

September Newsletter

[Subscribe/Unsubscribe](#)

Greetings all,

I hope that for those going to Orlando next week, you have an action filled Agenda and build lots of new networks. Sorry that I am unable to join you.

A few updates to share

- Updating fibre adapter firmware?
Read the details in the link, and be aware that it is recommended to stop any I/O flow, even if in a dual VIO Server configuration, your clients may be affected.
[Link](#)
- Important NIM updates
Support has published a short note covering the recent updates in NIM - including the important security fixes in nimsh.
[Link](#)
- NIM FAQ
The NIM FAQ page was also updated. This is primarily to handle the NIMSH fixes and includes a reminder that IBM recommends securing NIM communications.
[Link](#)
- Auto Zoning... really???
For Power Systems using IBM FlashSystems and Brocade Fabrics, the collection of WWPNs, creating and maintaining zones on the Fabric switches, creating the storage mappings and validation can now be automated!
[Link](#)
- Want to hop on the Bob-sled?
A beginner friendly hands-on workshop (sadly not very friendly times for AP!)
[Link](#)
- TechXchange Orlando '26
Don't forget to build your Agenda for the TechXchange at the end of this month.
- Warning about llu_mode=1 for AIX 7.3 TL3 & TL4
Background: The Live Library Update (LLU) function eliminates downtime for workloads when the AIX operating system is updated. The LLU function shifts the applications from using the old library to the updated library without any downtime. If the llu_mode tunable parameter is set to 1, the LLU function is enabled by default for all processes unless explicitly overridden by using the LLU program attribute. Christian and Chris have warned that there are some issues with the operation on AIX 7.3 TL3/4, so best not to set until these are resolved.



- Planning Power9, Power10 or Power11?

In order to plan correctly for system firmware maintenance, there is a need to know when new Releases are coming out along with when a Release is planned to go end of service pack support. Therefore, the below table was created to provide this detail. Things can come up that alters this schedule somewhat, but this is the current plan. As things change, the table is updated to reflect status.

[Link](#)

Quick bites

The August Power Service pack releases

Support have put out a number of useful updates around the critical fixes that came out last month (please check them out, if you haven't started planning your updates already), for example, do you wonder what is the order of the updates? See [Link](#) (Note a number of reboots are required)

VIOS Updates: Harnessing the Power of Alternate Disk Update

The `updateios` command updates the Virtual I/O Server to the latest maintenance level. With the release of VIOS 4.1.2.00, a significant enhancement has been introduced: support for updateios on an alternate disk. This feature is designed to improve system availability and reduce risk during updates. This means that the VIO server can stay online during update, minimising downtime with a clean/easy roll back.

[Link](#)

VIO Server 4.1.2.20

Note: the AIX level is 7.3.4.2, so the NIM Server has to be at this level or higher

Note 2: If you have installed packages from the VIO Server expansion pack, they will also need to be updated. If there is firmware to be updated, update it first!

Note 3: The update may fail if you have a loaded media repository virtual CD

Details:

- Existing VIOS systems with supported versions of 3.1.n.m can be upgraded to VIOS version 4.1.2.20 (DVD image) or 4.1.2.20 (Flash image) using viosupgrade tool. Although it is recommended to be on VIOS 3.1.4.30 or later SP level before upgrading to VIOS 4.1.x level.
- VIOS systems with SSP configuration must be on 3.1.4.50 or later level before upgrading to 4.1.x level or adding 4.1.x nodes into cluster.
- If Active Memory Sharing (AMS) is configured on the VIOS, it should be un-configured before upgrading.
- If updating from 4.1.1.00 or later and encounter a boot issue, it may be due to an incorrect combination of disk attributes – see the readme.
- If doing a NIM altdisk update, the operation may not return a completed successfully message – again check the readme.
- Customers using NVMe Over Fabric may experience a failure to boot – discover the LUN via the SMS menu and try to boot.



- Some adapters (such as the 32Gb PCIe4 2 port FC adapter must have their firmware updated prior to the upgrade – see readme).

New features

Live Partition Mobility (LPM) enhancements

You can bypass the port-level validation for the N_Port ID Virtualisation (NPIV) devices for logical partition (lpar) migration with Linux or IBM i operating systems. This action successfully completes the LPM operation if all the Logical Unit Numbers (LUN) are available across both the source and destination VIOS instances, irrespective of the storage ports through which they are accessible.

