Integration News

Q2 2023

IBM Sterling B2B Collaboration Suite: SECURITY NEWS.

In this issue:

- IBM Sterling B2B Integrator vulnerable to sensitive information exposure due to IBM MQ
- IBM Sterling Connect:Direct for UNIX is affected by security restriction bypass due to Spring Framework
- IBM Sterling External Authentication Server is vulnerable to multiple vulnerabilities due to IBM Java Runtime

IBM Sterling B2B Integrator

IBM Sterling B2B Integrator vulnerable to sensitive information exposure due to IBM MQ

CVEID: CVE-2022-42436

IBM Sterling Connect:Direct for UNIX

IBM Sterling Connect:Direct for UNIX is affected by security restriction bypass due to Spring Framework

CVEID: CVE- 2023-20860

IBM Sterling Secure Proxy

IBM Sterling External Authentication Server is vulnerable to multiple vulnerabilities due to IBM Java Runtime

CVEID: CVE-2022-21628
CVEID: CVE-2022-21626
CVEID: CVE-2022-3676
CVEID: CVE-2022-21426
CVEID: CVE-2022-21830

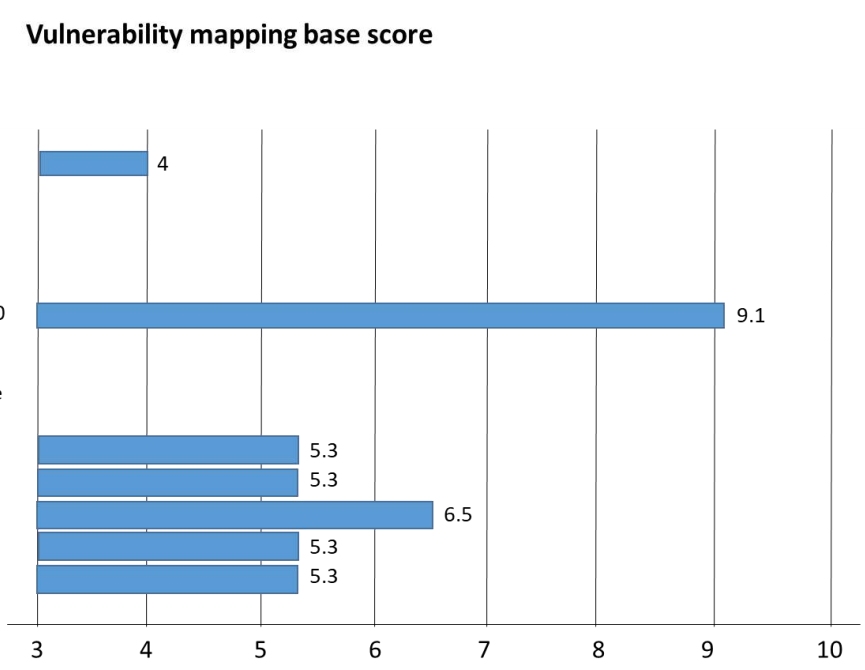

IBM Sterling B2B Integrator vulnerable to sensitive information exposure due to IBM MQ

Vulnerability Details

CVEID: CVE-2022-42436

Description: IBM MQ 8.0.0, 9.0.0, 9.1.0, 9.2.0, 9.3.0 Managed File Transfer could allow a local user to obtain sensitive information from diagnostic files. IBM X-Force ID: 238206

CVSS Base score: 4
CVSS Temporal Score: Click here.
CVSS Vector: (CVSS:3.0/AV:L/AC:L/PR:N/UI:N/S:U/C:L/I:N/A:N)N/UI:N/S:U/C:N/I:H/A:N)

Affected Products and Versions

Affected Product(s)	Version(s)
IBM Sterling B2B Integrator	6.0.0.0 - 6.0.3.7
	6.1.0.0 - 6.1.2.0

Remediation/Fixes

Product	IBM Sterling B2B Integrator	
Version	6.0.0.0 - 6.0.3.7	6.1.0.0 - 6.1.2.0
APAR	IT43073	
Remediation & Fix	Apply 6.0.3.8	Apply 6.1.2.1

Workarounds and Mitigations

None.

IBM Sterling Connect:Direct for UNIX is affected by security restriction bypass due to Spring Framework

IBM Sterling Connect:Direct for UNIX File Agent component is affected by security restriction bypass due to Spring Framework. Spring Framework has been upgraded in IBM Sterling Connect:Direct for UNIX File Agent component. [CVE-2023-20860]

Vulnerability Details

CVEID: CVE-2023-20860

Description: VMware Tanzu Spring Framework could allow a remote attacker to bypass security restrictions, caused by the use of an un-prefixed double wildcard pattern with the mvcRequestMatcher in Spring Security configuration. An attacker could exploit this vulnerability to create a mismatch in pattern matching between Spring Security and Spring MVC

CVSS Base score: 9.1
CVSS Temporal Score: Click here.
CVSS Vector:
(CVSS:3.0/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:N)

Affected Products and Versions

Affected Product(s)	Version(s)
IBM Sterling Connect:Direct for UNIX	6.2.0.0 - 6.2.0.6.iFix012

Remediation/Fixes

Product	IBM Sterling Connect:Direct for UNIX
Version	6.2.0
Remediation & Fix	Apply 6.2.0.6.iFix013

Workarounds and Mitigations
None.

IBM Sterling External Authentication Server is vulnerable to multiple vulnerabilities due to IBM Java Runtime

There are multiple vulnerabilities in IBM® Runtime Environment Java™ Version 1.8 used by IBM Sterling External Authentication Server. IBM Sterling External Authentication Server has addressed the applicable CVEs.

Vulnerability Details

CVEID: CVE-2022-21628

Description: Java SE is vulnerable to a denial of service, caused by a flaw in the Lightweight HTTP Server. By sending a specially-crafted request, a remote attacker could exploit this vulnerability to cause a denial of service condition.

