However, due to limitations with multi-

threading, the recommended number of concurrent sessions is *200 concurrent sessions* per PSR instance. If more than 200 sessions are active simultaneously, failures in Session Gateway related components can occur .

Uploader Threads

The uploader microservice is a multi-threaded process that pushes recordings from the recording stash to the PIM Session Recorder database. This value is set in the `ibm-session-gateway.rsp` file.

```
uploader_threads=80
```

The default number of threads is 5. Based on testing, *Security Performance recommends 80 threads to achieve the proper balance of CPU and memory utilization.*

Docker Storage Drivers

The PIM Session Gateway (PSG) is deployed as docker container, thus it consists of a read-write filesystem which is layered over multiple read-only filesystems that make up a Docker image. Each read-only filesystem contains:

- Any additional files added when the Docker image was built
- The base OS layer as the foundation

The PSG 's multi-tiered filesystem is implemented by a storage driver. As updates to the files from a read-only layer occur, they are copied up from the read-only layer to the read-write layer in a strategy known as copy-on-write.

File-System	Mount Point
/dev/dm-5	//var/lib/docker/devicemapper/mnt/05cf2a02e01c851fd927a77aa79e686f5ec428878afa86514c4
shm	//var/lib/docker/containers/7e336cf3eb25e162a8c14a288db0210e73ac4dd4391bb26e18a4b/shm

The PIM Session gateway ships with a default configuration of **devicemapper, loop-lvm mode**. The main advantages to the loop-lvm mode are:

- ease of operation
- rapid deployment

However, *loop-lvm-mode is not recommended for production use.* According to official Docker documentation."loopback devices are slow and resource-intensive, and require you to create file on disk at specific sizes. They can also introduce race conditions."

It is for this reason that Security Performance recommends loopback devices to be configured in *direct-lvm mode*. "Production hosts using the devicemapper storage driver must use `direct - lvm` mode. This mode uses block devices to create the thin pool. This is faster than using loopback devices, uses system resources more efficiently, and block devices can grow as needed. However, more set-up is required than `loop - lvm` mode." [Official Docker Container Documentation](#)

Notices

This information was developed for products and services offered in the U.S.A. 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 give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing

IBM Corporation

North Castle Drive

Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing

Legal and Intellectual Property Law

IBM Japan, Ltd.

1623-14, Shimotsuruma, Yamato-shi

Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law :

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 states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement might 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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you. Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation

11501 Burnet Road

Austin, TX 78758 U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases payment of a fee. The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

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

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

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 the names and addresses used by an actual business enterprise 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. 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 IBM's application programming interfaces.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year).

Portions of this code are derived from IBM Corp.

Sample Programs. © Copyright IBM Corp. _enter the year or years_. All rights

reserved.

If you are viewing this information in softcopy form, the photographs and color illustrations might not be displayed.

Trademarks

- IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., 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

www.ibm.com/legal/copytrade.shtml.

- Adobe, Acrobat, PostScript and all Adobe-based trademarks are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.
- Linux is a trademark of Linus Torvalds in the United States, other countries, or both.
- Microsoft, Windows, Windows Server 2008, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and its affiliates.
- Advanced Visualization Powered by IBM ILOG Elixir Enterprise.
- Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
- ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office. IT Infrastructure Library (ITIL) is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

- Cell Broadband Engine and Cell/B.E. are trademarks of Sony Computer Entertainment, Inc., in the United States, other countries, or both and is used under license therefrom.

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