VIOS command enhancements

The *fcstat* command displays the I/O latency statistics across multiple layers like Storage Area Network (SAN) and fiber channel adapter. For virtual AIX clients, you can also view additional statistics for VIOS, hypervisor and client Virtual Fibre Channel (VFC) adapter layers.

The *updatevios* command uses alternate disks to install new updates and retain the current state of the rootvg disk.

The following AIX commands can be accessed in VIOS through the Role Based Access Control (RBAC) security functionality:

mkitab command

lsitab command

rmitab command

chitab command

The *viosecure* command permits or denies firewall rules for an entire subnet range along with the IP address.

The *lsmap* command filters NPIV device mappings based on the partition name and operating system type. Use the *-cpname* and *-cpas* options along with the *-npiv* option to display the relevant device mappings.

Virtual Fibre Channel (VFC) adapter enhancements

The Quality of Service (QoS) feature in VIOS is a technological preview for the VFC stack. It enables you to manage and prioritise storage traffic between a virtual machine (VM) and a Fibre Channel (FC) storage area network (SAN) by setting performance limits on VFC host adapters.

During critical events, VIOS automatically sends notifications to client lpar so that the lpar can take action on these critical events. This rectification action reduces the disruption or delay in the I/O processing activity if other unaffected VFC paths are available. These critical events may occur due to the following disruptive operations:

Firmware (microcode) updates

Extended Enhanced Error Handling (EEH) events

Live dump operations

For more information, see the latest AIX and IBM i documentation.

Shared Storage Pool (SSP) enhancements

When the VIOS nodes in an SSP cluster is upgraded to VIOS version 4.1.2.20 or later, the Node Upgrade Status of the nodes changes to the UP_LEVEL state. After



all the nodes are upgraded to the 4.1.2.20 level, the rolling upgrade process starts automatically and the Node Upgrade Status level changes to the ON_LEVEL state. Log recording is enabled for the SSP- related commands on the VIOS command line interface (CLI).

Shared Ethernet Adapter (SEA) accounting enhancements

The SEA accounting functionality monitors the network traffic to and from the client virtual Ethernet adapters it serves. Enhancements have been made in the functionality to reduce the overhead in SEA accounting. This improves the overall SEA performance.

Known Capabilities and Limitations

Requirements for Shared Storage Pool

- System requirements per SSP node:
- Minimum CPU: 1 CPU of guaranteed entitlement
- Minimum memory: 4GB
- Storage requirements per SSP cluster (minimum): 1 fiber-channel attached disk for repository, 10 GB
- At least 1 fiber-channel attached disk for data, 10GB

SSP Configuration

Feature	Min	Max
Number of VIOS Nodes in Cluster	1	16*
Number of Physical Disks in Pool	1	1024
Number of Virtual Disks (LUs) Mappings in Pool	1	8192
Number of Client LPARs per VIOS node	1	250*
Capacity of Physical Disks in Pool	10GB	16TB
Storage Capacity of Storage Pool	10GB	512TB
Capacity of a Virtual Disk (LU) in Pool	1GB	4TB
Number of Repository Disks	1	1
Capacity of Repository Disk	10GB	1016GB
Number of Client LPARs per Cluster	1	2000

Support for additional VIOS Nodes and LPAR Mappings:

Over 16 VIOS Nodes requires that the SYSTEM (metadata) tier contains only SSD storage.

Over 250 Client LPARs per VIOS requires each VIOS have at least 4 CPUs and 8 GB memory.

Here are the new maximum values for each of these configuration options, if the associated hardware specification has been met:

Feature	Default Max	High Spec Max
Number of VIOS Nodes in Cluster	16	24
Number of Client LPARs per VIOS node	250	400

Other Notes:

- Maximum number of physical volumes that can be added to or replaced from a pool at one time: 64
- The Shared Storage Pool cluster name must be less than 63 characters long.