CVSS Base score: 5.3
CVSS Temporal Score: Click here.
CVSS Vector:
(CVSS:3.0/AV:N/AC:L/PR:N/UI:N/S:U/C:N/I:N/A:L)

CVEID: CVE-2022-21626

Description: An unspecified vulnerability in Java SE related to the Security component could allow an unauthenticated attacker to cause a denial of service resulting in a low availability impact using unknown attack vectors.

CVSS Base score: 5.3
CVSS Temporal Score: Click here.
CVSS Vector:
(CVSS:3.0/AV:N/AC:L/PR:N/UI:N/S:U/C:N/I:N/A:L)

CVEID: CVE-2022-3676

Description: Eclipse Openj9 could allow a remote attacker to bypass security restrictions, caused by improper runtime type check by the interface calls. By sending a specially-crafted request using bytecode, an attacker could exploit this vulnerability to access or modify memory.

CVSS Base score: 6.5
CVSS Temporal Score: Click here.
CVSS Vector:
(CVSS:3.0/AV:N/AC:L/PR:N/UI:N/S:U/C:L/I:L/A:N)

CVEID: CVE-2022-21426

Description: An unspecified vulnerability in Java SE related to the JAXP component could allow an unauthenticated attacker to cause a denial of service resulting in a low availability impact using unknown attack vectors.

CVSS Base score: 5.3
CVSS Temporal Score: Click here.
CVSS Vector:
(CVSS:3.0/AV:N/AC:L/PR:N/UI:N/S:U/C:N/I:N/A:L)

CVEID: CVE-2023-21830

Description: An unspecified vulnerability in Java SE related to the Serialization component could allow a remote attacker to cause a denial of service resulting in a low integrity impact using unknown attack vectors.

CVSS Base score: 5.3
CVSS Temporal Score: Click here.
CVSS Vector:
(CVSS:3.0/AV:N/AC:L/PR:N/UI:N/S:U/C:N/I:L/A:N)

Affected Products and Versions

Affected Product(s)	Version(s)
IBM Sterling External Authentication Server	6.1.0
	6.0.3

Affected Products and Versions

Product	IBM Sterling External Authentication Server	IBM Sterling External Authentication Server
Version	6.1.0.0	6.0.3.0
Remediation & Fix	iFix 03	iFix 07

Workarounds and Mitigations
None.

IBM Sterling B2B Integrator. Upgrade Compatibility.

This article describes upgrade scenarios and supported upgrade paths for IBM Sterling B2B Integrator.

Always check upgrade compatibility before planning an upgrade.

To upgrade the IBM Sterling B2B Integrator environment to a higher version, choose a version compatible with the existing version. In general, you must upgrade to a version that is released after the existing version.

For example:

- Upgrade from v5.2.6.3_6 (release date July-2018) to v6.0.0.0 (release date August-2018) – Compatible
- Upgrade from v5.2.6.3_9 (release date February 2019) to v6.0.0.0 (release date August-2018) – Not compatible

Note: The release date of each version is available in the Release Timeline illustration.

Refer to Sterling B2B Integrator System Requirements for detailed information on JDK versions in IBM B2B Integrator releases.

Previous Releases:
If you are on any of the following versions, you can upgrade to any of the latest versions as illustrated in the Upgrade Matrix below.

- 5.2.5_19 or earlier versions
- 5.2.6.0 all versions
- 5.2.6.1_9 or earlier versions
- 5.2.6.2_5 or earlier versions
- 5.2.6.3_12 or earlier versions

Upgrade Paths

- The chart below lists the compatible upgrade paths for IBM B2B Integrator releases.

Supported	Y	Indicates the upgrade is compatible.
Intermediate	I	Indicates an in-between upgrade and needs to upgrade further. When you upgrade to the latest Fix Pack release of a particular version, first upgrade to the Base release of that version. This is the Intermediate state. You must later upgrade to the Fix Pack release. Note: Do not use the system if the Base release on a higher version is older than the current release. The system remains in a transient state until you upgrade to the latest Fix Pack release. For example: If you want to upgrade to v6.0.0.1 from v5.2.6.3_9, first upgrade to v6.0.0.0 and later upgrade to v6.0.0.1. You must not use the instance at v6.0.0.0.
Not supported	X	Indicates the upgrade is not compatible.

