November 2016

Best Practices for NJE on z/VM: Security Configuration Steps for RSCS and VMBATCH

Introduction

The NJE communication protocol is primarily represented on z/VM® through the *Remote Spooling Communications Subsystem (RSCS).* RSCS is a special-purpose z/VM subsystem which allows for the sending and receiving of files, messages, jobs, and commands over a network. RSCS is built using, among other things, the NJE protocol and can communicate with other NJE-compatible subsystems such as VSE/POWER®, JES2 and JES3.

The purpose of this document is to highlight best practices for RSCS in order to facilitate subsystem operability without the risk of security exposures. It will also discuss configuration for z/VM batch facilities, as these may receive jobs from RSCS.

Configuration of TLS Encryption for RSCS

Note: this section assumes that the PTFs for APARs PI56474 (TCP/IP) and VM65788 (RSCS) have been applied to your z/VM 6.3 system.

Special consideration must be given when securing network connections that potentially have a public surface, or alternately exist outside of an isolated control plane. In cases such as these, encryption of network connections is strongly encouraged by IBM. This is especially true in cases where a network connection may traverse multiple nodes on multiple platforms – an NJE network is only as strong as its weakest link.

In order to enable encryption of TCPNJE connections for RSCS, one should configure the z/VM TLS/SSL Server to support general encryption for TCP/IP connections. This involves enabling particular protocols and cipher suites in the DTCPARMS file, as well as adding in appropriate digital certificates into a certificate database managed by the gskkyman application. The full set of instructions can be found in the *z/VM 6.3 TCP/IP Planning and Customization Guide* (SC24-6238-07), Chapter 16: Configuring the SSL Server (PDF).

Once the TLS/SSL Server is enabled, encryption between z/VM and another TCPNJE node (z/VM or otherwise) requires the addition of a certificate label on the PARM statement. This takes the format:

```
TLSLABEL=ZVMRSCS1
```

In the above case, the TLS label is the eight-character uppercase certificate label for the server certificate stored in the TLS/SSL Server's certificate database. A certificate is specified on a per-linkid basis, rather than on a per-node basis, so care must be taken to apply the new TLSLABEL= option for all pertinent TCPNJE links associated with the RSCS node. A TCPNJE connection (comprised of two unidirectional links) shall be established when the certificates of the local and remote ends of the connection validate trust based upon their analyses of certificates presented. This remains true whether the remote node is based on z/VM, z/OS® JES, or VSE/POWER. While the labels are not required to be the same on each node (see below), the certificates must be based on a common chain of trust (i.e., have a common root certificate or intermediary CA certificate).

Note that it is recommended that certificates for RSCS be deployed on a per-LPAR or per-node basis, to prevent a mass grant of authority. While a common certificate may be appropriate in test environments, it is discouraged for production environments. Also note that the z/VM TLS/SSL Server issues errors when self-signed certificates are used.

Note that the specific format of the certificate will influence the type of encryption used by the TLS/SSL Server. Algorithmic complexity and key values should be managed in accordance with enterprise security policy. When this document was published, the most common recommendation would be to use TLS 1.2, AES 256, RSA 2048 or greater, and SHA-256. This is expected to change over time.

Configuration for RSCS

RSCS should be configured with the security of your z/VM LPAR in mind. This means that the security context of RSCS should be considered before enabling communication:

- Examine the z/VM privilege classes for RSCS virtual machines. RSCS uses the CP MSGNOH command, which requires Class B authority. Since there are other commands in Class B which may not be pertinent to RSCS, consider remapping CP MSGNOH to a user-defined privilege class (e.g., “Privclass R”) which contains CP MSGNOH exclusively.
- Do not grant Class D authority to the RSCS virtual machines.
- If RACF® is enabled on your z/VM system, enable appropriate RACF protections for RSCS, including control of the TAG command.
- Exercise appropriate control of Diag x’F8’ for Secure Origin ID support.
- Examine which virtual machines shall be granted access to XAUTOLOG RSCS.
- Examine which virtual machines, if any, shall have alternate console access to RSCS.
- Examine which virtual machines shall have either secondary user access or observer access to RSCS. Note that secondary user access confers to a virtual machine all authorities to which RSCS is granted; therefore, secondary users should not be defined purely for auditing purposes.
- Accounting records should be created through Diag x’4C’ for RSCS.
- Access to the GCS segment should be restricted. RSCS should be authorized to the GCS segment.

In addition to z/VM configuration, RSCS should be tailored to meet local security policy requirements:

- Limit the number of AUTHorized users in the RSCS configuration file. Beware of 'blanket' authority. Consider the NOCP option so an authorized user can't issue RSCS CP commands
- Make use of the HIDECHAR character in the config file for information you don't want to be visible in a query command output.
- OPTION JOBName=Userid should be used to track communications. Consider the following additional parameters:
  - LISTPROC – limits the number of dataset headers a system can handle; used to mitigate potential spamming of large numbers of users
  - LOOPING – user LOOP checking to prevent files from running through the network continuously.
  - MAXHOPS – detects routing loops
  - QMSGLIM – prevents users from crashing system when querying very large systems and sending the output back through S&F nodes
  - SECORGID – use is recommended.
- Use RNPass and/or RLPass password. Requiring a destination-system password prevents unauthorized virtual machines from sending jobs through RSCS.
- Consider using Channel-to-Channel adapters instead of TCP/IP for RSCS communication paths. z/VM RSCS does not support the encryption of links, so isolated communication paths (especially those which offer hardware-native encryption) are preferred to open network traffic.
- Use RSCS Exits to implement additional security controls as pertinent to your security policy. RSCS Exits can help you monitor or restrict commands and files going through your network. Consult the References section for more details on which RSCS exits may be pertinent to security configuration.