- The Shared Storage Pool pool name must be less than 127 characters long.
- The maximum supported LU size is 4TB, however for high I/O workloads recommendation is to use multiple smaller LUs as it will improve performance. For example, using 16 separate 16GB LUs would yield better performance than a single 256GB LU for applications that perform reads and writes to a variety of storage locations concurrently.
- The size of the /var drive should be greater than or equal to 3GB to ensure proper logging.

Network Configuration

Uninterrupted network connectivity is required for operation. i.e., the network interface used for Shared Storage Pool configuration must be on a highly reliable network which is not congested.

A Shared Storage Pool configuration can use IPV4 or IPV6, but not a combination of both.

A Shared Storage Pool configuration should configure the TCP/IP resolver routine for name resolution to resolve host names locally first, and then use the DNS. For step by step instructions, refer to the TCP/IP name resolution documentation in the IBM Knowledge Centre.

The forward and reverse lookup should resolve to the IP address/hostname that is used for Shared Storage Pool configuration.

It is recommended that the VIOSs that are part of the Shared Storage Pool configuration keep their clocks synchronised.

Storage Configuration

- Virtual SCSI devices provisioned from the Shared Storage Pool may drive higher CPU utilisation than the classic Virtual SCSI devices.
- Using SCSI reservations (SCSI Reserve/Release and SCSI-3 Reserve) for fencing physical disks in the Shared Storage pool is not supported.
- SANCOM will not be supported in a Shared Storage Pool environment.

Shared Storage Pool capabilities and limitations

- On the client LPAR Virtual SCSI disk is the only peripheral device type supported by SSP at this time.
- When creating Virtual SCSI Adapters for VIOS LPARs, the option "Any client partition can connect" is not supported.
- VIOSs configured for SSP require that Shared Ethernet Adapter(s) (SEA) be setup for Threaded mode (the default mode). SEA in Interrupt Mode is not supported within SSP.
- Client LPARs are not supported if they use JFS as their filesystem. If JFS is used, there is a risk of data corruption in the event of a network outage. JFS2 and other file systems are unaffected by this issue.

VIOS upgrades with Third Party Software

When the user upgrades from 3.1.x.y to 4.1.0.00 level and above, third-party software is not packaged with the IBM supplied mksysb image. User needs to install the respective third-party software after the upgrade is complete and run `viosupgrade -o rerun` to restore the respective devices.



Warning: If VIOS rules have been deployed.

During update, there have been occasional issues with VIOS Rules files getting overwritten and/or system settings getting reset to their default values.

To ensure that this doesn't affect you, we recommend making a backup of the current rules file (home/padmin/rules/vios_current_rules.xml), capture the current settings (rules -o capture) and check that they match after the upgrade.

Check the release notes and readme for more details and the post upgrade steps / checking.

[Link](#)

Support has also published some best practice recommendations re the upgrade that also covers many of the above points.

[Link](#)

Configuring quantum safe analysis

You can configure a quantum safe analysis scan for a selected PowerSC endpoint. You can also turn on Port Scan to enable TLS port scanning. Note that this is only for AIX and Linux endpoints.

Email me if you would like a demo!

[Link](#)

IBM introduces IBM Concert Subscription for Power Version 3

IBM Concert Subscription for Power makes the Power specific subset of IBM Concert capabilities for Power Infrastructure Vulnerability Management available to Power customers as a stand-alone product which can be ordered through IBM Power sales channels under a subscription license.

This will be available after 8th October – contact your friendly IBM Rep.

Redbooks and Redpapers

- **IBM Power S1112 Technical Overview and Introduction**, Redbook, published 18 September 2026.
[Link](#)
- **The Definitive Guide to IBM Storage FlashSystem and AIX Integration**, Draft Redbook, 09 September 2026.
[Link](#)
- **Hybrid Cloud on IBM Power11**, Redbook, published 04 September 2026.
[Link](#)

IBM alerts and notices**AIX / PowerVM alerts:**

- **Multiple vulnerabilities impact AIX due to ISC BIND**
Vulnerabilities in ISC BIND could cause: Allowing forged authenticated NXDOMAIN responses (CVE-2026-10723), Denial of service (CVE-2026-10822, CVE-2026-11331, CVE-2026-11622, CVE-2026-12617, CVE-2026-13204), DNS



cache poisoning (CVE-2026-11721, CVE-2026-13321). AIX uses ISC BIND as part of its DNS functions.