Compatible Versions → Upgrade To

Upgrade From ↓

IBM B2B Integrator Releases	5.2.6.3_16	5.2.6.4	5.2.6.4_4	5.2.6.5	5.2.6.5_4	6.0.0.0	6.0.0.8	6.0.1.0	6.0.1.2	6.0.2.0	6.0.2.3	6.0.3.0	6.0.3.8	6.1.0.0	6.1.0.7	6.1.1.0	6.1.1.4	6.1.2	6.1.2.2
5.2.5_20	Y	I	Y	Y	Y	I	Y	I	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
5.2.6.1_10	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
5.2.6.2_6	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
5.2.6.3_14	Y	I	Y	I	Y	I	Y	X	X	I	Y	I	Y	Y	Y	Y	Y	Y	Y
5.2.6.3_15	Y	X	X	I	Y	I	Y	X	X	I	Y	I	Y	Y	Y	Y	Y	Y	Y
5.2.6.3_16	NA	X	X	I	Y	I	Y	X	X	X	X	I	Y	Y	Y	Y	Y	Y	Y
5.2.6.4	X	NA	Y	Y	Y	I	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
5.2.6.4_2	X	X	Y	Y	Y	I	Y	I	Y	I	Y	Y	Y	Y	Y	Y	Y	Y	Y
5.2.6.4_3	X	X	Y	I	Y	I	Y	I	Y	I	Y	Y	Y	Y	Y	Y	Y	Y	Y
5.2.6.4_4	X	X	NA	I	Y	I	Y	X	X	I	Y	I	Y	Y	Y	Y	Y	Y	Y
5.2.6.5	X	X	X	NA	Y	I	Y	I	Y	I	Y	Y	Y	Y	Y	Y	Y	Y	Y
5.2.6.5_2	X	X	X	X	Y	I	Y	X	X	I	Y	I	Y	Y	Y	Y	Y	Y	Y
5.2.6.5_3	X	X	X	X	Y	I	Y	X	X	I	Y	I	Y	X	Y	Y	Y	Y	Y
5.2.6.5_4	X	X	X	X	NA	I	Y	X	X	X	X	I	Y	X	Y	Y	Y	Y	Y
6.0.0.0	X	X	X	X	X	NA	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
6.0.0.6	X	X	X	X	X	X	Y	X	X	X	X	I	Y	X	Y	Y	Y	Y	Y
6.0.0.7	X	X	X	X	X	X	Y	X	X	X	X	I	Y	X	Y	Y	Y	Y	Y
6.0.0.8	X	X	X	X	X	X	NA	X	X	X	X	I	Y	X	Y	I	Y	Y	Y
6.0.1.0	X	X	X	X	X	X	X	NA	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
6.0.1.2	X	X	X	X	X	X	X	X	NA	I	Y	I	Y	Y	Y	Y	Y	Y	Y
6.0.2.0	X	X	X	X	X	X	X	X	X	NA	Y	Y	Y	Y	Y	Y	Y	Y	Y
6.0.2.1	X	X	X	X	X	X	X	X	X	X	Y	Y	Y	Y	Y	Y	Y	Y	Y
6.0.2.2	X	X	X	X	X	X	X	X	X	X	Y	I	Y	Y	Y	Y	Y	Y	Y
6.0.2.3	X	X	X	X	X	X	X	X	X	X	NA	I	Y	X	Y	Y	Y	Y	Y
6.0.3.0	X	X	X	X	X	X	X	X	X	X	X	NA	Y	Y	Y	Y	Y	Y	Y
6.0.3.6	X	X	X	X	X	X	X	X	X	X	X	I	Y	X	Y	I	Y	Y	Y
6.0.3.7	X	X	X	X	X	X	X	X	X	X	X	I	Y	X	Y	I	Y	I	Y
6.0.3.8	X	X	X	X	X	X	X	X	X	X	X	X	NA	X	Y	X	Y	I	Y
6.1.0.0	X	X	X	X	X	X	X	X	X	X	X	X	X	NA	Y	Y	Y	Y	Y
6.1.0.5	X	X	X	X	X	X	X	X	X	X	X	X	X	X	Y	I	Y	Y	Y
6.1.0.6	X	X	X	X	X	X	X	X	X	X	X	X	X	X	Y	I	Y	I	Y
6.1.0.7	X	X	X	X	X	X	X	X	X	X	X	X	X	X	NA	I	Y	X	X
6.1.1.0	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	NA	Y	Y	Y
6.1.1.1	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	Y	Y	Y
6.1.1.2	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	Y	I	Y
6.1.1.3	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	Y	I	Y
6.1.2	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	NA	Y
6.1.2.1	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	Y
6.1.2.2	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	NA

Note: From 6.1 onward, there is no intermediate upgrade path, which means customer can directly upgrade to fix pack without installing corresponding base/Mod.
For Example,
For a 5263_16 customer to upgrade to 6104, it's a one step upgrade. Customer can directly upgrade to 6104.
For 5263_16 customer to upgrade to 6035, they must first upgrade to 603 and then apply the 6035 fix pack.

Transformation Extender V10.1.2.0 Fix List

This article provides a complete list of fixes for Transformation Extender V10.1.2.0. The most recent fix is at the top of the table.

Note: An APAR link does not work if the APAR has not reached the publish phase yet.

APAR	Description
PH49130	*HIDDEN ITX LOG FILES (BEGINNING WITH A DOT) HAVE RW-RW-RW PERMISSION*
PH49094	*DESIGN STUDIO CRASH WHEN ADAPTER COMMAND LINE LENGTH EXCEEDS ~2048 CHARACTERS*
PH49043	*APP CONNECT ENTERPRISE (ACE) 11 ABEND WHEN THE ITX MAP NODE IS CALLED AND ANY OTHER ADAPTER (APART FROM FILE) IS CALLED*
PH48962	*USED AS INPUT AN UNBOUNDED TEXT TYPE INTERPRETED AS CHARACTER UTF 8, CAUSES MAP TO CRASH*
PH48692	*PATCH FOR APAR PH46687 CAUSES REGRESSION ISSUE WITH LAST FUNCTION*
PH48646	*ITX 10.1.1.X ZLINUX RUNTIME AND MONITORING INSTALLATION INCORRECTLY INCLUDES 31-BIT JRE INSTEAD OF THE REQUIRED 64-BIT JRE*