Finally, RSCS should be audited and adapted to track security pertinent events and potential threats to performance:

- Use SLOWDOWN to mitigate the possibility of flooding spool
- Use the RSCS EVENTS scheduler to do periodic checks of the system, spool, links, queued files, and any other items of interest specific to your policy.

Notes on Batch Processing for z/VM

RSCS can be used as an NJE communication path for the transmission of batch jobs to certain z/VM servers. The primary batch facility on z/VM is VMBATCH. If using VMBATCH, whether in an RSCS-enabled environment or not, consider the following steps before enabling it to process jobs.

- Thoroughly read Chapter 2 of the VMBATCH Installation, Customization, and Administration manual to understand how VMBATCH operates. Pay particular attention to the section on security.
- Do not disable Diagnose x'D4'. Diag x'D4' assigns spool files to the alternate userid and not to the task machine itself. If files spool files are owned by the task machine, it allows the next job access to those spool files should the previous job fail.
- Use an ESM such as RACF to protect userids and minidisks in your system. RACF also provides controls for Alternate Userid access (Diag x’D4).
- Recommend that users of VMBATCH **not** spool their job output but capture job results in others ways.
- Customize the DGRUCD EXEC which is the Command Screening user exit to protect which commands will be accepted and which will be rejected. Commands can be restricted on a per-user basis. The supplied sample exit is a good starting point but should be customized further.
- Customize the DGRUJB EXEC which is the Job Screening user exit to protect which jobs will and will not be accepted. You can also cancel a job, alter a job's job control options, permit a job that normally wouldn't be permitted, and directly set an alternate userid. The supplied sample exit is a good starting point but we recommend you customize this to your installation.
- Customize the DGRIDO EXEC which is the Job Problem exit. You will want to customize this to reflect changes you have made to the DGRUJB EXEC.
- Customize the DGRUMSG EXEC which is the Message Screening exit. This will allow you to monitor messages issued by the VMBATCH monitor to the users.
- Customize the DTRUAC EXEC which is the Job Accounting exit. This will allow you to create tracks for jobs run on your system.
- Use password protection for userids and minidisks in your system as jobs running in the VMBATCH task machines can link and access minidisks just like any other job.
- Limit the number of users which have authority to AUTOLOG and XAUTOLOG the task machines. Only the VMBATCH virtual machine should require this authority.
- Limit the number of ADMIN's authorized in the CONTROL file. Reducing the number of administrators will control privileged access to trusted users.
- Use LLS (Load Level Scheduling) to help keep malicious users from maxing out your system or a task machine.
- Carefully customize your CONTROL file. Some of the options for CLASS allow NONE as a parameter. For example, MAXCPU will allow none. A malicious user could use

this to max out a CPU. MAXPRT and MAXPUN will also accept NONE. Someone could flood spool with spool records if this is set to NONE. We suggest you carefully consider your settings and not accept NONE.

- Carefully consider if you want to allow remote users access to your local batch machines. If you do you must customize the DGRUCD and DGRUJB exits to allow this. You also must supply an alternate userid to be used for running jobs. If you do not have a default one you will have the opportunity to supply one with each job submission or reject the jobs at your discretion. Thoroughly read the section in Chapter 5 of the Installation, Customization, and Administration book entitled Enabling a Facility for Remote Processing.

In addition to VMBATCH, the CMS operating system supports a CMSBATCH functionality. IBM **does not recommend** the enabling and use of CMSBATCH in any context. This function is present for legacy purposes only and should not be used in production environments.

Reference

For more information on exit configuration, refer to the *IBM z/VM RSCS Networking Exit Customization Manual* (SC24-6224-00). In particular:

- Exits 11, 12, 13, 14, 15, 16 for NJE dataset headers and trailers
- Exit 19 for command screening
- Exit 21 for accepting or rejecting spool files (pertinent for performance of security functions)
- Exit 29 for Unknown Commands
- Exit 32 for NMR Reception
- Exit 31 for Sort Priority Change
- Exits 37, 38, 39, 40, 41, 42, and 43 for S&F files

Authors

Colleen Brown – z/VM Development and Support – browncol@us.ibm.com
Leslie Geer III – z/VM RSCS Support – lesgeer@us.ibm.com
Brian Hugenbruch – z Systems Virtualization Security – bwhugen@us.ibm.com

Best Practices for NJE on z/VM: Security Configuration Steps for RSCS and VMBATCH

IBM Systems
Route 100
Somers, New York 10589
U.S.A.
Produced in the United States of America,
11/2016

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

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.It is the user’s responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.