Vulnerability Details

CVE-2026-10723 - BIND may accept incorrect child-zone NSEC3 records as valid, which could allow an attacker to forge authenticated NXDOMAIN responses. This issue affects BIND 9 versions 9.18.0 through 9.18.50, 9.20.0 through 9.20.24, 9.21.0 through 9.21.23, 9.11.3-S1 through 9.18.50-S1, and 9.20.9-S1 through 9.20.24-S1.

CVE-2026-10822 - If BIND encounters a particular invalid data structure in a DNS record, it will accept the invalid data, and may subsequently abort and exit. BIND will first need to store a DNS record for a key (KEY, DNSKEY, etc.). That key must specify a PRIVATEDNS algorithm (253), and in the algorithm identifier, improperly give a length longer than the actual identifier data. The invalid identifier will be stored. If BIND later needs to render that record to text, it will use the invalid length during processing, leading to a consistency check failing. This issue affects BIND 9 versions 9.18.0 through 9.18.50, 9.20.0 through 9.20.24, 9.21.0 through 9.21.23, 9.18.11-S1 through 9.18.50-S1, and 9.20.9-S1 through 9.20.24-S1.

CVE-2026-11622 - A DNSSEC validating resolver that is under a random subdomain attack against a DNSSEC-signed zone can suffer from runaway memory usage. The attacker needs to be able to send queries faster than the resolver can perform validation. The increased memory usage can be orders of magnitude beyond the limit configured in the `max-cache-size` parameter. This issue affects BIND 9 versions 9.11.0 through 9.18.50, 9.20.0 through 9.20.24, 9.21.0 through 9.21.23, 9.11.3-S1 through 9.18.50-S1, and 9.20.9-S1 through 9.20.24-S1.

CVE-2026-11721 - It is possible for an attacker's zone to respond to a query with an RRSIG that has a smaller number of labels than the zone in which the RRSIG is contained. This causes `named` to produce a wildcard name for a zone that is shorter than the attacker's zone, which can result in cache poisoning. For this attack to have any effect, the resolver under attack must have set `synth-from-dnssec yes;` (which is the default). This issue affects BIND 9 versions 9.11.0 through 9.18.50, 9.20.0 through 9.20.24, 9.21.0 through 9.21.23, 9.11.3-S1 through 9.18.50-S1, and 9.20.9-S1 through 9.20.24-S1.

CVE-2026-13204 - If a provably insecure domain is covered by both an NSEC and NSEC3 record at the parent, and there exist an RRSIG for only one of these types, then BIND may exit unexpectedly with an assertion while validating this proof. This issue affects BIND 9 versions 9.11.0 through 9.18.50, 9.20.0 through 9.20.24, 9.21.0 through 9.21.23, 9.11.3-S1 through 9.18.50-S1, and 9.20.9-S1 through 9.20.24-S1.

CVE-2026-13321 - The BIND resolver accepts validly-signed NSEC records where the "Next Domain Name" field points outside the signer's zone. This issue affects BIND 9 versions 9.11.0 through 9.18.50, 9.20.0 through 9.20.24,



9.21.0 through 9.21.23, 9.11.3-S1 through 9.18.50-S1, and 9.20.9-S1 through 9.20.24-S1.

CVE-2026-11331 - An attacker who knows (or guesses) that a resolver uses RPZ with wildcard CNAME policies can craft query names long enough to trigger a NAMETOOLONG error condition during RPZ processing. This is not handled correctly and may lead to defeating the RPZ rule. It also may lead to an unexpected exit of the BIND 9 software. This issue affects BIND 9 versions 9.16.0 through 9.18.50, 9.20.0 through 9.20.24, 9.21.0 through 9.21.23, 9.16.8-S1 through 9.18.50-S1, and 9.20.9-S1 through 9.20.24-S1.

CVE-2026-12617 - The issue is unexpected program termination based on ordering and/or specific content in responses to queries for CNAME or DNAME, and A records. Specifically, if a client queries for a DNAME and A record below the DNAME to the resolver, and the authoritative server responds positively to the A query but delays the DNAME response and later responds negatively, `named` may quit unexpectedly. Or, if a client queries for a CNAME and A record for the same name to the resolver, and the authoritative server responds positively to the A query but delays the CNAME response and later responds with a self-referential CNAME, the same failure may occur. This issue affects BIND 9 versions 9.18.0 through 9.18.50, 9.20.0 through 9.20.24, 9.18.11-S1 through 9.18.50-S1, and 9.20.9-S1 through 9.20.24-S1.