PH48600	XSLT FUNCTIONS FAIL TO PRODUCE ANY OUTPUT WHEN MAP IS RUN UNDER ACE 11 ON AIX
PH48328	ACE MESSAGE FLOW ITX MAP NODE OUTPUT DEFINED AS XMLNSC DOMAIN ESCAPES APOSTROPHE AS &APOS; IN OUTPUT AND MESSAGE TREE
PH48118	TODATETIME FUNCTION CONVERTING INCORRECT MERIDIAN TIME ZONE
PH48105	AUDIT LOG SHOWS CONTENTRETURN="0" FOR A CARD WHERE WE SHOULD HAVE CONTENTRETURN="28"
PH47902	CRASH ENCOUNTERED USING GETANDSET WITH A LINE FEED CHARACTER
PH47645	ITX FAILING ON XML SHORTCUT TAGS THAT ARE OPTIONAL IN THE SCHEMA
PH47466	ADDITIONAL COMMA FOUND IN JSON DATA WHEN THE FIRST ELEMENT IS OPTIONAL AND EVALUATES TO NONE
PH47292	AN APP CONNECT ENTERPRISE FLOW BEING STOPPED RESULTS IN A CORE FILE GENERATION WHEN INTEGRATED WITH TRANSFORMATION EXTENDER
PH47285	ITX 10.1.1.X DESIGN SERVER INSTALLATION FAILS ON FRENCH LOCALE MACHINES WITH UNZIPDESIGNSERVER.BAT SCRIPT STEP
PH46687	CRASH USING INVALID JSON AS INPUT TO MAP
PH46400	?ERROR -12 : PREMATURE END OF JSON FILE?, INCONSISTENTLY OCCURS WHEN PARSING JSON DATA
PH46377	INCORRECT ERROR MESSAGE WHEN RUNNING A MAP IN LAUNCHER FOR A MISSING INPUT WHEN GETANDSET() IS USED ON THE MISSING INPUT
PH46235	AUDIT LOG SHOWS CONTENTRETURN="0" FOR A CARD WHERE WE SHOULD HAVE CONTENTRETURN="21"
PH46110	PAGE USAGE COUNT ERROR WHEN RUNNING A MAP
PH45920	RESULTS FROM ITX V1 REST API UNNECESSARILY DUPLICATED
PH45919	GOOGLE CLOUD STORAGE ADAPTER FAILS TO COMPLETE SUCCESSFULLY - ERROR MESSAGE '-305 JAVA EXCEPTION OCCURRED'
PH45887	FAILURE TO CALL A STORED PROCEDURE USING AZURE SQL ADAPTER IN A CARD WHEN USED WITH TEXT BASED ARGUMENT
PH45824	RESOURCECHANGES.BAT SCRIPT NOT WORKING AS IT APPEARS THERE ARE MISSING COMPONENTS ON ITX DESIGN SERVER IMPLEMENTATION
PH45778	LAUNCHER CRASH OCCURS DURING SHUTDOWN ON AIX WHEN A JAVA BASED ADAPTER LISTENER IS DEFINED AS A SOURCE EVENT WATCH
PH45593	CALL TO A STORED PROCEDURE FAILS TO RETRIEVE VARBINARY COLUMN DATA
PH45288	UNABLE TO CALL A MS SQL SERVER STORED PROCEDURE WITH ARGUMENT DEFINED AS VARBINARY DATATYPE AND DATA > 8000 BYTES
PH45111	LAUNCHERADMIN.SH "INVALID ARGUMENT FOR THE PORT RANGE OPTION. ENDPORT SHOULD BE AT LEAST 3 PORTS APART FROM STARTPORT."
PH45106	CHANGING CERTIFICATE PASSWORD RESULTS IN ERROR
PH45017	WTX MAP SERVICE INSTALL SCRIPT DOES NOT REMOVE THE PREVIOUS VERSIONS OF STERLING B2B INTEGRATOR MODULE M4SI.JAR
PH44942	MAP OUTPUT CARD BACKUP FEATURE CREATES ZERO BYTE OUTPUT FILE WHEN USING ITX 10.1.0.1 OR 10.1.1.0 VERSIONS
PH44889	LOG4J IS CONTAINED WITHIN COM.HCL.HIP.ADAPTERS.M4C OBOL/KOOPA.JAR
PH44795	SOAP ADAPTER NOT FOUND ERROR ON 2ND AND EVERY SUBSEQUENT MAP EXECUTION UNDER LAUNCHER WITH ITX 10.1.1
PH44400	ITX MAP ENDS WITH 'ASSERTION FAILED: 0, FILE UINVCHAR.C, LINE 208 IOT/ABORT TRAP(COREDUMP)' ON AIX
PH43824	ODBC ADAPTER INSERTS GARBLED CHARACTERS INTO COLUMN WHEN LARGER DATA IS BEING INSERTED AND BUFFER IS EXPANDED
PH43391	ITX 10.1.1 LAUNCHER CRASHES AT STARTUP ON AIX WITH RESULT 132
PH43353	DIFFERENCE IN OUTPUT BETWEEN WTX 8.4.1.3 AND ITX 10.1.0.0 WHICH NAY OCCUR DUE TO SERIES CONSUMING FUNCTIONS
PH42789	ITX USAGE OF LOG4J AND THE IMPACT OF VULNERABILITY CVE-2021-44228
PH42744	UNABLE TO GET A ZERO PRODUCED IN A MANDATORY XML FIELD WHEN NULL RECEIVED FROM DATABASE IN ITX MAP OUTPUT
PH42358	UNABLE TO CALL A MS SQL SERVER STORED PROCEDURE WITH ARGUMENT DEFINED AS VARBINARY DATATYPE AND DATA > 8000 BYTES
PH42293	THE HTTP LISTENER IS VULNERABLE TO THE SWEET32 BIRTHDAY ATTACK
PH42292	THE HTTP LISTENER IS VULNERABLE TO DOS (DENIAL OF SERVICE) ATTACKS
PH42034	UNABLE TO INSERT MORE THAN 8000 BYTES INTO A MICROSOFT SQL SERVER DATABASE TABLE COLUMN DEFINED AS VARBINARY DATATYPE
PH41127	INTERMITTENT GSKIT GSK_SECURE_SOC_INIT FAILED AND RETURNED 406 GSK_ERROR_IO ERROR MESSAGE WHEN USING HTTP ADAPTER

Release strategy for IBM Sterling B2B Integrator and IBM Sterling File Gateway v6.0.x and v6.1.x onwards

This article provides details of the release policy for IBM Sterling B2B Integrator and IBM Sterling File Gateway of versions v6.0.x and v6.1.x onwards.

IBM Sterling B2B Integrator and IBM Sterling File Gateway follow the IBM V.R.M.F release policy.

The release details are as follows:

- **Major/Minor:**
  - Each Major/Minor release has new features (enhancements) and defect fixes.

- It receives Fix Packs for the duration of 3 years (Standard Support) and 2 years (Extended Support).

- **Modification (Mod Pack):**
  - Each Mod Pack release has new features (enhancements) and defect fixes.
  - It receives Fix Packs for 6 months after the release of the subsequent Mod Pack. To receive further Fix Packs, you must upgrade to the latest Mod Pack.
  - For example: Mod Pack 6.0.X is released in March 2019 and the subsequent Mod Pack 6.0.X+1 is released in June 2019. Mod Pack 6.0.X receives Fix Packs till December 2019, which is 6 months after June 2019.
  - The last Mod Pack of a Major/Minor release continues to receive Fix Packs till the end of Extended Support of that particular Major/Minor release.
- **Fix Pack:** Each Fix Pack release has only defect fixes.
- **Interim Fix (iFix) Pack:** A specific fix for a single customer.

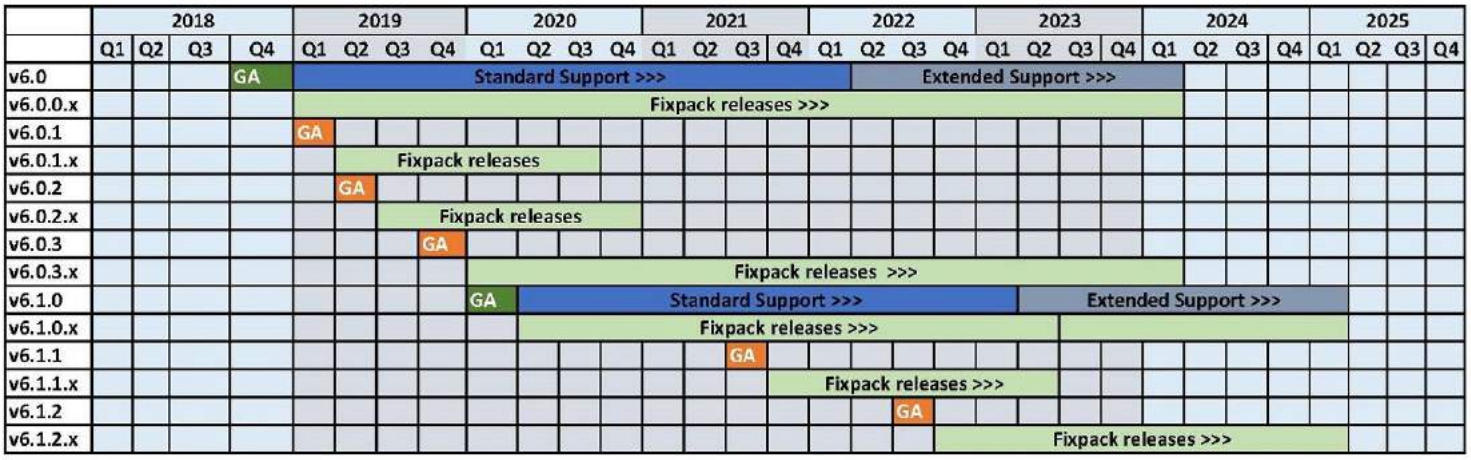

Troubleshooting

Sterling B2Bi node goes down abruptly due to com.ibm.crypto.fips.provider.FIPSRuntimeException

Problem

IBM Sterling B2Bi node goes down or dashboard throws HTTP ERROR 500 com.ibm.crypto.fips.provider.FIPSRuntimeException

Symptom

Error

Dashboard:

HTTP ERROR 500
com.ibm.crypto.fips.provider.FIPSRuntimeException
URI: /dashboard/
STATUS: 500
MESSAGE: com.ibm.crypto.fips.provider.FIPSRuntimeException
SERVLET: default
CAUSED BY:
com.ibm.crypto.fips.provider.FIPSRuntimeException

Caused by:
com.ibm.crypto.fips.provider.FIPSRuntimeException
at com.ibm.crypto.fips.provider.HASHDRBG.engineNextBytes(Unknown Source)
at com.ibm.crypto.fips.provider.SHA2DRBG.engineNextBytes(Unknown Source)
at java.security.SecureRandom.nextBytes(SecureRandom.java:471)
at java.security.SecureRandom.next(SecureRandom.java:494)
at java.util.Random.nextInt(Random.java:340)

Security.log

ERROR LM.refresh run caught Exception
ERROR [1634112649079] null ERRORDTL [1634112649079]com.ibm.crypto.fips.provider.FIPSRuntimeException
at com.ibm.crypto.fips.provider.X509Factory.engineGenerateCertificate(Unknown Source)
at java.security.cert.CertificateFactory.generateCertificate(CertificateFactory.java:407)
at com.sterlingcommerce.security.lc.LicenseSig.verify(LicenseSig.java:596)
at com.sterlingcommerce.security.lc.FeatureSet.load(FeatureSet.java:455)
at com.sterlingcommerce.security.lc.LM.loadMapFromFile(LM.java:2500)
at com.sterlingcommerce.security.lc.LM.loadMap(LM.java:2559)
at com.sterlingcommerce.security.lc.LM.refresh(LM.java:3188)
at com.sterlingcommerce.security.lc.LMThread.run(LMThread.java:210)
at

java.lang.Thread.run(Thread.java:818

Wf.log

ERROR [1658996685245] null ERRORDTL [1658996685245]com.ibm.crypto.fips.provider.FIPSRuntimeException
at com.ibm.crypto.fips.provider.HASHDRBG.engineNextBytes(Unknown Source)
at com.ibm.crypto.fips.provider.SHA2DRBG.engineNextBytes(Unknown Source)

System.log and noapp.log

ALL 000000000000 GLOBAL_SCOPE com.ibm.crypto.fips.provider.FIPSRuntimeException

Cause

The FIPSRunTimeException could be a case where the IBMJCEFIPS provider causes issues in non-FIPS mode. IBMJCEFIPS is the security provider libraries that are part of the JDK. By default , B2Bi is running in non-fips mode. This issue happens when the application gets many concurrent calls(high volume) and somehow application java ends up with a race condition between concurrent JDK calls.
This can happen randomly and is quite unpredictable. When this happens, it can leave the JVM in an inconsistent state therefore node get freeze.

Resolving The Problem

To resolve the issue, follow the below steps

1) Stop B2Bi node using ./hardstop.sh from <B2Bi_install>/bin folder.

2) Backup the current java.security file found under <B2Bi_Install>/jdk/jre/lib/security

3) Edit java.security file , make the below changes and save the file

Move com.ibm.crypto.fips.provider.IBMJCEFIPS down
e.g., From
security.provider.2=com.ibm.crypto.fips.provider.IBMJCEFIPS
security.provider.3=com.ibm.crypto.provider.IBMJCE
To
security.provider.2=com.ibm.crypto.provider.IBMJCE
security.provider.3=com.ibm.crypto.fips.provider.IBMJCEFIPS

4) Start B2Bi node using ./run.sh from <B2Bi_install>/bin folder

Perform these changes in all nodes if you are running B2Bi in Cluster environment.