Affected Products and Versions

Affected Product(s)	Version(s)
AIX	7.2
AIX	7.3
VIOS	4.1

The vulnerabilities are being addressed for the following filesets:

key_fileset = aix

Fileset	Lower Level	Upper Level	KEY
bind.rte	7.1.916.0	7.1.916.4800	key_w_fs
bind.rte	7.2.916.0	7.2.916.4801	key_w_fs
bind.rte	7.2.918.0	7.2.918.4900	key_w_fs
bind.rte	7.3.916.0	7.3.916.4800	key_w_fs
bind.rte	7.3.918.0	7.3.918.4900	key_w_fs

[Link](#)

- **Multiple vulnerabilities in Python affect AIX**

Vulnerabilities in Python: Improper Input Validation (CVE-2025-13462, CVE-2026-3644), Denial of Service (CVE-2026-15308, CVE-2026-7210, CVE-2026-66046), Use-After-Free / Memory Corruption (CVE-2026-76957), Out-of-Bounds Read / Infinite Loop (CVE-2026-72522, CVE-2026-76641), Insufficient Entropy vulnerability (CVE-2026-76956). Python is used by AIX as part of Ansible node management automation.

Vulnerability Details



CVE-2026-7210 - ``xml.parsers.expat`` and ``xml.etree.ElementTree`` use insufficient entropy for Expat hash-flooding protection, which allows a crafted XML document to trigger hash flooding.
Fully mitigating this vulnerability requires both updating libexpat to 2.8.0 or later and applying this patch.

CVE-2025-13462 - The "tarfile" module would still apply normalization of AREGTYPE (`\x00`) blocks to DIRTYPE, even while processing a multi-block member such as GNUTYPE_LONGNAME or GNUTYPE_LONGLINK. This could result in a crafted tar archive being misinterpreted by the tarfile module compared to other implementations.

CVE-2026-15308 - The incremental HTML parser (`html.parser.HTMLParser`) allows for CPU denial-of-service through repeated unterminated markup declarations when processing uncontrolled data.

CVE-2026-3644 - The fix for CVE-2026-0672, which rejected control characters in `http.cookies.Morsel`, was incomplete. The `Morsel.update()`, `|=` operator, and unpickling paths were not patched, allowing control characters to bypass input validation. Additionally, `BaseCookie.js_output()` lacked the output validation applied to `BaseCookie.output()`.

CVE-2026-76641 - Expat through 2.8.3 contains an out-of-bounds read vulnerability that allows attackers to trigger memory corruption by processing XML with external entity parsers created via

`XML_ExternalEntityParserCreate`. A struct size mismatch between `ELEMENT_TYPE` members causes `storeAtts` to read the `attIndex` member past allocated memory boundaries, resulting in failure to normalize whitespace in non-CDATA attributes or a wild pointer dereference causing a segfault. This vulnerability was introduced by the fix for CVE-2026-66046.

CVE-2026-76956 - In libexpat 2.8.2 and 2.8.3 before 2.8.4, misinterpretation of `getentropy`'s return code leads to insufficient entropy, which results in being vulnerable to hash flooding attacks, causing a denial of service via crafted XML content.

CVE-2026-72522 - libexpat before 2.8.3 has an out-of-bounds read and resultant infinite loop because low surrogates are treated the same as high surrogates during Unicode processing in the `*_toUtf16` functions.

CVE-2026-66046 - Expat through 2.8.3 contains a denial of service vulnerability caused by quadratic algorithmic complexity in the `storeAtts()` function in `xmlparse.c`, where processing `N` specified attributes with non-normalized values triggers an $O(N^2)$ linear scan of `elementType-defaultAtts` to determine CDATA status. A remote unauthenticated attacker can supply a single well-formed XML document of a few megabytes to an application parsing untrusted XML to cause excessive CPU consumption, resulting in denial of service without requiring authentication, external entity resolution, or non-default parser options.