Connect:Direct for UNIX: Running Behind a Load Balancer

Considerations for running Connect:Direct for UNIX behind a Load Balancer.

1. Introduction.

Customers may wish to use a connection load balancer to distribute incoming Snode sessions across multiple instances of Connect:Direct for UNIX. In such an arrangement the Connect:Direct for UNIX instances behind the load balancer appear as a single Snode to the outside world. A given incoming session is distributed to one of the instances based on criteria defined in the load balancer.

If an incoming session fails and is restarted by the remote PNODE, then the restarted session may be assigned to any of the instances behind the load balancer and will not necessarily be established with the original SNODE instance.

Generally, from the point of view of the nodes behind the load balancer only incoming or “SNODE” sessions are affected by the load balancer;

PNODE, or outgoing sessions operate the same way they normally do. Connect:Direct for UNIX includes an enhancement that allows COPY checkpoint/restart and RUN TASK resynchronization to work when Connect:Direct is set up in this way.

This document points out some extra considerations that come into play only when instances of Connect:Direct for UNIX are configured to operate behind a connection load balancer. These considerations can be categorized as:

- SNODE server characteristics
- SNODE Connect:Direct node setup
- Shared SNODE work area setup
- “Stranded” SNODE TCQ entries
- RUN TASK RESTART parameter
- Process Statistics
- Remote Snode alternate.comminfo parameter

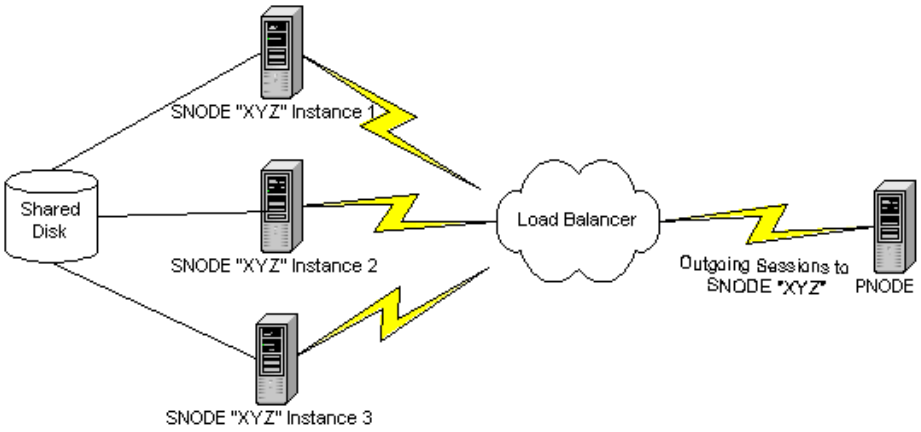

2. SNODE Server Considerations.

- The servers used for the Connect:Direct for UNIX instances behind the load balancer must all have access to a common shared cluster disk storage since any COPY statement source and destination files for SNODE processes must reside in directories accessible to all of the servers. All nodes must have access to a common SNODE work area. If you want to use NFS for the shared drive, the NFS version must be v4 or greater.
- The system clocks on all the servers must be synchronized for COPY checkpoint/restart and RUN TASK synchronization to work.
- The administrator user ID used to install Connect:Direct for UNIX must be defined the same on each server and must be the same user and group number on each server.
- The servers should all be of the same hardware platform type and running the same Operating System.

3. SNODE Connect:Direct Node setup.

- One Connect:Direct for UNIX node should be installed on each server behind the load balancer.
- Each node must be the same Connect:Direct for UNIX version and maintenance level.
- Each node must be installed by the same user ID.
- Each node must have the same C:D node name.
- Each node must have the same node-to-node connection listening port.
- Each node must specify the same path for the snode.work.path attribute of the ndm.path initialization parameter in the initialization parameter file.
- Each node must specify its local host name in the tcp.api listening address specification of the netmap local node entry.

4. Shared SNODE work area setup.

A directory should be established for the shared SNODE work area used by the Connect:Direct for UNIX nodes behind the load balancer. The path to this directory must be specified in the snode.work.path attribute of the initialization parameter file for each instance. Following is an example for this parameter.

Miscellaneous Parameters
ndm.path:path=/<local_install_path>/cdunix:\:snode.work.path=/<shared_disk_mount_point>/cdunix/shared:

SNODE return code files (steprc files) and COPY checkpoint information are created and stored in this area when the snode.work.path attribute is specified in the initialization parameters.

This directory should be owned by the Connect:Direct administrator ID and must be accessible to all of the servers behind the load balancer. It must be on a cluster file system. If you want to use NFS for the shared drive, you must use NFS v4 or greater.

The initial size of the shared work area is dependent upon the number of Connect:Direct for UNIX servers behind the load balancer and the total workload on the servers. You can start with 50Mb, or more. You will need to monitor the usage of this disk space from time to time and be prepared to increase the size if necessary.

5. "Stranded" SNODE TCQ entries.

As mentioned above, when Connect:Direct instances are setup behind a load balancer and an incoming session fails and is restarted by the remote PNODE, then the restarted session may be assigned to any of the SNODE instances behind the load balancer and will not necessarily be established with the original SNODE instance. In this scenario, each SNODE instance that receives a session for a given process creates a TCQ entry for the process. (Note that each SNODE instance has its own TCQ file; these are not shared among SNODE instances. Only the work files created in the shared work area are shared among instances.)

Consider this scenario based on the drawing depicted in the introduction. If a process P is submitted on the remote PNODE, a session may be established with SNODE instance 1, which creates an SNODE TCQ entry in its TCQ file. If process execution is interrupted and the process is requeued and restarted, the restart session may be established with SNODE instance 2, which also creates an SNODE TCQ entry in its TCQ file. If the process runs to completion on the restart session with SNODE instance 2, then that instance deletes its TCQ entry for process P and also deletes any SNODE work files for process P from the shared SNODE work area. However, SNODE instance 1 still retains its TCQ entry for process P. This TCQ entry is "stranded" since the process has completed and the work files have been deleted.

Stranded SNODE TCQ entries are checked automatically when Connect:Direct for UNIX is initialized and also when the TCQ is scanned

periodically during product execution. The stranded SNODE TCQ entries will be deleted according to the parameter value set for 'ckpt.max.age' in the "TCQ Information" section of the 'initparm.cfg' file. The default value is 8 days before automatic deletion occurs. No administrator action is required.

6. Run Task Restart parameter.

The Run Task Restart parameter works slightly differently when Connect:Direct for UNIX is configured with shared SNODE work areas.

Ordinarily when process execution is interrupted during a Run Task step and subsequently restarted, then the setting of the Run Task Restart parameter determines whether or not the task is restarted only if it is determined that the task is no longer active. Connect:Direct for UNIX determines whether or not the task is still active by checking if the task system process is still active on the server.

When shared SNODE work areas are configured and the Run Task is on the SNODE, then at restart time it is generally not possible for Connect:Direct for UNIX to determine whether the original task is still active or not since the restart session may be with a different SNODE instance on a different server machine.

Therefore, in this scenario, the task is either restarted or not restarted based solely on the setting of the Run Task Restart parameter. Because of this, caution should be used when specifying Run Task Restart = Yes. It is possible that a task could be restarted even though it may be active on another server machine.

7. SNODE Process Statistics.

Because of the fact that statistics files are not shared among the SNODE instances behind the load balancer, when a process is interrupted and restarted to a different SNODE instance, the statistics records for that process will be distributed between the two SNODE instances involved. Currently there is no provision for selecting all the statistics records for a given process in a single operation when the records are distributed across multiple SNODE instances.

8. PNODE to SNODE Considerations.

When the Connect:Direct for UNIX servers behind the load balancer are acting as Pnodes sending to an Snode, the remote Snode needs to have the Netmap record for the Pnodes configured with the "alternate.comminfo" parameter. This will allow Netmap checking to work successfully on the remote Snode.

CDU Example:
XYZ:\
:comm.info=9.1.2.2;1364
:\ <<This is the IP address
of the load balancer>>
:alternate.comminfo=9.1.
2.3, 9.1.2.4,
9.1.2.5:\ <<<These are
the IP addresses of the
individual nodes behind
the load balancer>>>

Support for distributing API connections through a load balancer is limited to being used for submit process commands without a maxdelay specification. When a process is submitted with maxdelay, the connection is idle until the process is completed. Load Balancers tend to kill idle connections. A "Best Practice" for submitting a long running process via a Load Balancer distributed API connection is to submit the process without maxdelay, keep the API connection open, and periodically poll the server for process results. For example, a C:D File Agent operating outside of the CDU LB cluster could have its connections distributed via the load balancer.

Query and configuration type commands should be submitted directly to each Connect:Direct for UNIX server via its own API Listening port.

9. FASP Considerations.

CDU uses both TCP and UDP for FASP connections. When the CDU environment is the Snode, destination address affinity persistence, also known as 'sticky persistence', needs to be enabled in the load balancer.

DAAP supports both TCP and UDP protocols. This will direct session requests to the same server based solely on the destination IP address of a packet.

10. Example Illustrations.

These illustrations contain examples of the configuration settings used to enable the connectivity between the two representative companies.

The first is an example of the logical representation of the connectivity of the load balancer environment. It´s an example only – your environment may be different.

The second is an example of the physical connectivity that one might use to setup this environment. It´s is an example only – your environment may be different.