CVE-2026-76957 - libexpat before 2.8.4 lacks handler call depth tracking with custom encoding callbacks. Thus, a use-after-free can occur. NOTE: this is similar to CVE-2026-50219, CVE-2026-56131 and CVE-2026-56412.

Affected Products and Versions

Affected Product(s)	Version(s)
AIX	7.3
VIOS	4.1

The vulnerabilities in the following filesets are being addressed:

key_fileset = osrcaix

Fileset	Lower Level	Upper Level	KEY
python3.9.base	3.9.0.0	3.9.24.0	key_w_fs
python3.11.base	3.11.0.0	3.11.15.1	key_w_fs

[Link](#)

- **Multiple vulnerabilities in OpenSSH affect AIX**

There are multiple vulnerabilities in OpenSSH used by AIX (CVE-2026-73281, CVE-2026-73282, CVE-2026-73283). OpenSSH is used by AIX for remote login.

Vulnerability Details

CVE-2026-73281 - In ssh-agent in OpenSSH before 10.5, some operations can occur remotely but were intended to occur only locally, including operations that add tokens or use keys. This is caused by misinteraction between agent locking and the session-bind@openssh.com extension.

CVE-2026-73282 - In ssh in OpenSSH before 10.5, a use-after-free for realloc data can occur if a certain pair of remote-forwarding operations are concurrent.

CVE-2026-73283 - In sshd in OpenSSH before 10.5, the restrict keyword (in authorized_keys) was supposed to be applicable to tunnel forwarding but was not.

Affected Products and Versions

Affected Product(s)	Version(s)
AIX	7.2
AIX	7.3
VIOS	4.1

The vulnerabilities are being addressed in the following filesets:

key_fileset = aix

Fileset	Lower Level	Upper Level
openssh.base.client	9.2.112.0	9.7.3013.1000
openssh.base.server	9.2.112.0	9.7.3013.1000
openssh.base.client	9.9.3015.1000	9.9.3015.2001
openssh.base.server	9.9.3015.1000	9.9.3015.2001

[Link](#)



Firmware alerts:

- **IBM Machine Microcode Services for IBM Power**

Effective on the dates shown in the linked document, select feature codes from the following program numbers will no longer be available for ordering. Entitlement and support for all other feature codes in the program will continue for the duration of your contract.

[Link](#)

[Power Announcements](#)

GPFS/Scale alerts:

- **FIFO (named pipe) read/write hangs on nodes that run IBM Storage Scale 5.2.3.8 - 5.2.3.9 and 6.0.1.0 - 6.0.1.1**

Read and write operations on FIFO (named pipe) files in an IBM Storage Scale file system hang indefinitely on nodes that run IBM Storage Scale 5.2.3.8 - 5.2.3.9 or 6.0.1.0 - 6.0.1.1. FIFO creation succeeds, but processes that attempt to access the FIFO for reading or writing remain blocked indefinitely.

Content

A FIFO file that is created with `mkfifo file.fifo` is successfully created and is visible in directory listings. However, any process that subsequently attempts to read or write to it, such as `cat file.fifo` or `echo data > file.fifo`, remains blocked indefinitely.

The following example demonstrates the issue:

```
[root@node ~]# mkfifo /ibm/fs1/test/node.fifo
[root@node ~]# ssh node2 "chmod 0600 /ibm/fs1/node.fifo && chmod 0666 /ibm/fs1/node.fifo"
[root@node ~]# echo HI > /ibm/fs1/test/node.fifo
# second session:
[root@node ~]# cat /ibm/fs1/test/node.fifo # hangs — no output
```

Users Affected

All IBM Storage Scale users on Linux client/compute nodes at the affected versions (5.2.3.8 - 5.2.3.9 and 6.0.1.0 - 6.0.1.1) and whose applications create or access FIFO (named pipe) files within an IBM Storage Scale file system.

Typical affected workloads include the following:

Applications that use named pipes for inter-process communication

Shell scripts and data pipelines that redirect output through FIFOs stored in an IBM Storage Scale file system

Any application that calls `mkfifo()` and then attempts to read or write to the resulting file in an IBM Storage Scale file system

[Link](#)

Keep safe and talk soon,
Red, Belisama