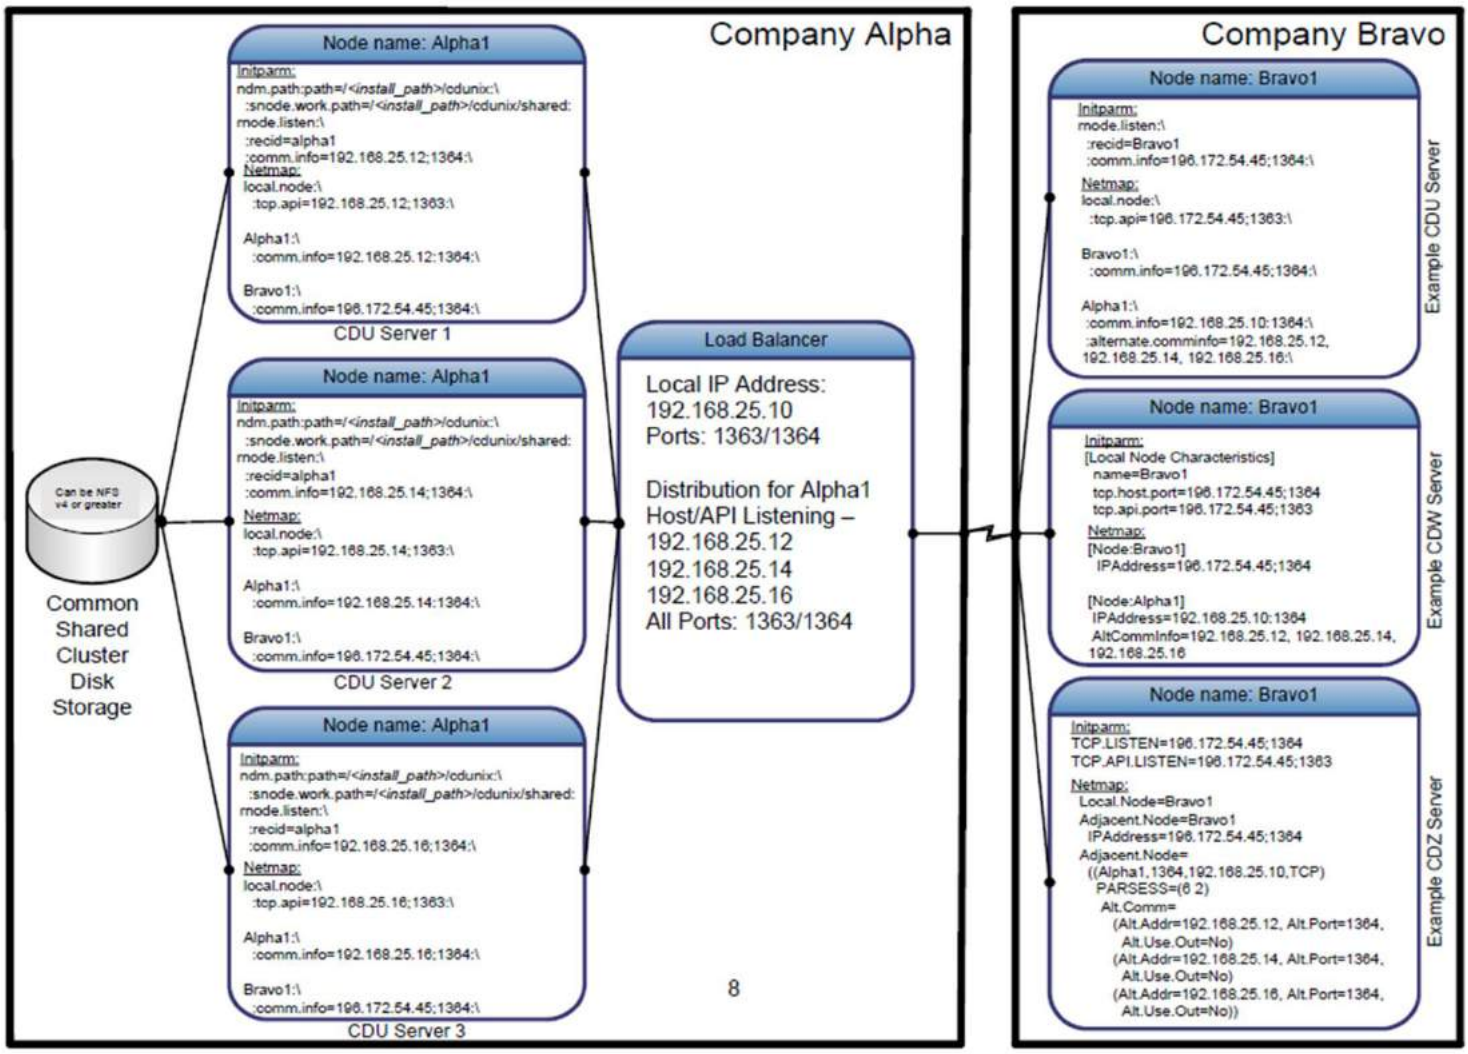

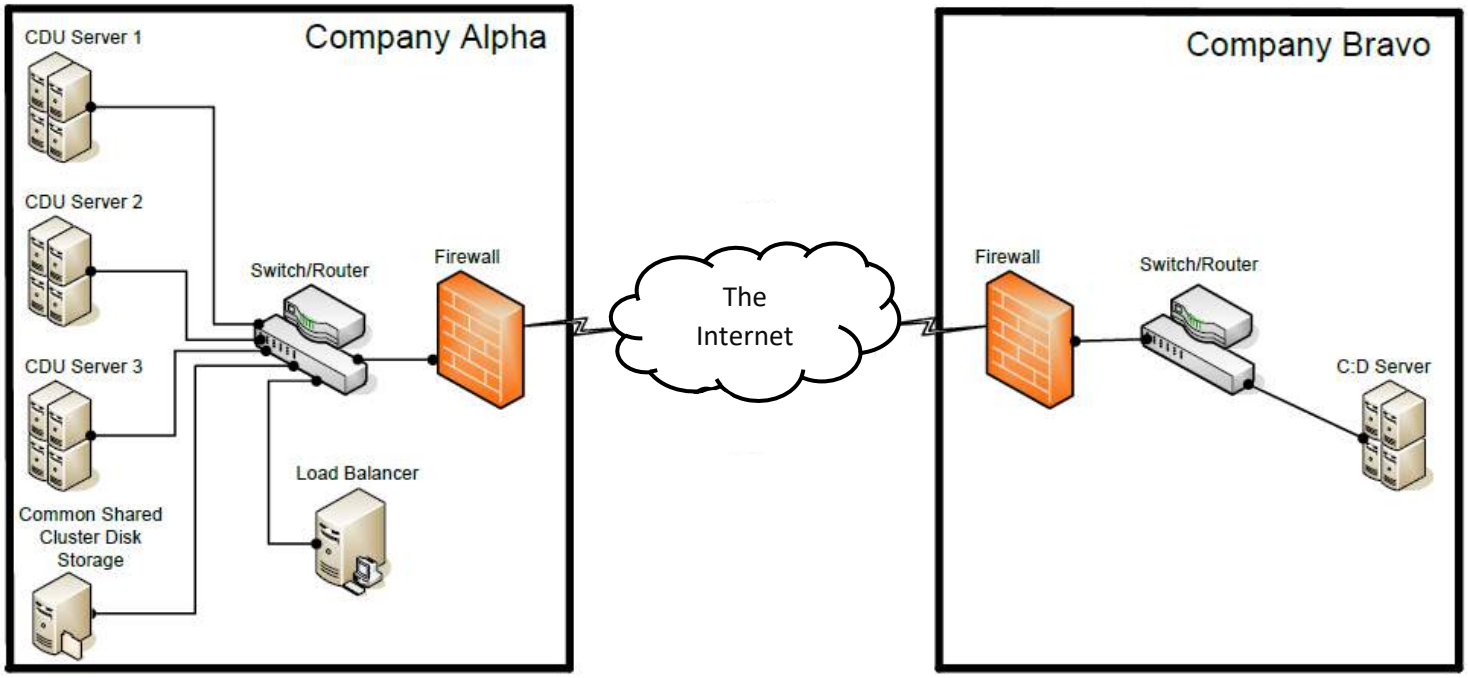

Communication

If you need more information about any of the contents of our newsletter, please do not hesitate to contact us. We will be happy to answer your questions.

info@b2b.solutions